SHAPE SKELETONS AND SHAPE SIMILARITY

By

ERICA JAN BRISCOE

A Dissertation submitted to the

Graduate School-New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Psychology

written under the direction of

Jacob Feldman

and approved by

____

____

____

____

New Brunswick, New Jersey

[October, 2008]

ABSTRACT OF THE DISSERTATION

Shape Skeletons and Shape Similarity

by ERICA JAN BRISCOE

Dissertation Director:

Jacob Feldman

Judgments of similarity play an integral role in the human cognitive system as they provide a means for extracting information about how objects in the world relate to each other. This similarity information is applied in various cognitive tasks, such as categorization, recognition, and identification. Previous work suggests that perceived objects are cognitively represented in a psychological space where similarity is preserved, allowing for an internal *structured* representation of objects in the world (Shepard, 1964). For an internal representation to be formed, information about an object must be extracted. Shape, a highly informative and salient property of an object, is often used. Judgments made about shape similarity reflect how humans functionally represent and utilize shape information from an object. Computational shape representation has been achieved with varying amounts of success (e.g. Blum, 1973;

Biederman, 1987). This variability is due, in part, to the complexity of mimicking the seemingly effortless human ability to make judgments about shape even in spite of numerous possible complications, such as sparse information and occlusions. This work presents the use of a Bayesian estimation of a shape's skeleton, the maximum *a posteriori* (MAP) skeleton (Feldman & Singh, 2006), as part of a generative model of shape that allows for the computation of a probabilistically-based similarity metric. This method of shape representation makes possible the prediction of similarity judgments reported by human subjects on collections of shapes that exhibit differences in both part structure and metric qualities and that have been generated by an unrelated process. It is argued that the derivation of a similarity metric from this model provides the previously unavailable relationship between shape representation and categorical judgments about shape.

TABLE OF CONTENTS

LIST OF FIGURES

1.0 INTRODUCTION

Judgments of similarity play an integral role in the human cognitive system as they provide a means for extracting information about how objects in the world relate to each other. At first glance, the term "similarity" seems too ambiguous to offer much in the way of a manageable application to cognitive modelling. As Nelson Goodman (1972) suggested, similarity could be an empty explanatory construct, as any two things can be as similar or dissimilar, depending on the respect in which their similarities are depicted. Objects with seemingly disparate qualities can become similar by virtue of, for example, being in the same room. Despite this ambiguity, the degree to which people perceive two things as similar fundamentally affects their rational thought and behavior (e.g. Tversky, 1977; Coombs, 1964). When similarity is considered in a psychological context under experimental conditions, it becomes constrained in such a way to allow for the evaluation of its role in cognition, most notably for its involvement in recognition (e.g. Ashby & Perrin, 1988), identification (e.g. Ashby & Lee, 1991), learning (Gentner, 1989), problem solving (Gick & Holyoak, 1980), and categorization tasks (e.g. Posner & Keele, 1968; Rosch, 1978).

Defining similarity in concrete terms has proven difficult. This difficulty arises not only because of its ambiguous nature, but also because similarity is not a static construct. Its determinants may change based on various factors, such as context.

Similarity sometimes means the degree to which two objects share physical or perceptual features, and sometimes the degree to which two objects share "deeper", non-perceptual features (Quine, 1977). Popular theories of similarity measurement from the psychological literature include Shepard's (1957) spatial account, Tversky's (1977) contrast model, and the more recent alignment model (Markman & Gentner 1993; Goldstone & Medin, 1994). The assessment of similarity for particular machine applications have also been mechanized by various techniques, such as neural networks, statistical methods, or productions rules (Bareiss & King, 1989). While these specifically designed methods may prove successful for their intended applications, they are often not broad enough to be applied within any context and are built on assumptions distant from the psychological considerations that motivate cognitive modelling.

The role of similarity as related to the human visual system is also ambiguous and requires an investigation of how judgments of similarity arise from the information provided by the visual system. It is often assumed in the study of vision that the human visual system should attempt to provide an accurate replication of the world to use in later stages of cognitive processing. To this end, computational models of vision are created to mechanize human visual and cognitive capabilities, and as a result, models of 2D shape are of critical value. While objects may be characterized in numerous ways, such as by color or texture, shape has been shown to be a primary source of information for object recognition (Biederman, 1987; Biederman & Ju, 1988). For a variety of reasons, shape representation has proven difficult to model, resulting in a wide assortment of approaches.

Mere representation, however, is not the only goal of shape modelling efforts. It is also important that models can be applied to solve various types of problems, especially those requiring systematic comparisons between shapes, as this is a critical function utilized for tasks such as recognition, matching, and classification (e.g. Siddiqi, Dickinson, & Zucker, 1999; Bai & Latecki, 2008). Comparisons of this nature necessitate that shape representation models allow for some account of the similarity between shapes; however, shape similarity is difficult to assess as it relies on the selection of discriminative shape features, which might change according to task, context, or due to the particular characteristics of the shape. Often shape representation models that *do* provide a simple metric for shape similarity, such as those based on simple descriptors, are not robust enough to capture important shape information. More complex models, such as chain codes (Freeman, 1961), capture more information about a shape (though not all, such as part structure), but do not readily provide a similarity metric by which to compare shapes. Because it is so critical to many cognitive tasks, such as categorization and overcoming changes in viewpoint, the determination of shape similarity should ideally arise from any theory of shape representation.

Humans are extremely capable of distinguishing subtle differences in shape for various tasks, such as similarity and categorization judgments. Shepard (1957, 1987) was first to recognize the importance of similarity in human judgments with his proposal of a *universal law of perceptual generalization*, which explained human similarity judgments in terms of the proximity between representations of objects in an internal representation space. The use of an internal shape representation space is valuable only if the patterns of proximities between objects there are reflective of the similarities among the real-world objects themselves. Under a "space" representation theory, visual perception should provide the qualities of shape that allow for placing it in the proper place in representational space (Cutzu & Edelman, 1996). In order to produce a veridical representation, an entire geometrical reconstruction is not always necessary, where often only representation of contrasts or dissimilarities between objects is adequate (Ashby & Perrin, 1988).

2.0 THE PROBLEM

Most of the work available on shape representation attempts to capture specific shape attributes (see section 4.1) exclusively, based on either contour or part structure. However, few, if any, shape representation techniques are so robust as to be able to represent the characteristics of shape that allow for a wide range of comparisons between different types of 2D contour shapes, as occur in human similarity judgments. For example, to group the four shapes in Figure 1 into two sets on the basis of similarity, the shapes might be arranged as two sets, one including shapes *a* and *b* and one including shapes *c* and *d* on the basis of part structure. Shapes *a* and *b* can be seen as exhibiting one part while *c* and *d* exhibit multiple parts. However, another grouping, where shapes *a* and *c* are one set and *b* and *d* form the other, may also be distinguished, as these groupings reflect the similar "roundness" of the contour. This example highlights the need for a shape representation method that is not solely based on either part structure or contour properties, as both elements may be influential in human similarity judgments. The model presented in this thesis attempts to place structural and metric shape properties in a common framework by considering their generative origins in a common formalism, a shape-generating skeleton.

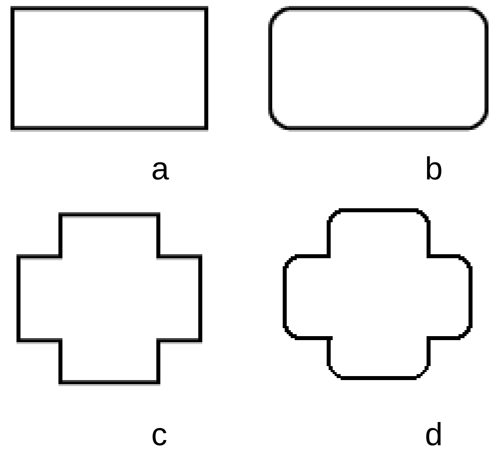

Figure 1: Example of shapes that may be grouped differently, either by general shape (rows) or roundness of boundary contour (columns) on the basis of similarity

A large portion of the available work on shape similarity, primarily within the field of computer science, is focused a *shape matching,* which involves shape *classification* by some method of label assignment. Shape matching usually concentrates on images rather than simple contours and is largely unmotivated by psychological constructs. Most often, shape matching is used for applications with a specific goal, such as for shape retrieval, classification and recognition by machine processes, which creates a dependence of the similarity computation on the particular application and the specific types of associated shapes. Shape matching is related to the work presented here, as it typically involves the measurement of the resemblance of one shape to another using a predefined measure, however, its motivation and scope is greatly disparate. While there are various successful matching algorithms, shape matching involves a limited classification into a well-defined set of disjoint categories, without regard to gradations of similarity either within or between the categories. A more broadly construed treatment of similarity, like the one presented here, emphasizes the gradations and attempts to find a similarity metric robust enough to support inductive inferences above and beyond predefined classification.

Hume (1975) suggested that arguments created on the basis of experience are founded on similarity, with our expectation that similar effects flow from similar causes, making similarity critical to inductive cognitive processes. A more theoretical difference between shape matching and psychologically motivated similarity metrics is that the similarity used in perception has a much more far-reaching impact towards the human ability to make generalizations about the world (e.g. Shepard, 1987; Tennenbaum & Griffiths, 2001). This difference is largely created from the diverse types of problems that are faced by computer science applications and the human visual system, where human similarity judgments represent a more comprehensive attempt to extract shape information that extends beyond mere classification.

In shape matching applications the determination of similarity between two shapes is largely static; within the same database of images, the comparison of two shapes will consistently produce the same result. This lack of variation is not, however, exhibited in human behavior, where a variety of influences can affect similarity judgments (Rips, 1989; Stewart & Chater, 2002; Ashby & Townsend, 1986). As an example of how human similarity judgments can vary, Figure 2 graphically depicts two categories *A* and *B*, where the size of each oval represents the category's variability. Both categories contain the objects *x* and *y*. When comparing only objects from category *A*, the judgment of dissimilarity between *x* and *y* will be much greater than a judgment between the two objects when the comparison is only between stimuli from category *B*. This difference results from the consideration of within-category variability, as the separation between *x* and *y* in category *A* covers more of the *total* “space” than it does in category B, resulting in a higher dissimilarity rating. Therefore a cognitively motivated representation of shape should provide a means of incorporating contextual effects, such as category variability, within the representatioal framework.

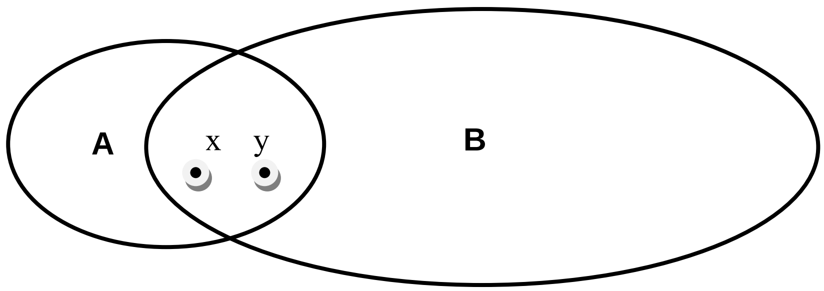

Figure 2: Two categories, *A* and *B*, whose variability affects the similarity rating of exemplars *x* and *y*.

3.0 OBJECTIVES

This thesis makes multiple contributions. First, it will demonstrate the need for a shape representation theory that bridges the gap between psychologically-based computer vision models of shape, which primarily concentrate on successful application, and cognitive models, wherein shape representation is presently either overly simplified or assumed without a comprehensive computational model. Second, a new model of shape similarity, called the Generative Similarity (GS) model is presented that provides a well-motivated *explanation* of shape by formalizing the processes that generate it. Third, this type of representation allows for the derivation of a meaningful metric of similarity that can be used in both computer vision and psychological models to predict experimentally collected data. Fourth, the model is presented as part of a larger, comprehensive probabilistic framework that is able to incorporate dynamic assumptions, primarily represented as probabilities, so that it may be applied within a variety of contexts and applications.

4.0 BACKGROUND

4.1 SHAPE REPRESENTATION

Shape is a critical feature of objects, and one of the primary means by which humans reason about objects in their world. Approaches to shape representation are numerous, largely because shape representation and description is such an enormously difficult task.

The most common classification of shape representation methods divides approaches into two types; *contour*- and *region*-based (Zhang & Lu, 2004). Contour-based representations (sometimes referred to as boundary-based, e.g. Pavlidis, 1978) are based on the use of shape boundary points rather than the interior of the shape. Region-based methods use both the boundary and interior points in the representational scheme. Within both the contour and region distinctions, methods may be identified as either *global* or *structural* (Zhang & Lu, 2004). The former, *global*, results in a representation of the entire shape, while the latter, *structural*, results in a shape represented through a collection of components or *primitives*.

Under the broad dichotomy of contour and region-based shape representation, there are various other mid-level classifications that vary among authors and fields. Costa and Cesar (2001) separate the contour-based representation into three types.

Parametric contours are those where the shape outline is represented as a parametric curve, implying a sequential order. The second type is merely a *set of contour points*, without any particular order. The third type of contour representation is *curve approximation*, where a set of geometric primitives (such as splines or segments) is fitted to the shape contour. Costa and Cesar also separate region-based shape representation into three classes. Region decomposition allows for the shape region to be divided and represented by simpler forms (e.g. polygons). Under a bounding regions method, the shape is represented as an approximation by some defined geometric primitive that is fitted to the shape. The third type of region-based representation uses internal features that represent the shape's internal region (e.g. skeletons).

The psychologist Fred Attneave first suggested that an important source of information in a shape is contained in its high curvature points, where "contours change direction maximally" (Attneave, 1957). Attneave suggested many of the important elements of current shape representation algorithms, such as the extraction of curves as contours of an image and the smoothing and encoding of curves into a robust representation. Other psychologists have also acknowledged the importance of information stored at maximum curvature points. Hoffman and Richards (1984) proposed that the visual system decomposes objects at points of high negative curvature. Feldman & Singh (2005) proposed a formalism to represent the information at areas of maximum curvature, incorporating sign of curvature as well as magnitude. Leyton's (1999) work recognizes the importance of curvature maxima with his symmetry-curvature theorem, which proposes that all shapes are basically circles which have changed as a result of various deformations caused by external forces like pushing, pulling, and stretching.

Some of the more psychologically prominent approaches for shape representation involve parsing a shape into a defined set of parts that are then stored in a structural description of propositional relationships. The significance of this type of representation is based primarily on findings that suggest the psychological importance of a shape's parts, (e.g. Biederman, 1987; Hoffman & Richards, 1984; Marr & Nishihara, 1978; Saiki & Hummell, 1998). For example, Tversky and Hemenway (1984) had subjects report features of various categories and found that they often listed the objects' parts. A challenge for this approach is the representation of the infinite variety of shapes with a small set of primitives. The ability to form a hierarchical relationship between parts also makes these methods especially powerful. These methods are also attractive because they provide symbolic descriptions of objects; however, some objects do not have a clear decomposition into generic parts. It may also be difficult to extract generic parts from images in a meaningful way.

4.2 SKELETON SHAPE REPRESENTATION

Blum (1973) introduced the first axis-based representation of two dimensional shapes, referred to as the symmetric or the medial axis. The medial axis of an object is defined as the set of centers of all maximally inscribed disks in a shape (see Figure 3). For two dimensional objects the medial axis is one-dimensional. The hierarchical structured formed by the medial axis is attractive from a computational point of view as it captures local symmetries of an object and provides a natural decomposition of the object into parts that correspond to branches in the one-dimensional structure. Blum's medial axis, as originally computed, is sensitive to contour noise, where small disruptions in a shape's contour can cause gross differences in the resulting medial axis (see Figure 4).

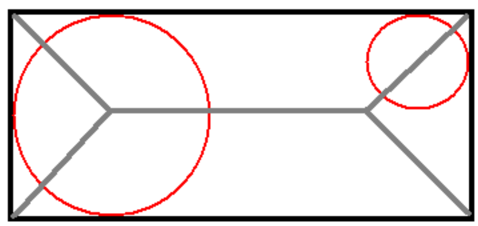

Figure 3: The medial axis of a rectangle, formed by connecting the centers of the maximally inscribed circles.

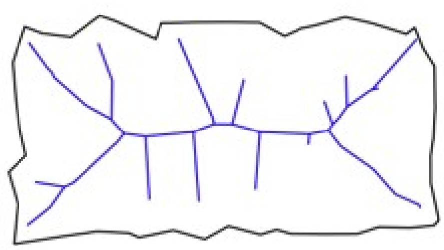

Figure 4: Medial axis from a rectangle with contour noise. The multiple axes demonstrate a weakness in Blum's medial axis approach.

In general, representations based on the medial axis are adequate to capture the geometry of 2D shapes. Medial axis models can also better capture natural deformations of objects compared to boundary models (Sebastian & Kimia, 2001). A closely related representation, called the shock graph (Siddiqi & Kimia, 1996), which uses the medial axis along with geometric and dynamic information at the points of discontinuities, has also met with much success and has been used to model deformation and changes in an object's structure (Sebastian, Klein & Kimia, 2001).

A recent effort to address many of the weaknesses of the traditional medial axis transform is a Bayesian estimation of the shape skeleton (Feldman & Singh, 2006). This approach assumes that shapes arise from a mixture of generative and random factors, specifically a skeleton and a "stochastic growth process". Here, a process estimates the skeletal structure most likely to have generated the shape. The model assumes (1) a prior over skeletons and (2) a likelihood function that models how shapes are generated from skeletons.

A key element of this approach, which will be expanded upon in the model presented below, is that it allows for alternative "explanations" for a given shape. In the approach to similarity developed in this thesis, the key idea is to question whether two shapes are better explained as resulting from a common generative source or from two independent sources; therefore, the Bayesian estimation of the shape skeleton allows for this evaluation.

4.3 SIMILARITY

Much of the historical study of similarity pertains to judgments used to determine whether an internal cognitive representation matches its reference in the "real world". The traditional view, attributed to Aristotle, is that an internal representation refers to an external object according to its resemblance or isomorphism between the two. For example, the representation of a mouse has the quality of being furry and small. Rather than believing that the internal representation of the mouse *is* small and furry, it is obvious that it is represented as a set of measurements that embody the visual qualities of a mouse. The form of this information in regards to shape, be it structural (e.g. Biederman, 1987), metric (e.g. Ullman, 1989), or as forwarded in this thesis, both, demonstrates the one to one correspondence between properties in the brain to those in the world.

Another importance of similarity arises when considering that the visual system attempts a *second-order isomorphism* (Shepard, 1968) between similarities of shape and similarities of the internal representations of shapes, as opposed to the first-order isomorphism between the shapes and their direct representations. For example, an internal representation of the shape of a cat may take various forms, but it should demonstrate some "closeness" or similarity to the internal representation of a rabbit than to the smell of burning tires, or any other such stored concept.

Utilizing this second-order isomorphism, similarity has been used to investigate the representation of shape and its veridicality to worldly objects (e.g. Shepard & Chipman, 1970; Shepard and Cermak, 1973; Cortese and Dyre, 1996; Cutzu and

Edelman, 1996). These studies have tested and demonstrated the consistency of both the correspondence of similarity judgments between subjects and, when using parametrically varied stimuli, between the subject-rating produced MDS patterns (see section 4.3.1) and those formed by the constructed stimulus space.

4.3.1 Multidimensional Scaling (MDS)

The main tool used to generate an internal shape space for similarity judgments is multidimensional scaling (MDS). In a typical MDS experiment, subjects are presented with pairs of stimuli and are asked to indicate their perceived similarity using a numerical rating. These ratings are then used to produce a geometric representation in which each stimulus is identified with a point in a multidimensional perceptual space with the property that similarity and the inferred psychological distance are inversely related. An assumption of MDS models is that subjects base their similarity judgments on the distance between the raw perceptual representations of the two stimuli. Thus, the computation of a distance metric is one way of integrating information from various perceptual dimensions, creating a psychological space, which can then be used to predict behavior on tasks such as categorization (Shin & Nosofsky, 1992). In this sense, MDS has value as a tool *after* similarity has already been evaluated. It is also necessary to be cautious in linking ”distance” in MDS space to true psychological similarity. Clark (1993) points out that although the distances within MDS space are monotonically related to similarity, it cannot be assumed that ratios of distances are interpretable and that there may not be a common unit to be used on each dimensional axis.

An MDS plot is a descriptive tool that can also be used to explain behavior in terms of a decision boundary model (Ashby & Townsend, 1986). According to this type of model, objects give rise to distributions of points in psychological space and a classifier partitions the space by forming boundaries. These boundaries create regions wherein ratings of similarity will be like. While MDS has an enormous literature, it should be noted that the goals of MDS do not include any type of predictive or explanatory computation of similarity, but focus instead on the understanding of the geometrical structure of empirically observed similarity judgments. In this sense, MDS does not by itself provide any means for determining the similarity between two stimuli.

4.3.2 Shape Similarity

Efforts toward creating efficient and robust methods of shape representation are often focused on shape matching, classification, or recognition. One of the primary means for performing these tasks is by utilizing shape similarity. For example, using shock graphs, in which a shape's contour is decomposed into a set of qualitative parts, captured in a directed acyclic graph, the similarity between two shapes is produced by comparing their shock graph topology and attributes (Sebastian, Klein, Kimia, 2002; Siddiqi, Shokoufandeh, Dickinson, & Zucker, 1998). Shape representation approaches that use a finite set of features may provide a measure of dissimilarity by finding the Euclidian distance between their features, where, for example, the features may be Fourier descriptors (Lin & Chellappa, 1987; Cortese & Dyre, 1996) or Zernike moments (Zhang & Lu, 2003). Local features may include local tangents, curvature, and other, qualitative descriptions of shape boundaries (Carlsson, 1999). In this case, computing the distance between shapes involves finding point-wise correspondences between the shapes' contours. These correspondences are often found by applying optimization techniques, particularly dynamic programming (e.g. Basri, Costa, Geiger, & Jacobs, 1998; Gdalyahu & Weinshall, 1999, Sebastian, Klien, & Kimia, 2003) or the fast marching method (Frenkel & Basri, 2003). A method developed by Pizer et al. (1999) represents shapes using a multiscale medial axis representation and has been successfully used by Yushkevich, Pizer, Joshi, and Marron (2001) for the examination of shape variability, in terms of growing and bending, in diagnostic classification for medical imaging.

Another more modern approach suggests that the similarity between two objects is a function of the "complexity" required to "transform" the representation of one object into another (Chater & Hahn, 1997; Hahn & Chater, 1997). The more simple this transformation process, the more similar the two objects are assumed to be. The idea that the involvement of a transformation is related to similarity has also been used to "map out" psychological space (Feldman & Richards, 1998) and for machine learning applications (Li & Vitanyi, 1997). This notion of transformation is highly relevant to the similarity metric proposed in this thesis as it introduces the role of *processes* in similarity measurement.

Another issue pertinent to shape similarity involves the integrality of stimulus dimensions, which is how different stimulus attributes combine and interact in perceptual processing. Traditionally, perceptual dimensions are characterized according to whether they are separable or integral (Attneave, 1950; Shepard, 1964). A pair of dimensions are said to be separable if it is easy to attend to one and ignore the other (Shepard, 1964). With an integral dimension, it is extremely difficult to attend to one dimension and ignore the other (Lockhead, 1966). Similarity judgments involve the assessment of similarity in multiple respects, often involving the combination of features, such as size, shape, number of parts, etc., which has been found to have a neural basis (Drucker & Aguirre, 2008). A well-founded similarity metric can reflect this integration if the shape representation and object comparison method on which it is based is capable of representing integrated information in a veridical fashion.

4.3.3 Shape Matching

As previously mentioned, shape matching uses the similarity between objects' shapes for particular tasks that are most often encountered within computer science applications. In many object recognition and content-based image indexing applications, object outlines are represented as curves and then matched. The matching relies on either aligning feature points using an optimal similarity transformation, which rely on matching feature points by finding the optimal rotation, translation, and scaling parameters (e.g. Umeyama, 1998), or on a deformation-based approach that aligns the properties of the two curves, such as by finding a mapping from one curve to the other that penalizes stretching and bending (e.g. Liu & Srinath, 1990; Basri, Costa, Geiger, & Jacobs, 1998).

Approaches to shape matching also utilize various other shape representation techniques, such as the *Hough Transform* (Ballard, 1981), *deformable templates* (Sclaroff & Pentland, 1995), *Fourier descriptors* (Loncaric, 1998), and *curvature scale space* (Mokhtarian, Abbasi, & Kittler, 1996). For example, by representing an image by its *Zernike moments* (Khotanzad & Hong, 1990), functions that are moment invariants can be defined for an image such that only a few low-order moments are needed for an adequate representation. These moment invariants can be put into a feature vector that may subsequently be used for matching (Veltkamp, 2001).

The actual similarity computation used in shape matching can take various other forms (see Veltkamp & Latecki (2006) for a review). The s*hape context* method (Belongie, Malik, & Puzicha, 2002) builds a shape representation for each contour point then uses statistics of other contour points "viewed" by each point to create distances that are used as the similarity measure. The *area of symmetric difference* measure, also called the *template metric*, is defined as the union of the difference of area of each of two shapes subtracted from the other (Alt, Fuchs, Rote, & Weber, 1996). More psychologically relevant is a similarity measure based on *convex parts correspondence* (Latecki & Lakaemper, 2000). This method uses the optimal correspondence of contour parts of two compared shapes, where the correspondence is computed on contours simplified by a discrete curve evolution using dynamic programming (Latecki & Lakaemper, 1999).

Recently, some of the more significant shape matching results are based on models that use a shape's skeletal structure as a basis for similarity. Shock graph matching has been used (Pelillo, Siddiqi, & Zucker, 1999; Sharvit, Chan, Tek, & Kimia,

1998) for object recognition and image indexing tasks. Zhu and Yuille (1996) have proposed a frame-work (FORMS) that decomposes 2D shapes into connected mid-grained skeletal parts. Matches are based on similarity between parts, computed as a joint probability by employing a grammar that "constructs" shapes out of their primitives.

A problem addressed by a few recent models is the fact that an object's shape may change as the result of part articulation. A model with reasoning similar to that of this thesis has been proposed by Bai and Latecki (2008), which develops a novel method for matching skeleton graphs by contrasting the geodesic paths between shape skeletons. Their primary motivation stems from the fact that similar shapes may exhibit grossly different topological skeletal structure, for example, the skeleton of a man with his arms down compared to the skeleton of a man with his arms raised. This "mobility" of parts creates the need for a technique that is articulation invariant. Their approach uses the similarity between the shortest paths between each pair of endpoints of a 2D shape's pruned skeleton. The similarity between the shortest paths establishes a correspondence relation between the endpoints for graphs from different shapes. The shapes are then matched by corresponding the shape descriptors of the skeleton endpoints. They find that their method is successful at classification of images by using standard image databases, outperforming shape matching based on shock trees (Siddiqi, Kimia, Tannenbaum, & Zucker, 1999) and shape contexts (Belongie, Malik, & Puzicha, 2002). A similar articulation invariant method uses inner-distances, the shortest paths between landmark points on a shape silhouette, as a descriptor of complex shapes (Ling & Jacobs, 2005; Ling & Jacobs, 2007);

4.3.4 Neural work on similarity

The mapping between "real world" stimuli and internal representation space leads to the search for neural evidence of such an internal representation. Specifically, investigators have looked for neural units that respond preferentially to certain objects, with response falling off monotonically with dissimilarity between the object and another stimulus. Several investigations have suggested that stimuli could be represented in the brain by a recoding of the visual input into distance functions (Edelman, 1999; Pouget & Snyder, 2000).

Tanaka et al., (1991, 1992, 1993) found selectivity for specific objects in recordings from the inferotemporal (IT) cortex of anesthetized monkeys. This work suggest a parallel between the functional organization of the IT cortex and the primary visual cortex, where the former has cells responding to similar shapes arranged in columns perpendicular to the cortical surface and the latter where the columnar structure reflects orientation selectivity. In a study by Sakai, Naya, and Miyashita's (1994), stimuli were varied parametrically by periodic 2D patterns. They found that the response of cells decreased monotonically with parameter space distance between the test stimulus and preferred pattern to which the cells were tuned. Other single-cell recordings in the IT cortex have also supported the notion of distance function for the representation of shape (Logothetis, Pauls, & Poggio, 1995; Op de Beeck, Wagemans, & Vogels, 2001), where at the individual level, most IT neurons respond maximally to the presentation of a particular shape in a stimulus group. Responses weaken for shapes located more distantly in the stimulus space.

A highly relevant study by Op de Beeck, Wagemans, & Vogels (2008) again found that neurons respond to particular shapes and display properties of distance functions, however, they also find that perceived shape differences may differ from experimentally manipulated parametric shape differences and therefore need to be independently quantified, as they are not necessarily equivalent to parametric differences in shape. They attribute the failures of spatial models of categorization that represent stimuli on only a low number of dimensions to fit behavioral data to subjects not using the dimensions explicitly included in the spatial representation. To address this failure, they cite the need for "richer" models of visual categorization that allow for the fact that visual stimuli may be represented at multiple levels.

5.0 THE GENERATIVE SIMILARITY MODEL

5.1 PREVIEW

The following section describes the details of the GS model. First, the motivation behind the model in terms of its generative structure is presented. Second, the skeleton creation for 2D shapes is presented as the precursor and critical groundwork for the GS approach. Next, the particular details of the model are presented according to the specific issues that they address, namely, the topology and metric aspects of the skeleton and the modelling of shape from the skeleton. Lastly, the computation of the similarity metric is presented and discussed.

5.2 A GENERATIVE MODEL

The primary goal of this thesis is to demonstrate that a probabilistic framework for shape representation is useful for modelling human shape similarity judgments. The basic idea is that similarity is determined by the *generative* processes that results in a shape. While "generative" can mean different things depending on the context, here “generative” describes the processes that, having occurred, produced a particular shape. The model described in this thesis allows for a computable similarity metric by providing an estimate of the likelihood that a particular shape has been *generated* by the model of the generative processes of another shape. Shapes that result from the same series of generative processes will be deemed as similar; those created from different generative processes will be gauged as dissimilar.

A generative approach to similarity suggests that even complex objects are the result of simple processes, a notion that spans disciplines, but is found most notably in biology (Thompson, 1961). This idea has also been applied to shape representation (Leyton, 1999), categorization (Feldman, 1997; Rehder, 2003) and syntax (Chomsky, 1965). Figure 5 gives an example of a hierarchical structure demonstrating the relationship between different shapes, their skeletons, and the generative history of their skeletons. Each skeleton is formed according to a generative path, where each point along the path results in subsequent transformation of the skeleton. Shapes that share a like skeletal structure, and thus a generative path, exhibit a high level of similarity. This organization, similar to biological taxonomies, reveals how similar shapes are likely to have common origins, or generative paths, that result in their observable similarity.

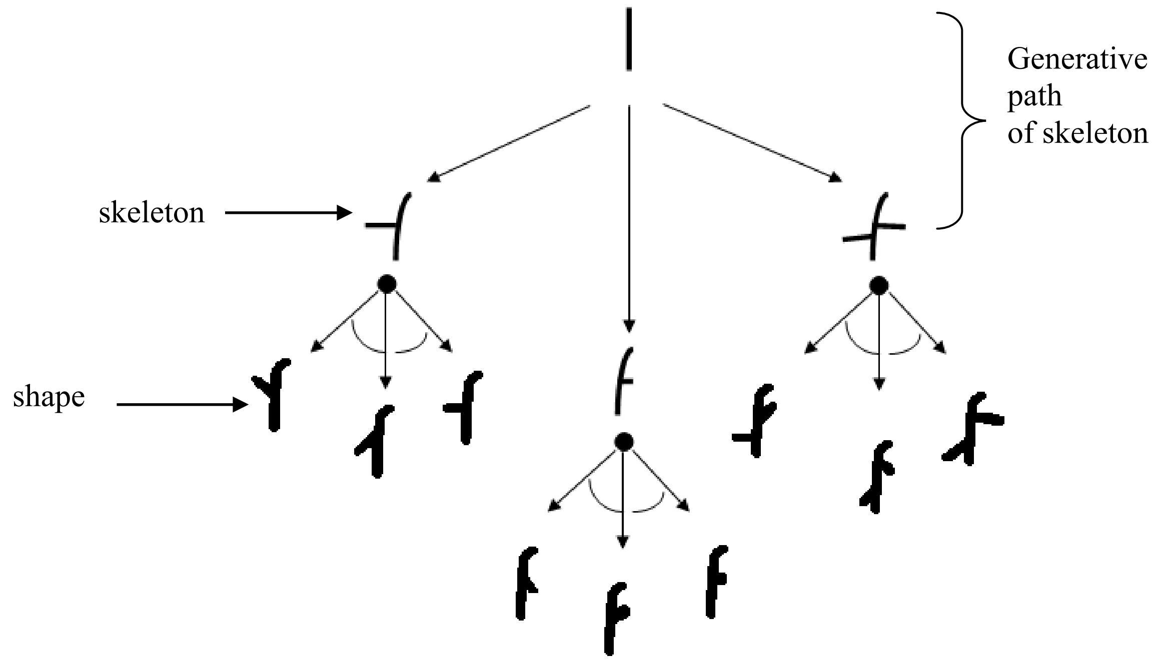

Figure 5: A hierarchical tree that depicts the relationship between shapes, their skeletons, and their generative history. Shapes with common origins are highly similar.

Kemp, Bernstein and Tenenbaum (2005) illustrate the usefulness of viewing similarity as related to generative processes by pointing out that given an object and asked to predict similar objects, there are two kinds of responses. Similar objects may result from small perturbations of one type of object or may be created from small perturbations of the *process* that generates the objects. They suggest that the latter response, the generative view, is much more feasible as the former would result in an object that is not created from a plausible generative process and would therefore lack a causal history. This idea suggests that the use of similarity is grounded in the real-world, where human cognitive processes are presumably designed to reflect natural processes.

5.3 SKELETON MODEL

The first step in evaluating the similarity between two shapes requires that each shape to be represented by its maximum *a posteriori* (MAP) skeleton, created from the skeleton formation technique described in Feldman & Singh (2006). The motivation behind this skeletal model is that shapes and their skeletons arise from generative processes. This approach lays the foundation for a larger, more comprehensive stochastic generative model that allows for comparisons of shape by virtue of their generative processes.

Feldman and Singh's model uses a probabilistic Bayesian framework to represent a shape with the "best" skeleton. This skeleton is determined from candidate skeletons by finding the one with the maximum posterior, that is, the one that best "explains" the shape. To judge the viability of candidate skeletons, the posterior

$$P(skeleton|shape) = \frac{P(shape|skeleton)P(skeleton)}{\sum_i p(shape|skeleton_i)p(skeleton_i)}$$

is calculated over all possible skeletons for a particular shape.

To create candidate skeletons, this approach begins with the medial axis transform (Blum, 1973) and then estimates the probability of that skeleton, *p(skeleton),* using a set of predefined priors, which serve as penalties involved in the skeleton formation. The first penalty derives from work on contour integration (Feldman & Singh, 2005) and penalizes increasing curvature of a skeleton axis. The skeletal axes are approximated by splines denoted by a series of *knotpoints*, creating smaller segments, over which there is a prior density that assumes successive points are generated by a density function that is centered on a zero curvature continuation of the axis with a deviation following a von Mises distribution. This means that relatively straight axes have a high probability that decreases with an increase in turning angle. Another prior penalizes increasing complexity, making complicated skeletons, those with more axes, more unlikely. Starting with the MAT for a particular shape, branches are pruned away if their explanatory power is outweighed by their complexity penalty, resulting in a skeleton that demonstrates a much decreased sensitivity to contour noise, a detrimental trait of the MAT (see Figure 6).

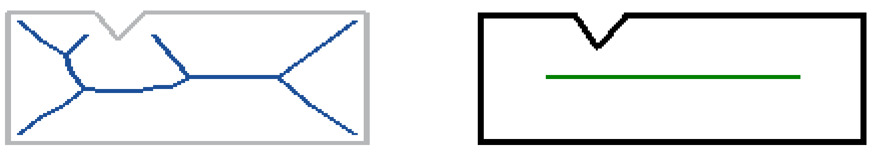

Figure 6: On the right, a medial axis transform skeleton that demonstrates sensitivity to noise. On the left is the MAP skeleton, which does not suffer from such sensitivity.

For a given a skeleton, the likelihood of a particular shape can then be found under a conditional probability density function, *p(shape|skeleton)*. For this likelihood portion of the posterior calculation, the model constructs "ribs" which radiate from the axes to the shape's contour as a means explaining contour points on the shape (see Figure 7). For each point on the skeletal axis, these ribs grow on both sides, approximately perpendicular to the axis but with random directional error (chosen from a von Mises density centered on zero). The likelihood of the shape point *p(x|skeleton)* is determined by the product of the expected rib length, directional error, and rib length error. The likelihood of the whole shape is then the product of the likelihoods of all of its points.

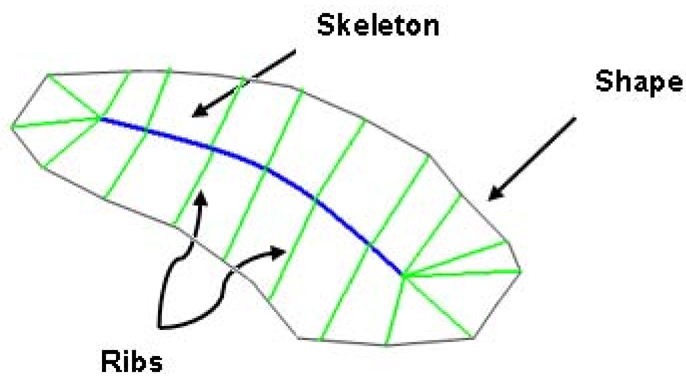

Figure 7: A depiction of the components of the MAP skeleton for a 2D shape.

The critical advantage of this model is that it generates a viable skeletal representation for shape by representing two integrated sources of information: the generative mechanism and a “noise” mechanism. It also provides skeletal components, namely a hierarchical set of axes and ribs, that can be further modelled for shape comparison tasks, as is described below.

5.4 STRUCTUAL ASPECTS OF SHAPE: THE TOPOLOGY OF THE SKELETON

In order to model the topological structure obtained through the skeletal process just described, a construct is needed that formalizes relationships exhibited in the structure of a skeleton. In linguistics, sentences can be modelled as composed of a noun phrase and a verb phrase, which again may be decomposed into simpler parts, such as a determiner and a noun and verb. The breakdown of these components signifies the hierarchical nature of the sentence and its parts. The set of allowed configurations of the parts is determined by a set of rules called a grammar. While there are many possible ways for the grammar rules to be expressed, a popular method is the use of string rewriting rules called productions.

Likewise, syntactic approaches to pattern recognition, usually referred to as rule-based grammars (Chomsky, 1957), can model complex objects by defining them in terms of their constituents; the constituents in turn are defined in terms of their subconstituents. At the bottom of this process lies a set of terminals for which no further expansion is possible. This leads to hierarchical descriptions of the objects of interest. A specific example, extracted from Fu (1989), is shown below, where the pictorial pattern of Picture P (Figure 8) can be described in terms of the hierarchical structures shown on the right of Figure 8.

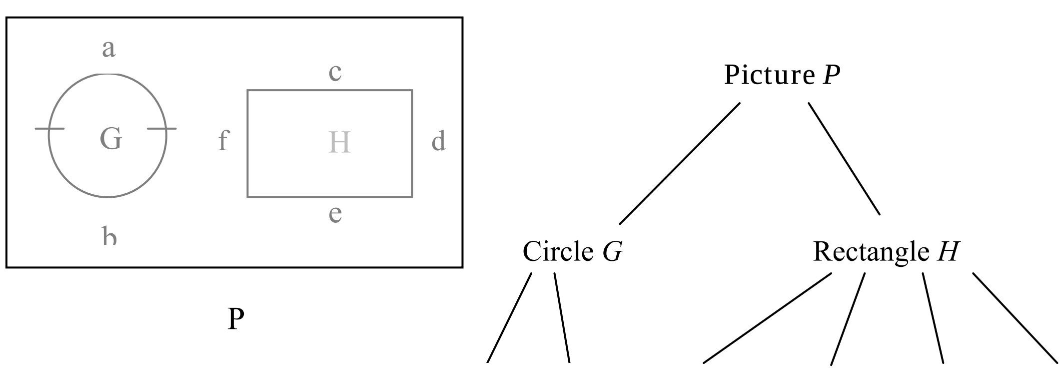

To use a syntactic approach to pattern recognition, objects must first be represented. The most common representation scheme for this type of application is strings of symbols, though arrays and graphs (Pavlidis, 1977) have also been used. Most notably, Fu (1982) developed a comprehensive *pattern description language,* influenced by Shaw's (1969) *pictorial description language*, as a primary pattern-modelling tool. The following is an example of a description language taken from Fu (1989).

Given a line representation of a triangle shown in Figure 9 (a) and the primitives shown in Figure 9 (b), the triangle may be represented by the string *abc*, using the "concatenation" operation from Fu (1974). Different sizes of the triangle are represented by the string $a^m b^m c^m$, where $m \geq 1$. An example of a larger triangle is shown in Figure 9 (d).

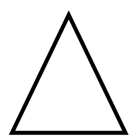

(a)

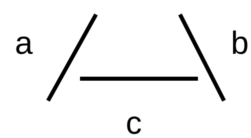

(b)

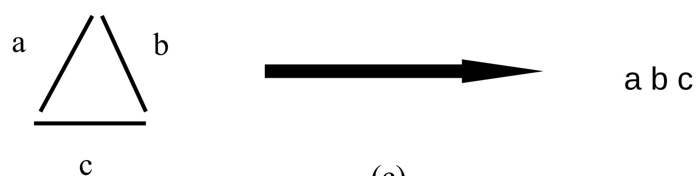

(c)

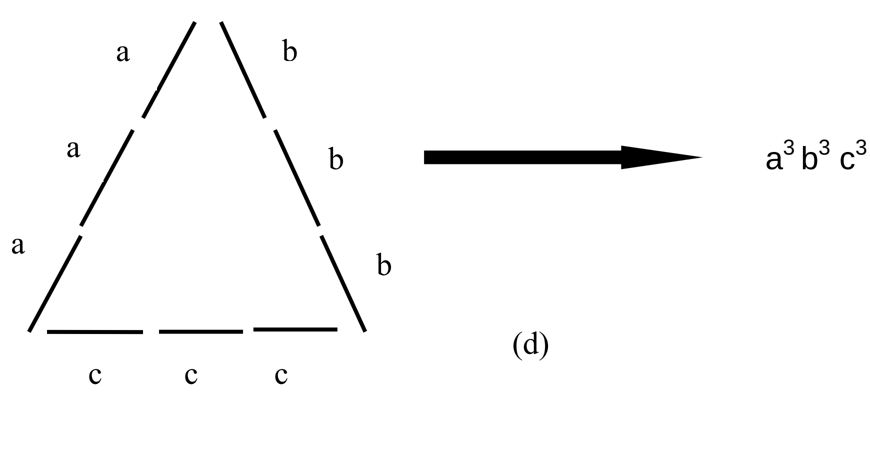

(d)

Figure 9: An example illustrating the structural or syntactic approach to shape representation (a) A line represenstation of a triangle (b) Primitives and their discrete symbol representation (c) Approximation of the triangle using primitives and the symbol string representation of the triangle (d) Approximation of a different size trinagle and its symbol strings representation

The triangle can be generated by the following grammar (Tremblay, 1975),

$$G = \langle \{S, X, Y\}, \{a,b,c\}, S, P \rangle$$

where P consists of the following production rules:

$$\begin{array}{lll} S & \rightarrow & aSXY \\ S & \rightarrow & aXY \\ YX & \rightarrow & XY \\ aX & \rightarrow & ab \\ bX & \rightarrow & bb \\ bY & \rightarrow & bc \\ cY & \rightarrow & cc \end{array}$$

Given a particular grammar and a collection of terminals, it is a natural question whether the grammar could have generated this collection of terminals. The process of determining the process by which a set of terminals was produced is known as *parsing*. Parsing is the basis of pattern recognition for syntactic methods. Typically, a set of raw data is pre-processed into a collection of terminals, which are then parsed again. Output of the parse is a list of rules that could have generated the collection of terminals, which may be singular or include multiple possibilities. The parse of the triangle in Figure 9 (d) is shown below.

$$
\begin{aligned}
S &\rightarrow aSXY \\
&\rightarrow aaSXYXY \\
&\rightarrow aaSXXYY \\
&\rightarrow aaaXYXXYY \\
&\rightarrow aaaXXYXYY \\
&\rightarrow aaaXXXYYY \\
&\rightarrow aaabXXYYY \\
&\rightarrow aaabbXYYY \\
&\rightarrow aaabbbYYY \\
&\rightarrow aaabbbcYY \\
&\rightarrow aaabbbccY \\
&\rightarrow aaabbccc
\end{aligned}
$$

When there exists more than one possible parse for a set of terminals, it is necessary to form a means of comparing or ranking these different interpretations. Probabilistic context-free grammars (PCFGs) are a modern variation to grammar creation in which probabilities are imposed on rewrite rules. Here, a context-free grammar is used to define the set of production rules with a corresponding probability distribution.

A similar scheme will be used in the work in this thesis; a novel compositional grammar will be used to define a set of rules with associated probabilities that will represent a “skeleton grammar”. A PCFG will be used to induce a probability distribution over shapes, with each shape skeleton seen as the result of a derivation according to the grammar. Though various primitives for representing the shape could be used (e.g. line segments (Fu, 1974), chaincode (Freeman, 1961), or polygonal and functional approximations (Pavlidis, 1977)), the approach in this thesis uses knotpoints as the primitives by which each axis can be represented (see Figure 10).

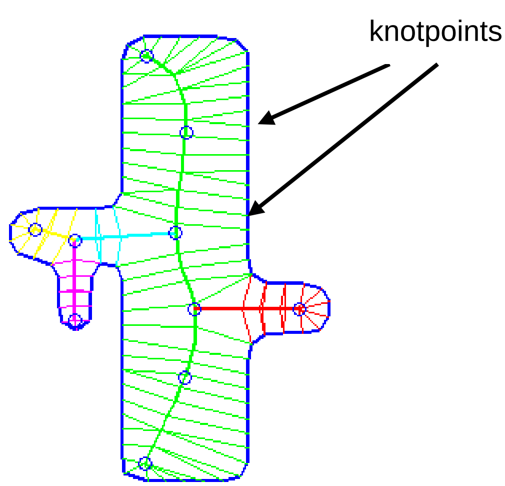

Figure 10: Example shape with its skeleton, ribs, and knotpoints, which are used as the terminals in the context free grammar model.

For a set of shapes, a general “shape grammar” is used that represent transformations that can occur to produce a given set of terminals (knotpoints) that represent each shape's skeleton (see Appendix A). Each individual rule can be seen as representing a kind of transformation. For example, the rule $K \rightarrow K\ K$ represents the replication of a knotpoint resulting in an expansion of the axis. The rule $K \rightarrow LB$ represents the creation of a left branching point on an axis. There are likewise rules that represent transformations such as right-branching and deletions of branches. As a probabilistic grammar, each rule has a corresponding probability. For example, the probability that a knotpoint on the main (root) axis of the skeleton will be right-branching could be assigned a value of 0.5. This would indicate that for each knotpoint on the main axis of a shape generated from this particular shape grammar, there is a 50% chance that it will result in a branch on the right side. Because it was not the specific goal of this thesis to incorporate meaningful priors, for example, those based on natural shape statistics (Wilder, Feldman, & Singh, 2008), the initial probabilities in the grammar were designed to be neutral, in the sense that particular production rules are not biased. For example, the production rules in the grammar are initially set with probabilities so that the occurrence of a right branch is equal to the probability of a left-branch.

Under this type of grammar there are multiple ways to model a specific shape. One approach would be to bias all of the rules of the grammar so that they resemble the entire process that could have been responsible for that shape. For example, to model the shape of a dog that has four legs branching from a torso (the root), the production rules responsible for left branching would reflect a higher probability of occurrence than that for right branching (assuming a right to left transversal down the main axis of the dog). This approach would be suitable for an iterative learning process, for example, a categorization model that iteratively creates a "definition" for a category (see section 7.1.1 for further discussion).

Another approach, and the one applied here, requires only one additional production rule (beyond the ones that are included for every shape) in the grammar. This single rule reflects the topology of the specific shape that is modelled, and in doing so, reflects a "shortcut" of the process that created that particular shape. When parsing occurs under a particular shape's grammar, the most likely "path" will be found. So, for instance, when parsing a shape similar to the one on which the grammar is modelled, it will likely use the rule that reflects the modelled shape's topology. If a very different shape is parsed, it will likely result from a process that does not utilize the topologically specific rule, as it is more likely to have been generated by a different process.

For example, the shape's skeleton on the left of Figure 11 is represented by the following knotpoints:

K1, LB1, K2, K1

where *S* stands for the starting point, *LB* stands for left-branch, *A* stands for axis, and *K* stands for a non-branching knotpoint. The number in each term indicates on which axis the knotpoint occurs. The transversal of the skeleton is from right to left and is depth-first, meaning that each encountered branch is followed until it reaches termination or another branching occurs. Under a grammar constructed for this particular shape, this series of knotpoints could have resulted from many "paths" or sequence of productions rules. Two example paths are as show below.

A: S → A1 → K1 LB1 K2 K1

B: S → A1 → K1 → K1 K1 →K1 K1 K1 → K1 LB1 K1 → K1 LB1 A2 K2 → K1 LB1 K2 K1

Path *A* shows the use of the rule that was specifically place in the grammar to represent the modelled shape's particular topology. Path *B* demonstrates how this shape could have "started from scratch", as each knotpoint is the result of a single production rule. Upon parsing this string of knotpoints, using the Natural Language Toolkit[1] developed by Loper and Bird (2002), a likelihood value for each path is found, with the first path achieving the highest likelihood. Because the probability of each step, a value that is less than or equal to one, is multiplied when the step is incorporated in the path, path *A's* higher likelihood results because it involves less production steps. If a specific grammar was created to model the shape on the right in Figure 11, then used to parse the shape on the left, which is highly dissimilar to the one on which the grammar was based, the most likely path would result from a "starting from scratch" path, as to begin with the other shape's topology would involve the too costly deleting of branches.

[1] The Natural Language Toolkit (NLTK) is a suite of Python modules distributed under the GPL open source license via nltk.org. The NLTK code supports corpus access, tokenizing, stemming, chunking, parsing, clustering, and language modeling, as well as other functionality.

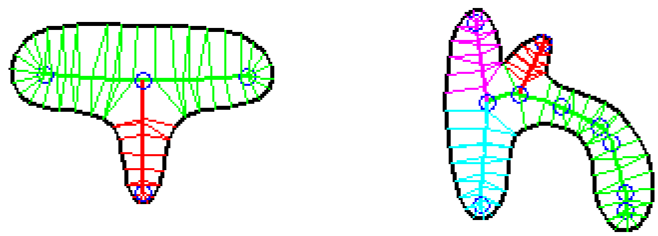

Figure 11: The shape on the left is represented by a specific derivation in the grammar. The shape on the right would not be well represented by the same derivation such that its most likely generative path would involve it "starting from scratch".

It should be noted here that this component of the model, representing shape skeletons with a grammar, is the most significant contributor to the computational complexity of the model. As the number of points (knots) used to represent the shape increases, so does the computational time for parsing each shape. Because the number of knots increases proportionally with the number of parts in a shape, the parsing of complicated shapes may take a significant amount of processing.

5.5 METRIC ASPECTS OF THE SKELETON

The axis depiction from the skeletal model represents the curvature of each specific axis as a series of angles and log ratios of line segment lengths. The angles are the measurements from each line segment to the tangent of the preceding segment. The log length ratios are formed between each pair of consecutive segments that form the axis. Both values are assumed to be normally distributed with the means equal to the modelled shape's values at each consecutive point.

The angle at which each axis "branches" from its parent is likewise represented under the assumption that a child branches from its parent according to a Normal distribution. In the GS model, the mean of this distribution is set to the value of the modelled shape's branching angle for that particular branch. The variance for the distributions representing the curvature and branching angle is set to a predefined value. Because this value is static and used across all shape comparisons, any influence it introduces is shared across comparisons and therefore does not cause any specific bias. Within the GS model, the curvature values and branching angles of a to-be-compared shape can then be assigned a probability according to the normal distribution of the modelled shape's corresponding axis and branch distributions.

5.6 METRIC ASPECTS OF THE CONTOUR: THE RIB MODEL

As explained in section 5.3, the skeleton representation produced includes a collection of ribs that explain points along the shape's contour. While the skeletal model captures many of the important properties of the shape, most significantly, the part structure, the ribs capture information about the contour at a finer level of detail. Within the GS model, the representational strategy employed for modelling the shape ribs is based on one created by Bicego and Murino (2004). Their model employs Hidden Markov Models (HMMs) for 2D shape classification by representing contours through their curvature coefficients along the boundary, allowing for the representation of a shape's boundary as a probabilistic sequence.

An HMM is a probabilistic model that represents a random sequence O=$O_1,O_2,O_{3,...,}O_T$ as the indirect observations of an underlying Markovian random sequence Q=$Q_1,Q_2,...Q_T$ (see Figure 12 for a graphical representation). An HMM λ is defined by the following components (Rabiner, 1989): a set of states, $S = \{S_1,S_2,...,S_k\}$; a transition matrix A=$\{a_{ij}\}$, where $a_{ij} \geq 0$ represents the probability of going from state S_i to state S_j; an emission matrix B = $\{b(o|S_j)\}$, the probability of an emission o from state S_j; and an initial state probability distribution $\pi = \{ \pi_i\}$, the probability of the first state $\pi_i = P[Q_1 = S_i]$.

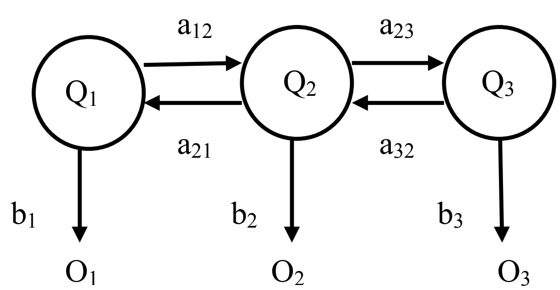

Figure 12: A graphical depiction of an HMM with states Q=$Q_1,Q_2,...Q_T$, observations O=$O_1,O_2,O_{3,...,}O_T$, transition probability matrix A=$\{a_{ij}\}$, and emission probability matrix B = $\{b(o|S_j)\}$.

With this model, given a sequence O, the *Baum-Welch* re-estimation procedure (Baum, Petri, Soules, & Weiss, 1970) can be used to determine the HMM parameters maximizing the probability $P(O| \lambda)$. To compute the probability $P(O| \lambda)$ given the sequence O and an HMM λ, the *forward-backward* procedure (Baum, 1970) is used.

The sequence of rib lengths that occurs about each axis of a shape in the GS model is captured by an HMM using the previously described HMM method. For the primary axis, the sequence of rib lengths starts from the right endpoint. For "child" branches, the sequence begins at the attachment-to-parent point. For both types of axes, the sequence of rib lengths is created by travelling around the axis until all rib lengths have been collected.

This obtained collection of rib lengths is then used to train a continuous HMM using the Baum-Welch method, where the emission probability of each state is represented by a one-dimensional Gaussian function. The number of states of each HMM is set as a function of the number of rib lengths collected (one tenth of the number of ribs), with a minimum of one and a maximum of six states. The HMM implementation used was developed by Kevin Murphy (2001) and uses a Gaussian Mixture Model clustering to initialize the emission matrix of the HMM by grouping the rib lengths into clusters and then using the Gaussian parameters to initialize the Gaussian of each state. At the end of the training, there is one HMM for each axis of the modelled shape's skeleton, so that the ribs for an entire shape are represented by the collection of HMMs for its component axes.

To produce the likelihood that one shape's ribs have arisen from another shape's model, the previously mentioned *forward-backward* procedure is run. Given the ribs from shape *A,* on which an HMM λ_A has been modelled (if shape *A* has multiple axes, λ_A will be multiple HMMs), and a shape *B* that is to be compared to *A*, this algorithm computes the probability $P(O_B| \lambda_A),$ where O_B is the sequence of rib lengths from shape

B. This probability represents the likelihood that shape *A's* model is the generative process that resulted in shape *B's* ribs.

5.7 COMPUTING SIMILARITY

Again, the goal of this thesis is to develop a comprehensive stochastic model for shape representation that allows for a computable similarity metric by comparing the processes that generate shapes. Using the shape representation previously described, an intuitive similarity metric between two shapes is defined by

$$\frac{P(ShapeA \mid ShapeModel\ B) + P(ShapeB \mid ShapeModel\ A)}{2}$$

For example, to compute the similarity for two shapes *A* and *B,* represented as a set of contour points, the first step is to find the skeleton for each shape by applying the MAP skeleton method (see section 5.3). This process results in a skeleton that is composed of a series of *n* axes with a hierarchical order consisting of one "root" and *n-1* children. Specific information about each axis of the skeleton is modelled separately, including: information about where branching off the root occurs (section 5.4), the curvature of each axis (section 5.5), the branching angle of each child axis (section 5.5), and the distribution of ribs that radiate from each axis to the contour of the shape (section 5.6). Taken together, these separate models result in the ability to compare shapes by finding the likelihood that one shape has been *generated* from another shape's model. The more similar two shapes are, in theory, the more their underlying generative processes will overlap, resulting in a higher probability for one shape to have been created from the other's generative process. To compute the similarity between two shapes, each shape is comprehensively modelled and then the probability that each shape was generated from the other's model is found and averaged.

5.8 FURTHER CONSIDERATIONS OF THE GS MODEL

An extremely important note should be made here: the GS model previously described contains many parameters, which through an optimization process, could be fit to a particular set of data. The following analysis represents the results of no effort to fit these parameters to the data. Initial parameter values were chosen either arbitrarily or to be as unbiased as possible. Presumably, better performance (for the following experimental analysis) could be achieved if data fitting through parameter optimization was incorporated.

Other important characteristics of the GS model that are not readily obvious are its invariance to rotation and scale. The model is approximately scale invariant, in that the representational theory used for shape representation is invariant, however, because implementation details possible make it slightly variant, if it was the intention to compare shapes that grossly differed in size, it would be necessary to account for their scale. The analysis presented in this thesis does not assume scale invariance and only uses shapes of approximately the same size. The GS model is not orientation invariant, in that orientation will be a factor in the similarity calculation. While this lack of invariance might be detrimental in a shape matching context, because subjects may be sensitive to orientation, it is assumed to be a significant factor in representing human similarity judgments.

6.0 EXPERIMENTS

The following four experiments were designed to test the theory behind the GS model explained in section 5.0. Each experimental analysis compares the generative similarity metric (discussed in section 5.7) with similarity judgments made by human subjects. The primary difference between the four experiments is in the specific types of 2D shapes that were presented to the subjects. Each set of shapes was chosen so as to evaluate the GS approach in a different way, allowing a full and comprehensive test of the model in a variety of settings. Much previous work on similarity within shape spaces (e.g. Cortese & Dyre, 1996) concentrates on a very narrow class of shapes, and then draws conclusions that seem strictly limited to that class. Very few, if any, psychological-based shape models effectively handle a wide variety of shapes, including both qualitative and quantitative differences, within the same theoretical framework.

Experiment 1 concentrates on shapes that show the emergence of a secondary part across the stimulus space. This examination serves to validate the importance of a shape model's ability to distinguish when noise along a shape's contour becomes meaningful enough to warrant perceptual significance. Experiment 2 and 3 build off of the work of Basri, Costa, Geiger and Jacobs (1998) who suggest that both part structure (qualitative) and metric (quantitative) differences should be considered in determining shape similarity. Experiment 2 uses shapes that vary according to their topology or part structure. This investigation tackles the difficult task of analyzing shapes that are highly perceptually variant. Experiment 3 offers a complimentary assessment of shapes that have identical topological structure, but whose metric differences, length and curvature of branches, are randomly varied. This experiment is designed to test subtle differences between shapes, those that, in general, are subjectively quite similar. The fourth experiment involves the use of shapes that have been created through a random unrelated process that intentionally deemphasizes part structure. This investigation attempts an unbiased challenge to the GS model in order to evaluate its comprehensive adequacy.

6.1 EXPERIMENT 1: PART EMERGENCE

Experiment 1 was designed to investigate the "emergence" of a part across a series of two-dimensional shapes. This emergence was created by presenting subjects with 25 shapes that varied parametrically over two dimensions, the length and width of a protrusion that arises from a curved base. For the shapes where the protrusion has lower values of length, the second "part" is virtually imperceptible, but becomes noticeable as the length increases (see Figure 13).

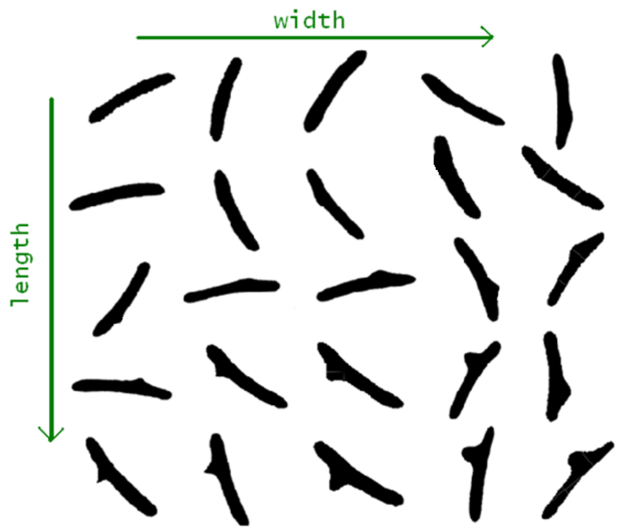

Figure 13: Experimental stimuli for experiment 1-shapes with an emerging part

This progression of shapes was designed to induce a boundary where, within this parameterized space, the protrusion qualifies as a separate part for both human subjects and the shape representation model. If the model distinguishes a difference in similarity between shapes with one part and two parts as do the subjects (without the model parameters being fit to the subject data), it will be evident by the correspondence between the similarity ratings provided by the subjects and the similarity metric produced by the model. Shapes with one part will be rated as similar to each other but dissimilar to the two part shapes.

6.1.1 Subjects:

Eleven undergraduates at Rutgers University received class credit for participation.

6.1.2 Stimuli:

Shapes were generated using the same major axis, but varied parametrically on the length and width of the secondary "part" or protrusion. The height of the protrusion varied in 12 equal unit increments starting from zero. The width of the protrusion varied in 20 percent increments, from 10 percent to 90 percent of the width of the major axis. To avoid having the subjects develop strategies focused on particular aspects of the shapes, rather than using a more holistic approach, Gaussian noise was applied to the contour of all the presented shapes. The shapes were solid black and presented on a white background. Shapes were approximately 10 degrees of visual angle in size at approximately 58cm viewing distance.

6.1.3 Procedure:

Each pair of shapes from the set of shapes displayed in Figure 13 was presented twice to subjects on an eMac computer. Each shape was presented in a random rotation on the display. Beneath the shapes a similarity scale was displayed, ranging from 1 to 7 (see Figure 14). The subjects indicated their perceived similarity of the two shapes by choosing the number corresponding to their rating with the mouse, where 7 represented "most similar" and 1 represented "least similar".

Figure 14: Example screenshot from similarity rating experiment.

6.1.4 Results

Figure 15 shows the MAP skeletons for each stimulus shape presented in Experiment 1. At some point within the parameterized space the protrusion becomes meaningful enough (in terms of the MAP skeleton procedure) to warrant a second axis in the shape's skeleton (indicated by the red boundary).

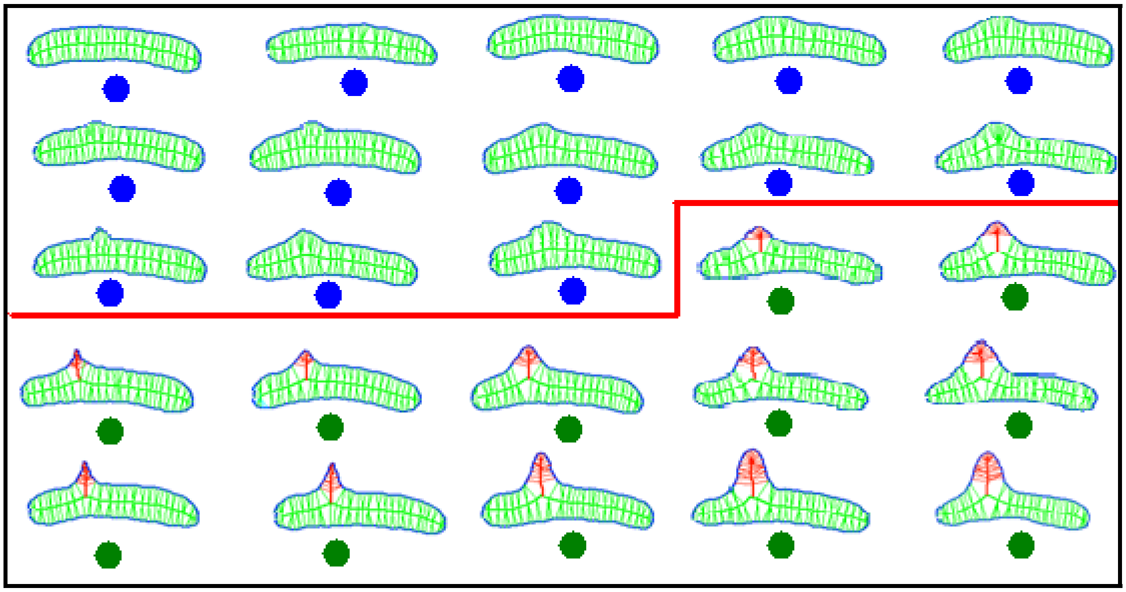

Figure 15: Experiment 1 shapes with their corresponding MAP skeletons and ribs. Color coding indicates separate axes in the MAP estimates; i.e. red ribs and axis constitute a secondary part. Red boundary indicates "gulf" in stimulus space separating one and two part shapes.

The GS approach presumes that structural aspects of shape (among others) play a key role in similarity judgments. This suggests that qualitative shape changes, as captured in the theory by topological changes to the inferred skeleton, will entail jumps in dissimilarity. To corroborate this prediction, a preliminary MDS analysis (ASCAL) was run to produce scaled inter-shape distances from the subject data, which represent the distances among shapes in "psychological space" derived from the similarity ratings. An MDS plot (see section 4.3.1 for an explanation) was created to visually represent how the subjects grouped the shapes on the basis of similarity. Figure 16 shows the MDS plot, where each shape is represented by a circle, whose color corresponds to the color circles below each shape in Figure x. Shapes grouped closely in the MDS plot indicates those shapes' high level of similarity as rated by subjects.

The MDS plot reveals a conspicuous “gulf” between two groups of shapes. This gulf corresponds to the same gulf as seen in Figure 15, between the one and two part shapes. Though only a visual aid, this plot indicates that a psychological difference exists between shapes with one part (as found by the skeletal model) and those with two.

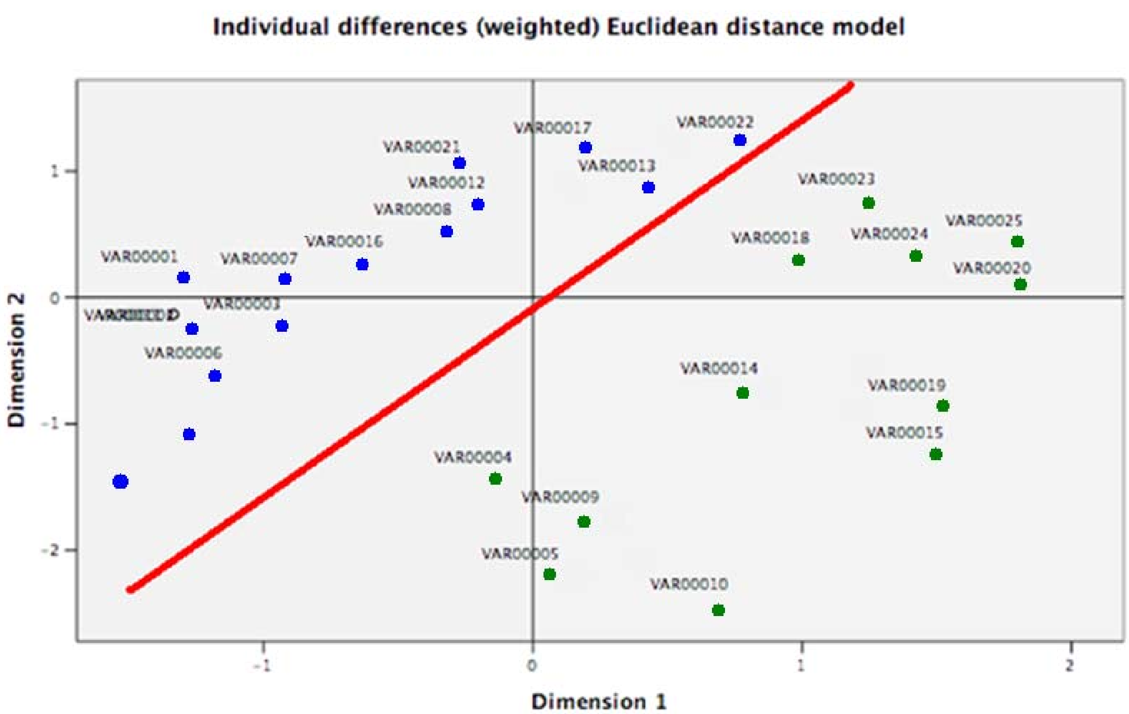

Figure 16: An MDS plot of subjects' similarity ratings for the shapes presented in Experiment 1. The red boundary represents the "gulf" that occurred between one and two part shapes determined by the skeletal representation.

To validate the “gulf” that is implied by the MDS plot, an additional analysis was run aiming to compare the magnitude of dissimilarity judgments within vs. between classes of topologically like shapes (i.e. those with one or two axis skeletons). To compensate for the actual physical differences between the shapes, each of the MDS scaled inter-shape distances from the subject data was divided by the corresponding Euclidean distance in the parameterized shape space. These distances were then separated into two groups. The first group was composed of the distances between shapes with either only one part (according to the skeletal model) and the distances between shapes with two parts (according to the skeletal model). This group can be seen as a list of distances among shapes with the same topology. The second group was composed of the distances between shapes with one and two parts, so that these distances were among shapes with differing topologies. A t-test was then used to analyze if there was a difference between the means of the two groups of distances, which resulted in a rejection of the null hypothesis that the means are equal (t=-4.346, df=245, p<.005). This result indicates that a psychological difference exists between shapes with one part (as found by the skeletal model) and those with two. The means and standard errors of the averaged distances for both one and two part shapes is displayed in Figure 17.

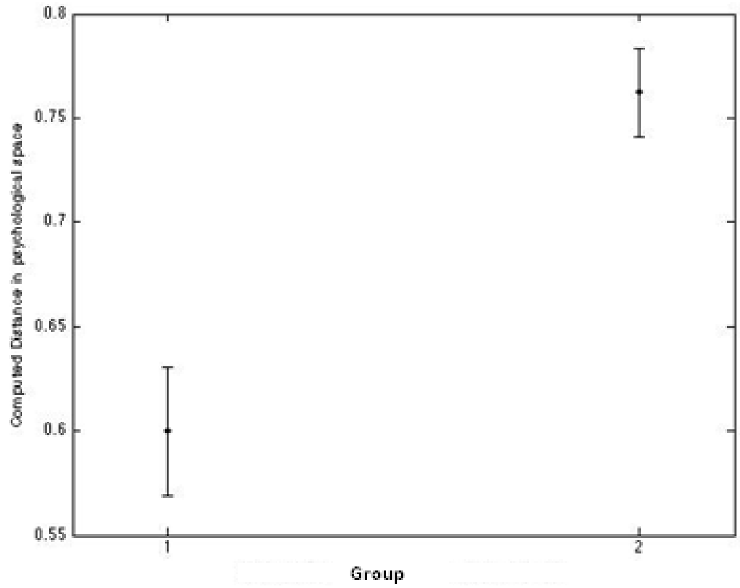

Figure 17: The means of the computed psychological distances in the shape space for two groups. Group one is made up of the distances between shapes with the same topology and group two is made up of the distances between shapes with different topologies (one part vs. two).

The similarity metric, derived from the process explained in section 5.7, was computed for each pair of shapes. Because the calculated similarities are small probabilities, the absolute value of the logarithmically scaled value was used for analysis. In order for the model and subject ratings to have a positive correlation, the subjects' ratings were inverted so that '1' became the highest value of similarity and '7' became the lowest. To determine if the model generated a similarity metric that corresponded to the subjects' similarity ratings, a linear regression model (Figure 18) was used. The model was found to adequately predict the similarity ratings of the subjects ($F(1, 623) = 80.033$, $p < .0005$).

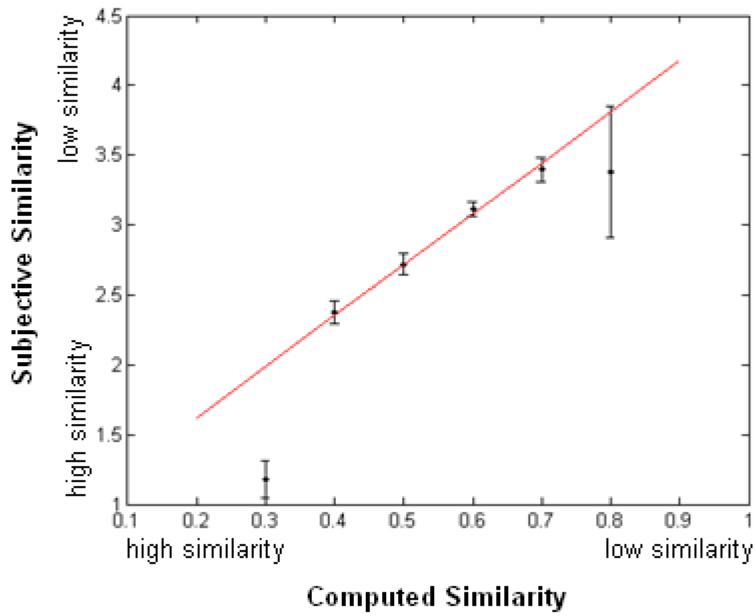

Figure 18: Regression graph of predicted and subjective similarity, with standard error bars, for all trials in Experiment 1-shapes with an emerging part.

In the graph in Figure 18, there is a noticeable nonlinearity in the data where similarity is higher. A post-hoc analysis has shown that this non-linearity pattern results from subjects' rating practices that occur when the stimuli presented is two of the exact same shape. For example, subjects may tend to keep their ratings around the middle range of the scale, but when they observe two of the same shape, they "jump" to the highest similarity rating. This 'jump' is indicative of the subjects doing something qualitatively different than the GS model. It is likely that this difference is results from the subjects performing something like identification, determining if the 2 shapes are exactly the same, and thus represents a separate type of process that the GS model does not intend to capture. If the trials where the same shape was shown twice are removed, the results remain significant (F (1, 598) = 51.122, $p < .0005$).

To forward the use of a similarity metric for further applications, a supplementary demonstration of shape matching based on the calculated similarity metric was also performed. Each shape in a corner of Figure 13 was chosen as a shape to be matched, which involved determining, out of the other 24 shapes, the shape that had the second highest similarity rating to the to-be-match shape. The first highest similarity to the to-be-matched shape is consistently itself and therefore this match is not considered. Figure 19 displays the to-be-matched shapes above their resulting matches. Results show that each shape was matched to an adjacent shape in the stimulus space.

Figure 19: The top row shows shapes that were matched to the shapes on the bottom row on the basis of the calculated similarity metric. Each match was with shapes that are adjacent in the stimulus space.

6.1.5 Discussion

From the regression analysis, it is evident that the subjective similarity is effectively predicted by the generative similarity metric produced by the model. Shapes that are judged to be dissimilar by the subjects are likewise found dissimilar according to the similarity metric. Because the skeletal model created 2-part skeletons for those shapes with a larger protrusion, according to the regression analysis, subjects are likely finding a similar distinction between shapes, where the size of the protrusion eventually determines it being deemed significant enough to warrant labelling it as a second part.

6.2 EXPERIMENT 2: SHAPES WITH TOPOLOGICAL DIFFERENCES

The second experiment was intended to test shapes with a strong part structure. The shapes used in this experiment were generated by running the grammar “forward”, where the same probabilistic rules used to parse a given shape (explained in section 5.4) were employed to generate new shapes (see Figure 20 for examples). These shapes were designed to test both metric and topological types of shape transformations, which is especially important because it emphasizes the unique strength of the GS model to handle differences in part configurations across shapes.

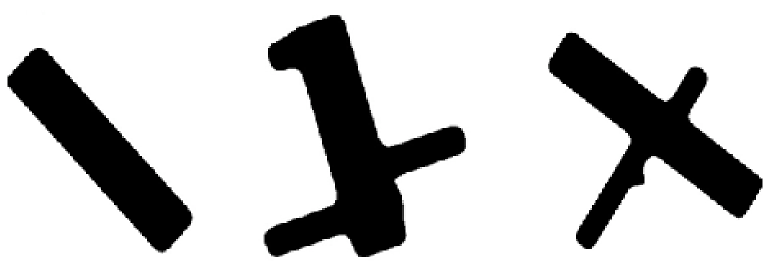

Figure 20: Example stimuli for Experiment 2-shapes with differing topology

6.2.1 Subjects

Subjects were 10 undergraduates at Rutgers University who participated for class credit.

6.2.2 Stimuli

As previously mentioned, the grammar contains production rules for the branching or topological structure of a shape. This set of 25 shapes was created by iteratively applying these rules according to the probability that they occur in generating a shape. Each shape was solid black and presented on a white background. Shapes were approximately 10 degrees of visual angle in size at approximately 58cm viewing distance.

6.2.3 Procedure

Each pair of shapes was presented twice to subjects on an eMac computer. Each shape was presented in a random rotation on the display. Displayed beneath the shapes was a similarity scale, ranging from 1 to 7 (see Figure 14). The subjects indicated their perceived similarity of the shapes by choosing the number corresponding to their rating with the mouse, where 7 represented “most similar” and 1 represented “least similar”.

6.2.4 Results

The similarity metric, derived from the process explained in section 5.7, was computed for each pair of shapes and compared to the GS model-produced similarity rating in a linear regression. Again, we find that the model accurately predicts subjects' similarity ratings ($F(1,623) = 28.71$, $p < .0005$).

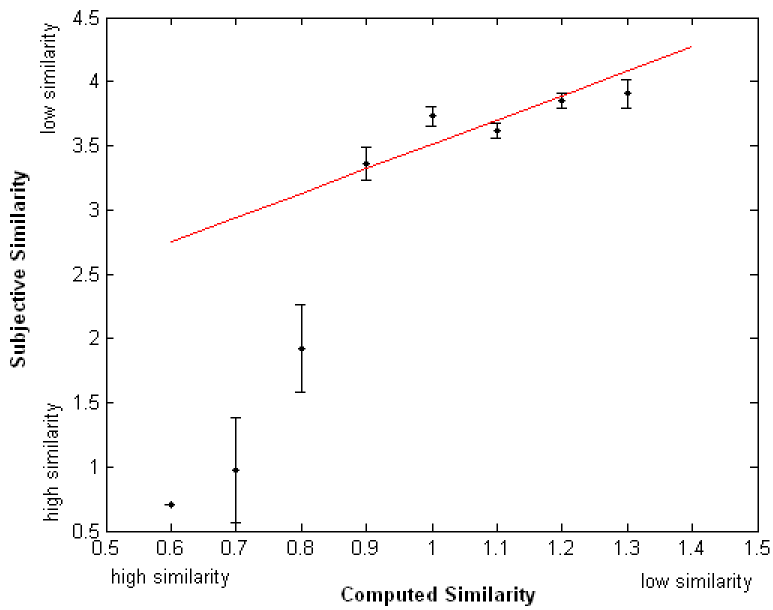

Figure 21: Regression graph of predicted and subjective similarity for all trials in Experiment 2-shapes with different topologies

Removing the trials where subjects viewed two of the same shape, because this causes the non-linearity discussed in section 6.1.4, still results in a significant regression (F (1, 598)=22.525, $p < .0005$).

Because the linear regression was run using averaged subject data, Figure 22 displays data from an individual subject. As can be seen, the exhibited pattern of the individual closely resembles that of the averaged data.

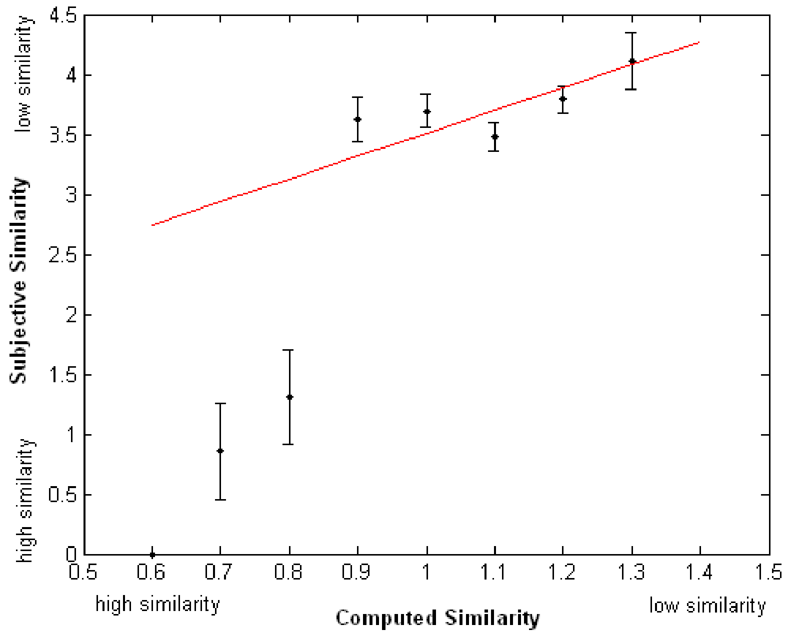

Figure 22: An individual subject's data from Experiment 2. The pattern closely resembles that of the averaged subject data used in comprehensive analysis.

As a shape matching exercise, two shapes were randomly chosen to find their most similar counterparts from the set. Figure 23 displays the two shapes along with the two shapes most similar to them according to the similarity metric (excluding self-similarity comparisons). These results correspond to the subjects’ ratings, as these shapes were also deemed as most similar to their “matched” counterparts.

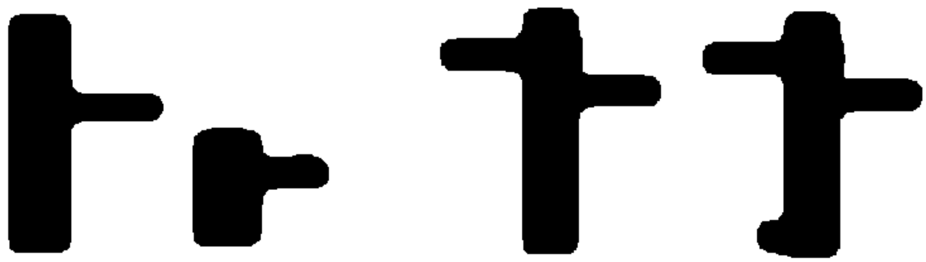

Figure 23: The shape matching results for two shapes chosen from the stimuli of Experiment 2. The two shapes on the right were found most similar as were the two on the left.

6.2.5 Discussion

Experiment 2 presented subjects with shapes that displayed a strong part structure that differed between each shape in the set. This specific variability is meant to focus on the use of a shape's topology in shape similarity comparisons by subjects. That subjects rate similarity of shapes differently purely on the basis of topological differences is not surprising as this corresponds to the adequacy of Marr and Nishihara's (1978) stick figures as simple but sufficient representations of shape. This experiment emphasizes the importance of a skeletal representation, but as can be seen in Experiment 3, this is but one source of information about a shape that a subject may utilize for similarity comparisons.

6.3 EXPERIMENT 3: SHAPES WITH METRIC DIFFERENCES

Complementary to Exp. 2, Exp. 3 was intended to test shapes that share a common topological structure, but that exhibit differences in their metric structure, namely the length and curvature of each branch of the shape (see Figure 24). These shapes can be seen as broadly belonging to the same qualitative type, but differing in their quantitative aspects of that type. Previous successes in shape representation, such as Biederman's geons (1987) or Richards and Hoffman's codons (1985) are unable to account for metric variations within a shape type. This experiment tests the GS model's capacity to represent much more subtle differences among shapes than were encountered in the previous experiments.

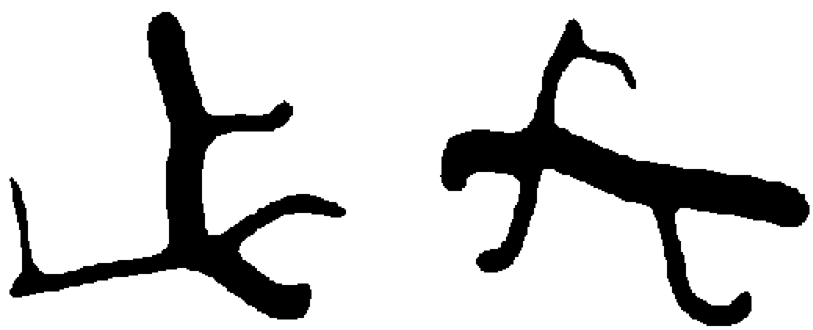

Figure 24: Experimental stimuli for Experiment 3-shapes with metric differences.

6.3.1 Subjects

Subjects were 10 undergraduates at Rutgers University who participated for class credit.

6.3.2 Stimuli

This set of 20 shapes was created by selecting an arbitrary "prototype" shape and then randomizing the branching angle, curvature and length of each branch according to a normal distribution. The width of each branch was set at 50% of its parent branch. These distributions were centered on the values corresponding to those of the prototype and used a predefined variance. Each shape was solid black and presented on a white background. Shapes were approximately 10 degrees of visual angle in size at approximately 58cm viewing distance.

6.3.3 Procedure

Each pair of shapes was presented twice to subjects on an eMac computer. Each shape was presented in a random rotation on the display. Displayed beneath the shapes was a similarity scale, ranging from 1 to 7 (see Figure 14). The subjects indicated their perceived similarity of the shapes by choosing the number that corresponded to their rating with the mouse, where 7 represented "most similar" and 1 represented "least similar".

6.3.4 Results

The similarity metric, derived from the process explained in section 5.7, was computed for each pair of shapes and compared to those rating given by subjects in a linear regression. Again, we find that the model accurately predicts subjects' similarity ratings (*F* (1, 398) = 64.461, $p < .0005$).

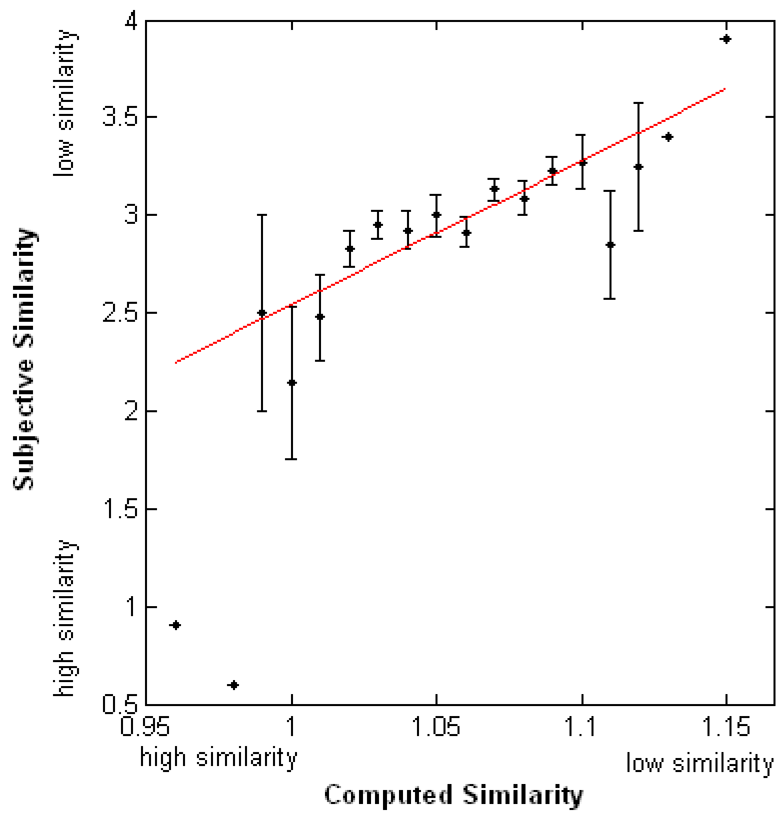

Figure 25: Regression graph of predicted and subjective similarity for all trials in Experiment 3-shapes with metric differences.

To account for the non-linearity in the subjects' responses that results from the trials with two of the same shape (causing a "jump" on the similarity scale discussed in section 6.1.4) the analysis was also run without the self-comparison trials. This too found the model as an accurate predictor (*F* (1, 378) = 18.841, $p < .0005$).

6.3.5 Discussion

This experiment attempts to further the notion that a shape possesses multiple sources of information that may be utilized for comparison tasks. Experiment 2 demonstrated that the GS model effectively accounts for the way that topological differences influence perceived shape similarity. Experiment 3 takes a complimentary approach by removing topological differences from the comparison, leaving only metric differences. Because topology is just one method of shape representation, this experiment demonstrates an additional source of variability between shapes, emphasizing that, by design, a purely skeletal representation would be insensitive to metric differences between shapes within a shape type. As the proposed model does utilize aspects of shape beyond its part structure, it is able to predict subjects' responses in this type of similarity task.

6.4 EXPERIMENT 4: ATTNEAVE SHAPES

In Exp. 2 and Exp. 3, shapes were generated via a method that in some ways mirrored the generative model underlying the skeleton estimation model, resulting in shapes with natural axes. Experiment 4 presents randomly generated shapes (see Figure 26) from a generation method specifically chosen because it does not initially generates shapes by first producing their axes, resulting in shapes without a predetermined and distinct part structure. Previous work has shown success in modelling the perceptual similarity of shapes (Cortese & Dyre, 1996) in the exact terms (amplitude and phase of a Fourier descriptor frequency) through which shapes were created. Conversely, the goal for this experiment was to challenge the GS model by using a very different non-axial shape generation model that creates shapes from a process *unrelated* to the shape representation model proposed in this thesis.

Figure 26: Example experimental stimuli for Experiment 4. The shapes were generated from Attneave's (1957) random shape algorithm. The two shapes on the right are in the same "family", as are the two shapes on the left.

6.4.1 Subjects

Subjects were 10 undergraduates at Rutgers University who participated for class credit.

6.4.2 Stimuli

The two sets of shapes (12 in each set) were created from two prototypes generated from the random shape generation model proposed by Attneave (1957). To initially create a shape, this method chooses random points about which a convex hull is formed. Random points are then removed and connected to form a prototype shape. Other shapes in the same "family" of the prototype are created by changing the position of some number of points in the prototype. This creates a set of shapes very similar to the prototype (see Figure 26). It is important to note that no axes or elongated structures are involved in this process. Each shape was solid black and presented on a white background. Shapes were approximately 10 degrees of visual angle in size at approximately 58cm viewing distance.

6.4.3 Procedure

Each pair of shapes was presented twice to subjects on an eMac computer. Each shape was presented in a random rotation on the display. Displayed beneath the shapes was a similarity scale, ranging from 1 to 7 (see Figure 14). The subjects indicated their perceived similarity of the shapes by choosing the number that corresponded to their rating number with the mouse, where 7 represented "most similar" and 1 represented "least similar".

6.4.4 Results

The similarity metric, derived from the process explained in section 5.7, was computed for each pair of shapes. A linear regression was run using the subjects' and the model's similarity ratings. The model predicts subjects' performance (F (1, 574) = 63.9, $p < .0005$)

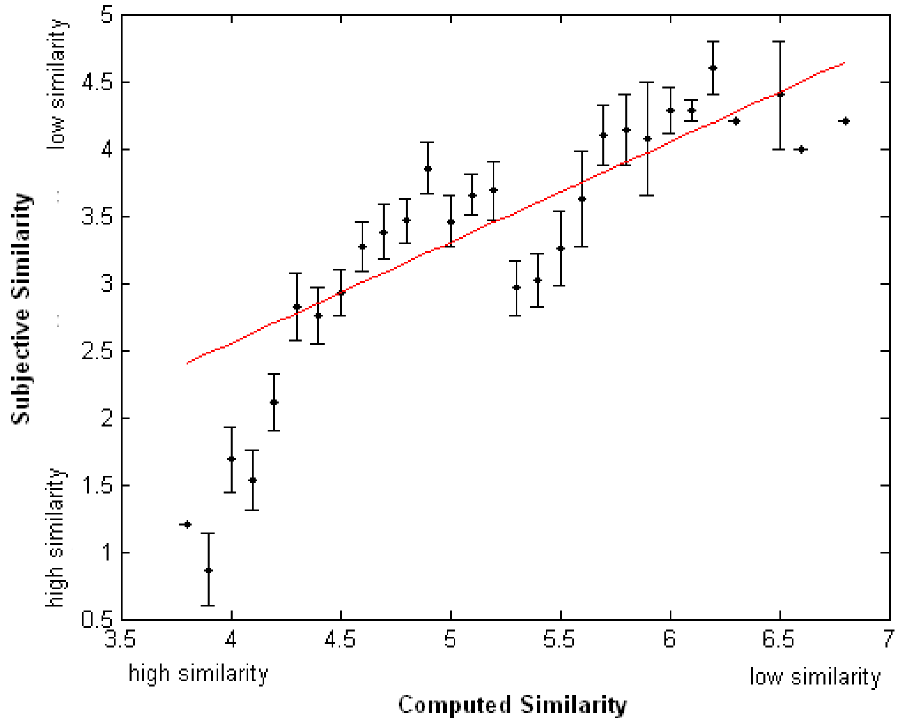

Figure 27: Regression graph of predicted and subjective similarity for all trials in Experiment 4- Attneave shapes.

As in previous experiments, there is again a noticeable nonlinearity in the data at the lower values of computed similarity. A post-hoc analysis has shown that this pattern results from a non-linearity in subjects' rating practices that occurs when the stimuli presented is two of the exact same shape (see section 6.1.4 for a discussion of why this occurs). Removing the trials where the same shape was shown twice does not affect the significance of the regression ($F(1,550) = 81.3$, $p < .0005$).

6.4.5 Discussion

The predictive value of the model in this experiment rests in its ability to model shapes that have been created from an arbitrary generative process. Where it might be expected that the presented model performs well with shapes that are specifically designed on the generative processes of the presented model (such as in Experiment 2), this particular set of shapes does not explicitly follow the same "history" of generation. That the model does perform well on these shapes implies that between shapes there exists a *common* notion of generation that sufficiently *explains* similarity between shapes.

7.0 GENERAL DISCUSSION

The studies presented in this work demonstrate that a 2D shape model implemented in a stochastic framework based on a skeletal representation is adequate for capturing critical shape information used in human shape similarity judgments. These judgments are predicted by a similarity metric that is based on the idea that similarity judgments reflect an inference of shared generative processes that represent the formation of shape through its skeletal structure. .

Experiment 1 showed that, for both subjects and the skeletal model, a part *emerges* as parametrically manipulated properties of the shape change; that is, dissimilarity undergoes a conspicuous “jump” or discontinuity when a new part emerges (see Figure 15). Because this topological change occurs as a function of a shape’s contour properties, this experiment emphasizes the need for a shape representation theory that incorporates both qualitative and quantitative properties. The importance of topological representation, in general, is demonstrated in Experiment 2. Here, the GS model was able to accurately account for the dissimilarity induced by gross changes in the shapes’ part structure. Experiment 3 revealed the role of the differences in shapes’ contours in absence of a disparity in topological structure. This experiment validated the rib and curvature components of the GS model, and proved their integral responsibility in representing subtle differences between shapes.

Experiment 4 provided the most comprehensive challenge to the GS model by confronting it with shapes that were formed using an independent method, one which did not intentionally introduce axial components. This investigation substantiates the use of the GS model in a general sense, proving that the components of the model adequately capture general shape qualities utilized by subjects in similarity judgments.

Taken as a whole, these experiments speak to the sufficiency of a similarity metric based on the idea that similarity results from generative processes. Shapes that share generative processes, those with similar shape models, are found to be perceptually similar. The accurate representation of these processes allows for a similarity metric that adequately predicts human shape similarity judgments. It should be noted that the presented model is but one method of implementation of the general theory that similarity arises from shared generative processes.

7.1.1 Categorization

By providing a shape-based similarity metric, the GS model provides a bridge between models of categorization and models of shape representation used for psychological inquiry (as opposed to models specifically used for shape matching). Visual processing often involves analyzing a shape that is not exactly like any shape that has been encountered before, which is why categorization is necessary. Categorization requires the processing of similarity between the stimulus object and representations of other previously encountered objects store in memory.

The most cited models of categorization describe the process in two steps (e.g. Kruschke, 1992; Nosofsky, 1984; Smith and Medin, 1981). First is the assumption of a psychological similarity space, with objects represented by points or regions in a space so that perceived similarity is reflected by spatial proximity. Second in this process, regions of this space are associated with particular categories. Thus, investigations of similarity are an important precursor to work in modelling categorization behavior. However, conventional treatments of similarity in psychology do not generally provide a method for its computation by which it can actually be computed from the stimuli.

An independent measure of similarity allows for an enlightened investigation of how categories are formed. The common view is that categories arise from regularity in the world (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). In light of the present work, the source of this regularity can be explained as the shared generative processes that result in similarity between objects of the same category. In representing these generative processes, the similarity metric from the GS model provides insight about where and when "gulfs" (such as those reported in Exp. 1) occur, which are indicative of the partitioning of the stimulus space into regions that map to category boundaries. This formation can also be seen through an iterative process of stimulus encounters, where, as shapes are observed, their relative similarity and related generative processes create a "definition" of the category.

Op De Beeck, Wagemans, and Vogels (2003) suggest that a fault in traditional categorization models is that they do not differentiate between perceptual and decisional processes. They find that the "attention parameter" present in many common categorization models may account for both changes in perception of stimuli and in decision strategies. Rather than having the attention parameter serve in both functions, the GS model would also allow for the "weighting" of particular dimensions in the perceptual process that the attention parameter has hereto provided. With the use of the GS model, perceptual bias of particular shape dimensions may be accounted for independent of the attention parameter, which then could serve only in a decision making capacity.

7.1.2 Additional experimental findings

Several subject behaviors were noted in the experiments that deserve further investigation. Unreported MDS plots of shapes similar to those in Experiment 2, those with intentionally different topologies, suggests that subjects make "clusters" of shapes based on attributes that are not represented by our shape model. Specifically, it is likely that subjects recognize both the symmetry and "special" placement of parts. Figure 28 shows an example of three shapes that subjects might place in different clusters in psychological space. Figure 28 (a) is separated from (b) because it demonstrates a symmetric quality about the vertical axis. This symmetry results in the perception of shape (a) having 2 parts (with one that continues through the main axis), rather than what in shape (b) would be considered as 3 parts. Shape (c) also exhibits this symmetric property but is further differentiated from shape (a) because its second part occurs at the bottom of the shape, somewhat like a foot or base, a placement which affords its distinctiveness. As the shape model presented in this thesis does not specifically model symmetry or any type of "special" placement of parts, it lacks the ability to capture the influence that these properties have on similarity judgments. Recognition of such "special" arrangements could, in principle, be added to the generative model, potentially rectifying these problems, but this possibility has not yet been tested.

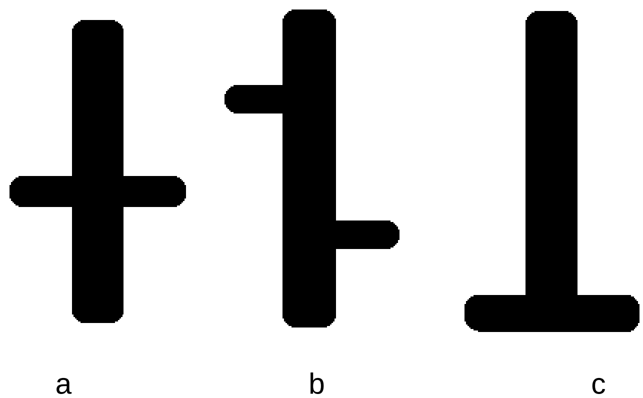

Figure 28: Example shapes that demonstrated qualities that subjects may utilize that are not included in the present version of the GS model. The shape on the left demonstrates a symmetry that the shape *b* lacks. Shape *c* demonstrates a part's "special" placement.

8.0 FUTURE WORK

This thesis represents the beginning of a comprehensive shape modelling and categorization project. As such, there are several fruitful directions that are logical next steps to this work.

As a theoretical extension, it would be valuable to relate the GS model's similarity metric to shape categorizations studies, wherein it could be employed in traditional categorization models. As previously discussed, most prominent categorization models utilize similarity for "dividing up" the stimulus space. An assumption in this type of application is that the salient dimensions of the space are known. With the GS model, however, exact dimensions are not assumed, making it a robust precursor for similarity-based categorization models that would allow for analysis of relatively complicated shapes.

It would also be of great interest to conduct categorization experiments of shape that manipulated category variability, addressing the issue brought up in section 2.0 (Figure 2). Here it could be determined whether the model accurately reflects how subjects represent category variability, for instance, by creating the need to adjust the probabilistic priors.

While section 6.0 presented some elementary shape matching exercises, it also seems important to further evaluate the GS model with this traditional computer science approach on a more complicated shape matching task. For example, with a given database of shapes, the similarity metric would be used to determine which shape a to-be-matched shape is most similar to in order to assign a classification label. The model's performance could then be compared to other leading shape classification models. This inherent classification approach also leads to a more interesting investigation- the clustering of shapes within a database on the basis of similarity. These groupings could be used to form taxonomies, which could then be compared to natural and scientific taxonomies, such as those formed from animals and plants.

An additional necessary evaluation involves determining the scalability of the GS model. With the relatively simple shapes presented here, running the model requires computation time that is negligible; however, this time increases dramatically with increasing complexity of the shapes that are compared (see section 5.4). To process much more complicated shapes will require the development of either more efficient algorithms or a more concise method for representing the shape skeletons.

8.1.1 Similarity and abstraction

Feldman (1997) argued that objects created by the same operations are likely considered to be in the same perceptual category. These shared processes create a categorical hypothesis in the same sense as the generative model of shape proposed in this thesis. The process that the GS model uses for the computation of similarity between shapes can be seen as exploiting the lattice structure described in Feldman (1997).

Lattices are partially ordered sets in which any two elements *x* and *y* share a *supremum*, the elements' least upper bound or *join* (denoted by $x \vee y$), and an *infimum*, the elements' greatest lower bound or *meet* (denoted by $x \wedge y$), of all finite subsets (Davey & Priestley, 2002). Lattices encode the algebraic behavior of the entailment relation and such basic logical connectives as conjunction and disjunction, which results in an adequate *algebraic semantics* for a variety of logical systems. Traditional applications of lattices vary, where they are currently most used in areas of computer science, social science, and operations research.

Figure 5 (in section 5.2) displayed a hierarchical structure representing possible generative paths of skeleton formation. Figure 29 shows a similar diagram, a lattice, which exhibits both a *join* and *meet* in a hierarchical arrangement of skeletal generation. Each node in the lattice can be seen as a skeletal model. Ascending the diagram provides the generative history, or a more abstract version, of the skeleton. Descending along a path explains the generation as it occurred, resulting in more constrained models. Each node inherits all properties of the nodes in its generative history. At point *d* a *meet* occurs, where the skeleton has two possible generative histories, or, perhaps more accurately, two different *sequences* in which the same generative operations could have occurred. Point *a* represents the join, or the common model, from which all skeletons in this diagram are derived.

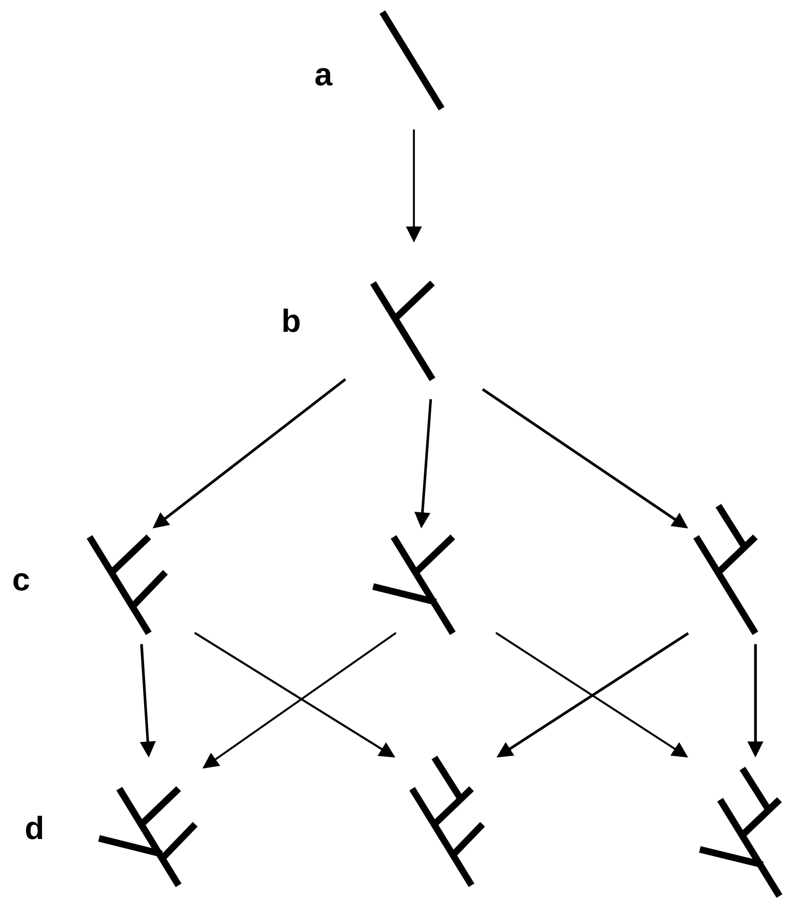

Figure 29: Example of a lattice structure that represents skeleton generation

The GS model takes advantage of this lattice structure by evaluating the paths over which a skeleton (and shape) can be generated. The similarity metric computed from the GS model is the probability that two shapes share a common generative path. If a shape, *x*, is better explained by a path that is different from the one that best explains shape *y*, then shape *x* will best be explained "from scratch" (not using shape *y's* model) and the shapes will likely be found as dissimilar. If the best explanations of both shapes fall along the same path, then the shapes will be found as similar.

In terms of the lattice configuration, the similarity metric can be expressed as

$$similarity(x, y) \sim \frac{p(x|x \vee y) \times p(y|x \vee y)}{2}$$

A more abstract similarity calculation that the lattice structure provides is the use of *any* common model $x \vee y$ (a *join*) instead of using each specific shape's model that the GS model uses (see section 5.0). Furthering exploiting the lattice structure, it would be a valuable pursuit to evaluate the similarity between two shapes on the basis of all of their shared common mid-level models (or *joins*) that occur within the lattice. This conception of similarity lends itself to the induction of category labels from observations (Richards, Feldman, & Jepson, 1992), where the lattice serves as a hypothesis "tester" for viewers seeking to align observed objects by aligning their generative processes.

9.0 CONCLUSION

In summary, this thesis makes the case that a generative approach to similarity, coupled with a suitably rich skeletal representation of shape, can provide an effective and psychologically valid similarity metric. Considering similarity as a function of an object's generative processes serves to bridge current disparate cognitive investigations of shape representation and shape categorization. This treatment of similarity offers many rich directions for both the improvement of the GS model and towards the enhancement of other similarity-based applications.

APPENDIX A

[GENERAL SHAPE GRAMMAR]

The following production rules represent the *general* shape grammar that was used in the similarity calculations presented in this thesis. To represent a particular shape, an additional rule was added to this grammar. The number to the right of each production rule is the probability associated with that rule.

S → A1 [1.0]
A1 → K1 [1.0]
A2 → K2 [1.0]
A3 → K3 [1.0]
A4 → K4 [1.0]
A5 → K5 [1.0]
A6 → K6 [1.0]
A7 → K7 [1.0]
LB1 → A2 [0.499995] | K1[0.0001] | 'LB1' [0.499995]
LB2 → A3 [0.499995] | K2[0.0001] | 'LB2' [0.499995]
LB3 → A4 [0.499995] | K3 [0.0001] | 'LB3' [0.499995]
LB4 → A5 [0.499995] | K4 [0.0001] | 'LB4' [0.499995]
LB5 → A6 [0.499995] | K5 [0.0001] | 'LB5' [0.499995]

LB6 → A7 [0.499995] | K6 [0.0001] | 'LB6' [0.499995]

RB1 → A2 [0.499995] | K1 [0.0001] | 'RB1' [0.499995]

RB2 → A3 [0.499995] | K2 [0.0001] | 'RB2' [0.499995]

RB3 → A4 [0.499995] | K3 [0.0001] | 'RB3' [0.499995]

RB4 → A5 [0.499995] | K4 [0.0001] | 'RB4' [0.499995]

RB5 → A6 [0.499995] | K5 [0.0001] | 'RB5' [0.499995]

RB6 → A7 [0.499995] | K6 [0.0001] | 'RB6' [0.499995]

K1 → K1 K1 [0.14] | RB1 [0.08] | LB1 [0.08] | RB2 [0.07] | LB2 [0.07] | RB3 [0.06] | LB3 [0.06] | RB4 [0.05] | LB4 [0.05] | RB5 [0.04] | LB5 [0.04] | RB6 [0.03] | LB6 [0.03] | 'K1' [0.2]

K2 → K2 K2 [0.12] | RB2 [0.08] | LB2 [0.08] | RB3 [0.07] | LB3 [0.07] | RB4 [0.06] | LB4 [0.06] | RB5 [0.05] | LB5 [0.05] | RB6 [0.04] | LB6 [0.04] |'K2' [0.28]

K3 → K3 K3 [0.10]| RB3 [0.08] | LB3 [0.08] | RB4 [0.07] | LB4 [0.07] | RB5 [0.06] | LB5 [0.06] | RB6 [0.05] | LB6 [0.05] | 'K3' [0.38]

K4 → K4 K4 [0.08] | RB4 [0.08] | LB4 [0.08] | RB5 [0.07] | LB5 [0.07] | RB6 [0.06] | LB6 [0.06] | 'K4' [0.5]

K5 → K5 K5 [0.06] | RB5 [0.08] | LB5 [0.08] | RB6 [0.07] | LB6 [0.07] | 'K5' [0.64]

K6 → K6 K6 [0.04] | RB6 [0.08] | LB6 [0.08] | 'K6' [0.8]

K7 → K7 K7 [0.02] | 'K7' [0.98]

BIBLIOGRAPHY

Alt, H., Behrends, B. & Blomer, J. (1995). Approximate matching of polygonal shapes. *Annals of Mathematics and Artificial Intelligence*, 251–265.

Alt, H., Fuchs, U., Rote, G. & Weber, G. (1996). Matching convex shapes with respect to the symmetric difference. In *Algorithms ESA '96, Proceedings of the 4th Annual European Symposium on Algorithms, Barcelona, Spain, September 1996*, (pp 320–333).

Ashby, F., & Lee, W. (1991). Predicting similarity and categorization from identification. *Journal of Experimental Psychology: General*, *120*, 150-172.

Ashby, F. & Perrin, N. (1988). Toward a unified theory of similarity and recognition. *Psychological Review, 95,* 124-150.

Ashby, F. & Townsend, J. (1986). Varieties of perceptual independence. *Psychological Review, 93*, 154-179.

Attneave, F. (1950). Dimensions of similarity. *American Journal of Psychology, 63*. 516–556.

Attneave, F. (1957). Physical determinants of the judged complexity of shapes. *Journal of Experimental Psychology, 53(4)*, 221-227

Bai, X. & Latecki, L. (2008). Path Similarity Skeleton Graph Matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 30(7).*

Ballard, D. (1981). Generalized Hough transform to detect arbitrary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 13(2),* 111–122.

Bandera, A., Urdiales, C., Arrebola, F. & Sandoval, F. (1999). 2D object recognition based on curvature functions obtained from local histograms of the contour chain code. *Pattern Recognition Letters, 20(1),* 49-55.

Bareiss, R. & King, J. (1989) Similarity assessment in a case-based reasoning. *Proceedings of a case-based reasoning workshop.* California: Morgan Kaufmann.

Basri, R., Costa, L., Geiger, D. & Jacobs, D. (1998). Determining the similarity of deformable shapes. *Vision Research, 38,* 2365–2385.

Baum, L. (1970). An Inequality and Associated Maximization Technique in Statistical Estimation for Probabilistic Functions of Markov Processes. *Inequality, 3,* 1-8.

Baum, L., Petrie, T., Soules, G. & Weiss, N. (1970). A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains. *Annals of Mathematical Statistics, 41(1)*, 164-171.

Belongie, S., Malik, J. & Puzicha, J. (2002). Shape Matching and Object Recognition Using
Shape Contexts. *IEEE Pattern Analysis and Machine Intelligence, 24(24),* 509-522.

Bicego, M. & Murino, V. (2004). Investigating Hidden Markov Models' Capabilities in 2D Shape Classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *26(3),* 281-286.

Biederman, I. (1987). Recognition-by-Components: a Theory of Human Image Understanding. *Psychological Review*, *94(2),* 115-147.

Biederman, I. & Ju, G. (1988). Surface versus edge-based determinants of visual recognition. *Cognitive Psychology, 20,* 38-64.

Blum, H. (1973). Biological shape and visual science, part I. *Journal of Theoretical Biology 38,* 205-287.

Carlsson, S. (1999). Order structure, correspondence and shape based categories. *International Workshop on Shape, Contour and Grouping, Springer Lecture Notes in Computer Science*, 1681-1999.

Chater, N. & Hahn, U. (1997). Representational distortion, similarity, and the universal law of generalization. Proceedings *of the Interdisciplinary Workshop on Similarity and Categorization, SimCat97* (pp. 31-36). Edinburgh: Department of Artificial Intelligence, University of Edinburgh.

Chomsky, N. (1957). *Syntactic structures (Janua linguarum).* The Hague: Mouton.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge: The MIT Press.

Clark, A. (1993). *Sensory qualities.* Clarendon Press: Oxford.

Coombs, C. (1964). *A theory of data*. New York: Wiley.

Cortese, J. & Dyre, B. (1996). Perceptual similarity of shapes generated from Fourier descriptors. *Journal of Experimental Psychology: Human Perception and Performance, 22(1),* 133-143.

Costa, L. & Cesar, T. (2001). *Shape Analysis and Classification: Theory and Practice*. Boca Raton: CR Press.

Cutzu, F. & Edelman, S. (1996). Faithful representation of similarities among three-dimensional shapes in human vision. *Proceedings of the National Academy of Science, 93,* 12046-12050.

Davey, B. & Priestley, H. (2002). *Introduction to Lattices and Order.* (2nd. ed.). Cambridge: Cambridge University Press.

Drucker, D. & Aguirre, G. (2008). Integral versus separable perceptual dimensional pairs are reflected in conjoint versus independent neural populations [Abstract]. *Journal of Vision*, *8*(6), 40.

Edelman, S. (1999). *Representation and Recognition in Vision*. Cambridge, MA: MIT Press.

Feldman, J. (1997). The structure of perceptual categories. *Journal of Mathematical Psychology,* 41, 145-170.

Feldman, J. & Richards, W. (1998). Mapping the mental space of rectangles. *Perception,* 27, 1191-1202.

Feldman, J. & Singh, M. (2005). Information along curves and closed contours. *Psychological Review, 112(1),* 243-252.

Feldman, J. & Singh, M. (2006). Bayesian estimation of the shape skeleton. *PNAS, 103(47),* 18014–18019.

Freeman, H. (1961). On the encoding of arbitrary geometric configurations. *IEEE Transactions on Electronics and Computing, (10)*, 260–268.

Frenkel, M. & Basri,R. (2003). Curve matching using the fast marching method. In A. Rangarajan, M. Figueiredo & J. Zerubia (eds.). *Lecture Notes in Computer Science, vol. 2683*. Springer.

Fu. K. (1974). *Syntactic Methods in Pattern Recognition.* New York: Academic Press.

Fu, K. (1982). *Syntactic Pattern Recognition and Applications*. Upper Saddle River, NJ: Prentice-Hall.

Gdalyahu,Y. & Weinshall, D. (1999). Flexible syntactic matching of curves and its application to automatic hierarchical classification of silhouettes. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 21(12)*, 1312-1338.

Geiger, D., Gupta, A., Costa, L. & Vlontzos, J. (1995). Dynamic-programming for detecting, tracking, and matching deformable contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 17(3),* 294–302.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). Cambridge: Cambridge University Press.

Gick, M. & Holyoak, K. (1980). Analogical problem-solving. *Cognitive Psychology, 12*, 306-255.

Goldstone, R. & Medin, D. (1994). Time course of comparison. *Journal of Experimental Psychology: Language, Memory and Cognition, 20*, 29-50.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (ed.), *Problems and Projects* (pp. 437-447). Indianapolis and New York: Bobbs-Merrill.

Hahn, U. & Chater, N. (1997). Concepts and similarity, In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts and categories* (pp. 43-92). Hove: Psychology Press.

Hoffman, D. & Richards, W. (1984). Parts of recognition. *Cognition*, *18,* 65-96.

Hume, D. (1975). *Enquiries concerning human understanding and concerning the principles of morals*. Oxford: Oxford University.

Kass, M., Witkin, A. & Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision, 1(4),* 321–331.

Kemp, C., Bernstein, A. & Tenenbaum, J. (2005). A generative theory of similarity. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society.*

Khotanzad, A. & Hong,Y. (1990). Invariant image recognition by Zernike moments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *12(5),* 489–497.

Kruschke, J. (1992). ALCOVE: An exemplar-based connectionist model of category learning.
Psychological Review, 99, 22-44.

Latecki, L. & Lakaemper, R. (1999). Convexity Rule for Shape Decomposition Based on
Discrete Contour Evolution. *Computer Vision and Image Understanding 73,* 441-454.

Latecki, L. & Lakaemper, R. (2000). Shape Similarity Measure Based on Correspondence
of Visual Parts. *IEEE Pattern Analysis and Machine Intelligence, 22,* 1185-1219.

Leyton, M. (1999). *Symmetry, Causality, Mind*. Cambridge MA: MIT Press.

Li., M, & Vintanyi, P. (1997). *An introduction to Kolmogorov complexity and its applications*, (2 ed.). New York: Springer-Verlag.

Lin, C. & Chellappa, R. (1987). Classification of partial 2-d shapes using Fourier descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 9(5),* 686–690.

Ling, H. & Jacobs, D. (2005). Using the Inner-Distance for Classification of Articulated Shapes. *IEEE Conference on Computer Vision and Pattern Recognition, 2,* 719-726.

Ling, H. & Jacobs, D. (2007). Shape Classification Using the Inner-Distance, *IEEE Transactions on Pattern Analysis and Machine Intelligence, 29(2),* 286-299.

Liu, H. & Srinath, M. (1990). Partial shape classification using contour matching in distance transformation. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 12(11)*, 1072-1079.

Lockhead, G. (1966). Effects of dimensional redundancy on visual discrimination. *Journal of Experimental Psychology, 72,* 95-104.

Logothetis, N., Pauls, J. & Poggio, T. (1995). Shape Representation in the Inferior Temporal Cortex of Monkeys. *Current Biology, 5,* 552-563.

Loper, E. & Bird, S. (2002). NLTK: The Natural Language Toolkit. In *Proceedings of the ACL Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*, 62–69.

Loncaric, S. (1998). A survey of shape analysis techniques. *Pattern Recognition*, *31(8),* 983–1001.

Markman, B. & Gentner, D. (1993). Splitting the differences: A structural alignment view of siiarlity. *Journal of Memory and Language, 32 (4*), 518-535.

Marr, D. & Nishihara, H. (1978). *Proceedings Royal Society of London, 200(B),* 269–294.

Mokhtarian, F., Abbasi, S. & Kittler, J. (1996). Robust and Efficient Shape Indexing through Curvature scale space. *British Machine Vision Conference.*

Murphy, K. (2001). *The Bayes Net Toolbox for MATLAB*. (Technical Report). Berkeley, CA: University of California at Berkeley.

Nosofsky, R. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 104-114.

Op de Beeck, H., Wagemans, J. & Vogels, R. (2001). Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nature Neuroscience, 4,* 1244-1252.

Op de Beeck, H., Wagemans, J. & Vogels, R. (2003). The effect of category learning on the representation of shape: Dimensions can be biased, but not differentiated. *Journal of Experimental Psychology: General, 132,* 491-511.

Op de Beeck, H., Wagemans, J. & Vogels, R. (2008). The representation of perceived shape similarity and its role for category learning in monkeys: A modeling study. *Vision Research, 48(4),* 598-610.

Pavlidis, T. (1978). Survey: A review of algorithms for shape analysis. *Computing Graphics and Image Processing, 7*, 243-258.

Pelillo, M., Siddiqi, K. & Zucker, S. (1999). Matching hierarchical structures using association graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 21(11),* 1105–1120.

Pizer, S., Fritsch, D., Yushkevich, P., Johnson, V. & Chaney, E. (1999). Segmentation, Registration, and Measurement of Shape Variation via Image Object Shape. *Transactions in Medical Imaging 18(10),* 851-865.

Posner, M., & Keele, S. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*, 353-363.

Pouget, A. & Snyder, L. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience. 3*, 1192-1198.

Quine, W. (1977). Natural kinds. In S.P. Schwarts (ed.). *Naming, necessity, and natural kinds*. Ithaca, NY: Cornell University Press.

Rabiner, L. (1989). A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of IEEE, 77(2),* 257-286.

Rehder, B. (2003). A causal-model theory of conceptual representation and categorization. *Journal Experimental Psychology : Learning Memory and Cognition*, *29*, 1141-1159.

Richards, W., Feldman, J. & Jepson, A. (1992). In D. Hogg D & R. Boyle (eds.). From features to perceptual categories. *British Machine Vision Conference.* Springer-Verlag, 99-108.

Richards, W. & Hoffman, D. (1985). Codon constraints on closed 2D shapes. *Computer Vision, Graphics, and Image Processing, 31(3),* 265-281.

Rips, L. (1989). Similarity, typicality, and categorization. In S. Vosnidou & A. Ortony (eds.), *Similarity and analogical reasoning.* (pp. 21-59). New York: Cambridge University Press.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (eds.). *Cognition and Categorization.* (pp.27-48). Hillsdale, N.J.: Erlbaum Associates.

Rosch, E., Mervis, C., Gray, W., Johnson, D. & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382-439.

Saiki, J., & Hummel, J. E. (1998). Connectedness and the integration of parts with relations
in shape perception. *Journal of Experimental Psychology-Human Perception and Performance, 24*(1), 227-251.

Sakai, K., Naya Y. & Miyashita, Y. (1994). Neuronal tuning and associative mechanisms in form representation. *Learning and Memory, 1,* 83-105.

Sclaroff, S. & Pentland, A. (1995). Modal matching for correspondence and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *17(6).*

Sebastian, T. & Kimia, B. (2001). Curves vs skeletons in object recognition. *IEEE International Conference of Image Processing, 3,* 22-25.

Sebastian, T., Klein, P. & Kimia, B. (2001). Recognition of shapes by editing shock graphs. IEEE International Conference on Computer Vision, (1), 755–762.

Sebastian, T., Klein, P. & Kimia, B. (2002). Shock-based indexing into large shape databases. In *Lecture Notes in Computer Science* (pp. 731 – 746). Springer-Verlag.

Sebastian, T., Klein, P. & Kimia, B. (2003). On aligning curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *25(1),* 116–125.

Sharvit, D., Chan, J., Tek, H. & Kimia, B. (1998). Symmetry-based indexing of image databases. *Journal of Visual Communication and Image Representation, 9(4),* 366-380.

Shaw, A. (1969). A formal description scheme as a basis for picture processing systems. *Information and Control, 14,* 9-52.

Shepard, R. (1957). Stimulus and response generalization: a stochastic model relating generalization to distance in psychological space. *Psychometricka*, 22, 325-345.

Shepard, R. (1964). Attention and the Metric Structure of the Stimulus Space. *Journal of*
Mathematical Psychology, 1, 54-87.

Shepard, R. (1968). Cognitive psychology: A review of the book by U. Neisser. *American Journal of Psychology, 81*, 285-289.

Shepard, R. (1987). Toward a universal law of generalization for psychological science. *Science, 237,* 1317–1323.

Shepard, R. & Cermak, G. (1973). Perceptual-cognitive explorations of a toroidal set of free-form stimuli. *Cognitive Psychology, 4,* 351-377.

Shepard, R. & Chipman, S. (1970). Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology, 1*, 1-17.

Shin, H. & Nosofsky, R. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General, 121(3),* 278-304.

Siddiqi, K. & Kimia, B. (1996). A shock grammar for recognition. *IEEE Conference on Computer Vision and Pattern Recognition,* 507–513.

Siddiqi, K., Kimia, B., Tannenbaum, A. & Zucker, S. (1999). Shapes, shocks, and wiggles. *Journal of Image and Vision Computing*, *17*, 365-373.

Siddiqi, K., Shokoufandeh, A., Dickinson, S. & Zucker, S. (1999). Shock Graphs and Shape Matching. *International Journal of Computer Vision, 35(1),* 13-32.

Smith, E., & Medin, D. (1981). *Categories and Concepts*. Cambridge, MA: Harvard University Press.

Stewart, N., & Chater, N. (2002). The effect of category variability in perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 893-907.

Tenenbaum, J. & Griffiths, T. (2001), Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*, 629-641.

Thompson, D. (1961). *On growth and form*. Cambridge: Cambridge University Press.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327-352.

Tversky, A. & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General, 113,* 169-193.

Ullman, S. (1989). Aligning pictorial descriptions: an approach to object recognition. *Cognition, 32*, 193-254.

Umeyama, S. (1998). An Eigendecomposition Approach to Weighted Graph Matching Problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *10 (5),* 695-703.

Veltkamp, R. (2001). Shape matching: Similarity measures and algorithms. *Technical Report UU-CS-2001- 03,* Utrecht University.

Veltkamp, R. & Latecki, L. (2006). Properties and Performance of Shape Similarity Measures. *10th IFCS Conference on Data Science and Classification*, Slovenia.

Wilder, J., Feldman, J., & Singh, M. (2008). Shape classification based on natural shape statistics [Abstract]. Journal of Vision, 8(6), 717.

Yushkevich, P. Pizer, S., Joshi, S. & Marron, J. (2001). Intuitive, Localized Analysis of Shape Variability. *Intelligent Platform Management Interface*, 402-408.

Zhang, D. & Lu, G. (2003). Evaluation of MPEG-7 shape descriptors against other shape descriptors. *Multimedia Systems.*, *9(1),* 15–30.

Zhang, D. & Lu, G. (2004). Review of shape representation and description techniques. *Pattern Recognition, 76(1),* 1-19.

Zhu, S. & Yuille, A. (1996). Forms: A flexible object recognition and modelling system.
International Journal of Computer Vision, 20(3), 187–212.

CURRICULUM VITA

Erica Jan Briscoe

Education

Ph.D. in Cognitive Psychology
Graduation Date: October 2008
Rutgers University, New Brunswick, NJ

M.S. in Cognitive Psychology – 2006
Rutgers University, New Brunswick, NJ

M.S. in Information Systems – 2001
Drexel University, Philadelphia, PA

B.S. in Industrial Engineering – 1998
Georgia Institute of Technology
Atlanta, Georgia

Experience

PublixDirect – eCommerce and Home Delivery: Atlanta, GA - Miami, Fl
(11/01-07/03) SENIOR ENGINEER

Drexel University - IST, Computer Resource Center: Philadelphia, PA
(3/01 – 9/01) GRADUATE ASSISTANT

Lockheed Martin - Naval Electronics and Surveillance Systems: Moorestown, NJ
(2/99 - 3/01) OPERATIONS AND SYSTEMS ANALYST, MEMBER ENGINEERING STAFF ***Secret Security Clearance**