Bayesian Knowledge-driven Ontologies

Representing and reasoning
about conceptual knowledge
with uncertainty and incompleteness

Jacob Jurmain
Adviser: Eugene Santos Jr.

Semantic Networks / AKA Ontologies

- Complex networks of concepts and their relationships in a domain
  - Asserts knowledge with subsumption (is-a) and relational operators (has-a)
  - Exist formally as description logic (Baader et al)
- Foundation of the effort to develop a “semantic web”
  - Embed deep contextual information in web pages
  - Search the web not just with keywords, but with general concepts and relationships

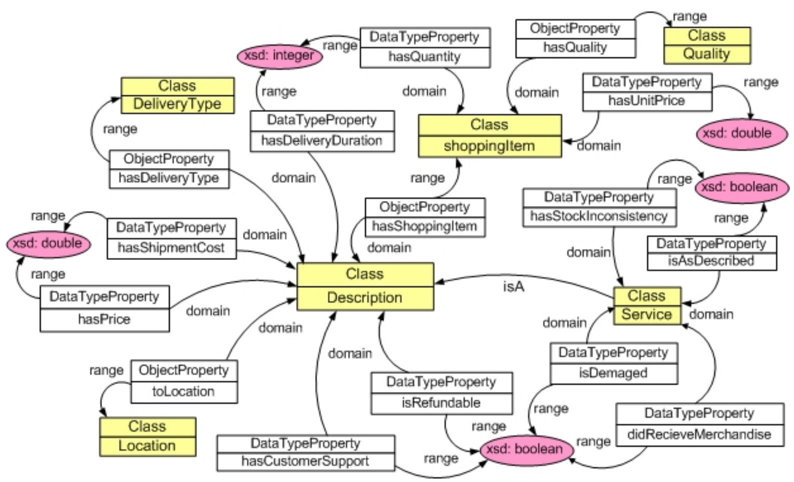

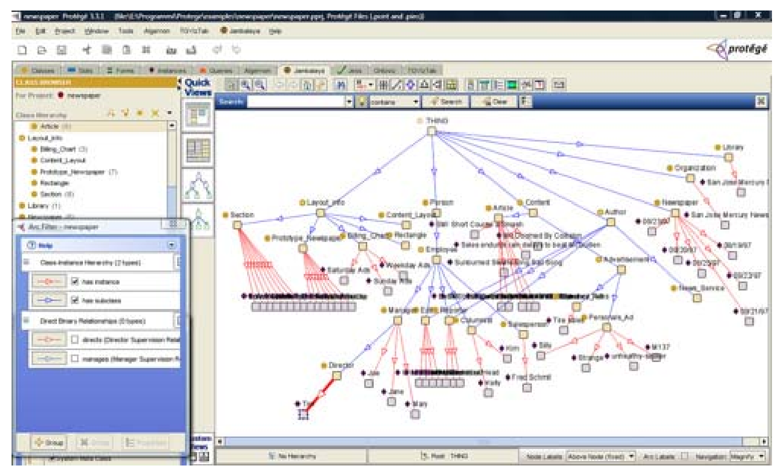

Img sources: http://mas.cmpe.boun.edu.tr/project/files/DomainlevelOntology.jpg, http://www.flickr.com/photos/saldatoccio/2623536667/

The Problem: Uncertainty

- Uncertainty is the existence of multiple conflicting possible instantiations (i.e. states) of a domain. The differences that distinguish domain instantiations are called variables.

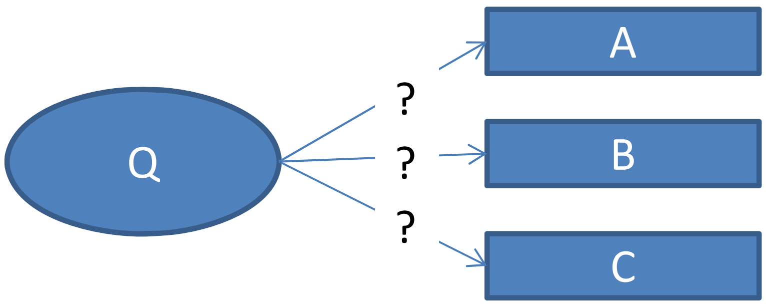

I1: Q→A
I2: Q→B
I3: Q→C
I4: Q→A,B
…

- The number of instantiations is exponential with respect to the number of variables. How can we narrow it down? We need to reason about the variables.

- We need a model of uncertainty that captures the interactions between variables.

Uncertainty Representation

- Uncertainty can be modeled through various theories:
  - Fuzzy logic models partial membership in classes. It measures ambiguity, such as "is the glass half full or half empty".
  - Possibility theory uses two metrics, possibility and necessity, to model uncertainty. It measures the level of "surprise" at a result.
  - Probability theory is the foundation of statistics. Its formulation is based on modeling outcome frequencies of repeated events, such as die rolls. It can be used to model degrees of belief and "if-then / what-if" variable interactions.
- We select probability theory for its strong capture of variable interactions.

Probabilistic Semantic Networks: State of the Art

- Two approaches dominate: those using probabilistic logic (Nilsson, 1986) and those using Bayesian Networks (Pearl, 1985). These approaches are fundamentally limited.
  - Probabilistic logic represents probability as bounds on an interval, so the frameworks using it experience a decay in relative precision over the course of reasoning as inferred probabilities move toward zero and one.
  - Bayesian Networks require complete definition of all variable interactions, so the frameworks using them must represent knowledge at a coarser granularity than conventional semantic networks.

- We want a powerful, fine-granularity probability theory for our framework.

Solution: Bayesian Knowledge Bases

- BKBs (Santos & Santos, 1999) are a generalization of BNs to admit incompleteness.
  - Fine granularity: represent knowledge as sets of "if-then" conditional probability rules among propositional variable instantiations.
  - Powerful reasoning: belief updating and revision like in BNs, but allowing for incompleteness.

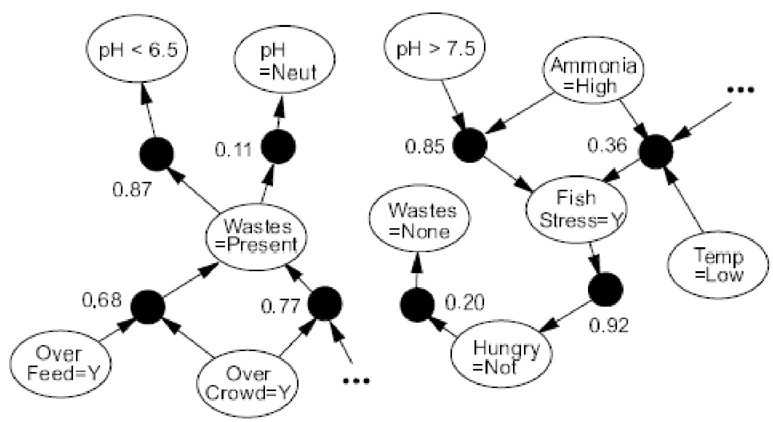

FIG. 2.2. A BKB fragment from fresh-water aquarium maintenance knowledge-base as a directed graph.

Img source: Santos, Eugene, Jr., Santos, Eugene S., and Shimony, Solomon Eyal., "Implicitly Preserving Semantics During Incremental Knowledge Base Acquisition Under Uncertainty," International Journal of Approximate Reasoning 33(1), 71-94, 2003.

Bayesian Knowledge-driven Ontologies

- Synthesis of semantic networks and BKBs
  - Represent knowledge as conditional probability rules with description logic assertions as arguments. Ex: $P(a \in C \mid b \in D) = 0.6$. These rules can be propositional (AKA assertional) or predicated (AKA terminological).

- Sets of these rules define a probability distribution over the possible instantiations of the domain.
  - We use description logic's reasoning functions to generate propositional rules from predicated ones.
  - We use BKBs' reasoning functions to calculate marginal probabilities for domain instantiations.

- Powered by two insights:
  - Generalizing the rule of universal instantiation to its probabilistic analog allows description logic to validly propagate uncertainty through its reasoning process.
  - Distilling the propositional knowledge from a BKO makes a BKB.

Key Formulation Excerpts

- **Notation:** A *probabilistic assertional axiom* (PAA) is a conditional probability rule of the form $P(Z|\{Y_1 \dots Y_n\}) = p$ where Z and $\{Y_1 \dots Y_n\}$ are propositional description logic statements and $p \in [0,1]$ is the conditional probability of Z given $\{Y_1 \dots Y_n\}$.

- **Notation:** A *probabilistic terminological axiom* (PTA) is a statement of the form $P(s \in D \mid s \in C) = p$ for any classes C and D and a probability $p \in [0,1]$ that any given individual s known to be a member of class C is also a member of class D.

- **Theorem:** The *probabilistic rule of universal instantiation* states the following: let “a” be a specific individual. Let C and D be classes. Let $Y_1 \dots Y_n$ be classical assertional axioms. Given a probabilistic assertional axiom $P(a \in C \mid Y_1 \dots Y_n) = p$ and a probabilistic terminological axiom $P(s \in D \mid s \in C) = q$, infer $P(a \in D \mid a \in C) = q$.

Computing Domain Instantiation Probabilities

- Start by reasoning over the model using the probabilistic rule of universal instantiation. Generate all inferrable propositions.

- Once all predicate reasoning has been done, the distilled propositional knowledge in the BKO describes a BKB. We can then perform BKB reasoning tasks:
  - Belief updating computes the marginal probability of a target variable instantiation.
  - Belief revision computes the marginal probabilities of domain instantiations.

- Both tasks can also be used to answer “what if” questions by setting evidence in the BKB, i.e. by temporarily holding some variable instantiations as true.

Potential Applications

- Semantic web
  - BKB's "Fusion" mechanic (Santos, Wilkinson, and Santos, 2009) applied to BKOs would facilitate ontology merging, a current major focus of semantic web research efforts.
- Robotics
  - Representation and reasoning with conceptual, uncertain, incomplete knowledge could facilitate robots' handling of more unpredictable environments.
- Machine learning
  - A long-term goal. In theory, machines could learn conceptually by performing variance analysis on their "memories" and adding the results to their knowledge base as new conditional probability rules. But many questions remain to be answered to make this happen.