

Indexing Point Triples Via Triangle Geometry

Charles B. Cranston and Hanan Samet *

Center for Automation Research, Institute for Advanced Computer Studies
Department of Computer Science
University of Maryland, College Park
{zben,hjs}@cs.umd.edu

Abstract

Database search for images containing icons with specific mutual spatial relationships can be facilitated by an appropriately structured index. For the case of images containing subsets each of which consist of three icons, the one-to-one correspondence between (distinct) point triples and triangles allows the use of such triangle attributes as position, size, orientation, and “shape” in constructing a point-based index, in which each triangle maps to a single point in a resulting hyperdimensional index space. Size (based on the triangle perimeter) can be represented by a single linear dimension. The abstract “shape” of a triangle induces a space that is inherently two-dimensional, and a number of alternative definitions of a basis for this space are examined. Within a plane, orientation reduces to rotation, and (after assignment of a reference direction for the triangle) can be represented by a single, spatially closed dimension. However, assignment of a reference direction for triangles possessing a k -fold rotational symmetry presents a significant challenge. Methods are described for characterizing shape and orientation of triangles, and for mapping these attributes onto a set of linear axes to form a combined index. The shape attribute is independent of size, orientation, and position, and the characterization of shape and orientation is stable with respect to small variations in the indexed triangles.

1 Introduction

Given a set of n pointlike objects, how can an index capture that set’s spatial interrelationships, so as to facilitate retrieval of similar sets in a position, scale, and orientationally independent manner? A companion paper [5] examined this

question for the case $n = 2$. This paper extends the analysis to the case $n = 3$.

Maps provide a means of maintaining a database of cartographic information. Although maps are usually thought of as capturing positional information, they are also used to capture spatial relation information. In particular, relevant objects or features are represented by symbols, and both the topological layout of these objects and their inter-object distances are significant. Examples of such maps include floor plans, blueprints, and satellite images. Examples of features include doorways and windows, component parts of an assemblage, and rivers, roads, and buildings.

A collection of digitized map images is an example of an image database. Searching such a database includes the determination of which images within the database contain a desired arrangement of symbols, such as a beach north of a city or a youth hostel near a train station. For such a search, the absolute position of the symbols within the database image is unimportant, only the relative spatial relationship of the symbols is significant. This is called a *position-independent search*. Search for symbols with a particular spatial relationship has applications in mobile computing, mobile data management, moving-object and moving-framework databases, and location-based services.

One method of specifying such a search is by constructing a *query image* containing the desired topological arrangement of symbols (*icons*). If database images can contain multiple instances of a particular kind of icon (such as multiple hotels in the same geographic area), an icon in a query image actually designates (a member of) an *icon class*, to be satisfied by a corresponding specific icon *instance* in database images matched. In addition to a topological arrangement of icon class designators, a query must also specify which spatial relationships between the classes of the query image are to be considered relevant in the search. This describes a spatial version of *retrieval by content*, which has been investigated in both the spatial and non-spatial domains. The MARCO (MAp Retrieval by Content) system of Samet and Soffer [22, 26] is an example

* The support of the National Science Foundation under Grants EIA-00-91474 and CCF-0515241, Microsoft Research, and the University of Maryland General Research Board is gratefully acknowledged.

of this search method.

Although much of the image database research has examined matching a single query icon to a database icon, some prior work [1, 2, 3, 6, 9, 10, 11, 14, 15, 17, 18, 19, 24, 25] does address queries containing several icons. One approach [3, 15, 17] involves determining the projections of image objects onto the coordinate axes and then encoding these projections as strings. The resulting encoded strings serve in turn as invariants for the index.

Another approach involves modeling the arrangement of image objects as an Attributed Relational Graph (ARG). In [2] a metric-space index is described, using a distance metric function based on the *optimal error-correcting (sub)graph isomorphism problem*.

Gudivada and Raghavan [10, 11] have investigated search based on the relative spatial relationships of icon pairs. The algorithm first determines the n icons common to the query and database images, and then compares the slopes of the $n(n-1)/2$ lines between icon pairs in each database image to the slopes of the lines between corresponding icon pairs in the query image, thus abstracting away the absolute positions of the icons and considering only the relative spatial relationships, as captured by these line slopes. Furthermore, by clustering the set of line slopes, their algorithm can perform a rotation-independent search.

In contrast, the search algorithms to be described here use an index created by examining every image in the database, and mapping each set of three icons into a single point in an abstract index space.

Any three distinct points uniquely determine a triangle. For the special case of three collinear points, a *collinear pseudotriangle* with overlapping sides (and with “angles” of 0° , 180° , and 0°) is determined. As such figures are not included in the classical definition of a triangle, this special case must be carefully considered when invoking any of the classical triangle properties.

Because of this one-to-one correspondence between point triples and triangles, point triples may be indexed using methods that effectively index these associated triangles. So-called “point-based” indexing methods [21, Section 3.4] operate by mapping each of the objects to be indexed into a single point within a multidimensional index space, with each dimension of that space representing a separable attribute of the object. For triangles such attributes might include position, size (based on total perimeter), and spatial orientation. Some abstract representation of “shape” would also be useful as an indexing attribute. To further the goal of attribute separability, the shape attribute should be independent of the triangle’s position, size, and orientation.

Indexing of the spatial position and size of a triangle is straightforward and will not be discussed in the remainder of this paper, except to note that spatial position must be referenced to some specific point in the triangle, and that




Figure 1. Euler line of a triangle

multiple candidates exist for the “center” of a triangle. This will be discussed further in Section 2.

The remainder of this paper is organized as follows: Section 2 briefly outlines some elementary and classical triangle properties, Section 3 introduces the terms *stability* and *selectivity*, which will be used in the remainder of the paper. Section 4 describes an inherently two-dimensional abstract space of triangle “shapes”, and examines methods for defining a basis for this space, including methods based on interior angles and methods based on normalized side lengths. Section 5 discusses the definition of the spatial orientation of a triangle, including the special problems posed by triangles possessing k -fold rotational symmetry. Section 6 presents several alternative schemes for construction of a point-based index, showing how the problems introduced by k -fold rotational symmetry can be addressed by innovative methods of integrating the shape and orientation dimensions. Section 7 contains concluding remarks.

2 Triangle Properties

Various standard terms are used to describe the shapes of triangles. When no two triangle sides have the same length (and thus by the law of sines no two angles are equal) a triangle is termed *scalene*. A triangle with three equal sides (and thus three equal 60° angles) is called *equilateral*. In the remaining case two sides of the triangle are equal; such triangles are called *isosceles*.

In this paper it will be useful to distinguish between two subclasses of isosceles triangles: *tall* isosceles, in which the third side is shorter than the other two, and *wide* isosceles, possessing a longer third side.¹

Classical triangle literature describes a large number of special points, some of which seem likely candidates for the “center” of a triangle. Figure 1 shows three such special points. The *centroid* (center of gravity, barycenter), by con-

¹While it might seem logical to call tall isosceles triangles *prolate* and wide isosceles triangles *oblate*, no support for this terminology could be found in the literature.

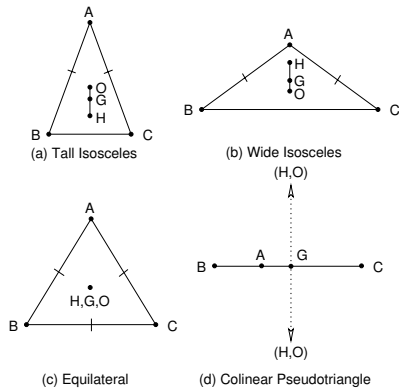


Figure 2. Euler lines of other triangles

vention labeled G, lies at the intersection of the three *medians* (lines from each vertex to the midpoint of the opposite side), which are shown as dotted lines. The *orthocenter*, labeled H, lies at the intersection of the *altitudes* (lines from each vertex perpendicular to the opposite side), which are shown as dashed lines. The *circumcenter*, labeled O, lies at the intersection of the three side perpendicular bisectors, which are shown as solid lines. These three points are always collinear, and define the *Euler line* of the triangle [4, page 17].

Figure 2 shows the Euler line of both tall (a) and wide (b) isosceles triangles. Note that the spatial order of the special points is reversed. While the Euler line will prove useful for defining the spatial orientation of a triangle, two anomalous cases must be considered. In an equilateral triangle (c) the three points coincide, so the classical Euler line is indeterminate. And in the degenerate pseudotriangle defined by three collinear points (d), both H and O lie at the point at infinity, so the polarity of the Euler line analogue becomes indeterminate.

Because it is always spatially enclosed within the triangle, the centroid seems the most logical candidate for being designated as the “center” of a triangle. This definition of “center” will be used in the remainder of this paper.

3 Index Stability and Selectivity

In order to support similarity search, the mapping from triangles to the index space should exhibit *stability*. Triangles differing only slightly in the indexed attributes should map to points in close proximity within the index space. Most “hashing” schemes do *not* exhibit absolute stability, in that cases exist where two similar data items map to widely distant points in the index space. If the index is not stable it may be the case that an object similar to the search target exists in the database but cannot be found.

From the standpoint of efficiency, another desirable in-

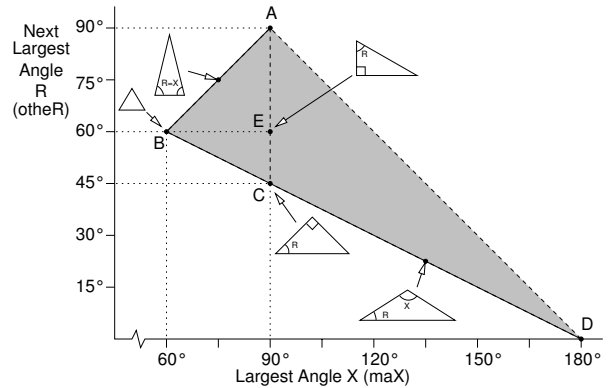


Figure 3. Shape space based on angles

dex property is *selectivity*, the ability to distinguish between different objects. As a simple example, a low-selectivity organization of a collection of screw-fasteners could be three cardboard boxes marked “small”, “medium”, and “large”, while a high-selectivity organization might be a much larger, multidimensional array of boxes, with one dimension for size (“#2” for smallest through “#10” for largest), another dimension for cap style (“flat-head”, “pan”, et. al.), and yet another dimension for application (“metal self-tapping”, “wood self-tapping”, “bolt” [to be used with a matching nut], etc.). The greater selectivity provided by such a multidimensional indexing scheme would allow access to quickly focus to the desired items only.

However, lack of selectivity does not cause a search to fail, it merely results in more “false positives” which must eventually be filtered out of the search results.

Considerations of stability and selectivity affect both the orientational and “shape” attributes, and will be specifically discussed in the corresponding sections.

4 Indexing on Abstract Shape

The three internal vertex angles of a (planar) triangle completely determine its shape: all triangles with the same three angles are *similar* (mathematically, all corresponding sides are in the same proportion). However, because the angles always sum to 180° , specifying any two of the angles completely determines the third. For this reason, the space in which each distinct triangle shape maps to a distinct point is inherently two-dimensional. A basis for such a space can be defined in many ways.

Figure 3 shows one possible basis, that of the largest angle X (maX) plotted against the next-largest angle R (otheR). The smallest angle will later be referred to as N (miN). The shaded area indicates where feasible triangles lie.

The largest angle in a triangle cannot be less than 60°

(if it were, the other two angles would sum to more than 120° , one of the other two angles would then necessarily be greater than 60° , and the first angle would not then be the largest). When the largest angle is 80° , the other two angles must total 100° , but by a similar argument the next-largest angle cannot be less than 50° (if it were, the remaining angle would be greater than 50°). In general, the size of the second-largest angle is bounded from above by either the largest angle itself (line A-B) or the remainder left by the largest angle from 180° (line A-D), and is bounded from below by half this remainder (line B-D).

All equilateral triangles map to point B, and all right triangles map to points along line A-C, with the 45° - 45° - 90° triangle at point C and the 30° - 60° - 90° triangle at point E. Wide isosceles triangles map to points along line B-D, while tall isosceles triangles map to points along line A-B. Points actually on line A-D represent infeasible triangles (two angles summing to 180° leaving no remainder for the third angle). All pseudo-triangles formed by three collinear points map to point D.

Mapping of triangles into this index space depends only on vertex angles, so it is position-independent, orientation-independent, and scale-independent. It is also invariant to a mirror reflection, as the largest and second-largest angles are selected for classification regardless of their clockwise or counterclockwise order in the triangle.

As this mapping is continuous (in the mathematical sense) the *shape stability* can only fail for the “corner cases” where a slight variation in triangle shape (for example, the situation in Figure 8) could cause *shape ambiguity*, a sudden shift in which angle is considered “largest”. For a scalene triangle, where the three angles materially differ, designation of angles as largest, smallest, and “other” is unambiguous. In the case of near-equilateral triangles, where the three angles are almost equal, such triangles will map to the general vicinity of point B, thus expanding the search to an appropriate neighborhood of point B will ensure that all such triangles are found.

In near-wide-isosceles triangles (when two angles are almost equal and the third angle is larger), choice of the largest angle will be unambiguous. Slight variations in the relative sizes of the other two angles will result in a set of index space points that approaches line B-D and then “bounces back” from it. As an example, if the largest angle is 90° , the other two angles will vary in the vicinity of 45° . The trajectory of points in the index space will be the small segment of line E-C in the vicinity of point C. All such triangles will map to the vicinity of point C, and expanding the search to an appropriate neighborhood of point C will ensure that all such triangles are found.

In near-tall-isosceles triangles (two angles are almost equal and the third is smaller), one of the two large angles will be chosen as largest, and the other will be chosen

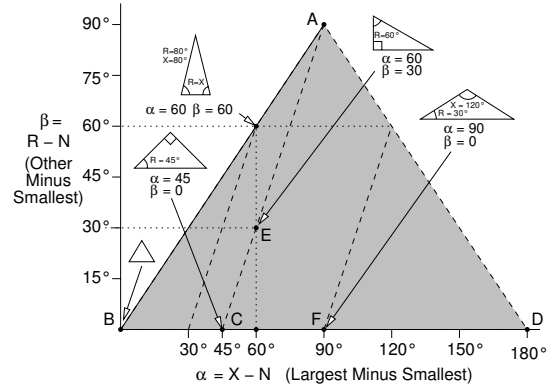


Figure 4. Shape space based on angle differences

as second largest. Slight variations in the relative sizes of these two angles will result in a set of index space points that approaches line A-B and then “bounces back” from it. Expanding the search to an appropriate neighborhood of the contact point with line A-B will ensure that all such triangles are found.

In both the tall and wide isosceles triangles, the invariance of the index to mirror reflection ensures stability in the mapping from triangles to the shape space. However, because the selection of which angles to index is based solely on largest vs. next-largest angle (regardless of their clockwise or counterclockwise order), no distinction is made between the left-handed and right-handed versions of a given triangle. Thus only half the inherent *shape selectivity* is used in the index. Other examples of the sacrifice of selectivity in order to secure stability will be presented later in this paper.

Because the space of triangle shapes is inherently two-dimensional, any pair of linearly independent axes that spans the space may be used as a basis. One such alternative basis, constructed using differences between pairs of angles, is shown in Figure 4, Coordinate α is the difference between the largest and smallest angles ($maX - miN$), and coordinate β is the difference between the other and smallest angles ($otherR - miN$). The original angles can be recovered from α and β as:

$$miN = \frac{180^\circ - \alpha - \beta}{3}$$

$$otherR = \beta + miN$$

$$maX = \alpha + miN$$

Using this basis, triangles with equal largest angles map to points along the lines shown as dashed. For example, all right triangles map to points along line A-C. Note that wide isosceles triangles map to points along line B-D,




Figure 5. Shape space based on normalized side lengths

The advantage to using this shape space is that when the reference orientation is defined as the direction from the center to the vertex with the smallest angle, all triangles with ambiguous reference orientations (equilateral, wide isosceles, and collinear) lie along line B-D. As will be seen in Section, 6, this enables one strategy for integrating the shape space and orientation dimension into the combined index space.

While angle-based methods for characterizing triangle shape are inherently size-independent, side-length-based methods can be made size-independent by basing them on appropriately normalized side lengths. For the remainder of this section it is assumed that triangle side lengths are scaled such that the total triangle perimeter length is 180 units. To compute values based on a perimeter scaled to 1.0, the numbers given should be divided by 180.

When the total perimeter is known, the triangle is completely determined by two side lengths (as the length of the third side can be directly computed). The space shapes generated by a side-based approach are similar to those generated by an angle-based approach.

Figure 5 shows a shape space based on normalized side lengths. The largest side of a triangle cannot be more than half the perimeter, nor can it be less than one third the perimeter. The next largest side is bounded from above by the largest side (line A-B) and from below by half the remaining perimeter (line B-D). In this space, tall isosceles triangles map to points along line A-B, wide isosceles triangles map to points along line B-D, and right triangles map to points along the dashed curve C-E-A.

There are two small advantages to using a side-length-based method over using an angle-based method. When point positions are given as x and y coordinates, side length computation is somewhat less expensive than angle computation, since evaluation of the computationally expensive arctangent function is not required.

Shape	30 60	45 45	60 30
Characterized	$\left[\begin{array}{c} \leftarrow 30 \rightarrow \\ \leftarrow 60 \rightarrow \\ \leftarrow 90 \rightarrow \end{array} \right]$	$\left[\begin{array}{c} \leftarrow 45 \rightarrow \\ \leftarrow 45 \rightarrow \\ \leftarrow 90 \rightarrow \end{array} \right]$	$\left[\begin{array}{c} \leftarrow 60 \rightarrow \\ \leftarrow 30 \rightarrow \\ \leftarrow 90 \rightarrow \end{array} \right]$
By Angle:	$0^\circ-180^\circ-0^\circ$	$0^\circ-180^\circ-0^\circ$	$0^\circ-180^\circ-0^\circ$
By Side Length:	90-30-60	90-45-45	90-60-30

Figure 6. Characterization of shape for collinear point triples (triangle perimeters normalized to 180).




Figure 7. Orientation in n-dimensional space

The second advantage of side-length-based methods is that pseudotriangles corresponding to collinear point triples map to points spread out along a line in the shape space, rather than all mapping to a single point. Figure 6 shows characterizations of three instances of point triples. An angle-based approach does not distinguish between these cases, while an approach based on normalized side length is capable of discriminating between them. In Figure 5 collinear pseudotriangles map to points spread out along line A-D, while in Figure 3 they all map to point D.

Avoiding the sacrifice of the shape information of these pseudotriangles results in a slight selectivity advantage for side-length-based methods.

As in the angle-based method, because the selection of which sides to index is based solely on largest and second-largest sides (regardless of their clockwise or counterclockwise order), only half of the shape selectivity inherent in the indexed triangles will be used in the shape index.

5 Indexing on Spatial Orientation

In the general n -dimensional case, characterization of an object's "spatial orientation" can be accomplished in three steps, as shown in Figure 7. First, a reference direction for the object itself is assigned. This can be visualized as a unit vector (a) that is rigidly attached to the object. If the object possesses any rotational symmetries an unambiguous assignment may not be straightforward. Second, a direction (b) in the embedding space is identified, in parallel

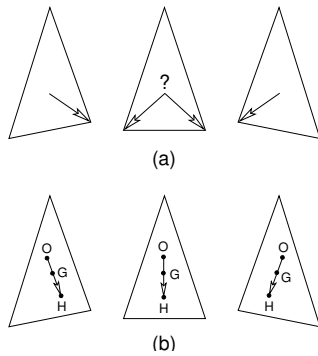


Figure 8. Stability of reference direction

with which the object's reference direction will be aligned. Third, the final remaining degree of freedom, that of the object rotating (c) around the parallel line, is characterized. This requires an additional reference direction assignment, which may be complicated by any rotational symmetries around the axis that the object might possess.

For this general case $n - 1$ parameters are required to define the direction in the embedding space, and $n - 2$ additional parameters are required to characterize the rotation. Thus $2n - 3$ parameters suffice to characterize the spatial orientation of the object. In three dimensions the convention for designating the direction in space (b) is by specifying two angles θ and ϕ . In two dimensions, however, orientation simplifies to rotation within a plane, and no further degree of freedom (c) is allowed (as the triangle cannot rotate up out of the plane). In this special case a single circular dimension is sufficient to characterize the orientation of a triangle.

The direction from the triangle center toward a *distinguished member* of the vertex set (such as the vertex with the smallest or largest interior angle) could be a candidate for a geometrically-based definition for a triangle's reference direction. However, instability in the designation of this vertex results in the loss of *orientational stability* in the index. Figure 8(a) shows the loss of orientational stability for triangles in the near-tall-isosceles region when the vertex with the largest angle is used as the distinguished member (In the case of designating the vertex with the smallest angle, an analogous situation arises with triangles in the near-wide-isosceles region.) As the triangle varies from near-isosceles-left to near-isosceles-right, the reference orientation snaps discontinuously from one vertex to another, and when the triangle is truly isosceles it is indeterminate. An alternative for defining the orientation of a triangle is to use the triangle's Euler line. As can be seen from Figure 8(b), the Euler line provides a definition of orientation that is stable for triangles in the near-tall-isosceles region.

As noted in Section 2, in the case of the collinear pseudotriangle the sense of the Euler line is indeterminate. While




Figure 9. Reference direction for collinear pseudotriangles

the definition of reference direction as the direction from the center to the vertex with the smallest angle suffers from instability in the near-wide-isosceles region, it does allow stable assignment of reference direction for asymmetric collinear pseudotriangles. Figure 9 shows the behavior of a triangle in the limit (a,b) as it approaches becoming an asymmetric collinear pseudotriangle (c). When the interior point is on the other side (d) the direction is reversed. However, when the interior point is exactly equidistant from the other two (the symmetric collinear pseudotriangle), the smallest angle method does not assign a stable reference direction.

No matter which method for assigning a reference direction is employed, two special cases remain a significant challenge. Equilateral triangles have an inherent 3-fold rotational symmetry. For these triangles it is not possible to geometrically define an unambiguous reference direction. Similarly, symmetric collinear pseudotriangles possess a 2-fold rotational symmetry. The common factor in these two special cases is that they possess a k -fold rotational symmetry. The inability to define a reference direction is not simply a result of a particular choice of definition. The figures themselves possess an innate rotational symmetry and do not admit an unambiguous geometric definition of reference direction.

The approach taken in this paper is to assign the value 360° as the inherent *orientational selectivity* of a scalene triangle. Triangles with a k -fold rotational symmetry are considered as possessing a k -fold *orientational ambiguity*, and $360^\circ \div k$ is assigned for their inherent orientational selectivity. Thus 120° of orientational selectivity is retained by an equilateral triangle, and 180° for a symmetric collinear pseudotriangle.

There are several possible approaches for handling objects possessing a potential k -fold rotational symmetry within the orientation dimension of a point-based index.

1. Each indexed object may result in several index entries, with differing index values. For example, a triangle could be entered in three different places, corresponding to directions from the center to each of the

three vertices. The drawback of this technique is that the index must contain three times as many entries.

2. The search algorithm can search more than one area of the index space. For example, directions from the center of the target triangle towards each of its three vertices can be determined, and the index space can be searched for all three directions. The drawback of this technique is that three times the area of the index must be accessed.
3. The symmetry can be broken by point tagging. If the points that form the triangle are tagged (say, one represents a hotel, another a railroad station, and the third a youth hostel), a predefined ordering of the tag values can be used to identify a distinguished member of the vertex set, and orientation can then be defined relative to that vertex. However, when a distinguished member cannot be determined (such as in cases when all points happen to be tagged identically) this technique fails to determine a reference direction.
4. The symmetry can be broken geometrically. Various schemes are possible for triangles, including largest-angle, smallest-angle, and Euler line. However, when the figure to be indexed has a k -fold rotational symmetry, this technique fails to determine a reference direction.
5. The orientation dimension itself can be folded by a factor that is a multiple of k , thereby bringing all points at which the object could potentially be mapped into a single position along the folded dimension. The set of all possible triangles includes both figures with a 3-fold rotational symmetry (equilateral triangles) and figures with a 2-fold rotational symmetry (symmetric collinear pseudotriangles). Folding the orientation dimension six times (that is, reducing the orientation angle modulo 60°) allows a direction to be defined for equilateral triangles (for example, the direction from the center to *any* vertex), with the property that regardless of which vertex is chosen, the triangle will map to the same coordinate in the orientation dimension. The same is true for symmetric collinear pseudotriangles. No matter which of the two possible senses is chosen, the pseudotriangle will map to the same coordinate in the orientation dimension. For example, the directions from the center of an equilateral triangle to its vertices might be 40° , 160° , and 220° , but when reduced modulo 60° all three values become 40° , providing an unambiguous characterization of the triangle's orientation. Similarly, a collinear pseudotriangle with an Euler line at an orientation of 30° or 210° (depending on the assigned sense) generates an unambiguous value of 30° when reduced modulo 60° . The drawback

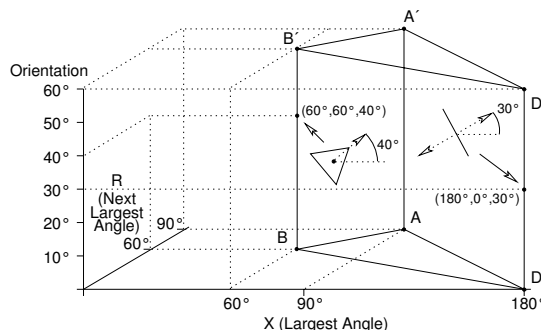


Figure 10. A60: Angle-based, folded at 60°

of this scheme is that only $1/k$ of the inherent orientational selectivity is used in the index. Furthermore, this sacrifice is made not only for the ambiguous triangles, but instead for every triangle in the index,

6. The orientation dimension can be combined with the other index space dimensions in a structure with the property that objects with ambiguity of reference direction map to regions of the index space where the orientation dimension is suppressed. The drawback of this technique is that *none* of the inherent orientational selectivity of such objects is used in the index.

The following section will give several examples of how to use these last two techniques, separately and in combination, to address the problem of orientation ambiguity in a point-based index space.

6 Index Space Construction

A point-based index for triangles can be constructed by mapping any chosen subset of the separable attributes of position, size, orientation, and shape onto orthogonal axes of a hyperdimensional index space. Index dimensions for position and size are straightforward. Several alternate methods of structuring the orientation and shape dimensions are presented here.

Figure 10 shows an index space (designated A60) formed by extending the angle-based shape space of Figure 3 along an orientational dimension with a 6-way folding (that is, the modulo 60° residue of the triangle's orientation determines the orientation component of the index space coordinates). Thus, for an equilateral triangle with possible orientations of 40° , 160° , or 280° , the triangle will be indexed at 40° up on the B-B' line. Similarly, a collinear pseudotriangle with orientation of 30° or 210° will be indexed at 30° up on the D-D' line. As the orientation of wide isosceles triangles is not suppressed, the Euler line method should be used to define their reference direction.




Figure 11. A360: Angle-based, folded at 360°

It should be noted that the orientational dimension is wrapped at 60°, that is, the top plane A'-B'-D' is actually identical to the bottom plane A-B-D. If this figure were redrawn to be analogous to the other figures in this section, it would resemble Figure 12, with the exception that the locus of point B during the rotation would be a small circle instead of a point.

This index space structure corresponds to using technique 5 from the previous section to address both the $k = 2$ and $k = 3$ cases of orientational ambiguity. Their least common multiple value of 6 is used to fold the orientational index dimension. The disadvantage of doing so is that the orientational selectivity of the index for all triangles is reduced by the same factor, resulting in the generation of a large number of false positives that must eventually be filtered out of the result set.

Figure 11 shows an index space (designated A360) formed by rotating the angle-based shape space of Figure 4 around the B-D axis. With this structure, all triangles with ambiguous reference orientation lie along line B-D and thus their orientational index information is suppressed. That is, no matter what their original spatial orientation, they map along line B-D in this index space. Because the orientation information for wide isosceles triangles is suppressed, the smallest-angle method can be used to define their reference orientation.

This index space structure corresponds to using technique 6 from the previous section to address both the $k = 2$ and $k = 3$ cases of orientational ambiguity. The disadvantage of doing so is that none of the inherent orientational selectivity of equilateral triangles, wide-isosceles triangles, and collinear pseudo-triangles is used in the index (because they all map to points along the line B-D, where the index structure does not capture orientation). However, the orientational space need not be folded, so there is no sacrifice of orientational selectivity for unambiguous triangles (the vast majority of the triangles indexed).

A compromise approach would preserve some of the orientational selectivity of wide isosceles triangles and of collinear pseudo-triangles, while sacrificing less of the orientational selectivity of the vast majority of triangles that possess no ambiguity.




Figure 12. A180: Angle-based, folded at 180°




Figure 13. S120: Side-based, folded at 120°

Figure 12 shows an index space (designated A180) formed by rotating the angle-based shape space Figure 3 around the equilateral triangle index point (B). Mapping all equilateral triangles, regardless of orientation, to this single point in the index space suppresses the orientational ambiguity due to 3-fold rotational symmetry. The rotational space is folded at 180° so the ambiguity in collinear pseudotriangles is suppressed, such figures map to points along circle D-D' at the bottom of the space. Wide isosceles triangles map to the cone D-B-D' and the Euler line method should be used to define their reference orientation. Right triangles map to the cylinder A-A'-C'-C within the space.

A similar structure could be generated by rotating Figure 4 around point B. In fact, any formulation of triangle shape space could be rotated around the equilateral triangle point to generate an analogous index space.

This index structure corresponds to using technique 5 to address $k = 2$ ambiguity and technique 6 to address $k = 3$ ambiguity. The disadvantage of doing so is that the inherent orientational selectivity of equilateral triangles is not used in the index. In addition, because the orientation space is folded at 180°, only half the inherent orientation selectivity of all other triangles is used in the index.

Figure 13 shows an index space (designated S120) formed by rotating the normalized-side-length based shape space Figure 5 around line A-D along which all the collinear pseudotriangles map. This index structure corresponds to using technique 5 to address $k = 3$ ambiguity and technique 6 to address $k = 2$ ambiguity. The disadvantage of doing so is that none of the inherent orientational selectivity of collinear pseudotriangles is used in the index. Because the




Figure 14. S180: Side-based, folded at 180°

orientational index is wrapped at 120°, all the orientational selectivity of equilateral triangles is used in the index. However, only one third of the inherent orientational selectivity of scalene and isosceles triangles is used in the index.

Figure 14 shows an index space (designated S180) formed by rotating the normalized-side-based shape space of Figure 5 around point B. In this index space all equilateral triangles map to point B, thus the 3-fold orientational ambiguity inherent in these equilateral triangles is accommodated, at the cost of sacrificing all their orientational selectivity. Right triangles map to points along the vase-shaped figure A-C-C'-A'. Tall isosceles triangles map to the cone A-B-A' and wide isosceles triangles map to the cone D-B-D'. Collinear pseudotriangles map to the cylinder A-D-D'-A'. The 2-fold orientational ambiguity of these pseudotriangles is accommodated by folding the orientation axis at 180°.

In common with the A180 structure, this index space corresponds to using technique 5 to address $k = 2$ ambiguity and technique 6 to address $k = 3$ ambiguity. The disadvantage of doing so is that only half of the inherent orientational selectivity of all triangles is used in the index. However, The orientational selectivity of asymmetric collinear pseudotriangles is used in the index. This is the advantage of using the side-based shape space.

The fraction of the inherent orientational selectivity (broken down for various classes of triangles) actually used for each of these index structures is shown in Figure 15.

7 Concluding Remarks

In an iconic image database, point triples can be indexed based on geometric attributes of their associated triangles. A size and orientation independent characterization of the triangle “shape” attribute can be based either on interior angles or normalized side lengths; in both cases the induced shape space is inherently two-dimensional.

In the general case of an object embedded in an n -dimensional space, characterization of spatial orientation requires n parameters, but in the special case of a planar (two-dimensional) space, orientation degenerates into sim-

	A60 S60	A120 S120	A180 S180	A360 S360
Scalene Triangle	1/6 1/6	1/3 1/3	1/2 1/2	All All
Tall Isosceles	1/6 1/6	1/3 1/3	1/2 1/2	All All
Wide Isosceles	1/6 1/6	1/3 1/3	1/2 1/2	None None
Asymmetric Collinear	1/6 1/6	None None	None 1/2	None All
Symmetric Collinear	1/3 1/3	None None	All All	None None
Equilateral Triangle	1/2 1/2	All All	None None	None None

Figure 15. Orientational selectivity used

ple rotation about an axis, and a single parameter suffices. Objects with rotational symmetry inherently possess ambiguity in the assignment of reference direction; no geometrical method exists to resolve this ambiguity.

In integrating the shape space and the orientation dimension into a combined index space, two techniques are available for addressing the orientational ambiguity of objects possessing k -fold rotational symmetry.

1. The points in the shape space to which the problematic objects map can be located at points in the combined index space where the coordinate value of the orientation dimension is suppressed (for example, at the zero point of a polar plot). The disadvantage of doing so is that the inherent orientational selectivity of such objects will not be used in the resulting index.
2. The orientation dimension itself can be folded at an angle less than 360°, such that all ambiguous orientation values map to the same point in the index space. The disadvantage of doing so is that only some fraction of the inherent orientational selectivity for every object will be used in the index, even for those objects that do not possess any orientational ambiguity.

The A60 index addresses both ambiguous cases, equilateral triangles (3-fold rotational symmetry) and symmetric collinear pseudotriangles (2-fold rotational symmetry) by folding the orientation dimension at $360^\circ \div 6 = 60^\circ$. The A360 index addresses both ambiguous cases by mapping them (and all wide isosceles triangles) to the zero point of a polar plot. The other example index spaces utilize combinations of these techniques, one technique to address each ambiguous case. Absent an application-specific knowledge of the distribution of triangles, the S360 index would appear to be optimal, as it uses all the available orientational selectivity of the majority of the triangles expected.

The methods described can be used to construct a point-based index with which any combination of position-independent, orientation-independent, and size-independent search can efficiently be accomplished.

References

- [1] D. Arkoumanis, M. Terrovitis, and E. Stamatogiannakis. Heuristic algorithms for similar configuration retrieval in spatial databases. In *Proceedings of the 2nd Hellenic Conference on Artificial Intelligence (SETN)*, pages 141–152, Thessalonica, Greece, 2002.
- [2] S. Berretti, A. D. Bimbo, and E. Vicario. Efficient matching and indexing of graph models in content-based retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(12):1089–1105, Oct. 2001.
- [3] S. K. Chang, Q. Y. Shi, and C. Y. Yan. Iconic indexing by 2-D strings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(3):413–428, May 1987.
- [4] H. S. M. Coxeter. *Introduction to Geometry*. John Wiley & Sons, New York, second edition, 1969.
- [5] C. B. Cranston and H. Samet. Efficient position-independent iconic search using an r-theta index. In preparation.
- [6] M. J. Egenhofer. Query processing in spatial-query-by-sketch. *Journal of Visual Languages and Computing*, 8(4):403–424, Aug. 1997.
- [7] J. Enderle, M. Hampel, and T. Seidl. Joining interval data in relational databases. In *Proceedings of the ACM SIGMOD Conference*, pages 683–694, Paris, France, June 2004.
- [8] C. Faloutsos and W. Rego. Tri-cell—a data structure for spatial objects. *Information Systems*, 14(2):131–139, 1989.
- [9] R. K. Goyal and M. J. Egenhofer. Similarity of cardinal directions. In *Proceedings of the Seventh International Symposium on Spatial and Temporal Databases.*, volume 2121 of *Lecture Notes in Computer Science*, pages 36–55, Los Angeles, California, July 2001.
- [10] V. Gudivada. $\Theta\mathfrak{R}$ string: A geometry-based representation for efficient and effective retrieval of images by spatial similarity. *IEEE Transactions of Knowledge and Data Engineering*, 10(3):504–512, May/June 1998.
- [11] V. Gudivada and V. Raghavan. Design and evaluation of algorithms for image retrieval by spatial similarity. *ACM Transactions on Information Systems*, 13(2):115–144, Apr. 1995.
- [12] A. Henrich, H.-W. Six, and P. Widmayer. The LSD tree: spatial access to multidimensional point and non-point data. In P. M. G. Apers and G. Wiederhold, editors, *Proceedings of the 15th International Conference on Very Large Databases (VLDB)*, pages 45–53, Amsterdam, The Netherlands, Aug. 1989.
- [13] K. Hinrichs and J. Nievergelt. The grid file: a data structure designed to support proximity queries on spatial objects. In M. Nagl and J. Perl, editors, *Proceedings of WG'83, International Workshop on Graphtheoretic Concepts in Computer Science*, pages 100–113, Haus Ohrbeck (near Osnabrück), West Germany, 1983. Trauner Verlag.
- [14] A. Holt. Spatial similarity and GIS: the grouping of spatial kinds. In P. A. Whigham, editor, *Eleventh Annual Colloquium of the Spatial Information Research Center (SIRC05)*, pages 241–250, Dunedin, New Zealand, Dec. 1999. University of Otago.
- [15] P. W. Huang and Y. R. Jean. Using 2D C^+ -strings as spatial knowledge representation for image database systems. *The Journal of the Pattern Recognition Society*, 27(9):1249–1257, Sept. 1994.
- [16] U. Lauther. 4-dimensional binary search trees as a means to speed up associative searches in design rule verification of integrated circuits. *Journal of Design Automation and Fault-Tolerant Computing*, 2(3):241–147, July 1978.
- [17] S. Y. Lee and F. J. Hsu. 2D C-string: a new spatial knowledge representation for image database systems. *Pattern Recognition*, 23(10):1077–1088, Oct. 1990.
- [18] D. Papadias, N. Karacapilidis, and D. Arkoumanis. Processing fuzzy spatial queries: A configuration similarity approach. *International Journal of Geographic Information Science*, 13(2):93–128, 1999.
- [19] D. Papadias, M. Mantzourogiannis, and I. Ahmad. Fast retrieval of similar configurations. *IEEE Transactions on Multimedia*, 5(2):210–222, June 2003.
- [20] J. B. Rosenberg. Geographical data structures compared: a study of data structures supporting region queries. *IEEE Transactions on Computer-Aided Design*, 4(1):53–67, Jan. 1985.
- [21] H. Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, Reading, MA, 1990.
- [22] H. Samet and A. Soffer. MARCO: MAp Retrieval by COntent. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(8):783–798, Aug. 1996.
- [23] B. Seeger and H.-P. Kriegel. Techniques for design and implementation of efficient spatial access methods. In F. Bachillon and D. J. DeWitt, editors, *Proceedings of the 14th International Conference on Very Large Databases (VLDB)*, pages 360–371, Los Angeles, Aug. 1988.
- [24] A. P. Sistla, C. Yu, and R. Haddad. Reasoning about spatial relationships in picture retrieval systems. In J. Bocca, M. Jarke, and C. Zaniolo, editors, *Proceedings of the 20th International Conference on Very Large Data Bases (VLDB)*, pages 570–581, Santiago, Chile, Sept. 1994.
- [25] S. Skiadopoulos, C. Giannoukos, P. Vassiliadis, T. K. Sellis, and M. Koubarakis. Computing and handling cardinal direction information. In *Proceedings of the 9th International Conference on Extending Database Technology (EDBT)*, volume 2992 of *Lecture Notes in Computer Science*, pages 329–347, Heraklion, Crete, Greece, Mar. 2004.
- [26] A. Soffer and H. Samet. Pictorial query specification for browsing through spatially referenced image databases. *Journal of Visual Languages and Computing*, 9(6):567–596, Dec. 1998. Also an abbreviated version in *Proceedings of the Second International Conference on Visual Information Systems (VISUAL97)*, pages 117–124, San Diego, CA, December 1997.
- [27] J.-W. Song, K.-Y. Whang, Y.-K. Lee, M.-J. Lee, and S.-W. Kim. Spatial join processing using corner transformation. *IEEE Transactions on Knowledge and Data Engineering*, 11(4):688–695, July/August 1999.