A Dynamic Logic of Subjective Belief

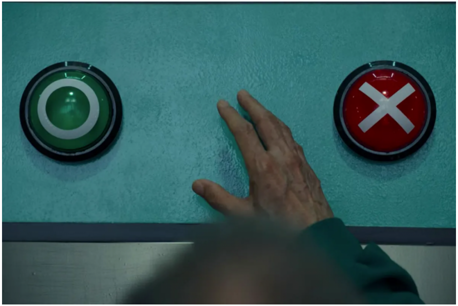

Tim French

The University of Western Australia

2024

Motivation

Elections, Conference Reviewing, impulse purchases, holiday destinations are subjective decisions, that are often informed by minimal experience.

This work provides a formal theoretical framework for examining how these minimal experiences impact decisions and understanding.

This framework is probabilistic system based on aleatoric logic.

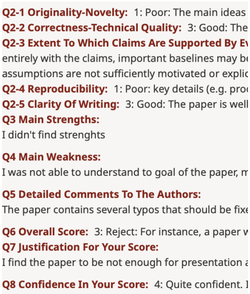

Q2-1 Originality-Novelty: 1: Poor: The main ideas
Q2-2 Correctness-Technical Quality: 3: Good: The
Q2-3 Extent To Which Claims Are Supported By Ev
entirely with the claims, important baselines may be
assumptions are not sufficiently motivated or explic
Q2-4 Reproducibility: 1: Poor: key details (e.g. pro
Q2-5 Clarity Of Writing: 3: Good: The paper is well
Q3 Main Strengths:
I didn't find strenghts

Q4 Main Weakness:
I was not able to understand to goal of the paper, m

Q5 Detailed Comments To The Authors:
The paper contains several typos that should be fixe

Q6 Overall Score: 3: Reject: For instance, a paper w
Q7 Justification For Your Score:
I find the paper to be not enough for presentation a

Q8 Confidence In Your Score: 4: Quite confident. I

Conference Submission

Suppose that you have written a paper and you can send it to the Mathematical Computing Conference (MCC) or the Computational Mathematics Conference (CMC).

$\mathfrak{MCM}$

You know that the program committee for MCC prefers theoretical contributions, while the program committee for CMC is more likely to accept empirical works.

CMC

You judge that your paper is more theoretical than empirical, but you also believe that you are likely to need two out of three MCC reviewers to like your paper for it to be accepted, whilst one out of two CMC reviewers liking your paper is usually enough.

A Logic for Subjective Belief

A basic elements of a language for describing subjective belief consist of:

- **Variables**, $x, y, \ldots \in \mathcal{V}$. These are labels for hypothetical *things*.
- **Predicates**, $p, q, \ldots \in \mathcal{P}$. These give properties to the the domain elements.
- **Atomic propositions**, $X, Y, \ldots \in \mathcal{A}$. The are probabilistic propositions, that reflect an element of chance.
- **Operators** to build complex hypothetical scenarios from the atoms, including an **Expectation operator**.

Propositions describe hypothetical stories that can be assigned an aleatoric probability.

I imagine a reviewer from CMC and I imagine a review they gave for my paper was not favourable.

Language and Syntax

The operators a consist of conjunction, and negation, as well as an *aleatoric or* operation, a *fixed point* operation and and *expectation* operation.

$$\alpha \ ::= \ \top \mid X \mid \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \bowtie \alpha \mid \mathbb{E}x.\alpha \mid \mathbb{F}X.\alpha^{*}$$

- Conjunction $\wedge$ and Negation $\neg$ are understood as usual.
- Aleatoric or, $\bowtie$ is a binary operator, understood as a random choice to evaluate one operand or the other.
- The fixed point operator, $\mathbb{F}X.$ is standard, in that it can be understood as replacing X with the fixed point formula.
- The expectation operation, $\mathbb{E}x.$ is understood as evaluating the formula with respect to an element that is randomly sampled from the domain.

*In $\mathbb{F}X.\alpha$, we require $\forall \beta_1 \wedge \beta_2 \subseteq \alpha$, X does not appear in both β_1 and β_2, and $\forall\ \mathbb{F}Y.\beta \subset \alpha$, X does not appear free in β

Semantic Interpretations

Definition

An *aleatoric interpretation* is given by the tuple:
$\mathcal{I} = (\mathcal{D}, \Sigma, \mu, \chi, \nu)$ where:

- $(\mathcal{D}, \Sigma, \mu)$ is a probability space.
- $\chi \in [0,1]^{\mathcal{A}}$ assigns a probability to each atomic proposition.
- $\nu \in \wp(\mathcal{D}^*)^{\mathcal{P}}$ assigns each predicate $P \in \mathcal{P}$ to a set of tuples over the $\mathcal{D}$.

Motivation for how experiences, imagined and real, inform predictions.

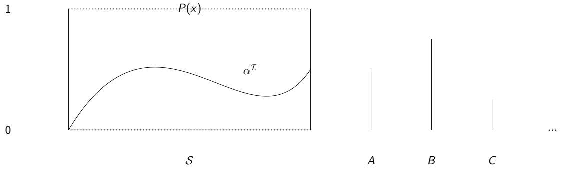

Semantic Definitions

The evaluation of a proposition is its probability of being satisfied by a random sampling.

Definition

Given some interpretation, $\mathcal{I} = (\mathcal{D}, \Sigma, \mu, \chi, \nu)$ and some $\alpha \in \mathcal{L}$, the *likelihood of* α *in* $\mathcal{I}$ is a function $\alpha^{\mathcal{I}} : \mathcal{D}^{\mathcal{V}} \longrightarrow [0,1]$, specified inductively as follows. Given some assignment $a \in \mathcal{D}^{\mathcal{V}}$:

$$
\begin{aligned}
X^{\mathcal{I}} &= \chi(X) \\
(P(\overline{x}))^{\mathcal{I}}(a) &= 1 \text{ if } (a(x_1), \ldots, a(x_{P\#})) \in \nu(P), \text{otherwise } 0 \\
(\neg\alpha)^{\mathcal{I}} &= 1 - \alpha^{\mathcal{I}} \\
\alpha \wedge \beta &= \alpha^{\mathcal{I}} \cdot \beta^{\mathcal{I}} \\
\alpha \bowtie \beta &= (\alpha^{\mathcal{I}} + \beta^{\mathcal{I}})/2 \\
(\mathbb{E}x.\alpha)^{\mathcal{I}} &= \int_{\mathcal{D}} \alpha^{\mathcal{I}} d\mu(x) \\
(\mathbb{F}X.\alpha)^{\mathcal{I}} &= \left\{ \begin{array}{l} 0.5 \text{ if } (\alpha^{\mathcal{I}})^{[X:0.5]} = 0.5, \text{ and} \\ p \text{ where } (\alpha^{\mathcal{I}})^{[X:p]} = p \text{ otherwise.} \end{array} \right\}.
\end{aligned}
$$

Example

I imagine a reviewer from CMC and I imagine a review they wrote for my paper was not favourable.

$$\mathbb{F}X.\mathbb{E}pc\mathbb{E}r\left[\begin{array}{c} (CMC(pc)\wedge wrote(pc,r)\wedge bad(r)\vee X \\ \bowtie \\ \neg(CMC(pc)\wedge wrote(pc,r)\wedge bad(r))\wedge X \end{array}\right]$$

This can be seen as a story, or *program*, of hypotheticals: :

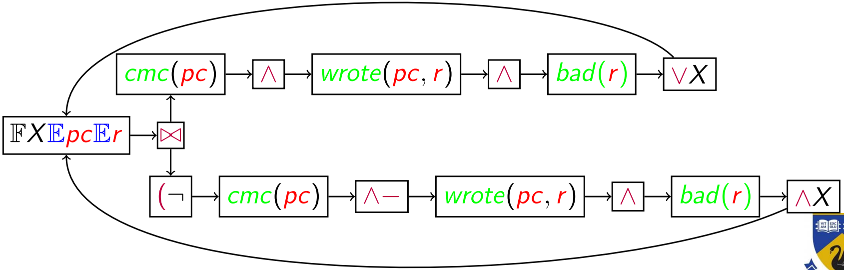

Intermezzo: Some Expressive Tricks

If-then-else

$$(\alpha \; ? \; \beta : \gamma) \equiv \mathbb{F}X. \left[\begin{array}{c} ((\alpha \wedge \beta) \vee X) \bowtie ((\alpha \rightarrow \beta) \wedge X) \\ \bowtie \\ ((\neg\alpha \wedge \gamma) \vee X) \bowtie ((\neg\alpha \rightarrow \gamma) \wedge X) \end{array} \right]$$

Existential Quantification

$$\begin{array}{rcl} \uparrow x\alpha & = & \mathbb{F}X.[\mathbb{E}x.(\alpha \vee X)] \\ \downarrow x\alpha & = & \mathbb{F}X.[\uparrow x \wedge X] \\ \exists x\alpha & = & \mathbb{F}X.(\uparrow x\alpha \wedge X) \bowtie (\downarrow x \vee X) \end{array}$$

Conditioning

$$(\alpha|\beta)_x = \mathbb{F}X.\mathbb{E}x.(\beta \; ? \; \alpha : X)$$

False

$$\perp = \mathbb{F}Y.(\mathbb{F}X.X \wedge Y)$$

Counting

$$\alpha^{\frac{0}{n}} = \perp$$
$$\alpha^{\frac{1}{1}} = \alpha$$
$$\alpha^{\frac{n}{m}} = (\alpha \; ? \; \alpha^{\frac{n-1}{m}} : \alpha^{\frac{n}{m-1}})$$

Choosing a Conference

We may suppose that there are five different "archetypes" of reviewers, and $\chi(M\,\mathrm{deg}) = 0.8$ and $\chi(C\,\mathrm{deg}) = 0.9$, and you need $2/3$ positive reviews for MCM and $1/2$ positive reviews for CMC.

Table: The reviewer archetypes

Type	*MPC*	*CPC*	*Like*	μ
Programmer	Yes	Yes	Yes	0.2
Statistician	No	Yes	No	0.3
Algebrist	Yes	No	Yes	0.3
Logician	No	Yes	Yes	0.1
Engineer	Yes	No	No	0.1

$$\mathrm{M} = ((\mathbb{F}X.\mathbb{E}x.(MPC(x)\ ?\ Like(x) : X))^{2/3}\ ?\ M\,\mathrm{deg} : \bot)$$
$$\mathrm{C} = ((\mathbb{F}X.\mathbb{E}x.(CPC(x)\ ?\ Like(x) : X))^{1/2}\ ?\ C\,\mathrm{deg} : \bot).$$

We find that $\mathrm{M}^{\mathcal{I}} = 0.4$ and $\mathrm{C}^{\mathcal{I}} = 0.459$.

Dynamics and Experiential Learning

This motivates the model as a state of mind, but how is that state achieved.

The agent focuses on some property, ϕ, and then makes an observation, β.

Then they consider what β tell them about ϕ; is it more or less likely?

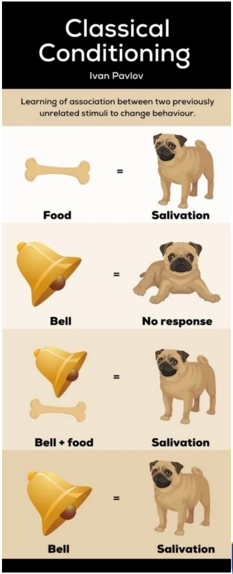

A Dynamic Operation

We would like to define an operation that transforms a belief model in response to observation.

The syntax is $[\alpha|\beta]_{\overline{x}}\phi$ where

- α is the proposition to be learnt, containing free variables $\overline{x}$
- β is the observation being learnt, containing free variables $\overline{x}$

Note α and β are often the same proposition, but may be different: learn $\mathit{infectious}(x)$ from $\mathit{sneezing}(x)$.

The observation is new information, like an announcement, but it is treated as evidence for a random event.

Sampling is an independent event, so conditioning doesn't work.

A game...

Suppose I have an urn of numbered marbles, and you win if you randomly draw one that is prime. **or...** we could flip a coin. If it is heads,

you need to draw two prime marbles in a row to win

but if it is tails, you get two chances to draw a prime marble. Which do you choose?

Sympathisers and Sceptics

If we want to learn α from observation(s), we can imagine filling two urns from the original urn:

1. To fill the α-sceptic urn, we draw a marble from the original urn, test $\neg\alpha$ for that marble. If it passes we add it to the sceptic urn and otherwise, we return it to the original urn and place a randomly drawn marble in the α-sceptic urn.
2. To fill the α-sympathetic urn, we draw a marble from the original urn, test α for that marble. If it passes we add it to the sympathetic urn and otherwise, we return it to the original urn and place a randomly drawn marble in the α-sympathetic urn.

Then α is more likely in one urn than the other, but the average distribution in the same as the original urn.

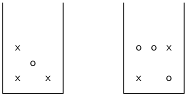

Sympathiser   Sceptic

Sceptics and Sympathisers: formalisation

Given an interpretation $\mathcal{I} = (\mathcal{D}, \Sigma, \mu, \chi, \nu))$ and $\sigma \in \Sigma$:

Definition

We let the *relativisation of* $\mathcal{I}$ *to* σ be the interpretation

$$\mathcal{I}^{\sigma} = (\sigma, \{\sigma' \in \Sigma | \ \sigma' \subseteq \sigma\}, \mu^{\sigma}, \chi, \nu)$$

where for all $\sigma' \subseteq \sigma$, $\mu^{\sigma}(\sigma') = \mu(\sigma')/\mu(\sigma)$

Definition

Given some property $\beta(x)$ to be learnt, the β-splitting is the pair (μ_ℓ, μ_r) where:

$$\begin{aligned} \mu_\ell(\sigma) &= (1 + (\mathbb{E}x.\beta)^{\mathcal{I}^{\sigma}} - (\mathbb{E}x.\beta)^{\mathcal{I}}) \cdot \mu(\sigma), \textit{and} \\ \mu_r(\sigma) &= (1 - (\mathbb{E}x.\beta)^{\mathcal{I}^{\sigma}} + (\mathbb{E}x.\beta)^{\mathcal{I}}) \cdot \mu(\sigma) \end{aligned}$$

The Observation Operator

Definition

Given an interpretation $\mathcal{I} = (\mathcal{D}, \Sigma, \mu, \chi, \nu)$ and some observation α, and some property β the α-β-*update* of $\mathcal{I}$ is the model $\mathcal{I}^\alpha_\beta = (\mathcal{D}, \Sigma, \mu^\alpha_\beta, \chi, \nu)$ where for all $\sigma \in \Sigma$

$$\mu^\alpha_\beta(\sigma) = \frac{(\mathbb{E}x.\alpha)^{\mathcal{I}^\ell} \cdot \mu_\ell(\sigma) + (\mathbb{E}x.\alpha)^{\mathcal{I}^r} \cdot \mu_r(\sigma)}{2 \cdot (\mathbb{E}x.\alpha)^{\mathcal{I}}}$$

where $\mathcal{I}^\ell = (\mathcal{D}, \Sigma, \mu^\ell, \chi, \nu)$, $\mathcal{I}^r = (\mathcal{D}, \Sigma, \mu^r, \chi, \nu)$, and (μ^ℓ, μ^r) is the β-splitting of $\mathcal{I}$.

We define $([\beta|\alpha]\gamma)^{\mathcal{I}} = \gamma^{\mathcal{I}^\alpha_\beta}$.

Example

Suppose we get the reviews back from CMC, and they're positive. We wonder if maybe this means logicians are more frequent in the CMC program committee

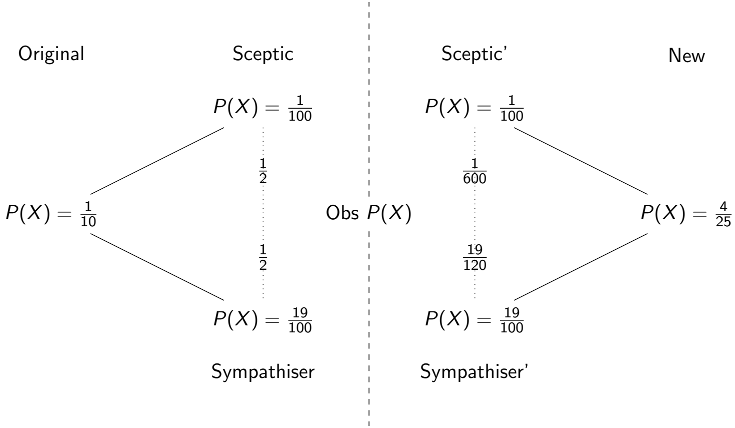

Syntactic Representation

We can show that the dynamic operator is expressible in the original language.

Definition

Given an observation of $\alpha(x)$ when learning β, the conditioning of ϕ by $\alpha(x)$ is the proposition:

$$\phi_\beta^\alpha = \phi[\mathbb{E}x.\gamma(x)\backslash\mathbb{F}X.(1/2 \ ? \ (\alpha_\beta^\ell \ ? \ \gamma_\beta^\ell : X) : (\alpha_\beta^r \ ? \ \gamma_\beta^r : X))]$$

where $\delta_\beta^r = \delta[\mathbb{E}x.\gamma\backslash\mathbb{E}x.(\beta \ ? \ \mathbb{E}x.\gamma : \gamma)]$ and $\delta_\beta^\ell = \delta[\mathbb{E}x.\gamma\backslash\mathbb{E}x.(\beta \ ? \ \gamma : \mathbb{E}x.\gamma)]$, for a given $\delta \in \mathcal{L}$.

Confidence

Could there be a better prediction for α?

The urn game is arbitrary, and there are many other ways to split an urn.

We can use a **confidence** to decide how many turns an agent should have to separate the sceptic/sympathiser urns.

By comparing predictions to frequencies of observations confidence can be learnt!

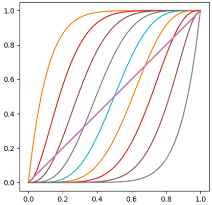

Conclusion

We have presented

- A first order aleatoric logic for representing subjective beliefs.
- Several expressivity results.
- A learning process, where an agents can condition beliefs by observations.
- A learning operator implementing this process.

Future work is to

- Develop the notion of confidence in learning.
- Give a complete axiomatization.

Thank you.