Bayesian Ontology Fusion

Formal reasoning about conflicting knowledge
from disagreeing sources

Jacob Jurmain
Adviser: Eugene Santos Jr.

RECAP RIP 2011: BKO THEORY

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 background context and conceptual relationships

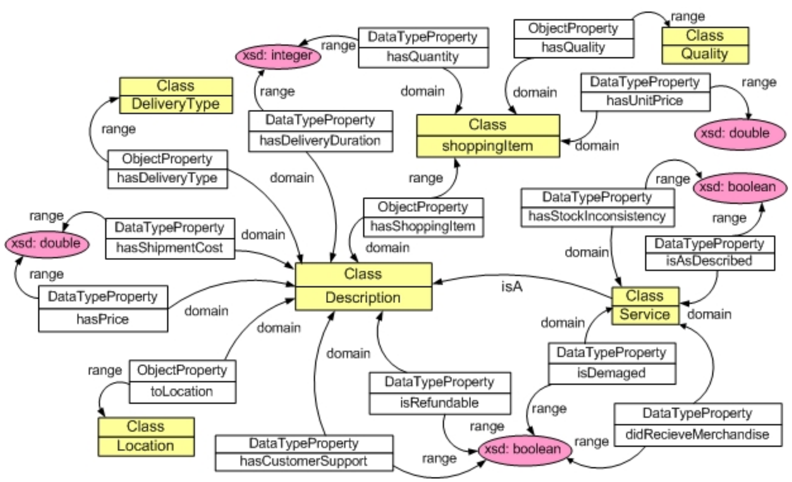

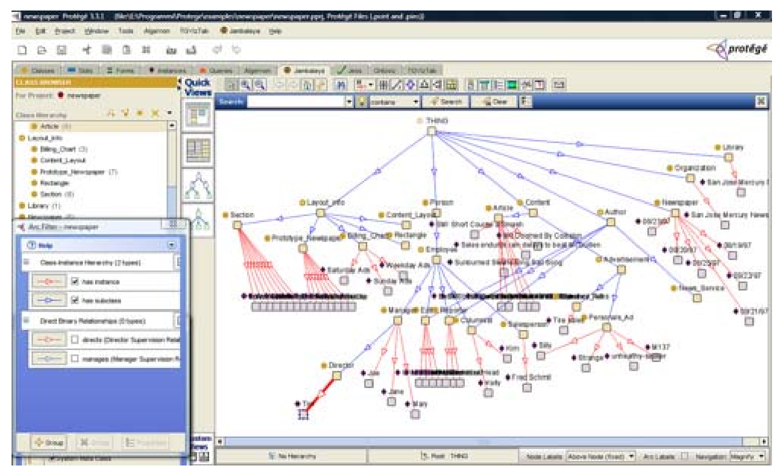

Img sources: http://mas.cmpe.boun.edu.tr/project/files/DomainlevelOntology.jpg, http://www.flickr.com/photos/saldatoccio/2623536667/

Uncertainty Reasoning on Ontologies

- Problem: No good way to reason on uncertain ontologies.
  - Want to answer question, “Given some evidence “E”, what is P(X|E)?”
  - Prior work either places unintuitive restrictions on what can be represented, or uses inadequate reasoning methods.

- Key insights
  - “Uncertainty” is just multiple possible ontologies.
  - Can model a probability distribution over them.
  - Can easily generalize ontology reasoning to work with it.
  - The result naturally ends up matching a powerful uncertainty reasoning theory, Bayesian Knowledge Bases.

Background: Bayesian Knowledge Bases

- BKBs (Santos & Santos, 1999) model probability distributions over possible states of the world.
  - Represent knowledge as sets of “if-then” conditional probability rules between variable states.
  - Allows incompletely defined relationships between variables.
  - Reasoning computes marginal probabilities and analyzes contributions.

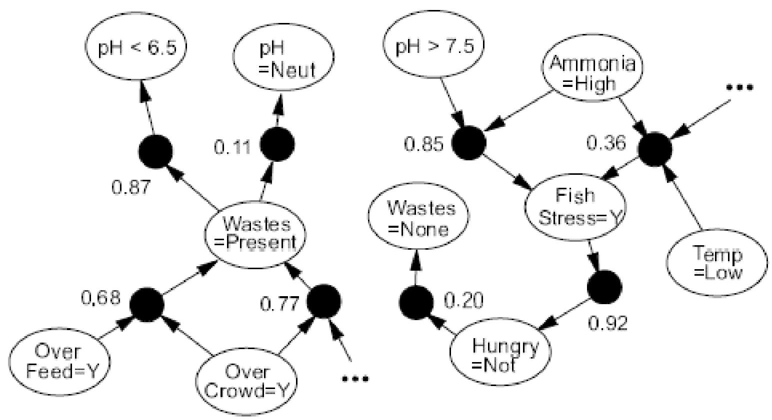

FIG. 2.2. A BKB fragment from fresh-water aquarium maintenance knowledge-base as a directed graph.

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.

Result: Bayesian Knowledge-driven Ontologies (Santos & Jurmain, 2011)

- Represent knowledge as conditional probability rules between DL assertions.
  - Ex: $P(a \in C \mid b \in D) = 0.6$.

- Reasoning turns a BKO into a BKB.
  - Ontology reasoning describes the states of the world as completely as possible.
  - The result is a BKB. Convenient!

- BKOs are provably a subclass of BKBs. Any reasoning that works for BKBs works for BKOs.

NEW GOAL: FUSING CONFLICTING KNOWLEDGE FROM DISAGREEING SOURCES

How Ontologies Fuse Knowledge

- If ontologies use the same interpretation, naïve fusion is fine.

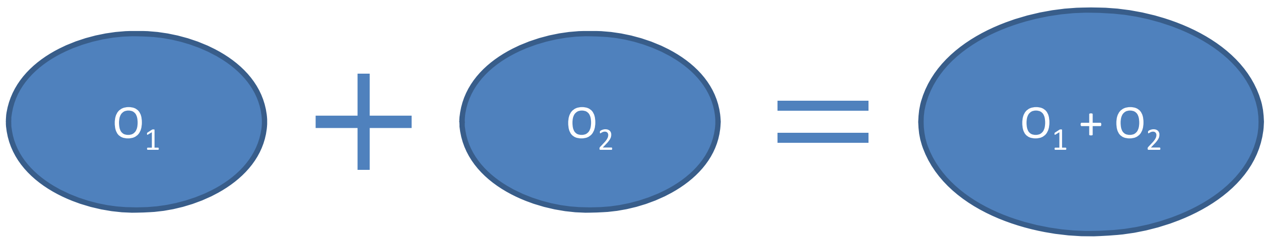

- If they don't, have to build a mapping ontology that translates, then fuse.

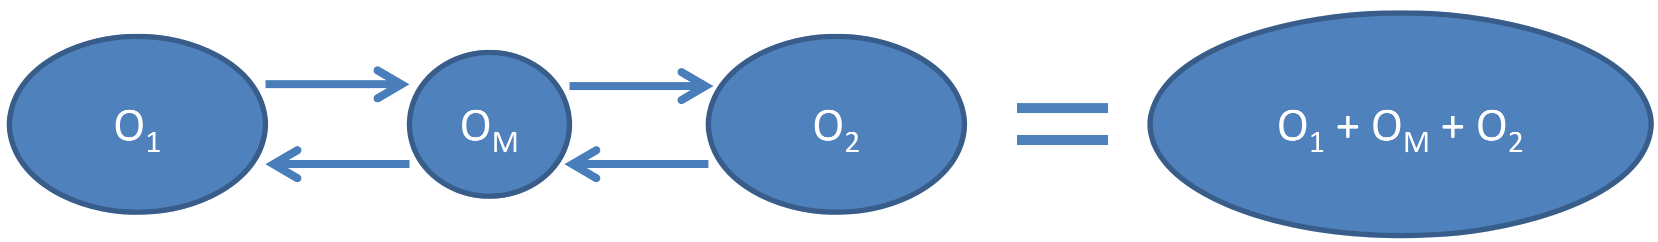

- If knowledge actually conflicts, either in the ontos or the mappings, have to discard something ad-hocly.

How BKBs Fuse Knowledge
(Santos et al, 2009)

- Create a probability distribution of source reliabilities.
- Sources can have different reliabilities for different rules.

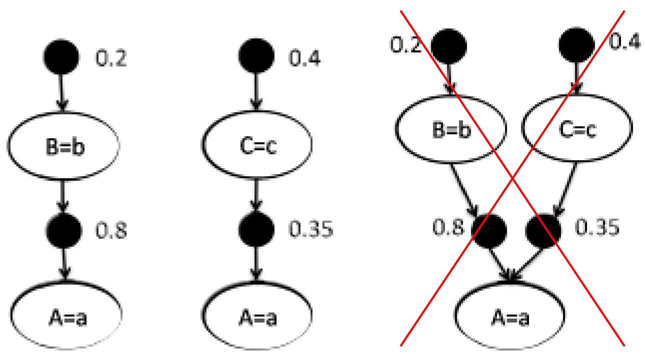

(1) (2)

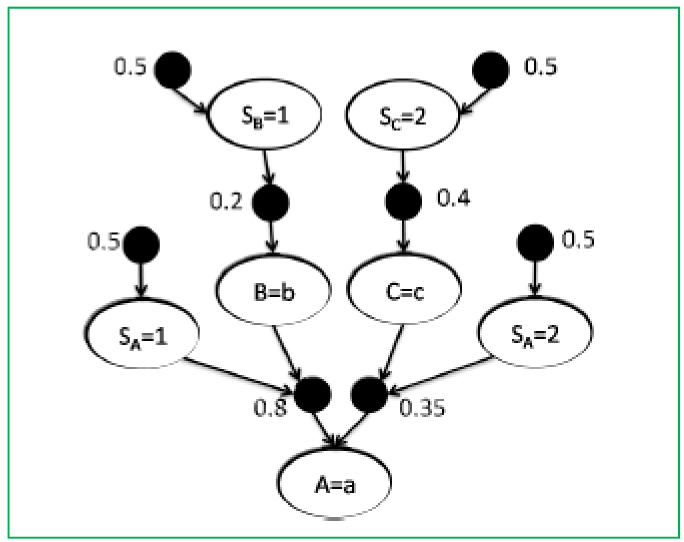

Naïve union of fragments (1) and (2) puts CPRs in conflict. Invalid.

Source variables S_x prevent rules from conflicting because they give the rules mutually exclusive conditions.

BKO Fusion combines both methods

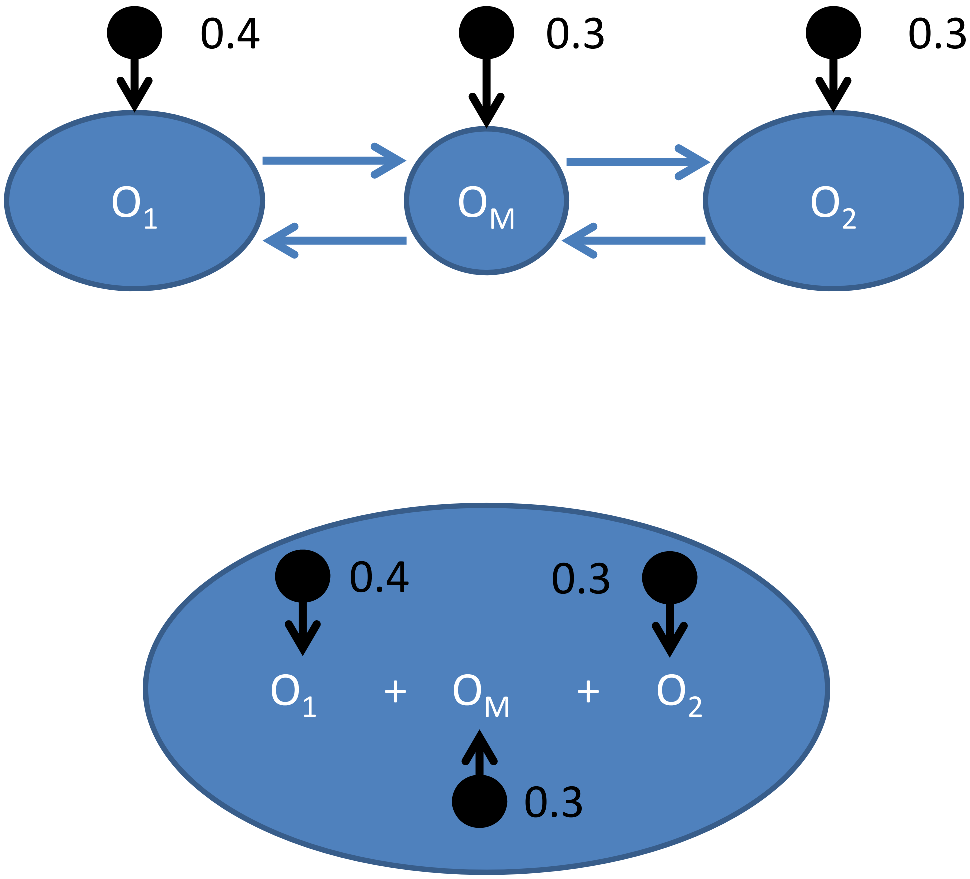

Ex: Which mapping is the right one?

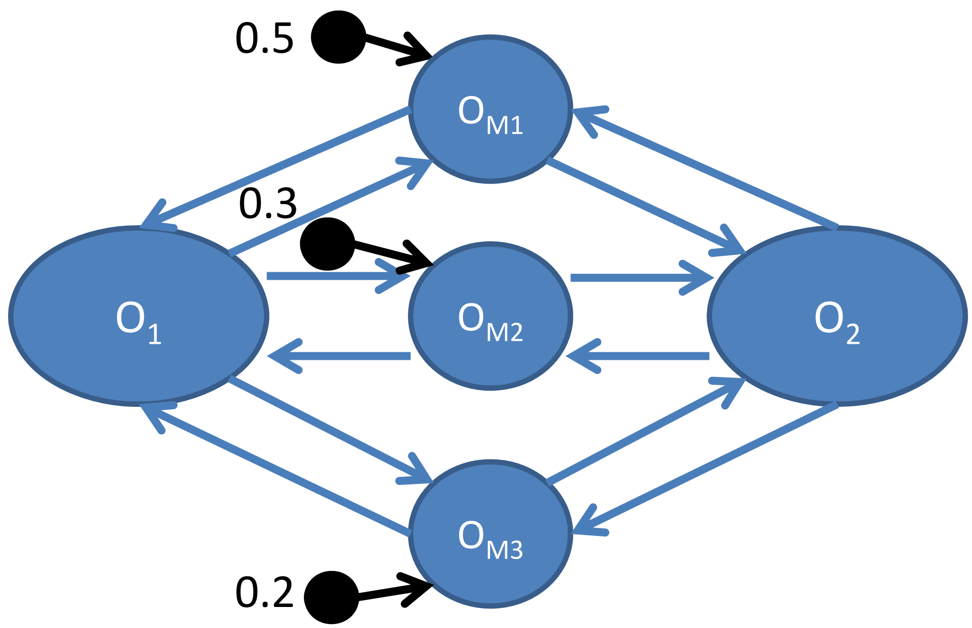

Approach

- Updating BKO theory to include source variables and define BKO fusion.

- Recall BKOs are a subclass of BKBs, so we have proof it will work.

Fusion Reveals New Insights

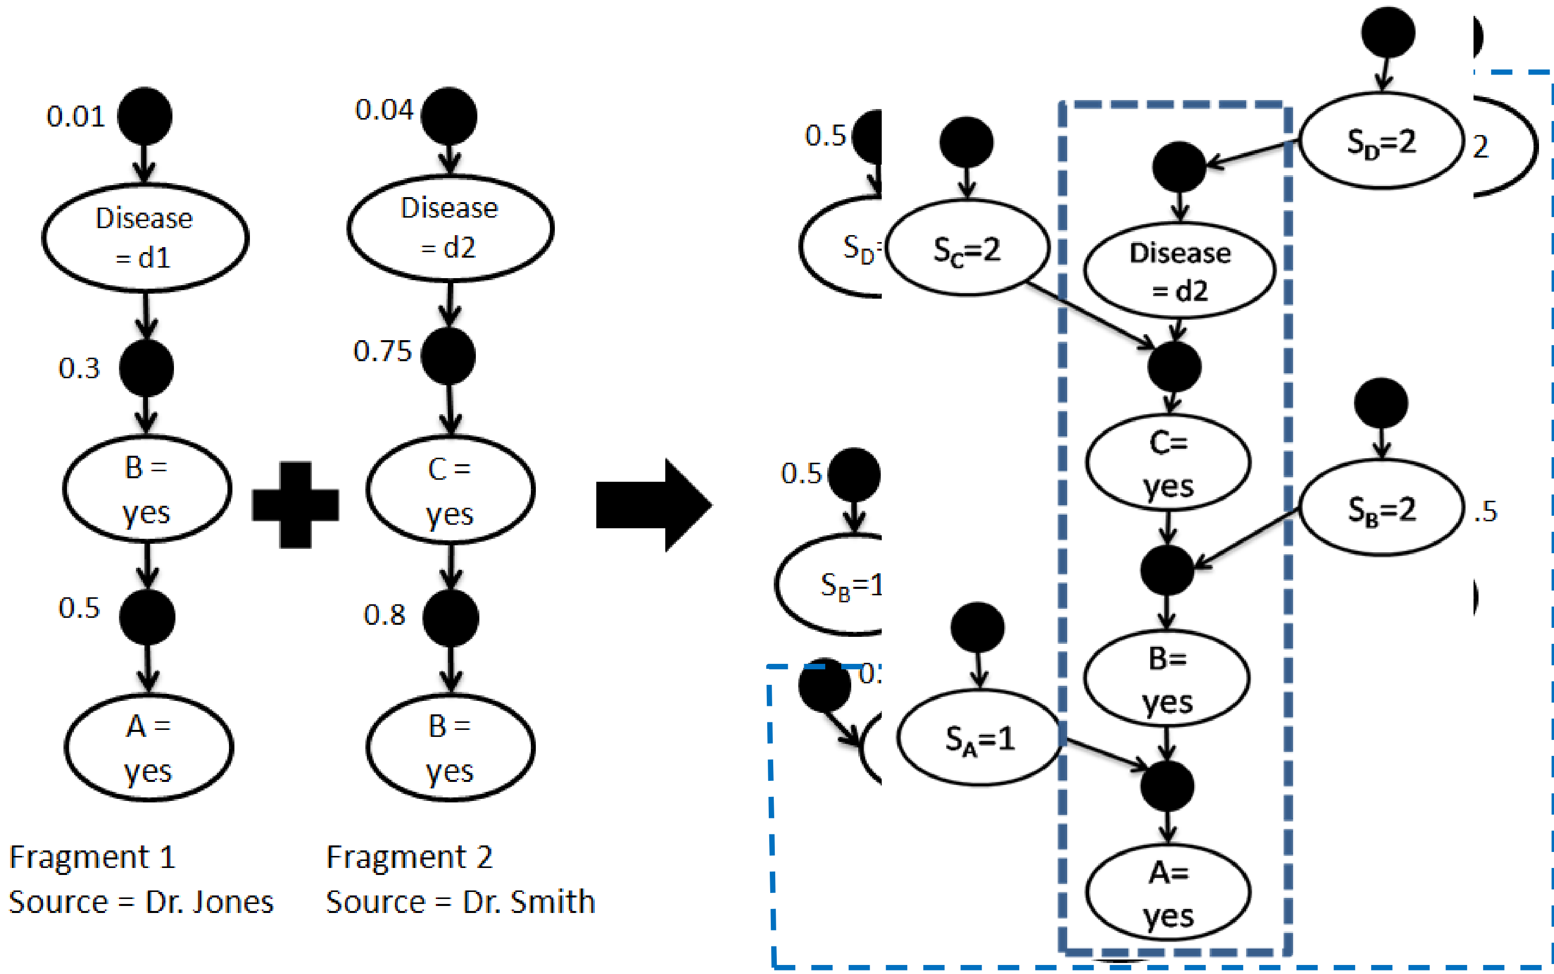

Bayesian Network / Ontology Syntheses

- Bayesian Networks (Pearl, 1985)
  - Restricted subclass of Bayesian Knowledge Bases that assumes complete information.
  - BNs require complete definition of “conditional probability tables” instead of working with individual rules like BKBs.
- PR-OWL (Costa and Laskey, 2005), BayesOWL (Ding et al, 2005), and P-CLASSIC (Koller et al, 1997) are representative works.

- Defines conditional probability tables using DL assertions as variables.
  - DL does not have BNs’ completeness requirement. Using BNs restricts the system’s expressiveness.
  - There are notions we can represent in DL that don’t work in BNs even when completely known.
  - Ex: Model probability distributions of gas mileage for various airplane models. What happens when one is a glider? Then any distribution, even context-specific independence (Boutilier et al, 1996), is unintuitive.

Fuzzy Description Logic

- Founded on fuzzy logic / fuzzy set theory (Zadeh, 1965)
- *Reasoning within fuzzy description logics* (Straccia, 2001) is a representative work.

- Extends DL to allow partial membership in concepts.
  - Coarse treatment of uncertainty with some information loss during reasoning. Does not intuitively capture if-then interactions like probability theory.
  - Ex: given the assertions

    a in C : 0.7   a in D : 0.4   C in E : 0.2   D in E : 0.6

    what is the membership of a in E?

    max(min(0.7, 0.2), min(0.4,0.6)) = 0.4

    Most of the numbers in the reasoning chain had no effect on the outcome. We usually don't think of causality as working this way.

Possibilistic Description Logic

- Founded on possibility theory (Zadeh, 1978) which extends fuzzy logic.
- *A possibilistic extension for description logics* (Qi et al, 2007) is a representative work.

- Models a DL assertion's uncertainty as two fuzzy numbers, possibility and necessity.
  - Possibility: to what degree could the assertion be true? Necessity: to what degree must the assertion be true?
  - Mathematically, possibility and necessity are simply two fuzzy description logic problems in parallel, with the axiom that possibility $\geq$ necessity.

  - As with fuzzy logic , this is a coarse treatment of causality.