A Simple Causal Model for Glucose Metabolism in Diabetes Patients

Kinzang Chhogyal[1], Abhaya Nayak[2], Rolf Schwitter[2], and Abdul Sattar[3]

[1] IIS, Griffith University, Brisbane, Australia
kinchhogyal@griffithuni.edu.au
[2] Macquarie University, Sydney, Australia
{abhaya.aayak | rolf.schwitter}@mq.edu.au
[3] IIS, Griffith University, Brisbane, Australia
NICTA Queensland Research Lab, Brisbane, Australia
a.sattar@griffith.edu.au

Abstract. Employing the theory of belief change, we study the implementation of a simple causal model that can capture how the blood sugar level changes in a diabetes patient. For this purpose we use distance measures between worlds as the underlying mathematical foundation. Using a simple example in the medical domain we investigate how an agent with initially incomplete and/or incorrect knowledge can iteratively develop a simple causal model by interacting with an oracle that represents the complete and correct model of a diabetic patient.

Keywords: Belief Revision, Belief Update, Causal Models, Glucose Metabolism

1 Introduction

The omnipresence of medical devices and the interconnectivity of the information age has overwhelmed the biomedical field with vast amounts of data. Researchers in the field of medical artificial intelligence are already harnessing the powers of traditional statistical and more novel data mining and (statistical) machine learning techniques to deal with this abundance of data. Nonetheless, the appropriateness of these approaches is still being debated. For instance, as Patel et al [9] argue:

> [such techniques only tend to uncover] relatively simple relationships in data and have not yet demonstrated the ability to discover the complex causal chains of relationships that underlie our human understanding from molecular biology to the complex multi-organism and environmental factor in the epidemiology of diseases such as malaria. Human expertise developed over centuries of experience and experimentation cannot be discarded in the hope that it will all be re-discovered (more accurately) by analyzing data'.

The primary objective of this paper is to develop computable causal models of domains of interest.

The earliest approaches to AI in medicine were used for the diagnosis of patients either by using *flow charts* or by using Bayesian models [12]. Flow charts had no place for any of the conceptual structures associated with medical reasoning. Bayesian models, though very successful, are highly sensitive to the data from which the prior probabilities in question are derived. Even expert systems like MYCIN are not quite satisfactory since, as [12] pointed out,

> The programs knowledge was mostly based on associations between diseases and their observable consequences, and lacked any explicit pathophysiologic model in terms of which those consequences could be grouped and explained. An adequate medical explanation, however, often demands that the associations which may have suggested the right diagnosis then be backed up by a consistent account of how the patients condition could have arisen from the suspected etiologies.

Thus we may conclude that causal modelling of physiology is an essential component of representing health information. A simple model with a few relevant parameters and states of clinical interest can provide more valuable information than a complex model with hundreds of parameters. Such models may be used both for diagnostic and therapeutic reasoning[12]. If the model is used for diagnosis, the observable parameters can be used as input to the model and the model predicts outcomes depending on the input.

In this paper, we attempt to understand how a simple causal model of the physiology, in particular the glucose metabolism, of a diabetes patient can be built incrementally by an agent using belief revision and update which have a well recognised theoretical foundation [1–3]. Although belief revision and update can be accomplished in several ways, our main approach is to use distance measure between sets of possible worlds to arrive at a stable causal model. We are interested in the use of distance measures to build causal models and the analysis of their effects on the revision and update process.

To gauge the effectiveness of the distance measures on the evolving model, we consider a simple scenario, namely that of a diabetic patient who may either be alert or non-responsive depending on her blood sugar-level. The agent starts with a preconceived model of the system and uses probing actions to elicit an output from the system. Available to us are two actions namely *administer insulin* and *administer glucose* that change the blood sugar level of her system. This discrepancy between the prediction and observation, if any, is used to successively revise the agent's model. It is hoped that the discrepancy between the predicted and observed outcome will eventually become stable. By evaluating the difference in the observation and expected output, the agent incrementally modifies its causal model so that after a number of iterations the model becomes stable.

2 Belief Change

2.1 Revision and Update

Revision In the AGM model [1, 3], epistemic states of an agent are represented by a belief set, which is a set of sentences or beliefs from a given language that is closed under classical logical consequence operation. In light of new information, a belief set may need to be modified. These modifications are generally classified as being *expansion* (the addition of new sentences to the belief set), a *contraction* (the removal of beliefs from the belief set) or a *revision* (incorporation of some information inconsistent with a belief set while maintaining consistency) .

In the case of contraction, a sentence must be removed along with other sentences that may logically entail it in the belief set. Sentences may collectively entail it and a decision must be made as to which other sentences should be removed. Similarly, with revision, if the new sentence to be added is inconsistent with the belief set, some sentences may first need to be removed in order to maintain consistency after adding the new sentence and this again presents us with a dilemma alike to that of contraction. Given this connection, it has been shown that contraction may be defined in terms of revision using the Harper Identity and revision may be defined in terms of contraction using the Levi Identity [3].

A guiding principle when devising a revision/contraction operation is to conform to the criterion of information economy, i.e., to retain as much of the old beliefs as possible. It is also vital that changes to the belief state be rational and this is guided by a set of rationality postulates for the given operation. Given a belief set K and a proposition α, a contraction function prescribes a method for choosing which sentences to delete from K so that α is not a logical consequence of the contracted belief set K. The largest subset of K that does not entail α and satisfies the two criteria above is called a *maximal subset* of K. In general, such maximal subsets do not purge sufficient information and exhibit undesirable behaviour [1, 3]. A way out of this problem is to use a method is called the *partial meet contraction*. This requires an ordering of the maximal subsets so that the selection function may select the best subsets. Though this and other methods describe general ways of constructing contraction functions, determining the content of the maximal subsets can be computationally costly.

An alternative method to forming contraction and revision functions is based on the notion of epistemic entrenchment [3, 4]. Some sentences may be believed to be more important than others in particular settings and hence are said to be more epistemically entrenched. When trying to decide between two sentences one of which should be given up during contraction, the less epistemically entrenched of the two is chosen to be discarded.

In our case, we are interested in viewing revision semantically as in Grove's account of system of spheres [5] where the beliefs are sentences in propositional logic and there are a set of possible worlds $[K]$ in which all the sentences in the the belief set K are true. If $[A]$ represents the possible worlds of a sentence A, then A is accepted in K only if $[K] \subseteq [A]$. If A is consistent with K, then $[A] \bigcap [K] \neq \{\}$ and if A is rejected in K (i.e. $\neg A$ is accepted in K), $[A]$ and $[K]$ are disjoint. In Grove's system of spheres $[K]$ constitutes the central innermost sphere and is surrounded by larger and larger concentric spheres S_i that are totally ordered and are analogous to the epistemic entrenchment of sentences. $[K], [A]$, each sphere S_i are all subsets of M which is the set of all possible worlds.

The revision of K by a sentence A is represented by the intersection of $[A]$ with the smallest sphere S_A (including $[K]$) that intersects $[A]$. This set, $[A] \bigcap S_A$ represents the set of closest elements in M in which A is true. Since there is a direct correspondence between a set of possible worlds and a belief set, the system of spheres may be used as a method for implementing revision functions.

Update In [7] it is argued that a new piece of information may be learned by an agent either when the world is static or when it is dynamic, and that the revision operation will not suffice when modifying the belief set in the latter case. The required operation in the second scenario is called an update and can be understood as follows: As seen in the system of spheres, if a belief set K is to modified by a sentence A, revision methods select from the models, i.e., set of possible worlds, for A that are closest to the set of models of K. In other words, given an ordering of the relationship between each and every model of A and K, an element of $[A]$ that is closest to $[K]$ is selected. On the other hand, while performing an update, for each element in $[K]$, it is assumed that there is a system of spheres centred on each world in [K], and the closest element in $[A]$ is selected and the union of all such models represents the updated belief set.

2.2 Motivation for distance measure

The ordering relationship between models over a language (or sentences in a belief set) may be defined by a distance. The way this distance is defined will affect the outcome of the resulting belief set after revision (or update). Revision and update are typically not one-step processes. There is a succession of these operations and therefore it is vital that the same operation must be applied during each iteration. In the system of spheres, each revised belief set is represented by a new system of spheres which is in general different from the preceding one. Similarly for epistemic entrenchment, for every revised belief set new epistemic entrenchment relations must be defined. In both cases, the number of spheres or epistemic entrenchment relations are exponential to the number of models. Distance measure uses only a polynomial number of distances in the number of models considered and furthermore it is coherent because the same revision/update functions are used during each operation [8, 11].

3 Causal Models

3.1 Causality

Causality in essence may be understood as the study of the relationship between two events, the first of which is the cause and the second, its consequence or effect. A preliminary analysis of causality shows that the two events have a temporal relationship where the cause is preceded by the effect. However, the effect may not be an immediate consequent of the cause but rather there may be subsequent unrelated events in between which makes it difficult to identify the real cause. Effects can also be associated with more than one cause and different combinations of these causes can in turn be seen as necessary or sufficient. Identifying and categorising causes presents another challenge. The study of causality is complex. Nevertheless given the importance of causal inference in learning about the world and in decision making, simple causal models can be used to improve one's understanding of the world.

Causal models may be idealised in the following way: initially a model of the system under consideration is posited and a hypothesis is formulated from which inferences are made about the expected outcomes that may result when the environment changes. The outcomes are compared with actual observations that are made and the model and the hypothesis are rectified so that the discrepancies between the expected and observed outcome is reduced. Rectifications can also occur in static domains when there is a transfer of knowledge from sources that can communicate with the learner. This is analogous to belief update and belief revision, and it is easy to see that both of these operations are necessary for developing causal models.

Belief sets along with the operations of revision and update provide a solid base that can be used to build causal models. Bayesian approaches to developing causal models are very common; they are probabilistic and face problems when encountering inconsistent observations. On the other hand belief revision and belief update with their ability to handle inconsistent information allow the development of a non-probabilistic account of causality. As distance measures serve as a fundamental tool for performing both revision and update, their choice and their influence on the correctness of rectifying causal models is of much interest.

3.2 Motivating Example

We consider a simple medical scenario of a diabetic patient whose blood sugar level can be *low*, *normal* or *high*, and she may be either *alert* or *not alert*. Accordingly, there are six possible states (worlds or models) denoted *S1...S6* as listed in the table below.

Table 1. States of the system

	Blood Sugar Level		
Patient Status ↓	*low*	*normal*	*high*
Alert	S1	S2	S3
Not Alert	S4	S5	S6

There is also an agent whose task is to develop a causal model which represents its knowledge of the system. By system, in this case, we mean the blood sugar level of the patient and how it is affected by various actions which we describe later. However, we assume the agent doesn't have access to any glucose measuring device and hence cannot observe the patient's blood sugar level. It can however *observe* whether the patient is *alert* or *not alert*. There are two actions that the agent can use to experiment with the system - *administer insulin* which has the direct effect of lowering the patient's blood sugar level (from *high* to *normal*, *normal* to *low*, and *low* to *low*), and we can also *administer glucose* which increases the blood sugar level (from *low* to *normal*, *normal* to *high*, and *high* to *high*). The causal model of the patient is called the *black box* since it is assumed the agent does not have direct access to the relevant causal mechanism that drives the patient's behaviour. The current knowledge of the agent is represented by its causal model that we call the *white box* since the agent has full access to this mechanism.

It is also assumed that the black box is a deterministic system and there is a measurable distance between the different states of the black box. For the sake of example, let Figure 1 represent the causal mechanism at work in terms of the real distance between the states of the black box. The distance between any two states is the sum of the segment lengths on the shortest path between those two states. For example, the distance between states $S5$ and $S3$ is $8 + 4 + 2 = 14$. State $S5$ which is highlighted represents the current state of the system where the patient's blood sugar level is *normal* but she is not alert. Now if the agent administers glucose to the patient, the immediate effect of this action is to increase the patient's sugar level from *normal* to *high*. There are two states, $S3$ and $S6$ in which the sugar level is high and of the two, the former is closer to the current state than the latter and hence the patient will be in state $S3$.

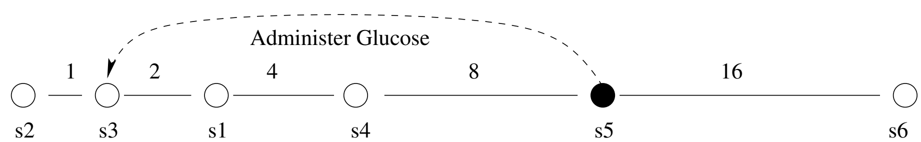

Fig. 1. The Black Box (Oracle) representing the system under observation

The agent meanwhile does not have access to the causal mechanism of the black box. Its causal knowledge evolves based on the observable outcomes of different actions performed on the black box. Figure 2 represents the evolving causal model of the agent. Whereas the system is really in state $S5$ as shown in Figure 1, the agent believes the real state is either state $S4$ or $S6$, consistent with its observation that the patient is not alert. Nevertheless, the agent makes a prediction based on its current causal model. The agent has administered glucose into the patient and therefore the agent makes a prediction by reasoning as follows:

1. The system is either in state $S4$ or $S6$.
   a) Consider the first case, $S4$ (*low* sugar level, *not alert*). If the patient is administered glucose, then it would move to one of the state $S2$ and $S5$

where the sugar level is *normal*. Since $S5$ is colser that $S2$ to $S4$, the new state will be $S5$.
   b) In the second case, $S6$, the new state be one in the set $\{S3, S6\}$ that is closer to $S6$, namely itself.
   c) The new state will be therefore be either $S5$ or $S6$.
2. In neither $S5$ or $S6$ is the patient *alert* so the agent predicts administering glucose will not result in any observable difference in the condition of the patient.

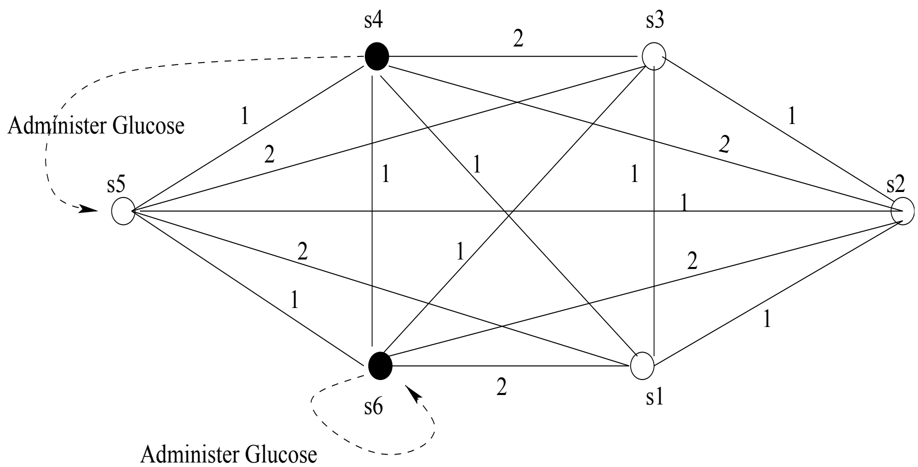

Fig. 2. The White Box representing the agent's knowledge

Steps 1 and 2 above comprise an update operation. Since the agent believes the patient is not alert which contradicts with the observation of the patient as actually being alert, it must perform a revision operation by taking into consideration that the patient may be in one of the three states in the set $\{S1, S2, S3\}$. The revision process goes as follows:

1. The agent believes the patient is in $S5$ or $S6$ and the observation requires that the system is either in one of the states in the set $\{S1, S2, S3\}$.
   a) From $S5$, it is closest to $S2$ with a distance of 1.
   b) From $S6$ it is closest to $S3$ also with a distance of 1.
   c) Since there is no unique state with a minimum distance, both $S2$ and $S3$ are what the agent now believes to be potential current states of the real system.

By thus iteratively performing an action on the patient, followed by an update and a revision operation, the agent tries to rectify its causal knowledge of the system.

4 Implementation

Since we are primarily interested in the use of distance measures by an agent to model this causal system, for illustrative purpose, we use the Dalal distance which is the Hamming distance between worlds [2]. We restrict ourselves to the use of propositional logic with a finitary language. The distance between states in Figure 2 above is indeed calculated using the Dalal distance. For instance, the Dalal distance between states $S5$*(normal, not alert)* and $S3$*(high, alert)* is 2 because these states differ in two propositional variables. A snapshot of a part our implementation in Java of the diabetes example is shown below in Figure 3.

Fig. 3. Snapshot of a part of the interface for the diabetes patient causal model application. The topmost figure shows the initial states of the black box ($S5$) and white box ($S4, S6$). The middle figure shows the transition states after *administering glucose*: ($S3$) and ($S5, S6$). The bottom most diagram shows the transition states after the agent performs a revision operation.

4.1 Results and Discussion

Since the number of states and the number of actions is small, the white box stabilises after an average of 5-6 iterations of the learning process when the actions are chosen randomly. In most instances, both the black box and white box arrive in either of states $\{S1, S2, S3\}$ where upon both models stabilise and under any action the transition states are identical.

Stability need not necessarily mean that the agent now has both a complete and correct knowledge of the system. Instead, the black box could be as in the case above, stuck in a cycle. If we take the black box to be a directed graph with actions as the arcs, this means that there is no path from any state in $\{S1, S2, S3\}$ to any state outside this set and by virtue of the distance measure the white box predicts the same outcome in these states.

The agent's choice of action may also give the impression that its model has stabilised. For example, if the black box is in $S4$ and the white box is in $\{S4, S5\}$, when the agent *administers glucose* the black box moves to $S5$ where the patient is *not alert*. The agent also updates its knowledge and believes that it is in either $\{S5, S6\}$. Since the patient is also *not alert* in both $\{S5, S6\}$, revision will not result in any noticeable change. *Administering insulin* will result in both the black and white box moving to their former states and again revision has no effect. In such circumstances, alternating the actions results in repeatedly identical results giving the false notion of stability. It is also worth noting that whenever the agent and system are both in state $S5$, upon *administering glucose* the agent will always believe it is in $S6$ where as the system will actually be in state $S3$ leading to a discrepancy between the prediction and observation.

5 Conclusion and Future Work

In this paper we presented a simple non-probabilistic causal inference model of the glucose metabolism in a diabetes patient based on belief revision and update. An action performed by an agent trying to model the causal system is followed by updates of its knowledge. Comparison of the predicted behaviour of the system and observed outcome leads to further rectification in its model of the system. The distance measure which is the underlying mechanism for improving the model itself may need to be corrected so that predicted and observed outcomes are identical.

Our experiment is a preliminary investigation into the use of a rudimentary distance measure for building causal models and the scenario we considered is simple. Nevertheless, it can be seen that distance measures can help the agent to reduce the discrepancy between its predicted and the system's actual outcome. Presently we are investigating other kinds of distance measures in order to be able to compare different approaches, with the aim of choosing the best among them. In our experiment, we only considered one observable variable namely whether the patient is *alert* or *not alert*. Presumably this limited ability to observe the system behaviour tends to get very quickly into a cycle in the process of modifying the causal model. We intend to study the effect of enhancing the agent's ability to observe. Similarly, we would also have to consider systems with more that just two actions that we have used. This would mean that it is important for the agent to adopt policies that can help choose actions judiciously so it can arrive at a stable causal model in a shorter time.

In our approach, the final belief state that results after revision may not be a necessary consequence under the given action because of incorrect distance measures. In an interesting but different approach [6] Hunter and Delgrande propose the use of action history trajectories to revise prior beliefs that are identified as the cause of the erroneous revised belief states given that the actions are infallible. It will be interesting to incorporate and exploit the advantages offered by this method in our context. It will also be interesting to explore these issues, keeping Sattar and Goebel's account of theory choice [10] in mind, while representing belief sets as Horn clause theories under appropriate integrity constraints. We also believe that the software developed in the process can be used as a pedagogical tool to help students and health care workers develop and rectify mental models of a given domain of interest.

References

1. Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
2. Mukesh Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 1988.
3. Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge Massachusetts, 1988.
4. Peter Gärdenfors. The dynamics of belief systems: Foundations vs. coherence theories. *Revue Internationale de Philosophie*, 44:24–46, 1990.
5. Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
6. Aaron Hunter and James P. Delgrande. Iterated belief change due to actions and observations. *J. Artif. Intell. Res. (JAIR)*, 40:269–304, 2011.
7. Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In Peter Gärdenfors, editor, *Belief Revision*, pages 183–203. Cambridge University Press, 1992.
8. Daniel J. Lehmann, Menachem Magidor, and Karl Schlechta. Distance semantics for belief revision. *J. Symb. Log.*, 66(1):295–317, 2001.
9. Vimla L. Patel, Edward H. Shortliffe, Mario Stefanelli, Peter Szolovits, Michael R. Berthold, Riccardo Bellazzi, and Ameen Abu-Hanna. The coming of age of artificial intelligence in medicine. *Artificial Intelligence in Medicine*, 46(1):5–17, 2009.
10. Abdul Sattar and Randy Goebel. Using crucial literals to select better theories. *Computational Intelligence*, 7:11–22, 1991.
11. Karl Schlechta, Daniel J. Lehmann, and Menachem Magidor. Distance semantics for belief revision. In Yoav Shoham, editor, *TARK*, pages 137–145. Morgan Kaufmann, 1996.
12. Peter Szolovits. How can medical computer programs reason? In *Proceedings of CARDIOSTIM*, 1988.