

Learning structural knowledge from the ECG

F. Wang¹, R. Quiniou², G. Carrault¹, M.-O. Cordier²

¹ LTSI, Campus de Beaulieu, 35042 Rennes Cedex FRANCE

² IRISA, Campus de Beaulieu, 35042 Rennes Cedex FRANCE

Abstract. We tackle the problem of discovering, without the “manual” aid of an expert, implicit relations and temporal constraints from a collection of dated events detected on temporally structured signals. The approach associates tightly signal processing and symbolic learning methods. It is illustrated on learning cardiac arrhythmias from ECGs.

1 Introduction

Signal processing specialists are often confronted to the problem of recognizing a sequence of events corresponding to the signature of a particular process state. This difficulty may be even greater in biomedical domains where processes behavior is poorly known and consequently the related scenarios that must be recognized. This is the case in neurology, for instance, where the relation between electroencephalogram peaks (transient paroxistic events associated to the pathology) and epilepsy crises is not well explained. This is also the case in cardiology, where precursor signs of arrhythmias in sinusal beats are poorly known. Then the question is: “if a relation exists between signals and process states, how can it be automatically qualified?”.

At an abstract level, the problem is to discover from a collection of events $E = \{E_1, \dots, E_n\}$ occurring respectively at time $\{t_1, \dots, t_n\}$ the implicit relations $\mathfrak{R}(E_j, E_k)$ and the temporal constraints $TC(t_j, t_k)$ between the events E_j et E_k . To illustrate the idea, look at the process trace in figure 1. The problem is to discover that events of type E_1 always precede events of type E_4 with a temporal delay $700 \text{ ms} < \Delta_t < 900 \text{ ms}$, independently of events of type E_2 and E_3 .




Fig. 1. Temporally structured events

We propose to associate signal processing and machine learning to discover such relations. The signal processing module abstracts the signal data into series of symbolic time-stamped events that an inductive logic programming (ILP) learner processes in order to discover relevant and discriminant high-level temporal signatures called chronicles. The resulting descriptions are used as reference patterns in an intelligent monitoring system which associates the above signal processing module and a specialized matcher [Dou96].

We illustrate our approach on the discovery of expert rules for recognizing cardiac arrhythmias from electrocardiograms (ECG) [CCQ⁺99]. Several attempts to learn knowledge for monitoring cardiac patients have already been made. They use formalisms such as expert rules [BML89], neural networks [Wat95], attributed grammars [Sko90,KAG97] or fuzzy rules [KNB98]. The originality of our approach lies in the fact that we use symbolic events and explicit temporal constraints that makes the knowledge more readable and understandable by clinical staff.

In the next section we develop the proposed approach and give some details about the signal processing and the machine learning techniques that are used. Then we present some results obtained in the field of cardiology.

2 Associating signal processing with machine learning

An important constraint in biomedical domain is to express knowledge at an abstraction level that is close to experts' own knowledge in order to let them assess the data. Very early, (expert) rules were recognized several advantages such as simplicity, uniformity, transparency, readability, etc. These properties make them desirable for applications such as on-line monitoring where signal processing techniques are involved but also where information presentation is an important aspect since they motivate the decisions. Rules are difficult to acquire from experts and machine learning techniques have been proposed to achieve this task. As first-order rules containing temporal constraints are the target, inductive logic programming (ILP) appears the most adequate [Mug99].

2.1 The signal processing module

The signal processing module aims at transforming the numerical input sequence into series of labeled symbolic events. More precisely, it aims *i*) at detecting and at identifying the markers of the cardiac activity (*P* waves, *QRS* complexes), and *ii*) at classifying waves in normal or abnormal classes. This information (occurrence time, wave duration, morphology, ...) is used to generate the symbolic event sequence that will be used by the chronicle recognition module.

The *QRS* detection is achieved by Gritzali's algorithm [Gri88] which can be seen as a generalization of several detectors combining different channels. Several versions of the algorithm were tried, varying the number of ECG channels and the combination method (distributed or centralized combination).

The *P* wave detection is far more complex. The new approach that is used here relies upon observing that the *P*, *QRS* and *T* waves in the ECG signal have different morphologies that enables a reliable classification by a neural

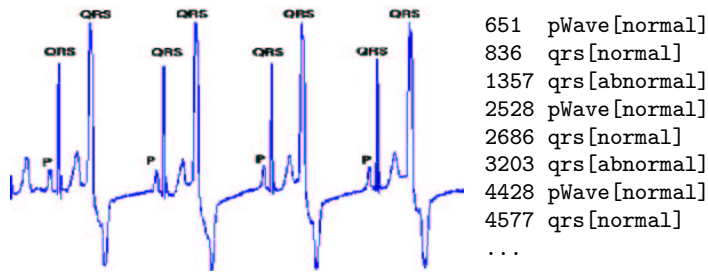


Fig. 2. Events extracted from the bigeminy ECG on the left.

network. After *QRS* complexes, high frequency noise and baseline wandering have been adaptatively eliminated [HC99], peaks above a certain threshold that may correspond to candidate *P* waves are looked for. Then, a multi-layered neural network classifies the candidate waves as being a real *P* wave or noise.

Wave qualification into normal or abnormal is achieved by wavelet decomposition, the results of which are fed into a probabilistic neural network. After this final step, interesting waves are translated into symbolic events (see figure 2), i.e a label, *p* or *qrs*, an occurrence time and a qualification is assigned to events.

2.2 Learning cardiac arrhythmias from symbolic events

Machine learning covers inductive techniques that aims at discovering definitions of a target concept from positive and negative examples of this concept. Inductive logic programming (ILP) is such a technique that induces concepts as first-order logical formulas. ILP has been successful at knowledge discovery in domains such as molecular biology or engineering [Mug99]. It can be viewed as searching a clause space for hypotheses that respect a particular syntax, the hypothesis language, specified by a learning bias. The selected theories, i.e. sets of clauses, must further be complete (they must cover all the positive examples) and correct (they must cover no positive example).

ICL [RL95] proposes a high-level concept specification language called DLAB in which the hypothesis language syntax can be defined. DLAB grammars are preprocessed in order to generate candidate hypotheses from the most general to the specific ones. We have used several DLAB specifications in order to induce set of rules enjoying such properties as readability, efficiency or robustness. A DLAB grammar consists in rule templates that fix the syntactic form of clauses defining the target concept. It makes use of operators *1-h*: meaning that from 1 to *h* elements must be taken from the list following the operator. *len* is the total length of the list.

The notion of cardiac cycle is commonly used to define rhythm disorders. We used this notion for bias specifications too. A rule expressing an arrhythmia is composed of several successive cardiac cycles. As specified in figure 3, each cycle is described as either a *P* wave followed by a *QRS* (lines 2-3), or an isolated *P* wave (line 6), or an isolated *QRS* (line 8). Additional constraints on inter-wave temporal intervals (*PP*, *PR*, *RR*) can be also added (lines 4-5, 7, 9).

```

1   1-1:[
2     len-len:[p_wave(P1, 1-1:[normal, abnormal], R0),
3       qrs(R1, 1-1:[normal, abnormal], P1),
4       0-len:[rr1(R0, R1, 1-1:[short, normal, long]),
5         pr1(P1, R1, 1-1:[short, normal, long])]],
6     len-len:[p_wave(P1, 1-1:[normal, abnormal], R0),
7       pp1(P0, P1, 1-1:[short, normal, long])],
8     len-len:[qrs(R1, 1-1:[normal, abnormal], R0),
9       0-1:[rr1(R0, R1, 1-1:[short, normal, long])]]]

```

Fig. 3. Syntactic specification of a cardiac cycle in DLAB

3 Results

The goal of the experiments was to demonstrate that understandable and useful arrhythmia specifications could be learnt from ECG temporal data. A second objective was to assess the flexibility of using a declarative bias for imposing desirable properties such as readability or robustness on induced concepts.

3.1 Data

Real ECGs were taken from the *MIT-BIH* database [MM90]. Normal sinus rhythm plus three classes of examples related to well-known arrhythmias have been selected: left bundle branch block (*lbbb*), Mobitz type II (*mobitz2*), bigeminy.

20 ECG lasting 10s each have been extracted for each class and translated into a symbolic form by the signal processing module described above. A symbolic ECG example is a set of events `wave(Evt, Type, Time, Qual, P_evt)` stating that `Evt` is related to a wave of type `Type` (`p` or `qrs`), which occurred at time `Time`, the shape of which is `Qual` (`normal` or `abnormal`) and `P_evt` just precedes `Event` on the ECG. This last argument provides a temporal structure to the set of events. It should be noted that this kind of specifications can be easily built from lists of symbolic events such as the list in figure 2.

3.2 Experiments

Learning shortest rules The first experiment was performed with a learning bias imposing only one mandatory cycle. The induced rules are supposed to be the most efficient ones for recognition as they involve the minimal number of events. The results obtained after a 10-fold cross validation (10 learning rounds on 90% of the examples and test on the remaining 10%) are shown in figure 4. The resulting accuracy (the ratio number of correctly recognized/total number of examples) was 100% for training as well as for test.

The rules are expressed as Prolog clauses. For instance, the first rule says that a bigeminy can be characterized by an abnormal *QRS* followed by a normal *P* wave followed by a normal *QRS* and the temporal interval between the two *QRS*s is long. Note, that this rule characterizes also trigeminy. As no examples were given for trigeminy there is no reason that the induced rule discriminates

```

class(bigeminy) :-    %[17, 0, 0, 0], [0, 18, 18, 19]
    qrs(R0, abnormal, _),
    p_wave(P1, normal, R0), qrs(R1, normal, P1), rr1(R0, R1, long).
class(lbbb) :-        %[0, 18, 0, 0], [17, 0, 18, 19]
    p_wave(P0, normal, _), qrs(R0, abnormal, P0),
    p_wave(P1, normal, R0), qrs(R1, abnormal, P1), rr1(R0, R1, normal).
class(mobitz2) :-    %[0, 0, 16, 0], [17, 18, 2, 19]
    p_wave(P0, normal, _), equal(P0, R0),
    p_wave(P1, normal, R0), qrs(R1, normal, P1), pr1(P1, R1, normal).
class(mobitz2) :-    %[0, 0, 2, 0], [17, 18, 16, 19]
    p_wave(P0, normal, _), equal(P0, R0),
    p_wave(P1, normal, R0), qrs(R1, abnormal, P1), pr1(P1, R1, normal).
class(normal) :-     %[0, 0, 0, 19], [17, 18, 18, 0]
    p_wave(P0, normal, _), qrs(R0, normal, P0),
    p_wave(P1, normal, R0), qrs(R1, normal, P1), rr1(R0, R1, normal).

```

Fig. 4. Rules induced with a bias imposing 1 mandatory cycle, 2 optional cycles

between bigeminy and trigeminy. This illustrates the fact that most general rules are induced. The lists following the rule heads give the number of examples covered by the rule and the number of examples not covered by the rule in each class (in the order bigeminy, lbbb, mobitz2 and normal).

Learning more readable rules As they represent recurrent phenomena, some arrhythmias are often described by experts on several cycles, e.g. 3 or 4. In this second experiment the number of mandatory cycles was set to 3. This bias causes the rules to be more specific and eventually they cover less examples. Sometimes, several rules are necessary to define a class, as for bigeminy:

```

class(bigeminy) :- % [13, 0, 0, 0], [5, 19, 17, 18]
    p_wave(P0, normal, _), qrs(R0, normal, P0),
    qrs(R1, abnormal, R0), rr1(R0, R1, short),
    p_wave(P2, normal, R1), qrs(R2, normal, P2).
class(bigeminy) :- % [5, 0, 0, 0], [13, 19, 17, 18]
    p_wave(P0, normal, _), qrs(R0, normal, P0),
    p_wave(P1, normal, R0), qrs(R1, abnormal, P1),
    p_wave(P2, normal, R1), qrs(R2, normal, P2).

```

The first rule corresponds to the usual bigeminy. The second rule illustrates an interesting phenomenon. The examples covered by this rule refers to ECGs where the ectopic beats occur between a *P* wave and the expected following *QRS*. This explains the presence of a large *QRS* which reflects an abnormal ventricle activation. This illustrates the power ILP that is able to discover such subtleties and use them for discrimination. Accuracy of 10-fold cross validation was still 100% for training and test in this experiment.

4 Conclusion

We have described how to associate signal processing and inductive logic programming to learn structural patterns from real signals and illustrated the ap-

proach on learning cardiac arrhythmias. Since the domain of cardiology is well known, the learning results could be assessed by specialists on different criteria such as relevance or readability. The learnt definitions discriminate between those classes for which examples are given and only those. This explains why they are sometimes a bit strange to clinicians used to definitions that discriminate between all known arrhythmias. However, all learnt definitions were rated good and relevant.

The flexibility of the ICL bias language was demonstrated. This makes it possible to tailor the learnt definitions to the particular context, explicative or efficient monitoring, for instance. The rules could be adapted to an individual patient by learning from his own ECGs or from arrhythmias he may suffer from.

A first extension of this work concerns taking into account multiple sources of information, i.e. multichannel ECGs as well as ECGs associated with hemodynamic signals. A second extension aims at applying the same approach to less known domains such as EEG analysis or processing signals during anesthesia.

References

- [BML89] I. Bratko, I. Mozetic, and N. Lavrac. *Kardio: A Study in Deep and Qualitative Knowledge for Expert Systems*. MIT Press, 1989.
- [CCQ⁺99] Guy Carrault, Marie-Odile Cordier, René Quiniou, Mireille Garreau, Jean-Jacques Bellanger, and Alain Bardou. A model-based approach for learning to identify cardiac arrhythmias. In *Proceedings of AIMDM'99*, volume 1620 of *LNAI*. Springer Verlag, 1999.
- [Dou96] C. Dousson. Alarm driven supervision for telecommunication networks. ii- on-line chronicle recognition. *Annales des Télécommunications*, 51(9-10):501–508, 1996.
- [Gri88] F. Gritzali. Towards a generalized scheme for QRS detection in ECG waveforms. *Signal Processing*, 15:183–192, 1988.
- [HC99] A. I. Hernández and G. Carrault. Annulation adaptative du QRS. In *colloque GRETSI*, pages 591–594, 1999.
- [KAG97] G. Kókai, Z. Alexin, and T. Gyimóthy. Application of inductive logic programming for learning ECG waveforms. In *Proceedings of AIME97*, volume 1211, pages 126–129. Springer-Verlag, 1997.
- [KNB98] 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.
- [MM90] G.B. Moody and R.G. Mark. The MIT-BIH arrhythmia database on cd-rom and software for use with it. In *Computers in Cardiology 1990*, pages 185–188. IEEE Computer Society Press, 1990.
- [Mug99] S. Muggleton. Inductive logic programming: issues, results and the challenge of learning language in logic. *Artificial Intelligence*, 114(1-2):283–296, 1999.
- [RL95] L. De Raedt and W. Van Laer. Inductive constraint logic. In *Proceedings of the 5th Workshop on Algorithmic Learning Theory*, Lecture Notes in Artificial Intelligence, 1995.
- [Sko90] E. Skordalakis. ECG analysis. In H. Bunke and A. Sanfeliu, editors, *Syntactic and Structural Pattern Recognition Theory and Applications*, pages 499–533. World Scientific, 1990.
- [Wat95] Raymond L. Watrous. A patient-adaptive neural network ECG patient monitoring algorithm. *Computers in Cardiology*, pages 229–232, 1995.