Probabilistic Belief Update via Mixing Endogenous Actions and Exogenous Events

Gavin Rens and Thomas Meyer

University of Cape Town,
Centre for Artificial Intelligence Research - CSIR Meraka, South Africa
{grens,tmeyer}@cs.uct.ac.za

Abstract. We investigate methods for accommodating both agent-actions and events outside the control of an agent for updating an agent's belief state. Update via the mixture of actions and events seems not to have enjoyed much attention in the literature. The framework we consider assumes given, a probability distribution over worlds thought possible, event likelihoods, state transition likelihoods and observation likelihoods. We present three methods: (i) trading off update due to agent actions, and update due to exogenous events via a measure of observation endogeny, (ii) a more direct mixture of expected action and event effects and (iii) employing a causal network involving actions and event.

Keywords: Probability · Belief Update · Agent Actions · Exogenous Events.

1 Introduction

When an agent performs an action, it expects its surroundings to change due to the effects of the action. An agent might also notice that its surroundings have changed, but not due to one of its own actions, but due to some exogenous event. Or the surroundings might have changed due to a chain of (indirect) effects of some action the agent performed in the past. The agent could have an idea as to the likelihood of particular actions and events happening in different situations. And the agent could have knowledge about the correctness of observations it makes in different situations.

Although agent-actions and events-from-outside (including actions performed by other agents) frequently occur simultaneously, there is very little literature attempting to formalize how an agent should update its beliefs in such cases. There is more literature discussing belief update for purely agent actions (e.g. the literature on Markov decision processes) or purely exogenous events (e.g. the literature on Bayesian networks).

In this work, we assume that an agent has a probability distribution over the possible states it can be in. We then want to answer the question, what is the agent's new distribution over states, given that it executed a particular action while some (exogenous) events simultaneously occurred and a particular observation is made shortly thereafter?

For instance, if monsters leave the forest when a fire is lit, and you can light a fire (agent action) and someone else can light a fire (event) and the monsters can leave the forest (event), then what should you believe after lighting a fire while knowing that there is a probability that someone else has also lit a fire, there is a probability that the monsters will leave the forest anyway and you think you hear the monsters leaving (observation)?

In this paper, we present three methods to combine knowledge of actions and events for probabilistic belief update. All of the uncertainties about an agent's environment models may be represented by probabilities of the occurrence of actions, events, effects and observations in the different possible situations. The first method computes a measure of the degree to which an agent believes that the observation it received is the effect of its own action. The agent then updates its beliefs by trading off action update and event update using this measure. The second method assumes that updating due to actions and updating due to events should not be separate but should be more integrated. The third method deals more realistically with the causal relationships between all actions and events. The other two methods do not allow for the explicit modeling of causal relationships.

We start off by covering the necessary technical background, including the basic framework: the language and probabilistic models an agent is assumed to have. Given the formal framework, we are then ready to specify an example scenario to be used throughout the rest of the paper. In Section 4, the three methods of update are presented. Section 5 discusses related work and Section 6 summarizes the findings and looks to future work.

2 The Agent Framework

In this work, an agent has the following vocabulary.

- F, a finite set of Boolean world features
- G, a finite set of Boolean environment events
- A, a finite set of agent actions

Each world feature and environment event may be either true (present in the world, resp., occurred) or false (absent from the world, resp., did not occur). A valuation function assigns *true* (1) or *false* (0) to each feature or event. Let W be the set of valuation functions (aka, possible worlds) over world features F, such that for every function/possible world $w \in W$, w assigns either *true* or *false* to each feature in F. And let X be the set of valuation functions (aka, atomic events) over environment events G, such that for every function/atomic event $e \in X$, e assigns either *true* or *false* to each (primitive) event in G. Let L be the classical propositional language induced from F using conjunction ($\wedge$), disjunction ($\vee$) and negation ($\neg$). Let $\Omega \subseteq L$ be a set of observations the agent is interested in. An agent has the following environment models.[1]

[1] Environment models will have to be learned by the agent or supplied by a knowledge engineer during the agent design process.

- $E : X \times W \to [0,1]$ s.t. $E(e,w)$ is the probability that event e occurs in world w.
- $T : W \times (A \cup X) \times W \to [0,1]$ s.t. $T(w, æ, w')$ is the probability of being in world w' after æ occurs.
- $O : \Omega \times (A \cup X) \times W \to [0,1]$ s.t. $O(\phi, æ, w)$ is the probability that observation ϕ occurs in world w, given action or event æ, s.t. if $\phi \equiv \psi$, then $O(\phi, æ, w) = O(\psi, æ, w)$.

Finally, an agent maintains its beliefs using a belief state (probability distribution) $B : W \to [0,1]$ over worlds, s.t. $\sum_{w \in W} B(w) = 1$.

We say that a sentence $\phi \in L$ is satisfied by a world $w \in W$ (denoted $w \Vdash \phi$) iff it is true when evaluated under w according to the classical laws of logic. The set of worlds satisfying ϕ is denoted as $Mod(\phi)$. We define satisfaction of primitive event $\epsilon \in G$ evaluated under atomic event $e \in X$ similarly. When e satisfies ϵ, we denote it as $e \Vdash \epsilon$. We may write $\overline{a}$ instead of $\neg a$, and ab instead of $a \wedge b$. A belief state $\{(w_1, p_1), (w_2, p_2), \ldots, (w_n, p_n)\}$ might be written more compactly as $\langle p_1, p_2, \ldots, p_n \rangle$, where for $F = \{q, r, s\}$, $w_1 \Vdash q \wedge r \wedge s$, $w_2 \Vdash q \wedge r \wedge \neg s$, $\ldots$, $w_8 \Vdash \neg q \wedge \neg r \wedge \neg s$.

3 An Example Scenario

Consider the following scenario as a running example.

There is a forest in which a wizard and a witch live. There are also orcs (monsters) living in the area. The orcs move in and out of the forest. Wherever they are, the orcs cause all sorts of mischief. When the orcs get close to the wizard, he sometimes lights a magical fire, at which time they usually leave the forest. The witch also has a magical fire for chasing the orcs away. When the orcs are chased out of the forest, they give an indignant cry. In this example, we are interested in the wizard's mindset.

To represent belief states more compactly, we use the form $(\overline{a}, b, \overline{c}, d, p) \in B$ instead of $(w, p) \in B$ such that $w \Vdash \neg a \wedge b \wedge \neg c \wedge d$. Let the wizard's vocabulary be

- $F = \{f_{wiz}, f_{wch}, i_o, c_o\}$ where the features have the following respective meanings: the wizard's fire is lit, the witch's fire is lit, the orcs are in the forest, the orcs have just cried out.
- $G = \{\ell_{wch}, n, x\}$ with respective meanings: the witch lit her fire, the orcs entered the forest, the orcs exited the forest.
- $A = \{\ell_{wiz}\}$ meaning the wizard has the single action to light his fire.

Atomic events ℓ_{wch}, n, x and $\overline{\ell_{wch}}, n, x$ are not considered; orcs cannot both enter and exit the forest.

The wizard uses only worlds in which features f_{wiz}, f_{wch} and i_o are considered. Let the wizard initially believe to a relatively high degree that no fires are lit and the orcs are in the forest. Let the wizard initially believe to a lesser degree that no fires are lit and the orcs are outside the forest. The wizard believes only to a low degree that only the witch's fire is lite. This initial belief state B can be written as

$$(w, 0.75) \in B : w \Vdash \overline{f_{wiz}} \wedge \overline{f_{wch}} \wedge i_o \qquad (w, 0.15) \in B : w \Vdash \overline{f_{wiz}} \wedge \overline{f_{wch}} \wedge \overline{i_o}$$

$$(w, 0.05) \in B : w \Vdash \overline{f_{wiz}} \wedge f_{wch} \wedge i_o \qquad (w, 0.05) \in B : w \Vdash \overline{f_{wiz}} \wedge f_{wch} \wedge \overline{i_o}$$

and for all other worlds $w \in W$, $B(w) = 0$. This implies that $B(f_{wiz}) = 0$, $B(f_{wch}) = 0.1$, $B(if) = 0.8$ and $B(c_o) = 0$. Let the possible observations be $\Omega = \{c_o, \neg c_o\}$.

All the environment models are defined in the following tables. Unfortunately, textual explanations are unlikely to clarify the definitions, because the functions were worked out separately for each combination of arguments. This was a knowledge engineering exercise, with probabilities assigned purely according to the first author's subjective idea about how such a scenario might behave and according to the axioms of probability.

The wizard's event likelihood function $E(e, w)$ is defined as follows.

e \ w	$f_{wiz}f_{wch}i_o$	$f_{wiz}f_{wch}\overline{i_o}$	$f_{wiz}\overline{f_{wch}}i_o$	$f_{wiz}\overline{f_{wch}}\,\overline{i_o}$	$\overline{f_{wiz}}f_{wch}i_o$	$\overline{f_{wiz}}f_{wch}\overline{i_o}$	$\overline{f_{wiz}}\,\overline{f_{wch}}i_o$	$\overline{f_{wiz}}\,\overline{f_{wch}}\,\overline{i_o}$
$\ell_{wch}n\overline{x}$	0	0	0	0	0	0	0	0.1
$\ell_{wch}\overline{n}x$	0	0	0.5	0	0	0	0.3	0
$\ell_{wch}\overline{nx}$	0	0	0.1	0.5	0.1	0	0.4	0.2
$\overline{\ell_{wch}}n\overline{x}$	0	0	0	0.2	0	0.4	0	0.3
$\overline{\ell_{wch}}\overline{n}x$	1	0	0.3	0	0.6	0	0.2	0
$\overline{\ell_{wch}}\,\overline{nx}$	0	1	0.1	0.3	0.3	0.6	0.1	0.4

The wizard's transition function $T(w, æ, w')$ is defined as follows. (We do not define it for w when $B(w) = 0$, because it is not necessary for the examples.) For all worlds such that $w \Vdash \overline{f_{wiz}} \wedge f_{wch} \wedge i_o$

$æ$ \ w'	$f_{wiz}f_{wch}i_o$	$f_{wiz}f_{wch}\overline{i_o}$	$f_{wiz}\overline{f_{wch}}i_o$	$f_{wiz}\overline{f_{wch}}\,\overline{i_o}$	$\overline{f_{wiz}}f_{wch}i_o$	$\overline{f_{wiz}}f_{wch}\overline{i_o}$	$\overline{f_{wiz}}\,\overline{f_{wch}}i_o$	$\overline{f_{wiz}}\,\overline{f_{wch}}\,\overline{i_o}$
ℓ_{wiz}	0.05	0.45	0.05	0.45	0	0	0	0
$\ell_{wch}n\overline{x}$	0	0	0	0	0.25	0.25	0.25	0.25
$\ell_{wch}\overline{n}x$	0.5	0	0	0	0	0.5	0	0
$\ell_{wch}\overline{nx}$	0.5	0	0	0	0.5	0	0	0
$\overline{\ell_{wch}}n\overline{x}$	0.25	0	0.25	0	0.25	0	0.25	0
$\overline{\ell_{wch}}\overline{n}x$	0	0.25	0	0.25	0	0.25	0	0.25
$\overline{\ell_{wch}}\,\overline{nx}$	0.25	0	0.25	0	0.25	0	0.25	0

For all worlds such that $w \Vdash \overline{f_{wiz}} \wedge f_{wch} \wedge \overline{i_o}$, $T(w, æ, w')$ is defined as

$æ$ \ w'	$f_{wiz}f_{wch}i_o$	$f_{wiz}f_{wch}\overline{i_o}$	$f_{wiz}\overline{f_{wch}}i_o$	$f_{wiz}\overline{f_{wch}}\,\overline{i_o}$	$\overline{f_{wiz}}f_{wch}i_o$	$\overline{f_{wiz}}f_{wch}\overline{i_o}$	$\overline{f_{wiz}}\,\overline{f_{wch}}i_o$	$\overline{f_{wiz}}\,\overline{f_{wch}}\,\overline{i_o}$
ℓ_{wiz}	0.25	0.25	0.25	0.25	0	0	0	0
$\ell_{wch}n\overline{x}$	0.5	0	0	0	0.5	0	0	0
$\ell_{wch}\overline{n}x$	0	0.5	0	0	0	0.5	0	0
$\ell_{wch}\overline{nx}$	0	0.5	0	0	0	0.5	0	0
$\overline{\ell_{wch}}n\overline{x}$	0.25	0	0.25	0	0.25	0	0.25	0
$\overline{\ell_{wch}}\overline{n}x$	0	0.25	0	0.25	0	0.25	0	0.25
$\overline{\ell_{wch}}\,\overline{nx}$	0	0.25	0	0.25	0	0.25	0	0.25

For all worlds such that $w \Vdash \overline{f_{wiz}} \wedge \overline{f_{wch}} \wedge i_o$, $T(w, æ, w')$ is defined as

$æ$ \ w'	$f_{wiz}f_{wch}i_o$	$f_{wiz}f_{wch}\overline{i_o}$	$f_{wiz}\overline{f_{wch}}i_o$	$f_{wiz}\overline{f_{wch}}\,\overline{i_o}$	$\overline{f_{wiz}}f_{wch}i_o$	$\overline{f_{wiz}}f_{wch}\overline{i_o}$	$\overline{f_{wiz}}\,\overline{f_{wch}}i_o$	$\overline{f_{wiz}}\,\overline{f_{wch}}\,\overline{i_o}$
ℓ_{wiz}	0.25	0.25	0.25	0.25	0	0	0	0
$\ell_{wch}n\overline{x}$	0.5	0	0	0	0.5	0	0	0
$\ell_{wch}\overline{n}x$	0	0.5	0	0	0	0.5	0	0
$\ell_{wch}\overline{nx}$	0.5	0	0	0	0.5	0	0	0
$\overline{\ell_{wch}}n\overline{x}$	0	0	0.5	0	0	0	0.5	0
$\overline{\ell_{wch}}\overline{n}x$	0	0	0	0.5	0	0	0	0.5
$\overline{\ell_{wch}}\,\overline{nx}$	0	0	0.5	0	0	0	0.5	0

For all worlds such that $w \Vdash \overline{f_{wiz}} \wedge \overline{f_{wch}} \wedge \overline{i_o}$, $T(w, æ, w')$ is defined as

æ \ w	$f_{wiz}f_{wch}i_o$	$f_{wiz}f_{wch}\overline{i_o}$	$f_{wiz}\overline{f_{wch}}i_o$	$f_{wiz}\overline{f_{wch}}\overline{i_o}$	$\overline{f_{wiz}}f_{wch}i_o$	$\overline{f_{wiz}}f_{wch}\overline{i_o}$	$\overline{f_{wiz}}\overline{f_{wch}}i_o$	$\overline{f_{wiz}}\overline{f_{wch}}\overline{i_o}$
ℓ_{wiz}	0.25	0.25	0.25	0.25	0	0	0	0
$\ell_{wch}n\overline{x}$	0.5	0	0	0	0.5	0	0	0
$\ell_{wch}\overline{n}x$	0	0.5	0	0	0	0.5	0	0
$\ell_{wch}\overline{n}\overline{x}$	0	0.5	0	0	0	0.5	0	0
$\overline{\ell_{wch}}n\overline{x}$	0	0	0.5	0	0	0	0.5	0
$\overline{\ell_{wch}}\overline{n}x$	0	0	0	0.5	0	0	0	0.5
$\overline{\ell_{wch}}\overline{n}\overline{x}$	0	0	0	0.5	0	0	0	0.5

The wizard's observation function $O(c_o, æ, w)$ is defined as

æ \ w	$f_{wiz}f_{wch}i_o$	$f_{wiz}f_{wch}\overline{i_o}$	$f_{wiz}\overline{f_{wch}}i_o$	$f_{wiz}\overline{f_{wch}}\overline{i_o}$	$\overline{f_{wiz}}f_{wch}i_o$	$\overline{f_{wiz}}f_{wch}\overline{i_o}$	$\overline{f_{wiz}}\overline{f_{wch}}i_o$	$\overline{f_{wiz}}\overline{f_{wch}}\overline{i_o}$
ℓ_{wiz}	0.6	0.1	0.4	0.05	0.1	0.05	0	0
$\ell_{wch}n\overline{x}$	0.4	0	0	0	0.3	0	0	0
$\ell_{wch}\overline{n}x$	0	1	0	0	0	0.7	0	0
$\ell_{wch}\overline{n}\overline{x}$	0.2	0.1	0	0	0.1	0	0	0
$\overline{\ell_{wch}}n\overline{x}$	0.2	0	0	0	0.1	0	0	0
$\overline{\ell_{wch}}\overline{n}x$	0	0.3	0	0.1	0	0.3	0	0
$\overline{\ell_{wch}}\overline{n}\overline{x}$	0.3	0.1	0.1	0	0.1	0	0	0

4 Combining Actions and Events for Update

4.1 Mixed Update via Strength of Endogeny

We say that an observation is *exogenous* (w.r.t. an agent) when it is caused by an event outside the control of the agent, and an observation is *endogenous* (w.r.t. an agent) when the observation is due to a recent agent action.

Suppose we know what observation to expect as an effect of an action just executed by an agent. Call that observation β. And suppose that the agent actually perceives α. Our reasoning is that the less similar α is to β, the more likely it is that α is an effect of some outside (exogenous) event, that is, not due to the agent's recent action. Intuitively, if the actual observation is completely unexpected, then the agent should update its beliefs as if some environmental event had just occurred. Also intuitively, if the actual observation completely matches what the agent expects due to its most recent action, then the observation should be fully associated with the action and update should proceed as in traditional state estimation (due to a known action).

We thus measure the degree of endogeny of observation α as the degree of similarity of α to the expected observation β. In this work, observations are logical propositions. The question thus becomes, how can one compute the degree of similarity between two propositions α and β?

We shall base the degree of similarity on the number of worlds on which the two proposition correspond. This amount is adjusted by the total number of worlds under consideration. For instance, assume that $Mod(\alpha)$ and $Mod(\beta)$ overlap with one world, and assume that $Mod(\alpha)$ and $Mod(\psi)$ also overlap with one world. Now assume further that $|Mod(\alpha)| = 2$, $|Mod(\beta)| = 2$ and $|Mod(\psi)| = 4$. Then we should consider β more similar to α than ψ.

Let the similarity between two proposition be define as follows.

Definition 1.

$$S(\alpha,\beta) \doteq \frac{|Mod(\alpha) \cap Mod(\beta)|}{|Mod(\alpha) \cup Mod(\beta)|}.$$

Notice that S is symmetrical, that is, $S(\alpha,\beta) = S(\beta,\alpha)$ for all $\alpha,\beta \in \Omega$.

Another approach towards defining similarity explicitly counts non-corresponding worlds, and normalizes the final count to fall within 0 and 1: The count is the number of β-worlds in $Mod(\alpha)$ minus the number of β-worlds in $Mod(\neg\alpha)$. The general approach for this kind of normalization is to subtract the minimum possible count (which could be negative in this case) from the total count, and then divide it by the difference between the minimum and maximum counts.

Definition 2. *Let β^{ψ} be the number of β-worlds in $Mod(\psi)$. Then*

$$S'(\alpha,\beta) \doteq \frac{\beta^{\alpha} - \beta^{\neg\alpha} - (-\beta^{\neg\alpha})}{|Mod(\alpha)| - (-\beta^{\neg\alpha})}.$$

Note that the maximum count is $|Mod(\alpha)|$ and the minimum is $-\beta^{\neg\alpha}$.

Proposition 1. *$S(\alpha,\beta) = S'(\alpha,\beta)$.*

Proof. Note that $S'(\alpha,\beta) = \frac{\beta^{\alpha}-\beta^{\neg\alpha}-(-\beta^{\neg\alpha})}{|Mod(\alpha)|-(-\beta^{\neg\alpha})} = \frac{\beta^{\alpha}}{|Mod(\alpha)|+\beta^{\neg\alpha}}$. And note that $\beta^{\alpha} = |Mod(\alpha) \cap Mod(\beta)|$ and that $|Mod(\alpha)| + \beta^{\neg\alpha} = |Mod(\alpha)| + |(Mod(\beta) \cap Mod(\neg\alpha))| = |Mod(\alpha) \cup Mod(\beta)|$.

It turns out that $\frac{|X \cap Y|}{|X \cup Y|}$, where X and Y are sets, was defined by Paul Jaccard in 1901, to measure the similarity of plants by the (normalized) number of common features they display [7].[2]

As an example, suppose the vocabulary is $\{f_{wiz}, f_{wch}, i_o\}$. Note that[3]

- $Mod(f_{wiz} \land f_{wch}) = \{111, 110\}$,
- $Mod(f_{wch}) = \{111, 110, 011, 010\}$,
- $Mod(i_o) = \{111, 101, 011, 001\}$.

Thus, $S(f_{wiz} \land f_{wch}, f_{wch}) = 2/4 = 1/2$ and $S(f_{wiz} \land f_{wch}, i_o) = 1/5$ and $S(f_{wiz} \land f_{wch}, \neg f_{wch}) = 0/4 = 0$ and $S(f_{wiz} \land f_{wch}, f_{wiz} \land f_{wch}) = 2/2 = 1$.

In the context of this work, an agent does not expect a particular observation, given an executed action in a belief state. Rather, an agent has different degrees of belief for each of the observations it could possibly receive. The probability of observation β is

$$Pr(\beta \mid a, B) = \sum_{w' \in W} \sum_{w \in W} B(w)T(w,a,w')O(\beta,a,w').$$

We are now ready to define the expected endogeny of observation ϕ received after executing action a in belief state B.

[2] The measure is apparently known by different names: *Jaccard index*, *Intersection over Union* and *Jaccard similarity coefficient* (https://en.wikipedia.org/wiki/Jaccard_index).

[3] Here, worlds are represented by strings corresponding to their truth assignments.

Definition 3.

$$EEndo(\phi, a, B) \doteq \sum_{\beta \in \Omega} Pr(\beta \mid a, B) S(\phi, \beta).$$

Expected *exogeny* of observation ϕ received after executing action a in belief state B will be defined as $1 - EEndo(\phi, a, B)$.

Definition 4. *Let B be a belief state, a an action in A and ϕ be a sentence in L. The* endogenous update operator *endo is defined as*

$$B^{endo}_{a,\phi} \doteq \Big\{(w', p) \mid w' \in W, p = \gamma\big[O(\phi, a, w') \sum_{w \in W} B(w) T(w, a, w')\big]\Big\},$$

where

$$\gamma := 1 / \sum_{w' \in W} O(\phi, a, w') \sum_{w \in W} B(w) T(w, a, w').$$

Here and in the rest of the paper, γ denotes a normalizing factor.

Definition 5. *Let B be a belief state and ϕ be a sentence in L. The* exogenous update operator *exo is defined as*

$$B^{exo}_{\phi} \doteq \Big\{(w', p) \mid w' \in W, p = \gamma\big[\sum_{e \in X} O(\phi, e, w') \sum_{w \in W} B(w) E(e, w) T(w, e, w')\big]\Big\}.$$

An agent uses *exo* when it knows that observation ϕ is not (directly) due to the effects of one of its (recent) actions. The agent assumes that ϕ is due to some event in the environment, but because of the uncertainty about which events occur in which worlds, it maintains a probability distribution E over event occurrences, given a world (for all worlds).

Given Definition 3, we can generalize (endogenous) state estimation to incorporate the effects of (exogenous) events occurring in nature. We define $B^{\Delta}_{a,\phi}$ as the mixture of endogenous update and exogenous update using expected endogeny as the trade-off factor.

Definition 6.

$$B^{\Delta}_{a,\phi} \doteq \Big\{(w', p) \mid w' \in W,$$
$$p = EEndo(\phi, a, B) \cdot B^{endo}_{a,\phi}(w') + (1 - EEndo(\phi, a, B)) \cdot B^{exo}_{\phi}(w')\Big\}.$$

We now compute $B^{\Delta}_{\ell_{wiz}, c_o}$ and $B^{\Delta}_{\ell_{wiz}, \neg c_o}$. It turns out that

- $EEndo(c_o, \ell_{wiz}, B) = 0.28$,
- $B^{endo}_{\ell_{wiz}, c_o} = \langle 0.52, 0.09, 0.34, 0.05, 0.00, 0.00, 0.00, 0.00 \rangle$,
- $B^{exo}_{c_o} = \langle 0.14, 0.44, 0.02, 0.03, 0.07, 0.31, 0.00, 0.00 \rangle$,
- $EEndo(\neg c_o, \ell_{wiz}, B) = 0.72$,
- $B^{endo}_{\ell_{wiz}, \neg c_o} = \langle 0.13, 0.32, 0.20, 0.34, 0.00, 0.00, 0.00, 0.00 \rangle$,
- $B^{exo}_{\neg c_o} = \langle 0.20, 0.02, 0.1, 0.14, 0.22, 0.07, 0.10, 0.16 \rangle$.

Therefore, we get

$$B^{\Delta}_{\ell_{wiz}, c_o} = \langle 0.24, 0.34, 0.11, 0.04, 0.05, 0.22, 0.00, 0.00 \rangle$$

and

$$B^{\Delta}_{\ell_{wiz}, \neg c_o} = \langle 0.15, 0.24, 0.17, 0.29, 0.06, 0.02, 0.03, 0.04 \rangle.$$

4.2 Hybrid Endogenous-Exogenous Update

Another approach is to think of actions and events truly occurring simultaneously. Let $T(w,a,e,w')$ be the probability of a transition from world w to world w', given action a is performed and event e occurs. And let $O(\phi,a,e,w')$ be the probability of perceiving ϕ in w', given action a is performed and event e occurs. Update could then be defined as follows.

Definition 7.

$$B^{\mathscr{H}}_{a,\phi} \doteq \Big\{(w',p) \mid w' \in W, p = \gamma\big[\sum_{e\in X} O(\phi,a,e,w') \sum_{w\in W} B(w)E(e,w)T(w,a,e,w')\big]\Big\}.$$

If one assumes that actions and events are independent, then $T(w,a,e,w')$ can be replaced with $T(w,a,w')T(w,e,w')$, and $O(\phi,a,e,w')$ can be replaced with $O(\phi,a,w')O(\phi,e,w')$, and normalizing factor γ is adjusted accordingly.[4]

With this method, using the crying orcs scenario, we get

$$B^{\mathscr{H}}_{\ell_{wiz},c_o} = \langle 0.60, 0.34, 0.05, 0.01, 0.00, 0.00, 0.00, 0.00\rangle$$

and

$$B^{\mathscr{H}}_{\ell_{wiz},\neg c_o} = \langle 0.26, 0.07, 0.19, 0.48, 0.00, 0.00, 0.00, 0.00\rangle$$

4.3 Using Causal Networks

The reader might have felt uncomfortable with the difference in the chain of effects of the wizard's fire and the witch's fire. The wizard's fire causes the orcs to leave the forest, which in turn cause them to cry out, but there is no 'exit forest' event (x) explicitly mentioned. On the other hand, the witch's 'fire lighting' (ℓ_{wch}) is an explicit event and so is 'exit forest', but without a causal relationship being modeled. One would expect that fire lighting causes forest exiting (at least stochasticly).

To resolve these inconsistencies, we propose to use a network of actions and *primitive* events (elements of G) as nodes, with edges indicating causal relationships. (See, for instance, [13] on causal modeling.) More precisely, we propose a directed acyclic graph with three properties:

1. action nodes may not have parent nodes (no causes),
2. there is a special *end* node representing the final world reached after all chains of events have occurred and
3. there is a path from every node to the end node.

The crying orcs scenario could be modeled as in Figure 1(a) or 1(b). The arcs show the causal relationships. Figure 1(b) is a made-up network capturing more elaborate causal relationships. Note that some actions (e.g. a_1) do not cause any events, and some events (e.g. ϵ_1) are not caused by an action.

Such causal nets are also more compact representations than modeling $E(e,w)$ for every $e \in X$. There is likely an opportunity to represent such networks even more

[4] Note that if B and C are independent, then $Pr(A \mid B,C) = Pr(A \mid B)Pr(A \mid C)/Pr(A)$.

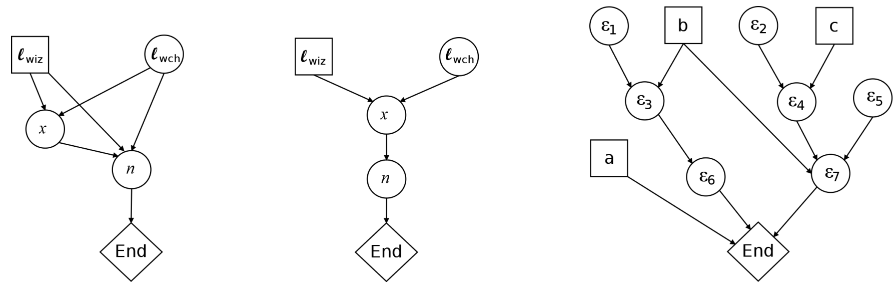

(a) One causal net for the crying orcs scenario.

(b) Another causal net for the crying orcs scenario.

(c) An example causal net.

Fig. 1. Causal networks including actions and events.

compactly by reasoning/conditioning over features (F) instead of worlds (W), which is left to do in future.

The causal net belief update operation finds the probability of each world w' in the new belief state by employing the given causal net to calculate the probability of the end node, given the action performed in the current belief state and the subsequent observation perceived in w'.

Definition 8.

$$B^{\mathscr{C}}_{a,\phi} \doteq \{(w',p) \mid w' \in W, p = \gamma \cdot \mathit{CausalNet}(a,\phi,B,\mathit{End},w')\}.$$

We define $\mathit{CausalNet}(a,\phi,B,\epsilon',w')$ with Algorithm 1. It requires that event likelihoods, transition probabilities and observation probabilities are defined in terms of primitive events $\epsilon \in G$. Hence, we define

- $E(\epsilon,w) := \sum_{e\in X, e\Vdash\epsilon} E(e,w)$,
- $T(w,\epsilon,w') := \sum_{e\in X, e\Vdash\epsilon} T(w,e,w')$,
- $O(\phi,\epsilon,w) := \sum_{e\in X, e\Vdash\epsilon} O(\phi,e,w)$.

These functions involving primitive events could be specified directly and compactly using factored representations [4, 3]. Compact specification is, however, not our focus in this paper. To keep the algorithm concise, we define $E(a,w) = 1$ for every action a and world w.

Using this method requires no more information than the causal network and the probabilistic models used in the other two methods. Note that conditional probability distributions/tables are not used as in Bayesian networks; probabilistic information is taken from the agent's environment models.

Our biggest concern with this procedure is that weighting the final probability (*sum*) by the observation probabilities (line 9) may not be the best approach. Our intuition is that ϕ is an aggregate of all effect signals caused by action a and all primitive events (to the degree that they occurred). One would thus like to weight

Algorithm 1: *CausalNet*

Input: $a, \phi, B, \epsilon', w'$
1 $sum \leftarrow 0$;
2 **for** ϵ *in Parents*(ϵ') **do**
3   **for** w *in* W **do**
4   **if** ϵ *is an orphan or* $\epsilon == a$ **then**
5             $sum \leftarrow sum + B(w)E(\epsilon, w)T(w, \epsilon, w')$
6   **else**
7             $sum \leftarrow sum + CausalNet(a, \phi, B, \epsilon, w)T(w, \epsilon, w')$
8 **if** $\epsilon' == End$ **then**
9   **return** $sum \cdot \prod_{\epsilon \in Parents(End)} O(\phi, \epsilon, w')$
10 **else**
11   **return** sum

the term being added to sum (lines 5 & 7) by the probability of some observation ϕ' which contributed to the final observation ϕ. But what are the ϕ'?

When using the causal net of Figure 1(a), we get

$$B^{\mathscr{C}_a}_{\ell_{wiz}, c_o} = \langle 0.00, 0.55, 0.00, 0.02, 0.00, 0.43, 0.00, 0.00 \rangle$$

and

$$B^{\mathscr{C}_a}_{\ell_{wiz}, \neg c_o} = \langle 0.00, 0.29, 0.00, 0.19, 0.00, 0.32, 0.00, 0.20 \rangle$$

When using the causal net of Figure 1(b), we get

$$B^{\mathscr{C}_b}_{\ell_{wiz}, c_o} = 0.00, 0.55, 0.00, 0.02, 0.00, 0.42, 0.0, 0.00$$

and

$$B^{\mathscr{C}_b}_{\ell_{wiz}, \neg c_o} = \langle 0.00, 0.27, 0.00, 0.22, 0.00, 0.30, 0.00, 0.22 \rangle$$

5 Related Work

Without it being labeled as endogenous, the state estimation function has been the basic belief state update function in POMDP theory since the 60's [1, 12, 11]. The *endo* operator is essentially the update operation used in partially observable Markov decision process (POMDP) theory, except that it takes propositional sentences as observations instead of elements from Ω, and POMDPs operate over primitive states, not logical worlds.

The first well-known formal definition of classical (qualitative/non-probabilistic) belief update was proposed by Katsuno and Mendelzon [8]. They propose a family of preorders $\{\leq_w \mid w \in W\}$ for their operator, where each $\leq_w$ is a reflexive, transitive relation over W. Each such relation is interpreted as follows: if $u \leq_w v$ then u is at least as plausible a change relative to w as is v; that is, situation w would more readily evolve into u than it would into v. They proposed a set of rationality postulates and proved that their update operator satisfies the postulates.

Shapiro and Pagnucco [17] propose an approach based on the situation calculus, where exogenous actions are presumed present, which may influence what is sensed. They assume sensors to be exact and actions to be deterministic.

Our framework assumes that every action performed by the agent is immediately followed by an observation. Their framework assumes that an (endogenous) action is either a sensing action (with only epistemic and no ontic effects) or a non-sensing action (with only ontic and no epistemic effects).

Agent actions (whether sensing or not) may be interspersed with exogenous actions. When what is sensed contradicts what the agent believes, the occurrence of exogenous actions is hypothesised to correct discrepancies between belief and sensing.

Whereas our framework assumes that endogenous and exogenous actions (and their corresponding effects/observations) occur more or less simultaneously (hence, the uncertainty in the source of the subsequent observation), in the framework of Shapiro and Pagnucco [17], actions always occur separately, in sequences. In their framework, update involves only the ontic effects of actions; and revision occurs when observations due to sensing is unexpected. So they use revision, even for belief change due to ontic (exogenous) actions. We categorize any belief change due to ontic actions as update. Here, we see again (as is still often the case) that the difference between revision and update is unclear.

Finally, the main difference between the two frameworks is that ours has a model for the likelihood of exogenous actions (which we call events) occurring. Due to our framework assuming that actions are stochastic and observations are noisy, whereas their actions are deterministic and observations exact, it makes a comparison of the two frameworks difficult.

Van Ditmarsch and Kooi [5] present of modal logic for reasoning about multiple agents performing epistemic and ontic actions. To quote them, "In the literature update models are also called action models. Here we follow [...] and call them update models, since no agency seems to be involved."

Agency of agents in their logic is not modeled, that is, events (including ontic events) may or may not be associated with the actions of agents. In a sense, in their framework. all events are exogenous with some uncertainty about what event actually occurred (each agent has a different uncertainty model). It might be possible to simulate the mixing of exogenous and endogenous events/actions in our framework with their *multi-pointed update* (update given a set of events).

However, besides the fact that Van Ditmarsch and Kooi [5] work with Kripke models to represent uncertainty (whereas we use probabilities), it is hard to see whether their logic could essentially represent the uncertainty about the source of ontic effects. In any case, they seem not to be interested or aware of the possibility. A comparison is left for future work.

There are many more logics, formalisms and frameworks which contain both notions of ontic actions and epistemic observation. It seems that in most of those, there is the potential to confuse or conflate the idea that an observation is epistemic in nature or has a physical source. But, as far as we know, none of them attempts to model the mixture of the source of an observation from physical events (exogenous or endogenous).

Rens [15] combines probabilistic belief (belief state) update and revision by trading off between the two via a measure of ontic strength (the agent's confidence that the observation is due to a physical happening). In that work, Rens assumes that the agent is passive, that is, that the agent does not perform actions, it only receives observations. His update operator is defined as

$$B^{\diamond}_{\phi} \doteq \Big\{(w',p) \mid w' \in W, p = \gamma O(\phi,w') \sum_{w \in W} \sum_{e \in X} T(w,e,w')E(e,w)B(w)\Big\},$$

which is inspired by Boutilier [2] who mentioned that event-driven update can be defined as $\big\{(w',p) \mid w' \in W, p = \sum_{w \in W} \sum_{e \in X} T(w,e,w')E(e,w)B(w)\big\}$. Note that Rens's observation function excludes an action/event argument. And Rens and Boutilier take the set of events to be atomic objects, not derived from a set of primitive events like G.

Rens [16] defined an operator very similar to $B^{\triangle}_{a,\phi}$ for use in a stochastic belief change framework:

$$B^{\triangleleft}_{a,\phi} \doteq \Big\{(w',p) \mid w' \in W, p = Eng(\phi,w',a)B^{endo}_{a,\phi}(w') + (1 - Eng(\phi,w',a))B^{\diamond}_{\phi}(w')\Big\},$$

where $\phi \in L$ is a propositional sentence and $Eng(\phi,w',a)$ is the agent's confidence that ϕ perceived in w' was caused by action a. In other words, *Eng* has the same meaning as *EEndo* but is not derived, but assumed given by the framework designer.

Dynamic Bayesian networks (DBNs) model physical change by specifying how a BN changes due to an action. A DBN is a pair $\langle \mathscr{B}_0, \mathscr{B}_{\rightarrow} \rangle$, where $\mathscr{B}_0$ is a BN representing an initial belief state and $\mathscr{B}_{\rightarrow}$ is a BN representing how probabilities of events change from one time-slice to the next [9]. We suspect that our causal networks will be more compact than DBNs, at least, graphically. A formal comparison of the two graphical models is left for the future.

Although Lang's work [10] is not directly applicable to ours in terms of 'mixing' endogenous and exogenous update, it does unpack and highlight several important characteristics of update. Lang also discusses the relationship between update and revision. His insights are on the fundamental nature of update and might guide our future efforts in this area.

The work of Halpern [6] and Pearl [14] should also be consulted when the causal net approach is developed.

6 Conclusion

In this paper, we have made a few suggestions on how to take into account both an agent's action and some other events that might occur simultaneously to update the agent's belief state. With respect to our crying orcs scenario, the update operation results in very different belief states of the wizard agent. Table 6 places the updated

Table 1. Comparison of belief states updated after doing ℓ_{wiz} and perceiving c_o.

	$f_{wiz}f_{wch}i_o$	$f_{wiz}f_{wch}\overline{i_o}$	$f_{wiz}\overline{f_{wch}}i_o$	$f_{wiz}\overline{f_{wch}}\,\overline{i_o}$	$\overline{f_{wiz}}f_{wch}i_o$	$\overline{f_{wiz}}f_{wch}\overline{i_o}$	$\overline{f_{wiz}}\,\overline{f_{wch}}i_o$	$\overline{f_{wiz}}\,\overline{f_{wch}}\,\overline{i_o}$
$B^{\Delta}_{\ell_{wiz},c_o}$	0.24	0.34	0.11	0.04	0.05	0.22	0.00	0.00
$B^{\mathscr{H}}_{\ell_{wiz},c_o}$	0.26	0.07	0.19	0.48	0.00	0.00	0.00	0.00
$B^{\mathscr{C}_a}_{\ell_{wiz},c_o}$	0.00	0.55	0.00	0.02	0.00	0.43	0.00	0.00
$B^{\mathscr{C}_b}_{\ell_{wiz},c_o}$	0.00	0.55	0.00	0.02	0.00	0.42	0.00	0.00

belief state distributions next to each other - updated due to the wizard's action of lighting his fire and perceiving the orcs' crying.

Questions we may ask are

- Do the three methods simply reflect different but reasonable update strategies?
- Are some of the methods fundamentally flawed?
- Is it possible or even worth it to combine ideas used in the different methods?

Developing the method employing causal networks seems especially promising. However, a major question with that method is where and how to employ the observation weight (probability) in the calculations.

Note that most terms in the methods defined in this paper take the form

$$B(w)E(æ,w)T(w,æ,w')O(\phi,æ,w'). \quad (1)$$

If any one of the functions is zero, the whole term is zero. Knowing this, there are two things to watch out for. One, for an update operator to be well behaved, there must be at least one world w' such that none of the functions in (1) evaluates to zero. Two, if either $E(æ,w)$ or $T(w,æ,w')$ is known to be zero at w, then the other function value needs not be known/specified at w, and if either $T(w,æ,w')$ or $O(\phi,æ,w')$ is known to be zero at w', then the other function value needs not be known/specified at w'.

Rationality postulates for qualitative belief update have been proposed [8, e.g.], but we have found it difficult to translate them into probabilistic versions. We would like to tackle this in future.

Acknowledgements

Gavin Rens was supported by a Clause Leon Foundation postdoctoral fellowship while conducting this research. This work is based on research supported in part by the National Research Foundation of South Africa (Grant number UID 98019). Thomas Meyer has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agr. No. 690974.

References

1. Aström, K.: Optimal control of Markov decision processes with incomplete state estimation. Journal of Mathematical Analysis and Applications **10**, 174–205 (1965)
2. Boutilier, C.: A unified model of qualitative belief change: A dynamical systems perspective. Artificial Intelligence **98**(1–2), 281–316 (1998)
3. Boutilier, C., Dean, T., Hanks, S.: Decision-theoretic planning: Structural assumptions and computational leverage. Journal of Artif. Intell. Research (JAIR) **11**, 1–94 (1999)
4. Boutilier, C., Poole, D.: Computing optimal policies for partially observable decision processes using compact representations. In: Proceedings of the Thirteenth Natl. Conf. on Artif. Intell. (AAAI-96). pp. 1168–1175. AAAI Press, Menlo Park, CA (1996)
5. Ditmarsch, H.V., Kooi, B.: Semantic results for ontic and epistemic change. In: G. Bonanno, W.v.d.H., Wooldridge, M. (eds.) Logic and the Foundations of Game and Decision Theory (LOFT 7), Texts in Logic and Games, vol. 3, pp. 87–117. Amsterdam University Press, Amsterdam, Netherlands (2008)
6. Halpern, J.: Actual Causality. MIT Press, Massachusetts/England (2016)
7. Jaccard, P.: Étude comparative de la distribution florale dans une portion des Alpes et des Jura. Bulletin de la Société Vaudoise des Sciences Naturelles **37**, 547–579 (1901)
8. Katsuno, H., Mendelzon, A.: On the difference between updating a knowledge base and revising it. In: Proceedings of the Second Intl. Conf. on Principles of Knowledge Representation and Reasoning. pp. 387–394 (1991)
9. Koller, D., Friedman, N.: Probabilistic Graphical Models: Principles and Techniques. The MIT Press, Cambridge, MA and London, England (2009)
10. Lang, J.: Belief update revisited. In: de Mantaras, R.L. (ed.) Proceedings of the Twentieth Intl. Joint Conf. on Artif. Intell. (IJCAI-07). pp. 2517–2522. AAAI Press, Menlo Park, CA (January 2007)
11. Lovejoy, W.: A survey of algorithmic methods for partially observed Markov decision processes. Annals of Operations Research **28**, 47–66 (1991)
12. Monahan, G.: A survey of partially observable Markov decision processes: Theory, models, and algorithms. Management Science **28**(1), 1–16 (1982)
13. Pearl, J.: Causality: Models, Reasoning, and Inference. Cambridge University Press, New York, NY, USA (2000)
14. Pearl, J., Mackenzie, D.: The Book of Why: The New Science of Cause and Effect. Basic Books, New York, USA (2018)
15. Rens, G.: On stochastic belief revision and update and their combination. In: Kern-Isberner, G., Wassermann, R. (eds.) Proceedings of the Sixteenth Intl. Workshop on Non-Monotonic Reasoning (NMR). pp. 123–132. Technical University of Dortmund (2016)
16. Rens, G.: A stochastic belief change framework with an observation stream and defaults as expired observations. In: Booth, R., Casini, G., Klarman, S., Richard, G., Varzinczak, I. (eds.) Proceedings of the Third Intl. Workshop on Defeasible and Ampliative Reasoning (DARe-2016). CEUR Workshop Proceedings, ceur-ws.org (August 2016)
17. Shapiro, S., Pagnucco, M.: Iterated belief change and exogenous actions in the situation calculus. In: Proceedings of the Sixteenth European Conf. on Artif. Intell. (ECAI-04). pp. 878–882. IOS Press, Amsterdam (2004)