



On the problem of adjacency relations in the Spatial Aggregation approach

Liliana Ironi, Stefania Tentoni

IMATI - CNR

via Ferrata 1, 27100 Pavia, Italy

email addresses: {ironi,tentoni}@imati.cnr.it

Abstract. Spatial Aggregation (SA) is a computational approach to the analysis of large spatial data sets. It differs from other tools for data analysis for its hierarchical strategy in aggregating spatial objects at higher and higher levels until the behavioral and structural information about the underlying physical phenomenon, that is required for performing a specific task, is extracted from the data set. This characteristic makes SA an interesting and versatile framework for the development of tools for automated reasoning about physical phenomena spatially represented. The SA approach has been successfully applied to different domains and tasks; but, its soundness strictly depends on the definitions of spatial adjacency relations at different levels of abstraction. The definitions given in [Huang and Zhao2000] may reveal to be not fully sound when performing the contouring task. In such a context, the found drawbacks are at least twofold: (1) metric-based adjacency relations are too difficult to be optimally defined, and aggregation may fail to extract isocurves; (2) the soundness of neighborhood relations depends on the density of the isopoint set, and aggregation may fail to extract physical structures as isopoint contiguity may be lost. This paper illustrates the problems above, and presents new definitions and algorithms for their solutions.

Keywords: Data analysis, spatial aggregation, spatial reasoning, contouring.

Introduction

Our interest in the SA computational paradigm [Yip and Zhao1996, Bailey-Kellogg *et al.*1996] has been motivated by our long-term goal to build a framework for the automated interpretation of cardiac potential maps gathered by endocavitary probes. The map interpretation task is essential for the localization of the anomalous excitation sites associated with ventricular arrhythmia and other electrical conduction pathologies. Automating such a task would facilitate the introduction of cardiac maps into the clinical practice with a consequent great impact on health care. Conventional contour analysis allows us to identify patterns of electrical potential distribution but it does not facilitate the automated extraction of general rules necessary to infer the correlation between pathophysiological patterns and wavefront structure and propagation. The hier-

archical structure of the objects abstracted by the SA computational approach to represent and interpret a numeric input field is a promising ground (i) to capture spatial and temporal adjacencies at multiple scales, (ii) to identify causal relations between events, and (iii) to generate high-level explanations for prediction. The potential of SA is confirmed by its application to different physical contexts [Joskowicz and Sacks1991, Yip1997, Bailey-Kellogg and Zhao1997, Bailey-Kellogg and Zhao2001, Zhao1995] among those we mention the weather data analysis problem [Huang2000] as it presents some similarities with the electrocardiographic map interpretation: global patterns and structures are identified by looking at the isocurves and reasoning about their spatial relations.

In outline, the SA computational framework allows us to transform a numeric input field into a symbolic, parsimonious description of the structure and behavior of the physical process the field is associated with. The desired output results from successive transformations of lower-level abstract objects, called *spatial aggregates*. Each spatial aggregate is built by applying a common data structure, the *neighborhood graph* (n -graph), and a small set of generic operators. The n -graph encodes the spatial contiguity between the objects, and the operators, which are identical at each level of abstraction, make explicit the neighborhood and equivalence relations. These latter relations deal with two distinct concepts, namely *strong* and *weak adjacency*, that express intra and inter-relations, respectively. The former one binds a set of contiguous spatial objects into a spatial aggregate in accordance with a distinctive property; the latter one highlights the relations between the spatial objects aggregated at the previous level.

In a contouring context, at the first level of aggregation, (i) the n -graph defines the topological contiguities between isopoints, (ii) the strong adjacency relation identifies isocurves, whereas (iii) weak adjacency allows us to discover and reason on different global patterns. The classification of contiguous isopoints having the same isovalue as strong adjacent ones is a rather delicate issue. The strong adjacency definition originally proposed [Huang and Zhao2000, Huang2000] is based on the metric distance of the points in accordance with a user-supplied threshold value. Such an

approach may turn out to be unsound in many situations due to the difficulty to choose the optimal threshold value: too large or too small values may be responsible for isocurve entanglement and/or segmentation phenomena. Another possible cause of failure deals with a sparse density of isopoints on isocurves. As a matter of fact, the n-graph is built by exploiting strategies which locally express spatial contiguity but ignore the values of the physical variables associated with the geometric domain, and, consequently, points that do belong to the same isocurve may not result to be connected. This problem has been highlighted in [Huang and Zhao2000,Huang2000] but no solution of practical value has been proposed.

This paper discusses the problems above, and presents definitions and algorithms for a sound aggregation and classification of isopoints. More precisely, it gives (1) an algorithm for the n-graph generation capable of checking and fixing contiguity loss situations, and (2) an equivalence predicate which defines a strong adjacency relation between isopoints by exploiting topological adjacency properties rather than a metric distance.

Contouring: problems with SA

Given an input field, the SA computation process follows an iterative procedure which terminates as soon as the required behavioral and structural information for performing a specific task has been achieved: spatial objects at a given level of abstraction are transformed into ones at a higher abstraction level using neighborhood and equivalence relations defined on that set of objects. The overall process goes through three main steps (Figure 1). The first step, called *aggregation*, builds a n-graph from field primitive objects to explicate their spatial contiguity. This step is initialized with primitive objects which result from a pre-processing operation on the input field aimed at organizing the available data in a form suitable to perform the task. Then, an equivalence relation is defined on contiguous objects to characterize them with respect to properties of interest for the reasoning task. This is the *classification* step upon which, through the *redescribe* operator, the abstraction of new spatial objects is eventually accomplished.

Since our interest for the application of SA to contouring does not merely lie in visualizing a map, that is what other graphical tools efficiently do [Lorenson and Cline1987], but we rather want to produce abstract spatial objects, hierarchically organized, upon which further high-level reasoning can be performed. To achieve this goal we cannot leave out a thorough analysis of all the steps SA performs to transform an input field into isocurves. In the following, we will focus our discussion on each step to highlight the related problems, and propose solutions.

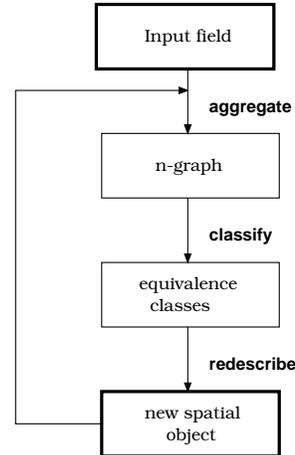


Figure 1: Main steps in SA.

Input field

Initial data for the 2D contouring problem consist of:

1. a geometrical domain discretized by a mesh made up of a set of elements $\{E_i\}_{i=1..NE}$, and of nodes, vertices of the elements, $\{\mathbf{x}_i\}_{i=1..N}$;
2. the scalar values $\{u_i\}_{i=1..N}$ of a function $u(\mathbf{x})$ either computed or measured at the mesh nodes;
3. a couple of scalar values $(\bar{u}_0, \Delta u)$ to uniformly scan the range of $u(\mathbf{x})$ and generate isocurves for the levels $\bar{u}_0, \bar{u}_0 + \Delta u, \bar{u}_0 + 2\Delta u, \dots, \bar{u}_0 + n\Delta u$.

Due to the frequent irregularities in the shape of realistic geometric domains as well as to the need to adequately capture the gradient of u , the given mesh is rarely uniform: its elements, usually triangles or quadrangles, can widely vary in size and shape. We here assume that mesh data can be given in either triangular or quadrangular form.

Let us observe that the optimal choice of the scan step Δu is usually suggested by domain-specific knowledge, and is critical in so far as it should be small enough to capture the function main changes over the geometric domain, but large enough to produce a contour map that is readable and not burdened with an unnecessary mass of minor, possibly artifactual, details.

Pre-processing The available input data are processed to derive and organize all the auxiliary knowledge that is needed to successfully perform the abstraction processes eventually aimed at the production of contour maps in the form of SA high level objects.

First, isopoints for the required levels are generated by assuming a piecewise-linear approximation $\tilde{u}(\mathbf{x})$ for the function $u(\mathbf{x})$, with $\tilde{u}(\mathbf{x}_i) = u(\mathbf{x}_i)$ for each mesh node \mathbf{x}_i . The isopoint set, \mathcal{I} , is generated by comparison of the values of u at the vertices of each element with the required levels, and by linear interpolation of the nodal values when necessary.

Isopoint aggregation - Possible failure of the n-graph to capture topological contiguities of isopoints

The spatial contiguity between isopoints is explicitly coded in the n-graph through the application on \mathcal{I} of the *aggregate* operator. Such operator uses a neighborhood specification, the definition of which is task-dependent. For example, a minimal spanning tree can be a n-graph well-suited to abstract features that rely on minimum-distance propagation pathways.

As regards the contouring task, the n-graph should reflect the local contiguity of isopoints. A natural way to express contiguity between points is by triangulation. A Delaunay triangulation [George and Borouchaki1998] is, for example, a rather good choice due to its higher regularity with respect to others possible: it guarantees that, for each element, the smallest angle has the largest possible width. But, isopoints are not merely geometrical points as they express a context-dependent property, namely the value of u . As a matter of fact, in our context, a Delaunay graph may fail, since it just connects points without taking into account the values of u . Then, it may happen - depending on the spatial course of u - that *ab initio* isopoints which truly belong to the same isocurve are not connected through the n-graph, whereas isopoints of two topologically non-adjacent isocurves are connected. Figure 2 illustrates this undesired situation: the isopoints A and B on the isocurve $\bar{u} + \Delta u$ are denied contiguity; whereas the edge CD states contiguity between isopoints which belong to topologically non-adjacent isocurves, namely \bar{u} and $\bar{u} + 2\Delta u$. As mentioned in [Huang2000, Huang and Zhao2000] this problem may realistically occur when isopoints are not “dense” enough on each isocurve. Since this latter property depends on the size of the elements, this is not a problem when the input data are from a uniform mesh with a sufficiently small element diameter, but it needs care when such conditions are not met, which often occurs with a generic mesh. It has been proved in [Huang2000] that if the set \mathcal{I} satisfies the *closeness condition*, then the n-graph \mathcal{D} generated by a Delaunay triangulation of \mathcal{I} guarantees that no contiguity loss occurs. But, unfortunately, it is objectively difficult to verify such a condition [Huang2000] within a contouring context as it explicitly relies on curves whose identification is the final goal.

The n-graph plays a crucial role in the classification phase as it establishes spatial contiguities upon which the strong adjacency relation is applied to generate isocurves. As a matter of fact, an inadequate n-graph might induce, by effect of possible isopoint contiguity loss, undesired *curve segmentation*. While it is objectively difficult to generate a triangulation that *a priori* guarantees avoiding the risk of topological contiguity losses, it is however necessary to identify risky situations and adopt some strategy to avoid the occurrence

of such events.

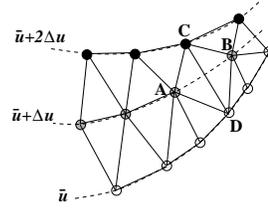


Figure 2: Isopoint contiguity loss after a triangulation. White, grey, and black nodes denote the values \bar{u} , $\bar{u} + \Delta u$, and $\bar{u} + 2\Delta u$, respectively; true isolines are represented by dashed lines; solid lines mark contiguities by the n-graph.

Isopoint classification - Possible failure of the metric-based strong adjacency relation

The classification step is based on the definition of an equivalence relation on contiguous elements of \mathcal{I} . Isocurves, which are abstracted as new higher level spatial objects through the *redescribe* operator, correspond to the classes resulting from the application of the *strong adjacency relation* on contiguous elements of \mathcal{I} . Such classes are redescribed as polylines, namely they instantiate the strong adjacencies holding between their constituent points.

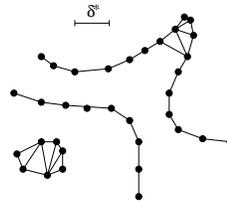


Figure 3: Metric-based strong adjacency (D1): curve entanglement may occur when the curve turns sharply with a high curvature.

The success of the abstraction process with respect to a given task relies on the appropriateness of both the n-graph, which is the device explicating spatial contiguity, and the definition of the strong adjacency relation, which implicitly identifies the new object “isocurve” through a distinctive property shared by its low-level components. To make this issue clear, let us consider the definition of strong adjacency between isopoints given in [Huang and Zhao2000, Huang2000], limited to the contouring task. Let $\mathcal{A}_{\mathcal{D}}(\mathbf{x})$ be the set of the contiguous nodes of $\mathbf{x} \in \mathcal{I}$ within the n-graph \mathcal{D} .

Definition D1. $\forall \mathbf{x}, \mathbf{y} \in \mathcal{I}$ and such that $\mathbf{y} \in \mathcal{A}_{\mathcal{D}}(\mathbf{x})$, \mathbf{x} and \mathbf{y} are said to be *strongly adjacent* if and only if $u(\mathbf{x}) = u(\mathbf{y}) \wedge d(\mathbf{x}, \mathbf{y}) < \delta^*$, where $d(\mathbf{x}, \mathbf{y})$ is the Euclidean distance between \mathbf{x} and \mathbf{y} , and $\delta^* > 0$ is a user-defined suitable nearness radius.

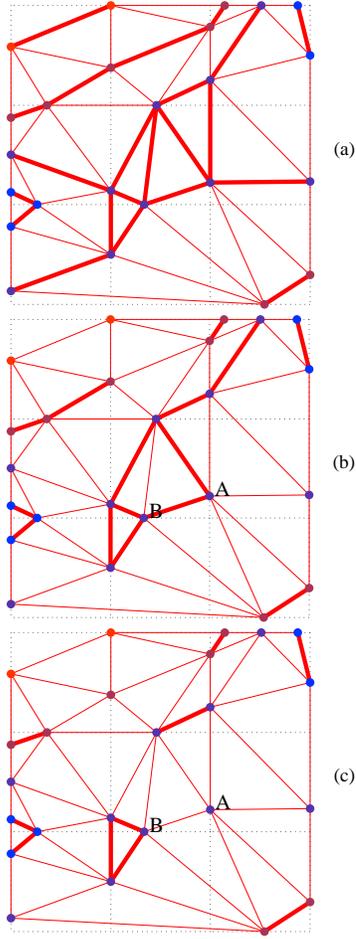


Figure 4: Metric-based strong adjacency (D1): curve entanglement (a), and (b); curve segmentation (b), and (c). Thick lines connect strongly adjacent isopoints; thin solid lines connect contiguous points within the n-graph. Thin dotted lines correspond to the mesh.

If definition D1 is used, a few problems may arise in the classification of isopoints, leading to incorrect isocurve abstraction, namely *curve entanglement* and *curve segmentation*.

A - Curve entanglement. The curve entanglement problem may appear where a single curve makes a sharp bend with a high curvature (Figure 3).

The same phenomenon, as illustrated by Figure 4a, may also occur when too large a nearness radius δ^* is used. In the case illustrated by the figure, the geometrical domain $[0,6] \times [0,6]$ is discretized by a uniform 3×3 square element mesh, and the matrix U of u values¹

¹For the sake of clearness, let us observe that the matrix element $U_{i,j}$ stores the value of u in the domain node $(2 * (i - 1), 6 - 2 * (j - 1))$, $i = 1..4$, $j = 1..4$.

on the 16 nodes is given:

$$U = \begin{bmatrix} 6.70 & 6.00 & 5.10 & 1.00 \\ 5.00 & 3.60 & 2.30 & 2.00 \\ 1.00 & 2.90 & 3.20 & 3.30 \\ 3.30 & 3.10 & 3.90 & 5.00 \end{bmatrix};$$

the required contour levels are $\{1.5, 3.0, 4.5, 6.0\}$, and δ^* is equal to the length of the diagonal of the square elements of the grid.

In the figure, entanglement regards the two curves of level 3.0 in proximity of the saddle point located within the element $[2,4] \times [2,4]$.

B - Curve segmentation. Another undesired classification result, symmetric with respect to entanglement, is *isocurve segmentation* that might derive by using definition D1 with a smaller radius, as illustrated by Figure 4b. In this case, δ^* is chosen equal to the length of the side of the square elements. As a consequence, curve segmentation appears, while curve entanglement is slightly reduced. Let us notice that, to completely avoid entanglement, δ^* should be - *a posteriori* - chosen less than the distance of the points A and B shown in Figure 4b, but this choice would unavoidably leave many points in singleton classes, and produce a diffuse unacceptable fragmentation of the isocurves (Figure 4c).

The previous examples highlight the criticality of the choice of δ^* , and the intrinsic inadequacy of the metric criterion in definition D1. The cited phenomena can appear when uniform regular meshes are given, as in the simple examples we have shown, but are far more frequent in the case of generic meshes which approximate more complex and realistic geometrical domains as those in Electrocardiology. The use of a very fine mesh, which makes isopoints spatially denser but at a higher computational cost, often avoids the problem in practice. But a really robust approach to such an issue is to make strong adjacency depend on topological relations rather than on metric distances.

Solutions

Our solutions to the isopoint aggregation and classification problems presented in the previous section are respectively provided by an algorithm for the identification of contiguity loss situations and their successive recovery, and by a new definition of strong adjacency relation (Figure 5). To tackle, satisfactorily and automatically, the most general case of non uniform meshes, both triangular and quadrangular ones, such a relation exploits the available knowledge related to the mesh data, and grounds on the inter and intra-element topological adjacency properties of isopoints.

Isopoint set generation

During the generation of the set \mathcal{I} in the pre-processing phase of the input field, we also incrementally build two maps, τ and γ , which state the inter

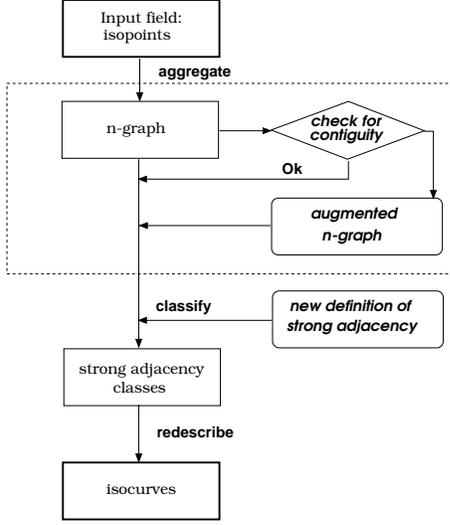


Figure 5: Main steps in the revised approach to SA-based contouring.

and intra-element topological adjacency, respectively. The former, τ , maps each isopoint \mathbf{x} that is added in \mathcal{I} to the set of mesh elements that share \mathbf{x} , that is:

$$\forall \mathbf{x} \in \mathcal{I}, \tau : \mathbf{x} \rightarrow \tau(\mathbf{x}) = \{E_i : E_i \ni \mathbf{x}, i = 1..NE\}.$$

This mapping plays a crucial role in the definition of a strong adjacency relation between isopoints, but is unable to treat complex and even ambiguous situations that may occur with a quadrangular mesh. To this end, we introduce the map γ that is implicitly defined by the assumed approximation for u . For each couple of newly generated isopoints \mathbf{x}, \mathbf{y} belonging to the same element E , γ flags their intra-element topological adjacency in accordance with the result of the comparison of their value $\bar{u} = \tilde{u}(\mathbf{x}) = \tilde{u}(\mathbf{y})$ with the nodal values u_i ($i = 1, \dots, 4$): it holds $\gamma(\mathbf{x}, \mathbf{y}) = 1$ if the segment $\mathbf{xy} = \{\mathbf{z} \in E \mid \mathbf{z} = \theta\mathbf{x} + (1 - \theta)\mathbf{y}, \theta \in [0, 1]\}$ divides E in two regions where the values of \tilde{u} are in agreement with the nodal values u_i . With reference to Figure 6, we have ²:

Cases 1,5: The element E contains either none or just one isopoint, and then, no intra-element relation applies to these cases.

Cases 2,3,6,8: The element E contains two isopoints \mathbf{x}, \mathbf{y} such that the segment \mathbf{xy} leaves the nodal values less than \bar{u} and greater than \bar{u} on opposite sides, respectively. Thus, $\gamma(\mathbf{x}, \mathbf{y}) = 1$.

Case 4: The element E contains four isopoints \mathbf{y}_{ij} , one of them on each side $\mathbf{x}_i\mathbf{x}_j$; then, set $\gamma(\mathbf{y}_{12}, \mathbf{y}_{23}) = \gamma(\mathbf{y}_{34}, \mathbf{y}_{41}) = 1$, and $\gamma(\mathbf{y}_{23}, \mathbf{y}_{34}) = \gamma(\mathbf{y}_{41}, \mathbf{y}_{12}) = 0$.

²in the following, (i_1, i_2, i_3, i_4) is a permutation of $(1, 2, 3, 4)$

Case 7: The element E contains three isopoints, one of them coincides with one of its vertices \mathbf{x}_i , $i = 1..4$. If both the isopoints \mathbf{x}, \mathbf{y} are different from \mathbf{x}_i , $\gamma(\mathbf{x}, \mathbf{y}) = 1$, otherwise $\gamma(\mathbf{x}, \mathbf{y}) = 0$.

Case 9: The opposite vertices, \mathbf{x}_{i_1} and \mathbf{x}_{i_3} , of the element E are isopoints, and both the other two opposite vertices, \mathbf{x}_{i_2} and \mathbf{x}_{i_4} assume a value either greater or lower than \bar{u} , thus $\gamma(\mathbf{x}_{i_1}, \mathbf{x}_{i_3}) = 0$.

Cases 10 - 12: The element E contains either three or four isopoints for which $\gamma(\mathbf{x}, \mathbf{y}) = 1$ in all cases. Let us note that case 12 describes an isoregion of value \bar{u} .

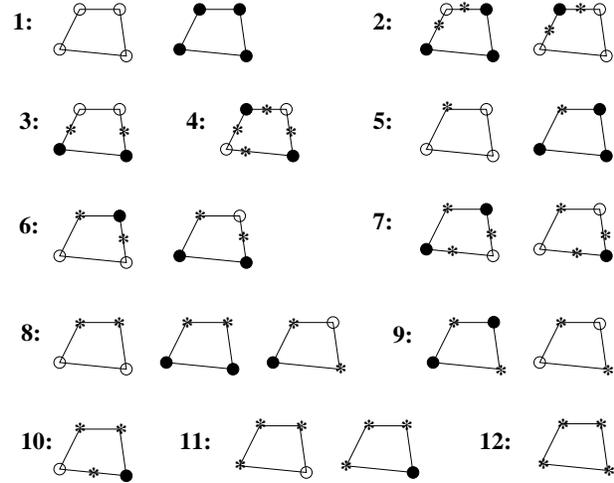


Figure 6: Quadranguluar element: possible locations of isopoints for level \bar{u} . A black circle, a white circle, and a star denote a value lower than \bar{u} , greater than \bar{u} , and equal to \bar{u} , respectively.

As a matter of fact, $\gamma(\mathbf{x}, \mathbf{y}) = 1$ basically means that the segment \mathbf{xy} belongs to the polyline which approximates the isocurve that crosses the element.

Cases from 1 to 4 are the most frequent ones; cases 5 to 7 are less frequent but not rare, while cases 8 to 12 are very unlikely. Case 4, $u_{i_k} < \bar{u} < u_{i_l}$, $k = 1, 3$ and $l = 2, 4$, which corresponds to the presence of a saddle point of u located inside the element, deserves particular attention. In fact this situation is intrinsically ambiguous with respect to which isopoints truly belong to the same curve. Figure 7 illustrates such ambiguity: four new isopoints, one on each side, are generated, and two distinct \bar{u} -valued isopolyines are expected to cross the element leaving the saddle point in between. The choice between a) and b), that marks the two possible ways of connecting the isopoints, is arbitrary at this level of abstraction, and is fixed by γ . We observe that such ambiguity might be naturally removed at a higher level of abstraction if an inconsistent interpretation would occur.

Let us notice that, in the case of a triangular mesh, intra-element adjacency of isopoints is never ambigu-

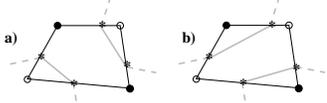


Figure 7: Saddle point: ambiguity in assessing intra-element isopoint adjacency.

ous since for any two isopoints belonging to the same element their associated segment divides the element into regions where the values of \tilde{u} are in agreement with the nodal values u_i ($\gamma = 1$).

Isopoint aggregation: An algorithm for the generation of a suitable n-graph

The adequacy of the candidate n-graph \mathcal{D} generated by Delaunay triangulation to represent the spatial contiguity between isopoints must be verified before isopoint classification is performed. The algorithm given below identifies possible contiguity loss situations, and locally amends \mathcal{D} by adding new connections where necessary.

Algorithm A1 (n-graph generation):

1. Generate \mathcal{D} by Delaunay triangulation of \mathcal{I} .
2. For each node $\mathbf{x} \in \mathcal{D}$
 - consider the set $\mathcal{A}_{\mathcal{D}}(\mathbf{x})$ of its contiguous nodes within \mathcal{D} ;
 - for each node $\mathbf{y} \in \mathcal{A}_{\mathcal{D}}(\mathbf{x})$
 - if $|u(\mathbf{x}) - u(\mathbf{y})| > \Delta u$ then
 - augment \mathcal{D} with any new connections between couples of nodes $\mathbf{z}, \mathbf{w} \in \mathcal{A}_{\mathcal{D}}(\mathbf{x}) \cap \mathcal{A}_{\mathcal{D}}(\mathbf{y})$ such that $|u(\mathbf{z}) - u(\mathbf{w})| \leq \Delta u$

The algorithm introduces new connections where “false” contiguities have locally taken the place of true ones (Figure 8). The data set used in this example results from a slight perturbation of the element U_{32} in the matrix U ($U_{32} = 2.92$). Figure 8 shows a “false” contiguity between the isopoints \mathbf{x} , \mathbf{y} in the graph. The algorithm detects such a situation, and locally adds a connection which re-establishes neighborhood between \mathbf{w} and \mathbf{z} . The resulting n-graph \mathcal{D} provides for a better spatial coverage of any region where adjacent points with unexpectedly far u values are detected. The algorithm introduces only a minimum number of extra connections, which prevents unnecessary computational overload. It can be proved that the overall algorithm keeps the same time complexity of the Delaunay triangulation algorithm, that is $O(n \log n)$ where n is the number of isopoints in \mathcal{I} . Let us remark that the algorithm may end up with a n-graph that does not correspond to a triangulation any more, but this is irrelevant with respect to the abstraction processes that are subsequently performed.

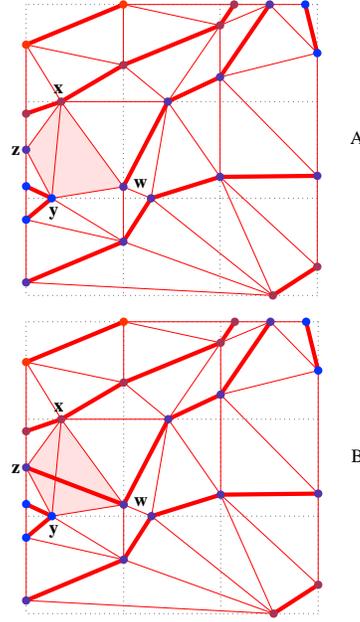


Figure 8: A - The Delaunay n-graph does not correctly explicate topological contiguities between isopoints; B - the n-graph fixed by algorithm A1 recovers contiguity between \mathbf{z} and \mathbf{w} .

Isopoints classification: A new definition of “strong adjacency” relation on \mathcal{I}

The definition of strong adjacency D1 given in the original version of SA establishes spatial adjacencies between aggregate objects using a metric distance, and spatial nearness is assessed when such a distance is within a user-predefined radius. The limits of this approach when applied to high level objects, such as curves, have been discussed in [Huang and Zhao2000]. However, as highlighted in the previous section of this paper, in a contouring context analogous limits can hold at the lowest level of abstraction as well, that is when the aggregate objects are points.

For the contouring task, given the n-graph \mathcal{D} , we define a *strong adjacency* relation on \mathcal{I} as follows:

Definition D2. $\forall \mathbf{x}, \mathbf{y} \in \mathcal{I}$ and such that $\mathbf{y} \in \mathcal{A}_{\mathcal{D}}(\mathbf{x})$, \mathbf{x} and \mathbf{y} are *strongly adjacent* if and only if

- (1) $u(\mathbf{x}) = u(\mathbf{y}) \quad \wedge$
- (2) $\tau(\mathbf{x}) \cap \tau(\mathbf{y}) \neq \emptyset \quad \wedge$
- (3) $\gamma(\mathbf{x}, \mathbf{y}) \neq 0$.

Conditions (2) and (3) in the definition D2 allow us to ground the relation on the topological adjacency properties of the isopoints at element level, which are available from the input mesh data and the pre-processing step. The definition of adjacency grounded on topological properties rather than on a metric distance makes the classification abstraction process more robust with respect to the given task as it does not depend any more on the particular mesh size and shape.

Figure 9 shows the result of classification according to definition D2 of strong adjacency with the same data set used in Figure 4: the use of a topological criterion allows us to a priori avoid the entanglement and segmentation problems which may occur when definition D1 is applied (Figure 4).

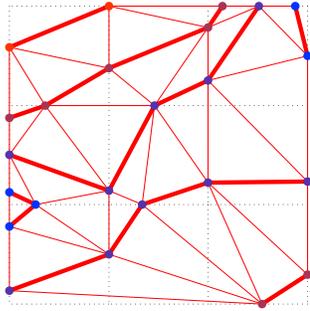


Figure 9: Classification of isocurves built in accordance with definition D2.

SA-based contouring in Electrocardiology

Electrocardiology is a stimulating and promising application domain for qualitative and spatial reasoning. Most research effort aimed at the automated interpretation of electrocardiograms (ECG's) [Bratko *et al.*1989, Weng *et al.*2001, Kundu *et al.*1998, Watrous1995], for which an interpretative rationale is well established. However, a major drawback of ECG's, which record electrical potential from nine sites only on the body surface, is their poor spatial covering. Thanks to the latest advances in technology, far more informative techniques are becoming available: epicardial/endocardial mapping, where electrical potential is either measured by endocavitary probing or obtained non invasively from body surface data through mathematical inverse procedures. Through these techniques a great deal of spatial information about the heart electrical activity is available, but the ability to relate visual features to the underlying complex phenomena still belongs to very few experts [Taccardi *et al.*1998]. Thus, the need for an automated tool for map interpretation to be used in a clinical context.

Figure 10A illustrates a 3D model ventricular geometry. Its discretization was carried out by 15 horizontal sections, 30 angular sectors on each section, and 6 radial subdivisions of each sector. A numerical simulation was carried out on this mesh to simulate the excitation wavefront propagation in the anisotropic ventricular tissue [Colli Franzone *et al.*1998], and for each node x_i the activation time u_i , i.e. the instant at which the excitation front reaches the node, was computed. To investigate the wavefront propagation, the contour maps of the activation time (isochrones) need to be built and analyzed not only on the ventricular surfaces, but also on sections of the wall (Figure10B). An

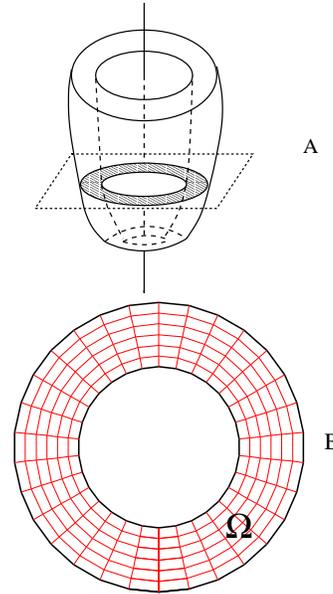


Figure 10: A - Geometry of a ventricle model; B - A horizontal section of the discretized ventricular wall.

activation isochrone map delivers a lot of information about the wavefront propagation: each contour line aggregates all and none but the points that share the same excitation state, and subsequent isocurves are nothing but subsequent snapshots of the travelling wavefront. Figure 11 highlights the performance of SA-based contouring on a horizontal section of this domain: Figure 11A and B show the isochrone maps obtained in accordance with the application of original SA and our revised approach, respectively.

Conclusions

The work here described focusses on the construction of isocurves from a numeric input field within the SA reasoning framework. It gives a significant contribution to original SA as to the contouring problem by providing new algorithms and definitions for soundly building adjacency relations upon 2D complex domain geometry.

An alternative way to solve the unsoundness problem of SA in performing contouring could be its integration with conventional contouring tools. More precisely, we could exploit the former ones to efficiently and soundly construct isocurves, and, afterwards, consider them as primitive spatial objects upon which the SA approach is applied to build spatial relations. But, in this way of proceeding, the hierarchical structure of the built spatial objects would be impoverished of the lower-level spatial aggregates with a consequent loss of information which could result to be relevant for a specific task, for example for the explanations of causal relations between events.

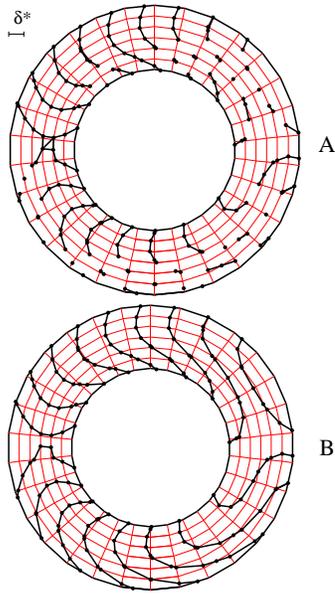


Figure 11: Activation isochrone map obtained by applying: A - the original SA; B - our revised approach.

Future work will deal with a thorough study of weak adjacency relations to identify and define spatial relations between isocurves, as well as between possible higher-level aggregated objects. Such relations should reveal adjacencies and interactions between the spatial objects, and, then, underlie the feature extraction process. To this end, a crucial issue deals with the specification of the neighborhood graphs that properly represent the spatial relations between geometrical objects at different levels of abstraction.

As far as the specific application problem, future work will deal with the interpretation of activation isochrone maps: contour shapes, the directions along which the front velocity is higher, minimum and maximum regions represent the features we are currently interested to extract from the input field. Such features can be correlated to the expected activation sequence, and used in a diagnostic context: deviations from the nominal patterns would highlight anomalous excitation pathways and possible ectopic sites associated with ventricular arrhythmias.

Acknowledgements. We would like to thank Feng Zhao and Piero Colli Franzone, who provided us with the SAL code and with the simulated data, respectively.

References

- C. Bailey-Kellogg and F. Zhao. Influence-based model decomposition for reasoning about spatially distributed physical systems. *Artificial Intelligence*, 130(2):125–166, 2001.
- C. Bailey-Kellogg and F. Zhao. Spatial aggregation:

Modelling and controlling physical fields. In L. Ironi, editor, *QR97*, pages 13–22. IAN-CNR, 1997.

C. Bailey-Kellogg, F. Zhao, and K. Yip. Spatial aggregation: Language and applications. In *Proc. AAAI-96*, pages 517–522, Los Altos, 1996. Morgan Kaufmann.

I. Bratko, I. Mozetic, and N. Lavrac. *Kardio: A Study in Deep and Qualitative Knowledge for Expert Systems*. MIT Press, Cambridge, MA, 1989.

P. Colli Franzone, L. Guerri, and M. Pennacchio. Spreading of excitation in 3-D models of the anisotropic cardiac tissue. II. Effect of geometry and fiber architecture of the ventricular wall. *Mathematical Biosciences*, 147:131–171, 1998.

P.L. George and H. Borouchaki. *Delaunay Triangulation and Meshing*. Hermes, Paris, 1998.

X. Huang and F. Zhao. Relation-based aggregation: finding objects in large spatial datasets. *Intelligent Data Analysis*, 4:129–147, 2000.

X. Huang. *Automatic analysis of spatial data sets using visual reasoning techniques with an application to weather data analysis*. PhD thesis, The Ohio State University, 2000.

L. Joskowicz and E. Sacks. Computational kinematics. *Artificial Intelligence*, (51):381–416, 1991.

M. Kundu, M. Nasipuri, and D.K. Basu. A knowledge based approach to ECG interpretation using fuzzy logic. *IEEE Trans. Systems, Man, and Cybernetics*, 28(2):237–243, 1998.

W.E. Lorensen and H.E. Cline. Marching cubes: a high resolution 3d surface construction algorithm. *Computer Graphics*, 21:163–169, 1987.

B. Taccardi, B.B. Punske, R.L. Lux, R.S. MacLeod, P.R. Ershler, T.J. Dustman, and Y. Vyhmeister. Useful lessons from body surface mapping. *Journal of Cardiovascular Electrophysiology*, 9(7):773–786, 1998.

R.L. Watrous. A patient-adaptive neural network ECG patient monitoring algorithm. *Computers in Cardiology*, pages 229–232, 1995.

F. Weng, R. Quiniou, G. Carrault, and M.-O. Cordier. Learning structural knowledge from the ECG. In *ISMDA-2001*, volume 2199, pages 288–294, Berlin, 2001. Springer.

K. Yip and F. Zhao. Spatial aggregation: Theory and applications. *Journal of Artificial Intelligence Research*, 5:1–26, 1996.

K. Yip. Structural inferences from massive datasets. In L. Ironi, editor, *QR97*, pages 215–220. IAN-CNR, 1997.

F. Zhao. Intelligent simulation in designing complex dynamical control systems. In Zfafestas and Verbruggen, editors, *Artificial Intelligence in Industrial Decision Making, Control, and Automation*. Kluwer, 1995.