Updating on Conditionals = Kulback-Leibler + Causal Structure

Stephan Hartmann* Soroush Rafiee Rad†

April 16, 2012

Indicative conditional statements of the form "if X, then Y" constitute a large part of the evidence that we obtain. But how should we change our beliefs in the light of such evidence? This question has prompted a huge literature. In a recent survey of the literature, Douven (2011a) points out that much work has been done in the context of belief revision (see Boutilier and Goldszmidt (1995) and Kern-Isberner (1999, 2001)). While these works are primarily concerned with categorical belief, an account of updating probabilistic belief by learning (probabilistic) conditional information is still to be formulated. The goal of this paper is to make some progress towards this goal.

Several updating procedures have been studied in relation to conditionals while non has provided a solution that can be applied to the problem in general. The most straight forward treatment of conditionals is to identify the indicative conditional with the material conditional, in which case we can replace A → B by ¬A ∨ B. This however, as was pointed out by Popper and Miller, will always leave the updated probability of the antecedent of the conditional at most as high as its prior probability, that is, the posterior probability of A is always smaller than its prior probability after having learned that A → B. In particular if $0 < P(A) < 1$ and $P(B|A) < 1$

$$P'(A) = P(A|\neg A \vee B) < P(A).$$

This however, is highly unintuitive in certain cases. An alternative updating rule proposed by David Lewis (1976) is called *imaging*. This account works in a possible world setting where a conditional is true if its consequence holds true in the closest possible world where its antecedent is true. Imaging on ϕ then transfers the probability of every world in which ϕ is false to its closest world where ϕ holds. It turns out, however, that this proposal also fails to do justice to some of our intuitive judgments. (see Douven (2011a) for a detailed study).

Another proposal for probabilistic updating after having learned an indicative conditional is to request that the posterior probability distribution minimizes the Kullback-Leibler distance to the prior probability distribution, taking the learned information as a constraint (expressed as a conditional probability statement) into account. This procedure, which we shall study in details, has been criticized in the literature based on several clever examples. The most famous one is perhaps the Fraassen's (1981) Judy Benjamin example. This example was for long thought not to have any Bayesian solution until recently Douven and Romeijn (2011) proposed a solution

*Tilburg Center for Logic and Philosophy of Science, Tilburg University, PO Box 90153, 5000 LE, Tilburg (The Netherlands) – http://www.stephanhartmann.org – s.hartmann@uvt.nl.

†Tilburg Center for Logic and Philosophy of Science, Tilburg University, PO Box 90153, 5000 LE, Tilburg (The Netherlands) – S.RafieeRad@uvt.nl.

based on Adam's conditioning rule defined as follows: let A and B be such that $P(B \mid A)$ does not have extreme values, and that the conditional probability is updated from $P(B|A)$ to $P'(B|A)$, then

$$P'(C) = P(A \wedge B \wedge C)\frac{P'(B \mid A)}{P(B \mid A)} + P(A \wedge \neg B \wedge C)\frac{P'(\neg B \mid A)}{P(\neg B \mid A)} + P(\neg A \wedge C).$$

This along with the assumption that upon learning the conditional $A \to B$, one sets the conditional probability $P(B|A) = 1$ gives an updating rule for learning a conditional $A \to B$ as

$$P'(C) = P(C \mid \neg A)P(\neg A) + P(C \mid A \wedge B)P(A).$$

Here however the updating rule is such that learning a conditional will have no effect on the probability of its antecedent since

$$P'(A) = P(A \mid \neg A)P(\neg A) + P(A \mid A \wedge B)P(A) = P(A)$$

which is unintuitive in many cases.

In this paper, we revisit three examples put forward in the literature, Fraassen's Judy Benjamin example and two example put forward by Douven and collaborators and show that one obtains intuitively correct results for the posterior probability distribution using the Kullback-Leibler distance minimization if the underlying probabilistic model reflects the causal structure of the scenario in question.

1 The Kullback-Leibler Distance and Probabilistic Updating

The Kullback-Leibler distance $D_{KL}(P'||P)$ measures the expected difference in the informativeness of two probability distributions P' and P (from the point of view of P'). Let $S_1, \ldots, S_n$ be the possible values for a random variable S over which distributions P' and P are defined. The Kullback-Leibler distance between P' and P is then given by

$$D_{KL}(P'||P) := \sum_{i=1}^{n} P'(S_i) \log \frac{P'(S_i)}{P(S_i)} \,. \tag{1}$$

It is interesting to note that $D_{KL}(P'||P)$ is neither symmetric nor does it satisfy the triangle inequality. It is nevertheless very popular as a measure for the distance between two probability distributions. See, for example, the works by Paris (1990, 1994), Paris and Venkovska (1997) and Williamson (2005, 2008).

The Kulback-Leibler distance has also been used to justify probabilistic updating (Diaconis and Zabell 1982). To make ourselves familiar with the Kulback-Leibler distance and to introduce the methodology which we use in this paper, we shortly show how this works. Consider two binary propositional variables.[1] The variable H has two values. H: "The hypothesis holds", and $\neg H$: "The hypothesis does not hold". The variable E has the values E: "The evidence obtains", and ¬E: "The evidence does not obtain". The prior joint distribution over H and E is given by

$$\begin{aligned} P(\mathrm{H}, \mathrm{E}) = \alpha \quad &, \quad P(\mathrm{H}, \mathrm{E}) = \beta \\ P(\mathrm{H}, \mathrm{E}) = \gamma \quad &, \quad P(\mathrm{H}, \mathrm{E}) = \delta \,, \end{aligned} \tag{2}$$

[1]Throughout this paper we follow the convention that propositional variables are printed in italic script, and that the instantiations of these variables are printed in roman script. See Bovens and Hartmann (2003).

with $\alpha + \beta + \gamma + \delta = 1$. Next, we learn that the evidence E obtains. This is a constraint on the posterior distribution P' which amounts to

$$P'(\mathrm{E}) = 1\,. \tag{3}$$

If we set

$$\begin{array}{lcl} P'(\mathrm{H},\mathrm{E}) = \alpha' & , & P'(\mathrm{H},\mathrm{E}) = \beta' \\ P'(\mathrm{H},\mathrm{E}) = \gamma' & , & P'(\mathrm{H},\mathrm{E}) = \delta'\,. \end{array} \tag{4}$$

then eq. (3) implies that

$$\alpha' + \gamma' = 1 \tag{5}$$

and (as $\alpha' + \beta' + \gamma' + \delta' = 1$) that $\beta' = \delta' = 0$. Hence, the posterior distribution is given by

$$\begin{array}{rcl} P(\mathrm{H},\mathrm{E}) = \alpha' & , & P(\mathrm{H},\mathrm{E}) = 0 \\ P(\mathrm{H},\mathrm{E}) = 1 - \alpha' & , & P(\mathrm{H},\mathrm{E}) = 0\,. \end{array} \tag{6}$$

Let us now calculate the Kulback-Leibler distance between P' and P:

$$\begin{aligned} D_{KL}(P'||P) &:= \sum_{H,E} P'(H,E) \log \frac{P'(H,E)}{P(H,E)} \\ &= \alpha' \log\left(\frac{\alpha'}{\alpha}\right) + (1-\alpha') \log\left(\frac{1-\alpha'}{\gamma}\right) \end{aligned} \tag{7}$$

To find the minimum of $D_{KL}(P'||P)$, we differentiate this expression by α' and obtain after some algebra:

$$\frac{\partial D_{KL}}{\partial \alpha'} = \log\left(\frac{\gamma}{\alpha} \cdot \frac{\alpha'}{1-\alpha'}\right) \tag{8}$$

To find the minimum, we set the latter expression equal to zero and obtain:

$$\alpha' = \frac{\alpha}{\alpha + \gamma} \tag{9}$$

Hence, using eq. (5),

$$\gamma' = \frac{\gamma}{\alpha + \gamma} \tag{10}$$

We then obtain for the posterior probability of H:

$$\begin{aligned} P'(\mathrm{H}) &= \alpha' + \beta' = \frac{\alpha}{\alpha+\gamma} = \frac{P(\mathrm{H},\mathrm{E})}{P(\mathrm{H})} \\ &= P(\mathrm{H}|\mathrm{E}) \end{aligned} \tag{11}$$

To complete the proof, we calculate

$$\frac{\partial^2 D_{KL}}{\partial \alpha'^2} = \frac{1}{\alpha'(1-\alpha')} > 0, \tag{12}$$

which shows that we have indeed found a minimum. Hence, we have shown that Bayes Rule follows from minimizing the Kulback-Leibler distance between the posterior distribution and the prior distribution, if one takes the learned information as a constraint on the posterior distribution into account. In short the Kullback-Leibler distance minimization updating procedure chooses the posterior distribution that accommodates the new information and otherwise stays information theoretically as close as possible to the prior distribution.

In this paper we will show that if one considers the complete causal structure of the problem, using the similar updating procedure, i.e. Kullback-Leibler distance minimization, one can account for the intuitive results that are expected from updating probabilistic belief by learning conditional evidence as manifested in the examples by van Fraassen and Douven et al..

References

Bovens, L. and S. Hartmann (2003). *Bayesian Epistemology.* Oxford: Oxford University Press.

Diaconis, P. and S. Zabell (1982). Updating Subjective Probability, *Journal of the American Statistical Association* 77: 822–830.

Douven, I. (2011). Learning Conditional Information. Forthcomning.

Douven, I. and R. Dietz (2011). A Puzzle About Stalnaker's Hypothesis, *Topoi*, in press.

Douven, I. and J. W. Romeijn (2011). A New Resolution of the Judy Benjamin Problem, *Mind*, in press.

Fraassen, B. van (1981). A Problem for Relative Information Minimizers in Probability Kinematics, *British Journal for the Philosophy of Science* 32: 375–379.

Fraassen, B. van (1988). The Problem of Old Evidence, in D. F. Austin (ed.): *Philosophical Analysis.* Dordrecht: Reidel, pp. 153–165.

Lewis, D. (1976). Probabilities of Conditionals and Conditional Probabilities, *Philosophical Review* 85: 297–315.

Popper, K.R. and D. Miller (1983). A Proof of the Impossibility of Inductive Probability, *Nature* 302: 687–688.