Uncertainty in AI

Mohsen Afsharchi

Fall 87

Some History

- Probabilistic models such as Bayesian nets are now accepted in AI

- It was not always so . . .

The problem with probability

- People currently vastly outperform any AI systems in inter- acting with the world
- People do not seem to use probabilistic inference to cope with uncertainty - do we?
- "The information necessary to assign numerical probabili- ties is not ordinarily available. Therefore, a formalism that required numerical probabilities would be epistemologically inadequate." (McCarthy and Hayes, MI 4, 1969)

The problem...(2)

- We may not know the probability of our train being canceled, but we still cope with the uncertainty that this possibility engenders.
- Intelligence often identified with symbolic reasoning, not numerical reasoning
- Therefore do not use probability

Logic

- Represent knowledge about the world with facts and rules
  - `bird(tweety).`
  - `fly(X) :- bird(X)`
- Use a theorem-prover to reason about the world. Prolog is a very simple one
  - `?- fly(tweety)`
  - `yes`
- If a conclusion follows from given premises A, B, C, then it also follows from any larger set of premises, as long as the original premises A, B, C, are included

non monotonic Logic

- This won't work outside of toy domains:non tautologous certain rules are hard to find
- In many instances of ordinary or everyday reasoning, people arrive to conclusions only tentatively, based on partial or incomplete information
- This reserves the right to retract those conclusions should they learn new facts
- non-monotonic, because the set of accepted conclusions can become smaller when the set of premises is expanded.

non monotonic Logic

- Replace “Birds fly” with “Birds normally fly”
  - Let $K = bird(X) \rightsquigarrow fly(X), penguin(X) \rightsquigarrow \neg fly(X)$
  - $K, bird(tweety) \vdash fly(tweety)$
  - $K, bird(tweety), penguin(tweety) \nvdash fly(tweety)$

- No one thing which is called “non-monotonic logic” but rather a family of different formalisms: i.e. default logic

Extending Logic: Rules and ...

How do we interpret $if A \longrightarrow^{m} B$

1. If you see A, then you are given the license to update the certainty of B by certain amount which is a function of the rule strength m. Certainty Factor
2. The set of worlds in which A and $\neg B$ hold simultaneously has low likelihood and hence should be excluded with probability m. Dempster-Shafer Theory
3. Among all worlds satisfying A, those that also satisfying B constitute an m percent majority. Bayesian Formalism

Extensional vs Intensional Approach

- Extensional:certainty of formula is defined to be unique function of the certainty of its subformulas
  - The certainty of $A \wedge B$ is given by some function of the certainty measure assigned to A and B individually.(e.g. the minimum or the product)
- Intensional: certainty measures are assigned to sets of worlds, and the connectives combine sets of worlds by set theory operations.
  - $P(A \wedge B) = P(A).P(B)$ True???
- A trade-off between semantic accuracy and computational feasibility.

Extensional Systems

Rules:

- **If *A* then *C* (*x*)**
- **If *B* then *C* (*y*)**
- **If *C* then *D* (*z*)**

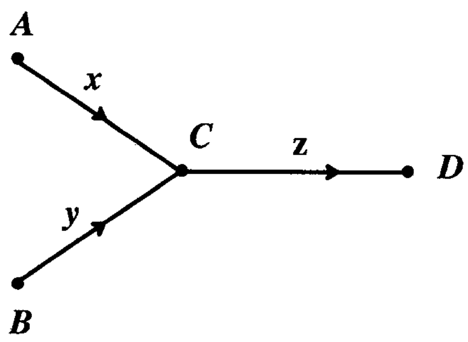

1. Parallel combination

$$CF(C) = \begin{cases} x + y - xy & x, y > 0 \\ (x + y) / (1 - \min(x, y)) & x, y \text{ different sign} \\ x + y + xy & x, y < 0 \end{cases}$$

2. Series combination

$$CF(D) = z \cdot \max(0, CF(C))$$

3. Conjunction, negation ...

Extensional Systems (2)

$A \rightarrow B \Longrightarrow$ if you see A in the knowledge base, then **regardless of what other things the knowledge base contains** and **regardless of how A was derived**, you are given the license to assert B and add it to the database.

- Locality:regardless of other things
- Detachment:regardless of how it was derived

Limits of Modularity

$R_1 =$ `If the ground is wet the assume it rained (with certainty` c_1`)`

- Does the truth of "`The ground is wet`" permit us to increase the certainty of "`It rained`"?
  - "`The sprinkler was on`"
  - "`The neighbor's grass is dry` "
- Imposing a connection between "`sprinkle was on`" and "`It rained`" defeats the spirit of modularity(i.e. locality) by forcing the rule writer to pack together items of information that are only remotely related.

Limits of Modularity (2)

$$(K \rightarrow P \wedge P \rightarrow Q) \Longrightarrow K \rightarrow Q$$

- Does chaining apply here?
  - "If the sprinkler was on then the ground is wet"
  - "If the ground is wet then it rained"
- Violation of detachment

Example

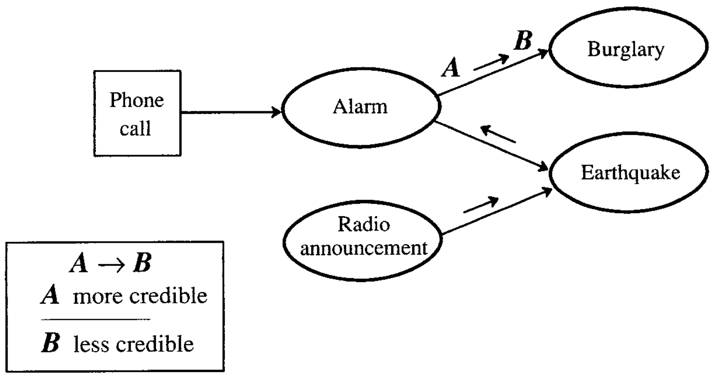

- From effect to cause
  - "Radio → Earthquake"
- From cause to effect
  - "Earthquake → Alarm "
- If A(Alarm) then B(Burglary); and A becomes more credible, then B becomes less credible.
- In contrast with local belief updating

Bidirectional Inferences

If $A \longrightarrow B$, then finding the truth value of B makes A more credible

- Reasoning in both ways, from A to B and from B to A.
  - "`Fire implies smoke`"
  - "`Smoke makes fire more credible`"
- Do we need two separate rules to perform these inferences??