Reasoning about Categories in Conceptual Spaces *

Peter Gärdenfors
Lund University Cognitive Science
Lund University
S-222 22 Lund, Sweden
peter.gardenfors@lucs.lu.se

Mary-Anne Williams
Business and Technology Research Lab
University of Newcastle
NSW 2308, Australia
maryanne@cafe.newcastle.edu.au

Abstract

Understanding the process of categorization is a primary research goal in artificial intelligence. The conceptual space framework provides a flexible approach to modeling context-sensitive categorization via a geometrical representation designed for modeling and managing concepts.

In this paper we show how algorithms developed in computational geometry, and the Region Connection Calculus can be used to model important aspects of categorization in conceptual spaces. In particular, we demonstrate the feasibility of using existing geometric algorithms to build and manage categories in conceptual spaces, and we show how the Region Connection Calculus can be used to reason about categories and other conceptual regions.

1 Introduction

Categorization is a fundamental cognitive activity. The ability to classify and identify objects with a high degree of exception tolerance is a hallmark of intelligence, and an essential skill for learning and communication. Understanding the processes involved in constructing categories is a primary research goal in artificial intelligence.

The conceptual space framework as developed by Gärdenfors [2000] provides a flexible approach to modeling context-sensitive categorization. Conceptual spaces are based on a simple, yet powerful, geometrical representation designed for modeling and managing concepts.

In this paper we show how algorithms developed in computational geometry, and the Region Connection Calculus (RCC) [Cohn *et al.*, 1997], a well known region-based spatial reasoning framework, can be used to model important aspects of categorization in conceptual spaces. In particular, we demonstrate the feasibility of using existing geometric algorithms to build and manage categories in conceptual spaces, and we show how the RCC can be used to reason about categories and other conceptual regions.

*This paper appeared in the *Proceedings of the Fourteenth International Joint Conference of Artificial Intelligence*, Morgan Kaufmann, 385 - 392, 2001.

2 Conceptual Spaces

Conceptual spaces provide a framework for modeling the formation and the evolution of concepts. They can be used to explain psychological phenomena, and to design intelligent agents [Chella *et al.*, 1998; Gärdenfors, 2000]. For the purposes of this paper conceptual spaces provide the necessary infrastructure for modeling the process of categorization.

Conceptual spaces are geometrical structures based on quality dimensions. Quality dimensions correspond to the ways in which stimuli are judged to be similar or different. Judgments of similarity and difference typically generate an ordering relation of stimuli, e.g. judgments of pitch generate a natural ordering from "low" to "high" [Gärdenfors, 2000]. There have been extensive studies conducted over the years that have explored psychological similarity judgments by exposing human subjects to various physical stimuli. Multi-dimensional scaling is a standard technique that can be used to transform similarity judgments into a conceptual space [Krusal and Wish, 1978]. An interesting line of inquiry is pursued by Balkenius [1999] who attempts to explain how quality dimensions in conceptual spaces could accrue from psychobiological activity in the brain.

In conceptual spaces objects are characterized by a set of attributes or qualities $\{\mathbf{q}_1, \mathbf{q}_2, ..., \mathbf{q}_n\}$. Each quality $\mathbf{q}_i$ takes values in a domain $\mathbf{Q}_i$. For example, the quality of pitch (or frequency) for musical tones could take values in the domain of positive real numbers. Objects are identified with points in the conceptual space $\mathbf{C} = \mathbf{Q}_1$ x $\mathbf{Q}_2$ x $\ldots \mathbf{Q}_n$, and concepts are regions in conceptual space.

In the definition above we use the standard mathematical interpretation of "domain". In [Gärdenfors, 2000] however, a domain is defined to be a set of *integral dimensions*, this interpretation is consistent with its use in the psychology literature. For example, pitch and volume constitute the integral dimensions of sounds discernible by the human auditory perception system. Integral dimensions are such that they cannot be separated in the perceptual sense. The ability to bundle up integral dimensions as a domain is an important part of the conceptual spaces framework. Domains facilitate the sharing and inheritance of integral dimensions across conceptual spaces.

For the purpose of this paper, and without loss of generality, we often identify a conceptual space $\mathbf{C}$ with $\mathbf{R}^n$, but hasten to note that conceptual spaces do not require the full rich-

ness of $\mathbf{R}^n$. Domains can be continuous or discrete[1]. They can also be based on a wide range of geometrical structures, for example, according to psychological evidence the human colour perception system is best represented using polar coordinates[2] [Gärdenfors, 2000].

For the purpose of problem solving, learning and communication, agents adopt a range of conceptualizations using different conceptual spaces depending on the cognitive task at hand.

Similarity relations are fundamental to conceptual spaces. They capture information about the similarity judgments. In order to model some similarity relations we can endow a conceptual space with a distance measure.

Definition 1 *A* distance measure d *is a function from* **C** x **C** *into* **T** *where* **C** *is a conceptual space and* **T** *is a totally ordered set.*

Distance measures lead to a natural model of similarity; the smaller the distance between two objects in conceptual space, the more similar they are. The relationship between distance and similarity need not be linear, e.g. similarity may decay exponentially with distance.

The properties of connectedness, star-shapedness and convexity of regions in conceptual spaces will prove useful throughout.

Definition 2 *A subset C of a conceptual space is:*

(i) connected *if for every decomposition into the sum of two nonempty sets $C = C_1 \cup C_2$, we have $\bar{C}_1 \cap C_2 \cup C_1 \cap \bar{C}_2 \neq \emptyset$ where $\bar{C}$ is the closure of C. In other words, C is connected if it is not the disjoint union of two non-empty closed sets.*

(ii) star-shaped with respect to a point **p** *(referred to as a* kernel point*) if, for all points* **x** *in C, all points between* **x** *and* **p** *are also in C.*

(iii) convex *if, for all points* **x** *and* **y** *in C, all points between* **x** *and* **y** *are also in C.*

Definition 3 *The* kernel of a star-shaped region C *is the set of all possible kernel points, and will be denoted* $kernel(C)$.

Connectedness is a topological notion, whilst star-shapedness and convexity rely only on a betweenness relation. A qualitative betweenness relation can be specified in terms of a similarity relation, $S(\mathbf{a}, \mathbf{b}, \mathbf{c})$, which says that **a** is more similar **b** than it is to **c**. Alternatively, a betweenness relation can be used as primitive, and axioms introduced to govern its behaviour [Borsuk and Szmielew, 1960]. In the special case where the distance measure is a metric, the betweenness relation can be defined as: "**b** is between **a** and **c**" if and only if $d(\mathbf{a}, \mathbf{b}) + d(\mathbf{b}, \mathbf{c}) = d(\mathbf{a}, \mathbf{c})$.

Convex regions are star-shaped, and in many topological settings star-shaped regions are connected. The kernel of a convex region is the region itself, and under the Euclidean metric kernels are convex.

[1]They can even be small and finite e.g. {male, female}.

[2]A scientific representation of colour would require a different representation however, one that captures important scientific features of the electromagnetic spectrum such that the wave properties of wavelength and amplitude constitute integral dimensions.

Constraints like connectedness, star-shapedness and convexity can be used to impose ontological structure on the categorization of the conceptual space, i.e. not any old region can serve as a category. In fact, there is compelling evidence that *natural properties* correspond to convex regions in conceptual space, and using the idea of a natural property in this way Gärdenfors [2000] is able to sidestep the enigmatic problems associated with induction.

In section 4 we show how categorization, the central theme of this paper, occurs in conceptual spaces, but first we briefly describe the RCC.

3 Region Connection Calculus

The RCC is a qualitative approach to spatial reasoning. It was developed in an attempt to build a commonsense reasoning model for space, and its remarkable utility has been illustrated in numerous innovative applications [Cohn *et al.*, 1997].

The RCC approach is region-based where spatial regions are identified with their closures. The RCC is based on a connection relation, $C(X, Y)$, which stands for "region X connects with region Y". The connection relation, C, is reflexive and symmetric. Despite the fact that the basic building blocks in the RCC are regions, C can be given a topological interpretation, namely $C(X, Y)$ holds when the topological closures of regions X and Y share at least one point.

The RCC framework comprises several families of binary topological relations. One family, the RCC5 fragment uses the following Jointly Exhaustive and Pairwise Disjoint base relations to describe the relationship between two regions (see Figure 1); DR (discrete), EQ (identical), PP (proper part), PP^{-1}(inverse PP), and PO (partial overlap).

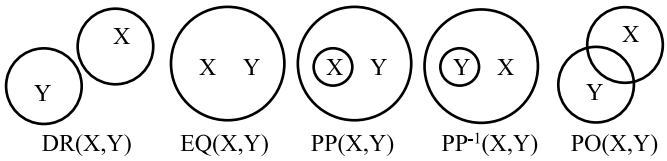

Figure 1: *The base relations in RCC5.*

Boundaries of regions are not distinguished in RCC5; there is no difference between two regions being totally disconnected and externally connected, and no difference between a proper part tangentially connected to the boundary and a proper part disconnected from the boundary. Another fragment, RCC8, possesses base relations that can make these distinctions.

In this paper our interests lie in similarity based categorization, and we use RCC5 to illustrate how the RCC can be used to represent conceptual regions. Extending to a more expressive mereotopological language to reason about the adjacency of categories, for example, is straightforward, and it can be done at no extra computational cost.

Transition tables can be used to perform reasoning about relationships between regions in RCC5 and RCC8. It is known that there is no complete first order finite axiomatization of topology. Nebel [1995] showed that propositional reasoning in RCC5 and RCC8 is NP-Hard, and Renz and Nebel

[1999] identified a maximal tractable subset of relations that contain all the base relations. Efficient implementations in Prolog have been constructed using a zero-order intuitionistic logic [Bennett, 1997].

The RCC can be used to represent spatial regions and to reason about them in any dimension, provided that all spatial entities possess precisely the same number of dimensions. It can also represent regions composed of multiple (spatially distinct) parts, but it cannot represent null regions.

In this paper we use the RCC to reason about spatial relationships between conceptual regions, and to model the indeterminacy of conceptual regions. The original formulation of the RCC concerned regions with crisp well-defined boundaries, but later Cohn and Gotts [1996] extended the RCC to handle indeterminate vague regions that could be "crisped" into less vague regions. We describe this crisping relation in more detail in section 5.2.

4 Categorization in Conceptual Spaces

4.1 Prototypes and Categorical Perception

A categorization results in a partitioning of a conceptual space into (meaningful) subregions. The geometrical nature of conceptual spaces coupled with representations for prototypes and the ability to manipulate dimensions independently of one another ensures that they provide a highly flexible and practical representation of context-sensitive categorization.

Within each category certain members are judged to be more representative than others [Rosch, 1975]. The most representative members of a category are *prototypes*. There is a wealth of psychological data supporting the existence of prototypes and their key role in categorization. Typically human performance experiments are used to determine how well, and how quickly humans can classify, label, rank or compare objects. Experimental results consistently show that the ease of classification varies with how similar an object is to a prototype. Furthermore, the more similar a nonmember is to the prototype, the more difficult it is to exclude. It has been shown, for example, that the reaction time recognizing that a robin is a bird is shorter than identifying a penguin as a bird regardless of whether the stimulus is the name or an image.

It is evident from experiments that humans make judgments about the degree of resemblance to a prototype during classification and identification; a robin is judged as a more prototypical bird than a penguin [Harnad, 1987]. A prototype consists of features of either a typical, or ideal category member, rather than invariant features common to every member. For example, a prototype of a category can be thought of as an amalgam of the characteristic attributes of its category's exemplars[3].

Classifying an object using prototypes is accomplished by determining its similarity to a prototype. Instances above some *threshold* of similarity to the prototype are taken as category members, all other instances are nonmembers.

Prototypes are central to the representation and processing of categories. People classify, generate, acquire, and reason about typical exemplars faster and more accurately than atypical exemplars. They also produce stronger inductive inferences with typical exemplars than with atypical exemplars [Hampton, 1993].

It is widely accepted on the basis of empirical psychological experiments that people sometimes judge membership of categories as graded. The existence of graded concepts supports the notion of *continuous perception*. The gradedness of category membership can be used to determine how closely an object resembles a prototype and can naturally be determined from the underlying similarity judgments.

Psychological evidence also suggests that people distinguish stimuli along a physical continuum much better when the stimuli are from different categories than when they are from the same category. This phenomenon is called *categorical perception* [Harnad, 1987], and is manifested in the ability to discriminate stimuli with more ease and accuracy between categories than within them. When categorical perception is at work, stimuli related to a specific category are perceived as indistinguishable, whereas stimuli from a "nearby"category are perceived to be entirely different. This phenomenon has been found in the way humans process sounds in speech. In color perception, for example, different shades of green are perceived to be more similar than green and yellow even though the wavelength differences are no larger. In other words, the psychophysical relationship between the physical intensity of a stimulus and the psychological intensity of the ensuing sensation is related to the categorization. Categorical perception has also be found in primates, other than humans [May *et al.*, 1989].

It is also well known that similarity judgments crucially depend on the context in which they occur for both continuous and categorical perception. It turns that certain features of objects and concepts are more salient for a particular categorization (for both classification and identification) depending on the context. The classic example is that a robin is a prototypical bird, but a canary is a prototypical pet bird.

In summary the key findings from psychological studies of categorization are (i) similarity judgments play a fundamental role in categorization and they are context sensitive, (ii) the degree of similarity is judged with respect to a reference object/region such as a prototype, (iii) category membership can be graded (discrete membership, if and when it exists, is considered to be a special case), and (iv) the psychophysical relationship between the stimulus and the response depends on the underlying categorization.

4.2 A Mechanism for Building Categories

In this section results in computational geometry are applied to categorization in conceptual spaces. We provide computational evidence for categorization based on prototypes, rather than appealing to the usual intuitive arguments found in the psychology literature derived from facts like the presentation of prototypes enhances learning. It can be shown, for example, that the conceptual space model predicts that it is easier to learn categories in which the natural prototype is central to a set of variations than it is to learn categories in which the prototype occurs as a peripheral member, as observed by Rosch [1975].

[3]Exemplars are previously perceived examples of objects in a category.

The main idea is that Voronoi tessellations around prototypes can be used to determine the threshold of similarity that forms category boundaries. In other words, the prototypes and the underlying similarity relation can be used to tessellate a conceptual space into categories.

Definition 4 *A Voronoi tessellation in conceptual space is given by the triple* $\Delta(\mathbf{P}, d, \mathbf{C})$ *where* $\mathbf{P}$ *is a set of distinct generator points* $\{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m\}$, d *is a distance measure on* $\mathbf{C}$ *a conceptual space. We define the tessellated regions* $c(\mathbf{p}_i)$ *to be* $\{\mathbf{x} | d(\mathbf{p}_i, \mathbf{x}) \leq d(\mathbf{p}_j, \mathbf{x})$ *for* $j = 1, 2, \ldots, m\}$, *and we call* $c(\mathbf{p}_i)$ *the category generated by* $\mathbf{p}_i$.

A Voronoi tessellation divides a conceptual space according to the *nearest-neighbor rule* which says each point/object in the space is associated with the prototype closest to it. This results in prototypes being centrally located in their category.

The Euclidean metric is a basic distance measure: the Euclidean distance $d(\mathbf{x}, \mathbf{p})$ between two points $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ and $\mathbf{p} = (p_1, p_2, \ldots, p_n)$ in $\mathbf{R}^n$ is calculated as $\sqrt{\sum_{j=1}^{n}(x_j - p_j)^2}$. If the underlying distance measure is the Euclidean metric, then the resultant categories $c(\mathbf{p}_i)$, are convex, and hence star-shaped with respect to $\mathbf{p}_i$. If the distance measure is the Manhattan or the supremum metric, then the generated categories are not necessarily convex, but are star-shaped with respect to $\mathbf{p}_i$. In fact, it can be shown that if $\mathbf{x}$ is between $\mathbf{y}$ and $\mathbf{p}_i$ is defined as $d(\mathbf{p}_i, \mathbf{x}) + d(\mathbf{x}, \mathbf{y}) = d(\mathbf{p}_i, \mathbf{y})$ and the distance measure satisfies the triangle inequality, then the generated categories are star-shaped with respect to $\mathbf{p}_i$.

Star-shapedness with respect to a prototype $\mathbf{p}$ is a desirable property for categories: if a category $c(\mathbf{p})$ is not star-shaped with respect to $\mathbf{p}$, then there is an object $\mathbf{x}$ that is between $\mathbf{p}$ and some $\mathbf{y} \in c(\mathbf{p})$ but $\mathbf{x} \notin c(\mathbf{p})$.

Definition 5 *A* well behaved categorization *in conceptual space produces regions which are star-shaped with respect to their prototype region and contain their central prototype.*

A Voronoi tessellation encapsulates the entire proximity information about the set of prototypes in a computationally compact fashion. Voronoi diagrams in the plane can be computed in $O(n\ log\ n)$ worst-case optimal time using $O(n)$ space [Okabe $et\ al$, 2000], and in d-dimensions for $d > 3$ in $O(n^{\lceil d/2 \rceil})$ worst-case optimal time [Klee, 1980].

Once constructed Voronoi tessellations can be used to: (i) identify the category of arbitrary objects in logarithmic query time without increasing the storage space - this is asymptotically optimal since it matches the information theoretical lower bound [Auberhammer, 1991], and (ii) compute the smallest enclosing sphere containing n prototype points in $O(n\ log\ n)$ worst-case optimal time [Auberhammer, 1987]. Furthermore, a prototype can be added or deleted to a Voronoi tessellation in $O(n)$ time, and two Voronoi tessellations can be merged in $O(n)$ time.

It is interesting to note that classical techniques and algorithms for information retrieval and cluster analysis are related to Voronoi tessellations. In fact, Voronoi tessellations have been used to construct robust approximate solutions to well known NP-complete information retrieval problems, i.e. acceptable approximate solutions can be found in $O(n\ log\ n)$ time [Okabe $et\ al$, 2000].

Much experimental psychological data concords with the idea of tessellating conceptual spaces into star-shaped (and sometimes convex) regions around prototypes or exemplars, e.g. stop consonants in phoneme classification [Petitot, 1989], other examples can be found in [Gärdenfors, 2000].

Not only do Voronoi tessellations generated by prototypes support the prototype model of categorization, but the generated boundaries provide a threshold of similarity and support a mechanism which can explain categorical perception. The precise mechanism involves crisping the distance measure and is described in Section 5.2.

4.3 Generalized Voronoi Tessellations

In this section we discuss some useful extensions of the basic Voronoi tessellation model. Ordinary Voronoi tessellations give rise to ideal categorizations, in the sense that crisp boundaries are generated from a single prototype. We generalize the definition so as to generate categories from conceptual regions rather than specific prototype points.

Definition 6 *A generalized Voronoi tessellation is given by the triple* $\Delta(\mathbf{P}, d, \mathbf{C})$ *where* $\mathbf{P}$ *is a set of generator regions* $\{P_1, P_2, \ldots, P_m\}$, d *is a distance measure on* $\mathbf{C}$ *a conceptual space. We define the tessellated regions* $c(P_i)$ *to be* $\{\mathbf{x} | d(P_i, \mathbf{x}) \leq d(P_j, \mathbf{x})$ *for* $j = 1, 2, \ldots, m\}$, *and we call* $c(P_i)$ *the category generated by the region* P_i.

We contrast two distance measures for generalized Voronoi categorizations; the *additively weighted distance* and the *power distance*[4]. The additively weighted Voronoi diagram is typically used to model the growth of biological cells, and can be used to model the growth of concepts also. The power distance, on the other hand, is best suited to handle indeterminacy and exemplar variability.

An additively weighted distance between a point $\mathbf{x}$ and a sphere $P \in \mathbf{P}$ in $\mathbf{R}^n$ with weight $w(P)$, denoted $d(\mathbf{x}, P)$, is defined as $d(\mathbf{x}, \mathbf{p}) - w(P))$ where $d(\mathbf{x}, \mathbf{p})$ is the Euclidean distance between $\mathbf{x}$ and $\mathbf{p}$ the center of P. A common way to define the additively weighted distance between a point $\mathbf{x}$ and a sphere P is to take $w(P)$ as the radius r_p of P, i.e. $d(\mathbf{x}, P) = d(\mathbf{x}, \mathbf{p}) - r_p$, which can naturally be interpreted as the shortest distance between the point $\mathbf{x}$ and the surface of the sphere P, see Figure 2(a). The resulting tessellation is called the Euclidean weighted Voronoi diagram. A point $\mathbf{x}$ lies on or inside the sphere P if and only if $d(\mathbf{x}, P) \leq 0$. Okabe *et al.* [2000] proved that the bisector of a Euclidean weighted Voronoi diagram is either a hyperbolic surface or a hyperplane, and that the generated regions are connected and star-shaped with respect to its generator sphere in $\mathbf{R}^n$. This result also holds for the Manhattan and the supremum metrics.

One way to obtain convex regions (with straight line bisectors) is to use the power distance (also known as the Laguerre distance): $d(\mathbf{x}, P) = \sqrt{d(\mathbf{x}, \mathbf{p})^2 - r_p^2}$. When $\mathbf{x}$ is outside the sphere P centered on $\mathbf{p}$ the distance from P to $\mathbf{x}$ is given by the length of the tangent from P to $\mathbf{x}$. The power distance and power bisector are illustrated in Figure 2(b) and (c), respectively.

[4] A related, but different, measure is used by Gärdenfors [2000].

The size of a particular prototype's radius relative to the surrounding prototype regions reflects its ability to influence its neighborhood. The magnitude of the radius can be related to the actual size of the category, the variability among the exemplars[5], or the correlation of qualities.

Definition 7 *We define a* power categorization *to be a generalized Voronoi tessellation* $\Delta(\mathbf{P}, d, \mathbf{C})$ *generated by a set of prototype regions* $\mathbf{P}$ *using the power distance.*

If the radii of the generator spheres are zero or equal in size, then the power categorization will be equivalent to the categorization based in the ordinary (point-based) Euclidean Voronoi tessellation.

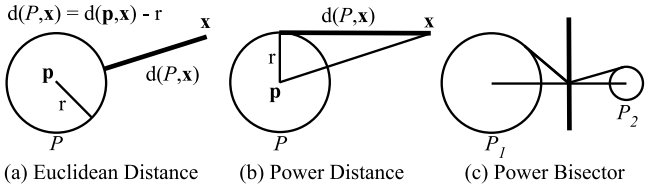

Figure 2: *The Euclidean and the power distance measures.*

The worst-case time complexity for the construction of the power categorization is no worse than that for the ordinary Voronoi diagram. Sometimes the structure of the space can be exploited, which means that the actual computational time can be dramatically lowered. Voronoi tessellations are used in a wide range of applications and domain constraints can be used to improve algorithms, typically linear time can be expected, e.g. linear time can be expected if the generating spheres/points are uniformly distributed [Dwyer, 1987].

Parallel algorithms have also been developed [Auberhammer 1991] which construct Voronoi diagrams in $O(log\ n)$ time using $O(n)$ processors.

It turns out that for power categorizations, if a generator sphere P_1 is a proper part of another generator sphere P_2 then the generated category $c(P_1)$ will not contain the center of the generator region P_1, and a well-behaved categorization will not be produced. For example, using the generic $bird$ and the $robin$ prototype regions in Figure 3 to generate a power categorization would not result in a well behaved categorization; since $PP(robin, bird)$, it turns out that $robin \notin c(robin)$ in the power categorization.

RCC5 can be used to ensure that generating spheres are not proper parts of other generating spheres, and hence can play a role in the categorization process itself, by determining the legitimate spheres to use as generators.

Finally, we define the notion of a bounded tessellation which provides a useful mechanism for selecting conceptual regions to focus on for conceptual spatial reasoning and categorization.

[5]The standard deviation of the exemplars from the prototype could be used [Gärdenfors, 2000]. In the bird conceptual space the standard deviation of birds is larger than that of emus, so we might expect the sphere that generates the bird category to be larger than that used for emus as in Figure 3.

Definition 8 *Given a generalized Voronoi tessellation* $\Delta(\mathbf{P}, d, \mathbf{C})$ *where the categories are generated by* $\mathbf{P} = \{\mathbf{p}_1, \mathbf{p}_2, ..., \mathbf{p}_n\}$ *define the Voronoi tessellation bounded by a region* S*, denoted by* $\mathbf{C}_{\cap S}$*, to be* $\{c(\mathbf{p}_1) \cap S, c(\mathbf{p}_2) \cap S, ..., c(\mathbf{p}_n) \cap S\}$*. We denote the bounded categories* $c(\mathbf{p}_i) \cap S$ *by* $c_{\cap S}(\mathbf{p}_i)$*.*

A bounded Voronoi diagram may be disconnected if every boundary region is not star-shaped with respect to its generator point [Okabe *et al.*, 2000].

5 Reasoning about Categories

Concept management involves categorization, concept acquisition, concept formation and conceptual change. Cognitive processes such as learning and communication impel and guide concept management. In the previous section we showed how conceptual spaces provide a rich and computationally effective representation for categorization based on prototype regions, and in this section we show that the RCC machinery can be used to reason about categories and to describe other aspects of concept management.

5.1 Determining Spatial Relationships

The RCC can be used to determine the relative configuration of conceptual regions such as categories, concepts, prototypes, and exemplars. For example we can determine: (i) if the smallest region containing all the prototypes is a proper part of a given category, (ii) if some category overlaps another category e.g. $PP(c(robin), c(bird))$, (iii) if the region containing all the prototypes contains all the exemplars, and (iv) if a category's kernel contains a specific region.

As an example let us consider the conceptual spaces described in Figure 3, below.

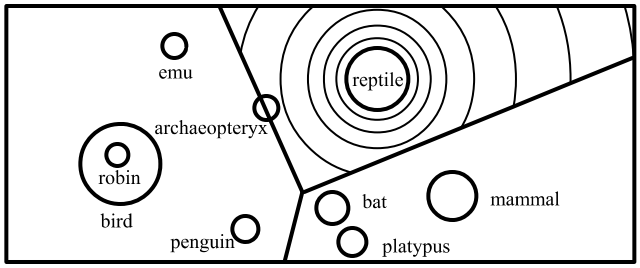

Figure 3: *Prototype regions in animal space, reptile crispings, & the power categorization of bird, mammal &reptile.*

Using RCC5 we can describe the following spatial relationships:

$DR(bird, penguin)$, $PP(robin, bird)$,
$PO(c(robin), bird)$, $PO(c(robin), kernel(c(bird)))$,
$PO(archaeopteryx, c(reptile))$,
$PO(archaeopteryx, c(bird))$,
$DR(c(archaeopteryx), c(mammal))$,
$DR(emu, penguin)$,
$DR(c(archaeopteryx), c(mammal))$,
$DR(c(robin), c(bat))$, and $DR(c(robin), c(platypus))$

The RCC, including the egg-yolk theory, allows disconnected regions, i.e. multi-pieced regions, so it supports the construction of arbitrarily complex concepts. For conceptual spaces that means one can juxtapose disconnected concepts to form eclectic ones, e.g. penguins and emus (nonflying birds), and build new concepts from existing ones. Within a single category the set of prototypes would in some applications be better modeled by a multi-pieced region rather than as a connected region such as the smallest surrounding sphere, or convex hull. In Section 5.3 we show that being able to model multi-pieced regions is important to support nonmonotonic reasoning.

For some applications it will be necessary to impose various ontological constraints on interrelated categories. In particular, it may be important to enforce consistency across the different levels of granularity so that the tessellated regions at one level are identical to the union of tessellated regions at lower levels: $c(\mathbf{p}) = \bigcup_{s \subseteq c(\mathbf{p})} c(s)$ where s are subcategories of $c(\mathbf{p})$. This constraint is present in many software engineering applications, and made explicit in data modeling techiques. One way to model this constraint in a computationally efficient manner (without distorting the underlying similarity relation) is to bound tessellations within categories. It is important to note here that $\Delta(\mathbf{P}, d, \mathbf{C}_{\cap S})$ is not identical to $\Delta(\mathbf{P}, d, S)$ in general, so bounding a conceptual space before or after the tessellation can, and typically will, result in a different categorization. For example, if the conceptual space parameters remain fixed then it would seem reasonable that the bird category region be the same regardless of whether the generator is the prototypical bird or the set of all prototypical birds at a lower level of granularity. In Figure 3 $c(bird)$ can be tessellated independently of $c(reptile)$ and $c(mammal)$, so that $c(bird) = c(robin) \cup c(penguin) \cup c(emu) \cup (c_{\cap c(bird)}(archaeopteryx))$. In other words, the generated subcategories of $c(bird)$ are bounded by $c(bird)$.

Other applications may possess weaker ontological requirements such as: If the prototype region P_1 is a subregion of the prototype region P_2 then $c(P_1) \subseteq c(P_2)$. This condition also places constraints on the way that a Voronoi tessellation can be generated across the levels of granularity, and is satisfied by $robins$ and $birds$ in Figure 3.

5.2 Crisping Conceptual Spaces

As noted in section 3 the RCC was extended by introducing an irreflexive, asymmetric and transitive binary relation $X < Y$ read as "X is crisper than Y", or "Y is a blurring of X". Cohn and Gotts also developed what has become known as the "Egg-Yolk Theory"for modeling indeterminate spatial regions. An *egg* is composed of two regions with definite boundaries; the *yolk* being a proper part of the egg. The egg and its yolk define the upper and lower bounds, respectively, on the range of indeterminacy of the region.

In this section we show how the crisper relation $<$ can be defined and the egg-yolk representation can be used to reason about categories and other conceptual regions.

A crisper relation in conceptual space can be constructed in a multitude of ways. One straightforward method is to use the proper part relation, PP, i.e. $X < Y$ if and only if $PP(X, Y)$. Figure 3 also illustrates some potential crispings of the $reptile$ category; a set of concentric spheres bounded by $c(bird)$ and $c(mammal)$. These crispings could be generated in numerous ways, e.g. using prespecified degradations in the distance measure, or by using the exemplars where each successive blurring captures another exemplar.

Given a conceptual space and a crisper relation we can build a vast range of useful queries using the RCC such as "Does a conceptual region constitute a crisping/blurring of another region?" "Does crisping a particular domain change the classification of a specific object?", "Does every category contain its prototype crisping?" and so forth.

As an example let us consider the conceptual space in Figure 3 where the crisper relation is based on PP. We have the following:

$penguin < c(penguin) < kernel(c(bird)) = c(bird)$
$robin < bird$ and $robin < c(robin) < c(bird)$
$emu < c(emu) < c(bird)$
$bat < c(bat) < c(mammal)$
$platypus < c(platypus) < c(mammal)$

Other relations describing the relationships between indeterminate regions can be constructed from the crisper relation $<$ such as $crisp(X)$ which is defined as "there does not exist a Y such that $Y < X$", and $MA(X, Y)$ which holds when X and Y are mutually approximate, i.e. they possess a common crisping [Cohn *et al.*, 1997]. From the conceptual space in Figure 3 we can say:

$Crisp(robin)$, $Crisp(mammal)$,
$Crisp(archaeopteryx \cap C(bird))$, $Crisp(penguin)$,
$\neg Crisp(bird), \neg Crisp(kernel(bird)), \neg Crisp(c(reptile))$,
$MA(c(penguin), c(bird))$, $MA(robin, bird)$,
$MA(archaeopteryx, c(bird))$,
$MA(archaeopteryx, c(reptile))$,

The RCC framework provides a number of axioms that govern the crisping relation in different kinds of applications. For example, there are axioms that ensure the existence of a complete crisping of any region, and the existence of alternative crispings and blurrings.

Crispings can play a role in the process of categorization itself; they can define regions to be used to generate tessellations. For example, in Figure 4, below, the bisector shifts towards the sphere P_1 with center $\mathbf{p}_1$ and radius r if P_1 is crisped to a smaller sphere P_1' with radius r'. This crisping can be modeled precisely; the bisector between P_1 and P_i moves by distance $(r - r')/2d(\mathbf{p}_1, \mathbf{p}_i)$ towards $\mathbf{p}_1$ in parallel to its previous location. So it is easy to show that $P_1' < P_1$ if and only if $c(P_1') < c(P_1)$, i.e. a local crisping (blurring) of a prototype region crisps (blurs) its category, and conversely.

In well behaved categorizations one can construct an egg yolk system using the kernel of each category. The yolk of the prototype region P_1 can be given by the largest sphere enclosed by $kernel(c(P_1))$, say $protoyolk$, and the egg can be given by the smallest sphere circumscribing $kernel(c(P_1))$, say $protoegg$. This egg-yolk prototype system can then be used to generate the corresponding egg-yolk category system: $c(protoyolk)$ and $c(protoegg)$.

Finally we extend the notion of crisping to distance measures to capture categorical perception; the observed phenom-

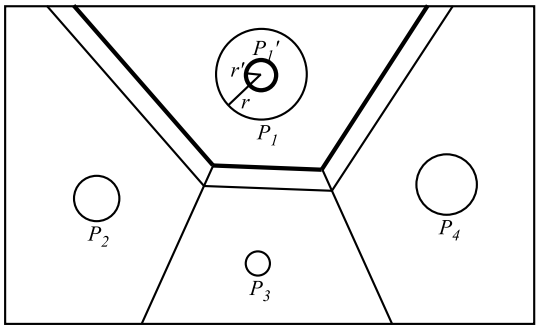

Figure 4: *Crisping a prototype crisps the category.*

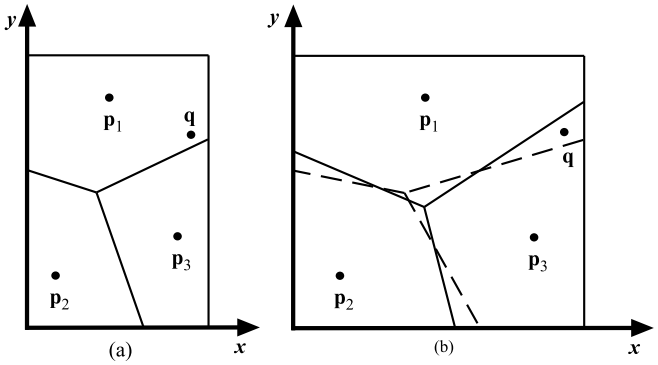

Figure 5: *Changed categorization due to a change in context.*

ena where category members are judged to be more similar than members and nonmembers immediately across category boundaries. For example, the *distance measure d' is a crisping of distance measure d* where

$$d'(\mathbf{a},\mathbf{b}) = \begin{cases} k(d(\mathbf{a},\mathbf{b})) \text{ if } \mathbf{a},\mathbf{b} \in c(P) \text{ for some } P \\ d(\mathbf{a},\mathbf{b}) \text{ otherwise} \end{cases}$$

for some $0 \leq k < 1$.

In the limit we have $d'(\mathbf{a},\mathbf{b})$ is 0 if **a** and **b** are in the same category and $d(\mathbf{a},\mathbf{b})$ otherwise. Categorical perception, and hence crisping a distance measure, represents a form of learning. Our definition can be extended in numerous ways and the similarity relation derived from a crisped distance measure can easily be given a wide variety of threshold behaviours.

5.3 Nonmonotonicity and Concept Management

In this section we highlight the nonmonotonic effects of changing context, and show how conceptual spaces can be used as an underlying model from which more traditional nonmonotonic reasoning formalisms can be derived.

Nonmonotonic changes to the categorization can arise in several ways: (i) by *focusing* on a region, (ii) by *modifying* the underlying conceptual space, or (iii) by *changing* the mapping of objects to conceptual regions.

Focusing on a Region

Focusing can be accomplished by changing the dimension weights, by crisping a region, bounding a region, or combining regions.

As noted earlier categorization is context-sensitive. In conceptual spaces context-dependence is modeled using weighted dimensions. For example, weighting the distance measure along the x-axis results in a different categorization via the Voronoi tessellation such that objects change categories. Technically this is achieved by multiplying the specific dimensions by a given weight where the weight reflects its salience. For instance, under the Euclidean metric a weight can be placed on dimension i as follows: $\sqrt{\sum_i w_i(x_i - p_i)^2}$. Weighting specific dimensions gives rise to a nonmonotonic crisping relation where $X < Y$ does not imply $PP(X,Y)$; some regions will contract in size others will dilate. In Figure 5, below, the quality dimension represented by the x-axis in (a) becomes more salient and is elongated causing the object "**q**" to be reclassified in (b).

Modifying the underlying Conceptual Space

The underlying conceptual space can be modified by adding and deleting exemplars or prototypes, changing the distance measure, changing the function relating the distance measure to similarity, or changing the generator prototype regions.

The introduction of exemplars and/or prototypes will shift the boundaries and create new categorical regions [Gärdenfors, 2000]. As noted earlier the addition and removal of prototypes is a fundamental concept management operation, and can be achieved in linear time in the worst case.

Changing the underlying distance measure or generating regions will shift boundaries. Merely crisping the distance measure or modifying the function relating the distance measure to similarity will not change the underlying categorization, but will affect the magnitude of the similarity judgments.

Changing the mapping of objects to regions

Nonmonotonic reasoning formalisms are typically logic based, and hence symbolic systems. The conceptual space framework can be used to model nonmonotonic information, and used to construct nonmonotonic inference rules via the RCC. Since the conceptual space model is based on the measure of similarity to prototypes the RCC's crisper relation can be used to build representations of minimal models or usual states of affairs.

Just as specific individual objects are points in conceptual space, generic (or under specified) objects are (possibly disconnected) regions. One might expect that the more generic or the more unspecified an object, the larger the region used to represent it.

The generic bird, Tweety, would be represented as a central region in bird space, e.g. the prototype $bird$ region. In which case we would expect Tweety to possess all the features common to birds in that region; we expect Tweety to possess feathers, two legs, wings for flight, a four chambered heart, and so forth. As we learn more about Tweety we adjust the target region used to represent him. If we learn he is a robin, then he could be remapped to the $robin$ prototype. On the other hand, if we learn that he is a nonflying bird, then we may remap him to the prototype $penguin$ region and the prototype emu region which in the example in Figure 3 is a disconnected region, and we still expect Tweety to have feathers, two legs, and a four chambered heart.

The RCC can represent conceptual regions which are re-

quired to support all concept management for nonmonotonic reasoning as described above, and as such it forms a natural bridge from the geometrical conceptual space representation to the symbolic representation in standard nonmonotonic formalisms.

6 Conclusion

We showed how algorithms in computational geometry and the RCC can be applied to the conceptual space framework. Categorization in conceptual spaces is achieved via (generalized) Voronoi tessellations based on a similarity relation which results in a prototype being centrally located in its category. An analysis of existing algorithms in computational geometry established that categorization based on an underlying similarity relation which is used to generate Voronoi diagrams is feasible. Furthermore, once a categorization is completed concept management tasks like determining the category of an arbitrary object, adding and deleting prototypes, and merging categories are computationally fast. For example, once a classification has taken place the time taken to identify the category of an arbitrary object is logarithmic, and the time required to add or delete a prototype is linear.

We demonstrated that the ability to reason about categories and other conceptual regions using the RCC enhances the conceptual space model. We showed that the RCC can be used to construct both monotonic and nonmonotonic reasoning systems from the information embedded in the categories of conceptual spaces, hence the RCC forms a natural bridge from the geometrical representation of information in conceptual spaces to its symbolic counterpart.

In addition, the RCC provides facilities to determine the prototype regions from which to generate categories so that well behaved categorizations are produced.

The crisper relation $<$ and egg-yolk systems can be generated in conceptual spaces using prototypes and exemplars, and hence can be used to model the indeterminacy of conceptual regions. Interesting properties followed from our constructions, for example, we were able to show that crisping a prototype region locally, leads to the crisping of its category under the power categorization, and that categorical perception can be explained by crisping the distance measure on the conceptual space.

References

[Aurenhammer, 1987] Aurenhammer, F., Power Diagrams: Properties, Algorithms and Applications, *SIAM Journal of Computing Surveys*, 16(1):78-96, 1987.

[Aurenhammer, 1991] Aurenhammer, F., Voronoi Diagrams: A Survey of a Fundamental Data Structure, *ACM Computing Surveys*, 23(3), 345 - 405, 1991.

[Balkenius, 1999] Balkenius, C., Are There Dimensions in the Brain? in *Spinning Ideas, Electronic Essays Dedicated to Peter Gärdenfors on His Fiftieth Birthday* online at http://www.lucs.lu.se/spinning.

[Borsuk and Szmielew, 1960] Borsuk, K., and Szmielew, W, *Foundations of Geometry*. Amsterdam: North Holland, 1960.

[Bennett, 1997] Bennett, B., Logical Representations for Automated Reasoning about Spatial Relationships, *PhD Thesis*, University of Leeds, 1997.

[Chella *et al.*, 1998] Chella, A., Frixione, M., and Gaglio, S. An architecture for autonomous agents exploiting conceptual representations, *Robotics and Autonomous Systems*, 25(3-4):231–240, 1998.

[Cohn *et al.*, 1997] Cohn, A. G., Bennett, B., Gooday, J. and Gotts. N.M., Qualitative Spatial Representation and Reasoning with the Region Connection Calculus, *Geoinformatica*, 1(3), 1-42, 1997.

[Cohn and Gotts, 1996] Cohn, A. G. and Gotts. N.M., Representing Spatial Vagueness: A Mereological Approach, in *Principles of Knowledge Representation and Reasoning: Proceedings of the 5th International Conference*, Morgan Kaufmann, San Francisco, 230-241, 1996.

[Dwyer, 1987] Dwyer, R.A., Higher-Dimensional Voronoi Diagrams in Linear Expected Time, in *Proceedings of the 5th Annual ACM Symposium on Computational Geometry*, 326 - 333, 1987.

[Gärdenfors, 2000] Gärdenfors, P., *Conceptual Spaces: The Geometry of Thought*, A Bradford Book, MIT Press, Cambridge Massachusetts, 2000.

[Harnad, 1987] Harnad, S., *Categorical Perception: The Groundwork of Cognition*, Cambridge University Press.

[Hampton, 1993] Hampton, J., Prototype Models of Concept Representation, *Categories and Concepts: Theoretical Views and Inductive Data Analysis*, Academic Press, London, 67 - 95, 1993.

[Krusal, 1978] Krusal, J.B., and Wish, M, *Multi-dimensional Scaling*. Beverley Hills, CA: Sage Publications, 1978.

[Klee, 1980] Klee, V., On the complexity of d-dimensional Voronoi diagrams, *Archiv de Mathematik*, 34: 74 - 80, 1980.

[May *et al.*, 1989] May, B. Moody, D.B., Stebbins, W.C., Categorical perception of conspecific communication sounds by Japanese macaques, Macaca fuscata , *J. Acoust. Soc. Am.*, Vol. 85, No. 2, Pages 837 - 847, 1989.

[Nebel, 1995] Nebel, B., Computational Properties of Qualitative Spatial Reasoning, in Wachsmuth, Rollinger and Brauer (eds), *Proceedings of the 19th German AI Conference*, Vol 981 LNCS, 233 - 244, Springer-Verlag, 1995.

[Okabe et al., 2000] Okabe, A., Boots, B., Sugihara, K., and Chiu, S.N., *Spatial Tessellations*, 2nd Ed, Wiley, 2000.

[Petitot, 1989] Petitot, J., Morphodynamics and the Categorical Perception of Phonological Units, *Theoretical Linguistics*, 15: 25 -71, 1989.

[Renz and Nebel, 1999] Renz, J. and Nebel, B., On the Complexity of Qualitative Spatial Reasoning: A Maximal Tractable Fragment of the Region Connection Calculus, *Artificial Intelligence*, 108 (1-2): 69 -123, 1999.

[Rosch, 1975] Rosch, E., Cognitive representations of semantic categories, *Journal of Experimental Psychology: General*, 104: 192 - 232, 1975.