Updating Beliefs

Material used

- Halpern: Reasoning about Uncertainty. Chapter 3

1 Updating knowledge
2 Probabilistic conditioning
3 Conditioning with sets of probabilities
4 Conditioning inner and outer measures
5 Conditioning possibility measures
6 Conditioning ranking functions

0 Introduction

- Agents continually obtain new information and then must update their beliefs to takes this new information into account.
- How this should be done depends in part on how uncertainty is represented. Each of the methods of representing uncertainty has an associated method of updating

By understanding the mechanism of updating we also get a better understanding of the different approaches to represent uncertainty.

1 Updating knowledge

Informally, if an epistemic space (W, W^0) is updated by some new information U (expressing that the actual world is in U), then we have to take the set of possible worlds to be $W^0 \cap U$.

Example: Tossing a die; $W = \{w_1, w_2, w_3, w_4, w_5, w_6\}$. Assume the initial epistemic space is (W, W) and then the agent learns that the die always lands on an even number: $U = \{w_2, w_4, w_6\}$. Then the system is updated to $(W, W \cap U)$.

Definition 1: Let (W, W^0) be an epistemic space, and $U \subseteq W$ a proposition. Then the operation of updating the initial epistemic space by the new information U is as follows:

$$(W, W^0)|U = (W, W^0 \cap U)$$

Example

Suppose that a world describes which of 100 people have a certain disease. A world u can be taken as a tuple of hundred 0s or 1s where the i^{th} component is 1 iff individual i has the disease. $W = \{0,1\}^{100}$

- Individual i has the disease: {u: u[i]=1}
- At least seven people have it: $\{u: |\{i: u[i]=1\}| \geq 7\}$
- More than 50% have it: $\{u: |\{i: u[i]=1\}| > 50\}$

$(W, W)|U_1 \ldots U_n = (W, W \cap U_1 \cap \ldots \cap U_n)$

Problems with this form of updating?

Four problems

- Explicit representation of 2^{100} possible worlds is beyond the capacity of any computational system. Implicit representation required
- Memory management (what about forgetting?)
- Inconsistent knowledge → belief revision
- Only one agent is involved. Normally, there are more than one agent; some of them have the relevant information, others not; some know who has the information, others not; … → multi agent systems

2 Probabilistic conditioning

Suppose that an agent's uncertainty is represented by a probability measure μ on W and he the agent observes or learns that U. How should μ be updated to a new probability measure $\mu' = \mu|U$?

(I) $\mu'(\neg U) = 0$

(II) $\mu(V_1)/\mu(V_2) = \mu'(V_1)/\mu'(V_2)$ [If all that the agent has learned is U then the relative likelihood of worlds in U should remain unchanged]

Fact 1: If $\mu(U)>0$ and $\mu' = \mu|U$ is a probability measure on W satisfying (I) and (II) then $\mu'(V) = \mu(V \cap U)/\mu(U)$.

Proof: Halpern, p. 72

The unique function $\mu'(V) = (\mu|U)(V)$ that satisfies the condition (I) and (II) is called *conditional probability*.

Following traditional practice we write $\mu(V|U)$, and we have

$$\text{(III)} \qquad \mu(V|U) = \mu(V \cap U)/\mu(U)$$

According to Halperns presentation this is not a definition but a consequence of the implicit definition given by (I) and (II). However, both formulations (I+II versus III) are equivalent.

Fact 3 (Bayes' rule): If $\mu(U), \mu(V) > 0$, then

$\mu(V|U) = \mu(U|V)\, \mu(V)/\mu(U)$

Proof: $\mu(U|V)\, \mu(V)/\mu(U) = \mu(V \cap U)\, \mu(V)/\mu(U)\, \mu(V) = \mu(V \cap U)/\mu(U)$.

Consequence

$\mu(V|U) = \mu(U|V)\, \mu(V)/ (\mu(U|V)\, \mu(V) + \mu(U|\neg V)\, \mu(\neg V))$

Proof: $\mu(U) = \mu(U \cap V) + \mu(U \cap \neg V) = \mu(U|V)\, \mu(V) + \mu(U|\neg V)\, \mu(\neg V))$

Standard Example

Suppose that Bob tests positive on a AIDS test that is known to be 99 % reliable. How likely is it that Bob has AIDS?

Reliability of 99%: 99% of the subjects with AIDS tested positive and 99% of the subjects that did not have AIDS tested negative.

$\mu(P|A) = 0.99$ and $\mu(\neg P|\neg A) = 0.99$

To calculate μ(A|P) we need the prior μ(A).

$$
\begin{aligned}
\mu(A|P) \quad &= \mu(P|A)\,\mu(A)/(\,\mu(P|A)\,\mu(A) + \mu(P|\neg A)\,\mu(\neg A)) \\
&= 0.99\,\mu(A)/(0.01 + 0.98\,\mu(A)) \\
&= 0.5 \;\text{ if }\; \mu(A) = 0.01 \\
&\approx 0.09 \text{ if } \mu(A) = 0.001 \;\; (\approx 0.98 \text{ if } \mu(A) = 1/3)
\end{aligned}
$$

Example with a biased coin

Suppose Alice has a coin and she knows that it has either bias 2/3 (*head is preferred*) or bias 1/3 (*tail preferred*). Further, she considers it much more likely that the bias is 1/3 than 2/3. Thus, she assigns a probability 0.99 to bias 1/3 and 0.01 to bias 2/3. Alice tosses the coin 25 times to learn more about its bias; she sees 19 heads and 6 tails. What is the result of updating her probabilities?

$W = \{1/3, 2/3\} \times \{\mathrm{h,t}\}^{25}$
e.g. (1/3, hhhhhhhhhhhhhhhhhhhtttttt)
or (2/3, thttthhthhtthhhttttthhthh)

Remark: the bias is taken here as a part of the possible world!

We have

- $\mu(B_{1/3}) = 0.99$ and $\mu(B_{2/3}) = 0.01$.
- $\mu(H^n|B_{1/3}) = (1/3)^n (2/3)^{25-n}$, where H^n denotes a particular sequence of tosses with n heads and 25–n tails.

We have to calculate $\mu(B_{1/3}|H^n)$, with $n = 19$

$\mu(B_{1/3}|H^{19}) =$
$\mu(H^{19}|B_{1/3})\ \mu(B_{1/3}) /(\mu(H^{19}|B_{1/3})\ \mu(B_{1/3}) + \mu(H^{19}|B_{2/3})\ \mu(B_{2/3}) =$
$0.99\ (1/3)^{19}\ (2/3)^{6} / (0.99\ (1/3)^{19}\ (2/3)^{6} + 0.01\ (2/3)^{19}\ (1/3)^{6}) =$
$99 / (99 + 2^{13}) \approx 0.01$

Consequence: Although initially Alice gives $B_{1/3}$ probability 0.99 and $B_{2/3}$ probability 0.01, after seeing the evidence of 19 heads and 7 tails she gives $B_{2/3}$ probability 0.99.

3 Conditioning with sets of probabilities

A biased coin:

- If we know the prior α for the bias, then we should include the bias into the possible world:

$$W = \{(1/3, h), (1/3, t), (2/3, h), (2/3, t)\}$$
$$\mu\,(1/3, h) = \alpha\cdot 1/3,\ \mu\,(2/3, h) = (1-\alpha)\cdot 2/3,\ \ldots$$

- If we do not know the prior for the bias, then we should not include the bias into the possible world and use a set of probabilities instead:

$$W = \{h, t\}$$
$$\mu_{1/3}\,(h) = 1/3,\ \mu_{2/3}\,(h) = 2/3,\ \ldots$$
$$\mathcal{P} = \{\mu_{2/3}, \mu_{1/3}\}$$

Definition 2: Let $\mathcal{P}$ be a set of probability measures. Then the updating of $\mathcal{P}$ with an event (proposition) U is as follows:

$$\mathcal{P}|U =_{\text{def}} \{\mu|U \in \mathcal{P},\ \mu(U) > 0\}, \text{ i.e.}$$
$$\mathcal{P}(V|U) =_{\text{def}} \{\mu(\mathrm{V}|U),\ \mu\in\mathcal{P},\ \mu(U) > 0\}$$

Example: with two tosses of a biased coin:

$W = \{\text{hh, ht, th, tt}\}$

$\mathcal{P} = \{\mu_{1/3}, \mu_{2/3}\}$, where $\mu_\alpha(hh) = \mu_\alpha(H^1)\ \mu_\alpha(H^2) = \alpha^2$, and

$\mu_\alpha(ht) = \mu_\alpha(H^1)\ \mu_\alpha(T^2) = \alpha\cdot(1-\alpha)$ (independence)

$\mathcal{P}|H^1\ (H^2) = \mathcal{P}(H^2|H^1) = \{\mu_{1/3}(H^2|H^1),\ \mu_{2/3}(H^2|H^1)\} = \{1/3,\ 2/3\}$

$\mathcal{P}|H^1\ (T^2) = \mathcal{P}(T^2|H^1) = \{\mu_{1/3}(T^2|H^1),\ \mu_{2/3}(T^2|H^1)\} = \{2/3,\ 1/3\}$

The three prisoner puzzle: classical treatment

W = {(a,b), (a,c), (b,c), (c,b)} where (x,y) represents a world where prisoner x is pardoned and the guard says that y will be executed.

Principle of Indifference

lives-a =	{(a,b), (a,c)}	1/3
lives-b =	{(b,c)}	1/3
lives-c =	{(c,b)}	1/3

$\mu(\textit{says-b}|\textit{lives-a}) = \delta$, a *fixed* value in the interval [0, 1].

Update the probability of *lives-a* after getting the information that b will be executed.

The calculation

We have

- $\mu(\textit{lives-x}) = 1/3$ for $x = a, b, c$.
- $\mu(\textit{says-b}|\textit{lives-a}) = \delta$
- $\mu(\textit{says-b}|\textit{lives-c}) = 1$ and $\mu(\textit{says-b}|\textit{lives-b}) = 0$

We have to calculate $\mu(\textit{lives-a}|\textit{says-b})$

$\mu(\textit{lives-a}|\textit{says-b}) = \mu(\textit{says-b}|\textit{lives-a})\ \mu(\textit{lives-a}) / C$
$C = \mu(\textit{says-b}|\textit{lives-a})\ \mu(\textit{lives-a}) + \mu(\textit{says-b}|\textit{lives-c})\ \mu(\textit{lives-c})$

$\mu(\textit{lives-a}|\textit{says-b}) = 1/3\ \delta / (1/3\ \delta + 1/3) = \delta/(\delta+1)$
If $\delta = ½$, then $\mu(\textit{lives-a}|\textit{says-b}) = 1/3$,
If $\delta = 1$, then $\mu(\textit{lives-a}|\textit{says-b}) = ½$!

The three prisoner puzzle: sets of probabilities

Set of probabilities: $\mathcal{P} = \{\mu_\delta : \delta \in [0, 1]\}$

Calculate $\mathcal{P}(\textit{lives-a}|\textit{says-b})$ and show
that the lower probability $\mathcal{P}_*(\textit{lives-a}|\textit{says-b}) = 0$
and the upper probability $\mathcal{P}^*(\textit{lives-a}|\textit{says-b}) = ½$.

You can use the result from the last slide in the form
$\mu_\delta(\textit{lives-a}|\textit{says-b}) = \delta/(\delta+1)$

4 Conditioning inner and outer measures

Given a probability space $(W, \mathcal{F}, \mu)$, we can extend μ by considering all algebras $(W, 2^W, \mu')$ such that $(W, \mathcal{F}, \mu)$ is a subalgebra of the classical algebra $(W, 2^W, \mu')$. Consequently, we can use the extension set $\mathcal{P}_\mu$ (defined as set of all extensions of μ to the classical algebra) and we condition with respect to sets of probabilities:

- For $U, V \subseteq W$, $\mu_*(\mathrm{V}|\mathrm{U}) = (\mathcal{P}_\mu)_*(V|U) = \mathbf{min}\{\mu'(V|U)\text{: } \mu' \in \mathcal{P}_\mu\}$
- For $U, V \subseteq W$, $\mu^*(\mathrm{V}|\mathrm{U}) = (\mathcal{P}_\mu)^*(V|U) = \mathbf{max}\{\mu'(V|U)\text{: } \mu' \in \mathcal{P}_\mu\}$

Fact about conditioning inner and outer measure

Fact 4: Let $(W, \mathcal{F}, \mu)$ be a finite probability space and suppose that $\mu^*(U)>0$. Then:

- $\mu_*(V|U)=\mu_*(V\cap U)/(\mu_*(V\cap U)+\mu^*(\neg V\cap U))$ if $\mu^*(\neg V\cap U)>0$
  $= 1$ if $\mu^*(\neg V\cap U)=0$

- $\mu^*(V|U)=\mu^*(V\cap U)/(\mu^*(V\cap U)+\mu_*(\neg V\cap U))$ if $\mu^*(V\cap U)>0$
  $= 1$ if $\mu^*(V\cap U)=0$

Proof: Halpern p. 90!

$W = \{(a,b), (a,c), (b,c), (c,b)\}$

lives-a : $\mu\{(a,b), (a,c)\} = 1/3$

lives-b : $\mu\{(b,c)\} = 1/3$

lives-c : $\mu\{(c,b)\} = 1/3$

says-b = $\{(a,b), (c,b)\}$ (μ not defined for it)

Calculate: $\mu_*(\textit{lives-a}|\textit{says-b})$ and $\mu^*(\textit{lives-a}|\textit{says-b})$!

- $\mu_*(\textit{lives-a} \cap \textit{says-b}) = \mu_*\{(a,b)\} = 0$ (since the only element with defined μ that is contained in $\{(a,b)\}$ is the empty set! Therefore $\mu_*(\textit{lives-a}|\textit{says-b}) = 0$
- $\mu^*(\textit{lives-a} \cap \textit{says-b}) = \mu^*\{(a,b)\} = 1/3$ and $\mu_*(\neg\textit{lives-a} \cap \textit{says-b}) = \mu^*(\neg\textit{lives-a} \cap \textit{says-b}) = \mu\{(c,b)\}=1/3$.
Therefore $\mu^*(\textit{lives-a}|\textit{says-b}) = (1/3)/(1/3+1/3) = 1/2$

5 Conditioning possibility measures

There are two approaches in the literature for defining conditional possibility measures:

Definition 3:

$\text{Poss}(V \| U) = \text{Poss}(V \cap U)/\text{Poss}(U)$ (assuming $\text{Poss}(U) > 0$)

Definition 4:

$$\text{Poss}(V|U) = \begin{cases} \text{Poss}(V \cap U) & \text{if } \text{Poss}(V \cap U) < \text{Poss}(U) \\ 1 & \text{if } \text{Poss}(V \cap U) = \text{Poss}(U) \end{cases}$$

It is not difficult to show that both $\text{Poss}(\cdot \| U)$ and $\text{Poss}(\cdot | U)$ are possibility measures.

6 Conditioning ranking functions

As we have seen, ranking functions are very similar in spirit to possibility measures.

Definition 5:

$\kappa(V|U) =_{\text{def}} \kappa(V \cap U) - \kappa(U)$

Notice that there is an obvious analogue of Bayes' rule for ranking functions:

$\kappa(U|V) = \kappa(V|U) + \kappa(U) - \kappa(V)$