UvA-DARE (Digital Academic Repository)

Arrangement: a spatial relation between parts for evaluating similarity of tomographic section

Tagare, H.D.; Vos, F.M.; Jaffe, C.C.; Duncan, J.S.

DOI

[10.1109/34.406653](https://doi.org/10.1109/34.406653)

Publication date

1995

Published in

IEEE Transactions on Pattern Analysis and Machine Intelligence

[Link to publication](#)

Citation for published version (APA):

Tagare, H. D., Vos, F. M., Jaffe, C. C., & Duncan, J. S. (1995). Arrangement: a spatial relation between parts for evaluating similarity of tomographic section. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17, 880-893. <https://doi.org/10.1109/34.406653>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Arrangement: A Spatial Relation Between Parts for Evaluating Similarity of Tomographic Section

Hemant D. Tagare, *Member, IEEE*, Frans M. Vos, Conrad C. Jaffe, James S. Duncan, *Member, IEEE*

Abstract—Medical tomographic images are formed by the intersection of the image plane and an object. As the image plane changes, different parts of the object come in view or drop out of view. However, for small changes of the image plane, most parts continue to remain visible and their qualitative embedding in the image remains similar. Therefore, similarity of part embeddings can be used to infer similarity of image planes. Part embeddings are useful features for other vision applications as well.

In view of this, a spatial relation called “arrangement” is proposed to describe part embeddings. The relation describes how each part is surrounded by its neighbors. Further, a metric for arrangements is formulated by expressing arrangements in terms of the Voronoi diagram of the parts.

Arrangements and their metric are used to retrieve images by image plane similarity in a cardiac magnetic resonance image database. Experiments with the database are reported which

- 1) validate the observation that similarity of image planes can be inferred from similarity of part embeddings, and
- 2) compare the performance of arrangement based image retrieval with that of expert radiologists.

Index Terms—Arrangement, image databases, recognition from parts, medical image processing, tomography, Voronoi diagram.

I. INTRODUCTION

GEOMETRICALLY speaking, a medical tomographic image is formed by the intersection of a subject's anatomy and the image plane (Fig. 1a). The same anatomical feature viewed in different image planes can have substantially different appearance. Consequently, most useful image comparisons are done with images obtained in similar planes.

This observation is crucial in creating a query-by-pictorial-example medical image database. A user queries such a database by providing an example image and indicating different anatomical features of interest. The database retrieves other images containing similar features. The user is expected to compare the anatomical features in the retrieved images with the example image and gain additional insight. For a meaningful comparison, the retrieved images should have image planes similar to the example image. Hence, a mechanism which can

Hemant D. Tagare is with the Department of Diagnostic Radiology, Yale University, New Haven, CT 06520-8042, and with the Department of Computer Science, Yale University, CT 06520-8285.

Frans M. Vos is with the Department of Mathematics and Computer Science, University of Amsterdam, Amsterdam, The Netherlands.

Conrad C. Jaffe is with the Department of Diagnostic Radiology, Yale University, New Haven, CT 06520-8042.

James S. Duncan is with the Department of Diagnostic Radiology, Yale University, New Haven, CT 06520-8042 and with the Department of Electrical Engineering, Yale University, New Haven, CT 06520-2157.

IEEECS Log Number P95060.

Fig. 1. Tomographic view and section.

infer similarity of image planes from image content can be put to good use in the database. In this paper, we propose such a mechanism and demonstrate its use in a cardiac image database.

The mechanism is based on an intuitive understanding of how a tomographic image changes as the imaging plane is changed. To explain it, we begin with a discussion of imaging geometry (with the help of Fig. 1). Fig. 1a shows the heart imaged in parallel planes. The planes are displaced from each other along the common normal. This is the standard imaging practice.

The image in each plane is the intersection of the plane and the heart (Fig. 1b shows the image formed in the plane labeled Section 1). In medical imaging terminology, the entire set of images are said to have the same “tomographic view,” or simply the same “view.” Any image in the set is said to have a particular “tomographic section,” or simply a “section.” Thus, view refers to the common normal of the image planes while section refers to the normal as well as the location of the planes.

If two image planes have similar sections (such as Section 1 and Section 2 in Fig. 1c), then most of the chambers of the heart visible in one image remain visible in the other. Furthermore, the chambers continue to surround each other in a similar pattern. Figs. 1b and 1d illustrate a case where the two ventricles and the two atria are visible and each is surrounded by exactly the same chambers. Intuitively, it is clear that the converse also holds. If two images contain same chambers, each surrounded similarly by neighbors, then we may conclude that the images have similar section.

The key idea of this paper is that the intuitive phrase "surrounded similarly by neighbors" can be made precise enough to design an algorithm which infers image plane similarity. In fact, two definitions are proposed to describe how chambers surround each other—that of the "local arrangement" around a chamber and that of an "overall arrangement" (or simply the "arrangement") in the image.

The local arrangement around a chamber is found by determining all its neighbors and noting the counterclockwise sequence in which they occur. For example, in Figs. 1b and 1d, the left ventricle has the right ventricle and the left atrium for neighbors (The right ventricle is too far away to be a neighbor. This is explained in Section II). The counterclockwise sequence in which they occur is the right ventricle followed by the left atrium. Thus the local arrangement around the left ventricle can be expressed by the string

left_ventricle: right_ventricle left_atrium.

The initial substring "*left_ventricle:*" indicates that the elements of remaining substring are neighbors of the left_ventricle. The substring "*right_ventricle left_atrium*" indicates the counterclockwise sequence in which the neighbors occur. Any circular shift of this substring is considered equivalent.

The overall "arrangement" is simply the set of all local arrangements in the image. The arrangement in Figs. 1b and also in 1d) is the set:

$$\left(\begin{array}{lll} \text{left_ventricle:} & \text{right_ventricle} & \text{left_atrium} \\ \text{right_ventricle:} & \text{right_atrium} & \text{left_atrium} \\ \text{right_atrium:} & \text{left_atrium} & \text{right_ventricle} \\ \text{left_atrium:} & \text{left_ventricle:} & \text{right_ventricle right_atrium} \end{array} \right)$$

Since arrangements are defined in terms of neighbors, they are intimately connected to Voronoi diagrams. We find it useful to present some facts about Voronoi diagrams in Section II before formally defining arrangements in Section III.

In Section IV, changes in arrangement are analyzed by considering changes in the Voronoi diagram. This leads to the diagonal exchange operation and the definition of the metric between arrangements in Section V.

We use arrangements in a query-by-pictorial-example cardiac magnetic resonance image (MRI) database. The query mechanism allows combinations of features (such as shapes and sizes) to be used during retrieval. Here, we only report the part of the mechanism which deals with similarity of image planes.

Experiments with the database are reported in Section VI. The experiments validate the intuition that images with similar arrangements have similar image planes.

Arrangements describe how parts are embedded in the image and are likely to be useful in other applications. A number of authors consider part embeddings as important image features for human as well as computer vision [4], [5], [6], [7], and [19]. Section VII contains a brief discussion of how arrangements may be used in other applications.

Finally, we note that, in theory, it is possible to construct an object whose arrangements are similar for vastly different sections. This is not a serious problem since it leads only to additional retrievals in the database which maybe rejected by specifying more visual features (such as the size or shape of chambers) in the query.

II. VORONOI DIAGRAMS

For generality, we refer to the chambers as parts. Suppose that $N - 1$ parts, denoted P_i , $i = 1, \dots, N - 1$, are embedded in an image. They are assumed to be nonintersecting points or regions. Each region is assumed to be a closed and bounded set, and to have a simply connected interior. All parts are assumed to lie inside a region of interest O (Fig. 2a). O' , the complement of O , is the "outside," and is an open set. Let $\tilde{O}'$ be the closure of O' (Fig. 2a). It is convenient to treat $\tilde{O}'$ as another part, and it is denoted P_N .

(a) Parts and the region of interest

(b) The Voronoi diagram

Fig. 2. Parts, Voronoi diagrams, and arrangements.

The Voronoi domain of a part P_i is the set of all points in the plane that are closer to P_i than to any other part. The Voronoi domain of P_i is denoted V_i and is given by

$$V_i = \{x \mid d(x, P_i) < d(x, P_j), \forall j \neq i\},$$

where, $d(x, P_i)$ is the Euclidean distance of the point x from the part P_i . This distance is the minimum of the Euclidean distance from x to points in P_i .

$$d(x, P_i) = \min\{\|x - y\|, y \in P_i\}.$$

Points that are equidistant from two or more P_i do not belong to any Voronoi domain. The set of all such points, which is the complement of $\bigcup_{i=1,N} V_i$ is the Voronoi diagram (also called the exoskeleton or the skiz [13]) of P_i [1], [12].

Voronoi domains and diagrams have the following properties [13] (Fig. 2b):

- 1) Each Voronoi domain V_i is a connected open set.
- 2) The Voronoi diagram is a closed set, but it is not necessarily connected. Fig. 2b shows an example where it is not connected.
- 3) Locally, the Voronoi diagram is a finite union of simple arcs [13]. The point of intersection of two arcs is a vertex (see Fig. 2b). The arcs minus the vertices are the edges.
- 4) The only structurally stable configurations for Voronoi diagrams are those with every vertex having exactly three edges incident on it. By this we mean the following: If P_i are placed in such a manner that some vertex has more than three edges incident on it, then arbitrarily small displacements of P_i immediately rearrange the Voronoi diagram so that the vertex under consideration splits into a number of vertices with only three edges incident on each resulting vertex. Conversely, given a Voronoi diagram with only three edges incident on every vertex, small changes of part positions do not change the number of edges incident at any vertex.
- 5) There are no dangling arcs in the Voronoi diagram [13], i.e., arcs either begin and end at a vertex (such as "B" in Fig. 2b) or are closed curves (such as "A" in Fig. 2b). From this, it follows that any connected component of a Voronoi diagram is either a single closed arc (such as Arc A in Fig. 2b), or has more than one vertex (such as Arc B in Fig. 2b), with each vertex having exactly three edges incident on it. We call the second type of connected component a cubic component since all of its vertices have degree three. In the rest of this paper, it is assumed that the Voronoi diagram has a single cubic connected component. In most cases of practical interest, the Voronoi diagrams do have a single cubic component. In fact, the Voronoi diagrams in all of our experiments have single cubic connected components.
- 6) Two parts P_i, P_j are neighbors if and only if their Voronoi domains are separated by an edge of the Voronoi diagram. This is the definition of neighbors.
- 7) Since the Voronoi diagram is a plane graph it has a dual graph [18].

A. Representation for the Voronoi Diagram

We represent the Voronoi diagram by a modified incidence matrix. The matrix is constructed thus:

- 1) For every vertex of the diagram, the three parts P_i, P_j , and P_k that created the vertex are noted in a counterclockwise manner. The vertex is labeled as $P_i P_j P_k$. Labels $P_i P_j P_k, P_k P_i P_j$, and $P_j P_k P_i$ are considered the same. It is shown in the appendix that each vertex has a unique label.

- 2) Every edge of the Voronoi diagram is labeled $P_i P_j$ if it separates part P_i from P_j .
- 3) A modified incidence matrix is created whose entry in the p th row and q th column is the edge that connects the p th vertex to the q th. If two vertices are not connected by a single edge, the entry is 0, and all diagonal entries are 1. For example, the modified incidence matrix of the cubic connected component of the Voronoi diagram of Fig. 2b is shown in Table I.

TABLE I
MODIFIED INCIDENCE MATRIX

	$P_1 P_3 P_6$	$P_2 P_1 P_6$	$P_2 P_6 P_4$	$P_4 P_6 P_3$	$P_2 P_4 P_3$	$P_2 P_3 P_1$
$P_1 P_3 P_6$	1	$P_1 P_6$	0	$P_3 P_6$	0	$P_1 P_3$
$P_2 P_1 P_6$	$P_1 P_6$	1	$P_2 P_6$	0	0	$P_1 P_2$
$P_2 P_6 P_4$	0	$P_2 P_6$	1	$P_4 P_6$	$P_2 P_4$	0
$P_4 P_6 P_3$	$P_3 P_6$	0	$P_4 P_6$	1	0	$P_3 P_4$
$P_2 P_4 P_3$	0	0	$P_2 P_4$	$P_3 P_4$	1	0
$P_2 P_3 P_1$	$P_1 P_3$	$P_1 P_2$	0	0	0	1

Two Voronoi diagrams are considered label-isomorphic if there is a one-to-one correspondence between their vertices and edges, so that the corresponding labels are the same, and with that correspondence, the modified incidence matrices are identical.

III. ARRANGEMENT

The boundary of the Voronoi domain of P_i can be traversed by keeping its Voronoi domain to the left. In doing so, if $P_{j_1}, P_{j_2}, \dots, P_{j_n}$ is the sequence of neighbors that form the boundary, then the local arrangement around P_i is

$$P_i : P_{j_1}, P_{j_2}, \dots, P_{j_n},$$

where any circular shift of $P_{j_1}, P_{j_2}, \dots, P_{j_n}$ is considered equivalent. The convention results in a counterclockwise traversal of the boundary of each part except the "outside." For the "outside" the boundary is traversed clockwise.

Repeating the procedure for the entire set of parts we get the set of strings $A(P_1, \dots, P_N)$

$$A(P_1, \dots, P_N) = \begin{pmatrix} P_1 : \dots \\ P_2 : \dots \\ \dots \\ P_N : \dots \end{pmatrix},$$

where, each string gives the local arrangement around a part. The set $A(P_1, \dots, P_N)$ is the arrangement of $P_1, \dots, P_N$.

Two arrangements $A_1(P_1, \dots, P_N)$ and $A_2(P_1, \dots, P_N)$ are identical if and only if the sets $A_1()$ and $A_2()$ are equal (any circular shift of the substring following " $P_i :$ " is considered equivalent).

In terms of the Voronoi diagram, the arrangement $A(P_1, \dots, P_N)$ is the list of cycles that make the boundaries of its faces. These cycles completely define the plane embedding of the graph structure of the Voronoi diagram.

Two arrangements $A_1(P_1, \dots, P_2)$ and $A_2(P_1, \dots, P_N)$ can be shown to be identical iff their Voronoi diagrams are label-

isomorphic (see the Appendix). This completely characterizes equality of arrangements.

When arrangements are not equal, their (dis-)similarity can be computed by measuring the "effort" it takes to change one arrangement into another using the diagonal exchange operation discussed below. The diagonal exchange can be repeatedly applied to a given Voronoi diagram to transform it to another and number of applications of the diagonal exchange measure the "effort" required to transform one arrangement into another. In the physics literature, the diagonal exchange is known as the 70 operator [1], [11].

IV. THE DIAGONAL EXCHANGE

The diagonal exchange is illustrated in Fig. 3. Consider the edge e in Fig. 3a. If the four Voronoi domains shown in the figures arise from four distinct parts P_1 , P_2 , P_3 , and P_4 , the diagonal exchange is the transformation that changes the configuration of Fig. 3a to the configuration of Fig. 3b. The edge e which separated the Voronoi domains of P_2 and P_4 before the exchange, separates the Voronoi domains of P_1 and P_3 after the exchange. The rest of the Voronoi diagram is unaffected.

Fig. 3. The diagonal exchange.

Fig. 3 also shows the change in the dual graph of the Voronoi diagram. Before the exchange, the dual edge of e forms the diagonal of the quadrilateral shown in the figure. After the exchange, this diagonal is replaced by the other diagonal of the quadrilateral. The rest of the dual graph remains the same. The replacement of the diagonal is the reason for calling this operation the diagonal exchange.

The diagonal exchange arises because the structurally stable vertices of the Voronoi diagram have only three edges incident on them. Referring to the Voronoi diagram of Fig. 3a, consider what happens to the vertex v as P_1 moves to the right. The vertex v , for the most part remains structurally the same as before, but the distance between it and v' decreases. At some point, the vertices v and v' merge, creating a vertex V which has four edges incident on it. However, this vertex is unstable. Small motion of the part P_1 to the left or right either gives back the original configuration or gives a configuration obtained by the diagonal exchange.

The diagonal exchange can be seen to have the following properties:

- 1) It maintains the connectivity of the Voronoi diagram as well as the dual graph.
- 2) The number of edges and vertices of the Voronoi diagram as well as the dual are conserved during the operation.
- 3) It is reversible.

Exchanging the same edge twice gives back the original Voronoi diagram and the dual graph.

If a sequence of diagonal exchanges on the edges $e_n, e_{n-1}, \dots, e_1$ transforms one Voronoi diagram D_1 into another diagram D_2 , then the sequence of exchanges on $e_1, \dots, e_{n-1}, e_n$ transforms D_2 back into D_1 , i.e., every sequence Γ of diagonal exchanges has an inverse sequence Γ^{-1} such that the composite sequence $\Gamma^{-1} \Gamma$ leaves the Voronoi diagram unaltered.

- 4) Successively exchanging the same edge an odd number of times gives the same result as exchanging the edge once. Successively exchanging it an even number of times gives the same result as not exchanging it.

The diagonal exchange is useful because it can be repeatedly applied to a Voronoi diagram to change it to any other Voronoi diagram caused by an embedding of the same parts. This is not immediately apparent because the diagonal exchange is defined only when the parts P_1 , P_2 , P_3 , and P_4 associated with the edge are distinct. In a given Voronoi diagram it is quite likely that this condition is not met for some edges, and to transform the diagram a sequence of exchanges which do not use such edges has to be found. Below, we demonstrate a constructive algorithm that achieves this [14]. The algorithm shows that any cubic connected Voronoi diagram can be reduced by a sequence of diagonal exchanges to a canonical diagram called the flower. The flower is shown in Fig. 4. In it, the domains of $P_2, \dots, P_{N-1}$ appear as "petals" attached to a central domain in a counterclockwise manner. The central domain is the Voronoi domain of P_1 . In the flower, the local arrangement of P_1 is " $P_1 : P_2 P_N P_3 P_N \dots P_{N-1} P_N$," the local arrangement of P_N is " $P_N : P_1 P_{N-1} P_1 P_{N-2} \dots P_2$," and the local arrangement of all other parts is " $P_j : P_N P_1$."

Suppose that Γ_1 is the sequence of diagonal exchanges that convert one Voronoi diagram into the flower and that Γ_2 is a sequence that converts another Voronoi diagram into the flower, then by property 3 of diagonal exchanges Γ_2^{-1} exists and the sequence $\Gamma_2^{-1} \Gamma_1$ converts the first Voronoi diagram into the second (Fig. 4a).

The sequence of diagonal exchanges that reduce a Voronoi diagram to a flower repeatedly uses certain subsequences of exchanges. These subsequences are called contact reduction, neighbor reduction, and domain propagation. Below, these are explained first and the complete algorithm is explained later.

Fig. 4. The canonical transformation.

The following discussion assumes that the number of parts is greater than or equal to 4 ($N \geq 4$). If the number of parts is three, then the cubic connected graph is already a flower (Fig. 4b) and there is no need to apply diagonal exchanges to reduce it to one.

A. Contact Reduction

Suppose P_i and P_j are neighbors whose Voronoi domains are separated by more than one edge. Let e be one such edge and let it be incident on vertices v_1 and v_2 (Fig. 5). If P_i, P_j, P_k are parts whose Voronoi domains create v_1 , and P_i, P_j, P_k are parts whose Voronoi domains create v_2 , then from the connectivity of Voronoi Domains it follows that $P_k \neq P_k$. Thus the edge e can be operated upon by the diagonal exchange. Before the exchange, the edge separates the Voronoi domains of P_i and P_j . Therefore, after the exchange, the Voronoi domains of P_i and P_j are not separated by the edge e , and the number of edges separating them reduces by one. By repeatedly applying this to other edges shared between P_i and P_j , the number of edges separating the two parts can be reduced to one.

The process can be repeated between other neighbors of P_i so that at the end of this step, P_i has as many neighbors as before, but is separated from each by exactly one edge (Fig. 5).

Fig. 5. Contact reduction.

B. Neighbor Reduction

Suppose that P_i ($P_i \neq P_N$) has three or more neighbors and that exactly one edge in the Voronoi diagram separates P_i from each of its neighbors. Then the boundary of the Voronoi domain of P_i has as many edges as P_i has neighbors (Fig. 6). Since this number is greater than or equal to three, every three consecutive edges separate three distinct neighbors. Hence, the diagonal exchange can be applied to any edge belonging to the boundary of the Voronoi diagram of P_i . After the exchange, P_i loses the neighbor that was originally separated by the edge and the number of its neighbors reduces by one.

This can be repeated until P_i is left with exactly two neighbors, each of which is separated from P_i by exactly one edge in the Voronoi diagram. The configuration is shown in Fig. 6 and is called a lens.

Fig. 6. Neighbor reduction.

C. Domain Propagation

Suppose P_i has a Voronoi domain which is a lens. Its boundary has two edges and two vertices. Let v be any one of the vertices. Since v is cubic it has an edge other than the two belonging to the boundary of the Voronoi domain of P_i . Suppose this edge is incident at the other end on a vertex v' . If v' happens to be the other vertex of the boundary of P_i then these two vertices along with the three edges form a connected component for exactly three parts (Fig. 7a). As mentioned above, the connected Voronoi diagram of three parts is already a flower and we need not consider any operations on it.

Fig. 7. Domain propagation.

If v' does not belong to the boundary of P_i , then the three faces that are incident at v' are P_j , P_k , and P_i , with $P_i \neq P_j$. (If $P_i = P_j$, then the Voronoi domain of P_i is not connected—a contradiction). Therefore, the diagonal exchange can be applied to the edge between v and v' with the result shown in Fig. 7b. Now the four domains P_i , P_j , P_k , and P_l are all different, any of the three edges of the boundary of P_i is exchangeable. Exchanging any one of them inserts the region P_i as a lens into one of the edges incident at v' (Fig. 7b).

By repeating this, the lens can be propagated along any sequence of edges and inserted into any edge by a series of diagonal exchanges.

D. Transformation to a Flower

The above sequences of diagonal exchanges can be combined to reduce any Voronoi diagram to a flower. This is as follows: assume that there are at least four parts. Let parts P_1 and P_m ($m \neq N$) be neighbors and e be one edge from vertex v to v' separating the domain of P_1 from P_m . For the moment also assume that $m \neq N-1$ (this case is dealt with below).

By contact reduction, the Voronoi diagram can be changed so that e is the only edge separating P_1 and P_m . Next, the Voronoi domain of the part labeled P_2 can be reduced to a lens by contact and neighbor reduction. The Voronoi diagram remains connected during this process and domain propagation can be used on the lens to propagate and insert it in the edge e (Fig. 8b). This splits e into two segments.

By a similar process, P_3 can be inserted as a lens in the segment of e that occurs in a counterclockwise travel of the boundary of the domain P_1 from lens P_2 . Continuing this, we insert P_4 , P_5 , ..., P_{m-1} , P_{m+1} , ..., P_{N-1} and obtain the graph of Fig. 8c.

The next series of exchanges reduces the domain of P_m to a lens between the domains of P_{m-1} and P_{m+1} . The sequence of exchanges has two stages. The initial steps of the first stage are

Fig. 8. Reduction to a Flower.

shown in Fig. 8d and Fig. 8e. Note that after these two steps, the domain P_m has lost contact with P_{N-1} . Iterating these steps causes the domain of P_m to lose contact with P_{N-2} , P_{N-3} , ..., P_{m+1} . At the end of this stage, the Voronoi diagram looks like Fig. 8f. Repeating this process from the other end as shown in the figure reduces P_m to a lens situated on the boundary of the Voronoi domain of P_1 , located between the Voronoi domains of P_{m-1} and P_{m+1} (Fig. 8g). The Voronoi diagram is thereby reduced to a flower.

The case when $m = N-1$ is similar to Fig. 8 except that the exchanges going from (c) to (d) to (e) to (f) are not required. The exchanges from (f) onwards may be directly applied to the graph obtained at (c).

Using this procedure, every cubic, connected Voronoi diagram can be transformed by a sequence of diagonal exchanges to a flower. Consequently, any two connected Voronoi diagrams formed by the embedding of the same parts can be transformed into each other by diagonal exchanges.

V. THE ARRANGEMENT METRIC

The above algorithm demonstrates the existence of one sequence of diagonal exchanges for transforming a Voronoi diagram into another. The sequence of exchanges is not unique. If an edge e_1 is exchanged in the sequence, then substituting that exchange with any odd number of consecutive exchanges of e_1 will have the same effect (by property 4 of diagonal exchanges). Therefore, there are infinitely many sequences of diagonal exchanges that transform one Voronoi diagram into another.

If $d(A_1, A_2)$ is the minimum number of diagonal exchanges required to transform the Voronoi diagram of arrangement A_1 to the diagram of arrangement A_2 , then

- 1) $d(A_1, A_2) \geq 0$,
 - 2) $d(A_1, A_2) = d(A_2, A_1)$. This is because $d(A_1, A_2)$ is the minimum number of diagonal exchanges and because diagonal exchanges are reversible.
 - 3) $d(A_1, A_2) = 0$ iff the Voronoi diagrams of A_1 and A_2 are label-isomorphic.
 - 4) $d(A_1, A_2) \leq d(A_1, A_3) + d(A_3, A_2)$, with equality iff A_3 occurs in the minimum exchange sequence from A_1 and A_2 . Since $d(A_1, A_2)$ is the sequence with the minimum number of exchanges, all subsequences of the sequence also have the minimum number of exchanges for transforming the starting diagram of the subsequence to the final diagram of the subsequence.
- Therefore, $d(A_1, A_2)$ is a metric.

The metric can be easily extended to compare arrangements of sets of parts that are not identical. Let $\{P_i\}$ and $\{P_i^*\}$ be two different sets of parts that have some parts in common. Parts that are not common to both sets cannot be compared by the relation "arrangement" since the change in those parts is not of position but of existence—a change which the relation is not designed to consider. On the other hand, parts that are common to both sets remain embedded and can be compared.

Therefore, the embeddings of $\{P_i\}$ and $\{P_i^*\}$ can be compared by computing the metric between arrangements of the parts that are common to both.

A. Computing the Metric

The numerical value of the arrangement metric can be computed with the help of the diagonal exchange tree (Fig. 9a). The nodes of the tree represent arrangements and the branches represent exchangeable edges. Traversing the tree from a node to one below it involves starting from the arrangement at the first node, exchanging the edge corresponding to the branch in the tree, and obtaining the arrangement at the node at the end of the branch.

Fig. 9. The exchange tree.

If an arrangement A_1 is the root of the tree and $d(A_1, A_2) = k$, then A_2 first appears in the tree at depth k . Therefore, $d(A_1, A_2)$ can be obtained by a breadth first search of the exchange tree with A_1 at the root of the tree.

The following strategies effectively prune the exchange tree and reduce the complexity of search:

- 1) From the discussion of the diagonal exchange, we know that exchanging an edge twice in succession gives the original Voronoi diagram back. Hence, beginning with arrangement A and exchanging an edge e_i twice in succession (see Fig. 9b) results in being at a node with exactly the same arrangement as A . The subtree of this node is identical to the subtree of the initial node, and since that subtree is already being searched, this subtree need not be searched.
- 2) If e_1 and e_2 are two exchangeable edges in an arrangement, and e_1 and e_2 are incident on different vertices,

then the arrangement obtained by exchanges e_1e_2 and e_2e_1 are identical. Therefore, the two subtrees shown in Fig. 9c are identical and only one of them need be searched.

Note that every edge in an arrangement has at most four other edges which share its two vertices. Therefore, most pairs of edges in a Voronoi diagram do not have a vertex in common and this rule can be applied to all such pairs.

We use breadth first search with the pruning strategies given above to compute the arrangement metric in the experiments.

VI. EXPERIMENTS

As described in Section I, we expect two cardiac images whose arrangements are similar to have the same or similar sections. Retrieval by similarity of section is one of the modules in a query-by-pictorial-example image database system that we are currently developing.

In this database, as each image is entered, the outlines of the chambers are manually traced, labeled, and stored along with the image. The outlined chambers are any of left ventricle, right ventricle, left atrium, and right atrium that may be present in each image. The ascending aorta, the pulmonary artery, and the region of interest are also outlined. The Voronoi diagrams of the outlined parts are computed and saved along with the database. Currently, there are about a hundred images in the database.

The outlines are created only once and are viewed as a fixed data entry cost for every image. The outlines are traced with a mouse without the help of any refinement such as active contours or deformable templates. Three operators were involved in outlining and the outlines reflect the varying accuracies of the operators. Although we did not make any attempt to quantify operator variation, the results reported below use all of the outlines and are not based on using only the best outlines.

Once the images have been entered into the database, the user selects an example image at run time and asks to retrieve all images with similar section. Every image in the database is compared with the example image in the following way:

- 1) If P_{ex} is the set of parts in the example image and P_{im} is the set of parts in the image being compared, then the sets

$$Diff_1 = P_{ex} - P_{ex} \cap P_{im},$$

$$Diff_2 = P_{im} - P_{ex} \cap P_{im},$$

are computed. $Diff_1$ gives the parts in P_{ex} that are not present in P_{im} and $Diff_2$ gives the parts in P_{im} that are not present in P_{ex} .

If $Diff_1 \geq 2$ or $Diff_2 \geq 2$, the image being compared is discarded since there are far too many differences in the parts that are present.

- 2) Otherwise, the arrangement metric is computed between the parts that are common to the example image and the image being compared (i.e., parts in the set $P_{ex} \cap P_{im}$). Since the metric is integer valued, the resulting image is binned according to whether the numerical value of the

metric is 0, 1 or ≥ 2 . This scheme was obtained after some experimentation with the database and is likely to differ for other applications.

- 3) After the images are binned, they are presented to the user in the sequence 0, 1, ≥ 2 . It is expected that images with low metric have the same or similar section as the example image.

Fig. 10 shows an example image. This image is a section from an "axial" view. Fig. 11 shows the outlined chambers and Fig. 12 shows the computed Voronoi diagram. Note that each face of the Voronoi diagram does enclose one part, and that each vertex of the Voronoi diagram has degree three. The results of binning after comparison with the example images are shown in Figs. 13-15. Fig. 13 shows typical images from metric = 0 bin, Fig. 14 shows typical images from metric = 1 bin, and Fig. 15 shows typical images from metric ≥ 2 bin¹. In each of the figures, the example image is presented in the top left corner for comparison. Also the Voronoi diagrams have been superimposed on the images. As can be seen from the figures, the majority of the images in bins 0 and 1 are of the same or similar section as the example image, while those in bin ≥ 2 are not.

Another example of retrieval is shown in Figs. 16-19. Fig. 16 shows the example image. Fig. 17 shows typical images from the metric = 0 bin, Fig. 18 shows typical images from metric = 1 bin and Fig. 19 shows typical images from metric ≥ 2 bin. The Voronoi diagrams are omitted for clarity. It should be noted that in spite of the large variation in the geometry of the heart and the large variation in the position of the heart in the chest, the arrangement metric robustly retrieves appropriate sections.

With regard to computational efficiency, we report that the current implementation of the algorithm computes the metric and bins the images in about two seconds on a SUN SPARC-Station 2.

The database has been implemented and experimented with for over a year. Different sections have been used as examples and images with similar section successfully retrieved from the database. One of the authors² is a radiologist and has verified the performance.

Besides subjective evaluations, two experiments were conducted to statistically compare the performance of the method with that of expert radiologists. These experiments are described below. Three expert cardiac radiologists participated in the two experiments which were conducted immediately in succession. These radiologists were unaware of the notion of "arrangement" and of any of the theory developed so far. A database of 100 cardiac MRI images were used in both experiments. Three images were chosen from this set to serve as query images. The query images had tomographic sections in the axial, sagittal, and coronal views.

1. Note that the rightmost image in the bottom row in Fig. 14 does not contain the left atrium, therefore its Voronoi diagram is compared to the Voronoi diagram of the example image with the left atrium in the example image deleted. The Voronoi diagram of the example image with the left atrium deleted is not shown in Fig. 12 or in Fig. 14.

2. C. Jaffe.

Fig. 10. Section in the axial view: Example image.

Fig. 11. Section in the axial view: Outlined chambers.

Fig. 12. Section in an axial view: Voronoi diagram.

Fig. 13. Result of binning using the axial section: Metric = 0.

Masked versions of all the database and query images were also produced by blanking out all of the image content except the heart and the four chambers, i.e., no background information was present in the masked images.

A. The First Experiment: Validation of Arrangement as Indicator of Tomographic Section

The first experiment was designed to validate the intuition that arrangements of chambers provide a strong clue about tomographic section. Two ideas were tested:

- 1) that most of the information regarding tomographic section is contained in the arrangement, i.e., masking out all of the information except the arrangement of chambers does not result in a significant loss of a radiologist's ability to deduce section, and
- 2) that the performance of the arrangements technique was similar to that of expert radiologists.

The unmasked and masked database images were mixed together in one group and the unmasked and masked query images were mixed together in a query group for this experiment. The three participating radiologists were asked to rank similarity of section in the following manner: Each radiologist was shown one query image at a time. Upon presentation of the query image the radiologist was asked to compare each of the database images with the query image and give it a similarity rating based on Table II.

TABLE II
SIMILARITY RATINGS FOR RADIOLOGISTS

Similarity Rating	Meaning
1	The two images have approximately the same view and section.
2	The two images have approximately the same view, but are from slightly different sections.
3	The intermediate between 2 and 4.
4	The two images have approximately the same view, but their sections are very different.
5	The two images are from different views

No definitions of the terms "slightly different" and "very different" in Table II were provided. We expected the radiologists to define these in terms of everyday radiology practice.

After the experiment, the similarity ratings for unmasked and masked images were separated and tabulated as shown in Table III (in this table $s_{...}$ indicates the similarity rating).

A T-test revealed that the two columns of Table III were statistically different ($P < 0.001$). However, the average difference between the ratings in the two columns was only 0.018 ± 0.061 , revealing that although masking the background does cause a difference in the performance of radiologists, the difference is quite small.

Next, the radiologists' ratings were compared to the arrangements method. First, the radiologists' ratings were used to create a "representative" rating. Since there were three radiologists, three ratings were available for every database image. The median of these ratings was chosen representative.

TABLE III
TABULATION OF SIMILARITY RATINGS

Similarity Ratings	
Unmasked Query Unmasked Dbase Image	Masked Query Masked Dbase Image
s_{1u}	s_{1m}
...	...
s_{100u}	s_{100m}
...	...
s_{900u}	s_{900m}

Table III had 900 entries (database images $\times$ query images $\times$ physicians).

Fig. 14. Result of binning using the axial section: Metric = 1.

Fig. 15. Result of binning using the axial section: Metric ≥ 2 .

Fig. 16. Retrieval for a coronal section: Example image.

Then, the arrangements method was used with the query images, and similarity ratings for it were given as shown in Table IV.

TABLE IV
SIMILARITY RATINGS FOR ARRANGEMENTS

Similarity Rating	Meaning
1	The two images are 0 diagonal exchanges apart.
2	The two images are 1 diagonal exchange apart.
3	The two images are 2 diagonal exchanges apart.
4	The two images are more than two diagonal exchanges apart.
5	There is more than one part not in common to the two images so they cannot be compared.

The representative and arrangement similarity ratings were tabulated in two $3 \times 100 \times 2$ tables (queries $\times$ images $\times$ similarity ratings for masked and unmasked images). Similarity ratings in both tables were compared using Spearman's rank-correlation test. The correlations are shown in Table V.

TABLE V
SPEARMAN'S RANK CORRELATION OF SIMILARITY RATINGS

Unmasked Query Unmasked Dbase Image	Masked Query Masked Dbase Image
0.64	0.66

Both correlations were significant ($P < 0.001$). A visual comparison of the similarity ratings revealed that most expert and arrangement similarity ratings either agreed or were off by a value of 1.

This experiment indicates that tomographic section can be inferred reliably from the embedding of the heart chambers in the image even when the rest of the image information is suppressed and that the similarity ratings of the arrangements technique are positively correlated to those of the radiologists.

B. The Utility of Arrangements in Image Retrieval

The second experiment evaluated the utility of the arrangements technique for retrieval by similarity of tomographic section. The masked query and database images were not used in this experiment since they are not representative of real-world MRI image databases.

The idea behind this experiment was to compare the performance of a radiologist acting as a database retriever to that of the arrangement technique. To achieve this, the set of 100 images was chosen as an image database and a set of three images chosen as queries.

- First, a "gold standard" was created which tabulated the images that a database engine ought to retrieve for each of the queries (this is described below).
- Second, an independent radiologist acting as a database engine ranked images which appeared to him to have sections similar to the query images.
- Third, the arrangements technique was used to rank sections similar to the sections of the query images. Accuracies of retrieval by the radiologist and the arrangements technique were computed and compared with respect to the gold standard.

The gold standard was established by the three radiologists who had participated in the first experiment. Upon presentation of a query image, the radiologists were asked to decide whether each image in the database ought to be retrieved for a "retrieval by similarity of section" query. If at least two radiologists agreed that an image should be retrieved, then it was considered appropriate to retrieve the image. This was taken to be the gold standard.

Fig. 17. Retrieval for a coronal section: Metric = 0.

Fig. 18. Retrieval for a coronal section: Metric = 1.

Fig. 19. Retrieval for a coronal section: Metric ≥ 2 .

TABLE VI
ERRORS BY THE EXPERT RADIOLOGIST

Expert	$\delta = 0$	$\delta = 1$	$\delta = 2$	$\delta = 3$	$\delta = 4$
Retrieval errors in a database of 100 images for three queries	25	42	56	58	59

TABLE VII
ERRORS BY THE ARRANGEMENTS METHOD

Arrangements	$\epsilon = 0$	$\epsilon = 1$	$\epsilon = 2$	$\epsilon = 3$
Errors in a database of 100 images for three queries.	57	69	77	76

TABLE VIII
CONTINGENCY TABLE OF ERRORS

	Desired Retrievals	False Retrievals	Desired Rejects	False Rejects	Accuracy (%Desired retr.+%Desired rejects)
Expert	25	8	250	17	92%
Arrangement	30	45	213	12	81%

Next, the independent radiologist was asked to interpret the database and the query images using the following scheme: Each image was labeled with a two-dimensional tag: (view, section). The view field v could take three values: sagittal, axial and coronal. The section field s could take five values (1, ..., 5), going from anterior to posterior, from superior to inferior or from left to right (depending on the view). After the radiologist had tagged all database and query images, a database image $I_{db}(v_{db}, s_{db})$ was considered desirable for retrieval for a query image $I_q(v_q, s_q)$ if both had the same view ($v_{db} = v_q$) and if the absolute difference of their section-values was smaller than a threshold δ ($\|s_{db} - s_q\| \leq \delta$), which was varied from 0 to 4 in steps of 1. For each value of δ two kinds of errors were measured by comparing the retrieved images with the gold standard. Errors of false retrievals were images that were retrieved by the radiologist that were not retrieved in the gold standard, and errors of false rejects were images retrieved by the gold standard that were not retrieved by the radiologist. Table VI summarizes the error (= false retrievals + false rejects) in retrieval made by the radiologist as a function of the threshold δ .

Next, the arrangements technique was used to compare each query image with all the database images and the database images were binned in the order of their metric. For a query image, an image in the database was considered to have been retrieved if the arrangement metric between the two was less than or equal to a threshold ϵ ($d(A_q, A_{db}) \leq \epsilon$). The threshold was varied in steps of 1 from 0 to 3 and by comparing the retrieved images at each step to the gold standard; the error (= %false retrievals + %false rejects) in retrieval was measured as above. Table VII shows the retrieval error as a function of ϵ .

Since the smallest errors were made for $\delta = 0$ and $\epsilon = 0$, these values of threshold were chosen for a more detailed comparison. The images which were retrieved at $\epsilon = 0$ and $\delta = 0$ were classified as desired retrievals (image was retrieved by the gold standard and the method), false retrievals (as defined before), desired rejects (image was not retrieved by the gold standard or the method), and false rejects (as defined before). The results are shown in Table VIII.

The χ^2 -test was used to compare the two rows of Table VIII presupposing (H_0) that the row and the column variable were independent. The test rejected the hypothesis ($\chi^2 = 30.1$

($P < 0.001$)) showing that the performance of the the expert and arrangements was dissimilar.

The last column of the table shows the accuracy (desired retrievals + desired rejects) of the expert and the arrangements method. The performance of the arrangements method is slightly worse than that of the expert. This should be considered in light of the fact that the expert uses the entire information in the image and that the arrangements method uses only relative position information. Adding further features like relative shapes and sizes of the chambers can be expected to improve the performance of the retrieval algorithm.

The main difference in the accuracy of the arrangements method and the expert is in the false retrievals column—the arrangements method has more false retrievals (15%) than the expert (2.6%). If the false positives column is deleted from the table, no significant difference is found between performance of the two ($\chi^2 = 1.77$, not significant at $P = 0.05$).

This experiment demonstrates the utility of including arrangements in a tomographic image database. Although the accuracy of the technique is lower than that of an expert radiologist, the decrease in accuracy is due to extra retrievals. In practice, such an increase in false retrievals is not detrimental to the performance of the database. Given the variability between experts, a small increase in false retrievals might in fact be desirable.

VII. REMARKS ON ARRANGEMENTS

Arrangements and the diagonal exchange operator are useful because together they completely specify the neighborhood system of embedded parts and can transform one neighborhood system into another. Any "qualitative" property of the embedding which depends on the neighborhood system can be expressed in terms of the arrangement.

An additional advantage of using arrangements to describe qualitative properties is that properties which may be intuitively clear for point-like parts can be easily extended to region-like parts since arrangements are defined for both. As an example, consider the notion of directional relations between neighbors. Intuitively, they are most easily defined for point-like parts. By expressing them in terms of the arrangement, they can be meaningfully extended to region-like parts.

Let $P_1, \dots, P_{N-1}$ be points embedded in the plane and P_N be the "outside." Suppose that P_1, P_2 , and P_3 are neighbors and P_3 lies to the left of the directed line P_1P_2 . Then, traversing the boundary of the Voronoi domain of P_1 (Fig. 20a), we encounter the sequence $\dots P_2 P_3 \dots$ so that one string in the arrangement is

$$P_1 : \dots P_2 P_3 \dots \quad (1)$$

If P_3 lies to the right of P_1P_2 , then the string is

$$P_1 : \dots P_3 P_2 \dots \quad (2)$$

Fig. 20. Directional relations.

It is clear that the converse also holds, i.e., the occurrence of a string such as that of (1) implies that P_3 is to the left of P_1P_2 and the occurrence of a string such as that of (2) implies that P_3 is to the right of P_1P_2 . Therefore, directional relations can be deduced from arrangement.

Having observed that directional relations between point parts can be deduced from arrangements, we can easily extend directional relations to parts which are not points. Consider Figs. 20b, 20c, and 20d. The figures show three parts in different configurations. In the progression shown in the figure, part P_3 starts in the configuration of Fig. 20b, extends itself around part P_2 , and pulls up to the configuration in Fig. 20d. Assuming that P_1 is not the "outside," the strings for P_1 in the three arrangements are

$$\text{Fig. 20b: } P_1 : \dots P_3 P_2 \dots$$

$$\text{Fig. 20c: } P_1 : \dots P_3 P_2 P_3 \dots$$

$$\text{Fig. 20d: } P_1 : \dots P_2 P_3 \dots$$

Using the definition of directional relations in terms of arrangement - the first string shows that P_3 is to the right of P_1P_2 , while the last string shows that P_3 is to the left of P_1P_2 , which is as expected. On the other hand, the middle string shows that P_3 is simultaneously to the right as well as to the left of P_1P_2 (the middle string is of the type $P_1 : \dots P_2 P_3 \dots$ as well as the type $P_1 : \dots P_3 P_2 \dots$). This makes sense, visually as well as semantically. Since P_3 is not a point, it can wrap around other parts in a variety of complex ways and can simultaneously appear to the left and right of other parts.

A point cannot be simultaneously to the left and right of another pair of points, whereas a region can. Using arrangements to express directional relations gives a consistent way of expressing this.

VIII. CONCLUSIONS AND DISCUSSION

We proposed a qualitative spatial relation called "arrangement" which describes how different parts are embedded in an image. Arrangements are closely related to Voronoi Diagrams, and based on this, a metric was defined between arrangements.

It was also proposed that similarity of part embeddings in medical tomographic images implies similarity of tomographic section. This observation, along with the arrangement metric, was used in a cardiac MRI database to retrieve images having sections similar to an example image.

Experiments which

- 1) demonstrate the utility of the metric in retrieving images with similar section, and
- 2) compare the effectiveness of this technique with that of expert radiologists were reported.

These experiments showed that the arrangements technique had a comparable though slightly inferior performance to that of expert radiologists.

APPENDIX

PROPERTIES OF LABELED VORONOI DIAGRAMS

This Appendix contains technical results used in the main paper. The main aim is to establish the equivalence of the arrangement and the labeled Voronoi diagram. We begin by showing that vertex and edge labels are unique.

A. Vertex Labels

Using Fig. 21 we present a geometrical demonstration that each vertex of the Voronoi diagram has a unique label. Each vertex is labeled by noting, in a counterclockwise manner, the parts whose domains define the vertex. Since stable vertices have degree three, the label for any vertex is of the form $P_i P_j P_k$.

The result is established by assuming that there are two vertices in a Voronoi diagram that are labeled $P_i P_j P_k$ and $P_i P_j P_l$ and showing that $P_k = P_l$. Since both vertices have P_i as well as P_j in their label strings, the vertices occur on the boundary of the Voronoi domain of P_i as well as P_j . Fig. 21a shows the boundary of the Voronoi domain of P_i with the two vertices in it. A closed path can be traced by starting from the vertex $P_i P_j P_k$ continuing along the boundary of the Voronoi domain of P_i to vertex $P_i P_j P_l$ then proceeding along the boundary of the Voronoi domain of P_j and arriving back at the starting vertex. There are two ways of doing this: either as shown in Fig. 21b or as shown in Fig. 21c (the distinction between the two is that in Fig. 21b P_l lies inside the path, while in Fig. 21c P_l lies outside the path). In both situations it can be seen that the Voronoi domains of P_i and P_j lie on different sides of the closed path, and therefore, the closed path separates them.

Fig. 21. Vertices have a unique label.

Since the path is always through the Voronoi diagram, it follows that the Voronoi diagram separates the Voronoi domains of parts P_i and P_j , i.e., the two are distinct parts.

Therefore, $P_1P_2P_1 \neq P_1P_2P_j$ and each vertex has a unique label.

B. Identifying Edges

Since vertices have unique labels, every edge in the Voronoi diagram may be identified by the labels of the vertices it is incident on. However, more than one edge may be incident on the same two vertices and it is necessary to distinguish between them.

Edges incident on the same vertices can be distinguished by the Voronoi domains they separate. Since a vertex is a point that is equidistant from three distinct parts, the three edges incident on any vertex separate three distinct parts, and each of the edges is uniquely identified by the two parts that it separates.

C. Arrangement and Label-Isomorphism

The arrangement and the modified incidence matrix of a Voronoi diagram convey the same information. To demonstrate this, consider the construction of the modified incidence matrix from the arrangement. The presence of a substring $P_i: \dots P_jP_k \dots$ or $P_i: P_k \dots P_j$ in the arrangement implies that P_i , P_j , and P_k are neighbors whose Voronoi domains meet at a vertex. If $P_i \neq P_N$, the vertex has a label $P_iP_jP_k$. If $P_i = P_N$, the vertex has a label $P_NP_kP_j$. Applying this rule to all the strings in the arrangement, the entire set of labeled vertices of the Voronoi diagram can be obtained. Of course, since each vertex occurs in the boundary of three Voronoi domains, applying the rule to every string in the arrangement generates each label three times. But since each vertex has a unique label, detecting repetition is straight-forward.

The set of labeled edges can also be obtained from the arrangement. If $P_i \neq P_N$ and the local arrangement of P_i is given by $P_i: P_jP_kP_l \dots P_m$ or $P_i: \dots P_mP_jP_kP_l \dots$, then there is an edge P_iP_j between vertices $P_iP_mP_j$ and $P_iP_jP_k$. If $P_i = P_N$ and the local arrangements are given as above, then the edge P_iP_j is still present except that the vertices that it joins are $P_NP_jP_m$ and $P_NP_kP_j$.

Applying this procedure to all strings in the arrangement, we obtain the entire set of edges. The algorithm counts each edge twice since each edge occurs on the boundaries of two

Voronoi domains. However, repetitions are easy to spot since an edge is uniquely determined by its label and the labels of the vertices it is incident on.

Constructing the arrangement from the modified incidence matrix is also straight forward. For every P_i a list of vertices whose labels contain P_i is constructed. All of these vertices, and only these vertices, lie on the boundary of the Voronoi domain of P_i . Suppose $P_iP_jP_k$ is a vertex in this list. If $P_i \neq P_N$, then we seek a vertex $P_iP_kP_l$ in the list which is connected to it by an edge P_iP_k . If $P_i = P_N$ and the original vertex is $P_NP_jP_k$, then we seek a vertex $P_NP_jP_l$ connected by an edge P_NP_j . Such vertices always exist since all vertices with the P_i in their label necessarily lie on a cycle of the Voronoi diagram.

Continuing the above, all vertices in the list are linked circularly. The circular list can be traversed to obtain the local arrangement of P_i . If $P_i \neq P_N$, and $P_iP_jP_k, P_iP_lP_j, \dots$ are consecutive vertices on the circular list, the local arrangement is $P_i: P_jP_kP_l \dots$. If $P_i = P_N$, and $P_NP_jP_k, P_NP_lP_j, \dots$ are consecutive vertices on the circular list, the local arrangement is $P_N: \dots P_kP_jP_l$.

ACKNOWLEDGMENTS

We gratefully acknowledge the discussions with Marjorie Senechal, Pradeep G. Mutalik, and Lisa R. Berlinger. We are also grateful to the three expert radiologists for participating in the experiments. We acknowledge the helpful suggestions made by the editors of PAMI and the reviewers of this manuscript.

This work was supported by the grant 1-R01-LM05007 from the National Library of Medicine.

REFERENCES

- [1] F. Aurenhammer, "Voronoi diagrams-A survey of a fundamental geometric data structure," *ACM Computing Surveys*, vol. 23, no. 3, Sept. 1991.
- [2] D.H. Ballard, C.M. Brown, *Computer Vision*. Prentice Hall, 1982.
- [3] J.W. Bruce, P.J. Giblin, *Curves and Singularities*. Cambridge University Press, 1984.
- [4] H.D. Ellis "Theoretical aspects of face recognition," *Perceiving and Remembering*, G. Davies, et al., eds. Academic Press, 1981.
- [5] M.A. Fischler, R.A. Eschlager, "The representation and matching of pictorial structures," *IEEE Trans. Computers*, vol. 22, no. 1, pp. 67-92, 1973.
- [6] R.L. Frantz, "The origin of form perception," *Scientific American*, vol. 204, no. 5, pp. 66-72, 1961.
- [7] E.C. Hildreth, S. Ullman, "The computational study of vision," *Foundations of Cognitive Science*, M.I. Posner, ed. M.I.T. Press, 1989.
- [8] B. Julesz, *Foundations of Cyclopean Perception*. The University of Chicago Press, 1971.
- [9] M.H.A. Newman, *Elements of the Topology of Plane Sets of Points*. Cambridge University Press, 1964.
- [10] D.I. Perret, A.J. Mistlin, M.H. Harries, "Seeing faces: The representation of facial information in temporal complex," *Seeing Contour and Color*, J.J. Kulikowski, et al., eds. Pergamon Press, 1989.
- [11] N. Rivier, "Continuous random networks: From graphs to glasses," *Advances in Physics*, vol. 36, no. 1, pp. 95-134, 1987.
- [12] M. Senechal, *A Brief Introduction to the Mathematics of Quasicrystals*, M. V. Jaric, ed. Academic Press, 1989.
- [13] J. Serra, *Image Analysis and Mathematical Morphology*, Academic Press, 1982.

- [14] H.D. Tagare, C.C. Jaffe, J.S. Duncan, "Arrangement: A Qualitative spatial relation between parts for image retrieval," Tech. Rep. 9602, Center for Systems Science, Dept. of Electrical Engineering, Yale University, April 1992.
- [15] H.D. Tagare, C.C. Jaffe, J.S. Duncan, "Arrangement: A spatial relation comparing part embeddings and its use in medical image comparisons," *Information Processing in Medical Imaging*, H.H. Barrett, A.F. Gmitro, eds. *Lecture Notes in Computer Science*, pp. 687. Springer-Verlag, 1993.
- [16] S. Truve, W. Richards, "From Waltz to Whitman," *Natural Computation*, W. Richards, ed., M.I.T. Press, 1988.
- [17] D. Weaire, N. Rivier, "Soap, cells and statistics-Random patterns in two dimensions," *Contemp. Phys.*, vol. 25, no. 1, 59-99, 1984.
- [18] R.J. Wilson, *Introduction to Graph Theory*. Longman Scientific and Technical, 1985.
- [19] A.W. Young, D.C. Hay, and K.H. McWeeny, "Right Cerebral Hemisphere Superiority for Constructing Facial Representations," *Neuropsychologia*, vol. 23, pp. 195-202, 1985.

Hemant D. Tagare (S'86-M'90) received the B.Tech degree from the Indian Institute of Technology, Bombay, in 1981, the MS degree from Rice University, Texas in 1983, and the Ph.D. degree from Rice University, Texas in 1990, all in electrical engineering.

Dr. Tagare is an assistant professor with the Department of Diagnostic Radiology at Yale University and is also an associate research scientist with the Computer Science Department at Yale University. His current research interests are

computer vision for medical images and mobile robots—in particular, the definition of nonrigid shape-based correspondence and its use in medical images, jitter resistant statistics for medical image data, and fast indexing schemes for place recognition.

Dr. Tagare is a member of IEEE.

Frans Vos obtained the M.Sc. degree in 1993 from University of Amsterdam in the areas of Medical Informatics and Computer Science. The research described in this paper was conducted at Yale University as a part of his master's thesis. He is now a Ph.D. student at the Free University of Amsterdam. His current research focuses on 3D shape reconstruction of the cornea.

James S. Duncan (S'72, M-'75, SM-'93) received the BSEE degree from Lafayette College, Easton, Penn in 1973, the MS degree in engineering from UCLA in 1975 and the PhD degree in electrical engineering from the University of Southern California in 1982.

In 1973, he joined the staff of Hughes Aircraft Company, Electro-Optical and Data Systems Group, and participated in research and development projects related to signal and image processing for forward looking infrared (FLIR) imaging systems until 1983. During this time, he held Hughes' Masters,

Engineer and Doctoral Fellowships. In 1983, he joined the faculty of Yale University, New Haven, Conn., where he currently is an associate professor of diagnostic radiology and electrical engineering. His research and teaching efforts have been in the areas of image processing, computer vision and medical imaging. His current specific research interests include the segmentation of deformable objects from both 2D and 3D data, the tracking of non-rigid object motion from 2D and 3D data, and the integration of processing modules in vision systems, all with a special interest in using these approaches for medical image analysis.

Dr. Duncan is a member of Eta Kappa Nu and Sigma Xi, is on the editorial board of the *Journal of Mathematical Imaging and Vision*, and is currently an associate editor for the *IEEE Transactions on Medical Imaging*. He was a Fulbright Research Scholar at the Universities of Amsterdam and Utrecht in the Netherlands during part of the 1993-94 academic year.

Conrad C. Jaffe received his B.S. degree in humanities and biology from MIT in 1965 and his MD degree from Columbia University College of Physicians and Surgeons in 1969. After postdoctoral work at MIT's Research Lab of Electronics, he joined Yale University School of Medicine where is now a professor of diagnostic radiology and medicine (cardiology). His research interests include mechanisms of diagnostic imaging techniques, medical image processing, observer variation, and computer-aided instruction. He is a specialist in diagnostic cardiac imaging, particularly echocardiography.