

Related Work

[Breese, 1992] – Construction of belief and decision networks

[Wellman, Breese, Goldman, 1992] – From knowledge bases to decision models

[Goldman & Charniak, 1993] – A language for construction of belief networks

[Poole, 1993] – Probabilistic Horn abduction and Bayesian networks

[Bacchus, 1993] – Using first-order probability logic for the construction of Bayesian networks

[Haddawy, 1994] – Generating Bayesian Networks from Probability Logic Knowledge Bases

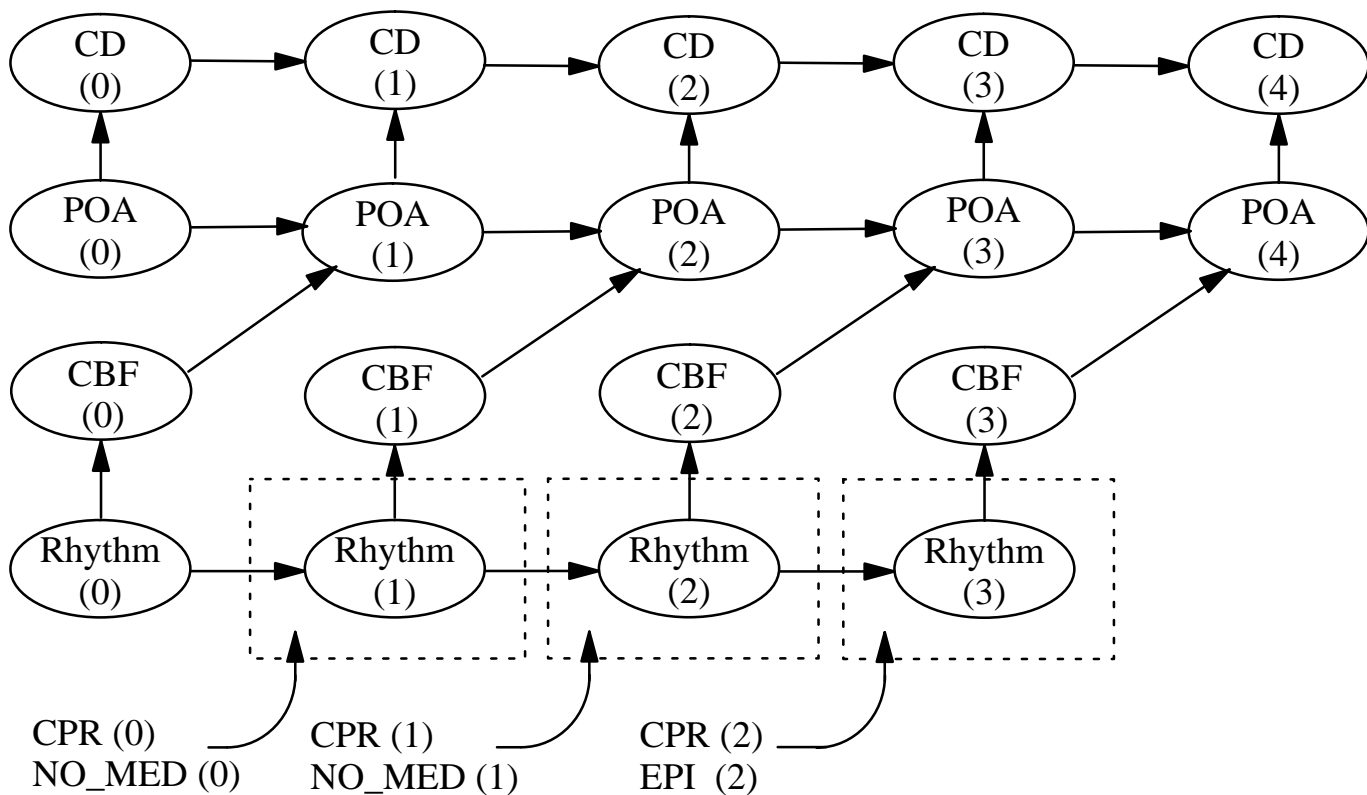
[Glesner & Koller, 1995] – Constructing Flexible Dynamic Belief Networks from First-Order Probabilistic Knowledge Bases

Drowning Victim

$E = \{\text{rhythm}(\text{john},0,a), \text{poa}(\text{john},0,5\text{min}), \text{cd}(\text{john},0,\text{none}), \text{cbf}(\text{john},0,\text{absent})\}$

$C = \{\text{CPR}(\text{john},0), \text{CPR}(\text{john},1), \text{CPR}(\text{john},2), \text{EPI}(\text{john},2)\}$

$Q = \text{cd}(\text{john},4,V)$



$P(\text{cd}(\text{john},4,\text{none})) = 0.84$

$P(\text{cd}(\text{john},4,\text{mild})) = 0.16$

$P(\text{cd}(\text{john},4,\text{moderate})) = 0.00$

$P(\text{cd}(\text{john},4,\text{severe})) = 0.00$

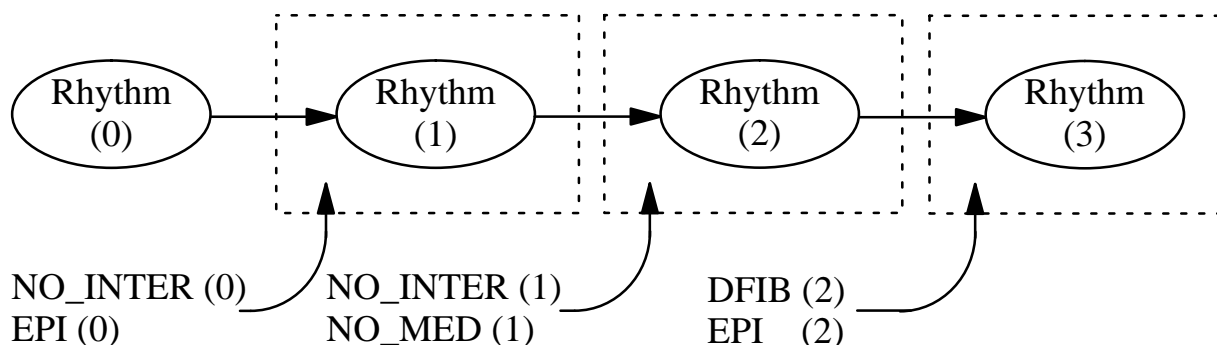
Medical Code Domain Examples

Heart Attack

$E = \{\text{rhythm}(\text{john},0,\text{vf}), \text{poa}(\text{john},0,\text{none}), \text{cd}(\text{john},0,\text{none}), \text{cbf}(\text{john},0, \text{present})\}$

$C = \{\text{EPI}(\text{john},0), \text{EPI}(\text{john},2), \text{DFIB}(\text{john},2)\}$

$Q = \text{rhythm}(\text{john},3,\text{V})$



$P(\text{rhythm}(\text{john},3,\text{nsr})) = 0.41$

$P(\text{rhythm}(\text{john},3,\text{vf})) = 0.09$

$P(\text{rhythm}(\text{john},3,\text{vt})) = 0.04$

$P(\text{rhythm}(\text{john},3,\text{af})) = 0.00$

$P(\text{rhythm}(\text{john},3,\text{svt})) = 0.01$

$P(\text{rhythm}(\text{john},3,\text{b})) = 0.00$

$P(\text{rhythm}(\text{john},3,\text{a})) = 0.44$

Soundness & Completeness

Theorem

Given a complete query $P(Q)=?$, where Q is at time t , $\iota \leq t \leq \tau$, a set of evidence E , a set of context information C , and a KB ,

if

- the COMBINE function always generates finite sets of sentences,
 - CRPB is (ι, τ) -bound completely quantified and (ι, τ) -bound consistent,
 - the proof procedure for $C \cup CB$ is sound and complete wrt any query generated by Q -procedure, and
 - Q -procedure stops after a finite amount of time,
- then* Q -procedure is sound and complete.

Theorem

In any framework $\langle P(Q)=?, E, C, KB \rangle$ where $P(Q)=?$ is a complete query and Q is at constant time t , $\iota \leq t \leq \tau$,

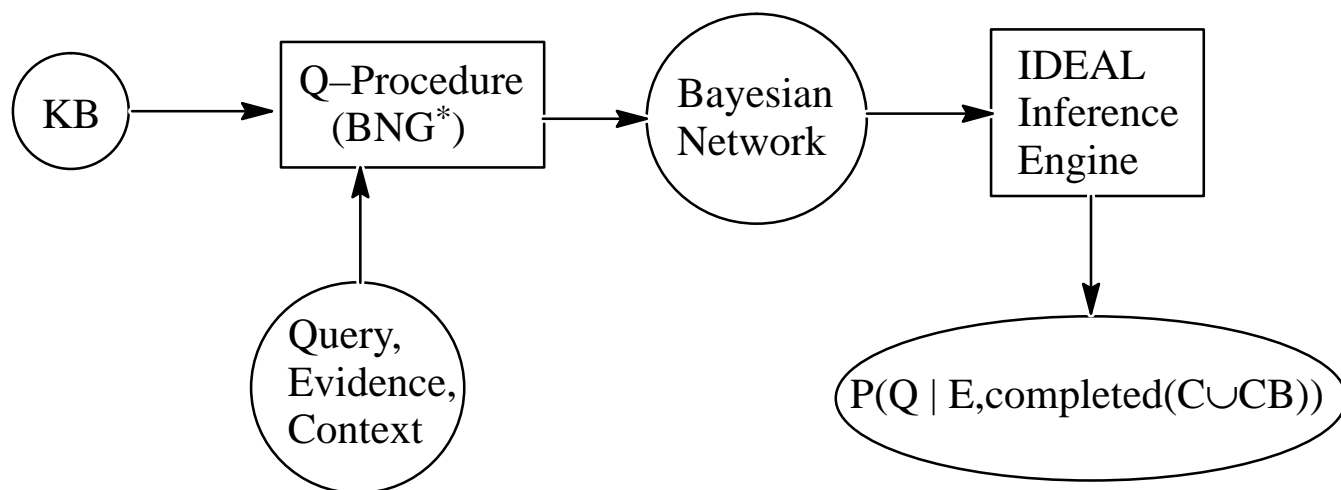
if

- PB and CB are both acyclic,
- $\langle E, C, KB \rangle$ is allowed,
- CRPB is (ι, τ) -bound completely quantified and (ι, τ) -bound consistent,

then Q -procedure is sound and complete.

Constructing the network by building supporting networks is justified by the fact that atoms that influence neither the evidence nor the query are irrelevant.

The supporting networks are constructed via calls to a BUILD-NET function. BUILD-NET receives as input an atom, whose supporting network needs to be explored. It augments the Bayesian network and returns a set of substitutions such that for each substitution there exists a supporting network for the ground instance of the atom corresponding to the substitution. BUILD-NET frequently calls SLDNF to answer queries on $C \cup CB$.



* Available via www at <http://www.cs.uwm.edu/faculty/haddawy>

Query Answering Procedure

A **complete query** is a query of the form $P(Q) = ?$, where the last argument of Q is a variable and the other arguments may also contain variables.

Q-Procedure

- 1) Build the necessary portion of the Bayesian network, each node of which corresponds to an $\text{obj}(A)$, where A is a ground p-atom in RAS.
- 2) Update the Bayesian network using the set of evidence E .
- 3) Output the updated beliefs of the query nodes.

Main idea: Build a supporting network for each random variable corresponding to a ground instance of an evidence atom or the query.

Let A be a ground p-atom and consider the set of all ground p-atoms B such that A is influenced by B in CRPB.

The **supporting network** for $\text{obj}(A)$ is a Bayesian network consisting of $\text{obj}(A)$ and the set of all $\text{obj}(B)$, with the relevant influenced-by relations represented as links or sequences of links.

Proposition

The set of (ι, τ) –models of CRPB corresponds to a partition of:
1) the set of all (t_0, t_1) –models of CRPB if $t_0 \leq \iota, \tau \leq t_1$, and
2) the set of all possible models of CRPB.

Induced Probability Distribution

A probability distribution which is (ι, τ) –bound induced by the set of evidence E , the set of context information C , and KB is a probability distribution on possible (ι, τ) –models of CRPB satisfying (ι, τ) –CRPB and the independence assumption implied by (ι, τ) –CRPB.

Consistent Completely Quantified CRPB

A completely quantified CRPB is consistent if

- 1) it contains no cycles, and
- 2) for all $P(A_0 \mid A_1, \dots, A_n) = \alpha$ in CRPB,
 $\sum \{ \alpha_i \mid P(A_0 \mid A_1, \dots, A_n) = \alpha_i \in \text{CRPB and } \text{obj}(A_0') = \text{obj}(A_0) \} = 1$

Theorem

If the (ι, τ) –RAS is finite and the (ι, τ) –CRPB is completely quantified and consistent, then there exists one and only one (ι, τ) –bound induced probability distribution.

Completely Quantified CRPB

A CRPB is completely quantified if

- 1) for all ground atoms A in RAS, there exists at least one sentence in CRPB with A as the consequent, and
- 2) for all ground sentences S in CRPB we have a specification of the probability of $\text{conse}(S)$ given all possible combinations of values of the atoms in $\text{ante}(S)$.

Probabilistic Independence Assumption

If $P(A_0 \mid A_1, \dots, A_n) = \alpha$ is in CRPB then for all ground p-atoms B which are not in $\text{Ext}(A_0)$ and not influenced by A_0 , A_0 and B are probabilistically independent given A_1, \dots, A_n .

Model Theory

We define the semantics by using possible worlds on the Herbrand base. Characterizing the semantics by canonical Herbrand models is widely used in work on logic programming.

Possible (ι, τ) -Model

Given a set of evidence E , a set of context information C , and a KB, a possible (ι, τ) -model M of the corresponding CRPB is a set of atoms in (ι, τ) -RAS such that for all A in (ι, τ) -RAS, $\text{Ext}(A) \cap M$ has exactly one element.

$E = \{\text{rhythm}(\text{john},0,\text{vf}), \text{poa}(\text{john},0,\text{none}), \text{cd}(\text{john},0,\text{none}), \text{cbf}(\text{john},0,\text{present})\}$

$C = \{\text{EPI}(\text{john},0), \text{EPI}(\text{john},2), \text{DFIB}(\text{john},2)\}$

$\iota = 0, \tau = 3$

RAS =

rhythm(john,0,nsr), rhythm(john,0,vf), rhythm(john,0,vt), ...
rhythm(john,1,nsr), rhythm(john,1,vf),rhythm(john,1,vt), ...
cbf(john,0,present), cbf(john,0,absent)
cbf(john,1,present), cbf(john,1,absent), ...
poa(john,0,none), poa(john,0,1min), poa(john,0,2min), ...
cd(john,0,none), cd(john,0,mild), cd(john,0,moderate), ...
...

RPB =

$P(\text{rhythm}(\text{john},0,\text{nsr})=0.001, P(\text{rhythm}(\text{john},0,\text{vf})=0.74, ...$
 $P(\text{poa}(\text{john},0,\text{none})=0.99, P(\text{poa}(\text{john},0,1\text{min})=0.005, ...$
 $P(\text{cd}(\text{john},0,\text{none})) = 0.99, P(\text{cd}(\text{john},0,\text{mild}) = 0.005, ...$
 $P(\text{cbf}(\text{john},0,\text{present}) = 0.99, P(\text{cbf}(\text{john},0,\text{absent})) = 0.01, ...$
...
 $P(\text{rhythm}(\text{john},1,\text{nsr})|\text{rhythm}(\text{john},0,\text{nsr}))=.05$
 $P(\text{rhythm}(\text{john},1,\text{nsr})|\text{rhythm}(\text{john},0,\text{vf}))=.01$
 $P(\text{rhythm}(\text{john},1,\text{nsr})|\text{rhythm}(\text{john},0,\text{vt}))=.01$
 $P(\text{rhythm}(\text{john},1,\text{nsr})|\text{rhythm}(\text{john},0,\text{af}))=.01$
...
 $P(\text{rhythm}(\text{john},1,\text{vf})|\text{rhythm}(\text{john},0,\text{nsr}))=.10$
...
 $P(\text{rhythm}(\text{john},2,\text{nsr})|\text{rhythm}(\text{john},1,\text{nsr}))=1.0$
 $P(\text{rhythm}(\text{john},2,\text{nsr})|\text{rhythm}(\text{john},1,\text{vf}))=.05$
...
 $P(\text{poa}(\text{john},1,3\text{min})|\text{cbf}(\text{john},0,\text{present}),\text{poa}(\text{john},0,2\text{min}))=0.0$
 $P(\text{poa}(\text{john},1,3\text{min})|\text{cbf}(\text{john},0,\text{absent}),\text{poa}(\text{john},0,2\text{min}))=1.0$
...
 $P(\text{cd}(\text{john},2,\text{mild})|\text{poa}(\text{john},2,3\text{min}),\text{cd}(\text{john},1,\text{mild}))=0.0$
 $P(\text{cd}(\text{john},2,\text{mild})|\text{poa}(\text{john},2,\text{sustained}),\text{cd}(\text{john},1,\text{mild}))=.98$
 $P(\text{cd}(\text{john},2,\text{moderate})|\text{poa}(\text{john},2,\text{sustained}),\text{cd}(\text{john},1,\text{mild}))=.02$

...

Relevant Probabilistic Sentences (RPB)

Given a set of evidence E , a set of context information C , and a KB , the set of relevant probabilistic sentences (RPB) is defined as the set of all $\text{prob}(S)$, where S is a ground probabilistic sentence such that $\text{completed}(C \cup CB) \models \text{context}(S)$, $\text{cons}(S) \in \text{RAS}$, and $\text{ante}(S) \subseteq \text{RAS}$.

Combined Relevant PB (CRPB)

The combined relevant PB (CRPB) is constructed by applying the corresponding combining rules to each maximally coherent set of sentences in RPB which have the same atom in the consequent.

(ι, τ) -Bounded RAS, RPB, CRPB

Given two integers ι, τ ($\iota \leq \tau$), a set of evidence E , a set of context information C , and a KB , the (ι, τ) -RAS is the set $\{A \mid A \in \text{RAS} \text{ and if } A \text{ is timed then it is timed at } t, \iota \leq t \leq \tau\}$. The (ι, τ) -RPB and (ι, τ) -CRPB are confined versions of RPB and CRPB on (ι, τ) -RAS.

We can condition the sentences in PB on completed($C \cup CB$) by simply eliminating those where $P(\text{context})=0$ and by eliminating the contexts from those where $P(\text{context}) = 1$.

$$P(\text{DFIB}(\text{john},1))=1$$

$$P(\text{ATRO}(\text{john},1))=1$$

$$P(\text{NO_INTER}(\text{john},1))=0$$

$$P(\text{rhythm}(X,t,\text{vf}) \mid \{\text{rhythm}(X,t-1,\text{af})\}, \{\text{DFIB}(X,t-1), \text{ATRO}(X,t-1)\})=.35$$

$$P(\text{rhythm}(X,t,\text{nsr}) \mid \{\text{rhythm}(X,t-1,\text{vt})\}, \{\text{NO_INTER}(X,t-1), \text{EPI}(X,t-1)\})=.01$$

$$P(\text{rhythm}(\text{john},2,\text{vf}) \mid \{\text{rhythm}(\text{john},1,\text{af})\}) = .35$$

Relevant P-Atoms (RAS)

Given a set of evidence E , a set of context information C , and a KB, the set of relevant p-atoms (RAS) is defined recursively:

- 1) $\text{ground}(E) \subseteq \text{RAS}$
- 2) if S is a ground instance of a probability sentence such that $\text{completed}(C \cup CB) \models \text{context}(S)$ and $\text{ante}(S) \subseteq \text{RAS}$, then $\text{conseq}(S) \in \text{RAS}$.
- 3) if a p-atom A is in RAS then $\text{Ext}(A) \subseteq \text{RAS}$.
- 4) RAS is the smallest set satisfying the above conditions.

This is similar to the construction of Herbrand least models for Horn programs.

Inference problem:

- Query – Q
- Set of evidence (p-atoms) – E
- Set of context information (c-atoms) – C
- Knowledge base – KB

Compute $P(Q|E)$ within context C.

Semantics

Let $S = P(A_0 \mid A_1, \dots, A_n) = \alpha \leftarrow C_1, \dots, C_n$

$\text{ante}(S) = A_1, \dots, A_n$

$\text{conse}(S) = A_0$

$\text{context}(S) = C_1, \dots, C_n$

Let $A = q(t_1, t_2, v)$

$\text{obj}(A) = (q, t_1, t_2)$

$\text{Ext}(A) = \{q(t_1, t_2, v_i), 1 \leq i \leq n\}$

PB: $P(\text{conse} \mid \text{ante}) = \alpha \leftarrow \text{context} \equiv$


$\forall x P(\text{conse} \mid \text{ante}, \text{context}) = \alpha$

CB:

Conceptually, the context information C is a set of observations which we elaborate with the context base CB.

Apply completion semantics to the Herbrand universe of $C \cup \text{CB}$ and take $\text{completed}(C \cup \text{CB})$ to hold with probability 1. So for every c-atom C_i either $P(C_i) = 0$ or $P(C_i) = 1$.

Now computing $P(Q|E)$ within context C amount to computing $P(Q \mid E, \text{completed}(C \cup \text{CB}))$.




$$P(\text{poa}(X,t,3\text{min})|\text{cbf}(X,t-1,\text{present}),\text{poa}(X,t,2\text{min}))=.0$$

$$P(\text{poa}(X,t,3\text{min})|\text{cbf}(X,t-1,\text{absent}),\text{poa}(X,t,1\text{min}))=.0$$

$$P(\text{poa}(X,t,3\text{min})|\text{cbf}(X,t-1,\text{absent}),\text{poa}(X,t,2\text{min}))=1.0$$

...



$$P(\text{cd}(X,t,\text{mild})|\text{poa}(X,t,3\text{min}),\text{cd}(X,t-1,\text{mild}))=.0$$

$$P(\text{cd}(X,t,\text{mild})|\text{poa}(X,t,\text{sustained}),\text{cd}(X,t-1,\text{severe}))=.0$$

$$P(\text{cd}(X,t,\text{mild})|\text{poa}(X,t,\text{sustained}),\text{cd}(X,t-1,\text{mild}))=.98$$

$$P(\text{cd}(X,t,\text{moderate})|\text{poa}(X,t,\text{sustained}),\text{cd}(X,t-1,\text{mild}))=.02$$

...

The rules in PB will typically not be a complete specification of a probability distribution over the random variables represented by the p-atoms.

The specification of the probability of a variable given combinations of values of two or more variables that influence it may not be given. For real applications this information can be difficult to obtain.

We associate a **combining rule** with each p-predicate.
e.g. generalized noisy-OR

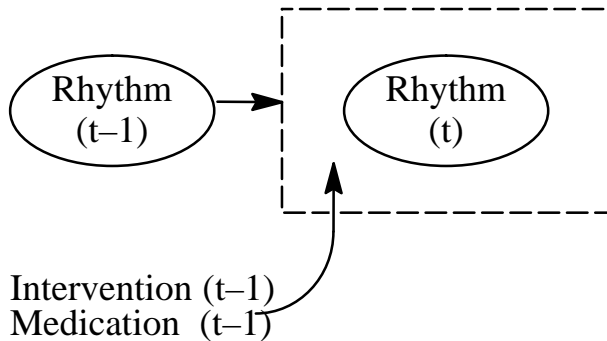
The relations among the p-predicates are defined in the **probabilistic base (PB)**:

$P(\text{rhythm}(X,0,\text{nsr})) = 0.001$, $P(\text{rhythm}(X,0,\text{vf})) = 0.74$, ...

$P(\text{poa}(X,0,\text{none})) = 0.99$, $P(\text{poa}(X,0,1\text{min})) = 0.005$, ...

$P(\text{cd}(X,0,\text{none})) = 0.99$, $P(\text{cd}(X,0,\text{mild})) = 0.005$, ...

$P(\text{cbf}(X,0,\text{present})) = 0.99$, $P(\text{cbf}(X,0,\text{absent})) = 0.01$



$P(\text{rhythm}(X,t,\text{nsr})|\text{rhythm}(X,t-1,\text{nsr})) = .05 \leftarrow \text{NO_INTER}(X,t-1), \text{EPI}(X,t-1)$

$P(\text{rhythm}(X,t,\text{nsr})|\text{rhythm}(X,t-1,\text{vf})) = .01 \leftarrow \text{NO_INTER}(X,t-1), \text{EPI}(X,t-1)$

$P(\text{rhythm}(X,t,\text{nsr})|\text{rhythm}(X,t-1,\text{vt})) = .01 \leftarrow \text{NO_INTER}(X,t-1), \text{EPI}(X,t-1)$


...

$P(\text{rhythm}(X,t,\text{vf})|\text{rhythm}(X,t-1,\text{af})) = .35 \leftarrow \text{DFIB}(X,t-1), \text{ATRO}(X,t-1)$

...

$P(\text{rhythm}(X,t,\text{a})|\text{rhythm}(X,t-1,\text{vf})) = .15 \leftarrow \text{NO_INTER}(X,t-1), \text{NO_MED}(X,t-1)$

...



$P(\text{cbf}(X,t,\text{present})|\text{rhythm}(X,t,\text{nsr})) = 1.0$

$P(\text{cbf}(X,t,\text{absent})|\text{rhythm}(X,t,\text{nsr})) = .0$

$P(\text{cbf}(X,t,\text{present})|\text{rhythm}(X,t,\text{vf})) = .0$

$P(\text{cbf}(X,t,\text{absent})|\text{rhythm}(X,t,\text{vf})) = 1.0$

...

The last attribute of a p–predicate represents the value of the corresponding random variable. The values must be mutually exclusive and exhaustive.

This is specified with **predicate declarations (PD)**:

rhythm(X: Person, t: time, V)

VAL(rhythm)={nsr, vf, vt, af, svt, b, a}

cbf(X: Person, t: time, V)

VAL(cbf)={present, absent}

poa(X: Person ,t: time, V)

VAL(poa)={none, 1min, 2min, 3min, 4min, 5min, sustained}

cd(X: Person, t: time, V)

VAL(cd)={none, mild, moderate, severe}

Representation Language

Two types of predicates:

- Context predicates – deterministic
- Probabilistic predicates – represent random variables

The relations among c-predicates are defined in the context base, an acyclic normal logic program.

C-predicates:

DFIB(X,t)
CPR(X,t)
NO_INTER(X,t)

LIDO(X,t)
ATRO(X,t)
EPI(X,t)
NO_MED(X,t)

CB:

$\text{NO_INTER}(X,t) \leftarrow \neg \text{DFIB}(X,t), \neg \text{CPR}(X,t)$
 $\text{NO_MED}(X,t) \leftarrow \neg \text{LIDO}(X,t), \neg \text{ATRO}(X,t), \neg \text{EPI}(X,t)$

Rhythms

- NSR – normal sinus rhythm
- VF – ventricular fibrillation
- VT – ventricular tachycardia
- AF – atrial fibrillation
- SVT – super ventricular tachycardia
- B – brady
- A – asytole

Interventions

- CPR – cardiopulmonary resuscitation
- DFIB – defibrillation

Medications

- LIDO – lidocane
Slows down heart; restores rhythm
- ATRO – atrophine
Speeds up heart
- EPI – epinephrine
Speeds up heart; increases blood pressure

Medical Code Domain

Biomedical Process

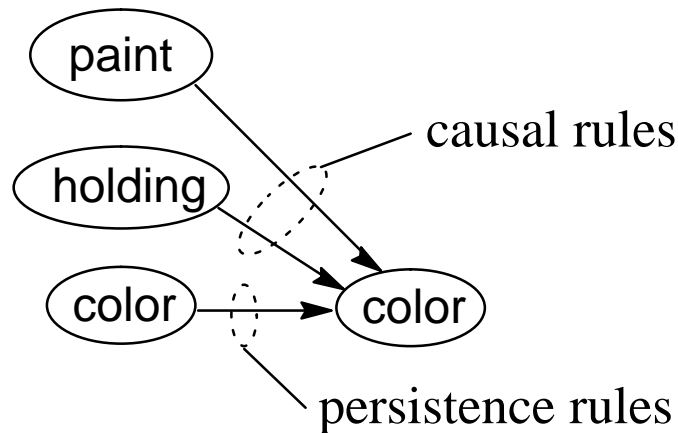
The cardiac arrest can present with several heart rhythms. The particular rhythm determines the amount of cerebral blood flow. The length of time without cerebral blood flow is known as the period of anoxia. Once the period of anoxia has exceeded five minutes, cerebral damage may begin to occur. The level of probable cerebral damage depends on the existing level of damage and the length of the period of anoxia.

The Projection Problem

Given a current observed cardiac rhythm, a set of planned medical interventions, a set of planned medications, infer the future rhythm and extent of cerebral damage.

Example: Plan Projection

Actions are typically represented as nodes in the network. This often results in networks with large numbers of nodes and large link matrices. For each node we need both causal rules and persistence rules.



But since when evaluating a plan, the performance of one's own actions is deterministic knowledge, actions can be used as context information.

$$(P(\text{painted}(x,t) \mid \text{holding}(x,t-1)) = .99) \leftarrow \text{paint}(x,t-1)$$

$$(P(\text{painted}(x,t) \mid \text{painted}(x,t-1)) = .95) \leftarrow \neg \text{paint}(x,t-1)$$

The link matrix for the network representing both action effect and persistence is divided between the two types of rules and only one type of rule is applicable at any time.

General Approach: Represent a class of temporal Bayesian networks with a KB of probabilistic rules augmented with context constraints.

Problems

- Inference remains NP–hard
- Cannot effectively do temporal reasoning with a static representation.

Solution

- Treat networks as sets of ground instances of Horn–clauses.

Knowledge–Based Model Construction

- Represent a class of networks with a set of schematic rules
- Construct networks dynamically

Two extreme approaches:

- Emphasis on practical model construction algorithms
[Breese, 1992], [Goldman & Charniak, 1993]
- Emphasis on formal semantics
[Poole, 1993], [Bacchus, 1993]

Temporal Probability Model Construction

Common approach: Represent time discretely and create an instance of each time–varying random variable for each point in time.

We can greatly reduce the size of the network models if we can identify some deterministic knowledge and use it as context to index the probabilistic information.

Introduction

Temporal Probabilistic Reasoning

- Diagnosis
- Prediction
- Planning

Bayesian Networks

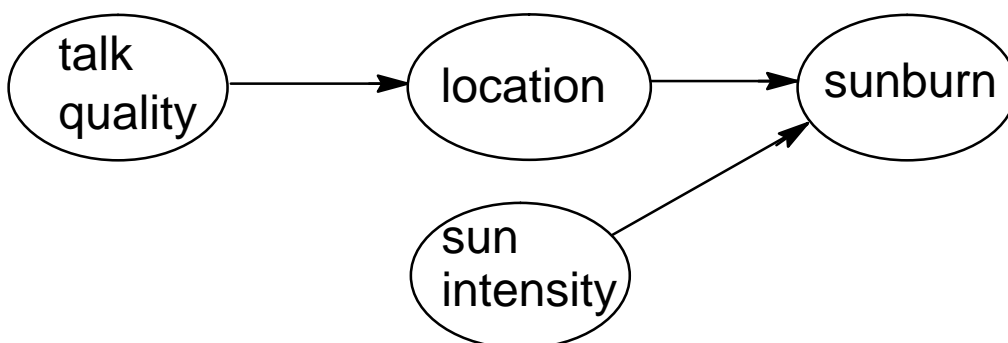
Graphical representation of a probability distribution:

A DAG in which nodes represent random variables and edges represent probabilistic dependencies.

With each node is associated a conditional probability or link matrix, that encodes the probability of the values of that node given the values of its immediate parents.

Probabilistic independencies are encoded in the network topology.

- Reduces complexity of model specification
- Facilitates model building
- Allows for efficient inference algorithms
- Facilitates explanation generation



Temporal Reasoning with Context–Sensitive Probability Logic

Peter Haddawy

Liem Ngo

Robert Krieger

James Helwig

*Dept. of Elect. Eng. & Computer Science
University of Wisconsin–Milwaukee*