A model-based approach for learning to identify cardiac arrhythmias

G. Carrault[1], M.-O. Cordier[3], René Quiniou[2], M. Garreau[1], J.J. Bellanger[1], and A. Bardou[1]

[1] LTSI, Campus de Beaulieu, 35042 Rennes Cedex FRANCE
[2] IRISA, Campus de Beaulieu, 35042 Rennes Cedex FRANCE
[3] IRISA/Université de Rennes 1, Campus de Beaulieu, 35042 Rennes Cedex
currently at INRA, Unité Rennes-Quimper, 35042 Rennes Cedex FRANCE

Abstract. ECG interpretation is used to monitor the behavior of the electrical conduction system of the heart in order to diagnose rhythm and conduction disorders. In this paper, we propose a model-based framework relying on a model of the cardiac electrical activity. Due to efficiency constraints, the on-line analysis of the ECG signals is performed by a chronicle recognition system which identifies pathological situations by matching a symbolic description of the signals with temporal patterns stored in a chronicle base. The model can simulate arrhythmias and the related sequences of time-stamped events are collected and then used by an inductive learning program to constitute a satisfying chronicle base. This work is in progress but first results show that the system is able to produce satisfying discriminating chronicles.

1 Introduction

Researchers have been paying attention to arrhythmia identification for a long time. Two main steps are clearly established i) the signal processing step which produces pertinent information such as P and QRS wave features, ii) the diagnosis step which attempts to explain the abnormal features discovered during the first step. Classically, deterministic tests, Bayesian networks or decision trees have been used for the last issue [Bla86]. More recently, artificial intelligence techniques have been proposed, such as knowledge-based approaches using expert rules [Shi85,Lon96], sometimes combined with fuzzy logic as in [KNB98]. In order to overcome the problem of expertise acquisition, model-based approaches, relying on a heart model instead of expertise knowledge, were experimented ([BML89,TW93,Gue96]). A typical study in this direction is described in [SCLB95] where a cardiac model such as those developed in theoretical cardiology [BABC96] is used for detection as well as diagnostic tasks.

Our work is part of this model-based framework. It relies on a model of cardiac electrical activity. The model can simulate a large number of arrhythmias and outputs the associated signal and its description as a sequence of time-stamped events. Arrhythmia features, learnt from these simulations, are used to analyze on-line ECG signals, and to classify them. A major strength of this project is to bring together specialists in the three following fields: signal processing, artificial intelligence (diagnosis, and machine learning) and theoretical cardiology. This is highly desirable when the objective is to build an operational system. Another originality is the fact that the heart model used is not an ad hoc one, developed only for the sake of our project. On the contrary, it was intended to be used as a cardiac simulator in pharmacology and teaching experiments [SCLB95]. Furthermore, the time occurrences of events for atria and ventricular activities in normal and disorder cases, have been validated by clinicians.

An overview of the architecture of CALICOT (Cardiac Arrhythmias Learning for Intelligent Classification of On-line Tracks) is presented in section 2 as well as the cardiac model on which the approach is based. The on-line analysis of the ECG signals is sketched in section 3. It relies on a chronicle base which is acquired off-line by using inductive learning programming techniques as explained in section 4. This work is in progress but the first results, mainly related to chronicle learning, are promising. They are analyzed in section 4.2. The last section is devoted to a comparison of our work with related approaches.

2 A model-based approach

2.1 CALICOT architecture

The architecture of CALICOT, given in figure 1, separates on-line and off-line treatments. The on-line part includes the signal processing module (1) which outputs the symbolic representation of the analyzed signal in terms of time-stamped events, and the chronicle recognition module (2) which analyses the stream of events and identifies arrhythmia disorders by detecting characteristic patterns.

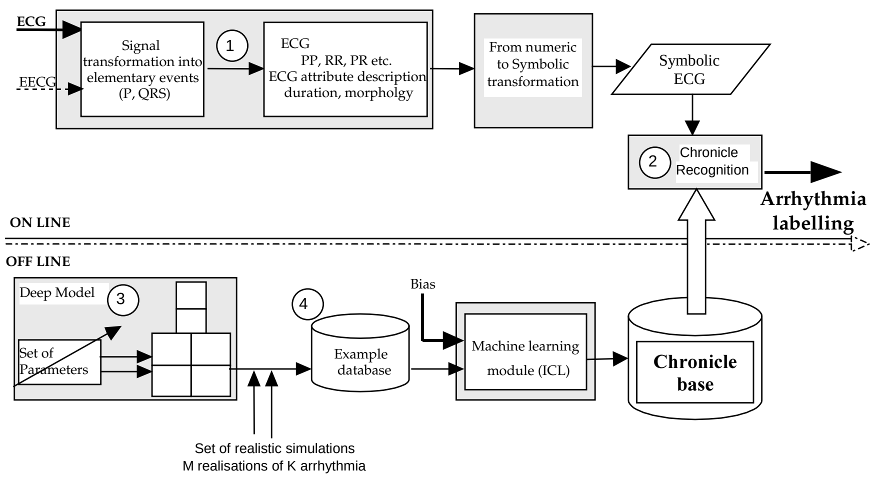

Fig. 1. Architecture of the proposed approach

The off-line part aims at building a set of relevant chronicles (the chronicle base) from cardiac model simulations. The cardiac model (3) is used to generate sequences of events corresponding to normal as well as pathological signals. The learning module (4) relies on inductive logic programming (ILP) to extract the most discriminating patterns which identify the normal and pathological ECGs.

2.2 The electrical heart model

The heart model Cardiolab [SCLB95] is central to our approach. It was originally designed to be a cardiac simulator and explicitly integrates anatomophysiological and electrophysiological knowledge. The model combines a qualitative and quantitative modeling of the electrical heart function and is based on the propagation of action potentials in myocardium. This approach is in fact close to the semi-quantitative approach proposed by [BK92] where numerical information is used in addition to qualitative constraints.

The heart behavior is described according to two abstraction levels. The first level provides a structural model based on macro-cells and the elementary interactions between them as receiving and transmitting delayed impulses to their neighboring cells. These cells (21 in total) correspond to the nodal and the myocardium tissue elements. A cell is mainly characterized by its states phase duration during the cardiac cycle namely resting state (muscle tissue only), depolarization, absolute refractory, relative refractory, slow diastolic depolarization (nodal tissue only). The second level is concerned with all the electrophysiological concepts such as conduction fronts, reentry, blocks. This level combines the elementary events produced at the first level and constructs more complex events. The resulting model is able to simulate a large variety of rhythms by tuning the parameters of the myocardium cells (mainly phase durations). The generic character of the simulation makes the model well adapted for intelligent monitoring.

3 On-line arrhythmia identification

As explained before, the ECG signals are analyzed on-line to identify arrhythmia phenomena. This is achieved by two main modules, the signal processing module and the recognition module.

3.1 The signal processing module

This module processes the signal and outputs its symbolic representation in terms of time-stamped events: i) it detects and identifies markers of the cardiac revolution (P wave, QRS complex) ii) it determines their descriptive attributes (duration, morphology amplitude...) iii) determines temporal relations between these events.

Figure 2 illustrates the functionalities of the module with the analysis of an ECG signal corresponding to a *ventricular couplet*. P_i and R_i are the detected time occurrences of the P and QRS waves. It can be noticed that the ventricular

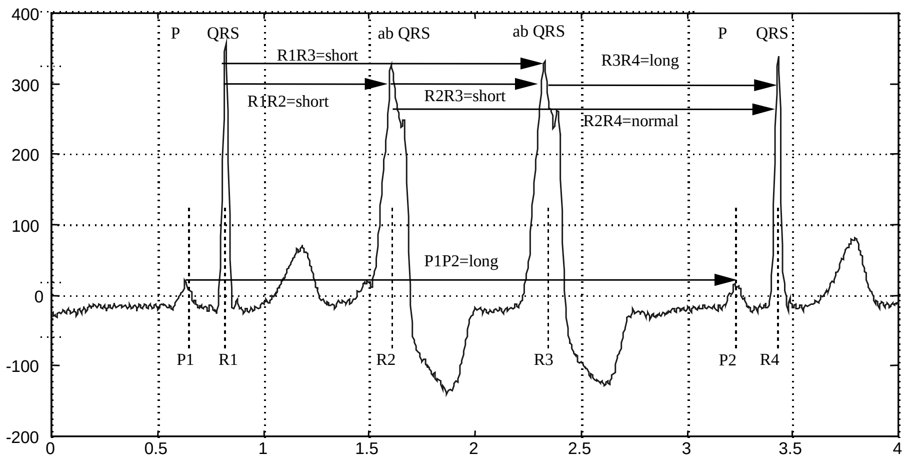

Fig. 2. A signal and its characterization as produced by the signal processing module (the actual coded description is similar to the example described in table 1)

activities $R2$ and $R3$ were not preceded by corresponding P waves. The qualitative values qualifying the wave duration (normal, short or long) and morphology (normal, *ab*normal) are computed by comparing their observed time occurrence to reference values found in [BH80]. In the example, $R2$ and $R3$ are premature (intervals $R1R2$ and $R2R3$ are short) and abnormal (*ab* feature of QRSs) while $P2$ occurs lately (interval $P1P2$ is long).

3.2 Chronicle recognition

The sequence of symbolic events produced by the signal processing module is then analyzed to identify potential arrhythmia troubles. The chronicle recognition approach was specifically designed to analyze a flow of time-stamped events under real-time constraints [Gha96,RDF97]. A chronicle is a set of events constrained by quantitative or qualitative temporal relations (e.g. *before*, *after*, *500 ms before*). It describes a temporal pattern which characterizes the phenomenon to be identified.

The following chronicle, expressed in Prolog, describes the *bundle branch block* disorder. `p_wave` and `qrs_complex` denote events occurring in the ECG. Their arguments stand respectively for the time occurrence and the qualification of the wave shape. The brace brackets specify the temporal constraints on the time occurrence of events. Here, the temporal constraints are qualitative and `pp`, `pr` and `rr` denote the intervals between two successive P waves, a P wave and the next QRS complex and two successive QRS complexes, respectively.

```
bundle_branch_block :-
    p_wave(T1, normal), qrs_complex(T2,abnormal),
    p_wave(T3, normal), qrs_complex(T4, abnormal),
    {pp(T1, T3, normal), pr(T1, T2, normal),
     rr(T2, T4, normal)}
```

This chronicle states that the four specified events must occur and satisfy the following constraints: a first *p_wave* at time T1 followed by a *qrs_complex* at time T2 and again a *p_wave* at time T3 followed by a *qrs_complex* at time T4; the shapes of *p_waves* are normal but those related to *qrs_complexes* are abnormal ; the duration of the *pp*, *pr* and *rr* intervals must be normal (close to the nominal rate). In the bundle branch block case, inter-wave durations are normal but the two QRS shapes are abnormal.

Chronicle recognition consists in skimming the flow of events coming from an observed process and detecting the specific events that belong to a chronicle. Once a chronicle has been evoked, the recognizer checks the presence of the remaining events and their temporal constraints. This process is very similar to pattern-matching associated with temporal constraint satisfaction.

4 Off-line chronicle base acquisition

On-line chronicle recognition is very efficient but chronicle acquisition from experts is often difficult and raises completeness and soundness problems. Deep models are usually simpler to build because they rely on the theoretical or practical knowledge relating to the process itself. Our approach consists in using such a model as a simulator and to generate the observed events resulting from determined situations (usual failures or pathologies). In our application, the cardiac model is used to generate sets of examples, one set per arrhythmia under interest. As these examples cannot be directly used as a rule base (they are in general too numerous and complex), they are processed by an inductive learning tool in order to produce a set of characteristic patterns, which identify the classes (here a set of discriminating chronicles[1]); these patterns are then used for on-line recognition of phenomena from the flow of observations.

4.1 Learning chronicles

ICL [LRD97] is a classification system that learns clausal theories (first-order representations) which discriminate, as well as possible, between several classes of examples. Positive examples have to be models of the target theory and negative examples must not be models of this same target theory. ICL also makes use of

[1] A chronicle is discriminating if it expresses a sufficient condition for the recognition of the related phenomenon with respect to the other chronicles.

a background theory which contains what is already known and useless to learn. Such a theory improves the efficiency of machine learning.

In our case, a class corresponds to a normal rhythm or to some disorder. Positive examples are instances of the particular class to be characterized and negative examples are all the other instances.

Multi-class learning in ICL is formally defined as follows:

- Given
  - B a background theory (containing only definite clauses) ;
  - E the whole set of examples ;
  - $C_P \subset E$ the set of positive examples related to class C ;
  - $C_N = E - C_P$ is the set of negative examples ;
  - L_H a language which defines the syntax of acceptable clauses for characterizing the classes
- For each class C, find a clausal theory $H_C \subseteq L_H$ such that
  - $\forall p \in C_P, M(B \cup p)$ is a true interpretation of H_C (completeness) ;
  - $\forall n \in C_N, M(B \cup n)$ is a false interpretation of H_C (consistency).

4.2 Coding sequence of events

Cardiac data coming from an ECG have interesting features that must be taken into account for representation. Firstly, cardiac data are highly structured: the normal cardiac activity may be seen as a sequence of beats, each of which is composed of a series of events corresponding to the electrical activation of different areas of the heart. Furthermore, these events are temporally constrained. Secondly, the number of beats necessary to represent and then to detect an arrhythmia is variable: some can be detected on one cycle, from one beat to the next, while some need several beats.

wave name	P or QRS	time in ms	dur.	1st P 1st R	PP1 RR1	2nd P 2nd R	PP2 RR2	1st R 1st P	PR1	2nd R	PR2
p1	p_wave	17	n	p2	n	p3	n	r1	n	r2	n
r1	qrs_complex	137	n	r2	n	r3	n	p2			
p2	p_wave	807	n	p3	n	p4	n	r2	l	r3	l
r2	qrs_complex	982	n	r3	l	r4	l	p3			
p3	p_wave	1591	n	p4	n	p5	n	r3	l	r4	l
r3	qrs_complex	1983	n	r4	l	r5	l	p4			
p4	p_wave	2374	n	p5	n	p6	n	r4	l	r5	l
p5	p_wave	3157	n	p6	n	p7	n	r4	n	r5	l
r4	qrs_complex	3277	n	r5	l	r6	l	p6			
p6	p_wave	3940	n	p7	n	p8	n	r5	l	r6	l
r5	qrs_complex	4237	n	r6	l	r7	l	p7			
p7	p_wave	4723	n	p8	n	p9	n	r6	l	r7	l

Table 1. A Wenckebach ECG coded as a positive input example for ICL

The ECG is coded as a sequence of P waves and QRS complexes. The duration of an ECG used as an example is about 5 seconds corresponding to 5 to 8 beats. Table 1 gives an example of coded ECG. The description attributes associated to waves are the following:

P **wave**
- duration qualified as short(s), normal(n) or long(l),
- pointer to the next P wave and length of the temporal interval between the current one and the next one (qualified as short, normal or long),
- pointer to the second next P wave and length of the related temporal interval,
- pointer to the next QRS wave and length of the temporal interval between the current P one and this next QRS one,
- pointer to the second next QRS wave and length of the related temporal interval.

QRS **complex**
- duration qualified as short(s), normal(n) or long(l),
- pointer to the next QRS wave and length of the related temporal interval between the current one and this next one,
- pointer to the second next QRS wave and length of the related temporal interval,
- pointer to the next P wave.

4.3 First results

As a first experiment, the approach has been used to learn 5 rhythm and conduction disorders: normal, Wenckebach 4:3 (wenckebach in the sequel), Wolff-Parkinson-White (kent), left bundle branch block (lbbb) and trigeminy. Due to space limitations, the language bias and the background theory used to process these examples is not given.

In order to demonstrate the feasability of the approach, an initial learning set containing 25 examples (5 examples per class, each modeling 5 seconds ECGs), has been provided to ICL. Below are the rules that were learnt from the set of simulated sequences. They are given under a disjunctive normal form. ICL proposed one rule per class except for the wenckebach class which is characterized by the disjunction of three rules. These rules are quite satisfying with respect to the restricted set of arrhythmias and the low size of the learning set.

```
normal: p_wave(P,normal), pr1(P,R2,normal), pr2(P,R3,normal),
   qrs(R,normal), rr2(R,R2,normal),
   p_wave(P1,normal), pr2(P1,R2,normal).
lbbb: p_wave(P,normal), pr1(P,R,normal), qrs(R,long).
kent: p_wave(P,normal), pr1(P,R,short).
trigeminy: p_wave(P,normal), pr2(P,R1,short).
wenckebach: p_wave(P,normal), pr2(P,R2,long), qrs(R2,normal).
wenckebach: p_wave(P,normal), pp2(P,P2,short).
```

```
wenckebach: p_wave(P,normal), pp1(P,P1,normal),
   succ_P(P1,P2,correct), pr1(P,R2,normal), pr2(P,R3,normal),
   qrs(R,normal), rr2(R,R2,normal).
```

5 Related work

Much work has been done in the field of qualitative cardiac modeling for ECG recognition. The earliest ones used a knowledge base as a model. More recently hybrid models using quantitative and qualitative temporal constraint propagation were proposed. Our own model, Cardiolab, belongs to this class. We did not retain approaches based on cellular model [BABC96] because we are interested in more abstract events such as the time sequence of the atria and ventricular activities.

[Shi85] presents a knowledge-based approach for the interpretation of arrhythmias. One of the interesting points is the use of causal knowledge about the cardiac conduction system. It can take missing information (missing waves for example) into account. A prototype system named CAA (for Causal Arrhythmia Analysis) using a representation language based on semantic networks, the PSN language, has been implemented.

Kardio ([BML89]) shares several features with our own approach. The goal of Kardio is to recognize cardiac arrhythmias from a coded ECG. The system relies on a purely qualitative electrical model of the heart which cannot actually be considered as a cardiac simulator. Our system, Cardiolab, has a model of electrical heart conduction which combines quantitative and qualitative simulation and is closer to cardiac reality than Kardio, actually it can generate pretty realistic cardiac sequence events. The qualitative model of Kardio has been simulated in order to generate surface (diagnostic and predictive) rules which were then processed by an inductive engine in order to produce compact rule databases (a diagnostic one and a predictive one). The main difference between Kardio and our own approach is that Kardio uses feature-based induction whereas we use ILP. Thus, Kardio can only learn predefined propositional structural relations. By using ILP, we can learn first-order relations which may not be given before induction [SMKS96]. The learnt rules are much simpler and better suited to the chronicle recognition approach. Also, as far as we can see, Kardio uses a hand-coded ECG. Our goal is to devise an integrated system in which the signal processing module that labels the ECG can interact with the chronicle recognizer in order to better detect arrhythmias.

Einthoven [TW93] was designed for ECG interpretation. It uses an extensive quantitative as well as qualitative knowledge base. Einthoven suffers from the classical drawbacks of a knowledge-based approach: the knowledge-base must be extensive in order to provide good diagnoses but it is very difficult to acquire and maintain - we used machine learning for chronicle generation to avoid this drawback; rule-based reasoning is not very efficient and hardly usable for on-line monitoring if not associated with methods such as chronicle recognition.

Ticker [HK95] aims at teaching non-specialists to interpret ECG. The cardiac qualitative model is a directed graph in which the nodes represent states (electrical state of a region) and the arcs possible transitions with associated qualitative temporal contraints (flow of electrical conduction between regions). Different abstraction levels can be defined in order to tailor the explanation to the user's knowledge. As mentioned by the authors, the Ticker model is closer to reality than Kardio's but is too abstracted to produce correct traces.

The heart model of Guertin's Holmes system [Gue96] uses quantitative temporal constraints on intra-beat events as well as inter-beat events and thus can be considered as an extension of Ticker. Temporal abduction is used for diagnosis and the combinatoric explosion of possible explanations is controlled by temporal constraint propagation. The system is not integrated into a monitoring system and processes hand-coded ECG. Though constraint propagation improves the efficiency of abductive inference, we think that such an approach cannot be used on-line. Chronicle recognition is much more efficient from this point of view.

6 Discussion and conclusion

We have presented an integrated approach for automatic ECG interpretation. Our proposition relies on two points: 1) an on line chronicle recognizer detects arrhythmias on the symbolic labeling of the surface ECG produced by the signal processing module; 2) the specification of the searched arrhythmias as chronicles are learnt from classified example ECGs representing typical arrhythmias obtained from simulation traces of a heart model. For efficiency reasons, this model is not used directly for (abductive) diagnosis as in model-based approaches.

The benefits of learning chronicles over expert associative rules are: an easier acquisition and the fact that discriminating chronicles are learnt, thus insuring an efficient recognition. Other benefits include easy knowledge maintenance which could occur, for example, when more precise diagnoses are expected and some classes of arrhythmias must be split into several. As ICL is based on a formal learning model showing soundness and completeness properties, it guarantees that if the learning set is good (relevant and representative) the set of learnt chronicles will be good (discriminating and not overgeneralised). In our approach, the fact that the model has been validated insures that examples will be correctly classified. But, this says nothing about the completeness or representativity of the learning set. We are now working on model parameter tuning in order to build realistic learning bases.

The work is still in progress. The first results are very promising. As assessed by experts, the features retained in the learnt chronicles are relevant and adjusted to the level of detail that is intended: if few classes are to be discriminated, the learnt chronicles are very simple; if more classes are desired, a new learning step can add discriminating details into the chronicles. We plan the addition of the following disorders encountered in coronary care unit: ventricular asystole, atrioventricular dissociation, complete heart block or accelerated idioventricular rhythm. Another perspective is to confront the learnt chronicles to a real ECG database such as the MIT-BIH.

Up to now our work was focused on the signal processing module and the learning module independently. A further step will be devoted to linking the signal processing module and the chronicle recognizer. It must be emphasized that current available monitoring devices are far from ideal. False alarms and breakdowns are inevitable. The idea is to use the predictions made by the chronicle recognizer to reinforce hypotheses on which the signal processing module works, in order to increase the accuracy of the labeling process.

References

BABC96. A. Bardou, P. Auger, P. Birkui, and J.L. Chasse. Modeling of cardiac electrophysiological mechanisms : From action potential genesis to its propagation in myocardium. *Critical Reviews in Biomedical Engineering*, 24:141–221, 1996.

BH80. M. Blondeau and M. Hiltgen. *lectrocardiographie Clinique*. Ed. Masson, 1980.

BK92. D. Berleant and B. J. Kuipers. Combined qualitative and numerical simulation with Q3. In B. Faltings and P. Struss, editors, *Recent advances in qualitative physics*, pages 3–16. MIT Press, 1992.

Bla86. R. Le Blanc. Quantitative analysis of cardiac arrhythmias. *CRC: Critical Review in Biomedical Engineering*, 14(1):1–43, 1986.

BML89. I. Bratko, I. Mozetic, and N. Lavrac. *Kardio: A Study in Deep and Qualitative Knowledge for Expert Systems*. MIT Press, 1989.

Gha96. M. Ghallab. On chronicles: representation, on-line recognition and learning. In *Fifth international conference on knowledge representation and reasoning (KR'96)*, pages 597–606, Cambridge (USA), 1996.

Gue96. M. Guertin. Abductive inference of events: Diagnosing cardiac arrhythmias. In *Proceedings of the Florida Artificial Intelligence Research Symposium*, May 1996.

HK95. J. Hunter and I. Kirby. Ticker: A qualitative model of the electrical system of the heart. In *Research and Development in Expert System XII (Proc. of Expert Systems'95)*, pages 293–307. Information Press, 1995.

KNB98. M. Kundu, M. Nasipuri, and D.K. Basu. A knowledge based approach to ECG interpretation using fuzzy logic. *IEEE Trans. Sytems Man and Cybernetics*, 28(2):237–243, 1998.

Lon96. W. J. Long. Temporal reasoning for diagnosis in a causal probabilistic knowledge base. *Artificial Intelligence in Medicine*, 8:193–215, 1996.

LRD97. W. Van Laer, L. De Raedt, and S. Dzeroski. On multi-class problems and discretization in inductive logic programming. In Zbigniew W. Ras and Andrzej Skowron, editors, *Proceedings of the 10th International Symposium on Methodologies for Intelligent Systems (ISMIS97)*, volume 1325 of *Lecture Notes in Artificial Intelligence*, pages 277–286. Springer-Verlag, 1997.

RDF97. N. Ramaux, M. Dojat, and D. Fontaine. Temporal scenario recognition for intelligent patient monitoring. In *Proc. of the 6th Conference on Artificial Intelligence in Medecine Europe (AIME'97)*, 1997.

SCLB95. P. Siregar, M. Chahine, F. Lemoulec, and P. Le Beux. An interactive qualitative model in cardiology. *Computers and Biomedical Research*, 28(6), 1995.

Shi85. T. Shibahara. On using causal knowledge to recognise vital signals : Knowledge-based interpretation of arrhythmia. In *Proc. of IJCAI'85*, pages 307–314, 1985.

SMKS96. A. Srinivasan, S. Muggleton, R. King, and M. Sternberg. Theories for mutagenicity: a study of first-order and feature based induction. *Artificial Intelligence Journal*, 85(1,2):277–299, 1996.

TW93. D. A. Tong and L. E. Widman. Model-based interpretation of the ECG: A methodology for temporal and spatial reasoning. *Computers and Biomedical research*, 26(3):206–219, 1993.