Bayesian Knowledge-driven Ontologies

A knowledge theory for automated construction of reasoning models for operations planning and event analysis

Jacob Jurmain
Thayer School of Engineering
Dartmouth College
jacob.c.jurmain@dartmouth.edu

This is joint work with Dr. Eugene Santos, Jr.

http://di2ag.thayer.dartmouth.edu/

Contents

What is a reasoning model?

- Not just what does someone think, but why?
  - Knowledge about self and others
  - Assumptions and how they're justified
  - Values and their cultural context
  - Goals and why they're sought
  - Actions and their expected effects
  - Dynamics of how agents react to each other

- Reasoning models need:
  - Representation. A system that captures an aspect of the world.
  - Reasoning. A mechanism for deriving new knowledge from inputs.
  - Grounding. A relevance to human interpretations of the world.

Reasoning models support stronger operations planning

- Know yourself, your friends, and your enemies
  - Model your adversaries to show what contingencies are credible and why. Get inside their heads and act preventitively. Precisely inform psyops or manipulate the situation to push their decisions your way.
  - Model yourself to defend your reasoning. Spell out what you know, what you want to achieve, and how your actions will get you there.
  - Model your friends to improve cooperation. Coordinate terminology. Ensure compatible goals and sound reasoning.

- Learn from the models you build
  - Build reasoning chains to validate answers. Trace back to root causes and influencing factors. Verify what is knowledge and what is conjecture.
  - Fuse conflicting opinions and derive new insights from the synthesis.
  - Ask “what if?”, plug it into the model, and see how the outcome changes.

Contents

Ontologies model general domain knowledge and its projection onto specific instances

- Ontologies model the terms and concepts that define a domain, and use it to make inferences about specific cases being considered. Use ontologies for…
  - Coordinating terminologies to avert miscommunications.
  - Contextualizing events and observations from other perspectives.
  - Diagnosing scenarios using background knowledge.

- The theory Description Logic (Baader et al, 2007) is research's gold standard and the basis for most applications.
  - Models a "semantic network", a graph of how things and the concepts that describe them are related.
  - Reasoning algorithms infer new relationships by applying general knowledge to case-specific knowledge.

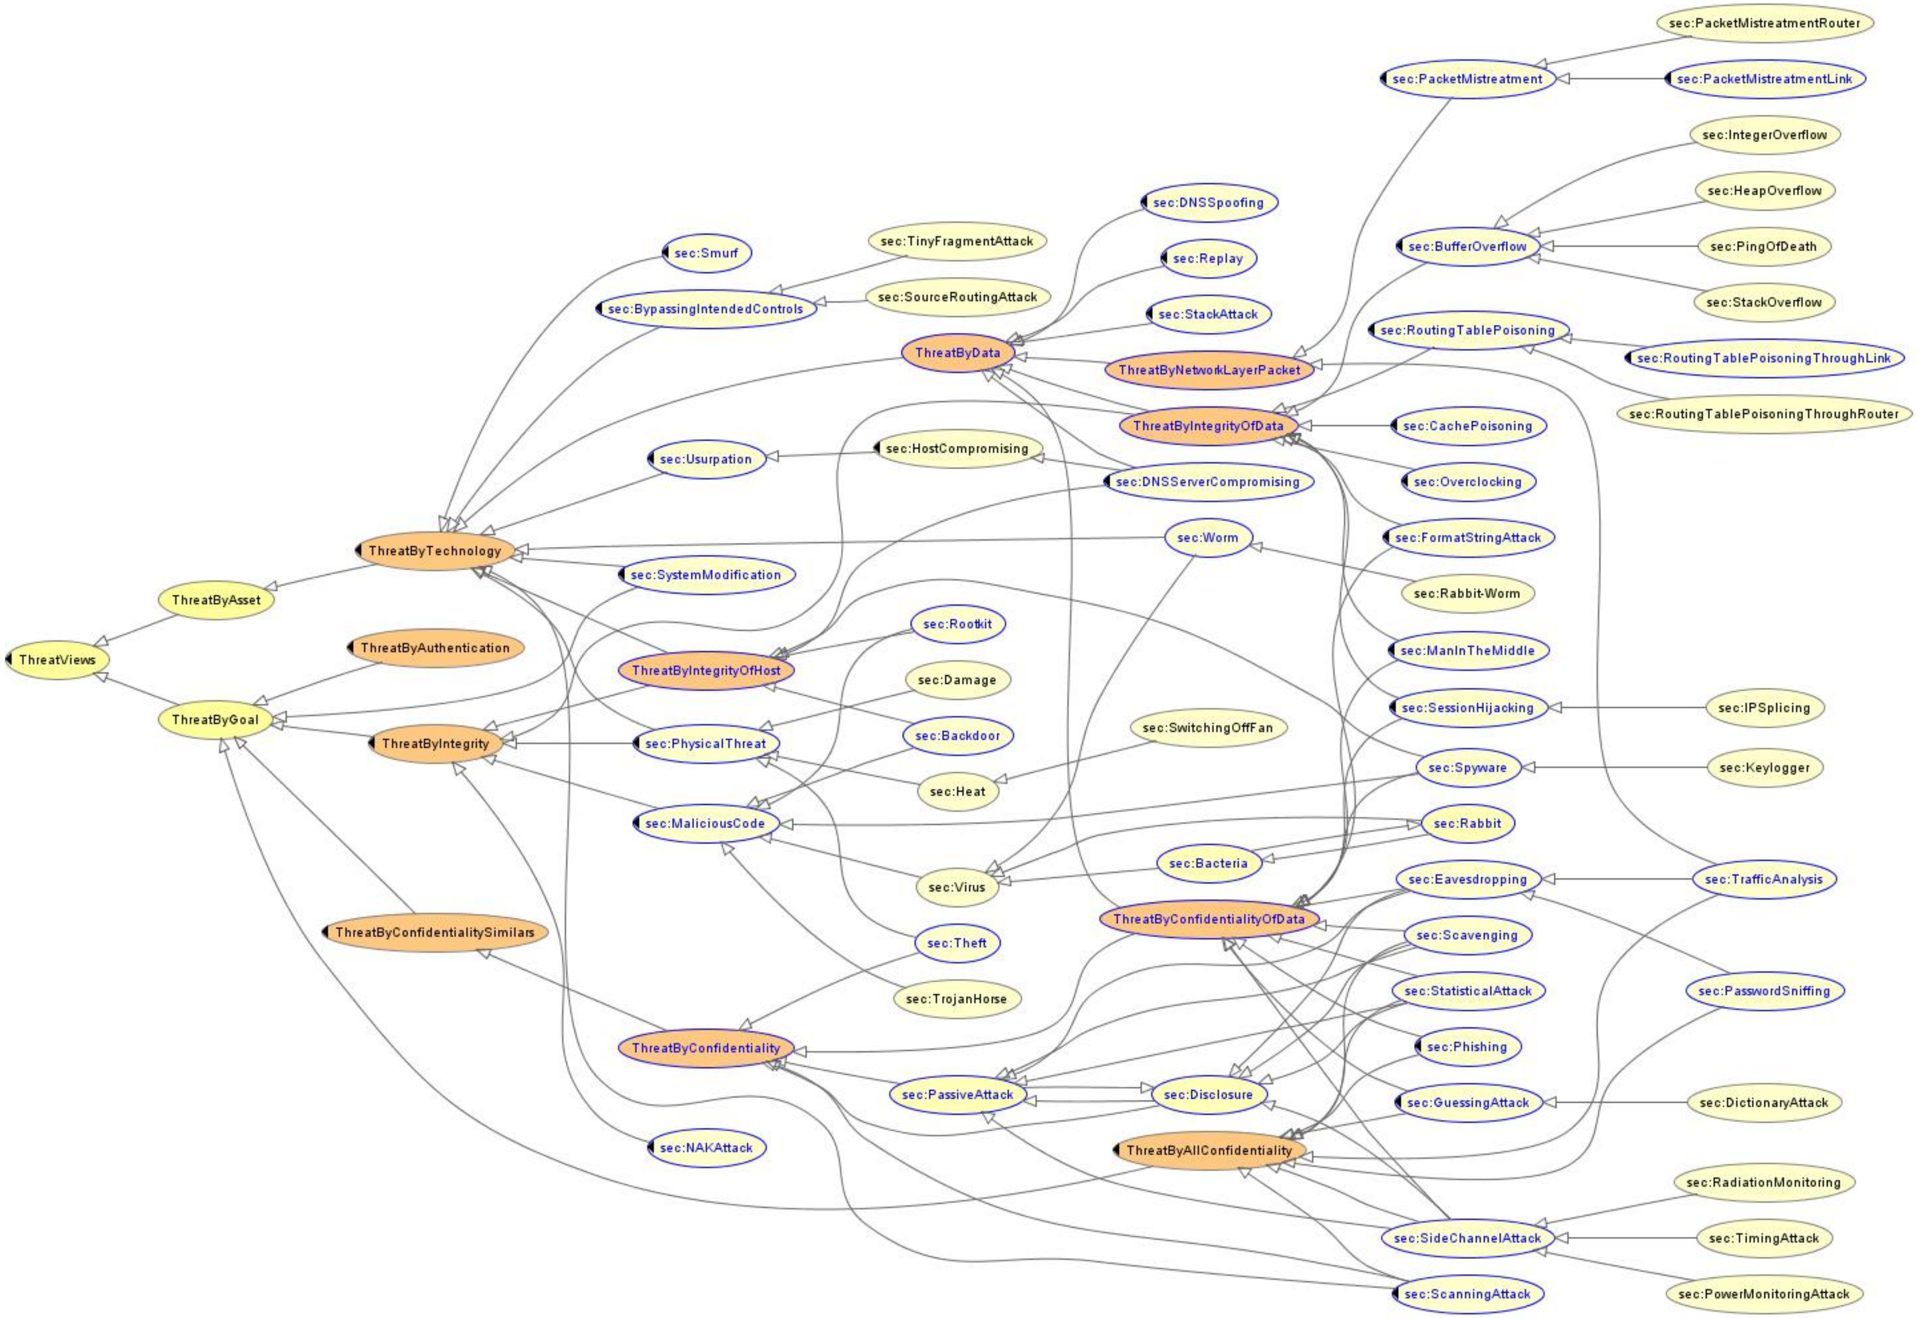

A. Herzog, N. Shahmehri, C. Duma, 'An Ontology of Information Security', International Journal of Information Security and Privacy, 1(4):1-23, 2007. Image retrieved from http://www.ida.liu.se/~iislab/projects/secont/

Contents

- Introduction to reasoning models
  – What are they?
  – Two complementary frameworks: ontologies and BKBs
    - Ontologies
    - BKBs
  – Merging ontologies and BKBs into BKOs
- Under the hood: Ontologies
  – Representation
  – Reasoning
- Under the hood: BKBs
  – Representation
  – Reasoning
  – Unique features
  – Compare/contrast with ontologies
- Prior synthesis attempts
- Under the hood: BKOs
  – Representation w/ example
  – Reasoning w/ example
  – Recap & next steps

Bayesian Knowledge Bases model uncertain variables and their interactions in domain instances

- Use BKBs for…
  - Explaining and predicting intent – getting inside the adversary's head.
  - Decision support models – justifying what's inside your head.
  - Modeling interacting variables and the uncertainty behind them.

- BKBs are a sound model of probabilistic relationships between variables. (Santos & Santos, 1999)
  - Very fine-grained semantics work like networks of "if-then" conditional probability rules. Inherently address incompleteness and cyclic dependencies.
  - Facilitates probabilistically sound fusion of knowledge from multiple, potentially conflicting sources.
  - Reasoning algorithm computes probabilities of variable states.

Bayesian Knowledge Fragment – part of a BKB

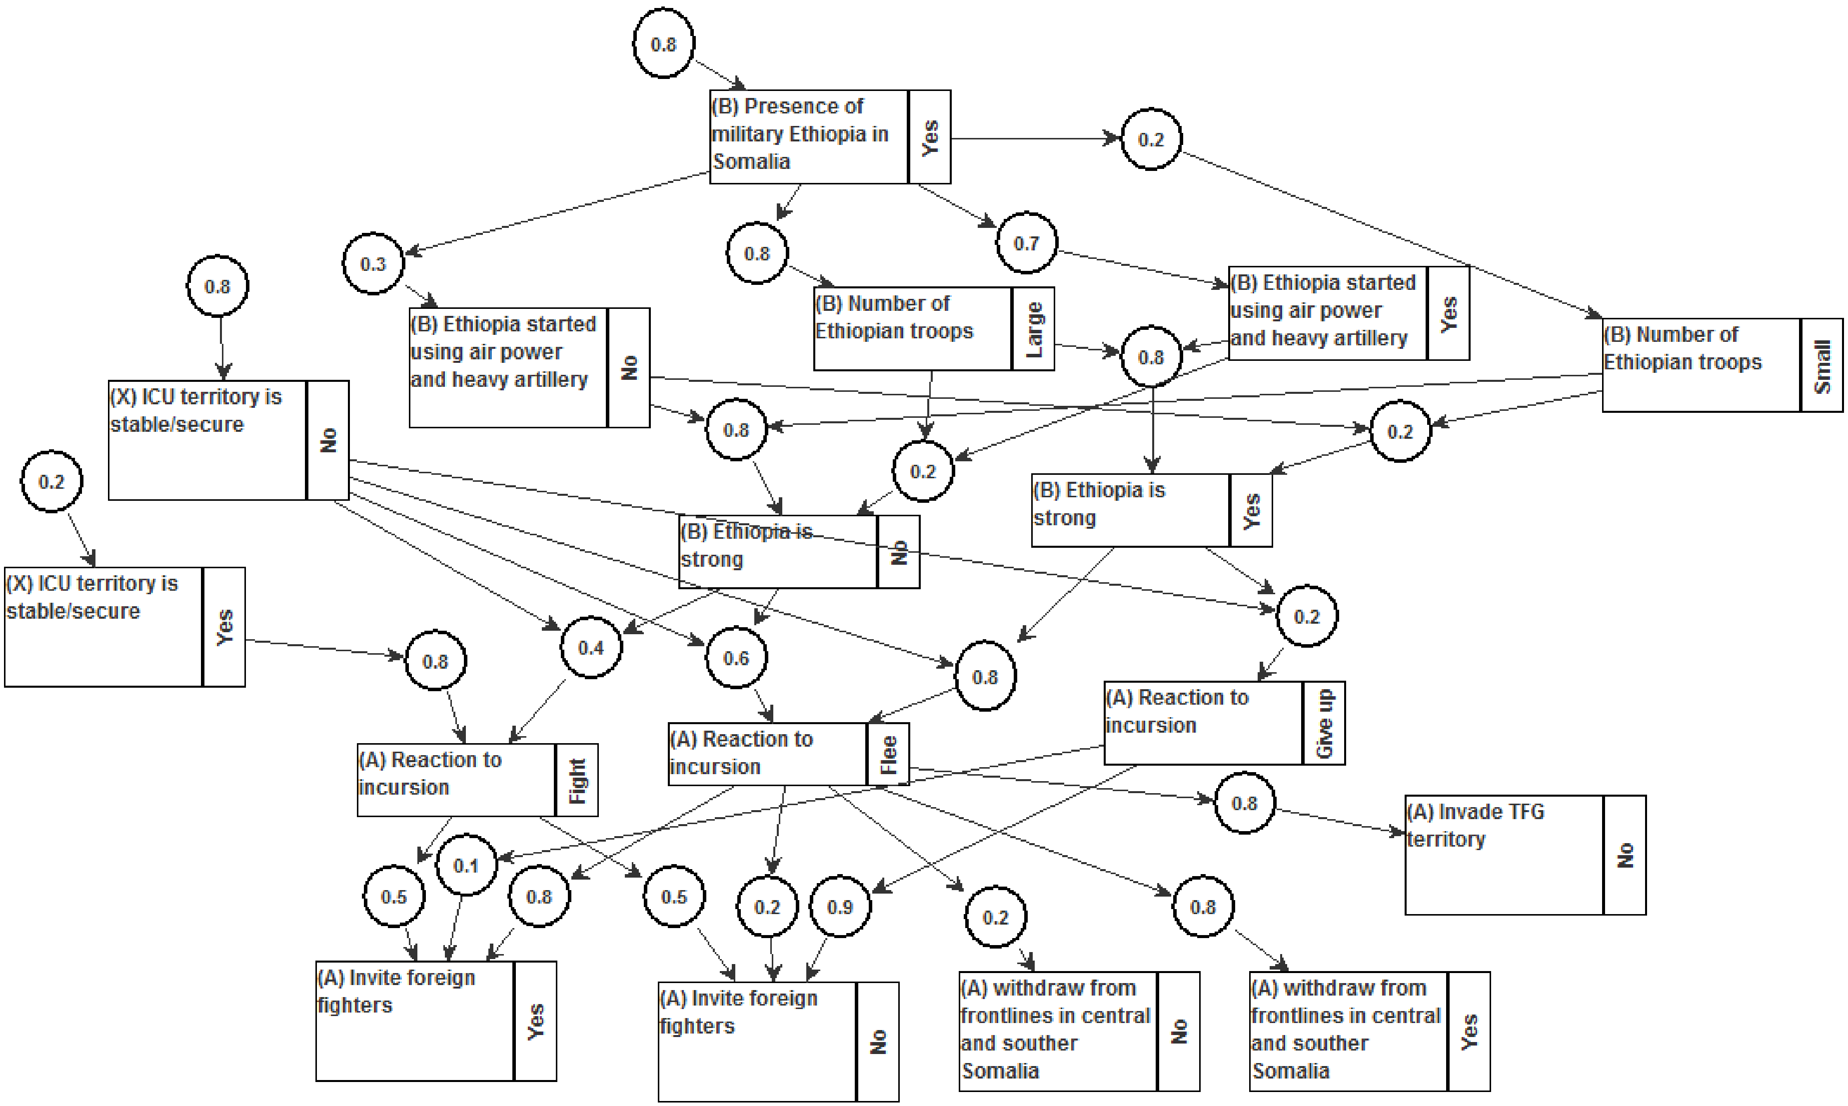

Bayesian Knowledge Bases in real-world applications

- War gaming
  - Modeling the Twentynine Palms base commander in the DENY FORCE scenario. (Lehman et al, 2005)

- Adversarial reasoning
  - Multi-agent simulation of soft factor interactions with the Dynamic Adversarial Gaming Algorithm. (Santos et al, 2007)

- Explaining opinion change
  - Model of voter opinion throughout the South Carolina Democratic Primary election campaign. (Santos et al, 2009)

- Explaining complex historical events
  - Recent project: modeling the 2006 rise and fall of the Islamic Courts Union in Somalia.

Contents

BKBs and ontologies are complementary…

- Ontologies are efficient to work with because they are fluidly reusable, but they can't capture real-world uncertainty.
  - Description Logic infers implicit knowledge, helping the knowledge engineer build and validate the model.
  - An ontology represents all possible states of the world without priority or order, so they are used mostly for domain descriptions.

- BKBs capture real-world uncertainty, but working with them can create duplications of effort.
  - Conditional probability rule semantics can capture subtle variable interactions, like variables which only affect each other under certain conditions, or distributions which are unintuitive with a completeness assumption.
  - BKBs' mechanism for reusing implicated knowledge operates at a coarse granularity, and it must be set up and controlled by a human.

…but difficult to coordinate

- BKBs and ontologies model different but potentially overlapping aspects of a domain. Reconciling them or having them work together is our question.

- Past attempts at making synthesis models using uncertainty theories other than BKBs fell short.
  - Probabilistic description logic’s precision decays during reasoning.
  - BN-ontology combinations contain an assumption conflict – BNs require completeness, but DL does not.
  - Fuzzy and possibilistic semantic networks treat variable interactions coarsely.

Bayesian Knowledge-driven Ontologies combine ontologies' reusability with BKBs' richness

- BKOs contextualize case models in general knowledge, and make inferences about the case from that context.
  - First build a library of general domain knowledge, including uncertain knowledge.
  - Then build case models with just the specifics of the scenarios you want to analyze. The reasoner automatically fills in all the relevant contextual information from the domain library.
  - Concentrate your domain experts on building the library. Analysts can work with the case models.

- BKOs work by inferring new probabilistic rules from general knowledge that admits uncertainty.
  - BKO rules are conditional probability rules between DL assertions.
  - BKO reasoning is an extension of DL reasoning to validly preserve probabilities and dependencies in inferences according to BKBs' semantics. The result of this reasoning is guaranteed to be a valid BKB, and can be analyzed as such.

Contents

Ontologies represent concepts and their relationships

- Formal basis is Description Logic (Baader et al, 2007), a decidable subset of first-order predicate logic.

- Define terminological (general) knowledge in terms of concepts. Can make new definitions or use expressions in terms of other concepts.
  - Common expression operators: intersection, union, negation, existential quantification, value restriction, more
  - Ex: a backdoor intrusion could be defined as the intersection of the concepts “malicious code” and “host integrity threat”:

    Backdoor $\subseteq$ **Intrusion**

    Backdoor $\equiv$ **Malicious_Code** $\cap$ **Host_Integrity_Threat**

Representation (cont'd)

- Define assertional (case-specific) knowledge in terms of individuals described by concepts and linked by relational operators called properties, AKA roles.
  - Ex: Describing a particular kinetic attack's weapon and delivery method:

    The_Attack $\in$ **Kinetic_Attack**

    The_Attack *hasWeapon* Improvised_ANFO

    The_Attack *hasDeliveryMethod* Truck_Bomb

Representation (cont'd)

- Powerful feature: concept constructors
  - Can create new concepts as functions of existing ones and use them in assertions.
  - A property assertion inherently constructs a concept of individuals with that same property and target. These can even be used in constructor expressions.
  - Ex:

    The_Attack ∈ (**Kinetic_Attack ∩ (hasWeapon Improvised_ANFO)**)

Contents

- Introduction to reasoning models
  - What are they?
  - Two complementary frameworks: ontologies and BKBs
  - Merging ontologies and BKBs into BKOs
- Under the hood: Ontologies
  - Representation
  - Reasoning
- Under the hood: BKBs
  - Representation & reasoning
  - Unique features
  - Compare/contrast with ontologies
- Prior synthesis attempts
- Under the hood: BKOs
  - Representation w/ example
  - Reasoning w/ example
  - Recap & next steps

Terminological reasoning makes inferences about general knowledge

- Can relate concepts with relationships of subsumption, disjointness, and equivalence, then try to infer more relationships.

  – Ex: Take "kinetic" and "informational" intrusions as disjoint concepts…

    Kinetic $\neq$ Informational

    both subsumed by the more general "intrusion" concept, which is equivalent to the "attack" concept.

    Kinetic, Informational $\subseteq$ Intrusion

    Intrusion $\equiv$ Attack

    Then infer that "kinetic" and "informational" are both subsumed by "attack" as well.

    $\rightarrow$ **Kinetic, Informational $\subseteq$ Attack**

Assertional reasoning applies general domain knowledge to describe an instance.

- Can describe individuals' concept membership with sufficient and necessary conditions.
  - Sufficient conditions: if an individual meets these, it is described by that concept.
  - Necessary conditions: if an individual is a member of a concept, these are also true for it.

- Uses the rule of universal instantiation: if something is true for the general case, it's true for all specific cases.
  - If an individual meets a concept's sufficient conditions, it's a member of that concept.
  - If a concept has necessary conditions, infer that they are met by all the concept's individuals.
  - Repeat until the ontology stabilizes.

Assertional reasoning (cont'd)

- Ex: The bomb used in a kinetic attack used C4, so it is described by the concept C4_Bomb (sufficient condition). C4_Bombs always come from an Insider_Supplier (necessary condition), so infer that this particular bomb implies the existance of an insider.

Terminological Knowledge (general)

Concept **C4_Bomb**:

Sufficient condition(s): **Bomb** $\cap$ **(uses_explosive C4)**

Necessary condition(s): *has_source* SOME **Insider_Supplier**

Assertional Knowledge (case-specific)

The_Bomb *uses_explosive* C4

Inferred Assertional Knowledge

→ The_Bomb *has_source* SOME **Insider_Supplier**

Advanced DL features for our future work

- More sophisticated DL theories support more capabilities.
  - Role hierarchies
  - Role inclusion expressions
  - Cardinality restrictions
  - Etc.

- Some ontology systems depart from strict DL for more flexibility.
  - Role attributes: roles can be transitive, reflexive, disjoint, inverse…
  - Role domains and ranges
  - Often cannot guarantee decidability

Contents

BKBs model variable interactions

- Describes possible states of the world with discrete random variables
- Links rv states with conditional probability rules (CPRs) to describe a distribution over possible states of the world. Rules are kept from conflicting by ensuring their conditions are mutually exclusive.
- Incompleteness: conditional probability distributions can be partially undefined.
- Represent intent as (B) Beliefs, (X) Axioms, (G) Goals, (A) Actions.

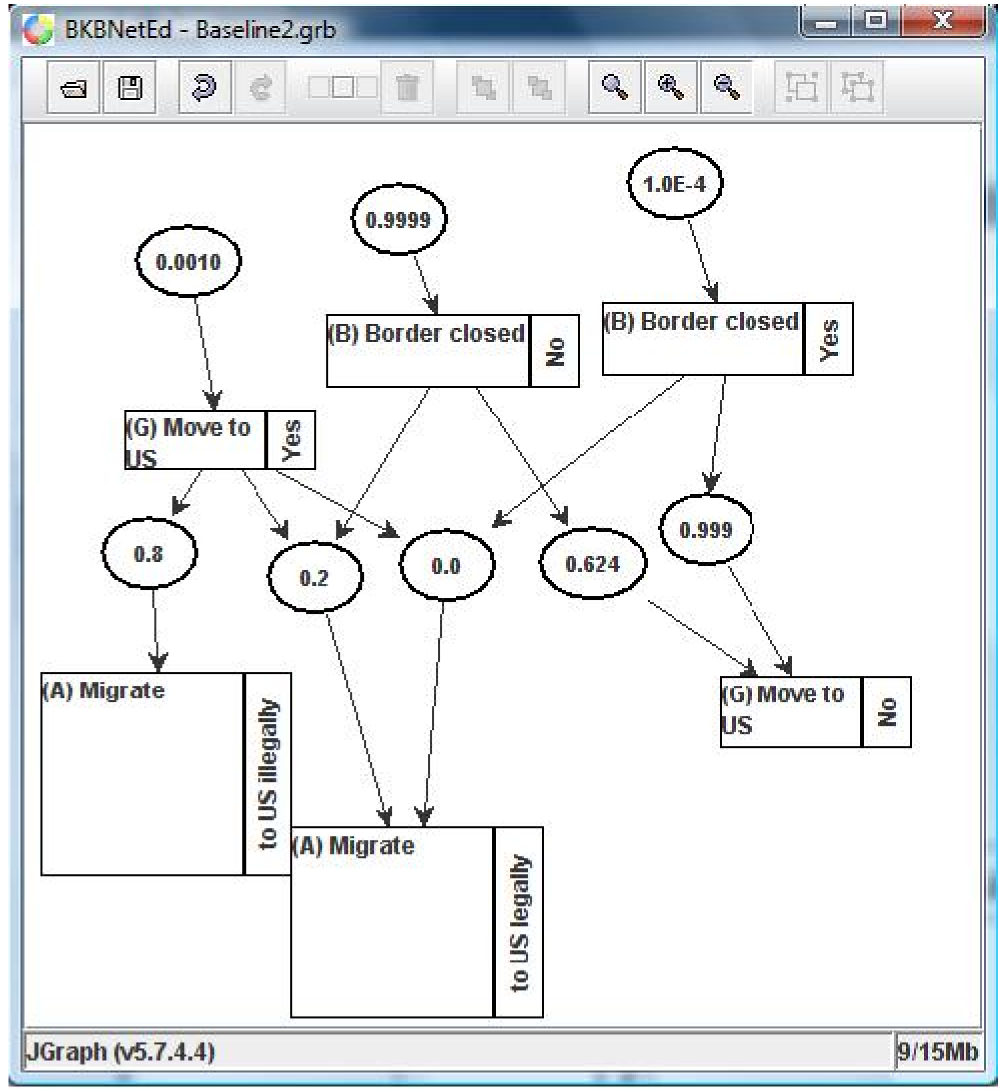

BKB reasoning computes probabilities of combinations of variable states

- Based on computing probabilities of inferences.
  - An inference represents one possible state of the world.
  - An inference is an acyclic chain of rules, such that any rv assignment is supported by only one chain, and none of the rules conflict.

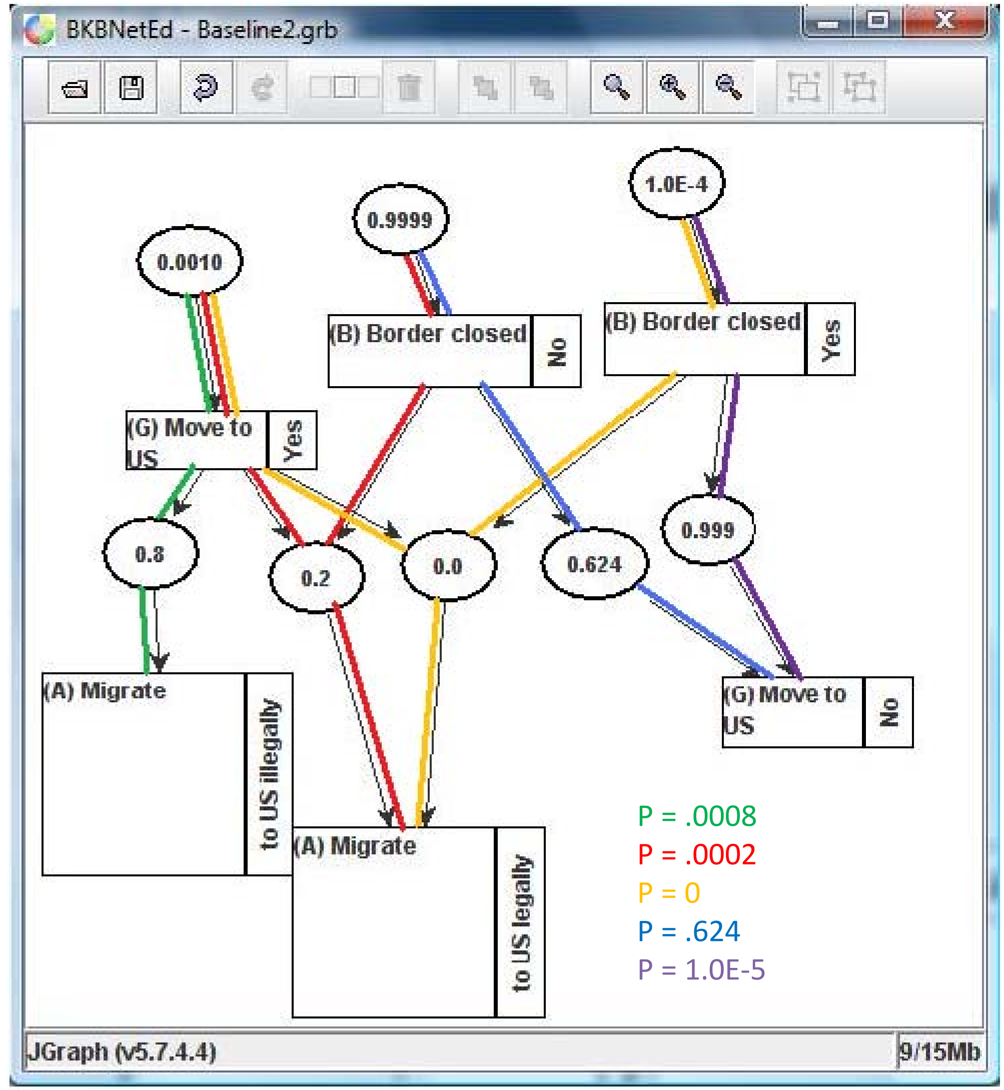

BKBs compute those combinations of variables in three ways

- Belief revision: determine most probable state of the world.
  - P = .624
- Belief updating: compute posterior probability of a single variable assignment.
  - Sum of probabilities of inferences that assignment appears in.
  - P((A) Migrate = to US legally) = 0 + .0002
- Contribution analysis: compute how much one random variable appears as a cause of another.
  - Sum of probabilities of inferences in which the hypothesized cause appears with the effect, divided by the effect's posterior probability from updating.
  - Contribution of (B) Border closed = Yes to (G) Move to US = No: 1.0E-5 / (.624 + 1.0E-5) = 1.6E-5

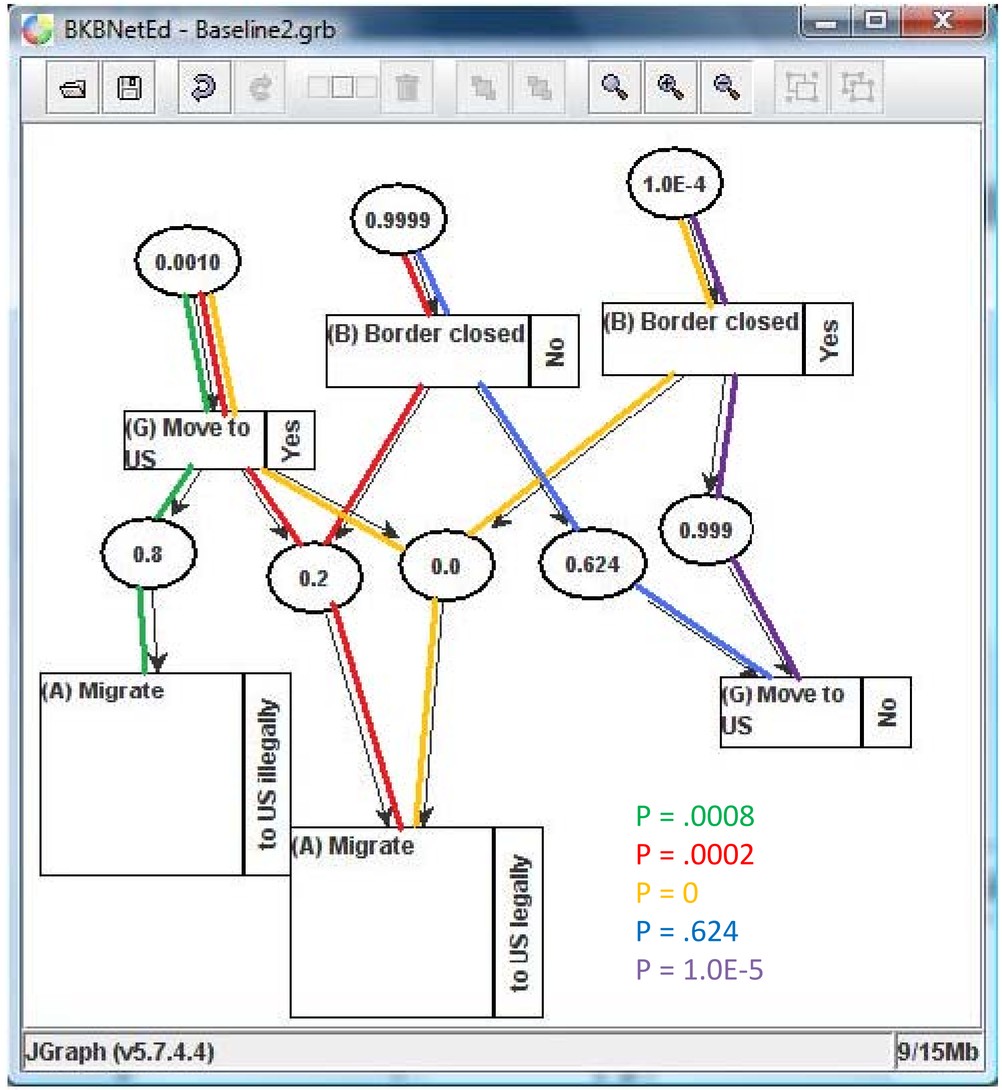

Contents

- Introduction to reasoning models
  - What are they?
  - Two complementary frameworks: ontologies and BKBs
  - Merging ontologies and BKBs into BKOs
- Under the hood: Ontologies
  - Representation
  - Reasoning
- Under the hood: BKBs
  - Representation & reasoning
  - Unique features
  - Compare/contrast with ontologies
- Prior synthesis attempts
- Under the hood: BKOs
  - Representation w/ example
  - Reasoning w/ example
  - Recap & next steps

BKBs capture subtle cyclic relationships that break other knowledge bases

- Can break cycles by having other variable states as additional conditions.

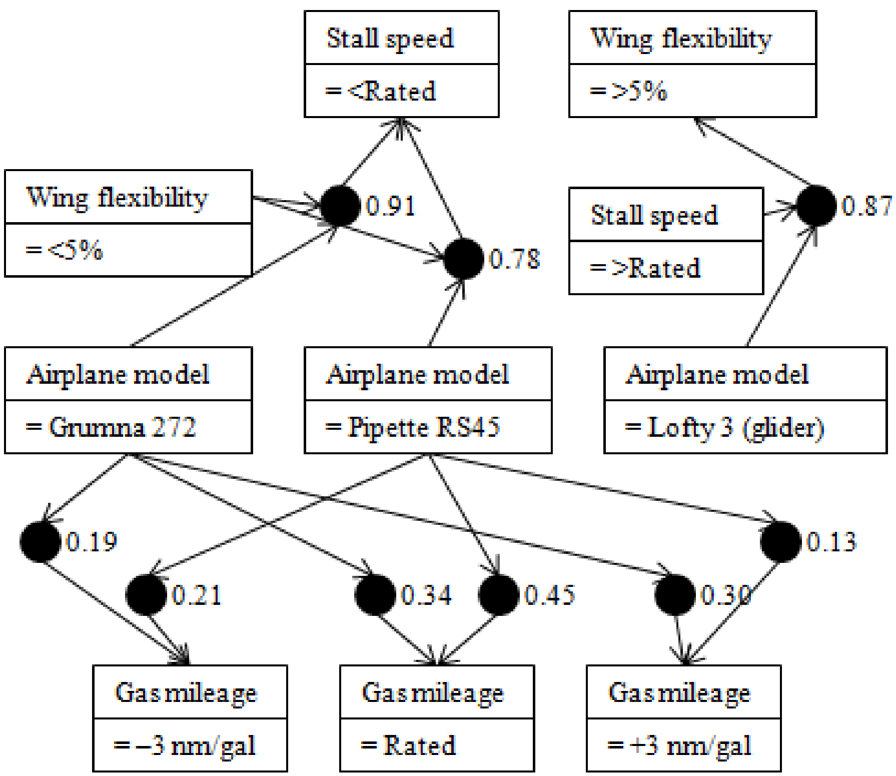

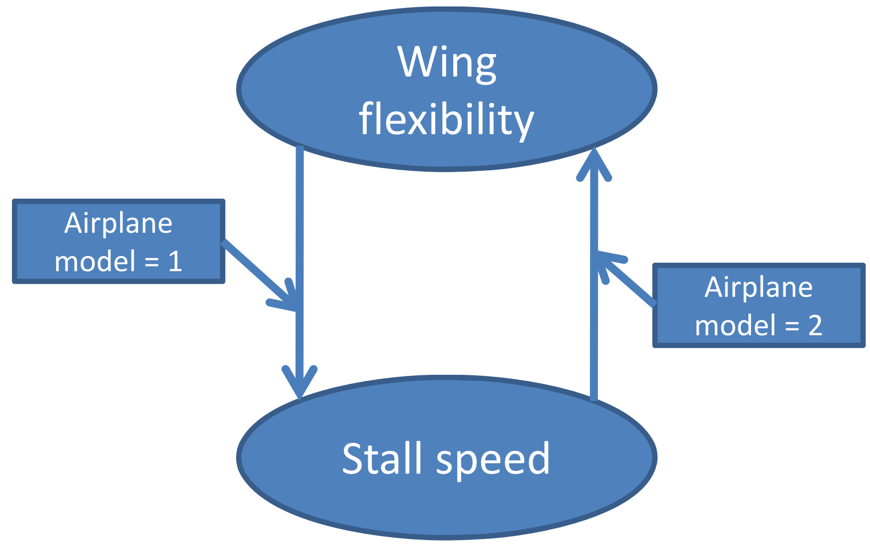

BKBs can fuse knowledge sources and find new insights in the fusion (Santos et al, 2009)

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

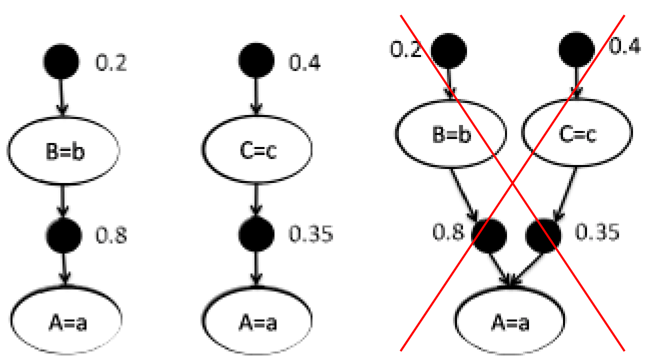

(1) (2)

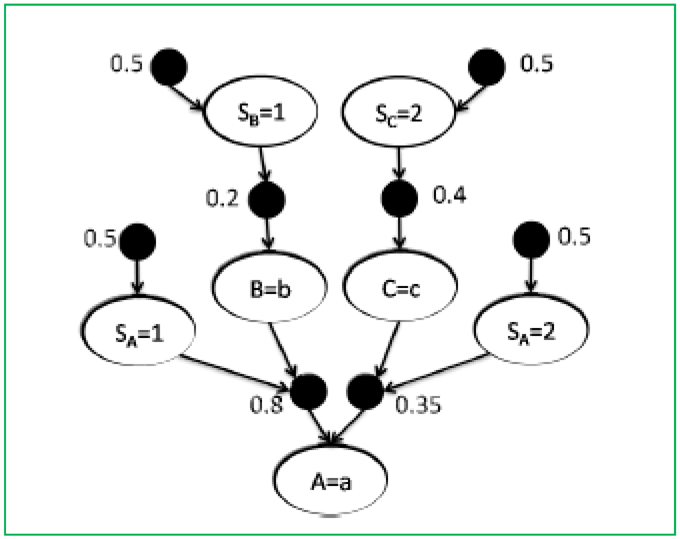

Naïve union of fragments (1) and (2) puts CPRs in conflict. Invalid.

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

Fusion cont'd

- Fusion can generate new inferences, and thereby new explanations of the world.

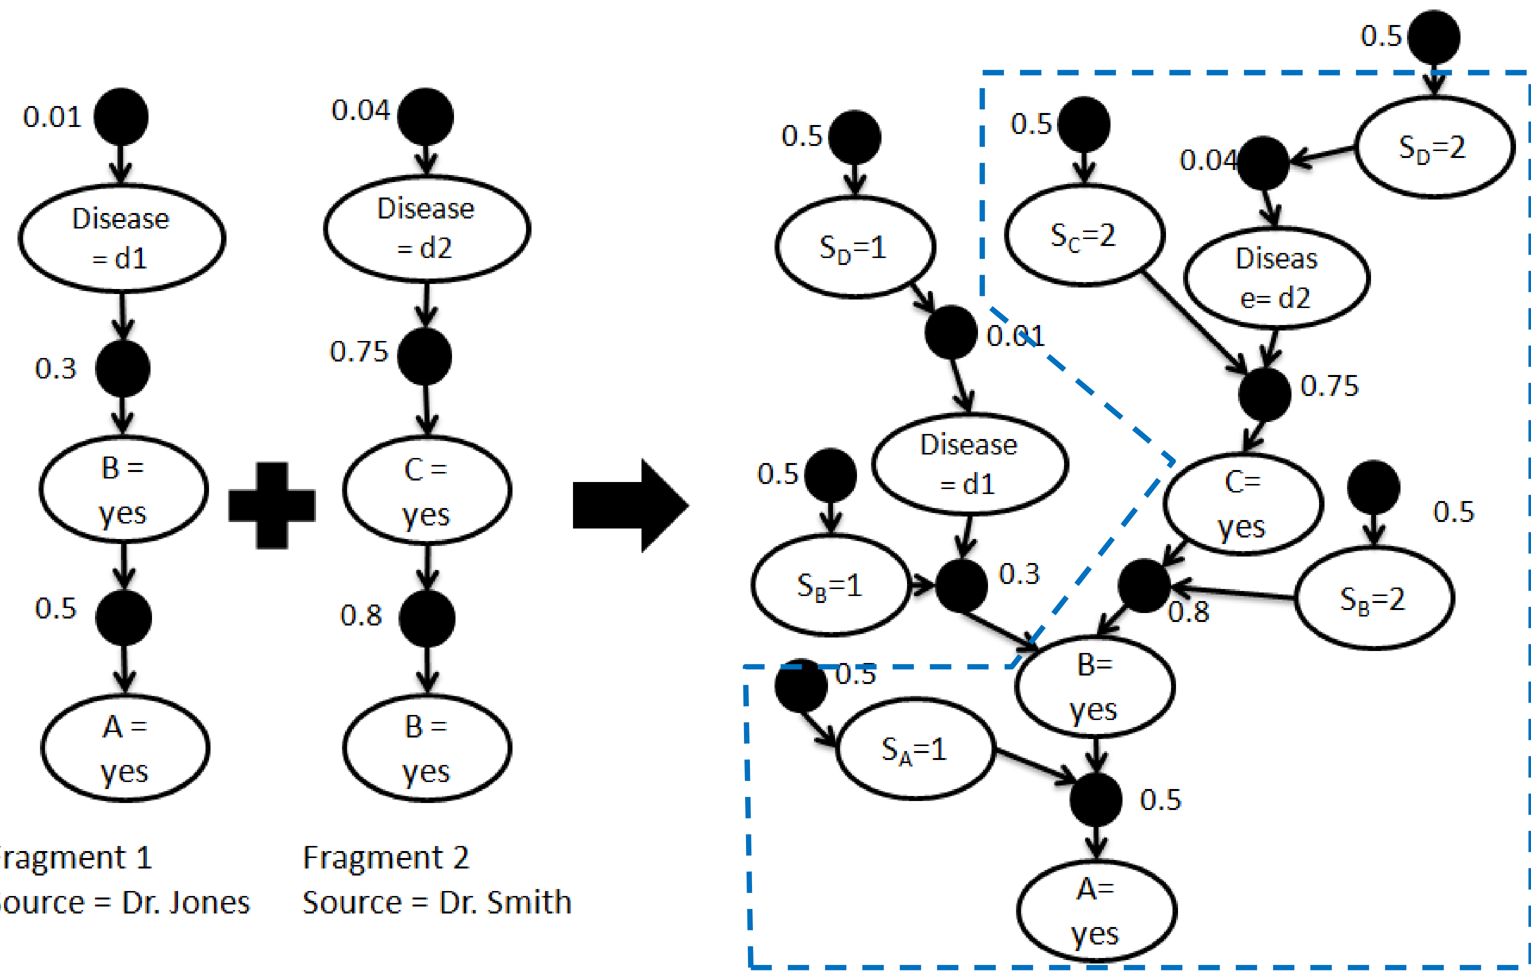

Fusion cont'd

- The fragments' knowledge is meaningfully combined even though they contain a conflict, and new insight into the domain is gained.

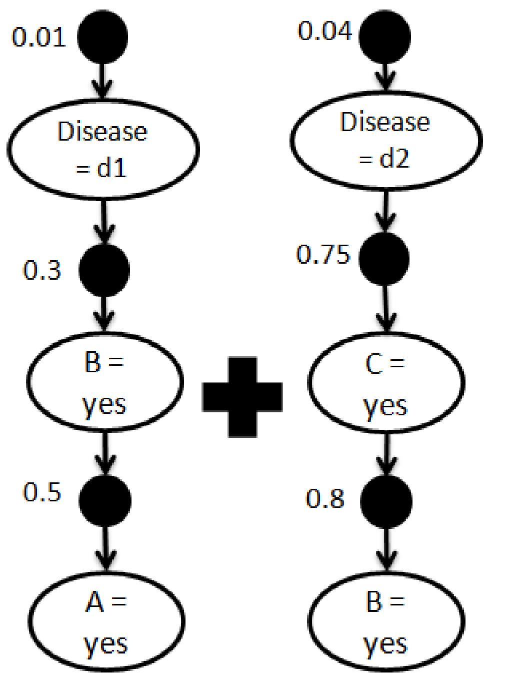

Fragment 1
Source = Dr. Jones

Fragment 2
Source = Dr. Smith

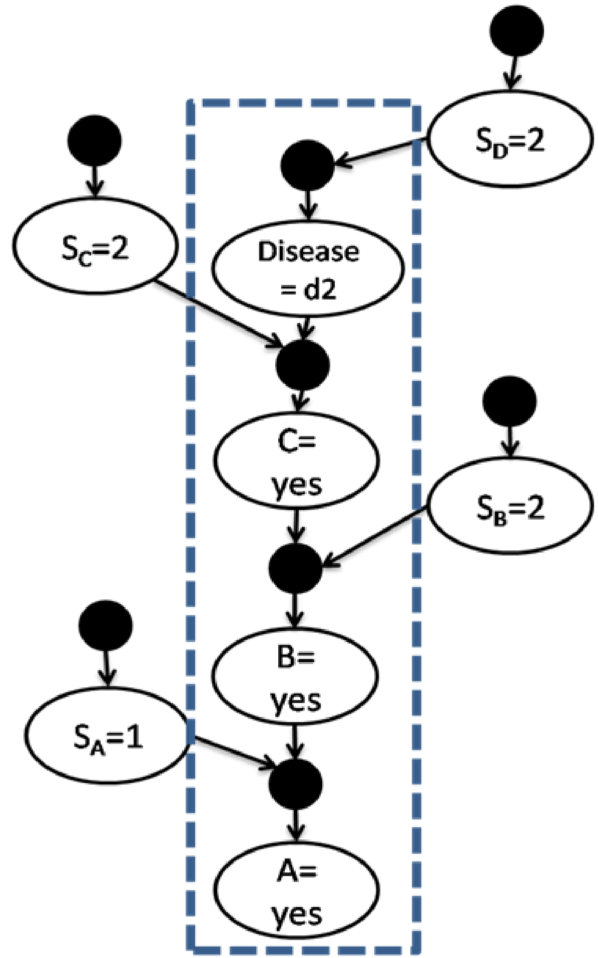

BKBs capture real-world complexity

- Correlation and causality
  - Can build complex webs of variable interactions even in the presence of incompleteness.
- Human intent
  - Our beliefs-axioms-goals-actions framework is an intuitive breakdown of intent.
- Decision processes
  - BKBs depict the reasons for decisions and how they interrelate, then verify that results are as they should be.

Contents

- Introduction to reasoning models
  - What are they?
  - Two complementary frameworks: ontologies and BKBs
  - Merging ontologies and BKBs into BKOs
- Under the hood: Ontologies
  - Representation
  - Reasoning
- Under the hood: BKBs
  - Representation & reasoning
  - Unique features
  - Compare/contrast with ontologies
- Prior synthesis attempts
- Under the hood: BKOs
  - Representation w/ example
  - Reasoning w/ example
  - Recap & next steps

	Ontologies: **Description Logic (DL)**	**Uncertainty:** **Bayesian Knowledge Bases (BKBs)**
Representation mechanism?	Network of lexical concepts and their relationships. Contextual grounding.	If-then conditional probability rules linking variable states. Permits incompleteness.
Reasoning mechanism?	First-order logical inferencing. Finds implicit, inferrable relationships and makes them explicit.	Belief updating and revision compute probability distributions across variable states and world states. Also, contribution analysis.
What do they capture?	General domain knowledge. Terminologies. Cultural perspectives. Snapshots of worldviews.	Individual instances from a domain's possible behaviors. Correlation and causality. Decision processes.
Common applications?	Widespread. Ex: Averting miscommunications in operations planning. Bioinformatics. Semantic web.	War gaming / adversarial reasoning. Explanatory analysis of events.
Limitations?	No uncertainty handling.	Only propositional contextual grounding of variables. No first-order reasoning with concepts.

What should the ideal synthesis do?

- We want a model that describes a probability distribution over possible instances of a domain.
  - Solve that problem of ontologies', that they represent all possible states of the domain simultaneously with no priority or weight.
- We want to be able to do DL reasoning under uncertainty.
  - Infer uncertain conclusions from uncertain premisses. "If the premisses, then (maybe) the conclusion."
  - An "uncertain conclusion from uncertain premisses" is really just a new CPR with the premisses as its conditions.
- We want to make BKBs' CPRs as fluidly reusable with general knowledge as ontologies manage to be.
  - This is actually the same capability as the bullet above, just turned around.

Contents

- Introduction to reasoning models
  - What are they?
  - Two complementary frameworks: ontologies and BKBs
  - Merging ontologies and BKBs into BKOs
- Under the hood: Ontologies
  - Representation
  - Reasoning
- Under the hood: BKBs
  - Representation & reasoning
  - Unique features
  - Compare/contrast with ontologies
- Prior synthesis attempts
- Under the hood: BKOs
  - Representation w/ example
  - Reasoning w/ example
  - Recap & next steps

Probabilistic Description Logic

- Founded on Probabilistic Logic (Nilsson, 1986)
- *Expressive Probabilistic Description Logics* (Lukasiewicz, 2008) is representative of the field.

- Assigns probability intervals to DL assertions.
  - Ex: 0.7 ≤ P(Tweety *is-a* **Bird**) ≤ 0.8
  - Not very intuitive. Uncertainty on an uncertainty metric?
  - Limitation: inferred probability intervals' relative precision (width/mean) decays during forward chaining. This cripples deep reasoning.

Ex:	0.7 ≤ P(Tweety *is-a* **Bird**) ≤ 0.8	RP: 0.1/0.75 = 0.13
	0.9 ≤ P(**Birds** *can* fly) ≤ 0.99	RP: 0.09/0.945=0.095
	→ 0.7 * 0.9 = 0.63 ≤ P(Tweety *can* fly) ≤ 0.8 * 0.99 = 0.79	RP: 0.16/0.71 = 0.22

Bayesian Networks and Ontologies

- Founded on 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 ≥ necessity.

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

Contents

Two technical insights allow DL and BKBs to merge

- Consider: for any individual a and any concept C, either $a \in C$ or $a \in \neg C$. Two exclusive states… sounds like a variable.
  - A state of the world in DL is an assignment of every individual to every class or its complement.
  - Can define a joint probability distribution over these atomic variables. Equivalently, that distribution is over the ontology's possible states.
  - This is the overlap between BKBs and ontologies, and it is the beginning of the answer to our question.

- Generalizing the rule of universal instantiation to its probabilistic analog lets DL reason with uncertainty.
  - "If P(X) = p for each member of a class, then P(X) = p for any particular member of that class."

Gory details (Santos & Jurmain, 2011)

Definition 5.1: For a domain description consisting of a finite set of individuals $\{a_1 \ \ldots \ a_m\}$ and a finite set of distinct atomic classes $\{C_1 \ \ldots \ C_n\}$ and their complements $\{\neg C_1 \ \ldots \neg C_n\}$, define the domain's *state distribution* as a discrete multivariate probability distribution over the variables (referred to as *atomic variables*) V_{ij} whose states are $r(V_{ij}) = \{a_j \in C_i,\ a_j \in \neg C_i\}$. The sample space of the state distribution is defined as the following cross-product:

$$\Omega = \prod_{j=1}^{m} \prod_{i=1}^{n} r(V_{ij}) = \prod_{j=1}^{m} \prod_{i=1}^{n} \{a_j \in C_i, a_j \in \neg C_i\}$$

The reasoner simplifies that verbose structure

- T/F atomic classes are the most expanded description of the domain possible. We can sometimes simplify it.

- Ex: A description of a ball that can have any one of three disjoint colors.
  - Using atomic classes, this needs three variables of six states total.
  - Instead, use a *constructed variable* of three states.

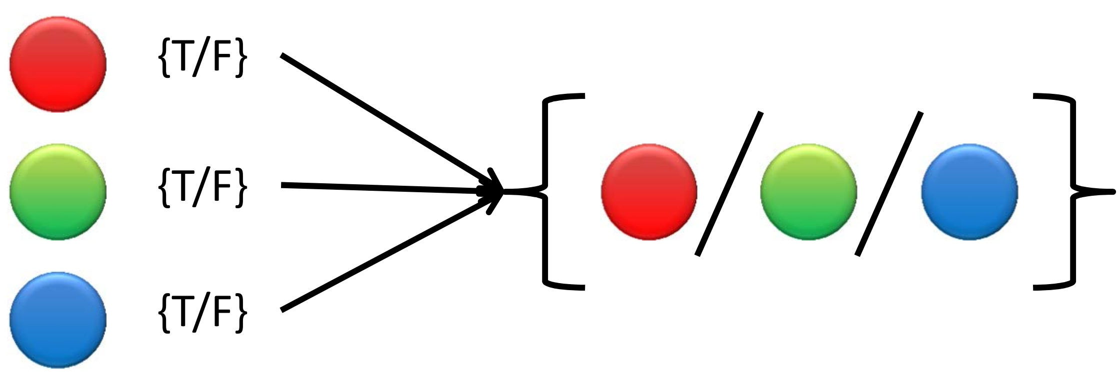

More gory details

Notation: A set of classes $\{C_1 \ \dots \ C_n\}$ is said to *span* a class D if $\cup \{C_1 \dots C_n\} = D$. $\{C_1 \ \dots \ C_n\}$ is said to be *world-spanning* if $\cup \{C_1 \dots C_n\} = \top$.

Definition 5.2: Let Q be a domain containing an individual, a, and a world-spanning set of constructed or atomic classes $\{C_1 \ \dots \ C_n\}$, and let the domain have a state distribution with sample space Ω. Then the set $\{C_1 \ \dots \ C_n\}$'s *constructed variable* is a variable V over Ω such that $r(V) = \{a \in C_1, \dots, a \in C_n\}$.

Asserting Knowledge

- Probabilistic equivalents of assertional (case-specific) and terminological (general) axioms.
  - Probabilistic assertional axiom: the probability of one individual's membership in a concept is p, given some other concept memberships [or not – rules can be unconditional].

    $$P(V_{i_n} = \{a_{i_n} \in B_{i_n}\} | [V_{i_1} = \{a_{i_1} \in B_{i_1}\} \wedge \ldots \wedge V_{i_{n-1}} = \{a_{i_{n-1}} \in B_{i_{n-1}}\}]) = p$$

    Short form:

    $$P(a_{i_n} \in B_{i_n} | [a_{i_1} \in B_{i_1} \wedge \ldots \wedge a_{i_{n-1}} \in B_{i_{n-1}}]) = p$$

  - Probabilistic terminological axiom: any member of one concept has some probability p of also being a member of another concept.

    $$P(x \in D \mid \text{any } x \in C) = p$$

    - C and D can be concept constructor expressions. Complex ones, even. That's the power of a PTA!

BKO example: Modeling Fish, Part 1

- *PAA: An individual tuna and an individual herring are both fish.*

  $\boxed{(1)\quad \text{Tuna, Herring} \in \text{Fish}}$

- *PAA: We are 99% sure we saw the tuna eat the herring.*

  $\boxed{(2)\quad P(\text{Tuna } \textit{ate} \text{ Herring}) = 0.99}$

- *PTA: If something ate a fish, we can be 90% sure it's a predator of fish.*

  $\boxed{(3)\quad P(x \in \text{Predator} \mid \text{any } x \in \textit{ate}_{(\textbf{some} \text{ Fish})}) = 0.9}$

- *PTA: A fish that preys on other fish has a 30% chance of having parasites.*

  $\boxed{(4)\quad P(x \in \text{Has\_Parasites} \mid \text{any } x \in \text{Fish} \cap \text{Predator}) = 0.3}$

Contents

BKO reasoning has two stages

- First, logical reasoning.
  - Reason on subsumption, equivalence, and disjointness as in conventional DL.
  - Reason on assertional information using the probabilistic rule of universal instantiation. (This is that “automatic self-construction” we’ve talked about.)
  - Prune unsupported inferences.

- When all logical reasoning is complete, the result is equivalent to a BKB.
  - Can now perform belief revision, updating, contribution analysis, and perhaps more, as desired.

BKO reasoning preserves conditional probabilities and prunes unsupported rules

- Applying the probabilistic rule of universal instantiation:
  - Generate *instantiations* of probabilistic terminological axioms.
    - PTAs are not CPRs. But they do generate CPRs when applied to specific individuals.

      Definition 5.5: The instantiation of a PTA $T: P(x \in D \mid \text{any } x \in C) = p$ for an individual a is defined as $T|_a: P(a \in D \mid a \in C) = p$.
  - This is where the domain library comes in. A big domain library, instantiated on a case model, can yield a lot of PTA instantiations. This is how contextual knowledge gets imported to an analyst's case model!

- Pruning:
  - In principle, BKO reasoning simply instantiates all PTAs for all individuals, then removes any of them with unsupported conditions. This would consume a lot of memory, so in practice it will check for a potential support chain as it goes.

Modeling Fish, Part 2

- Recall:

(1) Tuna, Herring $\in$ Fish
(2) P(Tuna *ate* Herring) = 0.99
(3) P(x $\in$ Predator | any x $\in$ $ate_{(\text{some Fish})}$) = 0.9
(4) P(x $\in$ Has_Parasites | any x $\in$ Fish $\cap$ Predator) = 0.3

- Instantiate PTAs:

(5) P(Tuna $\in$ Predator | Tuna $ate_{(\text{some Fish})}$) = 0.9
(6) P(Tuna $\in$ Has_Parasites | Tuna $\in$ Fish $\cap$ Predator) = 0.3
(7) P(Herring $\in$ Predator | Herring $ate_{(\text{some Fish})}$) = 0.9
(8) P(Herring $\in$ Has_Parasites | Herring $\in$ Fish $\cap$ Predator) = 0.3

- Problem: 5-8 don't quite link up with 1-4. But the connections are common sense.

Modeling Fish, Part 2 (cont'd)

(1) Tuna, Herring $\in$ Fish
(2) P(Tuna *ate* Herring) = 0.99
(3) P(x $\in$ Predator | any x $\in$ $ate_{(\text{some Fish})}$) = 0.9
(4) P(x $\in$ Has_Parasites | any x $\in$ Fish $\cap$ Predator) = 0.3
(5) P(Tuna $\in$ Predator | Tuna $ate_{(\text{some Fish})}$) = 0.9
(6) P(Tuna $\in$ Has_Parasites | Tuna $\in$ Fish $\cap$ Predator) = 0.3
~~(7) P(Herring $\in$ Predator | Herring $ate_{(\text{some Fish})}$) = 0.9~~
~~(8) P(Herring $\in$ Has_Parasites | Herring $\in$ Fish $\cap$ Predator) = 0.3~~

- We must infer "bridging" axioms from semantic knowledge to explicitly link 5-8 with 1-4:

(9) P(Tuna $\in$ $ate_{(\text{some Fish})}$ | Tuna *ate* Herring $\wedge$ Herring $\in$ Fish) = 1
(10) P(Tuna $\in$ Fish $\cap$ Predator | Tuna $\in$ Fish $\wedge$ Tuna $\in$ Predator) = 1
~~(11) P(Herring $\in$ $ate_{(\text{some Fish})}$ | Herring *ate* ? $\wedge$? $\in$ Fish) = 1~~
~~(12) P(Herring $\in$ Fish $\cap$ Predator | Herring $\in$ Fish $\wedge$ Herring $\in$ Predator) = 1~~

- (11)'s solution is not in the knowledge base. Pruning it renders (7) unsupported, then (12), then (8).

What’s left after pruning is a BKB

- After logical reasoning, all general knowledge has been turned into case-specific knowledge. It’s now redundant, so eliminate it.
- What’s left is pure CPRs, and BKO theory’s formulation guarantees they form a BKB. (Santos & Jurmain, 2011)
- Reason on this BKO-generated BKB like any other!

Modeling Fish, Part 3

- Now the unpruned PAAs form a BKB.

$$\left.\begin{array}{ll}(1) & \text{Tuna, Herring} \in \text{Fish} \\ (2) & P(\text{Tuna } \textit{ate} \text{ Herring}) = 0.99\end{array}\right\} \text{Original case model}$$

$$\left.\begin{array}{ll}(5) & P(\text{Tuna} \in \text{Predator} \mid \text{Tuna } \textit{ate}_{(\textbf{some} \text{ Fish})}) = 0.9 \\ (6) & P(\text{Tuna} \in \text{Has\_Parasites} \mid \text{Tuna} \in \text{Fish} \cap \text{Predator}) = 0.3\end{array}\right\} \text{PTA instantiations}$$

$$\left.\begin{array}{ll}(9) & P(\text{Tuna} \in \textit{ate}_{(\textbf{some} \text{ Fish})} \mid \text{Tuna } \textit{ate} \text{ Herring} \wedge \text{Herring} \in \text{Fish}) = 1 \\ (10) & P(\text{Tuna} \in \text{Fish} \cap \text{Predator} \mid \text{Tuna} \in \text{Fish} \wedge \text{Tuna} \in \text{Predator}) = 1\end{array}\right\} \text{Bridging axioms}$$

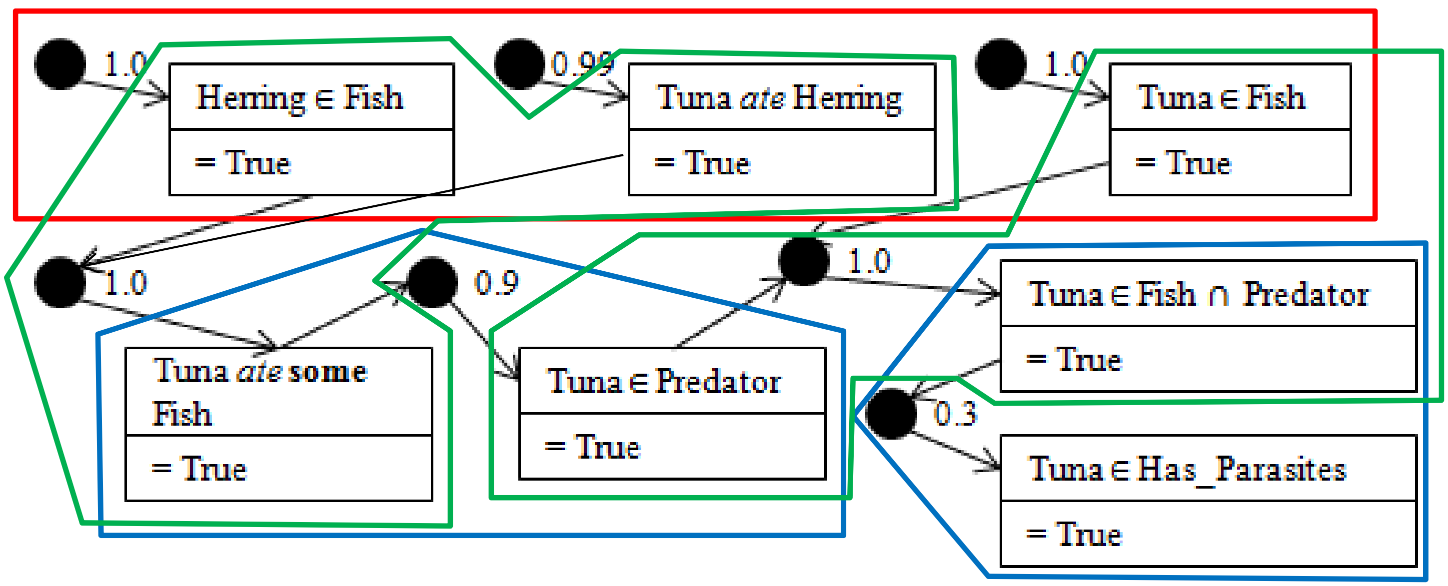

Contents

BKO Theory Summary

- Rich, flexible representation.
  - Formally grounded CPRs, AKA probabilistic DL assertions
  - Concept constructors facilitate easy expression of complex concepts.
- Powerful reasoning.
  - Logical reasoning can automatically assemble and fuse large amounts of case-relevant information starting from a small seed.
  - BKB reasoning facilitates several analyses for explaining complex variable interactions.
- Principled approach to instantiating any number of varied BKB behavior models from the domain.
  - Make small changes (various Twentynine Palms commanders)
  - Make large changes (change from an air battle to a ground battle)

Next Development Steps

- Additional theory development to incorporate advanced ontology and BKB capabilities.
  - Advanced ontology languages
  - BKB fusion, sensitivity analysis, interfacing with social networks…
- Optimization of reasoning algorithm
- Software development.
  - Representation language/format for domain libraries and case models
  - Reasoning application

Far Future Applications

- Robotics
  - Representation and reasoning with conceptual, uncertain, incomplete knowledge could facilitate robots' handling of more unpredictable environments.

- Machine learning
  - 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.

References

- F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi and P. F. Patel-Schneider. *The Description Logic Handbook. Cambridge: Cambridge University Press, 2007.*
- Santos, Eugene, Jr. and Santos, Eugene, S., "A framework for Building Knowledge-Based Under Uncertainty," Journal of Experimental and Theoretical Artificial Intelligence, 11, 265-286,1999.
- Lehman, Lynn A., Krause, Lee S., Gilmour, Duane A., Santos, Eugene, Jr., and Zhao, Qunhua, "Intent Driven Adversarial Modeling," Proceedings of the Tenth International Command and Control Research and Technology Symposium: The Future of C2, McLean, VA, 2005.
- Santos, Eugene, Jr., Zhao, Qunhua, Pratto, Felicia, Pearson, Adam R., McQueary, Bruce, Breeden, Andy, and Krause, Lee, "Modeling Multiple Communities of Interest for Interactive Simulation and Gaming: The Dynamic Adversarial Gaming Algorithm Project," Proceedings of the SPIE: Defense & Security Symposium, Vol: 6564, Orlando, FL, 2007.
- Santos, Eunice E., Santos, Eugene, Jr., Wilkinson, John T., and Xia, Huadong, "On a Framework for the Prediction and Explanation of Changing Opinions," Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, 1146-1452, San Antonio, TX, 2009.
- E. Santos, Jr., J.T. Wilkinson, and E. E. Santos, "Bayesian knowledge fusion," in Proc. 22nd International FLAIRS Conference. AAAI Press, 2009.
- N. J. Nilsson. “Probabilistic logic.” *Artif. Intell. vol. 28(1), pp. 71-87, 1986.*
- T. Lukasiewicz, “Expressive probabilistic description logics,” *Artif. Intell. vol. 172(6-7), pp. 852-883, 2008.*
- J. Pearl. “Bayesian networks: A model of self-activated memory for evidential reasoning.” (UCLA Technical Report CSD-850017). In *Proceedings of the 7th Conference of the Cognitive Science Society, 1985, pp. 329–334.*

References cont'd

- P. C. G. Costa and K. B. Laskey, "PR-OWL: A framework for probabilistic ontologies." in *Proc. FOIS, 2006, pp. 237-249.*
- Z. Ding, Y. Peng, R. Pan, Z. Ding, Y. Peng and R. Pan, "BayesOWL: Uncertainty modeling in semantic web ontologies," presented at Soft Computing in Ontologies and Semantic Web, 2005.
- D. Koller, A. Levy and A. Pfeffer, "P-CLASSIC: A tractable probabilistic description logic," presented at AAAI, 1997.
- Boutilier, C., Friedman, N., Goldszmidt, M. and Koller, D. (1996) Context-specific independence in bayesian networks, *Proc. Conf. Uncertainty in AI.*
- U. Straccia. "Reasoning within fuzzy description logics." *Journal of Artificial Intelligence Research, vol. 14, 2001.*
- G. Qi, J. Z. Pan and Q. Ji. "A possibilistic extension of description logics." In *Proc. DL, 2007 .*
- Zadeh, L.A. (1965). "Fuzzy sets", *Information and Control* 8 (3): 338–353.
- Zadeh, Lotfi, "Fuzzy Sets as the Basis for a Theory of Possibility", *Fuzzy Sets and Systems* 1:3-28, 1978. (Reprinted in *Fuzzy Sets and Systems* 100 (Supplement): 9-34, 1999.)