A Diagnosis Methodology for Continuous Time Measurements using Hierarchical Signal Representations

M. Borahan Tümer
Doctoral Candidate
Electrical and Computer Engineering
Marquette University
Milwaukee, Wisconsin, 53201
bora@spiff.eece.mu.edu

Lee A. Belfore, II
Assistant Professor
Electrical and Computer Engineering
Old Dominion University
Norfolk, Virginia, 23529
belfore@ece.odu.edu

Kristina M. Ropella
Associate Professor
Biomedical Engineering
Marquette University
Milwaukee, Wisconsin, 53201
k.ropella@marquette.edu

ABSTRACT

A methodology for automated diagnosis of systems characterized by continuous signals is presented. The methodology requires the definition and construction of several fuzzy automatons each capable of identifying a particular condition. When the diagnostic system is in operation, the time sampled system measurements are presented to all automatons simultaneously. The fuzziness in automaton operation enables input processing from several perspectives, consistent with the operation of the automatons, allowing for toleration of measurement noise and other ambiguities. The methodology is applied to the problem of automatic electrocardiogram diagnosis.

1 INTRODUCTION

Diagnosis of complex nonlinear systems is a challenging research area in science and engineering, playing an important role in both medical and industrial applications. The diagnostic accuracy provided by the diagnostic system is among the most important criteria for measuring the performance of the diagnostic system; yet a highly desirable feature is the applicability of the diagnostic system to a variety of systems. Complicating the problem is that signals from which the diagnosis is derived can be perturbed by noise, localized baseline wander, and measurement nonlinearities. It is for these reasons that hierarchical fuzzy automatons (HFA) are employed here as a tool in automated diagnosis. HFAs are fuzzy automatons [1,2] that process a signal at several levels of detail. Moving up each level in the hierarchy results in the identification of more complex and global structures. At the apex of the hierarchy, is one fuzzy automaton. The input to the HFA is the time sampled signal that has been tokenized into primitives using an adaptive resonance theory 2 (ART2) artificial neural network (ANN) [3] where token fuzziness has been extracted in an ad hoc fashion from the internal state of ART2. Nondeterministic operation of individual HFAs is an essential feature of its operation. The state machine model supports simultaneous transitions from any starting state to relevant next states. As the state machine operates, memberships within all states evolve until the state memberships along with the transition paths dominate. As these states are identified, the HFA collapses into a small number of states for any given transition. Once synchronization is achieved, the diagnosis is determined by examining the respective performance of several HFAs. An HFA is associated either with each different condition and perhaps with significant variations of a particular condition. Thus, several HFAs are operated simultaneously on the same signal. The condition associated with the HFA can be associated with the highest membership indicates the identified condition or by some other metric such as the number of transitions that were matched with input templates.

Research has been conducted on automated diagnosis using both decision-theoretic [4] and syntactic [5–9] approaches. In general, the syntactic approach is more sophisticated since this approach utilizes information specific to the temporal structure of signals themselves. In syntactic methodologies, the set of templates (input signal primitives) are determined followed by the construction of a state machine that identifies sequences of templates. In [5,6], the state machines attempt syntactic recognition using formal grammars. In [7–9], attributed state machines are employed. Attributed state machines using formal grammars proved to be more powerful since classical state machines cannot handle the complexity of medical signals [9]; however, attributed methods require many parameter attributes to control the parsing process. Studies based on attributed automata have shown that extensive computation is necessary to handle these parameters, and hence, attributed state machines have problems in the practical use of computer methods in medicine. Furthermore, more robust systems are needed to deal with the noisy and subject-varying and time-varying medical signals. The hierarchical state machines are not frequently used by syntactic methods (only reported in [9]). The HFAs used in this study also have several advantages including in the proposed method considering all cases simultaneously.

0-7803-4778-1/98 $10.00 © 1998 IEEE

The methodology incorporates several important processes, hierarchically, that include both artificial neural networks and fuzzy automata. At the lowest level in the hierarchy, the continuous time signal is parsed by an ART2 ANN to identify low level primitive signal shapes. The ART2 ANN is trained with recordings from both good and pathologic systems. Significantly, this process can be performed in a manner independent of the system being monitored. At the intermediate level of the hierarchy, analogous to lexical analysis, a fuzzy automaton is constructed to identify structural features from both good and failed systems. At the highest level, a fuzzy automaton is created that transitions as each structural feature at the intermediate level is identified. The signal characterizing the system being observed is assumed to have structure containing some degree of periodicity. The process of training ANNs and creating fuzzy automatons is repeated for each condition and, if necessary, for variants of particular conditions.

This paper is organized into six sections including an introduction, a diagnosis system overview and operation, a presentation of the methdology for constructing the diagnosis system, an example application of automatic ECG diagnosis, the results of the analysis on the example, and a summary.

2 SYSTEM OVERVIEW AND OPERATION

In our approach, an HFA is constructed for performing automatic diagnosis. A block diagram of the system is shown in Fig. 1. An ART2 ANN is used to tokenize the

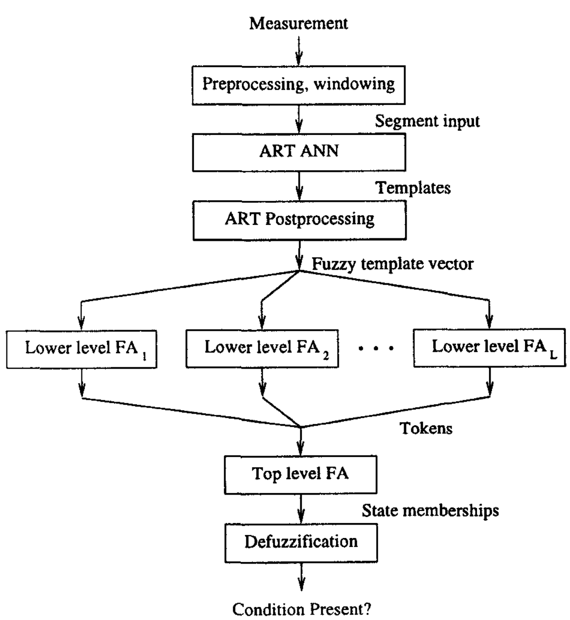

Figure 1. Overview of Diagnosis System Processing.

input signal for processing by the automatons that follow. The benefit of using the ART2 architecture is that learning is unsupervised and it has the ability to identify shapes in segments of the input signal. Furthermore, for each input segment, ART2 can also be modified to supply a measure which can be used as a fuzziness measure for membership to each of the pattern classified by ART2. The measures are passed, as a vector, to the lower level fuzzy automatons (LLFAs). Each LLFA identifies a syntactic structure from the input measurement. When an LLFA reaches an accepting state, a transition results in the top level fuzzy automaton (TLFA). A discussion of the theory and operation of fuzzy automatons can be found in [1,2]. In this work, the HFA acts as a nondeterministic finite automaton, with simultaneous membership in several states and transitions along several paths possible. Diagnosis is achieved when synchronization occurs within a LLFA and the TLFA by collapsing the state memberships to one or a few states within the HFA. Each condition for which diagnosis is desired requires the development of its own HFA system and training of the ART2 ANN. The full details of the design and operation of the system are presented in [10].

Operation of the HFA.

A key aspect of the diagnosis system is the operation of the HFA. Each HFA consists of a single TLFA with transitions driven by a set LLFAs, one for each token. Transitions within the LLFAs, in turn, are driven by the output of ART2 which determines the most likely set of templates for a given segment of the input. In operation, all fuzzy automatons operate and process inputs simultaneously. The operation of the HFA is discussed in a bottom up fashion, where at the bottom reside input templates, categorizations of the raw input, and at the top is the TLFA.

At the bottom of the hierarchy, the raw input is preprocessed and then classified into the template classes. Template classes are assigned a membership based on parameters related to the goodness of classification and several achieving good classifications are placed in the input template vector. In all these discussions, the superscript k is the segment index which is also the discrete time index. The input template vector can be represented as

$$\mathbf{S}^k = \{(\tau_j^k, \mu_j^k) | \mu_j^k > \mu_{\text{thres}}\}, \tag{1}$$

where $\mathbf{S}^k$ is the set of pairs consisting of the input templates produced by ART2 output, τ_j^k, and the memberships assigned for each template μ_j^k, and μ_{thres} some appropriately selected threshold. The measure of fuzziness is related to the classification process within ART2. The template membership is defined by

$$\mu_j^k = r_j^k, \tag{2}$$

where j is the template index, $r_j^k \in [0,1]$ is an internal parameter from ART and is larger for better matches with internally stored pattern templates. More specifically, r_j^k is the measure that ART2 compares with the vigilance as a part of the process for categorizing the input. Thus, for each input, a measure of fuzziness can be determined for all templates.

LLFAs recognize tokens, sequences of templates, that are representative structures within the signal being analyzed. All transitions are assigned a fuzziness of 0.5, allowing any relevant template to produce a transition. State membership fuzziness is the maximum of two quantities, the first being the fuzzy membership of the destination state. The second quantity is the minimum of the membership of the initial state and the transition membership. This relationship insures that certainty increases with unambiguous templates memberships and that ambiguity is maintained with sustained ambiguity. Mathematically, the state membership update rule is transition update rules are defined by

$$\mu_{\phi_j}^{k+1} = \begin{cases} \max\limits_{\phi_i \in I_s^k} (\mu_{\phi_j}^k, \min(\mu_{S_{\delta_{ij}}^k}, \mu_{\phi_i}^k)) & \text{if } \mu_{\delta_{ij}} \leq \mu_{S_{\delta_{ij}}^k} \\ \mu_{\phi_j}^k & \text{otherwise} \end{cases} \tag{3}$$

where ϕ_i is an initial state, ϕ_j is an destination state, δ_{ij} is the template that must be input to transition from state ϕ_i to ϕ_j, $S_{\delta_{ij}}^k$ a template derived ART2 classifications, μ_{ϕ_i} is the fuzzy state membership in state ϕ_i, $\mu_{\delta_{ij}}$ is the fuzzy transition membership from ϕ_i to ϕ_j, and

$$I_s^k = \{\phi_i | \phi_i \in I_j \wedge \delta_{ij} \in \mathbf{S}^k\} \tag{4}$$

where I_j is the state of states that can transition to state ϕ_j.

Before any input is presented, fuzzy state memberships are initialized so that LLFA starting states have the following memberships

$$\mu_{\phi_i} = \begin{cases} 1.0 & \text{Start state} \\ 0.5 & \text{Accepting state} \\ 0.0 & \text{Remaining states.} \end{cases} \tag{5}$$

The start state membership is 1 to ensure that each template receives full consideration in all LLFAs at any time implying that the receipt of an input template can initiate a transition in any LLFA. Zero memberships are assigned to states that are neither inputs nor accepting states so that a token is not accepted if processing begins in the middle of a token. Finally the membership of 0.5 assigned to accepting states guarantees that maximum ambiguity is reported until less ambiguous templates are matched beginning from the start state.

3 METHODOLOGY FOR CONSTRUCTING DIAGNOSIS SYSTEM

The diagnosis system described in the last section is constructed using some basic assumptions on the types of processing performed as well as on the characteristics of the signals in the class of systems for which automated diagnosis is desired. The diagnosis system is constructed in two stages representing the system independent and independent parts, respectively. First, preprocessing and template identification are performed. Next, HFAs are constructed for each condition which is a largely manual process and is system dependent.

System Independent Methodologies.

The systems under consideration are assumed to be measured and monitored through one or more time sampled analog signals. The signals are also assumed to have some structure. For example, a signal characterized entirely by signal statistics will not be a good candidate for the methodology described here.

Preprocessing is performed to remove baseline wander present in the signal if necessary. In addition, noise filtering and signal scaling operations may be performed. It is possible that nonlinear scaling may be applied because some signals may have components that are one or more orders of magnitude larger than other important features. For example, in ECGs, a large potential spike is associated with the QRS complex, which is an order of magnitude larger than other important signal features. In this work, however, it was not necessary to apply such techniques [10]. In addition, the input signal is sampled as a set of consecutive samples of the input. The sampling of the input may vary depending on whether or not training of other parts of the system are in process.

Alphabet identification serves as the boundary between the system independent and dependent methodologies. For different systems, the structural components may also vary as a function of the level of detail. At higher levels of signal detail, the components are weakly system dependent and general, while at lower levels of detail the components are strongly system dependent and specific. In addition, the true nature of signals observed at higher levels of detail are more ambiguous due to noise and other variations in the signal. The mechanism adopted in this work is to employ the ART2 ANN which can identify and then classify time sampled continuous measurements [3]. In this work, the Stuttgart Neural Network Simulator (SNNS) [11] and the accompanying ART algorithms have been modified to facilitate the creation of the fuzzy template vector used in subsequent processing. In learning, a sliding window is used to present exemplars to the ART2 network. Overlap between succeeding windows helps ART2 create exemplars for inputs at various phases with respect to the sampling window. When learning is complete, ART2 can classify input signals into one of the identified classes. Furthermore, ART2 can also give information useful in determining the ambiguity of the measured signal. In operation, the constraints on the sliding window can be relaxed so that succeeding windows have no overlap, for efficiency considerations. Creating an analogy with written text, the classes in the ART2 are simply like letters in the alphabet from which all words and sentences are constructed. Indeed, the output template identification will be a vector of templates, each with its associated fuzziness. By passing the template vector, no template with a strong similarity with the input is excluded from consideration. By examining a number of consecutive template vectors, an otherwise weaker template in the short term may result in a stronger overall result. Creation of template alphabets is performed for each different condition for which diagnosis is desired and the process of creating templates is relatively application independent.

System Dependent Methodologies.

System dependent processing takes information about the characteristics of the system for creating fuzzy automatons. Identifying tokens from sequences of alphabet templates moves down the next level of detail and up on the level of system specifics. Finally, strings of tokens are used to represent the different diagnosis conditions. Given a collection automatons for the different diagnosis conditions, defuzzification into a crisp diagnosis is achieved.

Tokens are used to identify signal features at an intermediate level of detail. Strings of alphabet templates are used to represent tokens, which in this work have been identified manually. Each token is identified by a fuzzy automaton that recognizes a string of alphabet templates. The LLFA is constructed by examining the stream of winning tokens that are emitted from ART2. The structure of the state machine is derived by examining the relationships between consecutive pairs of tokens emitted by ART2. Transition memberships were assigned ambiguous memberships of 0.5, but other information can be used to determine transition memberships. One challenge in developing lower level fuzzy automatons is selecting an appropriate partitioning between the LLFA and the TLFA. The construction of LLFAs requires some domain specific information knowledge, and the degree of domain specific information can vary depending on the application and desired complexity of the LLFAs. From a linguistic perspective, LLFAs recognize words in a language. The LLFAs work on the principal of being maximally ambiguous until a sequence of unambiguous templates are received beginning with the start state.

The TLFA describes transitions between states defined as a result of LLFA membership in accepting states. The fuzzy state measurement in the TLFA is computed with (3) where the token membership is computed with respect to the accepting state, ϕ_A^k, of the LLFA. The TLFA is constructed in a domain specific fashion and is based on both high level structure of the information and the structure of the LLFAs. Again from a linguistic perspective, the TLFA recognizes sentences.

In the last stage of processing, defuzzification is performed. While a fuzzy membership can be extracted from the HFA, it was found that a better measure of overall HFA success was a tally of the number of consecutive transitions within the state machine structure that are made over all input. This serves as a synchronization measure to measure how much of the input followed the HFA. Using this as a basis, if several machines are able to synchronize on the input, more transitions indicates a better match. In the event multiple HFAs give the same or very closely the same number of transitions, then the machine fuzziness is used to break the tie. In a sense, the diagnosis is the condition associated with the HFA that best synchronizes with the input recording.

4 EXAMPLE: ECG DIAGNOSIS

In this example, a demonstration of automatic diagnosis of the ECG is presented. Two HFAs are constructed to process, respectively, the normal sinus rhythm and the atrial fibrillation, a type of arrhythmia originating from the atria. A normal sinus rhythm is illustrated in Fig. 2.

In the normal sinus rhythm, the cardiac cycle is composed of the P wave, QRS complex, and the T wave followed by the constant isoelectric line separating two cardiac cycles. Fig. 3 shows an atrial fibrillation which is a disturbance of the regular heart rhythm. In atrial fibrillation, the ECG exhibits an undulation of the baseline, called the fibrillatory or the f-waves, accompanied by an irregular ventricular rhythm. While the QRS complex and the T wave keep the normal configuration, an irregular accompanying atrial fibrillation rhythm might look like $ffQRSTffffQRSTfQRSTfffQRSTfffffff$. The features of atrial fibrillation can be best detected in the lead V_1 or the standard lead II [12].

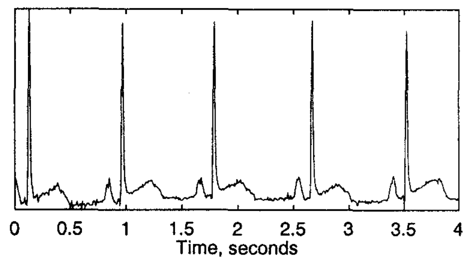

Figure 2. An ECG Illustrating Normal Sinus Rhythm

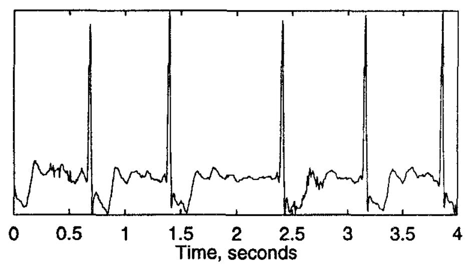

Figure 3. An ECG Illustrating Atrial Fibrillation

Template alphabet construction

In this example, two ART2 networks and template alphabets are created for the normal sinus rhythm and atrial fibrillation. The task of the preprocessor is the decomposition of the input ECG. The preprocessor accomplishes the decomposition with a moving window. The two parameters of the window, the window size w and the shift step s, are selected to be $w = 10$ and $s = 2$ during training. For the HFA from atrial fibrillation, an ART2 network is trained with 600 templates, 200 originating from three different patients. For the HFA from normal sinus rhythm, an ART2 network is trained with 200 templates originating from one patient. Each template consists of 10 consecutive time samples. For the normal sinus rhythm, an ART2 network with 83 categories (alphabet templates) is results. The ART2 network trained for atrial fibrillation contains 87 categories. Plots showing this template alphabet can be found in [10, 13].

FHA construction phase

The first task is to determine the structure of the TLFA and for each LLFA to recognize an ECG signal. The TLFA is used to recognize an entire heart beat. In the ECG signals analysis, each token represents either a complex or a wave. Each LLFA recognizes a part of the signal structure specific to that condition. For instance, in a normal sinus rhythm, a LLFA is designed to recognize each of the tokens (P wave, the PR interval, the QRS complex, the ST segment, the T wave, and the neutral line showing the inactive time period between two consecutive heart beats), while for the atrial fibrillation, LLFA are designed to recognize the tokens (the QRS complex, the ST segment, the T wave and the TQ interval) since atrial fibrillation changes the normal sinus rhythm and generates complexes with characteristics of atrial fibrillation. The transition diagram of the TLFA for the normal sinus rhythm is illustrated in Fig. 4. The LLFA showing the QRS wave

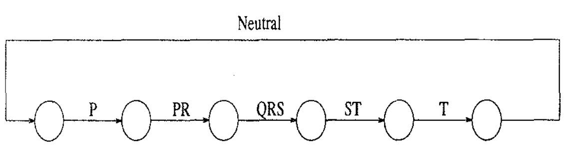

Figure 4. The TLFA Recognizing an ECG with Normal Sinus Rhythm

and the ST segment of the ECG is shown in Figs. 5 and 6.

The second task is to construct the LLFA for each condition from the condition-specific signal set . To do this, first the template sequences in the multi-category string, each of which corresponds to a token (this can be found by checking ECGs), are marked. Then, the LLFA for each token is constructed such that all alphabet templates that appear in the multi-category string as a part of the token are included as a possible transition initiater. Furthermore, alphabet templates that differ within a certain mean square error (MSE) from the first set of templates are included to transition the LLFA to the same state as template to which they are similar. The reason for including the second set of templates in LLFA is to provide a more complete syntactic representation of the token.

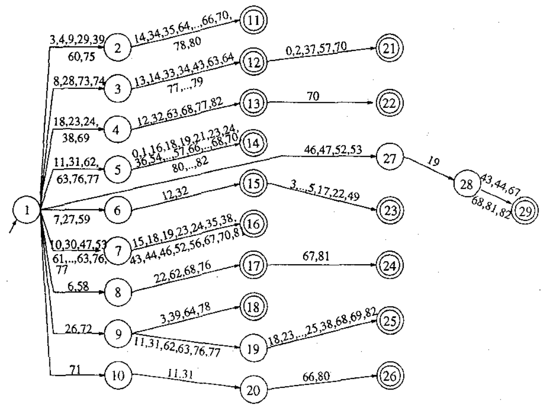

Figure 5. The LLFA Recognizing a QRS Complex in the Normal Sinus Rhythm

	Disease Present	Disease Absent	Total
Test positive	a=24	b=5	$a+b$=26
Test negative	c=1	d=5	$c+d$=9
Total	$a+c$=25	$b+d$=10	$a+b+c+d$=35

Table 1. Tabulations of Test Results

5 RESULTS

Twenty-five ECGs indicating atrial fibrillation and 10 ECGs indicating normal sinus rhythm are used to test the diagnosis methodology. Three out of 25 and one out of 10 ECGs were used to train the ART2 networks and to construct two HFA structures for the atrial fibrillation and the normal sinus rhythm, respectively. The standard medical diagnostic parameters are used to evaluate the results [14]. Definitions of these parameters and our results are summarized in Tables 1 and 2.

The pair of HFAs was able to distinguish correctly 24 out of 25 ECGs recorded from patients suffering from atrial fibrillation. Hence, the sensitivity of the HFA approach was 0.96. Furthermore, five out of ten normal sinus rhythms are correctly distinguished by the HFAs. Two of the ECGs with normal sinus rhythm are incorrectly diagnosed to have atrial fibrillation. Both HFAs (of the normal sinus rhythm and the atrial fibrillation) recognized the remaining three signals, able to accept all template sets of the three ECGs. For all three ECGs, the final minimum membership value of the atrial fibrillation was slightly higher than that of the normal sinus rhythm. It would be expected that developing the HFA for normal sinus rhythm from several patients would result in better overall performance.

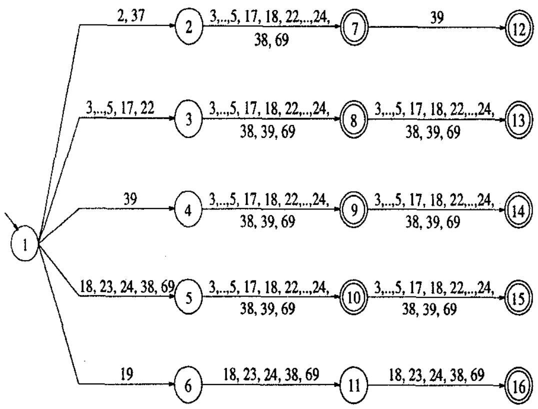

Figure 6. The LLFA Recognizing the ST Segment in the Normal Sinus Rhythm

Measure Name	Definition	Actual Value
Sensitivity	$\frac{a}{a+c}$	0.96
Specificity	$\frac{d}{b+d}$	0.50
False Negative Rate	$\frac{c}{a+c}$	0.04
False Positive Rate	$\frac{b}{b+d}$	0.50
Predictive Value positive	$\frac{a}{a+b}$	0.83
Predictive Value negative	$\frac{d}{c+d}$	0.83
False Alarm Rate	$\frac{b}{a+b}$	0.17
False Reassurance	$\frac{c}{c+d}$	0.17

Table 2. Standard Accuracy Criteria and Results

6 SUMMARY

This paper presents the use of HFAs as a diagnostic tool for nonlinear systems. In this paper, HFAs were defined in terms of structure and function. HFAs use a fuzzy syntactic approach for diagnosis of time-sampled signals. In operation, the HFA transforms the signal into a string of sets of elementary structures, templates. Then, examining the consecutive templates, the HFAs determine or not the string characterizes a particular condition. After the syntactic analysis is performedfor each desired condition, state synchronizing measures and state memberships from these HFAs are used to identify the condition. The HFAs were applied to the problem of ECG diagnosis with good results.

7 REFERENCES

[1] M. M. Gupta, G. N. Saridis, and B. R. Gaines, *Fuzzy Automata and Decision Processes*, Elsevier North-Holland, Inc., 1977.

[2] A. Kandel and S. C. Lee, *Fuzzy Switching and Automata*, Edward Arnold (Publishers) Ltd., 1979.

[3] G. A. Carpenter and S. Grossberg, "A massively parallel architecture for a self-organizing neural pattern recognition machine", *Computer Vision, Graphics, and Image Processing*, vol. 37, pp. 54–115, 1987.

[4] T. F. Yang, B. Devine, and P. W. Macfarlane, "Artificial neural networks for the diagnosis of atrial fibrillation", *Medical and Biological Engineering and Computing*, vol. 32, pp. 615–619, 1994.

[5] S. L. Horowitz, "A syntactic algorithm for peak detection in waveforms with applications to cardiography", *Communications of the ACM*, vol. 18, pp. 281–285, 1975.

[6] G. Belforte, R. DeMori, and F. Ferraris, "A contribution to the automatic processing of electrocardiograms using syntactic methods", *IEEE Transactions on Biomedical Engineering*, vol. 26, pp. 125–136, 1979.

[7] P. Trahanias and E. Skordalakis, "Syntactic pattern recognition of the ECG", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, pp. 648–657, 1990.

[8] E. Pietka, "Feature extraction in computerized approach to the ECG analysis", *Pattern Recognition*, vol. 24, pp. 139–146, 1991.

[9] A. Koski, M. Juhola, and M. Meriste, "Syntactic recognition of ECG signals by attributed finite automata", *Pattern Recognition*, vol. 28, pp. 1927–1940, 1995.

[10] M. Borahan Tümer, *A Fuzzy Syntactic Approach to Fault Diagnostics by Analysis of Time Sampled Signals*, PhD thesis, Marquette University, 1998.

[11] University of Stuttgart, *Stuttgart Neural Network Simulator User Manual Version 4.0*, 1995.

[12] H. J. L. Marriott, *Practical Electrocardiography*, Williams and Wilkins, 1983.

[13] M. Borahan Tümer, Lee A. Belfore, II, and Kristina M. Ropella, "Applying hierarchical fuzzy automatons to automatic diagnosis", in *Proceedings of the 1998 Meeting of the North America Fuzzy Information Processing Systems (NAFIPS'98)*, Pensacola, Florida, 1998.

[14] B. W. Jervis, M. R. Saatchi, A. Lacey, T. Roberts, E. M. Allen, N. R. Hudson, S. Oke, and M. Grimsley, "Artificial neural network and spectrum analysis methods for detecting brain diseases from the CNV response in the electroencephalogram", *IEE Proceedings: Science, Measurement and Technology*, vol. 141, pp. 432–440, 1994.