An Aleatoric Description Logic for Probabilistic Reasoning

Tim French, Thomas Smoker

The University of Western Australia

tim.french@uwa.edu.au, thomas.smoker@research.uwa.edu.au

Aleatoric Description Logic

A logic for representing *uncertainty* and *belief* where **concepts are coin tosses** and **roles are random samples**.

Introduction

Aleatoric Description Logics model uncertainty in the world as *aleatoric events*, i.e. by the roll of dice: concepts correspond to probabilistic events and roles are probability distributions of individuals. Reasoning in aleatoric description logic is akin to sampling these events and distributions to create an *imagined scenario*. Properties are not true or false, but rather have a probability of being true in a sampled scenario.

- We provide a syntax and semantics for aleatoric description logic that generalises $\mathcal{ALC}$, and corresponds to a probability space of DL interpretations
- We give a belief base formalism that can express axioms constraining the relative likelihood of concepts and determine whether the axioms are satisfiable
- We show how learning from observation can be formalised as marginalisation over possible individuals.

Syntax and Semantics

Syntax

The syntax for ADL describes probabilistic events, rather than truths and is necessarily different to standard DL syntax. The syntax is specified with respect to a set of *atomic concepts*, X and a set of *roles*, R:

$$\alpha \ ::= \ \top \mid \bot \mid A \mid (\alpha?\alpha{:}\alpha) \mid [\rho]\,(\alpha|\alpha)$$

where $A \in X$ is an atomic concept, and $\rho \in R$.

These operators describe events with the following intuitions: $\top$ is *always*; $\bot$ is *never*; A is some named concept that may hold for an individual; $(\alpha?\beta{:}\gamma)$ is *if* α *then* β, *else* γ; and $[\rho]\,(\alpha|\beta)$ is ρ *is* α *given* β.

There is a special role $\mathsf{id} \in R$ referred to as *identity*, which refers to different possibilities for the one individual,

Probabilistic Semantics

Definition

Aleatoric Belief Model An *aleatoric belief model* is specified by the tuple $\mathcal{B} = (I, r, \ell)$, where:

- I is a set of possible individuals.
- $r : R \times I \longrightarrow PD(I)$ assigns for each role and each individual a probability distribution over I (write $\rho(i,j)$ for $r(\rho,i)(j)$).
- $\ell : I \times X \longrightarrow [0,1]$ gives the likelihood of an individual i satisfying A, (write $A(i)$ for $\ell(A,i)$).

Definition

Semantics The probability $\mathcal{B} = (I, r, \ell)$ assigns α at i, $\mathcal{B}_i(\alpha)$, is:

$$\mathcal{B}_i(\bot) = 0 \qquad \mathcal{B}_i(\top) = 1 \qquad \mathcal{B}_i(A) = A(i)$$
$$\mathcal{B}_i((\alpha?\beta{:}\gamma)) = \mathcal{B}_i(\alpha).\mathcal{B}_i(\beta) + (1-\mathcal{B}_i(\alpha)).\mathcal{B}_i(\gamma)$$
$$\mathcal{B}_i([\rho]\,(\alpha|\beta)) = \frac{\sum_{j\in I}\rho(i,j)\mathcal{B}_j(\alpha)\mathcal{B}_j(\beta)}{E_i^\rho\beta} \ \text{if } E_i^\rho\beta > 0$$
$$\mathcal{B}_i([\rho]\,(\alpha|\beta)) = 1, \ \text{if } E_i^\rho\beta = 0$$

where $E_i^\rho\alpha = \sum_{j\in I}\rho(i,j)\mathcal{B}_j(\alpha)$, where $\rho \in R$.

Common Description Logic Operators

Table 1: Some abbreviations of operators in ADL.

term	formula	interpretation
$\alpha \sqcap \beta$	$(\alpha?\beta{:}\bot)$	$\mathcal{B}_i(\alpha).\mathcal{B}_i(\beta)$
$\alpha \sqcup \beta$	$(\alpha?\top{:}\beta)$	$\mathcal{B}_i(\alpha)+\mathcal{B}_i(\beta)-\mathcal{B}_i(\alpha).\mathcal{B}_i(\beta)$
$\neg\alpha$	$(\alpha?\bot{:}\top)$	$1-\mathcal{B}_i(\alpha)$
$\alpha \Rightarrow \beta$	$(\alpha?\beta{:}\top)$	$1-\mathcal{B}_i(\alpha)+\mathcal{B}_i(\alpha).\mathcal{B}_i(\beta)$
$E_\rho\alpha$	$[\rho]\,(\alpha|\top)$	$\sum_{j\in I}\rho(i,j).\mathcal{B}_j(\alpha)$
$\exists\rho.\alpha$	$\neg[\rho]\,(\bot|\alpha)$	1 if $E_\rho\alpha \neq 0$, 0 otherwise.

Example

Suppose three agents, Hector, Igor and Julia, who may or may not have a virus (*V*) or a fever (*F*). Each agent is represented by two possible individuals, one with the virus and one without. Each agent will occasionally come into contact with another agent, described by the role contact. A graphical representation is given in Figure 1.

Figure 1: A graphical example of the virus transmission scenario. There are two *possible individuals* for each of Hector, Igor and Julia: one with the virus and one without.

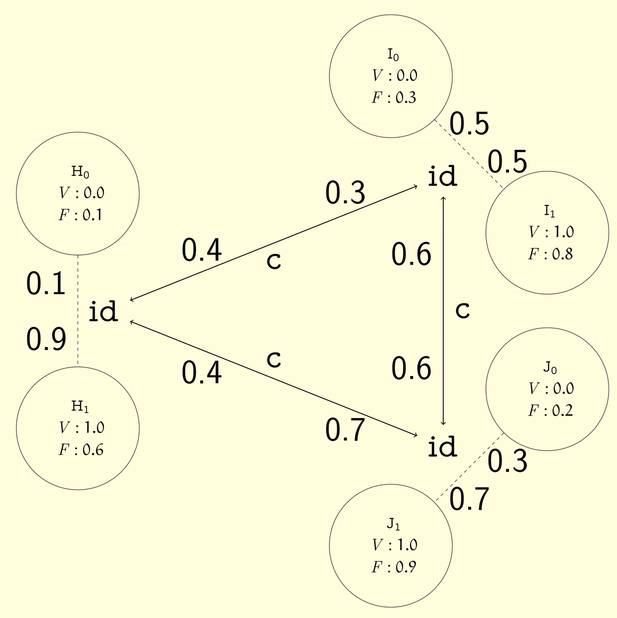

The probability of an agent being newly exposed to the virus, given that they went out and encountered someone with a fever, as:

$$\mathsf{exp} = E(\neg V \sqcap [c]\,(V|F))$$

Interpreting this for Hector, we see the probability Hector was newly exposed to the virus is approximately 0.7.

Aleatoric Knowledge Bases

Definition

The *aleatoric terminological axioms* have the form:

$$\alpha \preceq \beta \ (\alpha \text{ is no more likely than } \beta)$$

A *T-Book*, $\mathcal{T}$, is a set of aleatoric terminological axioms. The *aleatoric assertional axiom* have the form: $i :^p \alpha$, where i is an individual, $p \in [0,1]$ asserts that individual a satisfies concept α, with probability p. An *A-Book*, $\mathcal{A}$, is a set of aleatoric assertional axioms.

An *aleatoric knowledge base* is the pair $\mathcal{K} = (\mathcal{A}, \mathcal{T})$.

$\mathcal{B}$ *satisfies* $\mathcal{K} = (\mathcal{A}, \mathcal{T})$ iff it satisfies all the axioms in $\mathcal{A}$ and $\mathcal{T}$, so $\alpha \preceq \beta \in \mathcal{T}$ implies $\forall i$, $\mathcal{B}_i(\alpha) \leq \mathcal{B}_i(\beta)$ and $i :^p \alpha$ implies $\mathcal{B}_i(\alpha) = p$.

Results

The semantics compactly represent reasoning over a probability space of interpretations:

Theorem

Given a formula α of $\mathcal{ALC}$ and an aleatoric belief model $\mathcal{B}_i$, there exists: a formula α^* that is logically equivalent to α in $\mathcal{ALC}$; and probability space $(\Omega^{\mathcal{B}_i}, \mathcal{F}, \mathcal{P}^{\mathcal{B}_i})$ derived from $\mathcal{B}_i$ such that $\mathcal{P}^{\mathcal{B}_i}(\hat{\alpha}) = \mathcal{B}_i(\alpha^*)$.

Model checking is efficient:

Theorem

Given a pointed belief model $\mathcal{B}_i$ consisting of n possible individuals, and a formula α consisting of m symbols, the value $\mathcal{B}_i(\alpha)$ can be computed in time $O(n^2m)$.

Consistency checking for acyclic T-Books is decidable:

Theorem

Given a simple aleatoric knowledge base $\mathcal{K} = (\mathcal{A}, \mathcal{T})$ where $\mathcal{T}$ is acyclic, (so no concept appears in both the head and the tail of a chain of axioms) it is possible to determine if $\mathcal{K}$ is consistent with complexity PSPACE.

The process for the satisfiability theorem is to build a system of polynomial inequalities corresponding to the axioms in $\mathcal{K}$.

Learning from observations

Given an aleatoric belief model, if an event is observed the probability distributions can be updated via marginalisation:

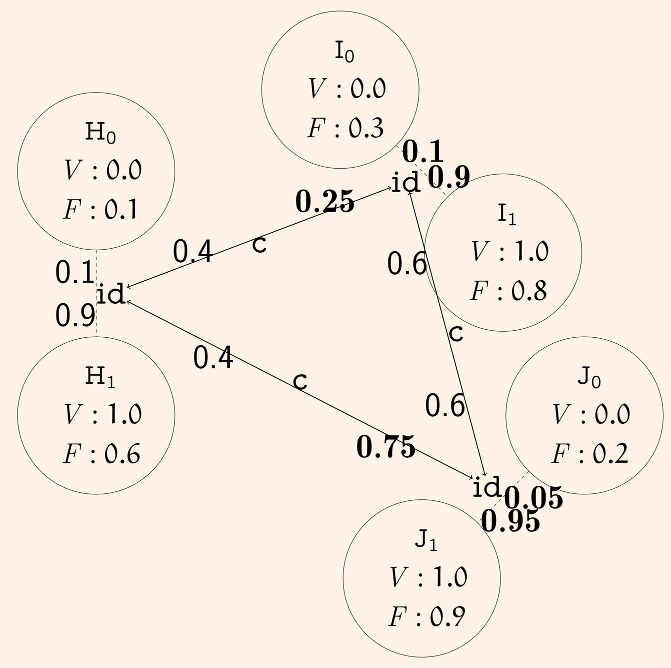

Figure 2: The update of the aleatoric belief model in Figure 1, after Hector is told a contact has tested positive for the virus. The updated values are bold.

Future work will examine the complexity of the satisfiability problem for non-acyclic T-Books, and investigate implementing a reasoning system for aleatoric description logics.

Find out More

SCAN ME