Knowledge Representation
A very brief intro

Jacek Malec
Dept. of Computer Science, Lund University

February 8, 2017

IBM Watson example

https://www.youtube.com/watch?v=DywO4zksfXw

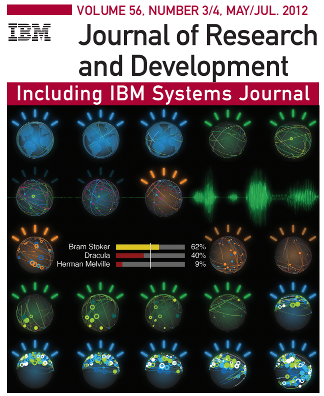

This Is Watson

Knowrob: Why is knowledge so important?

- if the robot does not know about the task, the environment, or the robot, then the programmer has to hardcode everything

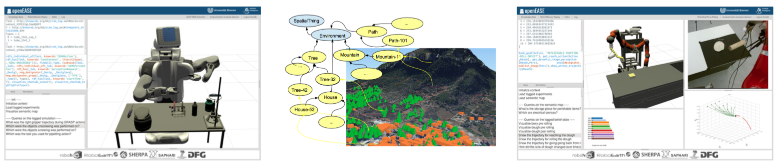

- programming/instructing at an abstract/semantic level
  - put the bolt into the nut and fasten it
  - pour water into the glass
  - ...

Knowrob: Ontology (knowrob.owl)

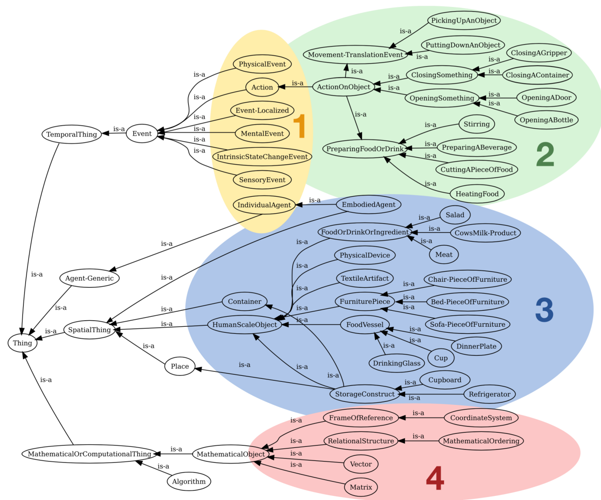

Knowrob: A task ontology

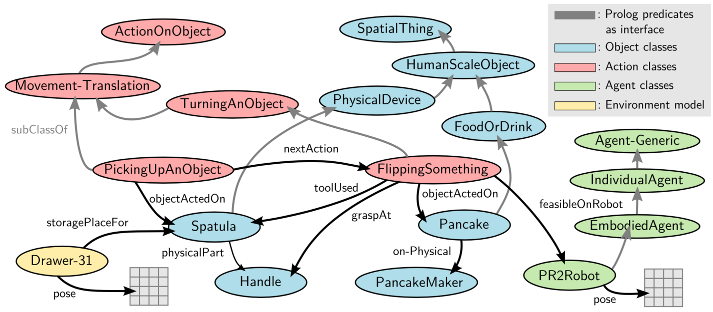

Knowrob: A task ontology

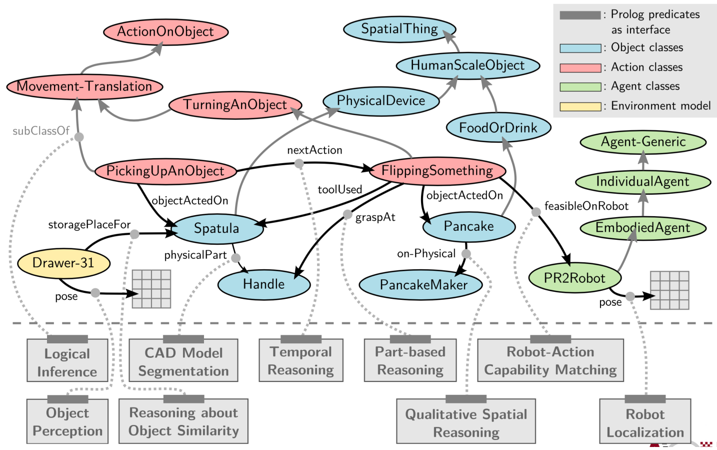

Knowrob: Knowledge types

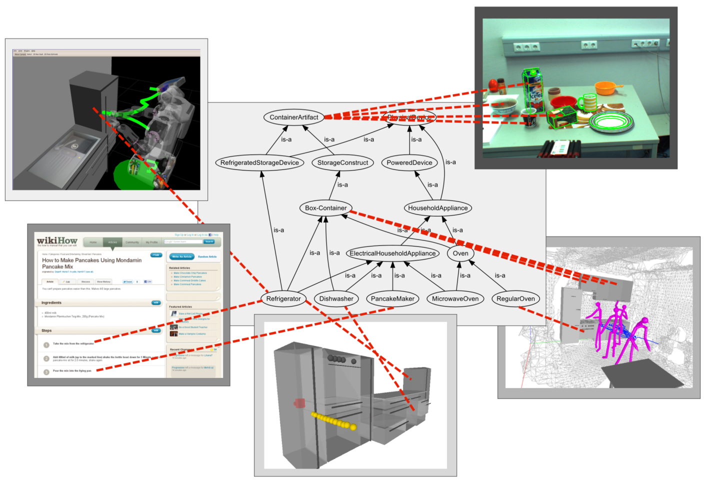

KnowRob Components

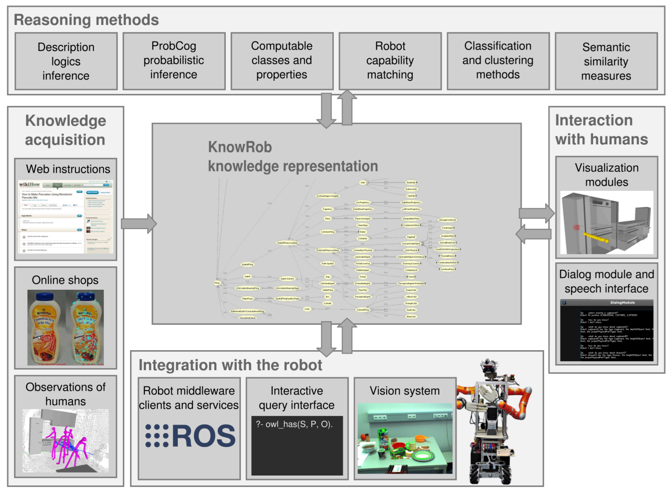

Knowrob: Procedural attachments

- Compute symbolic knowledge on demand from data structures that already exist on the robot by attaching procedures to semantic classes and properties
- Re-use existing information and make sure abstract knowledge is grounded

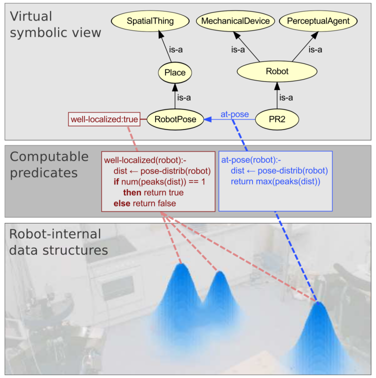

Knowrob: Inferring storage location

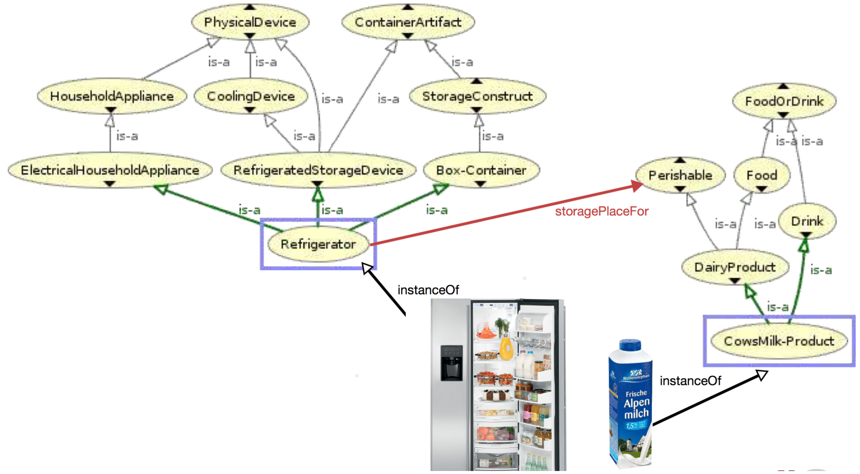

Knowrob: Summary

- declarative knowledge: ontologies
- procedural attachment
- logical inference
- multi-modal representation

Video (13 mins):
`https://www.youtube.com/watch?v=4usoE981e7I`

Plan for today

1. Knowledge-based systems
   - Tacit knowledge
   - Inferred knowledge
   - Domain-specific stuff
   - Changing premises
   - Uncertainty
   - Semantic anchoring
2. Architectures
3. Self-awareness

Tacit knowledge

Facts about:

Tacit knowledge

Facts about:

- objects

Tacit knowledge

Facts about:

- objects
- places

Tacit knowledge

Facts about:

- objects
- places
- times

Tacit knowledge

Facts about:

- objects
- places
- times
- events
- processes
- behaviours

Tacit knowledge

Facts about:

- objects
- places
- times
- events
- processes
- behaviours
- vehicle dynamics
- rigid body interactions
- traffic laws
- . . .

Tacit knowledge

Background knowledge for all this includes:

Tacit knowledge

Background knowledge for all this includes:

- ontologies

Tacit knowledge

Background knowledge for all this includes:

- ontologies
- theories

Tacit knowledge

Background knowledge for all this includes:

- ontologies
- theories
- physics
- mereology
- ...

Tacit knowledge

Background knowledge for all this includes:

- ontologies
- theories
- physics
- mereology
- ...

Not everything needs to be explicit, nor expressed in one monolithic formalism

Inferred knowledge

(or: turning implicit into explicit)

1. logics (language)
2. theorem proving (mechanics)
3. modes of reasoning

Logics: modal

1. take a logical language, let α be a wff
2. $\Box\alpha$ is a wff
3. $\Diamond\alpha$ is a wff
4. normally $\Box\alpha \leftrightarrow \neg\Diamond\neg\alpha$

Intended meaning?

Knowledge Representation

Logics: modal

1. take a logical language, let α be a wff
2. $\Box\alpha$ is a wff
3. $\Diamond\alpha$ is a wff
4. normally $\Box\alpha \leftrightarrow \neg\Diamond\neg\alpha$

Intended meaning?

1. $\Box\alpha$ means **Necessarily** α

Knowledge Representation

Logics: modal

1. take a logical language, let α be a wff
2. $\Box\alpha$ is a wff
3. $\Diamond\alpha$ is a wff
4. normally $\Box\alpha \leftrightarrow \neg\Diamond\neg\alpha$

Intended meaning?

1. $\Box\alpha$ means **Necessarily** α
2. $\Box\alpha$ means **Agent knows** α

Knowledge Representation

Logics: modal

1. take a logical language, let α be a wff
2. $\Box\alpha$ is a wff
3. $\Diamond\alpha$ is a wff
4. normally $\Box\alpha \leftrightarrow \neg\Diamond\neg\alpha$

Intended meaning?

1. $\Box\alpha$ means **Necessarily** α
2. $\Box\alpha$ means **Agent knows** α
3. $\Box\alpha$ means **Agent believes** α

Logics: modal

1. take a logical language, let α be a wff
2. $\Box\alpha$ is a wff
3. $\Diamond\alpha$ is a wff
4. normally $\Box\alpha \leftrightarrow \neg\Diamond\neg\alpha$

Intended meaning?

1. $\Box\alpha$ means **Necessarily** α
2. $\Box\alpha$ means **Agent knows** α
3. $\Box\alpha$ means **Agent believes** α
4. $\Box\alpha$ means **Always in the future** α

Knowledge Representation

Logics: modal

1. take a logical language, let α be a wff
2. $\Box\alpha$ is a wff
3. $\Diamond\alpha$ is a wff
4. normally $\Box\alpha \leftrightarrow \neg\Diamond\neg\alpha$

Intended meaning?

1. $\Box\alpha$ means **Necessarily** α
2. $\Box\alpha$ means **Agent knows** α
3. $\Box\alpha$ means **Agent believes** α
4. $\Box\alpha$ means **Always in the future** α
5. $G\alpha$ means **Always in the future (or: Globally)** α

Jacek Malec, Computer Science, Lund University 16(39)

Logics: Kripke semantics

Actually, meaning of modal formulae is defined on graph structures

Nodes: possible worlds

Edges: reachability relation

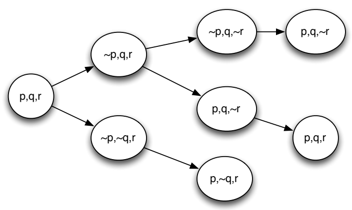

Logics: temporal

1. Globally (always):

   $\Box\Phi$

2. Finally (eventually):

   $\Diamond\Phi$

3. Next:

   $\bigcirc\Phi$

4. Until:

   $\Psi U \Phi$

Logics: temporal

1. Globally (always):
$$\Box\Phi$$
2. Finally (eventually):
$$\Diamond\Phi$$
3. Next:
$$\bigcirc\Phi$$
4. Until:
$$\Psi U \Phi$$

Cf. Richard Murray's verification of autonomous car controller:

$$(\Phi^e_{init} \wedge \Box\Phi^e_{safe} \wedge \Box\Diamond\Phi^e_{prog}) \rightarrow (\Phi^s_{init} \wedge \Box\Phi^s_{safe} \wedge \Box\Diamond\Phi^s_{prog})$$

Logics: description

Earlier known as semantic networks. Formal version of semantic web languages (OIL, DAML, OWL).

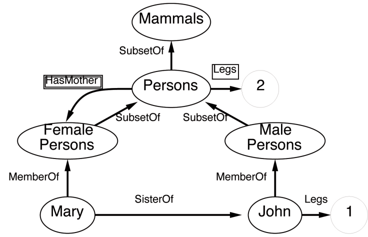

Effective reasoning:

- inheritance via SubsetOf (SubClass) and MemberOf (isA) links
- intersection paths
- special meaning of some links (e.g. cardinality constraints)
- classification, consistency, subsumption

Representation: ontologies

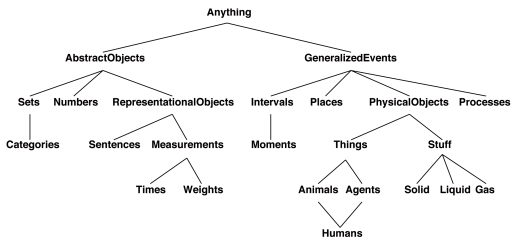

Lots of robot-related ontologies:
knowrob, IEEE CORA (Standard 1872-2015), intelligent systems ontology (2005, NIST), ...

Modes of reasoning: Deduction

$$RedLightAt(intersection1)$$
$$\forall(x) RedLightAt(x) \rightarrow \bigcirc StopBefore(x)$$

thus

$$\bigcirc StopBefore(intersection1)$$

General Pattern:

1. prior facts
2. domain knowledge
3. observations

Modes of reasoning: Deduction

$$RedLightAt(intersection1)$$
$$\forall(x) RedLightAt(x) \rightarrow \bigcirc StopBefore(x)$$

thus

$$\bigcirc StopBefore(intersection1)$$

General Pattern:

1. prior facts
2. domain knowledge
3. observations
4. conclusions

Sound.

Modes of reasoning: Deduction

$$RedLightAt(intersection1)$$
$$\forall(x) RedLightAt(x) \rightarrow \bigcirc StopBefore(x)$$

thus

$$\bigcirc StopBefore(intersection1)$$

General Pattern:

1. prior facts
2. domain knowledge
3. observations
4. conclusions

Sound. But note:
Birds fly. Tweety is a penguin. Penguins are birds.

Modes of reasoning: Induction

$OnDesk(monitor1) \wedge Monitor(monitor1)$,
$OnDesk(monitor2) \wedge Monitor(monitor2)$,
$OnDesk(monitor3) \wedge Monitor(monitor3)$,
$OnDesk(monitor4) \wedge Monitor(monitor4)$,
$OnDesk(monitor5) \wedge Monitor(monitor5)$
thus
$\forall(x) Monitor(x) \rightarrow OnDesk(x)$

General pattern:

1. Observe
2. Generalize

Fallible. Constructs hypotheses, not true facts. However, most of our practical reasoning, in particular learning, is of this kind.

Modes of reasoning: Abduction

General pattern:

1. prior facts
2. domain knowledge
3. observations

Modes of reasoning: Abduction

General pattern:

1. prior facts
2. domain knowledge
3. observations
4. explain the observation

Given a theory T and observations O
E is an explanation of O given T if
$E \cup T \models O$ and $E \cup T$ is consistent.
Usually we are interested in most plausible E, sometimes minimal E, most elegant E, ...

Probablilistic abduction: maybe Elin will mention it.

What do we want to represent?

- objects
- places
- times
- events
- processes
- behaviours
- vehicle dynamics
- rigid body interactions
- traffic laws
- ...

Qualitative spatial reasoning

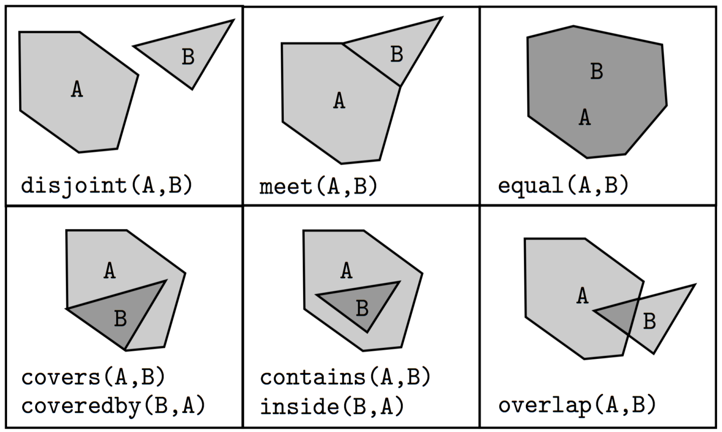

Qualitative spatial reasoning

	disjoint	meet	equal	inside	coveredby	contains	covers	overlap
disjoint	RCC8	disjoint meet inside coveredby overlap	disjoint	disjoint meet inside coveredby overlap	disjoint meet inside coveredby overlap	disjoint	disjoint	disjoint meet inside coveredby overlap
meet	disjoint meet contains covers overlap	disjoint meet equal coveredby covers overlap	meet	inside coveredby overlap	meet inside	disjoint	disjoint meet	disjoint meet inside coveredby overlap
equal	disjoint	meet	equal	inside	coveredby	contains	covers	overlap
inside	disjoint	disjoint	inside	inside	inside	RCC8	disjoint meet inside coveredby overlap	disjoint meet inside coveredby overlap
coveredby	disjoint	disjoint meet	coveredby	inside	inside coveredby	disjoint meet contains covers overlap	disjoint meet equal coveredby covers overlap	disjoint meet overlap coveredby overlap
contains	disjoint meet contains covers overlap	contains covers overlap	contains	equal inside coveredby contains covers overlap	contains covers overlap	contains	contains	contains covers overlap
covers	disjoint meet contains covers overlap	meet contains covers overlap	covers	inside coveredby overlap	equal coveredby covers overlap	contains	contains covers	contains covers overlap
overlap	disjoint meet contains covers overlap	disjoint meet contains covers overlap	overlap	inside coveredby overlap	inside coveredby overlap	disjoint meet contains covers overlap	disjoint meet contains covers overlap	RCC8

Qualitative spatial reasoning

RCC8: **region connection calculus**
Given e.g.,
contains(A, B) $\land$ *covers*(B, C) we can conclude *contains*(A, C)

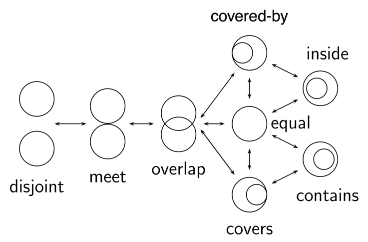

$$\Box(meet(A,B) \rightarrow \bigcirc(meet(A,B) \lor disjoint(A,B) \lor overlap(A,B)))$$

Juggling example (Apt)

- From some time on, at most one ball is not in the air:

$$\Diamond\Box(\forall b \in Balls.\ \forall h \in Hands.\ Q[b,h] = \text{meet} \rightarrow \forall b_2 \in Balls.\ b \neq b_2 \rightarrow \forall h_2 \in Hands.\ Q[b_2,h_2] = \text{disjoint}).$$

- A ball thrown from one hand remains in the air until it lands in the other hand:

$$\Box(\forall b \in Balls.\ \forall h_1, h_2 \in Hands.\ h_1 \neq h_2 \land Q[h_1,b] = \text{meet} \rightarrow Q[h_1,b] = \text{meet}\ \mathsf{U}\ (Q[h_1,b] = \text{disjoint} \land Q[h_2,b] = \text{disjoint} \land (Q[h_1,b] = \text{disjoint}\ \mathsf{U}\ Q[h_2,b] = \text{meet}))).$$

Interval calculus (Allen 1983)

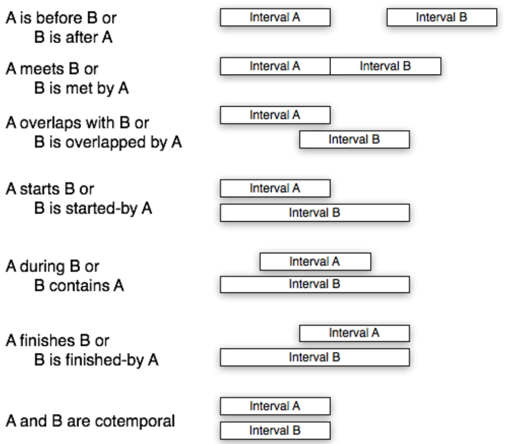

Invalidating conclusions

- Tweety is a bird.
- So it flies.

Invalidating conclusions

- Tweety is a bird.
- So it flies.
- But Tweety is a penguin.
- So it doesn't fly.

Invalidating conclusions

- Tweety is a bird.
- So it flies.
- But Tweety is a penguin.
- So it doesn't fly.

Non-monotonic reasoning.

Truth-maintenance systems.

Default reasoning. Circumscription. Closed World Assumption. Negation as failure. . . .

Uncertainty

Every perception is associated with uncertainty. Account for that. (Yesterday lectures. Perception module.)

Approaches:

- probabilistic representations
- fuzzy approaches
- multi-valued logics

Transformations between representations as needed.

Back to KnowRob

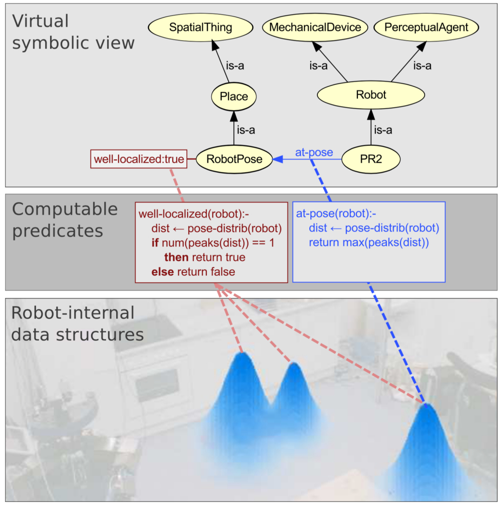

KnowRob lessons

Beetz and Tenorth, AIJ, 2016:

1. No fixed levels of abstraction, no layers, no “black boxes”;
2. A knowledge base should reuse data structures of the robot’s control program;
3. Symbolic knowledge bases are useful, but not sufficient;
4. Robots need multiple inference methods;
5. Evaluating a robot knowledge base is difficult.

Architectures of knowledge-based systems

AIMA agents (cf. introductory lecture)

1. Logical agents - declarative, compositional
2. Rule-based systems - compositionality on the rule level
3. Layered systems (distribution of concerns)
4. Blackboards - compositionality of reasoners (knowledge sources) (KnowRob, our SIARAS system)
5. Stream-oriented reasoning - Heintz@LiU

KnowRob as a blackboard

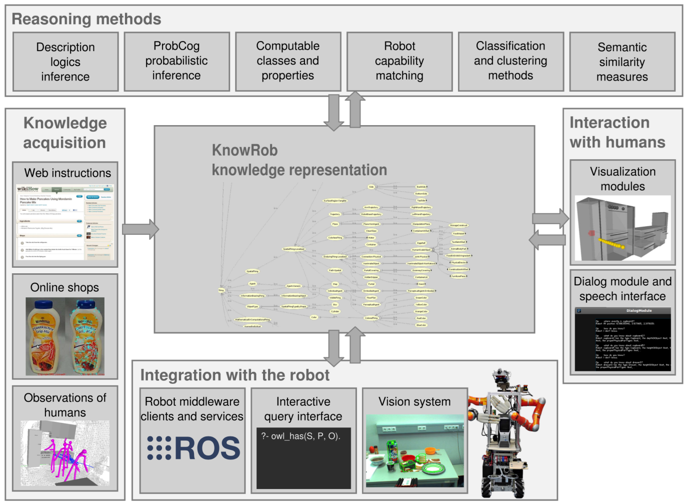

Self-awareness: Autoepistemic logic

1. Distribution axiom **K**:

$$(K\alpha \wedge K(\alpha \rightarrow \beta)) \rightarrow K\beta$$

2. Knowledge axiom **T**:

$$K\alpha \rightarrow \alpha$$

3. Positive introspection **4:**

$$K\alpha \rightarrow KK\alpha$$

4. Negative introspection **5:**

$$\neg K\alpha \rightarrow K\neg K\alpha$$

Self-awareness: motivation

1. true autonomy requires self-awareness
2. autoepistemic logic captures just one aspect: awareness of own knowledge
3. resource limitations: anytime algorithms, active logic
4. interaction: distributed knowledge
5. interaction: shared knowledge
6. explanation of own behaviour (trust)

References 1

https://www.youtube.com/watch?v=ymUFadN_MO4 (How Watson learns)
DOI: 10.1147/JRD.2012.2186519, Automatic knowledge extraction from documents, J. Fan, A. Kalyanpur, D. C. Gondek, D. A. Ferrucci, IBM J. RES. DEV. VOL. 56 NO. 3/4 PAPER 5, 2012
YAGO2: A Spatially and Temporally Enhanced Knowledge Base from Wikipedia, Johannes Hoffart, Fabian M. Suchanek, Klaus Berberich, Gerhard Weikum, Artificial Intelligence Journal, vol. 194, pp. 28-61, 2013
Representations for robot knowledge in the KnowRob framework, Moritz Tenorth, Michael Beetz, Artificial Intelligence Journal, in press, available on the journal site
Logics for Artificial Intelligence, Raymond Turner, Ellis Horwood, 1984

References 2

Logic In Action, Johan van Benthem, http://www.logicinaction.org, 2012
Rete: A Fast Algorithm for the Many Pattern/ Many Object Pattern Match Problem, Charles L. Forgy, Artificial Intelligence Journal, vol.19 (1982), pp. 17-37.
https://arxiv.org/pdf/1201.4089.pdf, A Description Logic Primer, Markus Kroetzsch, Frantisek Simancik, Ian Horrocks
Qualitative Spatial Representation and Reasoning, Anthony G Cohn and Jochen Renz, Handbook of Knowledge Representation, pp. 551-596, Elsevier, 2008