Australasian Medical Journal [AMJ 2012, 5, 9, 497-502]

A causal model for fluctuating sugar levels in diabetes patients

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

1. IIIS, Griffith University, Brisbane, Australia

2. Macquarie University, Sydney, Australia

RESEARCH

Please cite this paper as: Chhogyal K, Nayak A, Schwitter R, Sattar A. A causal model for fluctuating sugar levels in diabetes patients. AMJ 2012, 5, 9, 497-502. http//dx.doi.org/10.4066/AMJ.2012.1392.

Corresponding Author:
Kinzang Chhogyal
IIIS, Griffith University, 117 Kessels Road,
Nathan, QLD 4111
Email: kin.chhogyal@griffithuni.edu.au

Key Words
Belief Change, Belief Update, Belief Revision, Causal Models, Glucose Metabolism, Diabetes

Abstract

Background
Causal models of physiological systems can be immensely useful in medicine as they may be used for both diagnostic and therapeutic reasoning.

Aims
In this paper we investigate how an agent may use the theory of belief change to rectify simple causal models of changing blood sugar levels in diabetes patients.

Method
We employ the semantic approach to belief change together with a popular measure of distance called Dalal distance between different state descriptions in order to implement a simple application that simulates the effectiveness of the proposed method in helping an agent rectify a simple causal model.

Results
Our simulation results show that distance-based belief change can help in improving the agent's causal knowledge. However, under the current implementation there is no guarantee that the agent will learn the complete model and the agent may at times get stuck in local optima.

Conclusion
Distance-based belief change can help in refining simple causal models such as the example in this paper. Future work will include larger state-action spaces, better distance measures and strategies for choosing actions.

What this study adds:

1. This study explores the use of belief change to rectify causal models.
2. It employs a simple non-probabilistic causal inference system for modelling blood sugar level in a diabetes patient.
3. It provides a novel and theoretically well-founded account of knowledge evolution with potential implication for health care.

Background

It has been noted, for instance by Patel et al,[1] that most AI techniques tend to uncover only simple relationships in data, and that their efficacy at discovering complex causal chains of relationships that underlie our understanding of domains such as molecular biology or epidemiology is yet to be demonstrated. As they point out, "human expertise developed over centuries of experience and experimentation cannot be discarded in the hope that it will all be rediscovered (more accurately) by analysing data."[1]

Models of physiology have been extensively employed in expert systems. 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.[2] 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.

Method

We attempt to understand how a simple causal model of glucose metabolism in a diabetes patient can be built incrementally by an agent using belief revision and update which have a well recognised theoretical foundation.[3-5] Although belief revision and update can be modelled in

several ways, we employ a distance measure between sets of possible worlds to build causal models and analyse their effects on the revision and update process.

Belief change

In the AGM model,[3-5] an epistemic state of an agent is represented by a belief set, which is a set of sentences in a given language, closed under classical logical consequence operation, representing the beliefs of the agent. In light of a new piece of information, a belief set may need to be modified. These modifications are generally classified as being an expansion (addition of new sentences to the belief set), a contraction (removal of old beliefs from the belief set) or a revision (incorporation of some information inconsistent with the current belief set while maintaining consistency).

In the case of contraction, a sentence must be removed from a belief set, along with other sentences that logically entail it. Since a number of sentences may collectively entail the sentence being contracted, a decision must be made as to which other sentences should be removed as well. Similarly, in 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 before adding the new sentence, and this again presents us with a choice problem as in case of contraction. Given this connection, it has been shown that contraction can be defined in terms of revision using the Harper Identity,[4] and revision in terms of contraction using the Levi Identity.[4]

A guiding principle to follow when devising a revision/contraction operation is to conform to the criterion of information economy, i.e., to retain as much of the old information as possible. It is also vital that changes to the belief state are 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 no longer a logical consequence of the contracted belief set K-α. A subset of K that does not entail α and strictly satisfies the two criteria above is a maximal subset of K that does not entail α. In general, such maximal subsets exhibit undesirable behaviour.[3,4] A way out of this problem is to use a method called the partial meet contraction. This requires an ordering over such maximal subsets so that the best such subsets can be selected for this purpose. Though this and other methods describe general ways of constructing contraction functions, determining the content of the maximal subsets in question can be computationally costly.[6]

An alternative method to constructing contraction and revision functions is based on the notion of epistemic entrenchment.[4,5] Some sentences may be believed to be more important than others, 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,[7] which is taken to represent the plausibility ordering over different worlds, and is inter-translatable with epistemic entrenchment. The central sphere in this system, denoted [K], consists of the most plausible worlds, and represents the models of the current beliefs K.

An agent may learn of a new piece of information in an environment that may be static or dynamic. The revision operation will not suffice when modifying the belief set in the latter case.[8] The required operation in such a scenario is called an update operation. It can be understood as follows. If a belief set K is to be modified by a sentence A, revision methods select from [A], the models of A, those that are closest to the set of models of K. In other words, given the plausibility ordering over all the worlds, every 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 it, and the closest element in [A] is selected with respect to each such ordering, and the union of all such models represents the updated belief set.

The plausibility ordering can be defined by the presumed distance between different worlds. 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 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 a new epistemic entrenchment relations must be defined. In both cases, the number of spheres or epistemic entrenchment relations is 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.[9,10]

To gauge the effectiveness of belief change using 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. The discrepancy between the prediction and observation, if any, is used to successively revise the agent's model. By evaluating the difference in the observation and expected output, the agent incrementally modifies its causal model of the system so that after a number of iterations the model becomes stable.

Motivating example

Let us assume that the blood sugar level of our diabetic patient can be *low, normal* or *high*, and the patient 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

Patient Status	Blood Sugar Level		
	Low	*Normal*	*High*
Alert	S1	S2	S3
Not Alert	S4	S5	S6

There is also an agent whose task it is to develop a causal model (represented by its knowledge of the system) as to how the blood sugar level of the patient is affected by different actions. However, we assume the agent does not have access to any glucose-measuring device and hence cannot observe the patient's blood sugar level directly. It can however observe whether the patient is *alert* or *not alert*. The 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 *administer glucose*, which increases the blood sugar level (from *low* to *normal*, *normal* to *high*, and *high* to *high)*. It is important to note that there are two causal models. The causal model that really drives the patient's sugar level is called the black box since it is assumed the agent does not have direct access to it. The causal model it has constructed represents the current knowledge of the agent; we call the *white box* since the agent has full access to this mechanism. We note that even though in this case the agent is experimenting directly on the patient, in practice it may be based on the clinical records of the patient's response to such actions.

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.

Figure 1: The black box representing the system under observation

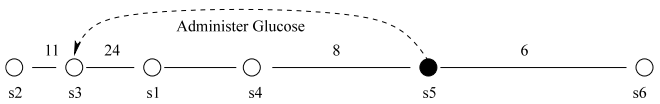

The distance between any two states is the sum of the segment lengths on the shortest path between those two states. For example, the distance between states S5 and S3 is 8 + 4 + 2 = 14. State S5 (highlighted) represents the current state of the system where the patient's blood sugar level is *normal* but the patient 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.

Figure 2: The white box representing the agent's knowledge

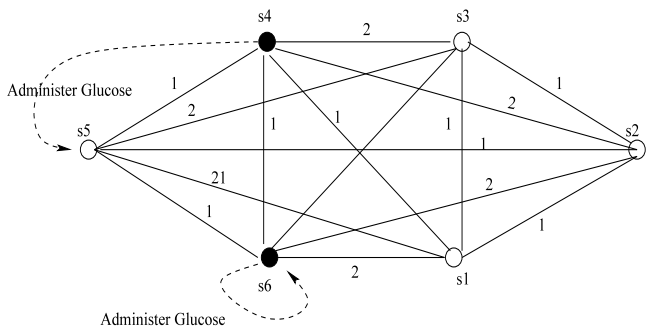

The agent meanwhile does not have access to the causal mechanism of the black box. Its causal knowledge evolves based on the observable outcomes of different actions performed on the black box. Figure 2 represents the evolving causal model of the agent. Whereas the system is really in state S5 as shown 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 states S2 and S5 where the sugar level is *normal*. Since S5 is closer than S2 to S4, the new state will be S5.
   b. In the second case for S6, the new state will be the one in the set {S3, S6} that is closer to S6, namely itself.
   c. The new state will therefore be either S5 or S6.
2. In neither S5 nor S6 is the patient *alert;* so the agent predicts administering glucose will not result in any observable difference in the condition of the patient.

Steps 1 and 2 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:

The agent believes the patient is in S5 or S6 and the observation requires that the system be 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.

Implementation

Since we are primarily interested in the use of distance measures by an agent to model this causal system, for preliminary investigation, we use the Dalal distance, which is the Hamming distance between worlds.11 The Hamming distance between two worlds (interpretations) is the number of propositional letters on which the two interpretations differ. We restrict ourselves to the use of propositional logic with a finite 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 of our implementation in Java of the diabetes example is shown below in Figure 3.

Figure 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 (S4, S6). The bottom-most diagram shows the transition states after the agent performs a revision operation

Results

Since the number of states and the number of actions is small, the white box stabilises after an average of five to six iterations of the learning process when the actions are chosen randomly. In most instances, both the black box and white box arrive in one of states in {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}. Furthermore, since the patient is not alert in either of {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 whereas the system will actually be in state S3 leading to a discrepancy between the prediction and observation. Under the current implementation, this problem will never be resolved by the agent.

Discussion

Our experiment is only 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 at least reduce the discrepancy between its predicted and the system's actual outcome.

Presently we are investigating other kinds of distance measure that can overcome the limitations of the Dalal distance. Furthermore, 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 leads 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 than just two actions. 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.

Hunter and Delgrande[12] have proposed 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. An interesting avenue for future work would be to incorporate and exploit the advantages offered by their method. If the belief sets can be represented in Horn clause theories under the right integrity constraints, the choice between theories may be made using crucial literals to test and eliminate falsifiable theories.[13]

In this paper, we presented a distance measure-based account of rectifying causal models. The Dalal distance was used primarily because it is very well studied in the literature, but we could employ other distance measures. One of the important avenues for future research is to investigate how one can not only choose the right distance measure, but also develop mechanisms that will allow the distance measure to be varied depending on the knowledge gained. An important question to ask is whether the distance measure should be changed just on account of the prediction being wrong, or whether more things should also be considered before changing it.

A challenge in the example that we used above is the presence of hidden variables (blood sugar level). Traditionally such problems are addressed using Bayesian inference, which can also deal with non-deterministic actions. The example that we used is very similar to Markov localization of robot in dynamic domains.[14] It would be worthwhile to explore if adapting ideas from Bayesian inference and Bayesian causal models[15] for a deterministic setting could be of use in our scenario.

Conclusion

We presented a simple non-probabilistic causal inference model of fluctuating blood sugar level in a diabetes patient using belief revision and update. An action performed by an agent who is trying to model the causal system is followed by an appropriate update of its knowledge. Comparison of the predicted behaviour of the system and the observed outcome leads to further rectification in the agent's causal model. We assumed a fixed distance measure as the underlying mechanism for improving the model. This measure itself may need to be corrected; we will examine this issue in our future work.

References

1. Patel VL, Shortliffe EH, Stefanelli M, Szolovits P, Berthold MR, Bellazzi R et al. The coming of age of artificial intelligence in medicine. Artificial Intelligence in Medicine 2009 May; 46(1): 5-17. Epub 2008 Sep 13.
2. Szolovits P. How can medical computer programs reason? In Furman S editor. Proceedings of CARDIOSTIM 88. Monaco, 1988.
3. Alchourron CE, Gardenfors P Makinson D. On the logic of theory change: Partial meet contraction and revision functions. Journal of Symbolic Logic. 1985; 50: 510-30.
4. Gardenfors P. Knowledge in Flux: Modelling the Dynamics of Epistemic States. Cambridge MA: Bradford Books, MIT Press; 1988.
5. Gardenfors P. The dynamics of belief systems:

Foundations vs. coherence theories. Revue Internationale de Philosophie. 1990; 44:24-46.
6. Gardenfors P. Belief Revision: An Introduction. In: Gärdenfors P editor. Belief Revision. Cambridge: Cambridge University Press; 1992; 1-20.
7. Grove A. Two modellings for theory change. Journal of Philosophical Logic. 1988; 17:157-70.
8. Katsuno H, Mendelzon AO. On the difference between updating a knowledge base and revising it. In Gardenfors P, editor. Belief Revision. Cambridge: Cambridge University Press; 1992; 183-203.
9. Lehmann DJ, Magidor M, Schlechta K. Distance semantics for belief revision. Journal of Symbolic Logic. 2001; 66(1): 295-317.
10. Lehmann DJ, Magidor M, Schlechta K. Distance semantics for belief revision. In Shoham Y, editor. TARK. De Zeeuwse Stromen, The Netherlands: Morgan Kaufmann; 1996; 137-145.
11. Dalal M. Investigations into a theory of knowledge base revision: Preliminary report. Proceedings of the Seventh National Conference on AI (AAAI-88). St. Paul, Minnesota, August 1988; 475-479.
12. Hunter A, Delgrande JP. Iterated belief change due to actions and observations. Journal of Artificial Intelligence Research. 2011; 40: 269-304.
13. Sattar A, Goebel R. Using crucial literals to select better theories. Computational Intelligence. 1991; 7: 11-22.
14. Fox D, Burgard W, Thrun S. Markov localization for mobile robots in dynamic domains. Journal of Artificial Intelligence Research. 1999;11: 391-427.
15. Pearl, J. Causation, Action and Counterfactuals. London: UNICOM; 1995.

ACKNOWLEDGEMENTS

We would like to thank the reviewers of AIH 2011, Perth, Australia and also the organizers for awarding a travel scholarship to the corresponding author.

PEER REVIEW

Not commissioned. Externally peer reviewed

CONFLICTS OF INTEREST

The authors declare that they have no competing interests.

FUNDING

Supported by the Safety Net Grant Number 9201101345: Eliciting and Refining Causal Knowledge, Macquarie University.

ETHICS COMMITTEE APPROVAL

N/A