On Using Causal Knowledge to Recognize Vital Signals: Knowledge-based Interpretation of Arrhythmias

Tetsutaro Shibahara

Dept. of Computer Science, Unir. of Toronto, Toronto, Ontario, CANADA

ABSTRACT

We have proposed a knowledge-based system for the recognition of time varying vital signals, such as electrocardiograms. This paper discusses its causal model approach.

A prototype system exhibits the efficacy of the method of knowledge base stratification, where each knowledge base (KB) represents a distinct perspective to the phenomenon, such as the observable waveform knowledge and the causal knowledge of the physiological entity. Projection links in our frame representation relate corresponding concepts in different KBs, *e.g.,* abnormalities in shape or temporality are mapped into abnormalities in physiological causality. The role of projections in the recognition process is to transduce established waveform hypotheses into corresponding event hypotheses and to form more global hypotheses using the causal model of entity.

Several types of one-shot causal links have been introduced to represent causal relationships among underlying physiological events. A causal link includes the existential dependency and the implicit temporal constraints between the effect and the cause events.

Using the causal knowledge with event statistics, the recognition system makes expectations for unseen events in relation to already-observed events when partial input information is given. Statistical information defined coherently through metaclasses of the PSN language supports a default reasoning process. The overall recognition framework is based on the hypothesize-and-test paradigm and the specialization-and-aggregation of hypotheses using similarity links in IS-A hierarchies and causal links in PART-OF hierarchies.

1 Introduction

The knowledge based systems approach has been applied to the recognition problem of time-varying vital signals such as electrocardiograms (ECGs). The developed recognition system uses a causal model of the physiological entity so that observed abnormalities of the temporality or morphology of the signal are explained by referring to the corresponding abnormalities of causal events and relationships in the entity model.

In the domain of electrocardiology, this causal reasoning process is especially important because the domain involves causal and temporal knowledge about the cardiac conduction system, with which cardiologists analyze clinical observations (ECGs) and thereby provide diagnostic interpretations of abnormal events in the underlying physiological mechanism of the heart. The recognition problem of ECG rhythm disorders, above all, is interesting because the overall performance of existing ECG programs (e.g., IBM Bonner's program) is at most 80% reliable for abnormal ECGs [Hagan79] and we believe a basic reason for this unreliability is that current systems lack underlying physiological knowledge to handle the complexity inherent in cardiac rhythms. The ECG wave identification is much complicated by its "antenna" nature of receiving only the aggregated of the electrical activity of the heart, i.e., there is no simple correspondence between signal features and individual electrical discharges in the heart.

Our approach to the problem of building such a system is to construct a knowledge base stratified by several distinct knowledge bases (KBs) from different perspectives of the domain. Its control structure, therefore, supports a guiding mechanism between corresponding concepts in different KBs as well as another guiding mechanism between causally related concepts in each KB. In our representational terms, the former is called projection links and the latter is causal links, and these links together contribute to the generation of hypotheses and the decision of overall interpretations in the recognition of ECG signals. This approach also integrates several established AI techniques. The system inherited the basic control framework from the ALVEN system [Tsotso8B0,Tsotsos85] such as the attention mechanism for specialization and aggregation, which is supported by the implementation of similarity links [Minsky75] and the exception handling mechanism. The hypothesize-and-test paradigm is used as in ALVEN and other systems like PIP [Szoiovits78] and HEARSAY-II [Mostow78]. The knowledge organizational method is based on the IS-A, PART-OF. and INSTANCE-OF hierarchies as used in the PSN (Procedural Semantic Network) formalism [Mylopoulos83].

To prove the efficacy of our methods, a prototype system called CAA (Causal Arrhythmia Analysis system) has been designed and implemented using a frame-input PSN system on Franz LISP (and UCI LISP) [Shibahara83]. The prototype with a Limited size of knowledge base is being tested and yielding so far satisfactory results.

2 Causality

2.1 Representation of Causal Connections

Causality may be viewed from its various facets. Rieger and Grinberg discerned the one-shot causality where the cause event(s) is required only at the start of the effect event(s) from the continuous causality where the continuous presence of the cause is required to sustain the effect. [Rieger76]

CAA causal links are based on two features of causal connections: first, they specify the existential dependency of an affected event on its causative event(s); second, they impose temporal constraints between causative and affected events. Thus, the affected events cannot occur without the occurrence of the corresponding causative events, with effects temporally following their causes. Interested in representing the dependencies of causal connections among events more precisely, we look at causality from the viewpoint of whether a causal influence is internal to a subject or it influences other distinct subject(s). One-shot causal links, therefore, are specialized into the following:

(1) Transfer: the subject of the event normally completes the current event and proceeds to the following event.
(2) Transition: the subject is forced to terminate its current event and proceed to a new event.
(3) Initiation: the causative event, due to a given subject, triggers a new event of another subject.
(4) Interrupt: the causative event, due to a given subject, interrupts and forces the termination of an event by another subject.
(5) Causal-block: the causative event of a subject, fails to influence an event of another subject due to a blockage of the causal flow.

The above CAA causal links include implicit temporal constraints; thus, causal structures are described more qualitatively without specifying time coordinate values.

Causal events are aggregated at several levels involving arbitrary numbers of causal links. However, causal links themselves remain atomic lest the semantics of causal connections become ambiguous.

2.2 Use of Causal links

To interpret real ECG signals, the knowledge base must contain the causal knowledge about normal and abnormal connections among cellular events, which produce particular ECG tracings in the observable signal domain. We represent such causal activities using CAA causal links. Fig. 1 illustrates a typical ECG tracing for a normal cardiac cycle in (a), its electrical conduction path in an anatomical diagram in (b), and the corresponding causal conduction model with causal links in (c). In this causal model, short symbols like EOa are used to denote one of four basic events (phases) in a small portion of the cardiac conduction system; these phases are "depolarization" [with symbol a], "under-repolarization" [with symbol b], "partial-rrpolarization" [with symbol c], and "full-repolarization" [with symbol d]. Such basic phase events are successively aggregated into "cycle", "activity", "beat", and "beat-pattern" events in the physiological event KB to describe more global and complex causal structures.

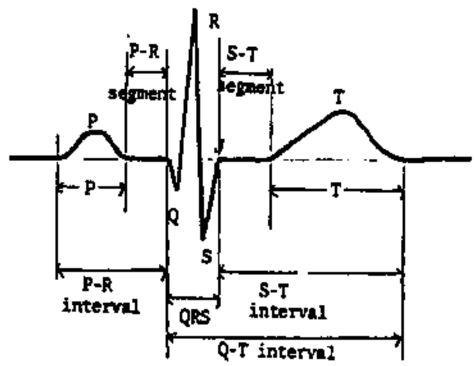

Fig. 1-(a)
Typical ECG Tracing for a Normal Cardiac Cycle

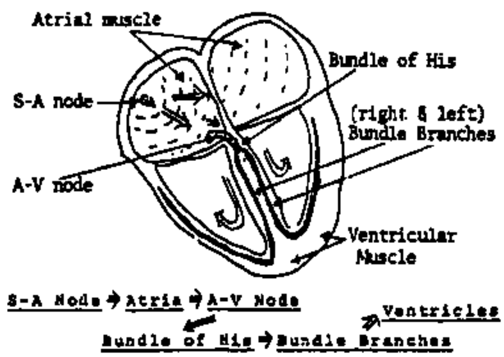

Fig. 1-(b)
Electric Conduction Path and Anatomy

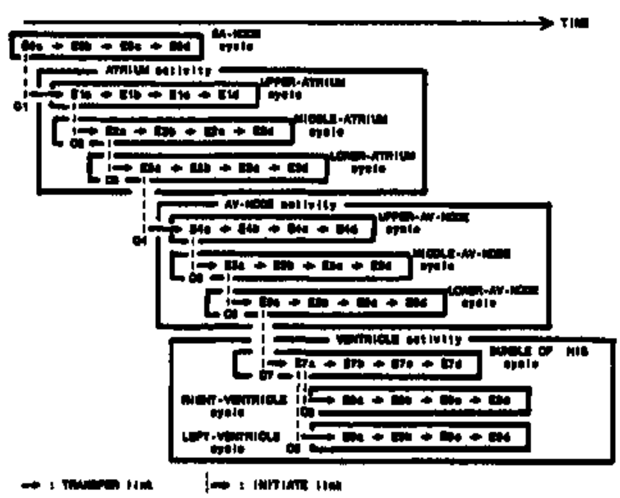

Fig. 1-(c) Internal Causal Structure of a Normal Cardiac Cycle

Note that causal links across beat events (not shown) are TRANSITIONS and INTERRUPTS except pace-making parts (normally, the SA-Node) because the overall oscillation of the conduction system is controlled (or triggered) by such self-oscillating cells. Also, since the current model is rather devoted to supraventricular arrhythmias, the bundle branches are included in the ventricles.

2.3 Recent Work on Causality

ABEL and CADUCEUS systems are recent medical expert systems that use causal notions. The ABEL system provides multiple levels of descriptions of medical hypotheses and hierarchically organizes disease structure [Patil81]. In the CADUCEUS system, analyzing differential diagnoses and causal graphs of diseases, Pople proposed sophisticated control links for efficient decision making [Pople82]. In spite of the sophistication in expressing causal mechanisms in ABEL and CADUCEUS, these systems do not seem to provide a means to construct a recognition system of time-varying signals due to the weakness in the representation of precise timing contexts among events.

Causality has been recently approached from the standpoint of "qualitative reasoning". In this regard, Long's work must be noted [Long83]. He introduced qualitative times to describe the causal relations that might or must have taken place. He interestingly proposed four causal templates that give an extension of "continuous causality" while our causal links are specialized in "one-shot causality". We have taken a different approach because original signals are given to the system as real-valued data and the use of some quantitative analysis is inevitable at the measurement level so that unnecessary ambiguity is avoided, as Kunz noticed in his AI/MM system [KunzB3].

Based on the methods of multivariate analysis Blum approached the problem statistically [BlumB2]. However, our domain includes mostly exact causal relationships. Therefore, we limit the use of statistical standards to the estimation of inherently spontaneous variables such as event durations.

3 Representation of Domain Knowledge

3.1 Frame Representation and Classes

Our knowledge representation is based on semantic networks, in particular, the PSN language, with IS-A, PART-OF, and INSTANCE-OF organizational relations. In our particular formalism, concepts such as events and waveforms are described by frames and called class-frames or classes. Fig. 2 exemplifies the use of a frame and causal links. (The dot "." notation is used to specify the component of the referred slot.) This normal activity of the ventricles is decomposed into three cycle events, *i.e.,* bundle-of-his-cycie-event, right-ventricle-cycle-event, and left-ventricie-cycle-event. Two INITIATE links represent the conductions from the bundle of His to the left and the right ventricles, respectively. Note that the information related to the class itself, in this case, the subject part name and the activation type, is given as the instantiation of a metaclass ACTIVITY-CONCEPT.

```
classVENT-ALL-MATURE-FORWARD-ACTJVITY
  is-a VENT-ACTIVITY
  instance-of   ACTIVITY-CONCEPT    instantiated-with
                subject: VENTRICLE;
                activation: FORWARD;;
with    components
 bundle-of-his-cycie-event: BHIS-MATURE-CELL-CYCLE;
 right-ventricle-cycle-event: RV-MATURE-CELL-CYCLE;
 left-ventricle-cycle-event: LV-MATURE-CELL-CYCLE;
 bhis-rv-delay:NUMBER-WITH-TOLERANCES
      calculate := /* delay set-up expression •/;
 bhis-lv-delay:NUMBER-WITH-TOLERANCES
      calculate := /* delay set-up expression •/;
causal-links
 bhis-rv-propagation: INITIATE
  causative-starting-event:
   bundle-of-his-cycle-event.depolarization-phase-event;
  initiate d- event:
   right-ventricle-cycle-event.depolarization-phase-event;
  delay: bhis-rv-delay;;
 bhis-lv-propagation: INITIATE
  c    ausative-starting-event:
   bundle-of-his-cycle-event.depolarization-phase-event;
  initiated-event:
   lcft-ventricle-cycle-event.depolarization-phase-event;
  delay: bhis-lv-delay;;
end
```

Fig. 2 Class Frame for Normal Activity of the Ventricles

3.2 IS-A and PART- OF Hierarchies in Knowledge Base

Let us examine how the IS-A and the PART-OF principles contribute to the organization of the CAA knowledge base. We take a look at the QRS and QRST waveforms in the ECG waveform KB as examples.

First, the QRST waveform consists of the QRS complex and the T wave; thus, the corresponding class QRST-COMPOSITE-WAVE-SHAPE has the generic PART-OF structure with major components shown in Fig. 3-(a). This generic QRST waveform is specialized into several QRST waveforms in Fig. 3-(b), along its IS-A hierarchy. Let us pick up one component from the STANDARD-QRST-COMPOSITE-SHAPE. NORMAL-QRS-COMPLEX is such a component and this class is itself included in the IS-A hierarchy of the QRS waveforms as in Fig. 3-(c). The orthogonality of IS-A and PART-OF hierarchies is shown in Fig. 3-(d) since STANDARD-R-WAVE-SHAPE is a component of STANDARD-QRS-COMPLEX-SHAPE and, at the same time, it is included in a local IS-A hierarchy of R-WAVE-SHAPE.

Similarly, various IS-A and PART-OF hierarchies are defined in the physiological KB. Such organizational hierarchies not only contribute to the clarification of the inter-dependency among domain concepts but also provide guiding knowledge for the recognition process as discussed later.

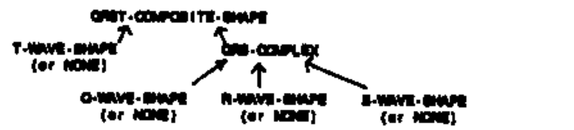

Fig. 3-(a) PART-OF Structure of QRST-COMPOSITE-SHAPE

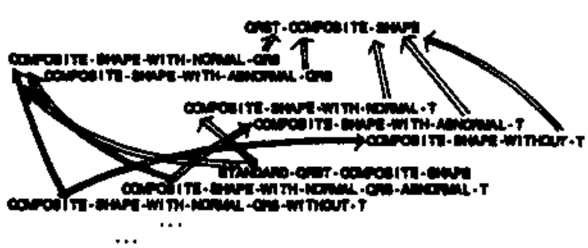

Fig. 3-(b) Specialization of QRST-COMPOSITE-SHAPE

Fig. 3-(c) IS-A Hierarchy of QRS-COMPLEX-SHAPE

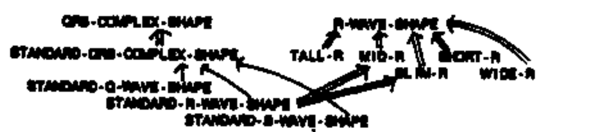

Fig. 3-(d) Orthogonality of IS-A and PART-OF hierarchies

3.3 Metaclasses and Statistical Standards

Statistical information, so commonly used in medical reasoning systems, has particular importance when insufficient information is available about the disease status of a patient [Szolovits78]. In our case, the recognition system uses statistical standards to make expectations for unknown attributes of events and to estimate consistencies (goodness-of-fit) of hypotheses. Since statistical standards about a class are not the attributes of any particular instance of the class but the attributes of the class itself, such standards could be defined in appropriate metaclasses and instantiated as properties of the class itself. In other words, event statistics are good examples of meta-knowledge or "knowledge 'about knowledge" and such knowledge is organized along the INSTANCE-OF axis. In fact, to provide "mean" and "standard-deviation" values to all the physiological phase events, CAA has a metadata CELL-PHASE-CONCEPT as shown in Fig. 4.

metaclass CELL-PHASE-CONCEPT
with *components*
subject: HEART-PORTION;
maturity: DEGREE-OF-MATURITY;
phase: PHASE-NAME;
mean: EXPRESSION *default* MEANFUNC;
deviation: EXPRESSION *default* DEVFUNC;
end

Fig. 4 Metaclass Definition for Statistical Information

In Fig. 4. default functions MEANFUNC and DEVFUNC are generic functions that are supposed to generate mean and standard deviation about durations of phase events. Such statistical standards about phases are function procedures of "subject", "maturity", "phase", and a state variable HR$ (heart rate). Therefore, such a standard, for example, a mean value is given by the expression "(mean subject maturity phase HR$)" in a particular phase event class. In the evaluation of this expression, the slot-names such as "mean" and "subject" are replaced by real properties of the class, such as "MEANFUNC" and "SA-NODE". This is considered as the tailoring process of the general "mean" expression to the definitional context of this event; *i.e.*. such statistics may change to fit into each event hypothesis. On the other hand, HR$ is a global variable that reflects the current state of the model, where hypotheses are being instantiated; in other words, such global variables are used to make statistical standards sensitive to the current recognition context. Heart rate, blood pressure and breathing rate are examples of dynamic or time varying global variables while age-group, sex, race, and types of medications are static global variables. Obviously, default functions, MEANFUNC and DEVFUNC, may be replaced by any ad-hoc functions if necessary.

3.4 Knowledge-base Stratification and Projection Links

Due to our causal model approach, we at least distinguish two subdomains, i.e., the ECG morphological (shape) domain and the electrophysiological domain. Therefore, the system's knowledge base is stratified by the ECG waveform KB and the physiological event KB. Our idea of stratifying a knowledge base resembles Rich's "overlays" since it provides different perspectives to the problem [Rich81]. In our method, however, the linking mechanism between different KBs is biased to recognition purposes.

Projection links have been introduced into the CAA system to relate corresponding concepts in distinct domain KBs. In our model based approach such links are essential since they relate temporal and/or morphological abnormalities in waveforms to corresponding abnormalities in physiological causal structures.

The diagram in Fig. 5 illustrates a projection link that defines the correspondence between the corner point information of a normal QRST waveform and the timings of a normal activity event of the ventricles. This

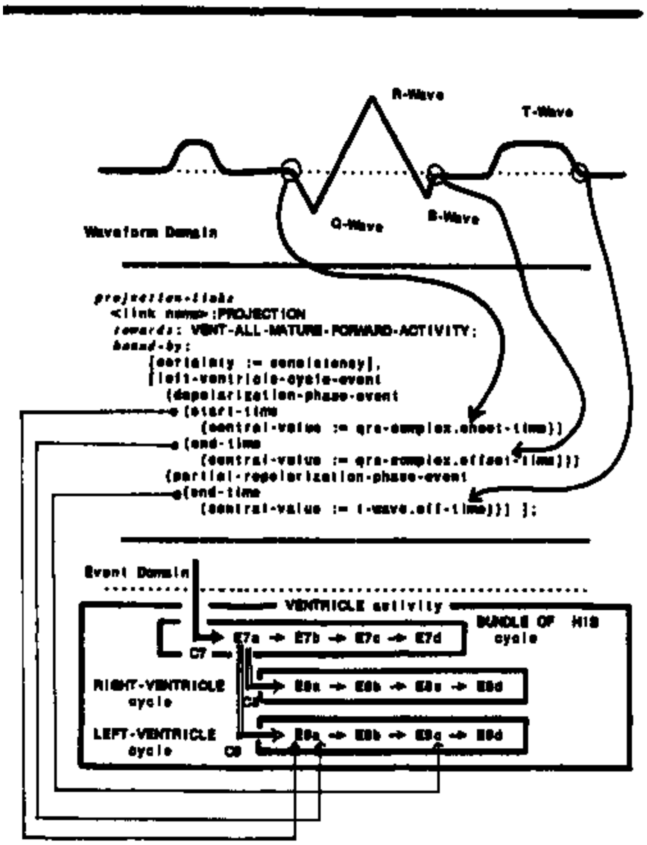

Fig. 5 Bindings by a Projection from Wave to Event

projection link must be defined in the class frame of the normal QRST waveform.

For recognition, the most important aspect of projection links is that they provide guiding paths to map concepts across differently organized KBs and support the synchronization of recognition activities in different domains. In our system, projections from established waveform hypotheses result in the basic data set (hypotheses) in the underlying event domain, on which the recognition of causal events works.

4 Recognition Strategies and Control System

4,1 Outline of Recognition Flow

Signals are processed by three functional modules in the following order:

(1) The peak-detection module extracts wave segments and slopes from sampled ECG input signals and emits peak tokens with the measured parameters. This module uses the syntactic method given by Horowitz [Horowitz75] based on piecewise linearization and parsing techniques using a context-free grammar.

(2) The waveform analyst module, for each cardiac cycle, forms waveform hypotheses on the peak tokens and refines the hypotheses to describe the given set of tokens best. Once established, such hypotheses are projected into the physiological event domain to form their corresponding event hypotheses.

(3) The errant analysis modulle accepts projected events as a starting data set and generates rhythm event hypotheses in a more global context of time to elucidate rhythm abnormalities in the underlying cardiac conduction system. Since most of physiological events do not have observable counterparts (waveforms), the event analysis module makes expectations on the attributes of unseen events using the causal knowledge of the conduction system and statistical standards of events. If the system encounters lack of information because of missing waves, it may request the peak-detection module to search for such missing tokens based on the expectation of such waves.

4.2 Specialization and Aggregation for Hypothesis Generation

Our recognition strategy is based on the hypothesize-and-test paradigm, in particular, the attention mechanisms of the ALVEN system. The focus-of-attention mechanism makes recognition (hypothesis formation) proceed from the generic to the specific along IS-A class hierarchies downward. When a ciaBS hypothesis succeeds, a focusing action is taken by choosing and hypothesizing an arbitrary specialized class of the succeeded class. When a current hypothesis failed, the change-of-attention mechanism chooses alternative hypotheses through similarity links, examining the similarity and the difference between classes.

Let us examine how the above specialization and aggregation process works for QRST waveforms (see Fig. 3). After all peaks are detected and measured, the waveform analysis module chooses groups of consecutive prominent peaks with high amplitude and steep slope as *anchoring shapes.* These anchoring shapes are candidates for QRST-COMPOSITE-SHAPE. The wave analysis for an anchoring shape starts with hypothesizing the class QRST-COMPOSITE-SHAPE on the prepared set of basic peak tokens, as the first step. This class is most generic for all the shapes composed of Q, R, S, and T waves and only requires the existenoe of any QRS complex wave as the sole component; thus, this component class, which is again the most generic class for QRS complex waves, is hypothesized and its instantiation follows using the prepared Q, R, and/or S wave tokens. If there is none of Q, R. or S wave tokens, the hypothesis of QRS-COMPLEX-SHAPE fails and so does QRST-COMPOSITE-SHAPE, too. As the second step, one of specialized QRST composite wave classes under QRST-COMPOSITE-SHAPE is hypothesized and all its attributes are tested, i.e., the slot tokens are tried to be instantiated. Since all the specialized classes are connected by similarity links, using exceptions raised by test results the system may choose the next appropriate hypothesis and finally reach the valid hypothesis for the given anohoring shape. The test procedure for each attribute slot, however, triggers an independent process to recognize the token of the slot. For example, class STANDARD-QRST-COMPOSITE-SHAPE has a slot named qrs-complex and this slot is defined by class NORMAL-QRS-COMPLEX which is an IS-A parent class to classes. STANDARD-QRS-COMPLEX-SHAPE, STANDARD-QR-COMPLEX-SHAPE. STANDARD-RS-COMPLEX-SHAPE and STANDARD-R-ONLY-COMPLEX-SHAPE; thus, the previous QRS wave slot token of the generic QRST-COMPOSITE-SHAPE must be specialized along the IS-A hierarchy of QRS-COMFLEX-SHAPE, and this process also uses the same procedure in order to reach the most refined QRS complex shape hypothesis. With such a specialized QRS wave token and a separately specialized T wave token, the second step decides the most appropriate hypothesis among QRST composite shapes for the given set of wave tokens.

Similarly but independently, in the physiological event domain, the specialization and aggregation process starts with the most generic beat pattern and eventually provides several specialized patterns as probable overall interpretations.

4.3 Projection Mechanism and Expectation Mechanism

The recognition starts with establishing hypotheses in the waveform domain. The projection mechanism maps such established hypotheses into the event domain, preparing a set of basic event hypotheses, which are treated like data in the event recognition process.

A beat pattern (rhythm) is a complex time-varying event aggregated from more local events such as beats, activities, cycles, and phases. Causal links in such an aggregated event imply connections among its component events. Thus, once projections are made to some of these components, the system can produce expectations of unknown components from the known components. Therefore, when the system hypothesizes such an aggregated event, it looks ahead or looks back for its component events which are causally linked to "already-established" component events. Most frequently, causal links are used to locate the temporal positions of "to-be-expected" events by their inherent temporal constraints. This expectation is made by the following basic equality implicitly imposed over starting or ending times of participating events:

<effect-time> = <cause-time> + <delay-period>.

Let us look at the above mechanisms in a clear small case where a QRST composite wave is seen but the P wave has not been recognized for the current wave group.

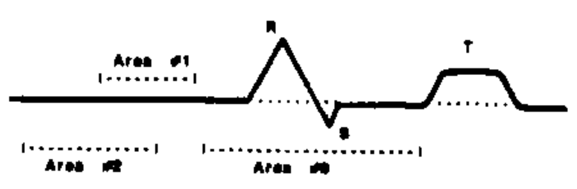

Fig. 6 Expected Areas of the P wave

Fig. 6 illustrates the case and the interval "Area #1" is the probable area where a P wave would appear if the beat is a normal sinus-pacing beat. To estimate such an area under a particular beat hypothesis is important since the peak-detection module may search for a P wave intensively in this area, again.

The area is estimated using the projection and the expectation mechanisms in the following fashion:

(0) A hypothesis of NORMAL-QRST-COMPOSITE-SHAPE is established.

(1) A projection to a normal ventricle activity event (Fig. 5):

(a) The onset and offset times of the QRS complex are bound to the starting and ending times of the depolarization phase of the left ventricle. The off-time of the T wave is bound to the ending time of the partial-repolarization phase. These phase events are generated immediately and two other phase events are expected by three TRANSFER causal links and event statistics. Thus, the left ventricle (LV) cycle event is generated.

(b) By the INITIATE causal link to the Bundle of His (BH1S) and subsequent TRANSFER links, the BHIS cycle event is generated. Also, by the INITIATE link from the BHIS to the right ventricle (RV). the RV cycle event is generated.

(c) With the above three cycle events, the projection to the normal ventricle activity event is completed.

(2) Expectation of the AV-Node activity, the Atrium activity, and the SA-Node activity under a hypothesis of the normal sinus-pacing beat (Fig 1-(c)):

(a) The INITIATE link C7 is invoked to expect the phase E8a. then E6b. E6c, and E6d phases are expected by three TRANSFER links, and finally, the lower AV-Node cycle event is generated. Similarly, using C8 and C5 INITIATE links, the middle and the upper AV-Node activity events are generated. Thus, the AV-Node activity event is formed with these component cycle events.

(b) Starting with the INITIATE link C4, the atrium activity event is expected in the same as above, and, next, the SA-Node cycle event is expected.

(c) A hypothesis of the normal sinus-pacing beat is completed.

Under this hypothesis, the on-time and the off-time of the P wave correspond to the starting time of the upper-atrium cycle and the ending time of the lower-atrium cycle, respectively. Therefore, the search area for a probable P wave is given as the interval between theses times *[e.g.,* from 110 ±16ms to 40 ±15ms before the QRS complex]. The request of the search for the P wave is fed back to the peak-detection module to repeat the detection with different sensitivity parameters.

The above CAA expectation mechanism is characterized by the following features:

(1) The expectation is made from the known to the unknown, forward or backward in time, and upward or downward in a PART-OF class structure.

(2) The expectation proliferates to make a closure of

temporal and/or structural dependencies and complete the PART-OF structure of the hypothesis.

Projections are made in the following fashion:

(1) Projections may be made between differently structured classes as seen in Fig. 5.

(2) To eliminate unnecessary instantiations of projections, any projected class is instantiated only when a current global hypothesis requests the class as a component.

4.4 Beat-pattern Analysis

To recognize a periodic or successive arrhythmia, its repetitive behavior is defined by the recursive definition of beat-pattern frames. By such a frame, recognition may proceed one beat to the next along the time axis instantiating successive beats to form the beat-pattern.

In the process of forming beat-patterns, causal links between adjacent beats allow the system to verify the causal relationship that govern the pace-making mechanism on a beat-to-beat basis. The overall consistency of a beat-pattern is calculated based on the consistencies of these causal links and beat components.

As well as the causal consistency among beats, overall characteristics and tendencies are observed and used to recognize individual arrhythmias. For this purpose, most beat-pattern classes include a component that monitors the changes of variables from one beat to another. A typical example is to monitor the change of the R-R interval or the P-R interval.

In arrhythmia beat-patterns, similarity links must also be defined to relate beat-patterns that have some features in common and handle situations where one or more matching exceptions have been raised. Fig. 7 shows ECG wave configurations that correspond to three different AV-Block arrhythmia patterns and the matching exceptions used by similarity links. Such similarity links between repetitive beat-patterns enable the system to switch beat-pattern hypotheses from one pattern to its alternatives according to the matching exceptions. For example, the class definition of Mobitz-1 second degree AV-Block contains a similarity link toward the Mobitz-2 second degree AV-Block and the first degree AV-Block for the situation where no progressive prolongation is seen in the atrium-ventricle-interval.

To recognize particular arrhythmia patterns, the specialization-and-aggregation process must be initiated with the most generic class for repetitive arrhythmias. The final interpretation, therefore, is given by a set of all survived beat-patterns with overall consistency factors. The consistency is is calculated using event statistics and a test-score function, which is similar to a fuzzy constraint in [Zadeh83].

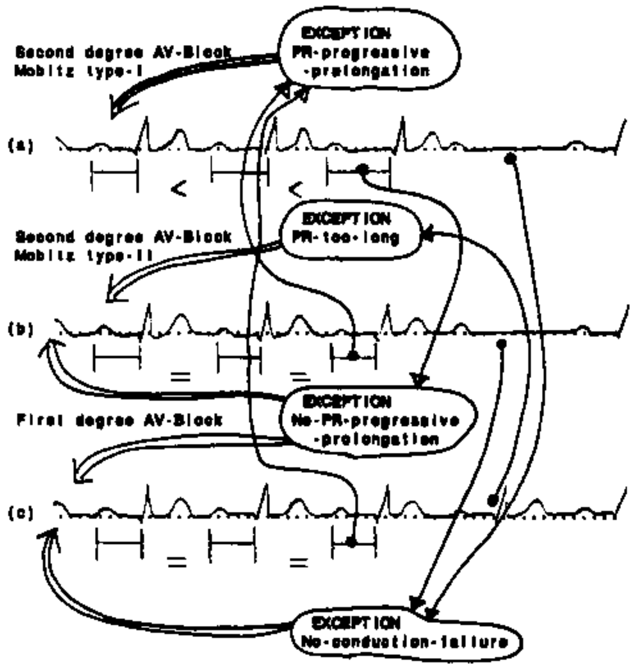

Fig. 7 Exceptions in Similarity Links among AV-Blocks

5 Concluding Remarks

We have discussed a recognition system guided by a causal model of the underlying physiological entity. We think a definite advantage of such a recognition system is the use of the causal knowledge and appropriate statistical knowledge about underlying events, enabling the system to make "expectations" of the event structure of the entity even when partial information is given. This approach is applicable in other medical applications because the ultimate purpose of recognition of vital signals is normally the elucidation of such causal abnormalities in the physiology, rather than merely the recognition of contour deformation or interval changes.

Our recognition system is effectively supported by various link constructs such as causal links, similarity links, and projection links, which are regarded as the distributed control knowledge. In fact, causal links support the expectation of unknown events using PART-OF hierarchies, similarity links guide the recognition in focusing or changing attentions using IS-A hierarchies, and projection links help hypothesis transduction across different KBs.

Our recognition method may not be successful when key waveforms are missing and no expected waves can be uncovered in spite of intensive search in the area of expectation. Therefore, we are designing the system that will interact with diagnostic interventions when the system faces lack of information. The purpose of such interventions is to perturb a patient's physiological state so that the state could move to a disease specific state.

ACKNOWLEDGEMENT

The CAA system has been developed as the author's Ph.D. work at the University of Toronto. A variety of people contributed to this work. As his supervisor, Prof. John Tsotsos of the Dept. of Computer Scienoe played a key role with his excellent insight and direction. Prof. John Mylopouios, who supervised the knowledge representational aspect, and his PSN group provided a fertile environment for this research. H. Dominic Covvey, director of Cardiovascular Computing at Toronto General Hospital (now at Clinicom International), contributed greatly to the development of the CAA system with his medical expertise and contacts. The author is especially thankful to Menashe B. Waxman, M.D., cardiologist at Toronto General Hospital for his helpful discussion on arrhythmia problems. Programming assistance from Ken Anderson, M.D. and Andrew Gullen is gratefully acknowledged as well as assistance from Brian Nixon in document preparation.

Financial assistance is provided by the Ontario Heart Foundation. During the course of the work, the author was a recipient of a scholarship from the Department for Foreign Affairs of Canada, and a Research Fellowship from Defense Research Establishment Atlantic.

References

[Blum82] R. L. Blum, in *Discovery and Rsprsssntation of Causal Relationships from a Largs Tims- Oriented Clinical Databass: Ths RX Project,* Springer-Verlag (1982). Lecture Notes in Medical Informatics.

[Hagan79] A. Hagan and others, "Evaluation of Computer Programs for Clinical Electrocardiography," in *Computer Tschniquss in Cardiology,* ed. D. Cady, Jr., Marcel Dekker, Inc., New York (1979).

[Horowitz75] S. L. Horowitz, "A Syntactic Algorithm for Peak Detection in Waveforms with Applications to Cardiography," *Communications of ths ACM* 18(5), pp.281-265, Association for Computing Machinery (May 1975).

[Kunz63] J. C. Kunz, "Analysis of Physiological Behavior using a Causal Model based on First Principles," *Procssdings of ths Third National Confsrsncs on Artificial Intslligsncs* 1, pp.225-228, American Association for Artificial Intelligence (August 1983).

[Long83] W. J. Long, "Reasoning about State from Causation and Time in a Medical Domain," *Procssdings of ths Third National Confsrsncs on Artificial Intelligence,* pp.251-254, American Association for Artificial Intelligence (August 1983).

[Minsky75]M. Minsky, "A Framework For Representing Knowledge," pp. 211-277 in *Ths Psychology of Computsr Vision,* ed. Patrick H. Winston, McGraw-Hill Book Company, New York (1975).

[Mostow78]D. J. Mostow and F. Hayes-Roth, "A Production System for Speech Understanding System," in *Pattsrn Directed Inference Systems,* ed. D. A. Waterman and F. Hayes-Roth, Academic Press, Inc. (1978).

[Mylopoulos83] J. Mylopouios, T. Shibahara, and J. K. Tsotsos, "Building Knowledge-Based Systems: The PSN Experience," *Computsr (IEEE)* 16(10), pp.83-89, Institute of Electrical and Electronics Engineers, Inc. (October 1983).

[Patil81] R. S. Patil, "Causal Representation of Patient Illness for Electrolyte and Acid-Base Diagnosis," Technical Report MIT/LCS/TR-2e7, Laboratory for Computer Science, Massachusetts Institute of Technology (1981). PhD Thesis.

[Pople82] H. E. Pople, Jr.. "Heuristic Methods for Imposing Structure on III-Structured Problems: The Structuring of Medical Diagnostics," pp. 119-190 in *Artificial Intelligence in Medicine,* Westview Press, Inc., 5500 Central Ave., Boulder, Colorado (1982). AAAS Selected Symposium Series 51.

[Rich81] C. Rich, "A Formal Representation for Plans in the Programmer's Apprentice," *Procssdings of the 7th International Joint Confsrsncs on Artificial Intelligence,* pp.1044-1052 (1981).

[Rieger76] C. Rieger and M. Grinberg, "The CausaJ Representation and Simulation of Physical Mechanics," TR-495. Univ. of Maryland (Nov. 1976).

[Shibahara83] T. Shibahara, J. K. Tsotsos, J. Mylopouios, and H. D. Covvey. "CAA: A Knowledge-Based System Using Causal Knowledge to Diagnose Cardiac Rhythm Disorders," *Proceedings of ths 8th International Joint Confsrsncs on Artificial Intslligsncs* (August 1983).

[Szolovits78] P. Szolovits and S. G. Pauker, "Categorical and Probabilistic Reasoning in Medical Diagnosis," *Artificial Intslligsncs* 11, pp. 115-144 (1978).

[Tsotsos80] J. K. Tsotsos, J. Mylopouios, H. D. Covvey, and S W. Zucker. "A Framework for Visual Motion Understanding." *IEEE Pattsrn Analysis and Machine Intslligsncs,* pp.563-573, Institute of Electrical and Electronics Engineers, Inc. (November 1980). Special Issue on Computer Analysis of Time-varying Imagery.

[Tsotsos85] J. K. Tsotsos, "Knowledge Organization and its Role in Representation and Interpretation for Time-varying Data: the ALVEN system," *Computational Intslligsncs* 1(1), pp. 16-32, National Research Council Canada (February 1986).

[Zadeh83] L. A. Zadeh, "Commonsense Knowledge Representation Based on Fuzzy Logic," *IEEE Computsr* 16(10), pp.61-65, IEEE (October 1983).