Model-Based Defeasible Reasoning

Twiddle Project

Overview

Humans use different but valid **patterns of reasoning** to draw conclusions from the same **knowledge base**. **Knowledge representation and reasoning** (KRR) encodes knowledge and formalises philosophical reasoning patterns, supporting ‘top-down’ forms of **artificial intelligence**.

Problem

Well-known patterns of reasoning are defined in two ways: **syntactically** (formula-based) and **semantically** (model-based). While there are formula-based algorithms for drawing conclusions, model-based **entailment** algorithms have not been explored.

Project Aims

We will investigate two reasoning formalisms: **rational closure** and **lexicographic closure**. Using model-based definitions, we aim to define **model representations**, **construction algorithms** and **entailment algorithms** for checking conclusions. We aim to **implement** and **compare** these algorithms with existing approaches.

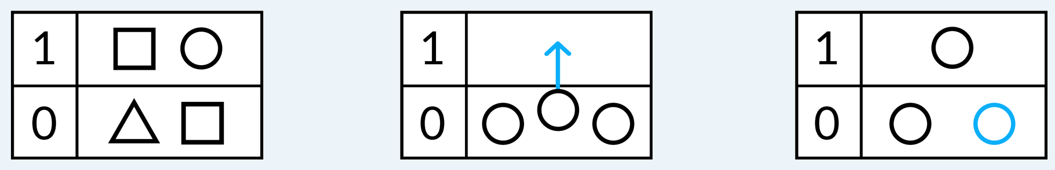

Representation *Construction* *Entailment*

KRR

KRR is concerned with representing information about the world through formalisms like mathematical logic. We focus on a foundational logic referred to as **classical propositional logic**.

Using propositional logic, we can encode the knowledge that mammals give birth to live young, platypuses are mammals and platypuses do not give birth to live young, in a **knowledge base**.

We can then conclude using **classical entailment** that ‘platypuses give birth to live young’.

$$K = \{m \rightarrow l,\ p \rightarrow m,\ p \rightarrow \neg l\}$$

$$K \models p \rightarrow l$$

Defeasible Reasoning

Propositional logic cannot express **typicality** where implications are usually true but may have exceptions. Conclusions cannot be retracted even with the addition of new conflicting knowledge.

Defeasible extensions of propositional logic and classical entailment address these shortcomings.

We can now reformulate our knowledge to produce more reasonable conclusions.

$$K = \{m \mathrel{|\!\sim} l,\ p \rightarrow m,\ p \mathrel{|\!\sim} \neg l\}$$

$$K \not\approx p \mathrel{|\!\sim} l$$

Results & Findings

For each form of entailment, we present **3 algorithms for constructing representations** and **2 algorithms for computing entailment** of a defeasible knowledge base.

1. The first construction algorithm produces a **ranking of worlds**. We define the first entailment algorithm to use this representation to answer queries.
2. The second algorithm uses **formulas to represent the ranked worlds**. We adapt the first entailment algorithm to use this representation.
3. The third algorithm constructs a **cumulative representation** of the ranked worlds using compact formulas. This representation is compatible with the second entailment algorithm.

We **implement** each of our algorithms and compare their performance with existing approaches. We find our algorithms trade construction time for entailment efficiency.

R_2	$\mathtt{lmp}$
R_1	$\overline{\mathtt{l}}\mathtt{m}\overline{\mathtt{p}}$, $\overline{\mathtt{l}}\mathtt{mp}$
R_0	$\mathtt{l}\overline{\mathtt{mp}}$, $\mathtt{lm}\overline{\mathtt{p}}$, $\overline{\mathtt{lmp}}$

1. Ranked Model

F_2	$(\neg(p \rightarrow m) \mathrel{|\!\sim} \bot) \land \neg F_0 \land \neg F_1$
F_1	$((\neg(p \rightarrow m) \mathrel{|\!\sim} \bot) \land (p \mathrel{|\!\sim} \neg l)) \land \neg F_0$
F_0	$(m \mathrel{|\!\sim} l) \land (\neg(p \rightarrow m) \mathrel{|\!\sim} \bot) \land (p \mathrel{|\!\sim} \neg l)$

2. Formula Model

F_2	$(\neg(p \rightarrow m) \mathrel{|\!\sim} \bot)$
F_1	$(\neg(p \rightarrow m) \mathrel{|\!\sim} \bot) \land (p \mathrel{|\!\sim} \neg l)$
F_0	$(m \mathrel{|\!\sim} l) \land (\neg(p \rightarrow m) \mathrel{|\!\sim} \bot) \land (p \mathrel{|\!\sim} \neg l)$

3. Cumulative Formula Model

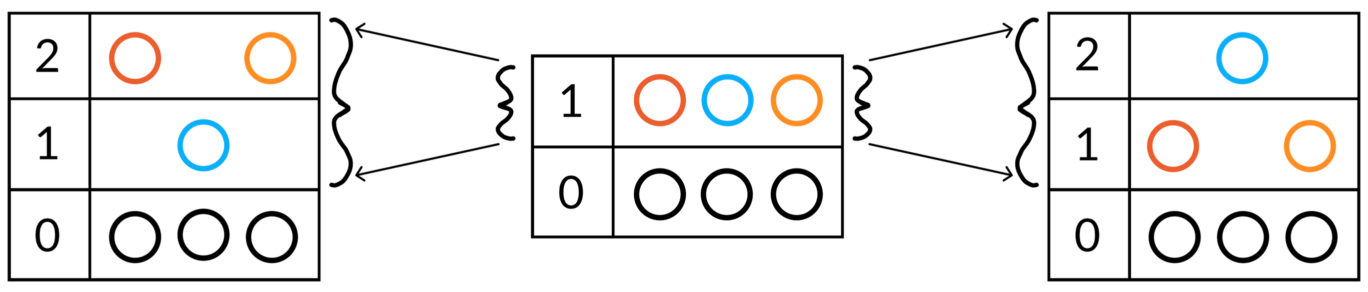

Different Lexicographic Refinements of Rational Closure

While formulating the cumulative lexicographic closure algorithm, we find and prove that there are **two different definitions of lexicographic closure** in the literature. Our algorithms correspond to a new count-based form of lexicographic closure. While both definitions are valid refinements of rational closure, they represent distinct forms of reasoning.

Conclusions

1. We defined **tractable** construction and entailment algorithms.
2. There are **similarities** between existing algorithms and our cumulative algorithms.
3. There are two **different forms of lexicographic closure** in the literature.

Future Work

1. Extend the algorithms to **general forms of defeasible entailment**.
2. Investigate model-based algorithms for **updating the representations**.
3. Investigate the **two forms of lexicographic closure** in the literature.
4. **Optimize** the algorithm implementations.

University of Cape Town · iYunivesithi yaseKapa · Universiteit van Kaapstad

SPES BONA

cair

CENTRE FOR ARTIFICIAL INTELLIGENCE RESEARCH

www.cair.org.za 021 650 2670

Authors

Jaron Cohen chnjar003@myuct.ac.za

Carl Combrinck cmbcar007@myuct.ac.za

Supervisor

Professor Thomas Meyer tmeyer@cair.org.za

www.sit.uct.ac.za

dept@cs.uct.ac.za

twiddleproject.com