Supplementing Belief Revision for The Aim of Truthlikeness

Simon D'Alfonso

April 30, 2011

Within the last few years there has been some interest in investigating the relationship between the truthlikeness (verisimilitude) and belief revision programs [2, 6] [1]. One prominent result of this investigation is that given any plausible account of truthlikeness and rational account of belief revision, expansions (+) and revisions (*) of a database (or belief state) **D** with true input A are not guaranteed to increase the database's truthlikeness. **D** here is a belief set (i.e. $\mathbf{D} = \mathrm{Cn}(\mathbf{D})$) and D stands for its propositional formula representation.

Using the classic hot-rainy-windy example logical space found in the truthlikeness literature (Table 1. w_1 is the actual world and each possible world has probability $\frac{1}{8}$), two examples are given below. For the sake of example and without loss of generality, the truthlikeness function (Tr) used for calculations throughout this article will be the Tichy/Oddie average measure [7, p. 34.]. The constructive model of AGM belief revision used is that based on possible world modelling as outlined in [6] [2].

State	h	r	w
w_1	T	T	T
w_2	T	T	F
w_3	T	F	T
w_4	T	F	F
w_5	F	T	T
w_6	F	T	F
w_7	F	F	T
w_8	F	F	F

Table 1: Truth table for example logical space

Example

$D = (h \wedge \neg r \wedge \neg w) \vee (\neg h \wedge r \wedge w)$

$A = h$

$D + A = h \wedge \neg r \wedge \neg w$

[1] Also, at the 2nd EPSA (European Philosophy of Science Association) conference a symposium titled 'Belief Revision Aiming At Truth Approximation' took place. Apparently a special issue of the journal Erkenntnis is also being dedicated to this topic.

[2] This corresponds to Dalal's update semantics.

$\text{Tr}(D + A) < \text{Tr}(D)$

Example

$D = h \wedge \neg r \wedge \neg w$

$A = (h \wedge r \wedge w) \vee (\neg h \wedge \neg r \wedge \neg w)$

$D * A = \neg h \wedge \neg r \wedge \neg w$

$\text{Tr}(D * A) < \text{Tr}(D)$

Thus, examples of decreases in truthlikeness are not hard to come by. In the first example it can be seen that the resulting increase in truth due to the input is offset by a greater decrease in truth somewhere else. In the second example, the original database content is closer to the completely false disjunct rather than the true disjunct, so minimal change favours selection of the former.

The possibility of true input decreasing truthlikeness is simply due to the fact that agents are accepting input under uncertainty and without knowledge of the complete truth (of course, if an agent already knew the complete truth then there would be no need to carry out belief revision).

These negative results prompt investigation into ways of supplementing the belief revision process with tools such as decision theory in order to make optimal decisions regarding the aim of increasing truthlikeness. To get the ball rolling on this matter, for the remainder of this article I would like to outline one simple idea, which combines non-prioritised belief revision with estimated truthlikeness calculations.

1 Screened Revision

With *non-prioritised* belief revision an agent weighs new input against the data it already holds and despite the input's novelty, it can be rejected [3, 4]. David Makinson's *screened revision* [5] extends the AGM framework and is one way to go about non-prioritised belief revision. It involves a set of potential core data $\mathbf{C}$ that is immune to revision. The database $\mathbf{D}$ is revised by the input A if A is consistent with the set $\mathbf{C} \cap \mathbf{D}$ of core data held by the agent; so with such a revision the elements of $\mathbf{C} \cap \mathbf{D}$ must remain.

An operation # on a database $\mathbf{D}$ is a *screened* revision if and only if:

$$\begin{aligned} \mathbf{D}\#A &= \mathbf{D} *_C A \text{ if } A \text{ is consistent with } \mathbf{C} \cap \mathbf{D} \\ &= \mathbf{D} \text{ otherwise} \end{aligned}$$

where using the Levi identity and defining revision in terms of contraction and expansion

$$\mathbf{D} *_C A = \text{Cn}((\mathbf{D} \div_C \neg A) \cup \{A\})$$

and $\div_C$ is a *contraction protecting* $\mathbf{C}$: $\mathbf{D} \div_C A = \cap\gamma(\mathbf{D}\bot_C A)$, where $\mathbf{D}\bot_C A$ is the set of all maximal subsets of $\mathbf{D}$ that do not imply A but do include $\mathbf{C} \cap \mathbf{D}$ and γ is a selection function, as in standard AGM [1].

2 Estimated Truthlikeness

The standard formula for expected utility in decision theory can be used to calculate the estimated truthlikeness of a statement A given prior evidence E [7, p. 180.]:

$$\mathrm{Tr}_{est}(A|E) = \sum_{i=1}^{n} \mathrm{Tr}(A, S_i) \times \mathrm{Pr}(S_i|E)$$

n stands for the number of possible worlds in the logical space, S_i stands for the state description corresponding to world i and $\mathrm{Tr}(A, S_i)$ stands for the truthlikeness of A given the actuality of world i.

3 Combining Screened Revision with Truthlikeness Estimation

Agents could use a combination of screened revision and truthlikeness estimation to help them decide whether or not to accept input. Since $\mathbf{C}$ is immune to revision, it could be treated as knowledge and used as an evidential basis from which estimated truthlikeness calculations are made. A piece of input A would be accepted if (1) it does not conflict with $\mathbf{C} \cap \mathbf{D}$ and (2) the estimated truthlikeness of $\mathbf{D}\#A$ is greater than the estimated truthlikeness of $\mathbf{D}$. Formally stated, the supplementary condition is:

$$A \in \mathbf{D}\#A \Rightarrow \mathrm{Tr}_{est}(D\#A|C) > \mathrm{Tr}_{est}(D|C)$$

Here is an example of this idea:

Example

$D = h \wedge (\neg r \vee \neg w)$

$C = h \vee r \vee w$

$\mathrm{Tr}_{est}(D|C) = 0.51$

Consider the inputs $A_1 = \neg h$ and $A_2 = r \wedge w$. Should the agent accept A_1, should the agent accept A_2?

A_1 is compatible with C and $D\#A_1 = \neg h \wedge ((r \wedge \neg w) \vee (\neg r \wedge w))$. So $\mathrm{Tr}_{est}(D\#A_1|C) = 0.48$.

A_2 is compatible with C and $D\#A_2 = h \wedge r \wedge w$. So $\mathrm{Tr}_{est}(D\#A_2|C) = 0.57$.

Thus according to these calculations, the agent should accept A_2 but reject A_1. Furthermore, this method gives an easy way to compare two inputs. If an agent could select only one of many inputs, then they should choose that input which results in the greatest estimated increase.

References

[1] P. Gardenfors. *Knowledge in Flux. Modeling the Dynamics of Epistemic States.* The MIT Press, Cambridge, Massachusetts, 1988.

[2] Cevolani Gustavo and Franceso Calandra. Approaching the truth via belief change in propositional languages. In *EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association*, pages 47–62, 2010.

[3] Sven Ove. Hansson. A survey of non-prioritised belief revision. *Erkenntnis*, 50(2-3):413–427, 1999.

[4] Sven Ove. Hansson. Logic of belief revision. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy.* Stanford University, spring 2009 edition, 2009. URL = <http://plato.stanford.edu/archives/spr2009/entries/logic-belief-revision/>.

[5] D. Makinson. Screened revision. *Theoria*, 63(1-2):14–23, 1997.

[6] I. Niiniluoto. Theory change, truthlikeness, and belief revision. In *EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association*, pages 189–199, 2010.

[7] G. Oddie. *Likeness to Truth.* D. Reidel, Dordrecht, Holland, 1986.