Reasoning with Probabilities
Dynamics

Joshua Sack

August 8, 2013

Dynamic Epistemic Logic

Action models

Definition (Action model)

Assume some language $\mathcal{L}$ (we have in mind *dynamic epistemic logic*, but its definition depends on what follows). An *action model* is a tuple $\mathbf{A} = (S, R_i, \text{pre})$, such that

- S is a set of states
- $R_i \subseteq S \times S$ is a binary relation
- $\text{pre} : S \to \mathcal{L}$ a precondition function.

Update product

Definition

Let

$$\mathbf{M} = (S^M, R^M, \|\cdot\|^M)$$

be an epistemic model. Let

$$\mathbf{A} = (S^A, R^A, \text{pre})$$

be an action model. Define $\mathbf{M} \otimes \mathbf{A} = (S, R, \|\cdot\|)$ by

- $S = \{(s, e) \in S^M \times S^A \mid (S, s) \vDash \text{pre}(e)\}$
- $(s, e) R_i (t, f)$ if and only if $s R_i^M t$ and $e R_i^A f$
- $(s, e) \in \|p|$ if and only if $s \in \|p\|^M$.

Syntax of dynamic epistemic logic

Add to epistemic logic formulas the form $[A, e]\varphi$ for every pointed action model (A, e).

- $M, s \vDash [A, e]\varphi$ if and only if $M, s \vDash \mathrm{pre}(e)$ implies $M \otimes A, (s, e) \vDash \varphi$

Probabilistic Dynamic Epistemic Logic

General goal

Describe how (properties of) a probabilistic epistemic model $\mathbf{M}$ transforms into (properties of) another model $\mathbf{M}'$ via an action, given by an action model $\mathbf{A}$:

$$\mathbf{M} \otimes \mathbf{A} \mapsto \mathbf{M}'$$

Different types of probabilities

- Prior Probabilities over States: These are the probabilities agents have in the prior probabilistic epistemic model.
- Occurrence Probabilities: These are (objective) probabilities of certain events taking place given certain preconditions.
- Observation Probabilities over events: These are the probabilities agents assign to certain events having taken place.

Both the observation probabilities and the occurrence probabilities are formalized in action models.

Probabilistic action model

$\mathbf{A} = (X, \{R_i\}, \Phi, \mathrm{pre}, \{\mu_i\})$

- X is a *finite* set
- $R_i \subseteq X \times X$ is a binary relation
- $\Phi = \{\varphi_1, \ldots, \phi_n\}$ is a finite set of pairwise inconsistent formulas called *preconditions*
- $\mathrm{pre} : \Phi \to (X \to [0, 1])$
- $\mu_i : X \to (X \to [0, 1])$ assigns for each $e \in X$, a probability function $\mu_{i,e}$ on X

- J. van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probability. *Studia Logica* 93:67–96, (2009).

Probabilistic action model

$\mathbf{A} = (X, \{R_i\}, \Phi, \text{pre}, \{\mu_i\})$

- X is a *finite* set
- $R_i \subseteq X \times X$ is a binary relation
- $\Phi = \{\varphi_1, \ldots, \phi_n\}$ is a finite set of pairwise inconsistent formulas called *preconditions*
- $\text{pre} : \Phi \to (X \to [0, 1])$
- $\mu_i : X \to (X \to [0, 1])$ assigns for each $e \in X$, a probability function $\mu_{i,e}$ on X

- J. van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probability. *Studia Logica* 93:67–96, (2009).

Update product

$M \otimes A = (X^{\otimes}, \{R_i^{\otimes}\}, \|\cdot\|^{\otimes}, \{\mathbb{P}_i^{\otimes}\}$

- $X^{\otimes} = \{(s,e) \in X^M \times X^A \mid \mathrm{pre}(s)(e) > 0\}$, (where $\mathrm{pre}(s)(e)$ is $\mathrm{pre}(\varphi_i)(e)$ for $\phi_i \in \Phi$ and $s \vDash \phi_i$).
- $(s,e)R_i^{\otimes}(t,f)$ iff $sR_i^M t$ and $eR_i^A f$
- $(s,e) \in \|p\|^{\otimes}$ iff $s \in \|p\|^M$
- $\mathbb{P}_i^{\otimes}$ is an updated probability that can be define differently depending on whether M is discrete (see following definitions)

Discrete case

$\mathbb{P}^{\otimes}_{i,(x,e)}$ is $\mu^{\otimes}_{i,(x,e)}$ defined by setting $\mu^{\otimes}_{i,(x,e)}(y,f)$ to be

$$\frac{\mu_{i,x}(y)\cdot \mathrm{pre}(y)(f)\cdot \mu_{i,e}(f)}{\sum_{\substack{z\in X^M\\ g\in X^A}} \mu_{i,x}(g)\cdot \mathrm{pre}(z)(g)\cdot \mu_{i,e}(g)}$$

if the denominator is non-zero, and $\mu^{\otimes}_{i,(x,e)}$ is the zero function otherwise.

- Prior probability
- Occurrence probability
- Observation probability
- Normalize

Continuous case

$\mathbb{P}^{\otimes}_{i,(x,e)} = (S^{\otimes}_{i,(x,e)}, \mathcal{A}^{\otimes}_{i,(x,e)}, \mu^{\otimes}_{i,(x,e)})$

- $S^{\otimes}_{i,(x,e)} = \{(y,f) \in X^{\otimes} \mid y \in S_{i,x}\} = X^{\otimes} \cap (S_{i,x} \times X^{A})$
- $\mathcal{A}_{i,(x,e)}$ is the smallest σ-algebra containing

$$\{(A \times B) \cap X^{\otimes} \mid A \in \mathcal{A}_{i,x}, B \subseteq X^{A}\}$$

- $\mu^{\otimes}_{i,(x,e)} : \mathcal{A}^{\otimes}_{i,(x,e)} \to [0,1]$ is defined for each $A' \in \mathcal{A}^{\otimes}_{i,(x,e)}$ by setting $\mu^{\otimes}_{i,(x,e)}(A')$ to be

$$\frac{\sum_{\substack{\varphi \in \Phi \\ e' \in X^{A}}} (\mu_{i,x})^{*}([\![\varphi]\!]^{\mathbf{M}} \cap \pi_1[A'_{e'}]) \cdot \mathrm{pre}(\varphi)(e') \cdot \mu_i(e)(e')}{\sum_{\substack{\varphi \in \Phi \\ e'' \in E}} (\mu_{i,x})^{*}([\![\varphi]\!]^{\mathbf{M}} \cap S_{i,x}) \cdot \mathrm{pre}(\varphi)(e'') \cdot \mu_i(e)(e'')}$$

if the denominator is strictly positive, and $\mu^{\otimes}_{i,(x,e)}(A) = 0$ otherwise,

Example

Suppose you are reading about some horrible disease on a website, and start wondering whether you have it. The chances of having the disease are very slight: 1 in 100,000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease the chance that this gland is swollen is 97%, while if you do not have the disease, the chance is 0% that it is swollen in this way. You immediately examine the gland. The problem is that it is hard to determine if it is swollen or not. It is the first time you actually examine the gland and — not being a physician — you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What chances should you assign to having the disease?

Example

Suppose you are reading about some horrible disease on a website, and start wondering whether you have it. The chances of having the disease are very slight: 1 in 100,000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease the chance that this gland is swollen is 97%, while if you do not have the disease, the chance is 0% that it is swollen in this way. You immediately examine the gland. The problem is that it is hard to determine if it is swollen or not. It is the first time you actually examine the gland and — not being a physician — you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What chances should you assign to having the disease?

Example

Suppose you are reading about some horrible disease on a website, and start wondering whether you have it. The chances of having the disease are very slight: 1 in 100,000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease the chance that this gland is swollen is 97%, while if you do not have the disease, the chance is 0% that it is swollen in this way. You immediately examine the gland. The problem is that it is hard to determine if it is swollen or not. It is the first time you actually examine the gland and — not being a physician — you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What chances should you assign to having the disease?

Example

Suppose you are reading about some horrible disease on a website, and start wondering whether you have it. The chances of having the disease are very slight: 1 in 100,000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease the chance that this gland is swollen is 97%, while if you do not have the disease, the chance is 0% that it is swollen in this way. You immediately examine the gland. The problem is that it is hard to determine if it is swollen or not. It is the first time you actually examine the gland and — not being a physician — you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What chances should you assign to having the disease?

Example

Suppose you are reading about some horrible disease on a website, and start wondering whether you have it. The chances of having the disease are very slight: 1 in 100,000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease the chance that this gland is swollen is 97%, while if you do not have the disease, the chance is 0% that it is swollen in this way. You immediately examine the gland. The problem is that it is hard to determine if it is swollen or not. It is the first time you actually examine the gland and — not being a physician — you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What chances should you assign to having the disease?

Example

Suppose you are reading about some horrible disease on a website, and start wondering whether you have it. The chances of having the disease are very slight: 1 in 100,000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease the chance that this gland is swollen is 97%, while if you do not have the disease, the chance is 0% that it is swollen in this way. You immediately examine the gland. The problem is that it is hard to determine if it is swollen or not. It is the first time you actually examine the gland and — not being a physician — you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What chances should you assign to having the disease?

Example

Suppose you are reading about some horrible disease on a website, and start wondering whether you have it. The chances of having the disease are very slight: 1 in 100,000. The website states that one of the symptoms of this disease is that a certain gland is swollen. If you have the disease the chance that this gland is swollen is 97%, while if you do not have the disease, the chance is 0% that it is swollen in this way. You immediately examine the gland. The problem is that it is hard to determine if it is swollen or not. It is the first time you actually examine the gland and — not being a physician — you do not know what its size ought to be. You are uncertain, but you think the chances are 50% that the gland is swollen. What chances should you assign to having the disease?

Let $(X, R, \Phi, \mathrm{pre}, \mu)$

- $X = \{s, n\}$, where s is the event that the gland is swells, and n is the event that it does not swell
- $R = X^2$ is your epistemic relation
- $\Phi = \{p, \neg p\}$, where p means that you have the disease.
- pre (giving the occurrence probabilities) is defined by

$$\mathrm{pre}(p) : \begin{cases} s \mapsto .97 \\ n \mapsto .03 \end{cases} \qquad \mathrm{pre}(\neg p) : \begin{cases} s \mapsto 0 \\ n \mapsto 1 \end{cases}$$

- μ (giving the observation probabilities) is defined by

$$\begin{tikzcd} s \arrow[loop left, ".5"] \arrow[rr, leftrightarrow, ".5"] & & n \arrow[loop right, ".5"] \end{tikzcd}$$

Prior model (giving prior probabilities):

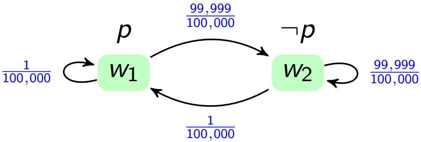

Updated model (Exercise: Find $\alpha, \beta, \gamma, x, y, z$.)

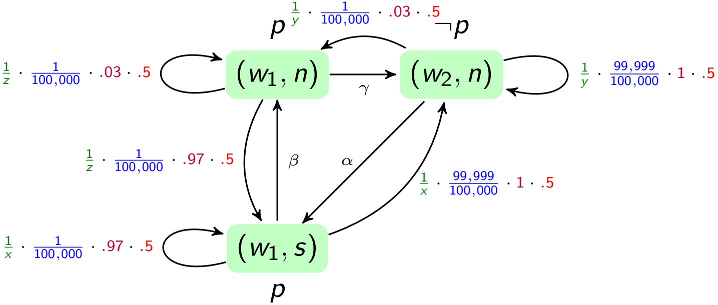

Proof system

- All axioms and rules from probabilistic epistemic logic
- $[A,e]p \leftrightarrow (\mathrm{pre}_{A,e} \to p)$
  $(\mathrm{pre}_{A,e} \equiv \bigvee_{\{\varphi\in\Phi \mid \mathrm{pre}(\varphi)(e)>0\}}\varphi)$
- $[A,e](\psi_1 \wedge \psi_2) \leftrightarrow [A,e]\psi_1 \wedge [A,e]\psi_2)$
- $[A,e][i]\psi \leftrightarrow (\mathrm{pre}_{A,e} \to \bigwedge_{(e,f)\in R_i}[i][A,f]\psi$
- probability reduction rule given by (next slide)
- $\vdash \varphi$ implies $\vdash [A,e]\varphi$

The axioms is blue are called reduction axiom, as repeated applications of them will eliminate occurrences of the action operators $[A,e]$.

Probability reduction axiom

$$
\begin{aligned}
&[A, e] \textstyle\sum_\ell a_\ell P_i(\psi_\ell) \geq b \\
&\leftrightarrow (\mathrm{pre}_{A,e} \to \Big(\Big(\sum_{\substack{\varphi\in\Phi \\ f\in A}} k_{i,e,\varphi,f} P_i(\varphi) = 0 \wedge 0 \geq b \Big) \\
&\qquad\qquad\qquad \vee \Big(\sum_{\substack{\varphi\in\Phi \\ f\in A}} k_{i,e,\varphi,f} P_i(\varphi) > 0 \wedge \chi \Big)
\end{aligned}
$$

where

$$k_{i,e,\varphi,f} \stackrel{\mathrm{def}}{=} \mathrm{pre}(\varphi)(f) \cdot \mu_i(e)(f) \in \mathbb{R}$$

and

$$\chi \stackrel{\mathrm{def}}{=} \sum_{\substack{\ell \\ \varphi\in\Phi \\ f\in A}} a_\ell k_{i,e,\varphi,f} P_i(\varphi \wedge \langle \mathrm{A}, \mathrm{e} \rangle \psi_\ell) \geq \sum_{\substack{\varphi\in\Phi \\ f\in A}} b k_{i,e,\varphi,f} P_i(\varphi)$$