The Semantics of Evidential Reasoning

Riccardo Pucella

Semantics Seminar

November 29, 2006

The Motivating Problem

Suppose Alice has two coins

- F Fair (Pr heads = 1/2)
- DH Double-headed (Pr heads = 1)

She chooses one of the coins and tosses it 100 times

- It lands heads every time

What is the probability that the coin is double-headed?

- No probability on choice!
- But still more "likely" to be the double-headed coin

Evidence

The outcome of an experiment (Heads, Tails) provides evidence for an hypothesis (Fair, Double-Headed)

Many approaches to encoding evidence are based on comparing likelihood functions

$\mu_h(ob)$ = probability of observing ob if h holds

- μ_{DH} (Heads) = 1
- μ_F (Heads) = 1/2

Shafer's Approach

Evidence space E = (H,O,**μ**):

- H: hypotheses $\{h_1, ..., h_n\}$
- O: observations
- **μ**: H-indexed likelihood functions

Weight of evidence of observation ob for hypothesis h

$$w_E(ob, h) = \frac{\mu_h(ob)}{\mu_{h_1}(ob) + \cdots + \mu_{h_n}(ob)}$$

Revisiting The Puzzles

E = ({ Fair, DH}, {100Heads,<100Heads}, **μ**)

μ	Fair	DH
100 Heads	2^{-100}	1
< 100 Heads	$1-2^{-100}$	0

w_E	Fair	DH
100 Heads	~ 0	~ 1
< 100 Heads	1	0

Justifying Evidence

Consider the motivating problem

Suppose we have a probability on Alice's choice

- Then probability of coin being DH after seeing Heads can be computed by Bayes' Theorem:

$$\Pr(DH \mid Heads) = \frac{\Pr(Heads \mid DH)\Pr(DH)}{\Pr(Heads)}$$

Evidence captures the information used in this update...

Updating Via Evidence

Can use weight of evidence function $\lambda x.w_E(ob,x)$ to update prior μ_0 to a posterior μ_{ob}:

$$\mu_{ob} = \mu_o \oplus (\lambda x.w_E(ob,x))$$

$$(f_1 \oplus f_2)(x) = \frac{f_1(x)f_2(x)}{\sum_{y \in X} f_1(y)f_2(y)}$$

Theorem: If $Pr(h) = \mu_0(h)$, then $\mu_{ob}(h) = Pr(h \mid ob)$

Application:

Rabin's primality test for a number n

- Polynomial-time computable predicate P(n,a):
  - n composite ⇒ P(n,a)=1 for ≥ n/2 choices of a
  - n prime ⇒ P(n,a)=0 for all a

Pick a number a at random between 0 and n
If P(n,a) = 1 return "composite"; otherwise return "prime"

w	prime	composite
"prime"	≥2/3	≤1/3
"composite"	0	1

A More Complex Problem

Suppose Alice has two coins

- B Biased towards heads (Pr heads = 3/4)
- DH Double-headed (Pr heads = 1)

Suppose Bob has a fair coin

- Alice and Bob each hand a coin to Zoe
- Zoe tosses one of the coins
- Outcome is heads

What is the evidence that Zoe chose Alice's coin?

One Approach

Consider all possible evidence spaces consistent with the information

- Agent ascribing weight of evidence does not know which is the right one
- Compute weights of evidence with respect to each evidence space
- Get a range of possible weights of evidence

Generalized Weight of Evidence

Generalized evidence space $G = (H,O,\Delta)$

- H: hypotheses
- O: observations
- Δ: set of indexed likelihood functions
  - E.g. $\{\boldsymbol{\mu}_{DH}, \boldsymbol{\mu}_{B}\}$

$$S(G) = \{ (H,O,\boldsymbol{\mu}) \mid \boldsymbol{\mu} \in \Delta \}$$

$$w_G = \{ w_E \mid E \in S(G) \}$$

Can reason with upper and lower weight of evidence

Updating Priors

Can update prior μ_0 by natural generalization

$$\mathcal{P}_{ob} = \{\mu_0 \oplus w(ob, -) \mid w \in w_G\}$$

This gives a set of all possible posteriors

Theorem: Bounds on posteriors agree with taking all possible results of Bayes' Theorem applied to the prior

What about Repeated Tosses?

Consider the Alice/Bob/Zoe scenario with 100 reps

- Zoe always makes same choice
- Difference is whether Alice gets to choose again
  - Same coin repeated?
  - Different coin chosen at every toss?

The First Experiment

Given hyp $\in$ {AliceCoin,BobCoin}

Alice chooses [DoubleHeaded, Biased]
Bob chooses Fair
if hyp is AliceCoin then
  Zoe chooses Alice's coin
else
  Zoe chooses Bob's coin
Zoe tosses coin 100 times

The Second Experiment

Given hyp ∈ {AliceCoin,BobCoin}

repeat 100 times:
  Alice chooses [DoubleHeaded, Biased]
  Bob chooses Fair
  if hyp is AliceCoin then
    Zoe chooses Alice's coin
  else
    Zoe chooses Bob's coin
  Zoe tosses coin

The Second Experiment

hyp ∈ {AliceCoin,BobCoin}

```
repeat 100 times:
  A := choose [DoubleHeaded, Biased]
  B := Fair
  if hyp is AliceCoin then
    Z := A
  else
    Z := B
  observe Z
```

The Second Experiment

$hyp \in \{1,2\}$

```
i := 0;
while (i < 100) do
  A := choose [DoubleHeaded, Biased];
  B := Fair;
  if (hyp=1) then Z := A else Z := B;
  observe Z;
  i := i + 1
```

The Second Experiment

hyp ∈ {1,2}

i := 0;
while (i < 100) **do**
  A := **choose** [{H:1,T:0} , {H:0.75,T:0.25}];
  B := {H:0.5,T:0.5} ;
  if (hyp=1) **then** Z := A **else** Z := B;
  observe Z;
  i := i + 1

A Language for Experiments

S ::=

$X := P$

$x := E$

observe X

$S_1;S_2$

if B **then** S_1 **else** S_2

while B **do** S

P ::=

$\{ob_1{:}r_1, \ldots, ob_n{:}r_n\}$

X

choose $[P_1, \ldots, P_n]$

B ::= Boolean expressions

E ::= Integer expressions

A program in this language is of the form:

$$hyp \in \{1,\ldots,n\}\ S$$

hyp is “unassignable”

Evidence-Based Semantics

Notation

- $\Pi(A)$: probabilities distributions on set A
- Σ : maps from variables to Int
- Γ : maps from observables to $\Pi(O)$

$[[\ B\ ]] : \Sigma \rightarrow \{\text{true},\text{false}\}$

$[[\ E\ ]] : \Sigma \rightarrow \text{Int}$

$[[\ P\ ]] : \Gamma \rightarrow \wp\left(\Pi(O)\right)$

$[[\ S\ ]] : \text{State} \rightarrow \wp(\text{State})$

where $\text{State} = \Sigma \times \Gamma \times \Pi(O^*)$

Evidence-Based Semantics (ctd)

$[[\ \{ob_1{:}r_1,\ldots,ob_n{:}r_n\}\ ]](\gamma) = \{\mu\}$ where $\mu(ob_i)=r_i$

$[[\ X\ ]](\gamma) = \{\gamma(X)\}$

$[[\ \textbf{choose}\ [P_1,\ldots,P_n]\ ]](\gamma) = \cup[[P_i]](\gamma)$

$[[\ X := P\ ]](\sigma,\gamma,\mu) = \{(\sigma,\gamma[X\mapsto v],\mu) \mid v\in[[P]](\gamma)\}$

$[[\ x := E\ ]](\sigma,\gamma,\mu) = \{(\sigma[x\mapsto[[E]](\sigma),\gamma,\mu)\}$

$[[\ \textbf{observe}\ X\ ]](\sigma,\gamma,\mu) = \{(\sigma,\gamma,\mu') \mid \mu'(obs;ob)=\mu(obs)\gamma(X)(ob)\}$

$[[\ S_1;S_2\ ]](\sigma,\gamma,\mu) = \cup\{\ [[S_2]](\sigma',\gamma',\mu') \mid (\sigma',\gamma',\mu')\in[[S_1]](\sigma,\gamma,\mu)\ \}$

$[[\ \textbf{if}\ B\ \textbf{then}\ S_1\ \textbf{else}\ S_2\ ]](\sigma,\gamma,\mu) =$

$\quad [[S_1]](\sigma,\gamma,\mu)$ if $[[B]](\sigma) = \text{true}$

$\quad [[S_2]](\sigma,\gamma,\mu)$ if $[[B]](\sigma) = \text{false}$

$[[\ \textbf{while}\ B\ \textbf{do}\ S\ ]](\sigma,\gamma,\mu) =$ ... usual definition ...

Evidence-Based Semantics (ctd)

Given an experiment hyp $\in$ {1,...,n} S:

$$E_S = (H, O^*, \Delta)$$

where $\boldsymbol{\mu} \in \Delta$ iff for all $h \in \{1,...,n\}$, there exists σ,γ s.t.

- $\sigma(hyp) = h$
- $(\sigma, \gamma, \boldsymbol{\mu}(h)) \in [[S]](\sigma_0[hyp \mapsto h], \gamma_0, \mu_0)$

and

- σ_0 : sets every variable to 0
- γ_0 : sets every observable to uniform probability
- μ_0 : $\mu_0(\langle\rangle) = 1$

Relationship to Probabilistic Semantics

Can we justify this semantics?

- Recall that evidence can be used to update a prior probability on hypotheses to a posterior probability on hypotheses, based on seeing a sequenece of observations

Let S be a deterministic experiment

- No occurrence of choose [...]

Standard probabilistic semantics for language:

- State = $\Sigma \times \Gamma \times O^*$
- P[[S]](s,A) : probability that program S starting in state s terminates in a state in A

The Correspondence

Let S be a deterministic experiment
Let μ be a prior probability distribution on hypotheses

The following probabilities are equal:

- $\mu \oplus w_{Es}$ (obs,-)
- $\lambda h.(\mu(h)\ P[[S]](\sigma_0[hyp \mapsto h], \Sigma \times \Gamma \times \{obs\}))$

What if S is not deterministic?

Conjecture: the above result holds if you extend the probabilistic semantics with nondeterminism