

Concept Representations and Nonmonotonic Inferences

Peter Grdenfors
Lund University Cognitive Science
Kungshuset. S-222 22 Lund, Sweden
e-mail: Peter.Gardenfors@fil.lu.se

How to Make your Logic Fuzzy
(PRELIMINARY VERSION)

D. M. Gabbay
Department of Computing, Imperial College
180 Queen's Gate, London SW7 2BZ (UK)
e-mail: dg@doc.ic.ac.uk

1 Nonmonotonic aspects of concepts

The basic idea of nonmonotonic inferences is that when more information is obtained about an object, some inferences that were earlier reasonable are no longer so. An important point that is often overlooked, is that information about an object may be of two kinds: *propositional* and *conceptual*. When the new information is propositional, one learns new facts about the object, e.g. that x is a penguin. When the new information is conceptual, one *categorizes* the object in a new way, e.g., x is seen as a penguin instead of as just a bird. It is important to notice that describing information as propositional or as conceptual does not mean that these kinds of information are in conflict with one another. On the contrary, they should be seen as different perspectives on how, e.g., observations are described (see Grdenfors 1994).

The theory of nonmonotonic inferences has focused on propositions; hence it has been seen as a nonmonotonic *logic*. However, in the examples discussed in the literature, the great majority derive from the nonmonotonicity of concepts. For example, the *default* rules studied by Reiter (1980) and his followers have been conceived of as inference rules, although a more natural interpretation of “defaults” is to view them as relations between *concepts*. When something is, for instance, categorized as a fruit, it will also, by default, be categorized as sweet.

It may be argued that there is no harm done in focusing on the propositional side of nonmonotonicity since information about categorization can be quite naturally transferred to propositional information: categorizing x as an emu, for example, can be expressed by the proposition “ x is an emu.” However, this transformation into the propositional form tends to suppress the *structure* of concepts. Once one formalizes categorizations of objects by *predicates* in a first order language, there is a strong tendency to view the predicates as primitive atomic notions and to forget that there are rich relations between concepts that disappear when put into standard logical formalism. Indeed, the fact that the concept of an emu is a subcategory of bird is often represented by an explicit axiom in the form of a universal sentence “ $(x)(Ex \rightarrow Bx)$.” However, if the structure of concepts were built into the predicates of the language themselves, such an axiom would be totally redundant.

The main purpose of this article is to bring out the role of conceptual structure in nonmonotonic inferences. First of all, I will argue that there are several, albeit related, kinds of nonmonotonic inferences that appear in the use of concepts. In order to explain these phenomena, I will go beyond propositional representations. The representational framework I will use is based on *conceptual spaces* (Grdenfors 1990a, 1990b, 1992, to appear).

To be sure, there is in the literature one well-known theory of nonmonotonic inferences that focuses on *conceptual relations*, namely *inheritance networks* (Touretsky 1986, Makinson and Schlechta 1991). However, in the theory of inheritance networks, concepts are represented by (non-structured) *points*, and their relations by two kinds of links: “is-a” and “is-not-a.” This form of representation is far too meagre to handle the structural relations between concepts that are exploited in inferences, monotonic as well as nonmonotonic. In contrast, I shall argue that a theory of conceptual structure is necessary in order to understand different kinds of nonmonotonic inferences with concepts.

As a challenge to any theory about nonmonotonic inferences, I would like to point out the following nonmonotonic aspects of concepts:

(a) *Change from a general category to a subordinate*

This is the most well-known nonmonotonic aspect of concepts. When we shift from applying a “basic” category (a term borrowed from prototype theory) like *bird* to an object x , to applying a “subordinate” category like emu, we often give up some of the (default) properties associated with the basic category: a bird is normally small and sings and flies, while an emu has none of these properties.

(b) *Contrast classes*

In standard uses of first order logic, *combinations* of concepts are expressed by (classical) *conjunctions* of predicates. However, there are many different combinations of concepts that cannot be analysed in this manner. Consider the seemingly innocent concept *red*. In the *Advanced Learner’s Dictionary of Current English*, it

is defined as “of the colour of fresh blood, rubies, human lips, the tongue, maple leaves in the autumn, post-office pillar boxes in Gt. Brit.” This definition fits very well with letting *red* correspond to the normal region of the color space (see Section 2). Now consider *red* in the following combinations:

- Red book
- Red wine
- Red hair
- Red skin
- Redwood

In the first example, *red* corresponds to the dictionary definition, and it can be combined with *book* in a straightforward extensional way that is expressed by a conjunction of predicates in first order logic. In contrast, *red* would denote *purple* when predicated of wine, *copper* when used about hair, *tawny* when of skin, and *pinkish brown* when of wood. Thus the class of objects that the concept is applied to, the so called *contrast class*, changes the meaning of *red* in a nonmonotonic fashion (Brostrm 1994).

(c) *Metaphors*

Even more drastic combination-effects occur in metaphorical uses of concepts. For example, when we talk about a *red newspaper*, we don't expect it to be printed on red paper, only to express a certain political viewpoint. The kind of conceptual change involved in a metaphor corresponds to a revision of the concept, and thus the inferences involved in metaphorical uses of concepts parallels *belief revisions* that are modelled in propositional systems (Grdenfors 1988; for a comparison between nonmonotonic inferences and belief revision, see Grdenfors and Makinson 1994).

Metaphors are notoriously difficult to handle within standard first order logic, since it is assumed that the reference of a predicate is fixed in advance. In most logical theories these linguistic figures have thus been treated as deviant phenomena and have been ignored or incorporated via special stylistic rules.

(d) *Context effects*

Combinations of concepts can result in nonmonotonic effects as exemplified by contrast classes and metaphors. Sometimes the mere context in which a concept is used may trigger different *associations* which lead to nonmonotonic inferences. Barsalou (1987, p. 106) gives the following example: “when *animals* is processed in the context of milking, *cow* and *goat* are more typical than *horse* and *mule*. But when *animals* is processed in the context of riding, *horse* and *mule* are more typical than *cow* and *goat*.”

Another example of how the context affects the application of concepts is the following, due to Labov (1973). He showed subjects pictures of objects like those

Figure 1: Cuplike objects (From Labov 1973).

in Figure 1 in order to determine how the variations in shape influence the names the subjects use. But he also wanted to see whether the functions of the objects also influences naming. In the “neutral” context, subjects were asked to imagine the object in someone’s hand. In a “food” context, they were asked to imagine the object filled with mashed potatoes; and in a “flowers” context, they were told to imagine the object with cut flowers in it.

Figure 2 shows the results, when the *width* of objects 1 to 4 (Figure 2a) and *depth* of objects 1 and 5 to 9 (Figure 2b) varied as is represented on the horizontal axes. The vertical axis represents the percentage of subjects that named the object with a particular word. As can be seen, the names for the objects were heavily influenced by the imagined context.

This example shows that even if the “core” or “prototype” of two concepts like *cup* and *bowl* remain unchanged, the context may change the *border* between the concepts. Such a change may clearly have nonmonotonic effects on the use of the concepts. I have now presented four different kinds of nonmonotonic aspects of concepts. It should be obvious that these aspects are interrelated. However, in the literature, the focus has almost exclusively been on what happens when one changes from a general category to a subordinate. I submit that other aspects are equally important and in need of a systematic explanation. In the following two sections, I will outline a theory that I believe can shed light on all of the nonmonotonic aspects of concepts discussed here.

Figure 2: Context effects on the borders between concepts (Labov 1973).

2 Conceptual spaces as a representational framework

As a framework for a geometric structure used in describing a cognitive semantics, I have proposed (Grdenfors 1990a, 1990b, 1992, to appear) the notion of a *conceptual space*. A conceptual space consists of a number of *quality dimensions*. Examples of such dimensions are: color, pitch, temperature, weight, and the three ordinary spatial dimensions. I have chosen these dimensions because they are closely connected to what is produced by our sensory receptors (Schiffman 1982). However, there are also quality dimensions that are of an abstract non-sensory character.

The primary function of the quality dimensions is to represent various “qualities” of objects. They form the “framework” used to assign properties to objects and to specify relations between them. The dimensions are taken to be independent of symbolic representations in the sense that we and other animals can represent the qualities of objects without presuming an internal language in which these qualities are expressed. The quality dimensions should be seen as abstract representations used as a modeling factor in describing mental activities of organisms.

The notion of a dimension should be understood literally. It is assumed that each of the quality dimensions is endowed with certain *topological* or *metric* structures. As a first example, I will take the dimension of *time*. In science, time is a one-dimensional structure which is isomorphic to the line of real numbers. If “now” is seen as the zero point on the line, the future corresponds to the infinite positive real line and the past to the infinite negative line. This representation of time is not universal, but is to some extent culturally dependent, so that other cultures

have a different time dimension as a part of their cognitive structure. There is thus no unique way of choosing a dimension to represent a particular quality, but in general one has a wide array of possibilities.

In order to separate different uses of quality dimensions, it is important to introduce a distinction between a *psychological* and a *scientific* interpretation. The psychological interpretation concerns how humans (or other organisms) structure their perceptions. The scientific interpretation, on the other hand, deals with how different dimensions are presented within a scientific theory. The distinction is relevant when the dimensions are seen as cognitive entities, in which case their topological or metric structure should not be determined by scientific theories which attempt at giving a “realistic” description of the world, but by *psychophysical* measurements which determine the structure of how our perceptions are represented.

A psychologically interesting example of a set of quality dimensions concerns *color perception*. In brief, our cognitive representation of color can be described by three dimensions. The first dimension is *hue*, which is represented by the familiar color circle. The topological structure of this dimension is thus different from the quality dimensions representing time or weight which are isomorphic to the real line. One way of illustrating the differences in topology is by noting that we can talk about psychologically *complementary* colors, that is, colors that lie *opposite* to each other on the color circle. In contrast it is *not meaningful* to talk about two points of time or two weights being “opposite” to each other.

The second psychological dimension of color is *saturation* which ranges from grey (zero color intensity) to increasingly greater intensities. This dimension is isomorphic to an interval of the real line. The third dimension is *brightness* which varies from white to black and is thus a linear dimension with end points. Together, these three dimensions, one with circular structure and two with linear, constitute the color space which is a subspace of our perceptual conceptual space.

This space is often illustrated by the so-called *color spindle* (see figure 3). Brightness is shown on the vertical axis. Saturation is represented as the distance from the center of the spindle towards its perimeter. Hue, finally, is represented by the positions along the perimeter of the central circle.

It is impossible to provide a complete list of the quality dimensions involved in the conceptual spaces of humans. Some of the dimensions seem to be *innate* and to some extent hardwired in our nervous system, as for example color, pitch, and probably also ordinary space. Other dimensions are presumably *learned*. Learning new concepts often involves expanding one’s conceptual space with new quality dimensions. *Functional properties* used for describing artifacts may be an example here. Even if we do not know much about the topological structures of these dimensions, it is quite obvious that there is some such non-trivial structure (see e.g., Vaina’s (1983) analysis of functional representation). Still other dimensions may be *culturally* dependent. Finally, some quality dimensions are introduced by *science*.

The theory of conceptual spaces will now be used to provide a definition of a concept. I propose the following criterion (Grdenfors 1990b, 1992) where the

Figure 3: The color spindle.

topological characteristics of the quality dimensions are utilized to introduce a spatial structure for properties:

Criterion P: A natural concept is a convex region of a conceptual space.

The motivation for the criterion is that if some objects which are located at v_1 and v_2 in relation to some quality dimension (or several dimensions) are both examples of the concept C , then any object that is located between v_1 and v_2 on the quality dimension(s) will also be an example of C . Criterion P presumes that the notion of *betweenness* is meaningful for the relevant quality dimensions. This is, however, a rather weak assumption which demands very little of the underlying geometrical structure. (For a different proposal based on topological properties, see Mormann 1993).

Criterion P does not presume that one can identify sharp borders between concepts, but it can be applied also to *fuzzy* concepts or concepts that are defined by *probabilistic* criteria. What convexity requires is that if two object locations v_1 and v_2 both satisfy a certain membership criterion, e.g., has a certain degree or probability of membership, then all objects between v_1 and v_2 also satisfy the criterion.

Most properties expressed by simple words in natural languages seem to be natural properties in the sense specified here. For instance, I conjecture that all *color terms* in natural languages express natural properties with respect to the psychological representation of the three color dimensions. It is well-known that different languages carve up the color circle in different ways, but all carvings seems to be done in terms of convex sets (see Berlin and Kay 1969).

The predicates of a first order language correspond to several different grammatical categories in a natural language, most importantly those of adjectives, nouns and verbs. The main semantic difference between adjectives and nouns, on the one hand, is that an adjective normally refers to a *single domain* (which is a

single quality dimension or a few related dimensions as in the case of colors), while nouns normally contain information about *several domains* (for a similar idea, see Bk 1973). Verbs, on the other hand, are characterized by their temporal structure, i.e., they essentially involve the time dimension (see Langacker 1987 for conceptual representations of verbs). Using conceptual spaces, one can thus express the fundamental semantic differences between the most important grammatical categories. First order languages do not seem to be sufficiently rich to make these distinctions in a systematic manner.

As an example of a concept that is represented in several dimensions, consider *apple* (compare Smith et al. 1988). The first domains that we learn about when we encounter apples as children are presumably color, shape, texture and taste. Later, we learn about apples as (biological) fruits, about their nutritional value, and possibly about some further dimensions. Note that I don't require that a concept should be associated with a closed set of quality dimensions, but this set may be expanded as one learns about further aspects of a concept.

Taste space can presumably be represented by the four dimensions sweet, sour, salty and bitter (see Schiffman 1982). Other domains are trickier: it is difficult to say much about, e.g., the topological structure of "fruit space." For some ideas about how "shape space" should be modelled, see Grdenfors (1990b) and Marr and Nishihara (1978). However, let me represent the apple regions associated with each of these domains verbally as follows:

- | — | — _ — " —
- *Dimension* " *Region* —
- Color " Red-yellow-orange —
- Shape " Roundish—
- Texture " Smooth—
- Taste " Regions of the sweet and sour dimensions —
- Fruit " Specification of seed structure, flesh and peel type, etc. —
- Nutrition " Values of sugar content, vitamins, fibres, etc.—
- " —

In addition to the regions associated with each domain, the concept representation contains information about the *saliency* of the different domains. The saliency of different dimensions determine which *associations* can be made and thus which *inferences* can be triggered by a particular use of a concept (see the following section).

I am not assuming that the saliency is fixed, but the values can change with the context, and with the knowledge and interests of the user. For example, if you are eating an apple, its taste will be more salient than if you are using an apple as a ball when playing with an infant, which would make the shape dimension particularly salient.

Generalizing immediately from the representation for *apple*, I want to propose

the following definition of concept representation:

A concept is represented as a set of convex regions in a number of quality dimensions together with a salience assignment to the dimensions.

In this theory of concepts I have tried to bring in elements from other theories in psychology and linguistics. Some related ideas can be found in, among others, Barsalou (1992), Holmqvist (1993), Langacker (1987, pp. 154-166), and Smith et al. (1988). The main difference between these theories and the one presented here is that I put greater emphasis on the geometrical structure of the concept representations, in particular via the requirement of representing by convex regions of quality dimensions. As will be seen in the following section, these structures are crucial for the analysis of the nonmonotonic aspects of concepts.

3 Explaining the nonmonotonic aspects

In this section, the theory of conceptual spaces will be applied to outline explanations for the different kinds of nonmonotonic features of concepts.

(a) *Change from general category to subordinate*

A first observation is that describing properties as convex regions of conceptual spaces fits very well with the so called *prototype theory* of categorization developed by Rosch and her collaborators (Rosch 1975, 1978, Mervis and Rosch 1981, Lakoff 1987). The main idea of prototype theory is that, within a category of objects, certain members are judged to be more representative of the category than others. For example robins are judged to be more representative of the category *bird* than are ravens, penguins and emus; and desk chairs are more typical instances of the category *chair* than rocking chairs, deck-chairs, and beanbags. The most representative members of a category are called *prototypical* members. It is well-known that some properties, like *red* and *bald* have no sharp boundaries and for these it is perhaps not surprising that one finds prototypical effects. However, these effects have been found for most properties including those with comparatively clear boundaries like *bird* and *chair*.

If concepts are described as convex regions of a conceptual space, prototype effects are indeed to be expected. In a convex region one can describe positions as being more or less *central*. For example, if color concepts are identified with convex subsets of the color space, the central points of these regions would be the most prototypical examples of the color.

Subordinate concepts may move away from the prototypes of the general concept and thus result in atypical properties. Here the representation of concepts as convex regions in a conceptual space may be useful. If the first thing I ever hear about the individual Gonzo is that it is a bird, I will naturally locate it in the conceptual space as a more or less prototypical bird, i.e., at the center of the region representing birds. (The relevant conceptual space may be something like a many-dimensional hierarchical space of coordinates in the style of Marr and Nishi-

hara (1978)). And in that area of the conceptual space, birds do fly, i.e., almost all individuals located there also have the ability to fly. However, if I then learn that Gonzo is an emu, I must *revise* my earlier concept location and put Gonzo in the emu region, which is a subset of the bird region but presumably lies at the outskirts of that region. And in the emu region of the conceptual space almost none of the individuals fly.

This simple example only hints at how the *correlation* between different parts of a region representing a property and regions representing other properties can be used in understanding nonmonotonic reasoning. For this analysis, the spatial structure of properties is essential. Such correlations will only be formulated in an ad hoc manner if a propositional representation of information is used where the spatial structure cannot be utilized.

(b) *Contrast classes*

I have proposed that properties correspond to connected regions of a conceptual space. However, as was noticed in Section 1, a word like *red* has many uses that can result in nonmonotonic inferences:

- Red book
- Red wine
- Red hair
- Red skin
- Redwood

Red book accords with the dictionary definition, but the other uses don't fit with the standard color region assigned to red. How can we then explain that the same word is used in so many different contexts?

I don't see how this phenomenon can be analysed in a simple way using possible worlds or some other tools from logical semantics. However, the idea of a contrast class can quite easily be given a general interpretation with the aid of conceptual spaces. For each domain, for example skin color, we can map out the class of possible colors on the color spindle. This mapping will determine a subset of the full color space. The shape of this subset may be rather irregular. However, if the subset is embedded in a space with the *same dimensional structure* as the full space we obtain a picture that looks like Figure 4.

In this smaller spindle, the color words are then used in the same way as in the full space, even if the hues of the color in the smaller space don't match the hues of the complete space. Thus, *white* is used about the lightest forms of skin, even though white skin is pinkish, *black* refers to the darkest form of skin, even though black skin is brown, etc. The embeddings into smaller conceptual spaces will naturally result in nonmonotonic effects. For example, from the fact that x is

Figure 4: The subspace of skin colors embedded in the full color spindle.

a white person, one cannot conclude that x is a white object, even though *person* is subordinate to *object*. This analysis of contrast classes is presented in greater detail in Grdenfors (to appear).

(c) *Metaphors*

Metaphors have been notoriously difficult to handle within traditional semantic theories, let alone first order logic. In contrast, they are given key positions within cognitive semantics. Not only poetic metaphors but also everyday “dead” metaphors are seen as central semantic features and are given systematic analyses. One of the first works in this area was Lakoff and Johnson (1980).

In an earlier work (Grdenfors, to appear), I have given an analysis of metaphors within the theory of conceptual spaces. The core hypothesis is that *a metaphor expresses a similarity in topological or metrical structure between different quality dimensions*. A concept that corresponds to a particular structure in one quality dimension can be used as a metaphor to express a similar structure about another dimension. In this way one can account for how a metaphor can *transfer knowledge* about one conceptual dimension to another.

As a simple example, let us consider the expression “the peak of a career.” The literal meaning of peak refers to a structure in physical space, namely the vertically highest point in a horizontally extended (large) object, typically a mountain. This structure thus presumes two spatial dimensions, one horizontal and one vertical

(a) The peak of a mountain (b) The peak of a career

Figure 5: The literal and a metaphorical meaning of “peak.”

(see Figure 5a).

A career is an abstract entity without location in space. So how can a career have a *peak*? What happens when we metaphorically talk about the peak of a career is that the same geometrical structure is applied to a two-dimensional space which consists of the *time* dimension (of the career) and a dimension of *social status* (see Figure 5b). The latter dimension is normally conceived of as being vertical: we talk about somebody having a “higher” rank, “climbing” in the hierarchy, etc (see Lakoff and Johnson 1980).

It can now be seen that the role of different contrast classes for the same concept, as described above is closely related to that of metaphorical uses of a word. A metaphor expresses a similarity in topological or geometrical structure between *different* quality dimensions. Now, in the case of contrast classes, one set of dimensions is not really mapped onto a different set, but it is mapped onto a subspace of itself retaining the same topological structure.

(d) Context effects

The main effect of applying a concept in a particular context is that certain dimensions of the concept are put into focus by the context. In relation to the model given in the previous section, this means that the context determines the *relative saliency* of the dimensions. For example, in a context of moving furniture, the weight dimension becomes highly salient. Hence, the concept *piano* may lead to an inference of *heavy*. In contrast, in a context of musical instruments, the weight dimension is much less salient and an application of the concept *piano* will probably not become associated with *heavy* (Barclay et al. 1974).

Another effect of changes in context is that change in saliency of certain dimensions may result in a shift of the *borders* between different concepts. As was seen in Section 1, once the functionality was put in focus in Labov’s study of the

concept *cup*, the border between *cup* and *bowl* changed considerably. Such changes of borders naturally lead to nonmonotonic effects when the concepts are applied in different contexts.

As a matter of fact, also the nonmonotonic effects of contrast classes discussed above can be seen as a context effect. By introducing a concept like *skin* or *wine*, a context is set up in which the contrast class of the objects falling under the concept is focused. Hence it is quite natural that the application of, for example, color concepts become restricted to the contrast class.

4 Conclusion

The aim of this paper has been to present various nonmonotonic aspects of reasoning about concepts. I have outlined a theory of conceptual spaces as a tool for representing the relevant conceptual structures. The theory emphasizes *topological* and *geometrical* structures rather than logical and linguistic. The theory of conceptual spaces provides a fruitful framework for analysing various aspects of concept formation. In particular, it can be used to explain different kinds of nonmonotonic inferences. I have tried to show that the theory provides better and richer tools for this purpose than do inheritance networks or proposition-based nonmonotonic logics.

Acknowledgments

An earlier version of this paper has appeared in preprint form in *Logic for a Change*, Uppsala Prints and Preprints in Philosophy, Uppsala 1995. My research has been supported by the ESPRIT Project 6156, *Defeasible Reasoning and Uncertainty Management Systems II*. I wish to thank the participants of the belief revision workshop in Lund, September 29 - 30, 1995 and the participants of the workshop "Aspects of Mechanizing Inference", Naples, October 30 - November 2, 1995, for very helpful comments.

References

- J. R. Barclay, J. D. Bransford, J. J. Franks, N. S. McCarrell and K. Nitsch (1974): "Comprehension and semantic flexibility," *Journal of Verbal Learning and Verbal Behavior* 13, 471-481.
- L. W. Barsalou (1987): "The instability of graded structure: implications for the nature of concepts," in *Concepts and Conceptual Development*, U. Neisser, ed., Cambridge University Press, Cambridge, 101-140.
- L. W. Barsalou (1992): "Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols," in *Theories of*

Memory, A. F. Collins, S. E. Gathercole, M. A. Conway, and P. E. Morris, eds., Lawrence Erlbaum Associates, Hillsdale, 29-89.

B. Berlin and P. Kay (1969): *Basic Color Terms: Their Universality and Evolution*. University of California Press: Berkeley, CA.

L. Bk (1973): "Zum Problem der Konstitution der wissenschaftlichen Objektivität zur der Intersubjektivität der Lebenswelt," manuscript, Universität Konstanz.

S. Broström (1994): *The Role of Metaphor in Cognitive Semantics*, Lund University Cognitive Studies 31, Lund.

P. Gärdenfors (1988): *Knowledge in Flux*, MIT Press, Cambridge, MA.

P. Gärdenfors (1990a): "Induction, conceptual spaces and AI," *Philosophy of Science* 57, 78-95.

P. Gärdenfors (1990b): "Frameworks for properties: Possible worlds vs. conceptual spaces," in *Language, Knowledge and Intentionality*, L. Haaparanta, M. Kusch, and I. Niiniluoto, eds. (Acta Philosophica Fennica, 49), 383-407.

P. Gärdenfors (1992): "A geometric model of concept formation," in *Information Modelling and Knowledge Bases III*, S. Ohsuga et al., eds., IOS Press, Amsterdam, 1-16.

P. Gärdenfors (1994): "Three levels of inductive inference," in *Logic, Methodology, and Philosophy of Science IX*, D. Prawitz, B. Skyrms and D. Westersth, eds., Elsevier Science, Amsterdam, 427-449.

P. Gärdenfors (to appear): "Mental representation, conceptual spaces and metaphors", to appear in *Synthese*.

P. Gärdenfors and D. Makinson (1994): "Nonmonotonic inference based on expectations," *Artificial Intelligence* 65, 197-245.

K. Holmqvist (1993): *Implementing Cognitive Semantics*, Lund University Cognitive Studies 17, Lund.

W. Labov (1973): "The boundaries of words and their meanings," in *New Ways of Analyzing Variation in English*, J. Fishman, ed., Georgetown University Press, Washington, DC, 340-373.

G. Lakoff (1987): *Women, Fire, and Dangerous Things*, University of Chicago Press, Chicago, IL.

G. Lakoff and M. Johnson (1980): *Metaphors We Live By*, University of Chicago Press, Chicago, IL.

R. W. Langacker (1987): *Foundations of Cognitive Grammar*, Vol. 1, Stanford University Press, Stanford, CA.

D. Makinson and K. Schlechta (1991): "Floating conclusions and zombie paths: two deep difficulties in the 'directly skeptical' approach to defeasible inheritance nets," *Artificial Intelligence* 48, 199-209.

D. Marr and H. K. Nishihara (1978): "Representation and recognition of the spatial organization of three-dimensional shapes," *Proceedings of the Royal Society in London*, B 200, 269-294.

C. Mervis and E. Rosch (1981): "Categorization of natural objects," *Annual Review of Psychology* 32, 89-115.

T. Mormann (1993): "Natural predicates and the topological structure of conceptual spaces," *Synthese* 95, 219-240.

S. E. Palmer (1978): "Fundamental aspects of cognitive representation," in *Cognition and Categorization*, E. Rosch and B. B. Lloyd, eds., Lawrence Erlbaum Associates, Hillsdale, NJ, 259-303.

R. Reiter (1980): "A logic for default reasoning," *Artificial Intelligence* 13, 81-132.

E. Rosch (1975): "Cognitive representations of semantic categories," *Journal of Experimental Psychology: General* 104, pp. 192-233.

E. Rosch (1978): "Prototype classification and logical classification: The two systems," in *New Trends in Cognitive Representation: Challenges to Piaget's Theory*, E. Scholnik, ed., Lawrence Erlbaum Associates, Hillsdale, NJ, 73-86.

H. R. Schiffman (1982): *Sensation and Perception*, 2nd ed., John Wiley and Sons, New York, NY.

E. E. Smith, D. N. Osherson, L. J. Rips, and M. Keane (1988): "Combining prototypes: A selective modification model," *Cognitive Science* 12, 485- 527.

D. S. Touretsky (1986): *The Mathematics of Inheritance Systems*, Morgan Kaufmann, Los Altos, CA.

L. Vaina (1983): "From shapes and movements to objects and actions," *Synthese* 54, 3-36.