Induction and the Evolution of Conceptual Spaces

Peter Gärdenfors
Department of Philosophy
Lund University
Sweden

Abstract

Peirce notes (CP 2.753) that there are certain forms of evolutionary constraints that delimit the vast class of human inductive inferences. The paper focuses on the problem of projectible properties which is analyzed in relation to modern cognitive science. In particular, the notion of a conceptual space is introduced and it is shown how this can be used to identify projectible properties. Furthermore, I outline how the so called prototype theory of concept formation can be interpreted in terms of conceptual spaces. The evolutionary origin of conceptual spaces is discussed. It is argued that the emerging picture of which predicates can be used in inductive inferences can be seen as an elaboration of Peirce's views.

1. Peirce's constraint on inductive inference

Humans do not perform inductions in an arbitrary manner. At the beginning of his discussion of inductive inferences, Peirce notes that there are certain forms of *constraints* that delimit the vast class of possible inferences. As he puts it:

> "Nature is a far vaster and less clearly arranged repertory of facts than a census report; and if men had not come to it with special aptitudes for guessing right, it may well be doubted whether in the ten or twenty thousand years that they may have existed their greatest mind would have attained the amount of knowledge which is actually possessed by the lowest idiot. But, in point of fact, not man merely, but all animals derive by inheritance (presumably by natural selection) two classes of ideas which adapt them to their environment. In the first place, they all have from birth some notions, however crude and concrete, of force, matter, space, and time; and, in the next place, they have some notion of what sort of objects their fellow-beings are, and how they will act on given occasions. (CP 2.753)

Here, Peirce hints at an *evolutionary* explanation of why "the human intellect is peculiarly adapted to the comprehension of the laws and facts of nature" (CP 2.750). My aim in this paper is to elaborate some of his ideas about innate constraints for a theory of induction. I shall focus on the problem of *projectibility*, i.e. the problem of which predicates may be used in inductive inferences. My approach to this problem will be a combination of Peirce's evolutionary perspective and some ideas from recent cognitive science. In particular, I shall introduce the notion of *conceptual spaces* as a framework for selecting projectible properties and outline how this notion can throw light on the so called *prototype theory* of concept formation.

2. Logical positivism and the problems of induction

The most ambitious project of analyzing inductive inferences during this century has been that of the logical positivists. As a background to my analysis of Peirce's constraints on inductive inferences, I shall start by commenting briefly on this research program.

Inductive inferences were important for the logical positivists, since such inferences were necessary for their verificationist aims. The basic objects of study for them were sentences in some more or less regimented language. Ideally, the language was a version of first order logic where the atomic predicates represented observational properties. These observational predicates were taken as *primitive*, unanalysable notions. The main tool used when studying the linguistic expressions was logical analysis. In its most pure form, logical positivism allowed only this tool. A consequence of this methodology was that all observational predicates were treated in the same way since there were no *logical* reasons to differentiate between them. For example, Carnap (1950, Section 18B) requires that the primitive predicates of a language be logically independent of each other.

However, it became apparent that the methodology of the positivists led to serious problems in relation to the problem of induction. The most famous one's are Hempel's (1965) "paradox of confirmation" and Goodman's (1955) "riddle of induction". I will not repeat these well-known problems, but only state my diagnosis of what causes them (for a more detailed account, cf. Gärdenfors (1990a)). What I see as the root of the troublesome cases is that if we use *logical* relations alone to determine which inductions are valid, the fact that all predicates are treated on a par induces *symmetries* which are not preserved by our understanding of the inductions: "Raven" is treated on a par with "non-raven", "green" with "grue" etc. What we need is a *non-logical* way of distinguishing those predicates that may be used in inductive inferences from those that may not.

There are several suggestions for such a distinction in the literature. One idea is that some predicates denote "natural kinds" or "natural properties" while others don't, and it is only the former that may be used in inductions. Natural kinds are normally interpreted realistically, following the Aristotelian tradition, and thus assumed to represent something that exists in reality independently of human cognition. However, when it comes to inductive *inferences* it is not sufficient that the properties exist out there somewhere, but we need to be able to grasp the natural kinds by our minds. In other words, what is needed to understand induction, as performed by humans, is a *conceptualistic* or *cognitive* analysis of natural properties. It is one of the aims of the present paper to outline such an analysis.

Another notion that is relevant here is that of "similarity". Quine (1969) discusses this notion and its relation to that of a natural kind. He notes that similarity "is immediately definable in terms of kind; for things are similar when they are two of a kind" (p. 117). Furthermore, he says about similarity that we

> "cannot easily imagine a more familiar or fundamental notion than this, or a notion more ubiquitous in its application. On this score it is like the notions of logic: like identity, negation, alternation, and the rest. And yet, strangely, there is something logically repugnant about it. For we are baffled when we try to relate the general notion of similarity significantly to logical terms." (Quine 1969:117)

To substantiate this claim he discusses, and rejects, several attempts to define similarity or natural kinds in logical or set theoretical terms. Again, what is needed for the present project is not a realistic analysis of similarity (like Lewis's (1973) similarity between possible worlds), but a conceptualistic or cognitive. I will return to this kind of similarity in my discussion of prototype theory.

3. Conceptual spaces

As mentioned earlier, Peirce says that men and animals "have from birth some notions ... of force, matter, space, and time" (CP 2.753) I propose to call these notions *quality dimensions* (cf. Gärdenfors (1988, 1990a, 1990b)). Apart from those mentioned by Peirce one can mention other perceptual qualities like color, temperature, and weight. When we talk and think about the qualities of objects it is in relation to such dimensions. A *conceptual space* consists of a class of quality dimensions (cf. Quine's (1960) notion of "quality space").

The dimensions should be given a *conceptualistic* or *cognitive* interpretation. Some of the dimensions are closely related to what is produced by our sensory receptors, but there are also quality dimensions that are of a more abstract character. Furthermore, the dimensions are *prelinguistic* in the sense that we (and other animals) can think about the qualities of objects, for example when planning an action, without presuming a language in which these thoughts can be expressed.

The notion of a dimension should be taken seriously. Each of the quality dimensions is endowed with a certain *topological* or *metrical* structure. For example, we normally think of the time dimension as being isomorphic to the line of real numbers; and the weight dimension is isomorphic with the positive real numbers (there are no negative weights). Since the dimensions are cognitive entities, their topological or metrical 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. For example, as regards the hue of colors, the psychological or cognitive representation is given by the well known *color circle* which thus determines the basic topology of this quality dimensions, while a more 'scientific' approach would use the wavelength of the light as the underlying dimension.

Figure 1 about here

In general, representing information about the world in terms of quality dimensions involves using *geometrical* or *vectorial* notions. In contrast the research program of the logical positivists (and much of current AI) is based on *symbolic* or *logical* modes of representation (cf. Sloman's (1971) distinction between 'analogical' and 'Fregean' representation and P.S. Churchland's (1986) criticism of the 'sentential paradigm'). The move to a new form of representation is, I believe, necessary; the problems that plagued the logical positivists indicate that we have to go beyond logic and language to find a solution to the problem of projectibility in inductive reasoning. Quine notes that something like a conceptual space is needed to make *learning* possible:

> "Without some such prior spacing of qualities, we could never acquire a habit; all stimuli would be equally alike and equally different. These spacings of qualities, on the part of men and other animals, can be explored and mapped in the laboratory by experiments in conditioning and extinction. Needed as they are for all learning, these distinctive spacings cannot themselves all be learned; some must be innate." (Quine 1969:123)

This concludes my general presentation of conceptual spaces. I believe it can be seen as a generalization of the state space approach, advocated among others by P.M. Churchland (1986) and P.S. Churchland (1986), and of the vector function theories of Foss (1988). To some extent conceptual spaces also function like the domains in Langacker's (1986) semantic theory. The theory of conceptual spaces is a theory for representing information, and not primarily an empirical psychological or neurological theory. I believe the theory can be applied to a number of philosophical problems in epistemology and semantics. Here, my aim is to show its viability as a foundation for an analysis of projectible properties.

4. Natural properties

Once the notion of a conceptual space is available, it is quite easy to give a cognitivistic analysis of "natural properties" and "similarity" (for a more detailed presentation, cf. Gärdenfors (1990a)). The key idea is that a natural property can be identified with a *convex region* of a given conceptual space. Via the notion of 'convexity' this utilizes the topological properties of the quality dimensions. A convex region is characterized by the criterion that for every pair o_1 and o_2 of points in the region, all points between o_1 and o_2 are also in the region. This definition presumes that the notion of 'between' is meaningful for the relevant dimensions. This is, however, a rather weak assumption which demands very little of the underlying topological structure.

As an application of the criterion of a natural property, I conjecture that all color terms in natural languages express natural properties with respect to the psychophysical representation of the quality dimensions of colors. In other words, the conjecture predicts that if some object o_1 is described by the color term C in a given language and another object o_2 is also said to have the color C, then any object o_3 with a color which lies between the color of o_1 and the color of o_2 will also be described by the color term C. It is well-known that different languages carve up the color circle in different ways, but all carvings seem to be done in terms of convex sets. Strong support for this conjecture can be gained from Berlin and Kay (1969), although they do not treat color terms in general but concentrate on basic color terms.

Using the present criterion on natural properties, it is now possible to formulate a more precise version of a constraint on induction, which I believe is in line with the quotation from Peirce presented at the beginning:

(C) *Only properties corresponding to a convex region of the underlying conceptual space may be used in inductive inferences.*

It is only proposed that convexity is a necessary condition, but perhaps not sufficient, for a property to count as natural and thus projectible. Anyhow, the criterion is strong enough to avoid many Goodman-type problems. The motivation for why convexity is an appropriate condition will become clearer in the next section.

5. Prototype theory and similarity

Apart from giving intuitively plausible solutions to the old riddles of induction (cf. Gärdenfors (1990a)), the definition of natural properties in terms of convex regions of conceptual spaces derives independent support from the *prototype theory* of categorization developed by Rosch and her collaborators (Rosch 1975, 1978, Mervis and Rosch 1981, Lakoff 1987). This theory can also be used to provide a reasonable explication of the notion of similarity that is presupposed in inductive inferences. (Quine (1969:119-120) argues that 'natural kind' is definable in terms of 'similarity' and he actually proposes a precursor to prototype theory.)

The main idea of prototype theory is that within a category of objects, like those instantiating a property, 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 beanbag chairs. 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'.

Now if the traditional definition of a property is adopted it is very difficult to explain such prototype effects. Either an object is a member of the class assigned to a property or it is not and all members of the class have equal status as category members. Rosch's research has been aimed at showing asymmetries among category members and asymmetric structures within categories. Since the traditional definition of a property does not predict such asymmetries, something else must be going on.

In contrast, if natural properties are defined 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 properties are identified with convex subsets of the color space, the central points of these regions would be the most prototypical examples of the color. Rosch has, in a series of experiments, been able to demonstrate the psychological reality of such 'focal' colors.

For more complex categories like 'bird' it is perhaps more difficult to describe the underlying conceptual space. However, if something like Marr and Nishihara's

(1978) analysis of shapes is adopted, we can begin to see what such a space would look like. Their scheme for describing biological forms uses hierarchies of cylinder-like modeling primitives.

Figure 2 about here.

Each cylinder is described by two coordinates (length and width). Cylinders are combined by determining the angle between the dominating cylinder and the added one (two polar coordinates) and the position of the added cylinder in relation to the dominating one (two coordinates). The details of the representation are not important in the present context, but it is worth noting that on each level of the hierarchy an object is described by a comparatively small number of coordinates based on lengths and angles. Thus the object can be identified as a hierarchially structured vector in a (higher order) conceptual space.

It should be noted that even if even if different members of a category are judged to be more or less prototypical, it does not follow that some of the existing members must represent 'the prototype' (Rosch has been misunderstood on this point. Cf. Lakoff (1987), Ch. 2). If a category is viewed as a convex region of a conceptual space this is easily explained, since the central member of the region specifies a possible individual but need not be among the existing members of the category.

It is possible to argue in the converse direction too and show that if prototype theory is adopted, then the representation of properties as convex regions is to be expected. Assume that some quality dimensions of a conceptual space are given, for example the dimensions of color space, and that we want to partition it into a number of categories, for example color categories. If we start from a set of prototypes p_1, ..., p_n of the categories, for example the focal colors, then these should be the central points in the categories they represent. One way of using this information is to assume that for every point p in the space one can measure the distance from p to each of the p_i's. If we now stipulate that p belongs to the same category as the *closest* prototype p_i, it can be shown that this rule will generate a partitioning of the space that consists of convex areas (convexity is here defined in terms of the assumed distance measure). This is the so called Voronoi tessellation.

Figure 3 about here.

Thus, assuming that a *metric* is defined on the subspace that is subject to categorization, a set of prototypes will by this method generate a unique partitioning of the subspace into convex regions. Hence there is an intimate link between prototype theory and the analysis of this article where properties are defined as convex regions in a conceptual space. Furthermore, such a metric is an obvious candidate for a measure of *similarity* between different objects. In this way, the Voronoi tessellation gives an explicit answer to how a similarity measure determines a set of natural kinds.

It is interesting to note that Quine formulates a precursor to this kind of prototype theory:

> "One may be tempted to picture a kind, suitable to a comparative similarity relation, as any set which is 'qualitatively spherical' in this sense: it takes in exactly the things that differ less than so-and-so much from some central norm. If without serious loss of accuracy we can assume that there are one or more actual things (*paradigm cases*) that nicely exemplify the desired norm, and one or more actual things (*foils*) that deviate just barely too much too be counted into the desired kind at all, then our definition is easy: *the kind with paradigm a and foil b* is the set of all things to which *a* is more similar than *a* is to *b*. More generally, then, a set may be said to be a *kind* if and only if there are *a* and *b*, known or unknown, such that the set is the kind with paradigm *a* and foil *b*. (Quine, 1969,:119-120)

Quine notes that, as it stands, this definition of a kind is not satisfactory:

> "Thus take red. Let us grant that a central shade of red can be picked as norm. The trouble is that the paradigm cases, objects in just that shade of red, can come in all sorts of shapes, weights, sizes, and smells. Mere degree of overall similarity to any such paradigm case will afford little evidence of degree of redness, since it will depend also on shape, weight, and the rest. If our assumed relation of comparative similarity were just comparative chromatic similarity, then our paradigm-and-foil definition of kind would indeed accommodate redkind. What the definition will not do is distill purely chromatic kinds from mixed similarity." (Quine, 1969:120)

The problem for Quine is that he does not assume anything like a conceptual space to help structuring the relations of similarity. However, if such a structure is given, we need not rely on actual objects as paradigm cases, but can use focal points on a particular quality dimension, like the color dimension, as a basis for comparing chromatic similarity. The shape, weight, and the rest of the qualities of objects will simply not be relevant for such comparisons.

6. The origin of quality dimensions

How do we know that the inductions generated from the natural properties determined by a conceptual space will be successful? In order to answer this question we must first consider the *origins* of the quality dimensions we use in classifying kinds and judging similarities.

There does not seem to be a unique origin of our quality dimensions. Some of the dimensions are presumably *innate* and to some extent hardwired in our nervous system, as for example color, pitch, and probably also ordinary space. These subspaces are obviously extremely important for basic activities like finding food and getting around in the environment.

Other dimensions are probably *learned*. Learning new concepts often involves expanding one's conceptual space with new quality dimensions. 'Volume' is an example here. According to Piaget's 'conservation' experiments with five year olds, small children do no make a distinction between the height of a liquid and its volume. The conservation of volume, which is part of its conceptual structure, is something that must be learned. Another example is provided by the functional properties which are used for describing artifacts. Such properties are characterized by a set of associated actions. (See e.g. Vaina's (1983) analysis of functional representation). Even if we do not know very much about the conceptual dimensions underlying actions, it is quite obvious that there is some non-trivial such structure.

Still other dimensions may be *culturally dependent*. 'Time' is a good example: in contrast to our linear conception of time, some cultures conceive of time as circular so that the world keeps returning to the same point in time, and in other cultures it is hardly meaningful at all to speak of time as a separate dimension. A sophisticated time dimension is needed for advanced forms of planning and coordination with other individuals, but is not necessary for the most basic activities of an organism. Although Peirce, in the quotation given earlier, counts some concepts related to our fellow-beings to be among the innate ideas, I believe that many social dimensions are culturally dependent. For example, kinship relations, which are fundamental to social concepts, are specified differently in different cultures.

Finally, some quality dimensions are introduced by *science*. As an example, let me mention the distinction between temperature and heat, which is central for thermodynamics, but which has no correspondence in human perception. (Human perception of heat is basically determined by the amount of heat transferred from an object to the skin rather than by the temperature of the object).

7. The relativism of inductive inferences

Natural properties have here been defined in relation to a given conceptual space. But isn't the choice of a conceptual space *arbitrary*? Since a conceptual space may seem like a Kuhnian paradigm, aren't we thereby stuck with an unescapable *relativism*? -- Anyone can pick his own conceptual space and in this way make his favorite properties come out natural in that space. For example, it is not difficult to construct a conceptual space where 'grue' would correspond to a convex region and thus be a natural property in that space.

I want to argue that freedom in choosing a conceptual space is rather limited and thus that the relativism inherent in this choice in not as problematic as it may first seem. As was noted in the previous section, many quality dimensions are innate. This is the main reason why human beings, to a remarkable extent, agree on which properties are projectible, in particular if the properties are closely connected to what is provided by our senses. For instance, different cultures show high agreement in identifying species of animals and plants, at least what concerns basic categories.

But even if such an agreement exists, we must answer the question of why our way of identifying natural properties accords so well with the external world as to make our inductions tend to come out right. Peirce notes that "it can no longer be denied that the human intellect is peculiarly adapted to the comprehension of the laws and facts of nature" (CP 2.750). Quine (1969, p. 126) formulates the problem in the following way:

"... why does our innate subjective spacing of qualities accord so well with the functionally relevant groupings in nature as to make our inductions tend to come out right? Why should our subjective spacing of qualities have a special purchase on nature and a lien on the future?"

The answer, it seems to me, comes from evolutionary theory. As is indicated by the earlier quotation, Peirce (CP 2.753) seems to take the same line. Natural selection has made us all develop a conceptual space that results in inductions that are valid most of the time and thus promote survival. It is not that "there is a general tendency toward uniformity in the universe" (CP 2.749), but our quality dimensions are what they are because they have been selected to fit the surrounding world. In Peirce's words: "... there is a special adaptation of the mind to the universe, so that we are more apt to make true theories than we otherwise should be" (CP 2.749). Now, as indicated in the previous section, not all of our quality dimensions have a genetic origin, but the agreement between the conceptual spaces of different individuals of our species is still high enough to produce very similar conceptions of, at least, the perceptual categories of the world. The upshot is that that there will exist rather limited freedom for humans in choosing a conceptual space for the most fundamental inductive inferences. The relativism that will occur only applies to the more advanced learned and culturally dependent quality dimensions.

However, a far-reaching consequence of this evolutionary account of conceptual spaces is that *our inductive capacities will be dependent on the ecological circumstances under which they have evolved*. Human conceptual spaces are attuned to the environment of thousands of years of hunting and gathering. Consequently, we cannot expect our intuitions about which properties are projectible to be successful in environments that wildly diverge from those present during our evolutionary history.

8. Conclusion: The role of science

The conclusion to be drawn is that we should only trust our capacities for inductive reasoning in situations which are "ecologically valid", to borrow a concept from J.J. Gibson. What should we do about the increasing proportion of artificial objects and situations that surround us?

This is where *science* enters on the scene. By introducing theoretically precise, non-psychological quality dimensions a scientific theory may help us find new inductive inferences that would not be possible on the basis of our subjective conceptual spaces alone. A scientific break-through is often made when a new quality dimension is introduced to a theory. Witness, for example, the distinction between temperature and heat that was mentioned earlier. This distinction was of crucial importance for the development of thermodynamics. Or witness Newton's distinction between weight and mass. This is the very first definition introduced in the *Principia*; without it his celestial mechanics would have been impossible.

The quality dimensions of scientific theories, and their associated measurement procedures, help us in producing new successful inductive conclusions in environments which are completely different from that of our evolutionary cradle. Furthermore, the precise metrics of the scientific quality dimensions make it possible to formulate functional laws which in turn enable us to compute more sophisticated predictions. Once a scientific conceptual space has been established, the formulation of such laws is a kind of inductive inference. Clearly, such laws supersede inductions of the form "All Fs are Gs" which are the only possible inductive inferences on quality dimensions that have only a crude or rudimentary metric. As Quine (1969:137-138) notes, it is a sign of mature science that notions of similarity become less and less important, being replaced by theoretically more sophisticated concepts.

In this way science builds upon our more or less evolutionarily determined conceptual spaces, but in its most mature form becomes independent of them. It seems appropriate to conclude with the following quotation from Peirce which supports this general thesis:

> "Side by side, then, with the well established proposition that all knowledge is based on experience, and that science is only advanced by the experimental verifications of theories, we have to place this other equally important truth, that all human knowledge, up to the highest flights of science, is but the development of our inborn animal instincts." (CP 2. 754)

Acknowledgements

Research for this paper has been supported by the Swedish Council for Research in the Humanities and Social Sciences and by the Erik Philip-Sörensen Foundation. The article draws heavily on material in Gärdenfors (1990a) and some of the ideas on prototype theory are borrowed from Gärdenfors (1990b).

References

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

Carnap, Rudolph (1950). *Logical Foundations of Probability*. Chicago University Press: Chicago, IL.

Churchland, Patricia S. (1986). *Neurophilosophy: Toward a Unified Science of the Mind/Brain*. Bradford Books, MIT Press: Cambridge, MA.

Churchland, Paul M. (1986) Some reductive strategies in cognitive neurobiology, *Mind 95*, no. 379, 279-309.

Foss, Jeff (1988). The percept and vector function theories of the brain, *Philosophy of Science 55*, 511-537.

Gärdenfors, Peter (1988). Semantics, conceptual spaces and the dimensions of music, in *Essays on the Philosophy of Music*, ed. V. Rantala, L. Rowell, and E. Tarasti, (*Acta Philosophica Fennica*, vol. 43), Helsinki, 9-27.

Gärdenfors, Peter (1990a). Induction, conceptual spaces and AI, to appear in *Philosophy of Science*.

Gärdenfors, Peter (1990b). Mental representation, conceptual spaces and metaphors, to appear in *Synthese*.

Hempel, Carl G. (1965). *Aspects of Scientific Explanation, and Other Essays in the Philosophy of Science*. Free Press, New York, NY.

Goodman, Nelson (1955). *Fact, Fiction, and Forecast*. Harvard University Press, Cambridge, MA.

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

Langacker, Ronald W. (1986) *Foundations of Cognitive Grammar (Vol. 1)*. Stanford University Press, Stanford, CA.

Lewis, David K. (1973). *Counterfactuals*. Blackwell: Oxford.

Marr, David 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*, pp. 269-294.

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

Peirce, Charles Saunders (1932). *Collected Papers of Charles S. Peirce*, ed. Hartshorne, Weiss and Burks, Harvard University Press: Cambridge, MA.

Quine, Willard V.O. (1960). *Word and Object*. MIT Press: Cambridge, MA.

Quine, Willard V.O. (1969). Natural kinds, in *Ontological Relativity and Other Essays*, Columbia University Press: New York, NY, 114-138.

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

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

Sloman, Aron (1971). Interactions between philosophy and AI -- the role of intuition and non-logical reasoning in intelligence. *Proceedings 2nd IJCAI*, London, reprinted in *Artificial Intelligence*, vol. 2.

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

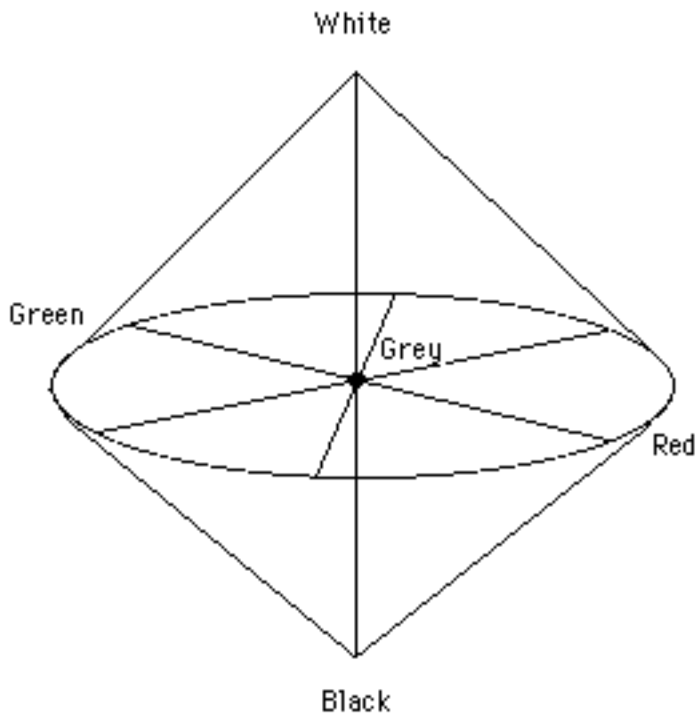

Figure 2.
The full color space

Figure 1. The three-dimensional color space

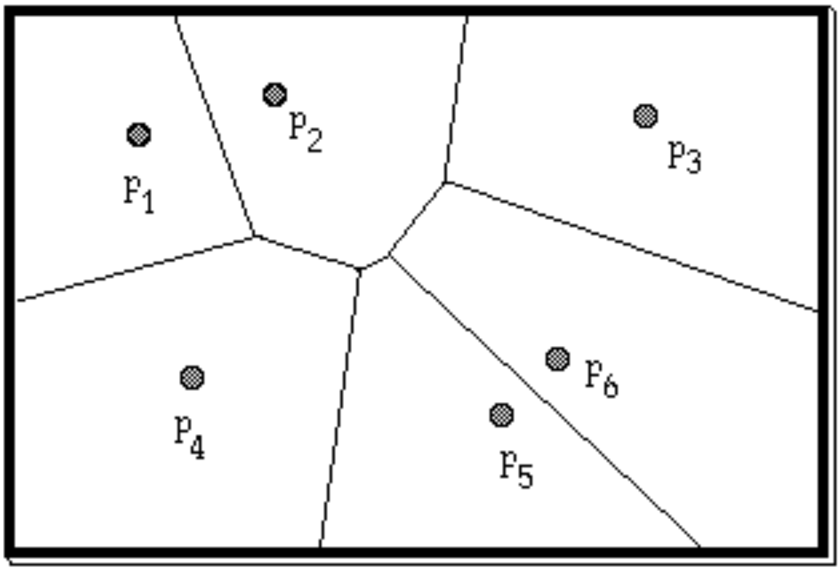

Figure 3. Voronoi tessellation of the plane into convex sets.