

Indexing Planar Point Quartets Via Geometric Attributes

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@cs.umd.edu, hjs@cs.umd.edu

ABSTRACT

An index is devised to support position-independent search for images containing quartets of icons. Point quartets exist that do not unambiguously map to quadrilaterals, however, four points do unambiguously determine a set of six interpoint line segments. Values for the “size”, “shape”, and “orientation” attributes of an icon quartet can be derived as functions of this interpoint line set, and can be used to construct a point-based index, in which each point quartet maps to a single point in the resulting hyperdimensional index space. Orientation can be represented by a single, spatially closed dimension. However, assignment of a reference direction for quartets possessing a k -fold rotational symmetry presents a significant challenge. Methods are described for determining shape and orientation attributes for point quartets, and for mapping these attributes onto a set of attribute axes to form a combined index. The orientation computation supplies, as a byproduct, one component of the shape attribute. All attributes are continuous with respect to small variations in the indexed point quartets.

1. INTRODUCTION

Given a set of n pointlike objects, we are interested in constructing an index that can capture that set’s spatial interrelationships, so as to facilitate retrieval of similar sets in a position, scale, and orientationally independent manner. Position-independent search for arrangements of point-like objects has also been investigated within the context of *astrometry* [6, 3, 4]. Previous work [1, 2] examined this question for the cases $n = 2$ and $n = 3$. This paper extends the analysis to the case $n = 4$, which leads to a number of interesting issues..

The search algorithms to be described here use a *point-based indexing* method [5, Page 1], which operates by mapping each of the objects to be indexed into a single point

* The support of the National Science Foundation under Grants EIA-00-91474, CCF-0515241, and IIS-07-13501, Microsoft Research, Nvidia, and the University of Maryland General Research Board is gratefully acknowledged.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM GIS '08, November 5–7, 2008, Irvine, CA, USA

Copyright 2008 ACM ISBN 978-1-60558-323-5/08/11 ...\$5.00.

within a multidimensional index space, with each dimension of that space representing a separable attribute of the object. For point configurations such attributes might include absolute spatial position, size, “shape”, and spatial orientation. We choose to define *shape* as that which remains when size and orientation have been removed.

2. INDEXES AND HASHING

In a point-based index, members of the set of objects to be indexed are mapped by an indexing function into points of a multidimensional index space. Ideally this mapping would be one-to-one; that is, two different objects would never map to the same point. Relaxation of this constraint results in *hashing*; the case where two different objects map to the same index point is called a *hash collision* (it is equally valid to view hashing as the general process, and to refer to the ideal case as *perfect hashing*).

A *false positive* is the penalty for a hash collision. A target object is used to determine an index space point, but some database objects indexed in the vicinity of this target point (in the same *hash bucket*) are not similar to the target object. When appropriate filtering is used to weed out false positives, the only effect hash collisions cause in a database search application is the degradation of response time, caused by both the increased amount of data retrieved from external storage and the added costs of the filtering process itself.

The goal of similarity search is to find an object similar to (but not necessarily identical to) a target object. To support similarity search, an indexing function must be mathematically *continuous*—that is, objects that are “similar to” each other must map to proximate points in the index space.

In some cases these goals of one-to-one and continuous mapping are mutually exclusive. For example, in order to ensure continuity in the regions corresponding to symmetric collinear point triples, the indexes described in [2] map mirror image triangles into the same index space point. For the case of four or more points, the configuration space is of a much higher dimensionality, and continuous one-to-one mappings between spaces of differing dimensionality are not possible. In such cases, it may be necessary to tolerate the degradation of response time caused by relaxation of the one-to-one requirement for the index function.

3. EMERGENT PROPERTIES

It is well known that when attempts are made to generalize based on a known case, complications (termed *emergent properties*) not previously present sometimes arise. An




Figure 1: Non-convexity

example of such a complication is the emergence of non-Abelian properties in the group of rotations in k -dimensional space when k increases from 2 to 3. Thus it is not surprising that several such complications arise when generalizing from point triples to point quartets. For example, the four points may be non-coplanar, or may fail to unambiguously determine a convex quadrilateral. Moreover, even if a quadrilateral is determined, it will generally be non-rigid, and in some cases may scale in non-isotropic ways.

With the exception of the one special case of collinear points (which has been suitably handled in [2]), configurations of three points uniquely determine triangles, which are always convex. However, configurations of four points fall into two distinct classes. In particular, the convex hull of the four points may contain all four points, as in Figure 1(a), in which case a convex quadrilateral is uniquely determined. Alternatively, as in Figure 1(b), it may be the case that the convex hull contains only three of the points, with the fourth point being interior to the triangle formed by the other three points, and in this case, an unambiguous definition of a traditional quadrilateral is not possible, as there are three equally valid possibilities, shown in Figure 1(c), (d), and (e). Furthermore, regardless of which of the three possibilities is selected, the resulting quadrilateral will fail to be convex.

In all cases, a set of six *interpoint lines* (as shown in Figure 1(f)) is determined. However, in the case where the fourth point is interior to the triangle, it is not possible to partition these six line segments into a set of four “sides” and a set of two “diagonals”.

A quadrilateral (or other higher-order polygon) is in general not rigid. Even if the side lengths are completely determined, there remain unconstrained degrees of freedom. Loosely, the quadrilateral can be said to “flop around”.

Another interesting emergent property is that in many cases the shape of a quadrilateral is not completely determined by the set of vertex angles. In particular, whenever two (necessarily opposite) sides of a quadrilateral are parallel, the set of side lengths is not constrained by the vertex angles, but can vary anisotropically. We call such behavior “tromboning”.

Figure 2 shows three figures with identical vertex angles but different shapes. The equivalence of the corresponding angles can be seen by observing that each figure consists of an inner rectangle flanked by a pair of congruent 3-4-5 right triangles (the dotted-line altitude delimits the boundary of one such triangle). Thus the doubly-slashed angle is $\arctan \frac{4}{3}$ or approximately 53.13° . This example has been constructed so that the three figures have equal perimeters (52). While this will not generally be true, it does demonstrate the existence of cases where the perimeter fails to distinguish the shapes.

4. SIZE

There are a number of different ways to characterize a configuration’s size. One would be to sum the lengths of the six interpoint line segments. In many index designs the size




Figure 2: Non-isotropic Scaling




Figure 3: Anomalies in side length characterization

attribute will be the only unbounded one. Orientation is inherently limited to 360° , and shape can be defined using a set of bounded measures. While the application domain of the index may impose an independent limit on the maximum size configuration to be indexed (and thus the maximum value of the size attribute), the size attribute is the only one that is inherently unbounded.

5. SHAPE

In this section we analyze the concept of the “shape” of a configuration of four points, by first examining the case of a traditional convex quadrilateral, and then extending the analysis to the more general case.

As we have seen, the abstract shape of a traditional quadrilateral cannot (due to non-rigidity) be unambiguously characterized by a set of side lengths, nor (due to tromboning) can it always be unambiguously characterized by a set of vertex angles. Therefore, the values for a shape attribute must be derived from a combination of both side lengths and vertex angles.

To accommodate uniform scaling, the perimeter of the quadrilateral is normalized (say to 1), with the normalization factor becoming one of the index attributes. Equivalently, the “size” of the quadrilateral becomes one axis of the index space. After this normalization, any three side lengths unambiguously determine the length of the fourth side.

The method of organizing the set of side lengths also affects the continuity of the index (recall the definition of continuity in Section 2). Figure 3(a-c) shows the results of an attempt to characterize the side lengths starting with the longest side and proceeding clockwise (note that in spaces of more than 2 dimensions even the notion of “clockwise” breaks down). When the longest side is unambiguous (cases (a) and (c)), the ordering of the lengths is well defined, but at the seam (b) where two sides are nearly the same length, continuity from both directions cannot be obtained.

One way to avoid this discontinuity is to construct the index based on the sorted set of side lengths, irrespective of their relative position in the quadrilateral, as shown in Figure 3(d,e). While this method avoids the discontinuity described above, it does introduce the possibility for additional hash collisions. In particular, this demonstrates that two very different shapes can generate the same sorted set of side lengths, and so would collide on this coordinate of a point-based index.

One vertex angle is also required to complete the shape characterization. This can be the largest (or the smallest) of the vertex angles. Although choosing the largest angle, regardless of its geometric relationship to the sides, gener-

	Points	-2	-1	-1	remain
	x, y	pos'n	scale	rot'n	“shape”
Pair	4	-2	-1	-1	0
Trio	6	-2	-1	-1	2
Quartet	8	-2	-1	-1	4
Quintet	10	-2	-1	-1	6

Figure 4: Degrees of Freedom (2-D Space)

ates the same kind of ambiguity as does a histogram of the side lengths, it does assure continuity of the index function at seams where the geometric position of the largest angle suddenly shifts.

As discussed in Section 3, when the four points do not all lie on their mutual convex hull (as in Figure 1(b)), it is not possible to unambiguously define a traditional quadrilateral. But for any arbitrary configuration of four distinct points, the six interpoint distances (Figure 1(f)) are well defined (and nonzero), and thus could be used to construct an alternative shape component of an index. Shape changes that for a traditional quadrilateral would be described as “flopping” also change the lengths of the diagonals, thereby allowing a shape index to be defined based on lengths only, without requiring the consideration of vertex angles.

Figure 4 contains an analysis of the degrees of freedom that are involved in the characterization of the “shape” of arbitrary planar configurations of n points, and demonstrates that it is inherently $2n - 4$ dimensional. This can be seen by observing that each of the n points in the configuration introduces two degrees of freedom, in the form of x and y coordinates. Two degrees of freedom are subsumed by the absolute spatial position of the configuration as a whole (e.g., the x and y coordinates of the center of gravity), one degree of freedom is subsumed in uniform scaling, and one in rotation (e.g., orientation) in the plane. For example, the six initial degrees of freedom of a trio of points (e.g., a triangle), are reduced to two inherent shape dimensions. This makes sense, as the “shape” of a triangle is completely determined by two vertex angles, or (given a normalized perimeter) by two side lengths (rigidity was discussed in Section 3).

From this analysis, as the shape space for configurations of four points is inherently four-dimensional, the shape component of a point-based index could be based on any of a number of schemes such as the largest four of the six interpoint distances, the largest two and the smallest two, etc. Monte Carlo experiments could investigate the performance of these schemes under differing data distribution statistics.

Note that for a collinear configuration the length of the longest segment (between the two extremal points) remains unchanged as the positions of the interior points vary. Thus, in order to fully utilize the shape variety of such configurations, the technique chosen to derive the shape attribute should ensure that segments other than the longest segment are included.

As we shall see in the next section, use of an “orientation” attribute can also contribute a component to the set of “shape” attributes.

6. ORIENTATION

While rotation in a plane involves only a single degree of freedom, determining a reference direction for an arbitrary point quartet is not straightforward. The method described in this section for deriving a value for the orientation attribute is based on the concept of an *eccentricity vector*.




Figure 5: Mapping a line segment in π -space

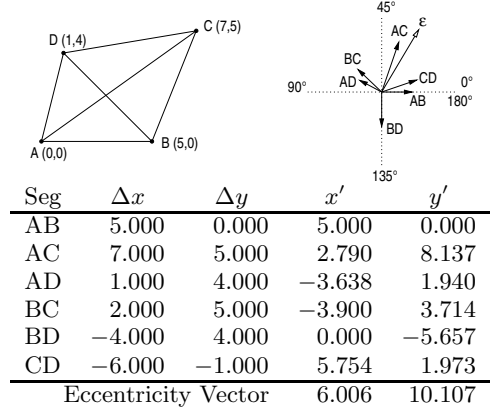


Figure 6: Point quartet with six lines

Intuitively, our aim is to capture a measure of the deviation of a given configuration from a symmetric one. When the vector is nonzero, its direction component provides an absolute value for the “orientation” of the configuration.

The ambiguity associated with defining the direction of a line segment can be removed by mapping the length and direction of the line onto a space in which a rotation through 180° leaves the direction invariant. Mapping a line segment as a vector in such a π -space is equivalent to moving one endpoint to the origin and doubling the angle corresponding to its slope. The mathematics involved follow directly from the classical double-angle formulas. This can be seen with the aid of Figure 5. Substituting the double-angle formulas $\sin 2\alpha = 2 \sin \alpha \cos \alpha$ and $\cos 2\alpha = \cos^2 \alpha - \sin^2 \alpha$ and simplifying yields:

$$x' = \frac{\Delta x^2 - \Delta y^2}{\sqrt{\Delta x^2 + \Delta y^2}} \quad y' = \frac{2\Delta x \Delta y}{\sqrt{\Delta x^2 + \Delta y^2}} \quad (1)$$

Note that reversal of the direction of the line segment (corresponding to the negation of both Δx and Δy) leaves Equation 1 unchanged, as the two negations cancel out in the numerator of the y' term and all other appearances of Δx and Δy are squared. In this method, angles were multiplied by two in order to accommodate a two-way rotational symmetry. The general technique of multiplication of the angle by k to accommodate a k -way rotational symmetry is called *normalization*, and is beyond the scope of this paper.

In order to accommodate the cases where a point quartet does not correspond to a classical convex quadrilateral, the six interpoint lines can be mapped into π -space and summed, to yield a π -measure for the rotation of the quartet. As a concrete example, Figure 6 shows a quartet of points that actually does form a classical quadrilateral, with the six interpoint lines mapped into π -space, and the resulting sum computed using Equation 1.

The use of this approach to define a reference direction may fail when the quartet possesses a k -fold rotational sym-




Figure 7: Rotational symmetry

metry. For example, consider Figure 7(a), where three points are arranged at the vertices of an equilateral triangle and a fourth point is located at the center. This figure possesses a 3-fold rotational symmetry. Figure 7(b) shows the result of mapping the six interpoint lines in π space and we see that the vector sum in this space is zero, which means that in this case no reference direction is determined.

As another example, consider Figure 7(c) where four points are arranged at the corners of a square. This figure possesses a 4-fold rotational symmetry. In Figure 7(d) the six interpoint lines are mapped into π space, and again, the vector sum is zero (as a convention, dashed lines indicate vectors drawn off-origin for clarity). However, in the case of a figure with two-fold rotational symmetry (e.g., Figure 7(e)) the π space vector sum (f) is non-zero. Thus this scheme does allow the definition of reference direction for figures with a two-fold rotational symmetry.

In general, recall from Figure 4 the configuration space of point quartets is inherently four-dimensional, and these methods map it to a two-dimensional (polar coordinate) space. Those shapes which map to the polar origin constitute the *kernel* of this mapping, and the Rank-nullity Theorem can be used to show that the dimensionality of this kernel is two. Thus the shapes which result in a zero eccentricity vector form some two-dimensional subspace of the four-dimensional shape space.

The choice of eccentricity vector used in any particular index should be influenced by the statistical properties of the expected data. Symmetric configurations have the potential of generating a zero eccentricity vector. We want to make sure that whatever eccentricity vector method that we use generates a zero vector only for configurations (symmetric and non-symmetric) that are not frequently expected. The reason for doing so is to avoid the loss of orientational selectivity for frequently encountered configurations that is a consequence of mapping such configurations to the axis of a cylindrical index space. This is described in Section 7.

7. INDEX CONSTRUCTION

After determining the shape and orientation methods to be used, the attribute axes must be integrated into an index space. As the orientation attribute is inherently circular, it is natural to use some form of polar coordinates, as in the index spaces described in [2], which are structured as the projection of a set of polar coordinates into a four-dimensional hypercylinder. This scheme can be extended quite naturally to configurations of four or more points.

For each indexed configuration, the value of the eccentricity vector determines both the angular and radial polar coordinates. The magnitude of the eccentricity vector is used

as one component of a basis for the shape space. $2n-5$ other basis components (to be selected as described in Section 5), in addition to one size component, become the linear axes of the hypercylindrical index space.

The drawback of this scheme is that all configurations for which the eccentricity vector is zero are mapped to points along the hypercylinder axis, and no information on their orientation will be represented in the index.

8. CONCLUDING REMARKS

A point-based index for four-point configurations can be constructed with attribute axes for size, “shape”, and orientation. Although orientation-independent search for point configurations can be supported by an index based only on size and shape, extending the index structure by adding a closed (e.g., circular) orientation axis requires some method to determine an orientation attribute value for each configuration. A method was presented for deriving an “eccentricity vector” that represents the deviation of the configuration from a symmetrical one. Using direction component of this vector as the orientation attribute value and the magnitude component as one element of the “shape” attribute value, configurations can be mapped into a hypercylindrical index space. However, there exists a set of pathological configurations which result in computation of a zero-length eccentricity vector. Such configurations map to the central axis of the cylindrical hyperspace, and the index will not exhibit orientational selectivity with respect to them. For applications that expect to index large numbers of such figures, the technique of rotation attribute normalization can be used, at the cost of a severe reduction in orientational selectivity for all indexed configurations.

9. REFERENCES

- [1] C. B. Cranston and H. Smote. Efficient position-independent iconic search using an r-theta index. In *Proceedings of the 14th ACM International Symposium on Advances in Geographic Information Systems*, pages 27–34, Arlington, VA, Nov. 2006.
- [2] C. B. Cranston and H. Samet. Indexing point triples via triangle geometry. In *Proceedings of the 23rd IEEE International Conference on Data Engineering*, pages 936–945, Istanbul, Turkey, Apr. 2007.
- [3] E. J. Groth. A pattern-matching algorithm for two-dimensional coordinate lists. *The Astronomical Journal*, 91(5):1244–1248, May 1986.
- [4] S. Roweis, D. Lang, K. Mierle, D. Hogg, and M. Blanton. *Making the Sky Searchable: Fast Geometry Hashing for Automated Astrometry*. Available as http://cosmo.nyu.edu/hogg/research/2006/09/28/astrometry_google.pdf on the web.
- [5] H. Samet. *Foundations of Multidimensional and Metric Data Structures*. Morgan-Kaufmann, San Francisco, 2006.
- [6] P. B. Stetson. In B. Barbuy, E. Janot-Pacheco, A. M. Magalhães, and S. M. Viegas, editors, *V Advanced School of Astrophysics, Image and Data Processing/Interstellar Dust*. Instituto Astrômico e Geofísico, Universidade de São Paulo, Brazil, 1989. An alternative description is at http://nedwww.ipac.caltech.edu/level15/Stetson/Stetson5_2.html.