

Fall Diagnosis using Dynamic Belief Networks

A. E. Nicholson

Department of Computer Science, Monash University,
Clayton, VIC 3168, Australia,
annn@cs.monash.edu.au.

Abstract. The task is to monitor walking patterns and give early warning of falls using foot switch and mercury trigger sensors. We describe a dynamic belief network model for fall diagnosis which, given evidence from sensor observations, outputs beliefs about the current walking status and makes predictions regarding future falls. The model represents possible sensor error and is parametrised to allow customisation to the individual being monitored.

1 Introduction

The task is to monitor the stepping patterns of elderly people, or recovering patients. Not only are actual falls to be detected causing an alarm to be raised, but irregular walking patterns, stumbles and near falls are to be monitored, and early warning of possible falls made in time for giving assistance. The monitoring is performed using two kinds of sensors: foot-switches which provide information about a foot step; and a mercury sensor which is triggered by a change in height such as going from standing upright to lying horizontal, and hence indicates a fall has occurred. Timing data for the observations is also given.

Previous work in this domain performed fall diagnosis with a simple state machine [3], however this does not allow representation of either degrees of belief as to the person's ambulatory status, or of the uncertainty in the sensor readings. *Dynamic belief networks* integrate a mechanism for inference under uncertainty with a secure Bayesian foundation, and are suitable for domains, such as the fall diagnosis problem, where the world changes and the focus is reasoning over time. In this paper we present a dynamic belief network model for the fall diagnosis problem, an interesting practical application of an AI approach to the real world problem of medical monitoring.

The organisation of this paper is as follows. The fall diagnosis problem is described in detail in Sect. 2. Sect. 3 gives an introduction to dynamic belief networks. In Sect. 4 we develop a complete belief network model for the fall diagnosis problem, with results given in Sect. 5. Extensions to the basic network are described in Sect. 6.

2 The Fall Diagnosis Problem

Davies [3] describes a project with Prof. Ian Brown at Monash University, Dept. of Electrical Engineering, for monitoring the stepping pattern of elderly people

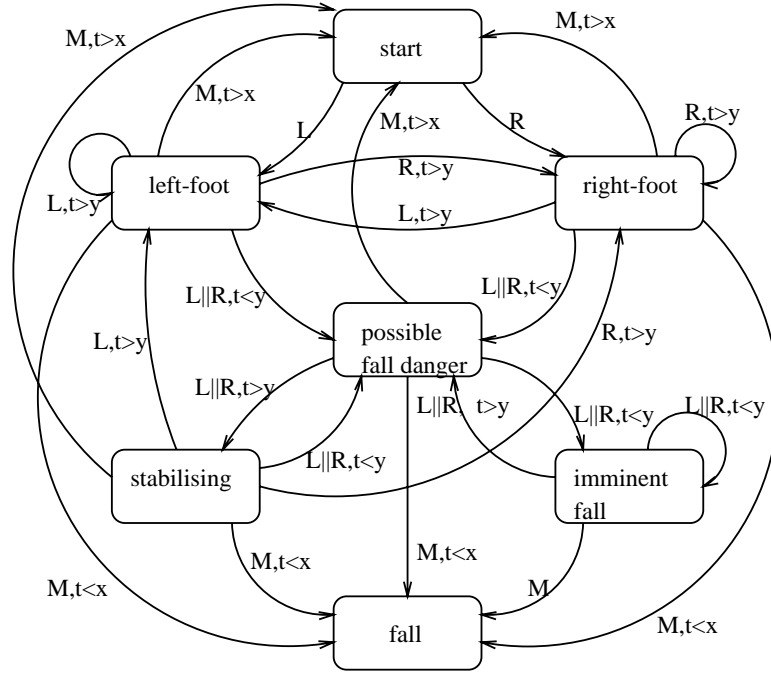


Fig. 1. Davies' State Machine for the fall diagnosis problem [3]

and patients. Step data is obtained using foot-switches and sent via a mobile data network to a remote monitoring station, which attempts to detect falls and near falls by using a state transition diagram, shown in Fig. 1. This model was developed by Davies in conjunction with expert medical practitioners.

The sensor observations are as follows: **L**: data from the left foot switch; **R**: data from the right foot switch; **M**: data from a mercury switch indicates a change in height. Each sensor observation is accompanied by a time, which is the time duration of the sensor observation. This timing information is crucial in performing fall diagnosis: y is the threshold time below which a foot switch reading is considered a stumble; x is the threshold time below which the mercury trigger is taken to indicate a fall (a slow change in height would be consistent with intended sitting or lying down).¹

The states of the state machine are as follows. **start** is a state of ignorance, entered when a slow mercury trigger is recorded, or when the machine is restarted after a fall alarm. **left-foot**, **right** are walking states, indicating which foot is currently forward in the process of walking. Normal stepping patterns, indicated by an observation time interval of $> y$, should see the state machine alternate between these two states. **possible-fall-danger** is an intermediate waiting state, entered after any abnormally fast step which may have been a

¹ Davies used the threshold times of $x = 2s$, $y = 0.8s$.

stumble, as the next input will determine if the patient is really stumbling or if the reading was a lone occurrence. The **stabilising** state will be reached after **possible-fall-danger** if a slow-controlled step is observed. The **imminent-fall** state will be reached after **possible-fall-danger** if another quick step is observed. The system currently increments a counter storing the number of near falls detected in the day. The **fall** state will be reached from **imminent-fall** with any triggering of the mercury switch, in which case the system sounds a local alarm and places an emergency call to the base station. This fall state is also reached from states other than **imminent-fall**, however in these cases the time data for the mercury switch must be $< x$ seconds.

This state machine model has a number of limitations. First, there is no representation of degrees of belief in the current state of the person's ambulation. Second there is no distinction between actual states of the world and observations of that state, for example, the fall state is really a **fall-alarm** state. That is, there is no explicit representation of the uncertainty in the sensor observation [8]. Possible sensor errors include: false positives, where the sensor wrongly indicates that an action (left, right, lowering action) has occurred (also called clutter, noise or false alarms); false negatives, where an action occurred but the sensor was not triggered and no observation was made (also called missed detection); wrong timing data. Also, there is no difference between a sequence of alternate foot steps, and a sequence of same foot steps (hopping).

3 Dynamic Belief Networks

Belief networks are directed acyclic graphs, where nodes correspond to random variables, which we assume to take discrete values (although in general they need not be discrete). In this paper the variables pertain to the world state or the sensor observations. The relationship between any set of state variables can be specified by a joint probability distribution. The nodes in the network are connected by directed arcs, which may be thought of as causal or influence links. The connections also specify the independence assumptions between nodes. Each node has associated with it a *probability distribution*, which, for each combination of the variables of the parent nodes (called a *conditioning case*), gives a probability of each value of the node variable. The probability distribution for a node with no predecessors is the prior distribution. Evidence can be specified about the state of *any* of the nodes in the network — root nodes, leaf nodes or intermediate nodes. This evidence is propagated through the network affecting the overall joint distribution (as represented by the conditional probabilities). There are a number of exact and approximate inference algorithms available for performing belief updating [11]; in this paper we are not concerned with the particular algorithm.

Belief networks have been used in various applications, such as medical diagnosis and model-based vision which initially were more static, i.e. essentially the nodes and links do not change over time. Such approaches involve determining the structure of the network; supplying the prior probabilities for root nodes

and conditional probabilities for other nodes; adding or retracting evidence about nodes; repeating the inference algorithm for each change in evidence.

More recently researchers have used belief networks in *dynamic* domains such as the fall diagnosis problem, where the world changes and the focus is reasoning over time [4, 6, 9]. Such dynamic applications include robot navigation and map learning based on *temporal* belief networks [4], monitoring robot vehicles [7], oil forecasting [2], [12], forecasting sleep apnea [1], automated vehicle control [5] and traffic plan recognition [13]. For such applications the network grows over time, as the state of each domain variable at different times is represented by a *series* of nodes. These dynamic networks are Markovian, which constrains the state space to some extent, however it is also crucial to limit the history being maintained in the network.

A generic dynamic belief network structure for monitoring application is shown in Fig. 2 [9]. The types of nodes are: **World** nodes, which describe the central domain variables (for example, position, heading, velocity) variables; **Event** nodes, which represent a change in the state of a world node; **Observation** nodes, which represent direct observations of world nodes, or the observable effects of an event. Time is discretised at irregular intervals, usually divided by the occurrence of discrete events. Each time slice within the network represents the static environment during that time interval. The structure within time slices is often regular. These networks are typically highly connected, particularly between adjacent time slices. The conditional probability distributions (CPDs) are shown in rectangular boxes. The CPDs of nodes with parents in the previous time slice are usually a function of the time interval. After addition of sensor observations as evidence to the DBN (indicated by dark shading), belief updating is performed, providing prediction for the values of the world nodes at time slice $T + 1$.

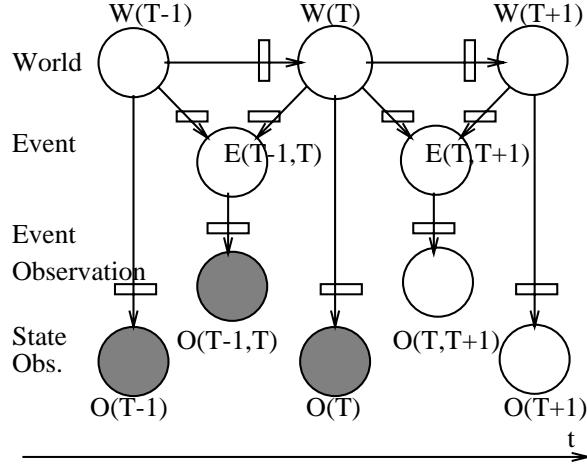


Fig. 2. Generic structure for a dynamic belief network

4 Basic DBN Model for Fall Diagnosis

Davies' state machine essentially defines the fall diagnosis problem as a set of *if-then-else* rules. When developing the DBN model, a key difference is that we focus on the causal relationships between domain variables, and make a clear distinction between observations and actual states of the world. A DBN for the fall diagnosis problem is given in Fig. 3. In the rest of this section, we describe the various features of this network in such a way as provide an insight into the network development process.




Fig. 3. Dynamic Belief Network for fall diagnosis problem including interslice arcs. The smaller nodes with thicker outlines are the sensor observations.

Nodes

When considering how to represent a person's walking situation, possibilities include whether the person is stationary on both feet, on a step with either the left or right foot forward, or has fallen and hence is no longer on their feet. We call the main world node representing this **F**, which may take 4 possible values: **both**, **left**, **right**, **off**. The event node **Fall**, which is a boolean, indicates whether a fall has taken place between time slices.

Fall warning and detection relies on an assessment of the person's walking pattern. The node **S** maintains the person's status, and may take the possible values **ok** and **stumbling**.

The action variable, **A**, may take the possible values **left**, **right**, or **none**. The last value is necessary for the situation where a time slice is added because the mercury sensor has triggered (i.e. the person has fallen) but no step was taken, or a foot switch false positive observation was registered.²

² Note that we can easily extend the model to handle a person jumping, by adding an additional possible value for **A**, say **jump**.

There is an observation node for each of the two sensors. The foot switch observations are essentially observations on the step actions, and are represented by the **AO** node which contains the same values as the action node. The mercury sensor trigger is represented by the node **M**, which represents a boolean variable.

The time between sensor observations is given by the node **T**. Given the problems with combining continuous and discrete variables [14, p.465] and the limitations of the sensor, node **T** has discrete values representing tenths of seconds. In order to represent the uncertainty in the sensor data, we say it can take values within an interval around the sensor time reading that generates the addition of a new time slice to the DBN. If there is some knowledge of the patients expected walking speed, values in this range can be added also. The time observation node, **T0**, has the same state space as **T**.

There is a new copy of each node added for each time slice; we will indicate the time slice by the subscript. The possibility of adding more time slices is shown by the dashed arcs to the right.

Note that there is no need to explicitly include **imminent-fall** or **fall** in the status node. The belief of a fall in the current time slice i is given by the posterior obtaining after adding evidence and running the inference algorithm, that is, $\text{bel}(\text{Fall}_i = \text{T})$,³ and a warning about an imminent fall can be based on the predictions for the next time slice, that is whether $\text{bel}(\text{Fall}_{i+1})$ is greater than some warning threshold.

Structure and Conditional Probability Distributions

The CPDs for the nodes **A**, **F**, **Fall** and **S** are given in Table 1. The model for walking is represented by the arcs from \mathbf{F}_i to \mathbf{A}_i , and \mathbf{F}_i , \mathbf{A}_i and \mathbf{S}_i to \mathbf{F}_{i+1} .

We assume that normal walking involves alternating left and right steps. Where the left and right are symmetric, only one combination is included in the table. We have priors for starting on both feet (r) or already being off the ground (s). Because we have restricted the possible actions to moving either feet or neither, there is no way for this model to reflect a person getting to their feet; we are assuming use of the model will begin with the person upright, and if not, they stay off their feet. Looking at the CPD for \mathbf{F}_{i+1} , we can see that a left step can have the walker finish on one foot or both feet, depending on whether it is a half or full step. By definition, if a person finishes on a particular foot, it rules out some actions; for example, if $\mathbf{F}_{i+1} = \text{left}$, the action could not have been **right**. These zero conditional probability are omitted from the table.

The CPD for \mathbf{F}_{i+1} for the conditioning cases where $\mathbf{S}_{i+1} = \text{stumbling}$ is exactly the same as for **ok** except the p and q probability parameters will have lower values, representing the higher expectation of a fall; that is, $p'_i < p_i$, $q'_i < q_i$, for all relevant i .

If there are any variations on walking patterns for an individual patient, for example if one leg was injured, the DBN can be customised by varying the

³ Also given by $\text{bel}(\mathbf{F}_i = \text{off})$; the redundancy is useful for describing the problem, but could be removed to improve computational efficiency.

Table 1. CPDs for step action node **A**, the foot node **F**, the **Fall** node and the walking status node **S**

$P(F_0=\text{left} \text{right})$	$= (1-r-s)/2$	
$P(F_0=\text{both})$	$= r$	
$P(F_0=\text{off})$	$= s$	
$P(A=\text{left} F=\text{right})$	$= u$	alternate feet
$P(A=\text{right} F=\text{right})$	$= v$	hopping
$P(A=\text{none} F=\text{right})$	$= 1-u-v$	stationary
$P(A=\{\text{left} \text{right}\} F=\text{both})$	$= w/2$	start with left or right
$P(A=\text{none} F=\text{both})$	$= 1-w$	stationary
$P(A=\text{none} F=\text{off})$	$= 1$	can't walk when off feet
$P(F_{i+1}=\text{left} F_i=\text{right}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= p_1$	succ. alternate step
$P(F_{i+1}=\text{both} F_i=\text{right}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= q_1$	half-step
$P(F_{i+1}=\text{off} F_i=\text{right}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= 1-p_1-q_1$	fall prob
$P(F_{i+1}=\text{left} F_i=\text{left}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= p_2$	succ. hop
$P(F_{i+1}=\text{both} F_i=\text{left}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= q_2$	half-hop
$P(F_{i+1}=\text{off} F_i=\text{left}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= 1-p_2-q_2$	fall prob
$P(F_{i+1}=\text{left} F_i=\text{both}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= p_3$	succ. first step
$P(F_{i+1}=\text{both} F_i=\text{both}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= q_3$	unsucc. first step
$P(F_{i+1}=\text{off} F_i=\text{both}, A_i=\text{left}, S_{i+1}=\text{ok})$	$= 1-p_3-q_3$	fall prob
$P(F_{i+1}=\text{left} F_i=\text{left}, A_i=\text{none}, S_{i+1}=\text{ok})$	$= p_4$	
$P(F_{i+1}=\text{off} F_i=\text{left}, A_i=\text{none}, S_{i+1}=\text{ok})$	$= 1-p_4$	fall when on left foot
$P(F_{i+1}=\text{right} F_i=\text{right}, A_i=\text{none}, S_{i+1}=\text{ok})$	$= p_5$	
$P(F_{i+1}=\text{off} F_i=\text{right}, A_i=\text{none}, S_{i+1}=\text{ok})$	$= 1-p_5$	fall when on right foot
$P(F_{i+1}=\text{both} F_i=\text{both}, A_i=\text{none}, S_{i+1}=\text{ok})$	$= p_6$	
$P(F_{i+1}=\text{off} F_i=\text{both}, A_i=\text{none}, S_{i+1}=\text{ok})$	$= 1-p_6$	fall when on both feet
$P(F_{i+1}=\text{off} F_i=\text{off}, A_i=\text{left}, S_{i+1}=\text{any})$	$= 1$	no "get up" action
$P(\text{Fall}=\text{T} \mid F_{i+1}=\text{off}, F_i=\{\text{left} \text{right} \text{both}\})$	$= 1$	from upright to ground
$P(\text{Fall}=\text{F} \mid F_{i+1}=\text{any}, F_i=\text{off})$	$= 1$	can't fall if on ground
$P(S_{i+1}=\text{ok} T_i=t)$	$= 1$	if $t \geq y$
$P(S_{i+1}=\text{stumbling} T_i=t)$	$= 1$	if $t < y$

probability parameters, s , r , p_i , q_i , u , v and w , and removing the assumption that left and right are completely symmetric. For example, we can relax the assumption that the person is equally likely to start on the left foot as the right. Note that having different p parameters indicates different expectations of a fall when the person is walking compared to hopping. Also, a person can end up off their feet even if the status node **S** is indicating **ok**.

The fall event node **Fall** has **F_i** and **F_{i+1}** as predecessors; a fall only occurs when the subject was on his or her feet to start with (**F_i** \neq **off**), and finishes off

their feet ($F_{i+1} = \text{off}$).⁴

The value of walking status node S is determined solely by the time between sensor readings (see next section for an extension which takes into account status history). In this DBN model, the T node has no predecessors. One possible model is to have uniform priors, or the prior can also be modified, based on sensor observations over time, to reflect an individual's ordinary walking speed.

When constructing the conditional probability distributions for the various observation nodes, the confidence in the observation is given by some value based on a model of the sensor's performance and is empirically obtainable; pos is the sensitivity of the positive sensor data, neg is the specificity of the negative sensor data (or, $1-neg$ is the probability of ghost data). We make the default assumption that missing or wrong data are equally likely — this need not be the case and can be replaced by any alternative plausible values.

Each observation node has a single predecessor: the mercury trigger observation node M has predecessor F ; the foot-step action observation node AO has predecessor A ; the time observation node TO has predecessor T . The conditional probability distributions for M , AO and TO are shown in Table 2. Note that the CPD for the case where the sensor is defective is uniform over the other time values; this could easily be changed to cluster around the true time interval. If the timing sensor fails and no data is obtained, fall diagnosis becomes impossible, so we do not handle the case of missing time data.

Note that when the monitoring begins, we do not need to have a known start state; we need only have a prior over the possible starting positions. Because the standard left foot, right foot, walking model, is represented by the conditional probability distribution between F_i and F_{i+1} , if the first data received is a **left** S_i , then after belief updating, the belief vector will include $\text{bel}(F_i=\text{off}) = 0$, $\text{bel}(F_i=\text{left}) < 0.25$, $\text{bel}(F_i=\text{left}) > 0.25$ and $\text{bel}(F_i=\text{left}) > 0.25$. The DBN presented is one possible model for the fall diagnosis problem; many other variations are possible. For example, the DBN does not handle the case where both foot switches provide data at the same time. See [10] for a comparison of the use of action nodes in this model with other monitoring and planning applications.

5 Results

The results described in this section were obtained using the Lisp-based IDEAL belief network development environment [15] on a GNU Common Lisp platform. We present results of a Fall Diagnosis network modelled for a given set of parameters: $s = 0.0$, $r = 0.9$, $u = 0.7$, $v = 0.2$, $w = 0.1$, $p_1 = 0.6$, $q_1 = 0.3$, $p'_1 =$

⁴ We do not model the situation Davies described where the mercury trigger data is ignored if the time is $\geq x$; this would be more correctly modelled by: adding an additional value, **sitting** to the state F ; adding an additional value, **sit**, to the action A ; adding another alternative, **sat**, to the fall event **fall**; adding a connection from T_i to A_i ; changing the CPD for A_i to say that if the time is above the threshold, then the **sit** action is possible.

Table 2. CPDs for observation nodes **M** (mercury trigger), **A0** (foot switch), **T0** (time data)

$P(\mathbf{M}=\mathbf{T} \mathbf{Fall}=\mathbf{T})$	$= pos_1$	ok
$P(\mathbf{M}=\mathbf{F} \mathbf{Fall}=\mathbf{T})$	$= 1-pos_1$	missing
$P(\mathbf{M}=\mathbf{F} \mathbf{Fall}=\mathbf{F})$	$= neg_1$	ok
$P(\mathbf{M}=\mathbf{T} \mathbf{Fall}=\mathbf{F})$	$= 1-neg_1$	false alarm
$P(\mathbf{A0}=\mathbf{left} \mathbf{A}=\mathbf{left})$	$= pos_2$	ok
$P(\mathbf{A0}=\mathbf{right} \mathbf{A}=\mathbf{right})$	$= pos_2$	ok
$P(\mathbf{A0}=\mathbf{right} \mathbf{A}=\mathbf{left})$	$= (1-pos_2)/2$	wrong
$P(\mathbf{A0}=\mathbf{left} \mathbf{A}=\mathbf{right})$	$= (1-pos_2)/2$	wrong
$P(\mathbf{A0}=\mathbf{none} \mathbf{A}=\mathbf{left})$	$= (1-pos_2)/2$	missing
$P(\mathbf{A0}=\mathbf{none} \mathbf{A}=\mathbf{right})$	$= (1-pos_2)/2$	missing
$P(\mathbf{A0}=\mathbf{none} \mathbf{A}=\mathbf{none})$	$= neg_2$	ok
$P(\mathbf{A0}=\mathbf{left} \mathbf{A}=\mathbf{none})$	$= (1-neg_2)/2$	false alarm
$P(\mathbf{A0}=\mathbf{right} \mathbf{A}=\mathbf{none})$	$= (1-neg_2)/2$	false alarm
$P(\mathbf{T0}=\mathbf{x} \mathbf{T}=\mathbf{x})$	$= pos_3$	ok, $y \neq x$, T and T0 have m values.
$P(\mathbf{T0}=\mathbf{y} \mathbf{T}=\mathbf{x})$	$= 1 - pos_3/m-1$	ok, $y \neq x$, T and T0 have m values.

0.5, $q'_1 = 0.4$, $p_2 = 0.6$, $q_2 = 0.3$, $p'_2 = 0.5$, $q'_2 = 0.4$, $p_3 = 0.6$, $q_3 = 0.3$, $p'_3 = 0.5$, $q'_3 = 0.4$, $p_4 = 0.95$, $p'_4 = 0.85$, $p_5 = 0.95$, $p'_5 = 0.85$, $p_6 = 0.9$, $p'_6 = 0.8$, $pos_1 = 0.9$, $pos_2 = 0.9$, $pos_3 = 0.9$, $neg_1 = 0.95$, $neg_2 = 0.95$. The **T** and **T0** time nodes had 4 possible values, \mathbf{t}_1 , \mathbf{t}_2 , \mathbf{t}_3 , and \mathbf{t}_4 ; the lowest, \mathbf{t}_1 was below the threshold **y** and meant the subject was considered to be stumbling.

After constructing the DBN, we entered a sequence of evidence, that is simulated observations from the sensors, and performed belief updating after every new piece of evidence was added. Table 3 shows the posterior probabilities, or beliefs, of the values of nodes in the network across this sequence of data. For reasons of space, we left out the initial **S0** node and the **T2** and **T02** nodes from the model, and do not give all the beliefs, especially if they are uniform or otherwise obvious. Probabilities have been rounded to 4 decimal places. Evidence added results in a 1.0 belief for that value, shown in bold in the table; also bolded are the beliefs described below in the text. The evidence sequence added, and the effect on the beliefs, was as follows.

No evidence added: All beliefs are based on the parameters. Belief in an immediate fall is small, $\text{bel}(\mathbf{Fall}_0 = \mathbf{T}) = 0.1194$, but chance of being off feet in 2 steps is higher, $\text{bel}(\mathbf{F}_0 = \mathbf{T}) = 0.2238$.

T0 set to \mathbf{t}_1 : This increases the probability that the person is stumbling, that is, $\text{bel}(\mathbf{S}_1 = \mathbf{stumbling}) = 0.9$, which in turn slightly increases the belief in a fall, $\text{bel}(\mathbf{Fall}_0 = \mathbf{T}) = 0.1828$.

A0 set to left: Foot switch information leads to a change in the belief in the initial starting state; $\text{bel}(\mathbf{F}_0 = \mathbf{right})$ has increased from 0.05 to 0.2550, reflecting the model of alternate foot steps.

- M₀ set to false:** The negative mercury trigger data makes it very unlikely that a fall occurred, $\text{bel}(\mathbf{Fall}_0=\mathbf{T})=0.0203$.
- T₀ set to t₂:** “Resetting” of the original timing data makes it less likely the person was stumbling, reducing the belief in a fall, $\text{bel}(\mathbf{Fall}_0=\mathbf{T}) = 0.0098$.
- M₀ set to true:** However, resetting the mercury trigger data makes a fall most probable, $\text{bel}(\mathbf{Fall}_0=\mathbf{T})=0.6285$, although there is still the chance that the sensor has given a wrong reading.
- M₁ set to false, T₀ set to t₄, A₀ set to none:** No action, and no mercury trigger data confirms the earlier fall, $\text{bel}(\mathbf{Fall}_0=\mathbf{T})=0.7903$, since if the person is already on the ground they won’t take a left or right step.

6 Extensions to the DBN

The states **imminent-fall**, **possible-fall**, and **stabilising** in the original state machine are an attempt to capture the idea that the history beyond the current time interval gives information about the likelihood of a fall soon. This is represented in a DBN by the use of a history node [9], which maintains a count of how long the agent has been exhibiting one type of behaviour. For our domain, this would be a status history node, \mathbf{H}_i , for each time slice; its predecessors are the previous and current walking status nodes, \mathbf{S}_{i-1} and \mathbf{S}_i . \mathbf{H} then becomes a predecessor of \mathbf{F}_{i+1} , and the CPD entries are changed so that the probability of falling is a function of the stumble count.

We can also improve the model of what a person’s ordinary walking pace is by adding an arc from \mathbf{T}_i to \mathbf{T}_{i+1} , which would allow a representation of the expectation that the walking pace should remain fairly constant.

The DBN described in the previous section provides a mechanism for handling (by implicitly rejecting) certain inconsistent data. It represents adequately the underlying assumptions about the data uncertainty, however it does not provide an explanation of *why* the observed sensor data might be incorrect. We can represent the most usual source of incorrect data, namely a defective sensor, by the addition of a sensor status node \mathbf{SS} [8] for each sensor. Each sensor status node becomes a predecessor of the corresponding observation node, and there is a connection between sensor status nodes across time slices. See [10] for details of the extensions described in this section.

7 Conclusions

We have shown the development of a dynamic belief network model for fall diagnosis which overcomes the limitations of early work. Given evidence from sensor observations, the model outputs beliefs about the current walking status and makes predictions regarding future falls. The model represents possible sensor error, and is parametrised to allow customisation to the individual being monitored.

Table 3. Changing beliefs as evidence is added or changed.

Node	Val	None	T0=t ₁	A0=left	M ₀ =F	T0=t ₂	M ₀ =T	SET
T ₀	t ₁	0.25	0.9000	0.9000	0.8914	0.0305	0.0535	0.0616
	t ₂	0.25	0.0333	0.0333	0.0361	0.9026	0.8812	0.8736
	t ₃	0.25	0.0333	0.0333	0.0361	0.0334	0.0326	0.0323
	t ₄	0.25	0.0333	0.0333	0.0361	0.0334	0.0326	0.0323
T0 ₀	t ₁	0.25	1.0	1.0	1.0	0.0	0.0	0.0
	t ₂	0.25	0.0	0.0	0.0	1.0	1.0	1.0
F ₀	left	0.05	0.05	0.0870	0.0860	0.0856	0.0964	0.0911
	right	0.05	0.05	0.2550	0.2717	0.2515	0.2792	0.2767
	both	0.90	0.90	0.6581	0.6422	0.6628	0.6244	0.6322
	off	0.0	0.0	0.0	0.0	0.0	0.0	0.0
A ₀	left	0.09	0.09	0.6403	0.6483	0.6453	0.6047	0.5427
	right	0.09	0.09	0.0356	0.0360	0.0359	0.0336	0.0302
	none	0.82	0.82	0.3241	0.3156	0.3188	0.3617	0.4271
A0 ₀	left	0.1265	0.1265	1.0	1.0	1.0	1.0	1.0
	right	0.1265	0.1265	0.0	0.0	0.0	0.0	0.0
	none	0.7470	0.7470	0.0	0.0	0.0	0.0	0.0
Fall ₀	True	0.1194	0.1828	0.1645	0.0203	0.0098	0.6285	0.7903
	False	0.8806	0.8173	0.8355	0.9797	0.9902	0.3715	0.2096
M ₀	True	0.1515	0.2053	0.1898	0.0	0.0	1.0	1.0
	False	0.8485	0.7947	0.8102	1.0	1.0	0.0	0.0
S ₁	ok	0.75	0.1	0.1	0.1086	0.9695	0.9465	0.9383
	stum'g	0.25	0.9	0.9	0.8914	0.0305	0.0535	0.0617
F ₁	left	0.0638	0.0425	0.2737	0.3208	0.5120	0.1921	0.0340
	right	0.0638	0.0425	0.0168	0.0197	0.0303	0.0114	0.0020
	both	0.7530	0.7322	0.5451	0.6391	0.4478	0.1680	0.1736
	off	0.1194	0.1828	0.1645	0.0203	0.0098	0.6285	0.7903
T ₁	t ₁	0.25	0.25	0.25	0.25	0.25	0.25	0.0326
	t ₄	0.25	0.25	0.25	0.25	0.25	0.25	0.9006
T0 ₁	t ₄	0.25	0.25	0.25	0.25	0.25	0.25	1.0
A ₁	left	0.0950	0.0749	0.0938	0.1099	0.1461	0.0548	0.0035
	right	0.0950	0.0749	0.2222	0.2605	0.3869	0.1451	0.0092
	none	0.8090	0.8502	0.6841	0.6296	0.4670	0.8001	0.9872
A0 ₁	left	0.1308	0.1137	0.1297	0.1434	0.1741	0.0966	0.0
	right	0.1308	0.1137	0.2389	0.2714	0.3788	0.1734	0.0
	none	0.7383	0.7730	0.6315	0.5851	0.4671	0.7301	1.0
Fall ₁	True	0.1044	0.0975	0.0959	0.1124	0.1099	0.0412	0.0024
	False	0.8956	0.9025	0.9041	0.8876	0.8901	0.9588	0.9976
M ₁	True	0.1387	0.1329	0.1315	0.1455	0.1434	0.0850	0.0
	False	0.8612	0.8671	0.8685	0.8545	0.8566	0.9150	1.0
S ₂	ok	0.75	0.75	0.75	0.75	0.75	0.75	0.9673
	stum'g	0.25	0.25	0.25	0.25	0.25	0.25	0.0327
F ₂	left	0.0673	0.0531	0.0898	0.1053	0.1472	0.0552	0.0258
	right	0.0673	0.0531	0.1335	0.1565	0.2291	0.08594	0.0076
	both	0.6415	0.6136	0.5164	0.6055	0.5040	0.1891	0.1740
	off	0.2238	0.2802	0.2603	0.1327	0.1197	0.6698	0.7927

Acknowledgements We wish to thank Dr. Jon Oliver for suggesting this fall diagnosis problem as an application for dynamic belief networks.

References

1. P. Dagum and A. Galper. Forecasting sleep apnea with dynamic network models. In *Proceedings of the Ninth Conference on Uncertainty in AI*, pages 64–71, 1993.
2. P. Dagum, A. Galper, and E. Horvitz. Dynamic network models for forecasting. In *Proceedings of the 8th Conference on Uncertainty in Artificial Intelligence*, pages 41–48, 1992.
3. James Davies. Fall diagnosis with a mobile data network. Unpublished BCSE Honours Report, Dept. of Electrical Engineering, Monash University, 1995.
4. Thomas Dean and Michael P. Wellman. *Planning and control*. Morgan Kaufman Publishers, San Mateo, Ca., 1991.
5. Jeff Forbes, Tim Huang, Keiji Kanazawa, and Stuart Russell. The BATMobile: Towards a bayesian automated taxi. In *Proceedings of the 14th Int. Joint Conf. on Artificial Intelligence (IJCAI'95)*, pages 1878–1885, 1995.
6. U. Kjærulff. A computational scheme for reasoning in dynamic probabilistic networks. In *Proceedings of the 8th Conference on Uncertainty in Artificial Intelligence*, pages 121–129, 1992.
7. A. E. Nicholson and J. M. Brady. The data association problem when monitoring robot vehicles using dynamic belief networks. In *Proc. of the 10th European Conf. on Artificial Intelligence (ECAI-92)*, pages 689–693, 1992.
8. A. E. Nicholson and J. M. Brady. Sensor validation using dynamic belief networks. In *Proceedings of the 8th Conference on Uncertainty in Artificial Intelligence*, pages 207–214, 1992.
9. A. E. Nicholson and J. M. Brady. Dynamic belief networks for discrete monitoring. *IEEE Systems, Man and Cybernetics*, 24(11), 1994.
10. A.E. Nicholson. A case study in dynamic belief networks: monitoring walking, fall prediction and detection. Technical Report Technical Report 96/251, Department of Computer Science, Monash University, 1996.
11. Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, Ca., 1988.
12. Gregory M. Provan. Tradeoffs in constructing and evaluating temporal influence diagrams. In *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence*, pages 40–47, 1993.
13. David Pynadeth and Michael P. Wellman. Accounting for context in plan recognition, with application too traffic monitoring. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, pages 472–481, 1995.
14. Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, 1994.
15. Sampath Srinivas and Jack Breese. Ideal: Influence diagram evaluation and analysis in lisp. Technical Report Technical Memorandum No. 23, Rockwell International Science Center, 1989.