COMP4418: Knowledge Representation and Reasoning

Commonsense Reasoning: Non-Monotonic Reasoning

Maurice Pagnucco
School of Computer Science and Engineering

COMP4418, Week 3

Strictness of FOL

To reason from $P(a)$ to $Q(a)$, need either

- facts about *a* itself
- universals, e.g. $\forall x(P(x) \rightarrow Q(x))$
  - something that applies to all instances
  - all or nothing!

But most of what we learn about the world is in terms of *generics*

- e.g., encyclopedia entries for ferris wheels, violins, turtles, wildflowers

Properties are not strict for all instances, because

- genetic / manufacturing varieties
  - early ferris wheels
- borderline cases
  - toy violins
- imagined cases
  - flying turtles
- cases in exceptional circumstances
  - dried wildflowers
- ...

Generics vs Universals

Violins have four strings
  vs.
All violins have four strings
  vs.
All violins that are not E_1 or E_2 or ... have four strings
  (exceptions usually cannot be enumerated)
Similarly, for general properties of individuals
  Alexander the great: ruthlessness
  Ecuador: exports
  pneumonia: treatment
Goal: be able to say a P is a Q *in general*, but not necessarily
  reasonable to conclude $Q(a)$ given $P(a)$ unless there is a good reason not to
Here: qualitative version (no numbers)

Varieties vs Defaults

General statements

- statistical: Most *P*'s are *Q*'s.
  - People living in Quebec speak French.
- normal: All normal *P*'s are *Q*'s.
  - Polar bears are white.
- prototypical: The prototypical *P* is a *Q*.
  - Owls hunt at night.

Representational

- conversational: Unless I tell you otherwise, a *P* is a *Q*.
  - default slot values in frames
  - disjointness in IS-A hierarchy (sometimes)
  - closed-world assumption (below)

Epistemic rationales

- familiarity: If a *P* was not a *Q*, you would know it.
  - an older brother
  - very unusual individual, situation or event
- group confidence: All known *P*'s are *Q*'s.
  - NP-hard problems unsolvable in polynomial time.

Persistence rationale

- inertia: A *P* is a *Q* if it used to be a *Q*.
  - colours of objects
  - locations of parked cars (for a while!)

Nonmonotonic Reasoning

- Suppose you are told “Tweety is a bird”
- What conclusions would you draw?
- Now, consider being further informed that “Tweety is an emu”
- What conclusions would you draw now? Do they differ from the conclusions that you would draw without this information? In what way(s)?
- Nonmonotonic reasoning is an attempt to capture a form of *commonsense* reasoning

Nonmonotonicity

Closed World Assumption

Predicate Completion

Circumscription

Default Logic

Nonmonotonic Consequence

- KLM Systems

Nonmonotonic Reasoning

- In classical logic the more facts (premises) we have, the more conclusions we can draw
- This property is known as *Monotonicity*

$$\text{If } \Delta \subseteq \Gamma, \text{ then } Cn(\Delta) \subseteq Cn(\Gamma)$$

  (where *Cn* denotes classical consequence)
- However, the previous example shows that we often do not reason in this manner
- Might a nonmonotonic logic—one that does not satisfy the Monotonicity property—provide a more effective way of reasoning?

Why Nonmonotonicity?

- Problems with the classical approach to consequence
    - It is usually not possible to write down all we would like to say about a domain
    - Inferences in classical logic simply make implicit knowledge explicit; we would also like to reason with tentative statements
    - Sometimes we would like to represent knowledge about something that is not *entirely* true or false; uncertain knowledge
- Nonmonotonic reasoning is concerned with getting around these shortcomings

Makinson's Classification

Makinson has suggested the following classification of nonmonotonic logics:

- Additional background assumptions
- Restricting the set of valuations
- Additional rules

David Makinson, *Bridges from Classical to Nonmonotonic Logic*, Texts in Computing, Volume 5, King's College Publications, 2005.

Nonmonotonicity

- Classical logic satisfies the following property
- Monotonicity: If $\Delta \subseteq \Gamma$, then $Cn(\Delta) \subseteq Cn(\Gamma)$ (equivalently, $\Gamma \vdash \phi$ implies $\Gamma \cup \Delta \vdash \phi$)
- However, we often draw conclusions based on 'what is normally the case' or 'true by default'
- More information can lead us to retract previous conclusions
- We shall adopt the following notation
  - $\vdash$ classical consequence relation
  - $\mid\!\sim$ nonmonotonic consequence relation

Consequence Operation Cn

Other properties of consequence operation Cn:

Inclusion $\Delta \subseteq Cn(\Delta)$

Cumulative Transitivity $\Delta \subseteq \Gamma \subseteq Cn(\Delta)$ implies $Cn(\Gamma) \subseteq Cn(\Delta)$

Compactness If $\phi \in Cn(\Delta)$ then there is a finite $\Delta' \subseteq \Delta$ such that $\phi \in Cn(\Delta')$

Disjunction in the Premises $Cn(\Delta \cup \{a\}) \cap Cn(\Delta \cup \{b\}) \subseteq Cn(\Delta \cup \{a \vee b\})$

Note: $\Delta \vdash \phi$ iff $\phi \in Cn(\Delta)$
alternatively: $Cn(\Delta) = \{\phi : \Delta \vdash \phi\}$

Example

Suppose I tell you ‘Tweety is a bird’
You might conclude ‘Tweety flies’
I then tell you ‘Tweety is an emu’
You conclude ‘Tweety does not fly’

$$bird(Tweety) \mathrel{|\!\sim} flies(Tweety)$$
$$bird(Tweety) \land emu(Tweety) \mathrel{|\!\sim} \neg flies(Tweety)$$

The Closed World Assumption

- A *complete* theory is one in which for every ground atom in the language, either the atom or its negation appears in the theory
- The *closed world assumption* (CWA) completes a base (non-closed) set of formulae by including the negation of a ground atom whenever the atom does not follow from the base
- In other words, if we have no evidence as to the truth of (ground atom) P, we assume that it is false
- Given a base set of formulae Δ we first calculate the *assumption* set
  $\neg P \in \Delta_{asm}$ iff for ground atom P, $\Delta \not\vdash P$
- $CWA(\Delta) \; = \; Cn\{\Delta \cup \Delta_{asm}\}$

Example

$$\begin{aligned}\Delta &= \{P(a), P(b), P(a) \rightarrow Q(a)\}\\ \Delta_{asm} &= \{\neg Q(b)\}\end{aligned}$$

Theorem: The CWA applied to a consistent set of formulae Δ is inconsistent iff there are positive ground literals $L_1, \ldots, L_n$ such that $\Delta \models L_1 \vee \ldots \vee L_n$ but $\Delta \not\models L_i$ for $i = 1, \ldots, n$.

- Note that in the example above we limited our attention to the object constants that appeared in Δ however the language could contain other constants. This is known as the *Domain Closure Assumption* (DCA)
- Another common assumption is the *Unique-Names Assumption* (UNA).

  If two ground terms can't be proved equal, assume that they are not.

Predicate Completion

Idea: The only objects that satisfy a predicate are those that must

- For example, suppose we have $P(a)$. Can view this as

  $$\forall x.\ x = a \rightarrow P(x)$$

  the *if*-half of a definition
- Can add the *only if* part:

  $$\forall x.\ P(x) \rightarrow x = a$$
- Giving:

  $$\forall x.\ P(x) \leftrightarrow x = a$$

Predicate Completion

- **Definition:** A clause is *solitary* in a predicate P if whenever the clause contains a postive instance of P, it contains only one instance of P.
  - For example, $Q(a) \vee P(a) \vee \neg P(b)$ is not solitary in P
    $Q(a) \vee R(a) \vee P(b)$ is solitary in P
- Completion of a predicate is only defined for sets of clauses solitary in that predicate

Predicate Completion

- Each clause can be written:
  $\forall y.\ Q_1 \wedge \ldots \wedge Q_m \rightarrow P(t)$ (P not contained in Q_i)
  $\forall y.\ \forall x.\ (x = t) \wedge Q_1 \wedge \ldots \wedge Q_m \rightarrow P(x)$
  $\forall x.(\forall y.\ (x = t) \wedge Q_1 \wedge \ldots \wedge Q_m \rightarrow P(x))$ (normal form of clause)
- Doing this to every clause gives us a set of clauses of the form:
  $\forall x.\ E_1 \rightarrow P(x)$
  $\ldots$
  $\forall x.\ E_n \rightarrow P(x)$
- Grouping these together we get:
  $\forall x.\ E_1 \vee \ldots \vee E_n \rightarrow P(x)$
- Completion becomes: $\forall x.\ P(x) \leftrightarrow E_1 \vee \ldots \vee E_n$
  and we can add this to the original set of formulae

Example

- Suppose $\Delta = \{\forall x.\, Emu(x) \rightarrow Bird(x),$
  $Bird(Tweety),$
  $\neg Emu(Tweety)\}$
- We can write this as
  $$\forall x.\, (Emu(x) \vee x = Tweety) \rightarrow Bird(x)$$
- Predicate completion of P in Δ becomes
  $$\Delta \cup \{\forall x.\, Bird(x) \rightarrow Emu(x) \vee x = Tweety\}$$

Circumscription

- **Idea:** Make extension of predicate as small as possible
- Example:

  $\forall x. Bird(x) \land \neg Ab(x) \rightarrow Flies(x)$

  $Bird(Tweety), Bird(Sam), Tweety \neq Sam, \neg Flies(Sam)$

- Want to be able to conclude $Flies(Tweety)$ but $\neg Flies(Sam)$
- Accept interpretations where *Ab* predicate is as “small” as possible
- That is, we *minimise abnormality*

Circumscription

- Given interpretations $\mathbf{I}_1 = \langle D, I_1 \rangle$, $\mathbf{I}_2 = \langle D, I_2 \rangle$, $\mathbf{I}_1 \leq \mathbf{I}_2$ iff for every predicate $P \in \mathbf{P}$, $I_1[P] \subseteq I_2[P]$.
- $\Gamma \models_{circ} \phi$ iff for every interpretation $\mathbf{I}$ such that $\mathbf{I} \models \Gamma$, either $\mathbf{I} \models \phi$ or there is a $\mathbf{I}' < \mathbf{I}$ and $\mathbf{I}' \models \Gamma$.
- ϕ is true in all minimal models
- Now consider

$$\forall x. Bird(x) \land \neg Ab(x) \rightarrow Flies(x)$$
$$\forall x. Emu(x) \rightarrow Bird(x) \land \neg Flies(x)$$
$$Bird(Tweety)$$

Reiter's Default Logic (1980)

- Add default rules of the form $\frac{\alpha:\beta}{\gamma}$
  - "If α can be proven and consistent to assume β, then conclude γ"
- Often consider *normal* default rules $\frac{\alpha:\beta}{\beta}$
- Example: $\frac{\mathit{bird}(x):\mathit{flies}(x)}{\mathit{flies}(x)}$
- Default theory $\langle D, W\rangle$

  D – set of defaults; W – set of facts
- *Extension* of default theory contains as many default conclusions as possible and must be consistent (and is closed under classical consequence Cn)
- Concluding whether formula ϕ follows from $\langle D, W\rangle$
  - Sceptical inference: ϕ occurs in *every* extension of $\langle D, W\rangle$ Credulous inference: ϕ occurs in *some* extension of $\langle D, W\rangle$

Examples

- $W = \{\}$; $D = \{\frac{:p}{\neg p}\}$ – no extensions
- $W = \{p \vee r\}$; $D = \{\frac{p:q}{q}, \frac{r:q}{q}\}$ – one extension $\{p \vee r\}$
- $W = \{p \vee q\}$; $D = \{\frac{:\neg p}{\neg p}, \frac{:\neg q}{\neg q}\}$ – two extensions $\{\neg p, p \vee q\}, \{\neg q, p \vee q\}$
- $W = \{emu(Tweety), \forall x.emu(x) \rightarrow bird(x)\}$; $D = \{\frac{bird(x):flies(x)}{flies(x)}\}$ – one extension
- What if we add $\frac{emu(x):\neg flies(x)}{\neg flies(x)}$?
- Poole (1988) achieves a similar effect (but not quite as general) by changing the way the underlying logic is used rather than introducing a new element into the syntax

Default Theories—Properties

Observation: Every normal default theory (default rules are all normal) has an extension

Observation: If a normal default theory has several extensions, they are mutually inconsistent

Observation: A default theory has an inconsistent extension iff D is inconsistent

Theorem: (Semi-monotonicity)
Given two normal default theories $\langle D, W \rangle$ and $\langle D', W \rangle$ such that $D \subseteq D'$ then, for any extension $\mathcal{E}(D, W)$ there is an extension $\mathcal{E}(D', W)$ where
$\mathcal{E}(D, W) \subseteq \mathcal{E}(D', W)$
(The addition of normal default rules does not lead to the retraction of consequences.)

Nonmonotonic Consequence

- Abstract study and analysis of nonmonotonic consequence relation $\mid\!\sim$ in terms of general properties Kraus, Lehmann and Magidor (1991)
- Some common properties include:

  Supraclassicality If $\phi \vdash \psi$, then $\phi \mid\!\sim \psi$

  Left Logical Equivalence If $\vdash \phi \leftrightarrow \psi$ and $\phi \mid\!\sim \chi$, then $\psi \mid\!\sim \chi$

  Right Weakening If $\vdash \psi \rightarrow \chi$ and $\phi \mid\!\sim \psi$, then $\phi \mid\!\sim \chi$

  And If $\phi \mid\!\sim \psi$ and $\phi \mid\!\sim \chi$, then $\phi \mid\!\sim \psi \wedge \chi$
- Plus many more!

KLM Systems

- Kraus, Lehman and Magidor (1991) study various classes of nonmonotonic consequence relations

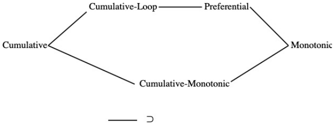

- This has been extended since. A good reference for this line of work is Schlechta (1997)

Summary

- Nonmonotonic reasoning attempts to capture a form of commonsense reasoning
- Nonmonotonic reasoning often deals with inferences based on defaults or ‘what is usually the case’
- Belief change and nonmonotonic reasoning: two sides of the same coin?
- Can introduce abstract study of nonmonotonic consequence relations in same way as we study classical consequence relations
- Similar links exist with conditionals
- One area where nonmonotonic reasoning is important is reasoning about action (dynamic systems)