

Automated detection of electrocardiographic diagnostic features through an interplay between Spatial Aggregation and Computational Geometry

Liliana Ironi and Stefania Tentoni

IMATI - CNR

via Ferrata 1, 27100 Pavia, Italy

{ironi,tentoni}@imati.cnr.it

Abstract

Within the medical domain, Functional Imaging provides methods for effectual visualization of diagnostically relevant numeric fields, i.e. of spatially referenced measurements of variables related to organ functions. Unveiling the salient physical events that underly a functional image is most appropriately addressed by feature extraction methods that exploit the domain-specific knowledge combined with spatial relations at multiple abstraction levels and scales. The identification of specific patterns that are known to characterize classes of pathologies provides an important support to the diagnosis of disturbances, and the assessment of organ functions. In this work we focus on Electrocardiographic diagnosis based on epicardial activation fields. This kind of data, which can now be obtained non invasively from body surface data through mathematical model-based reconstruction methods, can hit electrical conduction pathologies that routine surface ECGs may miss. However, their analysis/interpretation still requires highly specialized skills that belong to few experts. Given an epicardial activation field, the automated detection of salient patterns in it, grounded on the existing interpretation rationale, would represent a major contribution towards the clinical use of such valuable tools whose diagnostic potential is still largely unexplored. We focus on epicardial activation isochronal maps, which convey information about the heart electric function in terms of the depolarization wavefront kinematics. An approach grounded on the integration of a Spatial Aggregation (SA) method with concepts borrowed from Computational Geometry provides a computational framework to extract, from the given activation data, a few basic features that characterize the wavefront propagation, as well as a more specific set of diagnostic features that identify an important class of heart rhythm pathologies, namely reentry arrhythmias due to block of conduction.

Keywords: Biomedical imaging; functional imaging; image based diagnosis; spatial aggregation; computational geometry; electrocardiography; cardiac electrical function.

Introduction

One of the most important application domains where imaging has proved extremely useful is Medical Diagnosis. The process of identifying a pathological condition can be greatly supported by signs of deviations from normality

Copyright © 2009, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

that can be drawn from images. Within this context the term “imaging” usually refers to techniques to build images of anatomical districts of the human body (e.g. radiographies, CAT, NMR); more broadly, it can also include methods aimed at providing graphical representations of temporally/spatially referenced measurements of variables related to specific organ functions (e.g. EEG, ECG signals, activation maps). In this latter case the term “functional” imaging is to be preferred.

Many functional images are graphical representations of a physical field: a potential contour map, for instance, is the spatial representation of a potential field. Thereby, the task of analyzing such images is not adequately tackled by traditional Image Processing methods, which have been designed for raster images. The issue of unveiling the salient physical events underlying a functional image is more appropriately and effectively addressed through feature extraction methods that can exploit the domain-specific knowledge at different abstraction levels. Such an issue is particularly relevant in view of performing explanation and reasoning tasks.

Within the field of Qualitative Spatial Reasoning, Spatial Aggregation (SA) (Yip and Zhao 1996) provides the most appropriate conceptual framework for feature extraction at multiple levels, according to a powerful hierarchical abstraction strategy. In the direction of making the approach more robust and of integrating within the basic SA framework methods borrowed from quantitative research fields, several works have contributed to make it an even more attractive framework for the development of functional imaging tools (Ironi and Tentoni 2003a; 2003b). Any such tool would ground on domain-specific knowledge, as the inference mechanisms rely on a network of relations that, besides dealing with spatial properties, explicitly encode such knowledge.

With the present work we continue our research effort aimed at delivering novel tools to support the assessment of the electric cardiac function (Ironi and Tentoni 2007). Diagnosing the cardiac electric function has always been a hard task for the difficulty met in the identification of salient electrical events and their spatial association with specific epicardial sites. In the clinical context, diagnosis of conduction pathologies is still carried out on the ECG signals for which the interpretative rationale is well-established. Several tools

exist for automated ECG segmentation and classification. Most of these tools are based on the integration of wavelet transforms with neural/fuzzy-neural networks, to deal respectively with the signal decomposition and classification tasks (see for example (Clifford, Azuaje, and McSharry 2006)). Within AI, Qualitative Reasoning has also played an important role in providing a number of automated ECG interpretation tools (Bratko, Mozetic, and Lavrac 1989; Weng et al. 2001; Kundu, Nasipuri, and Basu 1998). Unfortunately some important rhythm disturbances may be incorrectly located or even missed by routine ECGs. Even body surface high resolution mapping may fail because signs of cardiac electrical events on the torso surface are weak.

In recent years, model-based numerical inverse procedures have made it possible to obtain non-invasively the epicardial activation field from body surface data. That has engaged researchers in the effort towards novel methods for electrocardiographic imaging (Oster et al. 1997; Ramanathan et al. 2004). However, the interpretative rationale for cardiac maps is only partially defined, and the ability to abstract the most salient visual features from a map and relate them to the complex underlying phenomena still belongs to few experts. Due to the extreme complexity of the physical system the task of automating diagnosis of conduction disturbances from a 2D/3D activation field is therefore hard, and necessarily limited to the current interpretation rationale. Within this field functional image-based diagnosis is at its beginning, and, in accordance with the available rationale, currently regards only a few classes of conduction disturbances. The potential of Qualitative Spatial Reasoning in contributing to its development is high: a tool for the automated extraction of spatially referenced features of the cardiac electrical function would bridge the gap between established research outcomes and clinical practice.

Our work fits into a long-term research project aimed at delivering an automated electrocardiac map interpretation tool. To detect salient spatiotemporal features in the epicardial activation field, we exploit the inference mechanisms provided by a computational tool grounded on Spatial Aggregation and on Computational Geometry concepts: from a given numeric field we extract spatial objects that, at different abstraction levels, qualitatively characterize spatiotemporal phenomena, and discover and abstract the skeleton of patterns relevant to the diagnostic task.

In this paper we focus on epicardial activation maps, which convey information about the heart electric function in terms of the depolarization wavefront kinematics. These kind of maps are very useful to diagnose rhythm disturbances. We describe how some basic spatiotemporal features that characterize the propagation of the electrical excitation can be abstracted from the given activation data, and in particular we define a set of distinctive features that identify an important class of rhythm disturbances due to blocks of conduction.

Feature abstraction from a numeric field

The comprehension of physical phenomena benefits from the visualization of the spatial course of relevant variables. A visual representation obtained from a given numeric field

can be further inspected, and searched for homogeneities and specific patterns that have a physical meaning. This “imagistic” reasoning activity, that goes beyond mere visualization, is performed at multiple levels through a sequence of abstractions and manipulations of spatial objects that capture key physical properties.

Spatial Aggregation

Spatial Aggregation (SA) is a general-purpose framework that provides a suitable ground to capture spatiotemporal adjacencies at multiple scales in spatially distributed data. It was designed to derive and manipulate qualitative spatial representations that abstract important features of the underlying data, for their use in automated reasoning tasks (Yip and Zhao 1996; Ironi and Tentoni 2003a; 2003b).

In outline, SA transforms a numeric input field into a multi-layered symbolic description of the structure and behavior of the physical variables associated with it. This results from iterating transformations of lower-level objects into more abstract ones through the exploitation of qualitative equivalence properties shared by neighbor objects.

SA abstraction mechanisms are based on three main steps, namely *Aggregation*, *Classification*, and *Redescription*, that exploit domain-specific knowledge and spatial adjacencies (see Fig.1):

1. *Aggregation*. Spatial adjacency of low-level objects is encoded within a neighborhood graph.
2. *Classification*. Neighbor objects are grouped by similarity, according to a domain-specific equivalence predicate that defines a feature of interest.
3. *Redescription*. Similarity classes are singled out as new high-level objects that provide an abstract representation of the feature.




Figure 1: Basic inference steps in Spatial Aggregation.

Step 1 mostly exploits geometrical properties, either metrical or topological. However, to ensure the robustness of the *Classification* step, it can also take into account available non-geometrical knowledge, associated with the objects to be aggregated and related to the physical context (Ironi and Tentoni 2003a). For example, with respect to the contouring task, the values of the variable associated with each field point can be exploited to generate an appropriate neighborhood graph that guarantees the abstraction outcome against both artifactual curve entanglement and segmentation (Ironi and Tentoni 2003a). Step 3 is crucial in that a non-effectual redescription of new objects may jeopardize subsequent abstractions stemming therefrom.

Such operators can be iterated over and over until the behavioral and structural information about the underlying physical phenomenon, that is required to perform a specific task, is extracted from the data set. The hierarchical structure of the whole set of the so-built objects defines a bi-directional mapping between higher and lower-level aggregates, and, consequently, it facilitates the identification of the pieces of information relevant for a specific task.

The role of Computational Geometry

Within the SA abstraction mechanism, *Redescription* instantiates visual features that may play a role in the spatial reasoning process. The geometric representation of the new objects must convey a meaningful effectual visual synthesis of the underlying similarity class.

Computational Geometry methods and concepts can play an important role in providing algorithms for the re-description of newly abstracted objects. An important class of objects for which the chosen representation format particularly needs to suit the reasoning task is that of 2D bounded regions. These latter ones can result, for example, from the application of a similarity relation grounded on interval values to a set of contiguous isopoints. The similarity classes correspond to regions that need to be instantiated as new geometrical objects for further treatment. Sometimes a centroid can do, while in other situations a more articulated - though compact - re-description may be necessary, for example when the qualitative topological structure of the region needs to be captured at multiple scales. The choice of the most appropriate format and scale for the redescription object is task-driven.

For qualitative reasoning tasks, a region descriptor should be:

- i) robust and stable with respect to noise and small perturbations of the boundary,
- ii) capable to roughly capture the location and global extent of the region,
- iii) capable to capture the topological structure of the region at an appropriate scale of details with respect to the task, and of course
- iv) computationally feasible.

An effectual representation of a region can be provided by its “gross skeleton”, as defined in the following.

Gross skeleton of a region

The concept of “gross skeleton”, that we are going to introduce, is derived from the “medial axis” of a bounded region. The medial axis is geometrically defined as the locus of the centers of circles that are internally tangent to the region’s boundary. That results in a set of curves roughly running along the middle of the region. Unfortunately, the medial axis is very sensitive to small perturbations of the boundary: noisy contours produce a number of secondary branches.

The medial axis can be thought of as a geometric skeleton of the figure, whose complexity, given by the number of branches, corresponds to the boundary complexity, defined as the number of its curvature extrema. For its instability

the medial axis is not immediately suitable as a figure descriptor in contexts affected by noise, and as such it is also inappropriate where finer scale details are irrelevant and to be ignored.




Figure 2: Panels A,B refer to two sample regions. For each panel, the Voronoi based medial axis (magenta) is shown when the region’s boundary is relatively smooth (left), and when it is affected by noise (right)

Exact computation of the medial axis is difficult in general. It is well-known that an approximation of the medial axis of a region can be obtained from the Voronoi diagram related to a finite set of points that sample the region’s boundary (Brandt and Algazi 1992): it consists in the subset of Voronoi edges that lie completely within the region’s interior. Such approximation, though, still suffers from the cited instability problem as it is illustrated in Fig.2.

The following algorithm builds a robust simplified topological skeleton of a given polygonal region, namely the gross skeleton, by exploiting a relevance measure (Sakai and Sugihara 2006), and selectively pruning the approximated Voronoi medial axis. Instability is removed by dropping spurious/irrelevant branches that correspond to unneeded information about finer contour’s details.

Algorithm (gross skeleton construction).

Given $\{P_1, \dots, P_n\}$, vertices of a closed polyline bounding a connected region \mathcal{L} ,

1. Compute \mathcal{M} , Voronoi approximation of the medial axis of \mathcal{L} , as follows:
 - (a) Build the Voronoi diagram related to the set of vertices $\{P_1, \dots, P_n\}$,
 - (b) Retain only the edges that are completely internal to \mathcal{L} .
2. Compute the “index of relevance” $\beta(E)$ of each edge $E \in \mathcal{M}$, as

$$\beta : \mathcal{M} \rightarrow (0, 1) \quad \beta(E) = 2|l|/|\partial\mathcal{L}|$$

where if P_i, P_k are the generators of Voronoi edge E , $|l|$ is the length of shortest path connecting P_i with P_k along the region’s boundary $\partial\mathcal{L}$, and $|\partial\mathcal{L}|$ is the length of the region’s boundary (Fig.3).

3. (*Selective pruning*) Starting with $\mathcal{L}^* = \mathcal{M}$,

$$\forall E \in \mathcal{M}, \beta(E) < \beta^* \Rightarrow \mathcal{L}^* := \mathcal{L}^* \setminus \{E\}$$

where β^* is a given relevance threshold.




Figure 3: Steps in the construction of the gross skeleton of a polygonal region. Vertices P_i, P_k of the region's boundary (partially shown, blue line) generate Voronoi edge E (thicker line). Part of the Voronoi tessellation (green thin lines), and of the approximated medial axis \mathcal{M} (magenta thick line) are also shown.

Selective pruning of the medial axis \mathcal{M} is performed according to an edge relevance criterion by which irrelevant boundary details are dropped: edges with a very low β value have a negligible effect on the region's boundary. The result

$$\mathcal{L}^* = \{E \in \mathcal{M} \mid \beta(E) \geq \beta^*\}$$

is a connected linear structure that reflects the global topological structure of the region, as well as its rough location and spatial extent.

The choice of the relevance threshold β^* affects the complexity of the resulting gross skeleton \mathcal{L}^* , and adjusts the descriptor to the scale required by the reasoning task: as greater β^* is, as more simplification is required.

In Fig. 4 a series of perturbations of the smooth sample region shown in Fig. 2A are reported: in each case both the Voronoi medial axis and the gross skeleton are computed. The figure clearly shows how more robust the gross skeleton is with respect to the Voronoi medial axis approximation, and how the global shape of the region is captured.

Functional Imaging of the cardiac electric function

The heart is site of cyclic electrical activity which causes the muscle to rhythmically contract. The propagation of the electric excitation within the myocardium is a quite complex 4D spatiotemporal process that electrocardiologists explore on reference surfaces (epicardial, endocardial) by means of relevant variables, such as the electric potential, the activation time and the wavefront propagation velocity. Due to the difficulty of combining spatial and temporal aspects, exploring the potential $u(\mathbf{x}, t)$, a function of space and time, is a hard task. A more global and synthetic view on the spatiotemporal process of excitation is provided by the epicardial representation of the *activation time* $\tau(\mathbf{x})$, defined as the instant at which an epicardial site \mathbf{x} changes its electric state from resting to activated. Such an instant is commonly estimated as the point of minimum (time) derivative extracted from the electrogram $t \rightarrow u(\mathbf{x}, t)$. Therefore, the activation




Figure 4: Each panel refers to a distinct perturbation of the smooth region shown in Figure 2A. In each panel: the Voronoi based medial axis (left), and the gross skeleton obtained by pruning with $\beta^* = 0.25$ (right).

time is a sort of landmark variable which embeds a qualitatively significant event in the electric potential time course, and, when spatially represented on the whole epicardial surface, it holds a powerful diagnostic potential.

In imaging of the cardiac electric function, an important role is played by activation maps: such maps are contour maps of the activation time that convey information about the wavefront structure and propagation. In a previous work (Ironi and Tentoni 2007), in accordance with the existing rationale of interpretation, we tackled the problem of defining and abstracting, within the SA framework, a set of spatial objects that capture a few important basic features of activation: isochrones, whose spatial sequence depicts the spread of excitation by snapshots, wavefront breakthrough and exit locations, fast propagation pathways.




Figure 5: Activation map as obtained from noisy data.

As an example, Fig. 5 shows an activation map obtained from noisy simulated data related to a case of normal propagation elicited by single site pacing. Let us remark that the activation time field is actually related to a 3D model of the epicardium; in order to have a unique global planar




Figure 6: Main wavefront propagation features abstracted from the sample data of Fig. 5: activation isochrones, breakthrough (B) and exit sites (e/E), and fast propagation pathways (thick vectors).

view with minimal spatial distortion, we considered an axial cylindrical projection of this map. After preliminary noise removal, from the activation field the main wavefront propagation features are detected: the sequence of isochrones, the breakthrough and extinction sites, which respectively mark where excitation starts and ends on the epicardial surface, and the fast propagation pathways (Fig. 6).

In this work we focus on an important class of pathological conditions, namely *reentry ventricular tachycardia* (VT), and provide SA-based definitions and algorithms for the abstraction and spatial redescription of the features involved. Reentry VT usually follows a myocardial infarction, as the presence of scar tissue enhances resistivity and modifies the patterns of wavefront propagation. When conduction is abnormally slow across a region (≤ 0.1 m/sec, (Cranefield 1975)), an anomalous activation pattern, called “reentry”, can be triggered: the excitation wavefront travels in single/multiple circular patterns, and reenters the area where it arose from. Much research effort has been devoted to the study and characterization of this disorder (Burnes, Taccardi, and Rudy 2000; Burnes et al. 2001; de Bakker et al. 1993).

The key components of the reentrant VT pattern, in terms of wavefront kinematics, are (Fig. 7):

1. a *cul-de-sac*-like region (isthmus), bounded by lines of block;
2. a breakthrough site in the isthmus area;
3. a single/multiple-loop reentry propagation pattern;
4. an excitation ending site located proximal to the breakthrough, but outside the blocked area.




Figure 7: Schematic VT reentry circuit components.





Figure 8: Isopoints (black dots) on the epicardial surface: the surface mesh is shown. Top panel: 3D geometry. Bottom panel: 2D cylindrical projection.

Given the discretized epicardial geometry Ω_h and the activation time field $\tau = \tau(\mathbf{x}_i)$, $\mathbf{x}_i \in \Omega_h$, the main steps carried out to map it to a structural spatial representation of the salient propagation features, including the possible presence of a reentry VT pattern, are here very briefly summarized:

1. Breakthrough and exit sites, isopoints, and the time ordered sequence of the isochrones are first obtained (Ironi and Tentoni 2007); a planar projection of the 3D geometry is used to provide an overall representation of the features on the epicardial surface (Fig. 8: isopoints as they are mapped onto the projection).
2. The velocity field is computed as (Colli Franzone, Guerri, and Pennacchio 1998)

$$\mathbf{v}(\mathbf{x}) = \nabla\tau(\mathbf{x})/|\nabla\tau(\mathbf{x})|^2$$

where ∇ is the gradient operator. By mapping the velocity module range into a small set of qualitative values, e.g. *very-slow*, *slow*, *medium*, *high*, *very-high*, in accordance with threshold values suggested by the experts, the epicardial surface gets partitioned into homogeneous subregions, each of them labeled by the same qualitative value of the velocity module. In this context, the value *very-slow* corresponds to a pathological condition. Then:

3. If the region labeled *very-slow*, \mathcal{L} , is not empty,

- (a) it gets redescribed by its gross skeleton, \mathcal{L}^* , which represents the abstracted “conduction block” line;
- (b) a set of propagation lines, obtained as stream lines of the vector field, is generated from a neighborhood of the ends of the block line. They get classified into “main propagation patterns” according to their ending site, which can only be either a wavefront exit site or the upper/lower border of the map;
- (c) check whether, among the ending sites associated with the main propagation paths, at least one is located close to the breakthrough site located nearest to the isthmus area (loop pattern).




Figure 9: The approximated medial axis \mathcal{M} (magenta thick line), and its pruned version \mathcal{L}^* (blue thick line) are shown within the very-low-velocity area bounded by $\partial\mathcal{L}$.




Figure 10: The conduction block is extracted as a complex of features: a line of block (simplified skeleton of the critical velocity area) which leaves a breakthrough and an extinction site at opposite sides.

Let us remark that very noisy data should be properly pre-processed to reduce noise to an acceptable level and allow reliable and robust feature extraction. Data smoothing actually corresponds to the way the expert approaches the visual reasoning task: getting rid of minor or spurious details to catch the main patterns.

Step 3 is aimed at discovering and abstracting a possible reentry circuit by singling out its key components.

Figure 9 shows, for the data set corresponding to Fig. 8, a detail of the area where isochrones are spatially denser: the boundary $\partial\mathcal{L}$ of a critical *very-slow* region is shown, as well as the Voronoi based medial axis \mathcal{M} , and the gross skeleton \mathcal{L}^* . Figure 10 shows the abstracted conduction block complex: a cul-de-sac region where isochrones get more crowded, bounded by a line of block which leaves a breakthrough and an extinction sites, spatially close to each other, at opposite sides. The line of block, so extracted as gross skeleton of the *very-slow* area, corresponds to merging the locally crowded isochrones.

Figure 11 shows the global outcome of the abstraction processes. It consists of: the sequence of activation isochrones, the breakthrough and exit sites, the discovered block of conduction, and the reentrant propagation patterns, starting at the ends of the block arc.




Figure 11: Outcome of the abstraction processes: activation isochrones (thin solid lines), breakthrough/exit sites (B/E labels), and the block of conduction (thick solid line). A couple of wavefront propagation lines, starting at the ends of the block arc, are shown (dashed thick lines).

Discussion

The approach herein proposed to automatically capture specific aspects of cardiac electrical activity is of broad methodological interest to electrocardiography, and more in general, to medical imaging. It results from the integration of standard computational geometry concepts with a spatial aggregation methodology. This latter, that aims at interpreting a numeric input field, allows us to capture structural information about the underlying physical phenomenon, and to identify its global patterns and the causal relations between them. Thanks to its hierarchical strategy in extracting objects at different scales, it facilitates the definition of inference rules that favor automated reasoning on spatiotemporal phenomena to perform a specific task.

In this work we focussed on algorithms for automated detection of diagnostic features from activation time fields. For the diagnosis of rhythm disturbances, activation maps are most appropriate as they provide information about the spatiotemporal course of the excitation wavefront. We showed how spatiotemporal features that characterize an important class of arrhythmias can be extracted from the given activation field.

As for the realization of a complete diagnostic tool for cardiac electric activity, further insight into the electric function could be drawn from the analysis of temporal sequences of potential data $t_i \rightarrow u(\mathbf{x}, t_i) \quad \mathbf{x} \in \Omega_h, \quad i = 1, \dots, n$, through the search for local current inflows/outflows identified by typical patterns of potential maxima and minima within isopotential maps. From potential data, especially intramural measurements, we could derive information about the electrical activity prior to its surface breakthrough that is complementary with respect to that obtainable from surface activation data. That would allow us to locate intramural components of reentry pathways associated with arrhythmogenic activity. However, the challenge of combining spatial and temporal aspects in a full 4D analysis goes with the still incomplete rationale of interpretation of such maps, and makes advances in this direction more remote.

At any rate, even if we limit our attention to epicardial activation data, additional work needs to be done:

- (i) To validate the robustness of the proposed methodology on measured data, and then clearly delimit its weaknesses and strengths when applied in a clinical context. To this regard, sensitivity to noise should be more deeply investigated.
- (ii) To deal with more complex phenomena, such as those involving the Purkinje network or multiple stimuli, and properly characterize and capture all of their propagation aspects. An example of such a complex situation is illustrated in Fig. 12: multiple wavefronts originate at distinct sites, after a few milliseconds they collide and merge into two fronts that advance in opposite directions. Wavefront collisions, which mark abrupt changes in the front topology and discontinuities in the velocity field, deserve both diagnostic attention and computational care.
- (iii) To define a strategy for the comparison of the features of a given map against those of a *nominal* one, with the aim to detect and explain possible deviations from the expected patterns.




Figure 12: Activation isochrones (thin solid lines), and breakthrough/exit sites (B/E labels) in a case of simulated Purkinje involvement. Multiple wavefronts break through at distinct sites, and then collide.

From a broader application perspective, the methodology we propose could contribute to a diagnostic tool specifically designed for arrhythmias, and could be used in a ther-

apeutic context to evaluate the efficacy of a drug therapy aimed at normalizing the rhythm, through the detection of its effects on the spatial activation patterns.

References

- Brandt, J. W., and Algazi, V. R. 1992. Continuous skeleton computation by Voronoi diagram. *55(3):329–338*.
- Bratko, I.; Mozetic, I.; and Lavrac, N. 1989. *Kardio: A Study in Deep and Qualitative Knowledge for Expert Systems*. Cambridge, MA: MIT Press.
- Burnes, J. E.; Taccardi, B.; Ershler, P. R.; and Rudy, Y. 2001. Noninvasive electrocardiogram imaging of substrate and intramural ventricular tachycardia in infarcted hearts. *Journal of the American College of Cardiology* 38(7):2071–2078.
- Burnes, J. E.; Taccardi, B.; and Rudy, Y. 2000. A noninvasive imaging modality for cardiac arrhythmias. *Circulation* 102(17):2152–2158.
- Clifford, G. D.; Azuaje, F.; and McSharry, P. E., eds. 2006. *Advanced Methods and Tools for ECG Analysis*. Boston/London: Artech House Publishing.
- Colli Franzone, P.; Guerri, L.; and Pennacchio, M. 1998. 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.
- Cranefield, P. F. 1975. *The Conduction of the Cardiac Impulse: the Slow Response and Cardiac Arrhythmias*. Mount Kisco NY: Futura Publishing Co.
- de Bakker, J. M.; van Capelle, F. J.; Janse, M. J.; Tasseron, S.; Vermeulen, J. T.; de Jonge, N.; and Lahpor, J. R. 1993. A noninvasive imaging modality for cardiac arrhythmias. *Circulation* 88(3):915–926.
- Ironi, L., and Tentoni, S. 2003a. On the problem of adjacency relations in the spatial aggregation approach. In Salles, P., and Bredeweg, B., eds., *Seventeen International Workshop on Qualitative Reasoning (QR2003)*, 111–118.
- Ironi, L., and Tentoni, S. 2003b. Towards automated electrocardiac map interpretation: an intelligent contouring tool based on spatial aggregation. In Berthold, M. R.; Lenz, H.-J.; Bradley, E.; Kruse, R.; and Borgelt, C., eds., *Advances in Intelligent Data Analysis V*, 397–417. Berlin: Springer.
- Ironi, L., and Tentoni, S. 2007. Automated detection of qualitative spatio-temporal features in electrocardiac activation maps. *Artificial Intelligence in Medicine* 39:99–111.
- Kundu, M.; Nasipuri, M.; and Basu, D. K. 1998. A knowledge based approach to ECG interpretation using fuzzy logic. *IEEE Trans. Systems, Man, and Cybernetics* 28(2):237–243.
- Oster, H. S.; Taccardi, B.; Lux, R. L.; Ershler, P. R.; and Rudy, Y. 1997. Noninvasive electrocardiographic imaging: reconstruction of epicardial potentials, electrograms, and isochrones and localization of single and multiple electrocardiac events. *Circulation* 96:1012–1024.

Ramanathan, C.; Ghanem, R. N.; Jia, P.; Ryu, K.; and Rudy, Y. 2004. Noninvasive electrocardiographic imaging for cardiac electrophysiology and arrhythmia. *Nature Medicine* 1–7.

Sakai, H., and Sugihara, K. 2006. A method for stable construction of medial axes in figures. *Electronics and Communications in Japan - Part II* 89(7):48–55.

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

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