AI Knowledge representation

- Core distinction between:
  - **Declarative** representation (**WHAT**)
    - All birds can fly
    - Ducks are birds
    - Penguins are birds, but cannot fly
    - Etc
  - **Inference** (**HOW** to compute)
    - Pengi is a penguin
    - Pengi is a bird
    - Can Pengi fly?

KR general intelligence

Structured object models
wheel(w1),wheel(w2),
frame(1),connected(w1,f)...

Occlusion relations
partially_behind..

Structured models of complete scenes
lawnmower(l),dustbin(b),sp
atial_relation(l,b)

Commonsense, dynamic prediction models
What will probably happen if I move this table?

Propositional vs. Relational

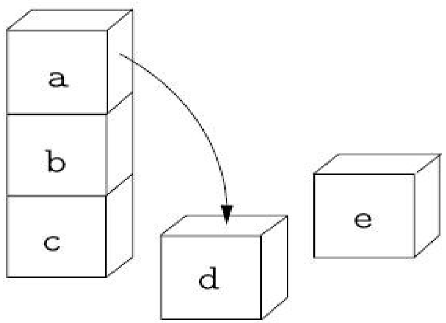

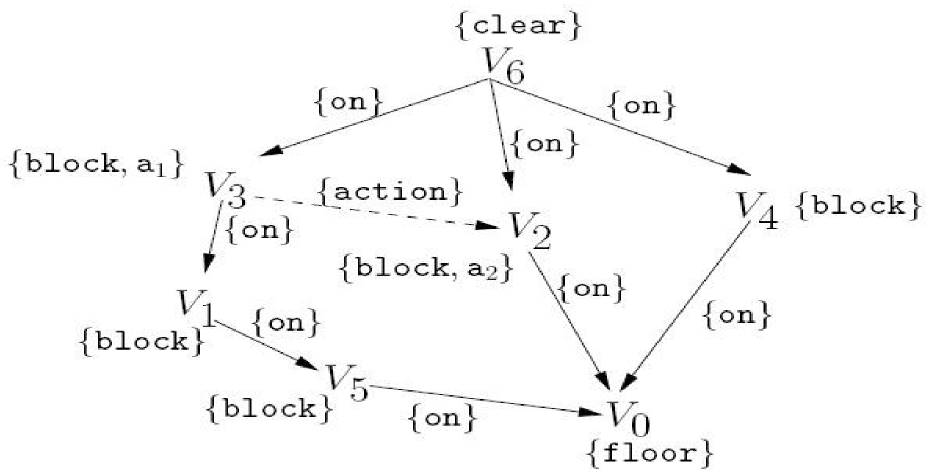

Finite worlds ↔ probabilistic graphical model, but

– Random variable for all ground relations, e.g. on_a_b=1

– Need to order all objects and relations

– No generalization over objects

Number of model parameters grows exponentially fast (in #objects and #relations) even if there is considerable structure

Again: Learning and representation (AIPSML)

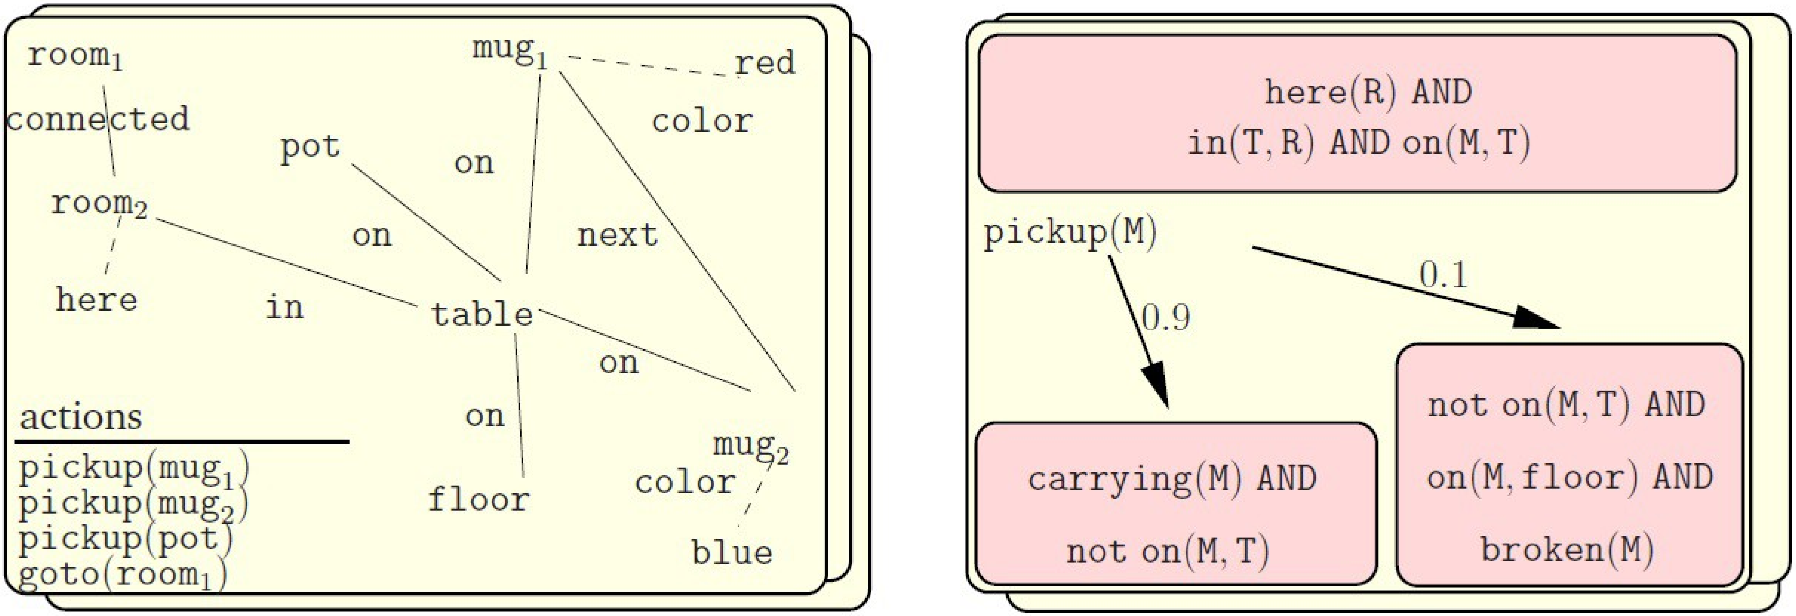

Figure 1.6: ***Data structures for*** **FREGE.** ***a)*** *A state and a set of applicable actions.* ***b)*** *Part of the transition model (T).*

States are described by objects and relations; powerful representations and learning algorithms need to support that (~Prolog-like)

An Logical
Machine Learning Classic

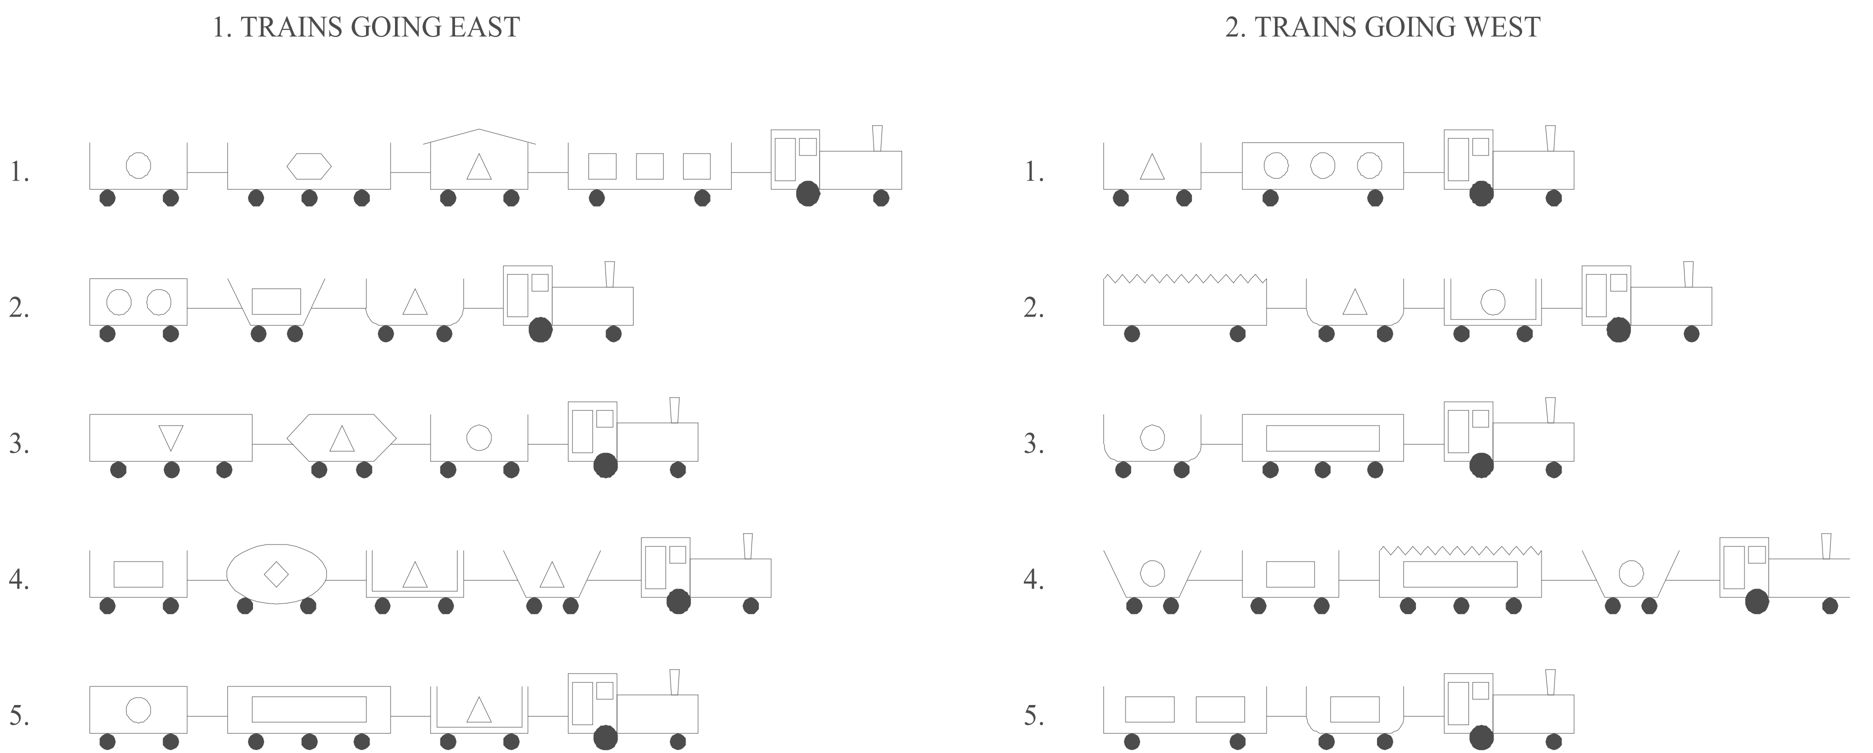

An Logical
(Machine Learning) Classic

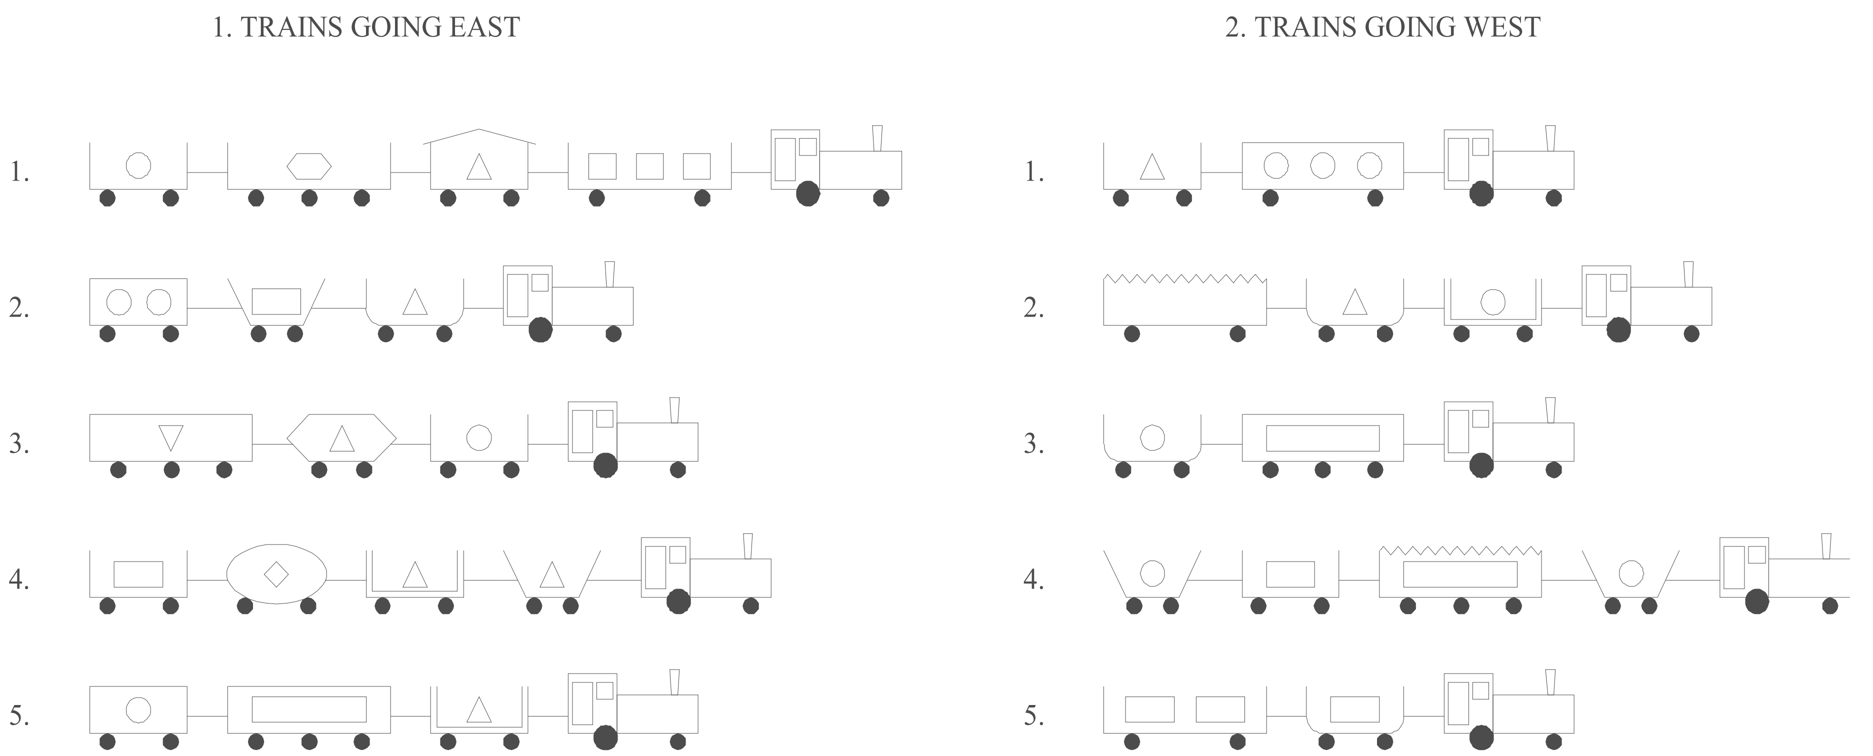

Relational Representation

Example:

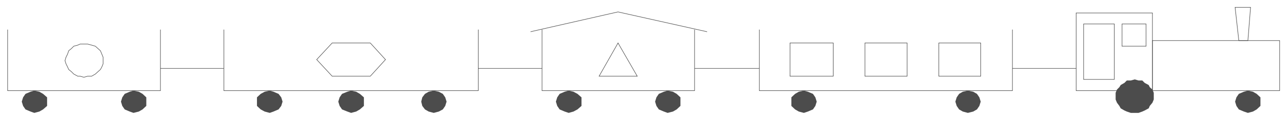

`eastbound(t1).`

Background theory:

```
car(t1,c1).
rectangle(c1).
short(c1).
none(c1).
two_wheels(c1).
load(c1,l1).
circle(l1).
one_load(l1).
three_loads(l4).

car(t1,c2).
rectangle(c2).
long(c2).
none(c2).
three_wheels(c2).
load(c2,l2).
hexagon(l2).
one_load(l2).

car(t1,c3).
rectangle(c3).
short(c3).
peaked(c3).
two_wheels(c3).
load(c3,l3).
triangle(l3).
one_load(l3).

car(t1,c4).
rectangle(c4).
long(c4).
none(c4).
two_wheels(c4).
load(c4,l4).
rectangle(l4).
```

Hypothesis:

`eastbound(T):-car(T,C),short(C),not none(C).`

AI started using relational KR

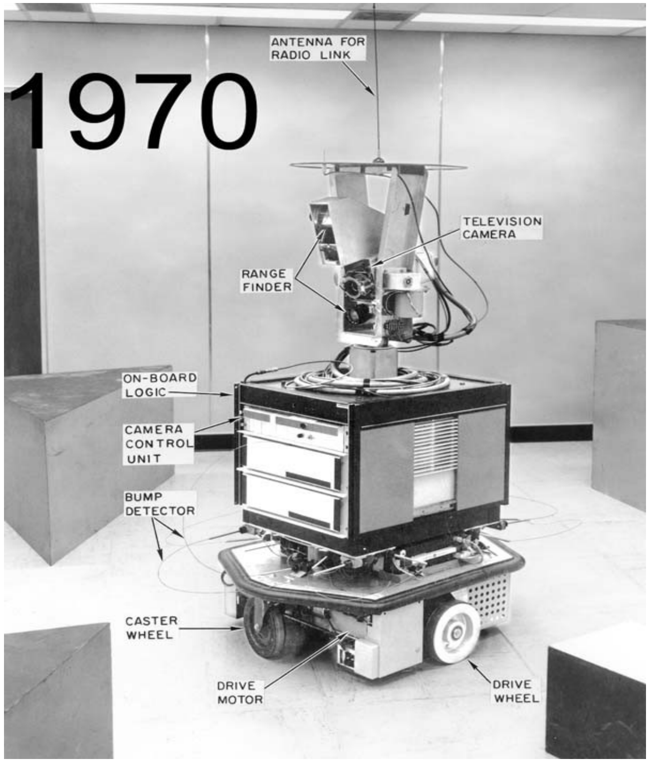

Shakey and
STRIPS planning

John McCarthy
Situation Calculus

Logic dominates in linguistics and spatial/temporal data

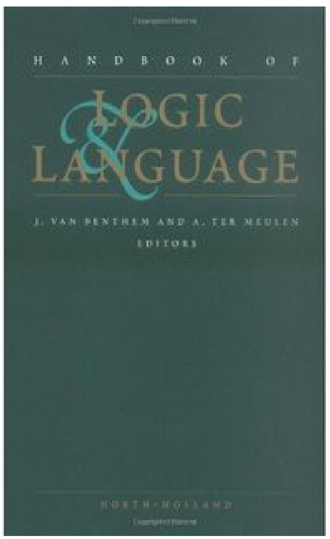

Semantics, discourses, grammars

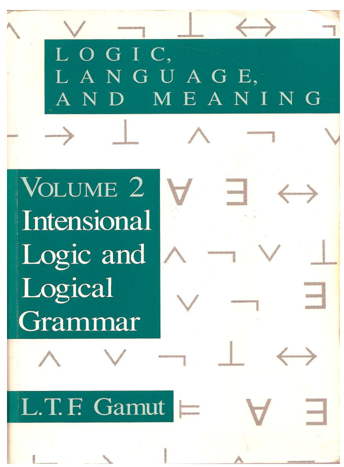

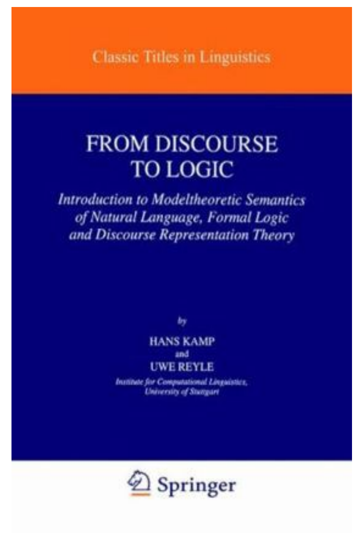

Spatial and Temporal formalisms

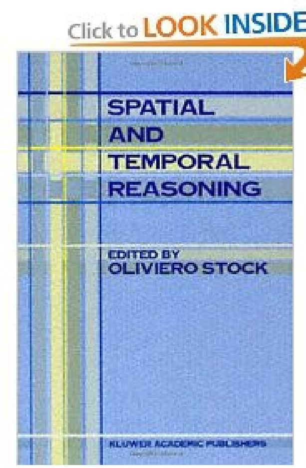

KR and AI

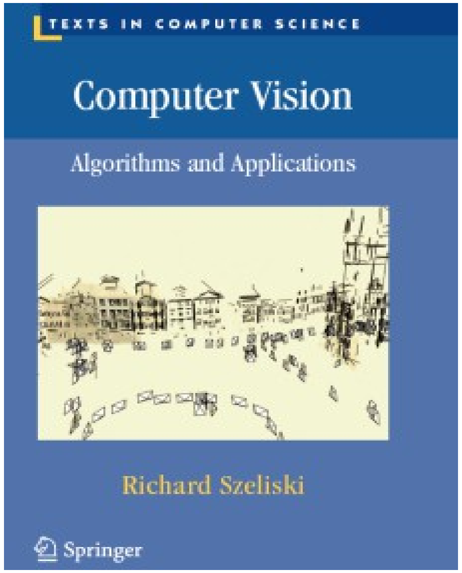

2010

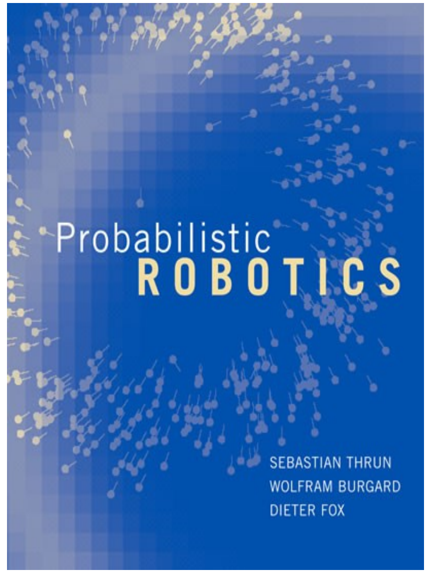

2005

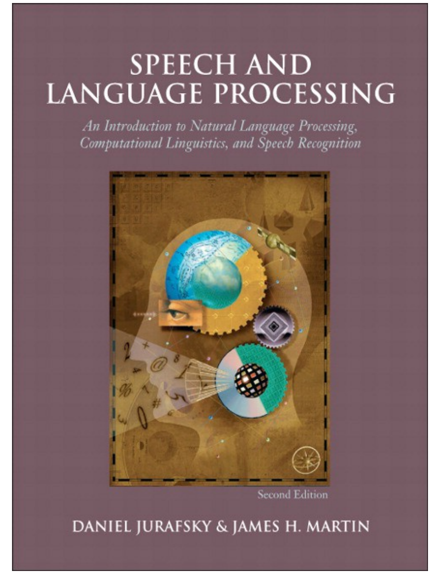

2010

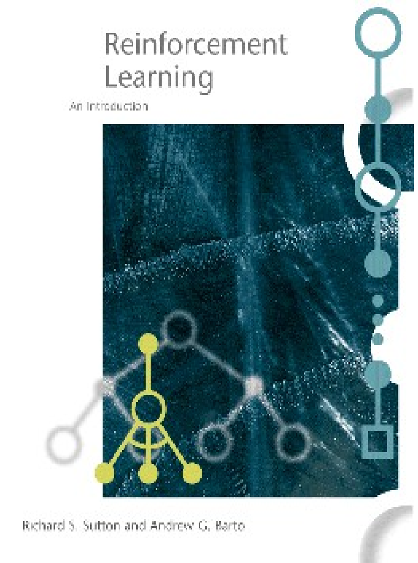

1998

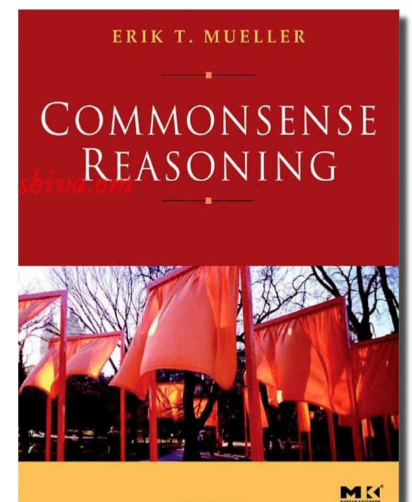

2006

State-of-the-art (learning) approaches:

Based on **propositional**, probabilistic graphical models

In: Computer vision, (cognitive) Robotics, Linguistics, and Reinforcement Learning

An AI classic: Bioinformatics, molecules, and proteins

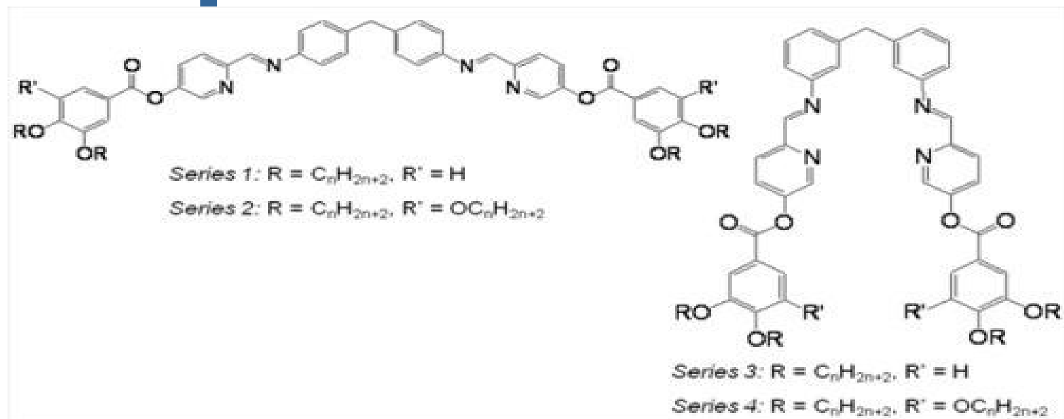

Examples **E**

pos(mutagenic(m_1))
neg(mutagenic(m_2))
pos(mutagenic(m_3))
...

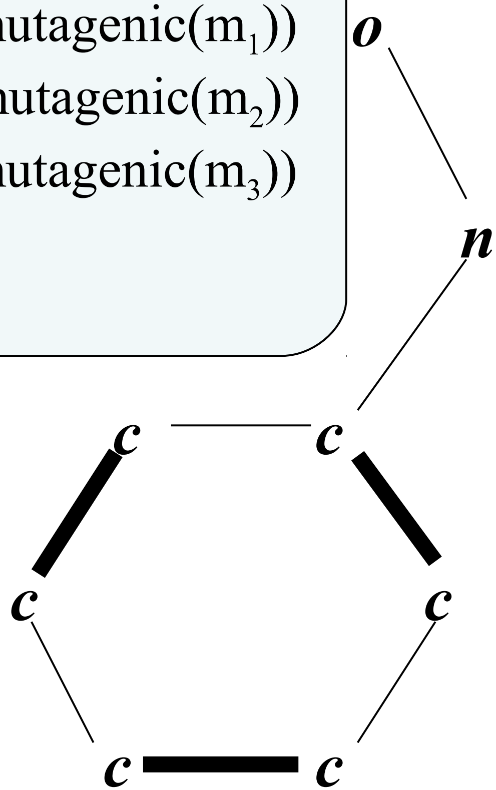

Background Knowledge **B**

molecule(m_1)	molecule(m_2)
atom(m_1,a_{11},c)	atom(m_2,a_{21},o)
atom(m_1,a_{12},n)	atom(m_2,a_{22},n)
bond(m_1,a_{11},a_{12})	bond(m_2,a_{21},a_{22})
charge(m_1,a_{11},0.82)	charge(m_2,a_{21},0.82)
...	...

Relational representations (graphs: objects and relations)

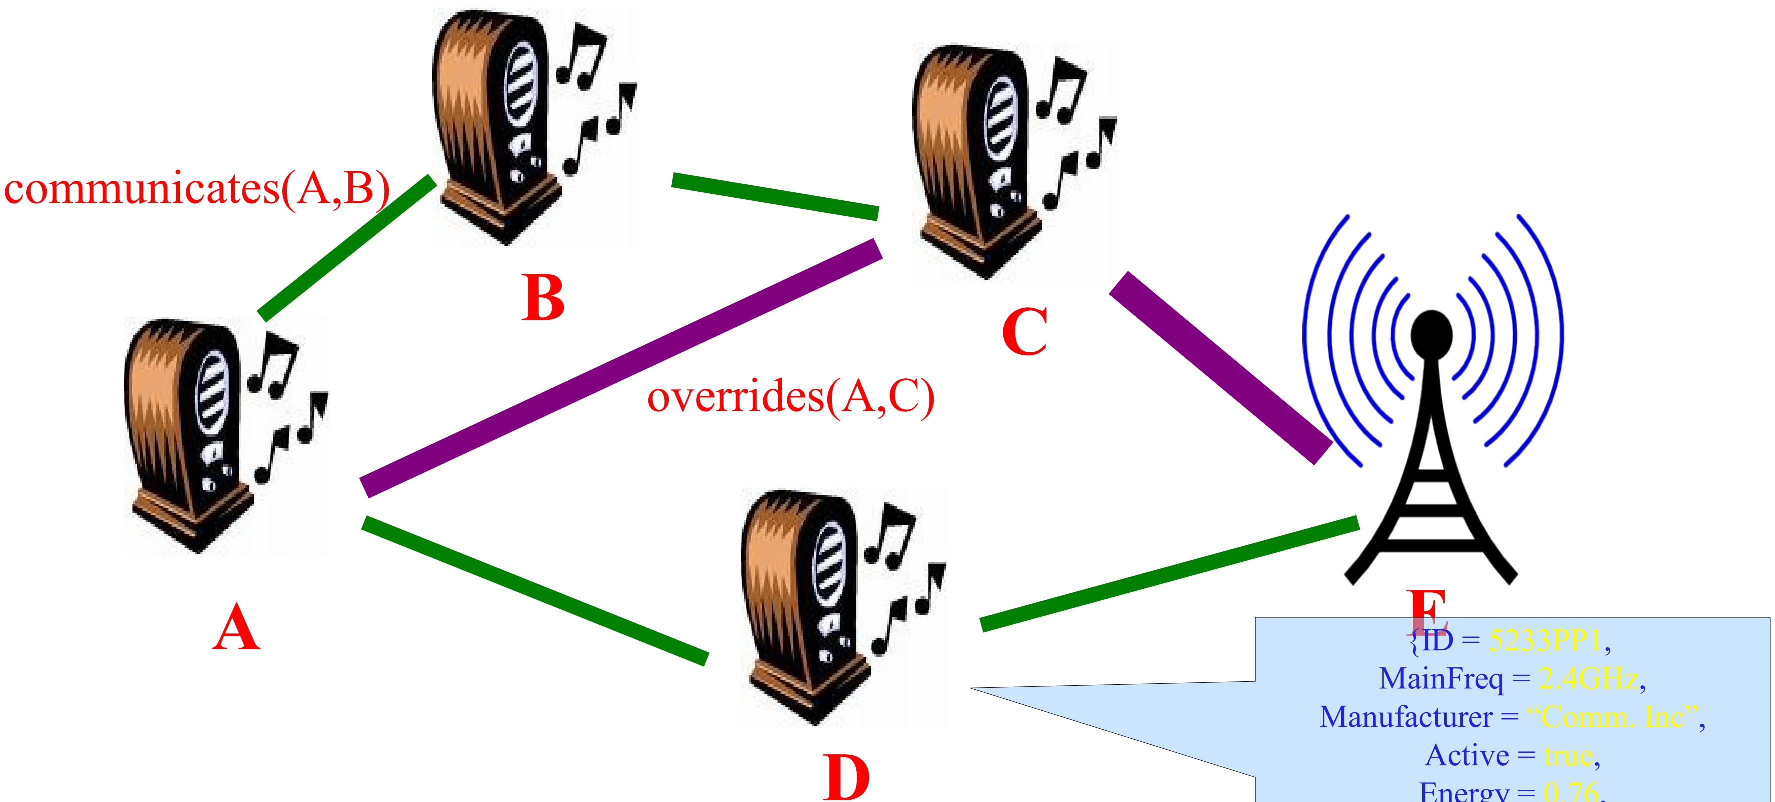

Tasks → subsumes attribute-value

Tools → rules, decision trees, etc

Real Time Strategy Games (RTS)

This subsumes all kinds of logistics domains!

StarCraft 2 (Blizzard)

And yes... Blocks World!

Uncertainty and Learning

- Core of this course:
  - Representation
  - Inference (deduction, abduction)
- Later
  - Some uncertainty
  - Some utility
  - No learning
- Current AI:
  - Statistical Relational Learning (SRL)
  - Probabilistic Logic Learning (PLL)
  - KR+learning+probability+utility+AI+...

Uncertainty?

- Real world domains
- Vision, Robotics, Linguistics
- Noisy sensor data
- Most things are uncertain!

- Use logic for knowledge representation
- Add uncertainty

- Often upgrades of probabilistic graphical models (e.g. Bayesian networks)

Upgrading: The Alphabet of SRL

Probabilistic logics (Nilsson, Halpern, Bacchus)
Knowledge-Based Model Construction (**KBMC**)
Stochastic Logic Programs (**SLP**)
Logic programs with annotated disjunctions (**LPAD**)
Causal Process logic (**CP-Logic**)
Probabilistic Relational Models (**PRM**)
Statistical Relational Models (**SRM**)
Bayesian Logic Networks (**BLN**)
Bayesian Logic Programs (**BLP**)
Relational Markov Models (**RMM**)
Markov Logic Networks (**MLN**)
Relational Decision Networks (**RDN**)
Relational Dependency Networks (**RDN**)
Bayesian Logic (**BLOG**)
...

```
1:s(A,B) :- n(A,C), v(C,D), n(D,B).
0.4:n([joe|T],T).  0.6:n([kim|T],T).
0.3:v([sees|T],T). 0.7:v([likes|T],T).
```

Many SRL systems Started as an attempt to **upgrade** a particular propositional learning /probabilistic model. (e.g. BN, CRF, MRF, ME, DBN, HMM, ..)

SCFG/SLP

Computational Logic + AI

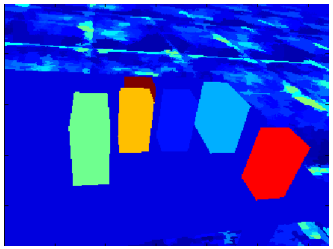

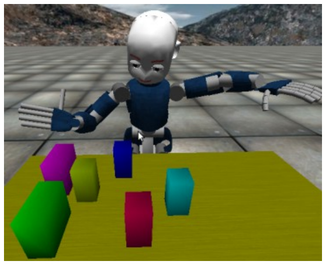

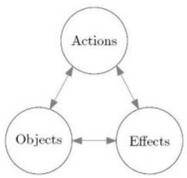

Inputs	Outputs	Function
(O, A)	E	Effect prediction
(O, E)	A	Action recognition/planning
(A, E)	O	Object recognition/selection

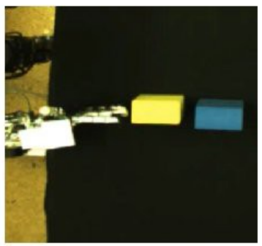

Now @ICRA'12

Computational, probabilistic logic + learning → high-level/cognitive **vision**, high-level **robotics**, Spatial/natural **language** grounding

Additional
Decision-theoretic planning, abductive logic
DT-Problog

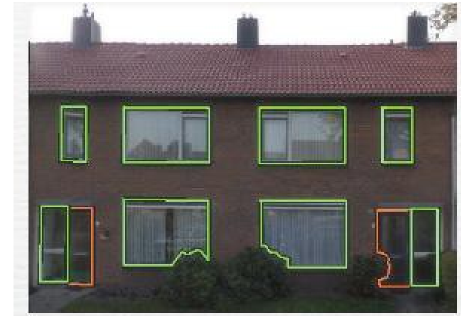

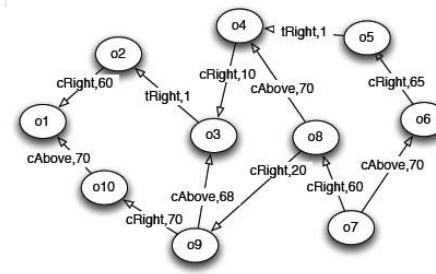

Best paper award @ICPRAM'12

In journal TCLP@ACM

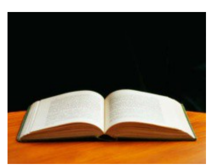

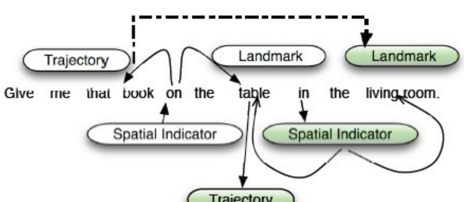

ON(book,table)
IN(table, living room)
IN(book,living room)?

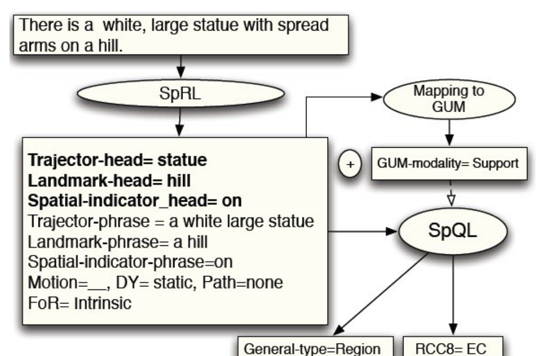

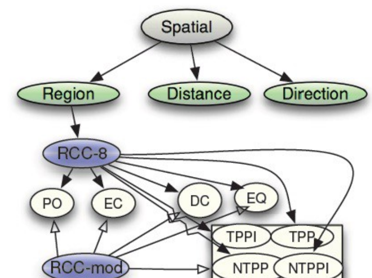

People Tracking

person(p1). ...
person(p5).

group(p4,p5).

KR for AI cognitive robotics

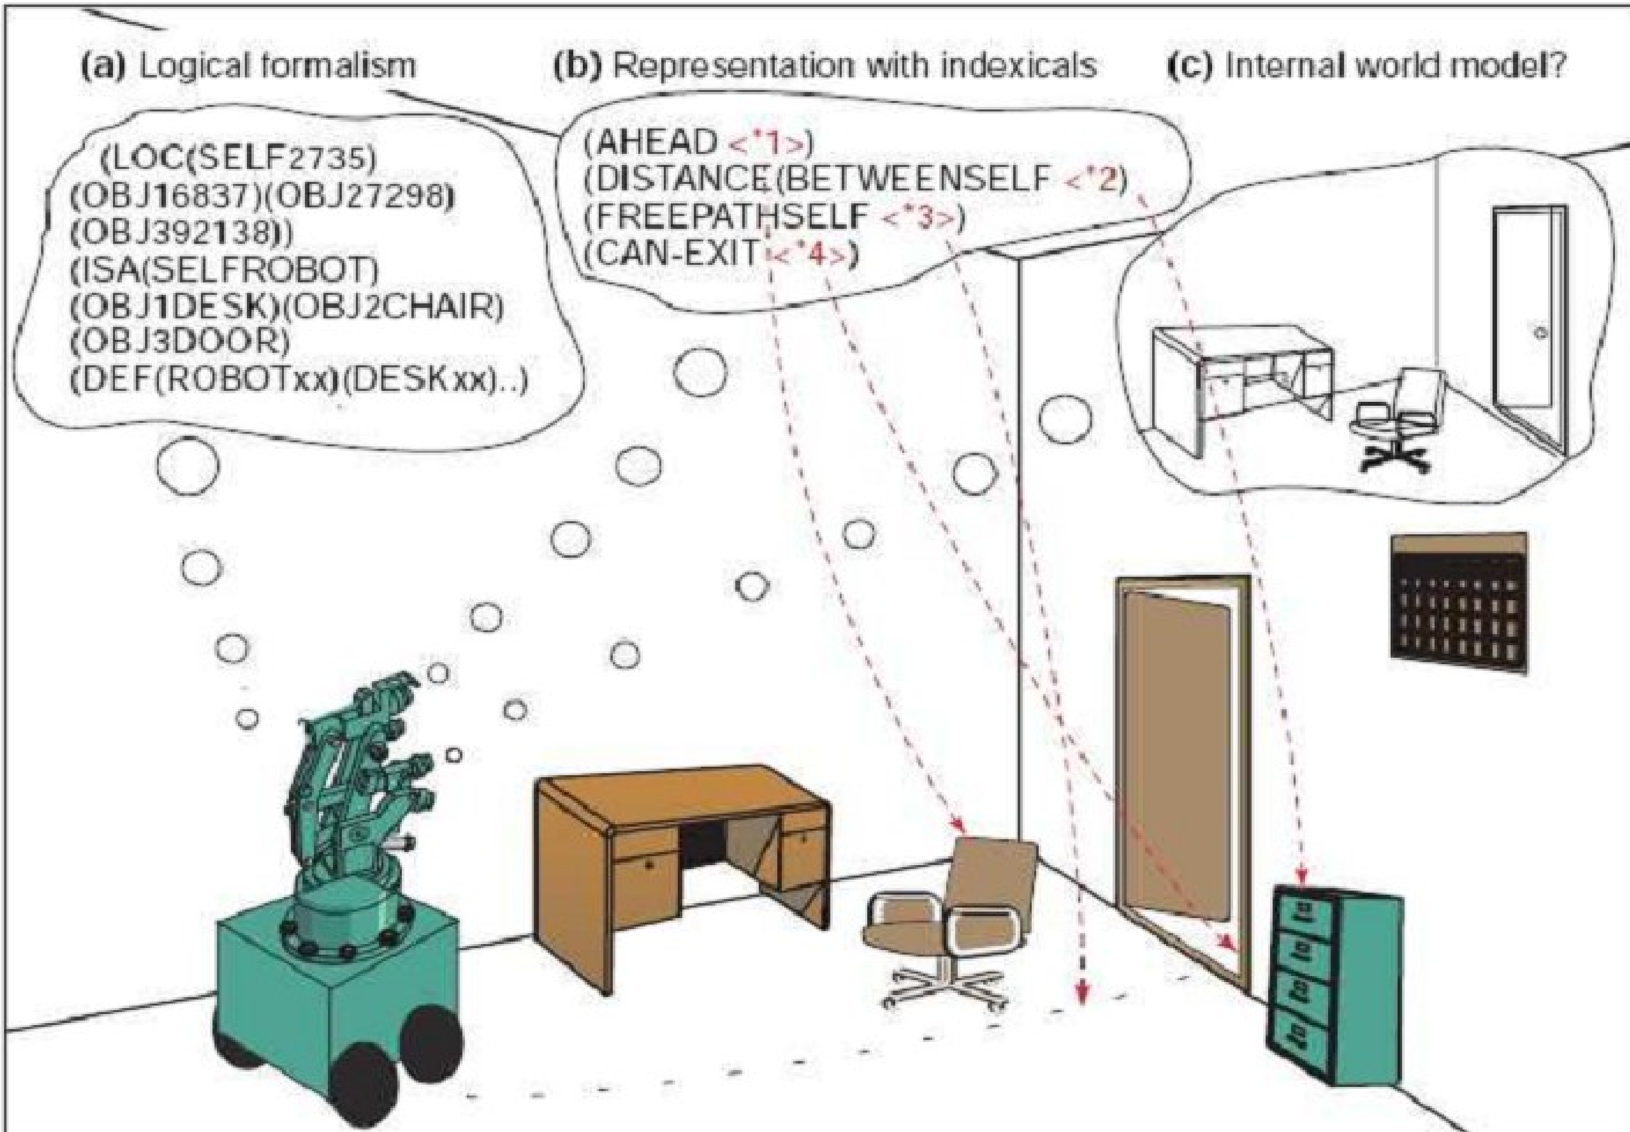

Relational/Cognitive Robotics

SEVENTH FRAMEWORK PROGRAMME

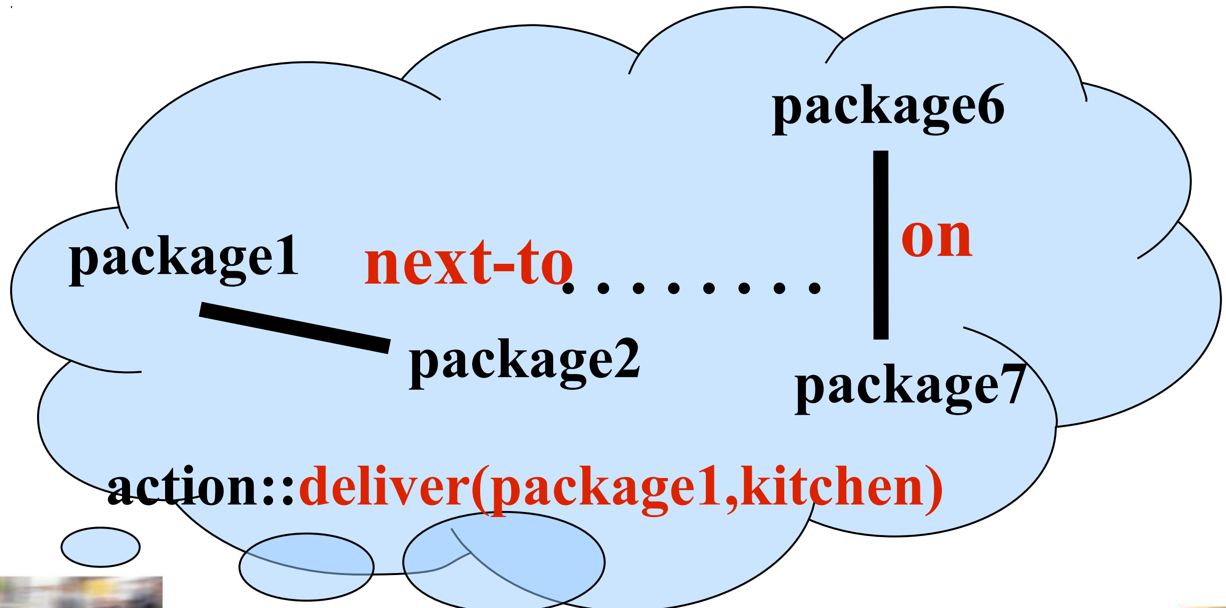

FIRST-MM

@KULeuven

KR for modern robotics

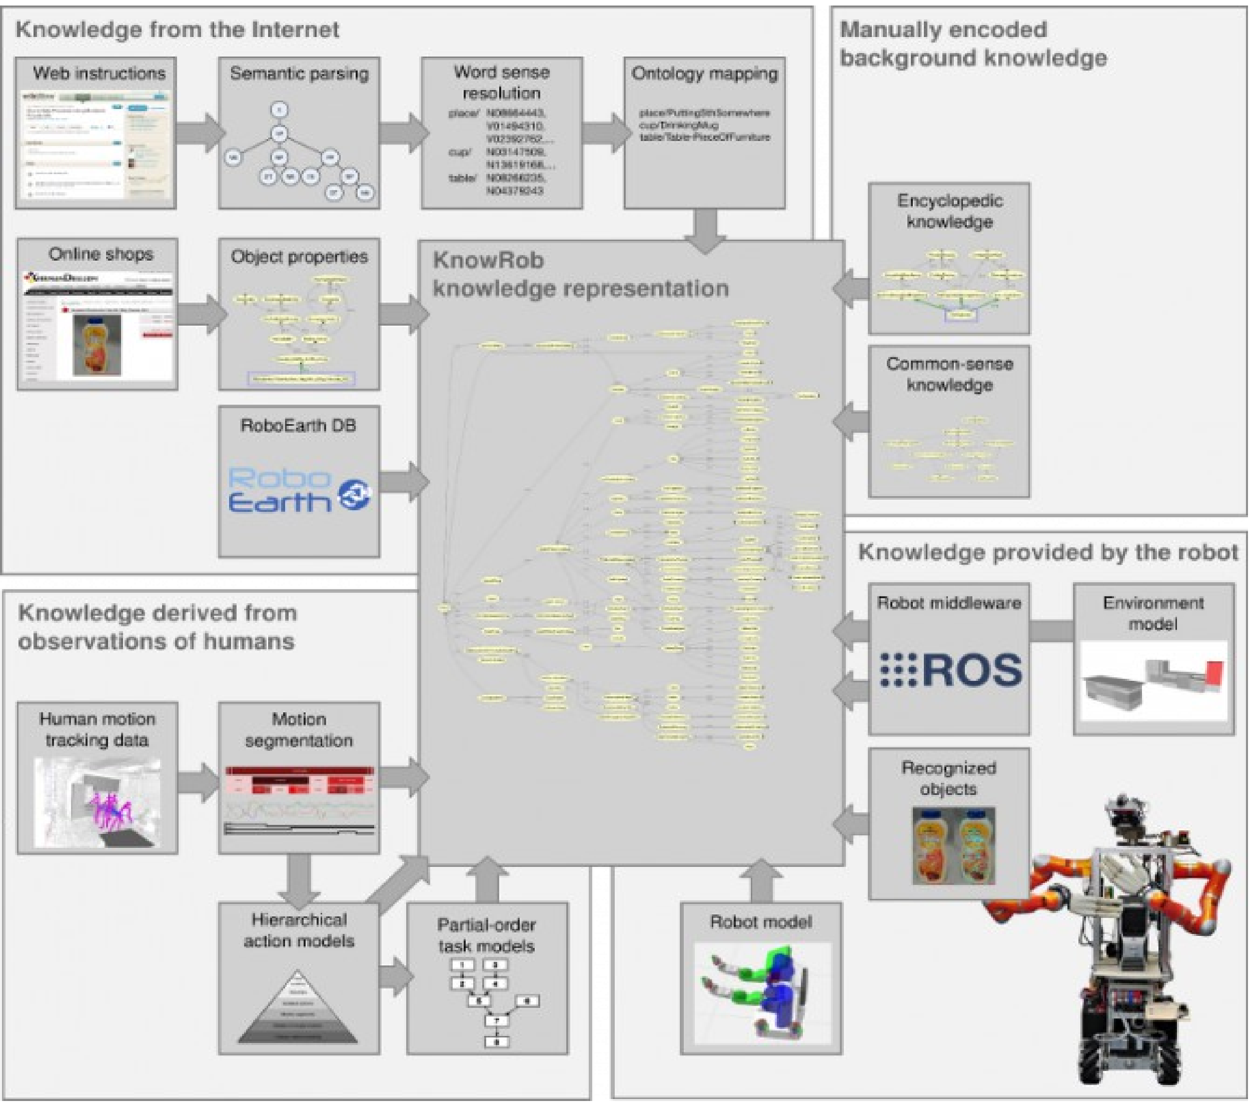

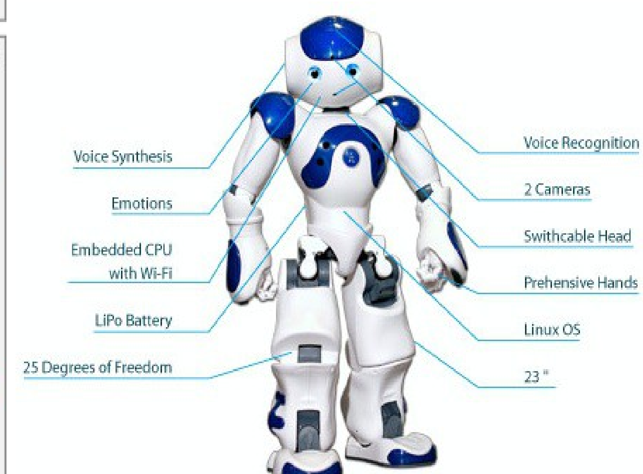

Relations in High-Level Vision

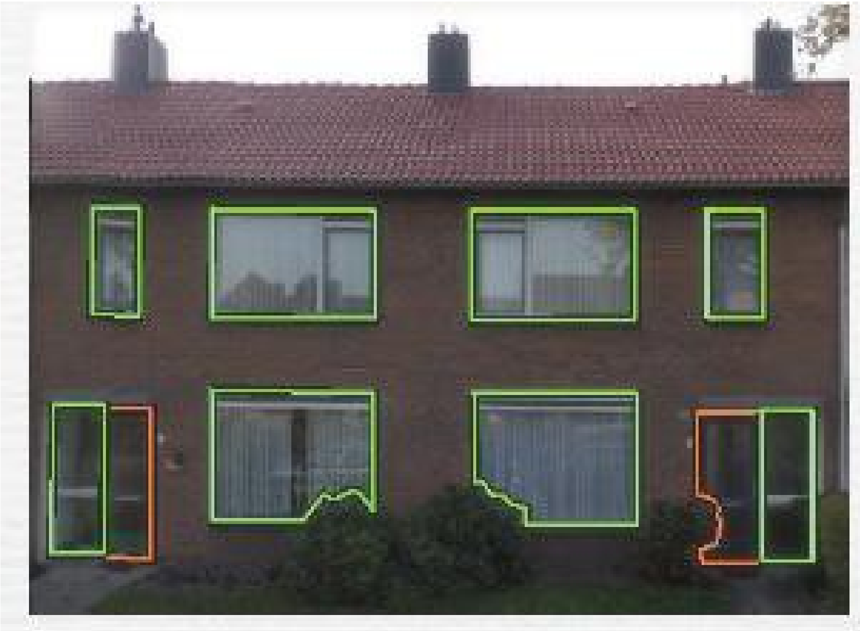

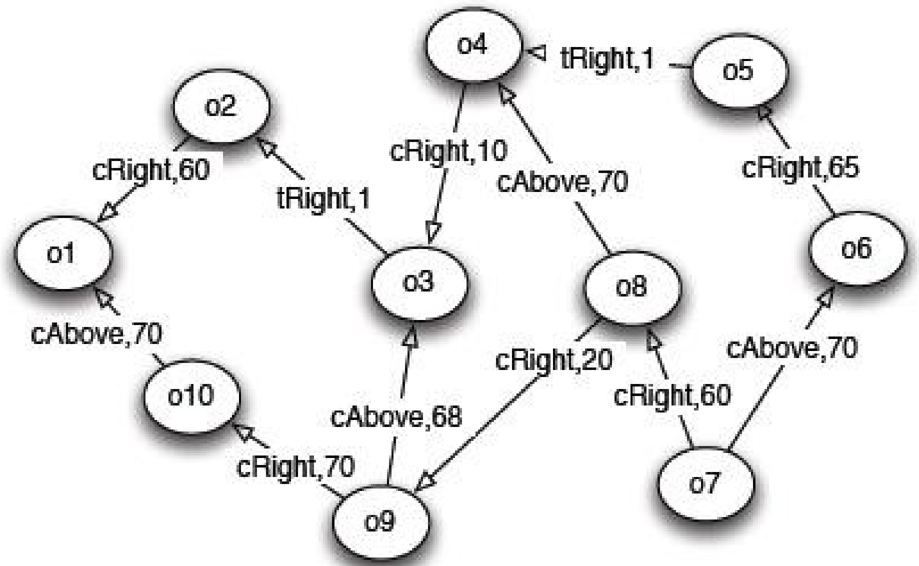

Bottom-up, hierarchical approach to understanding and semantically segmenting images of houses

(L. Antanas, M. van Otterlo, J. Oramas, T. Tuytelaars, L. De Raedt, ILP-2010)

Spatial (language)

Spatial surroundings of the robot:
near(arm, table),
distance(robot,door,10.3),
P(succ|forward)=0.9, ...

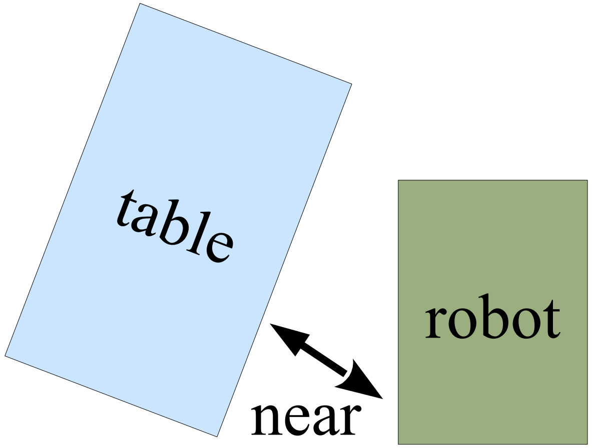

Instructions:
"Move to the first table on your left, and pick the object nearest to the edge of it"

Involves high-level vision, planning, manipulation, etc.

In addition to relations: Uncertainty

- Many domains (inc. linguistics) exhibit
  - Structure in terms of objects and relations
  - Highly statistical and ambiguous problems

I saw **the man with a telescope**
I saw **the man with a telescope**

The book is on the table
The book lies on the table
The book is supported by the table
The book is about a table
The book has a table of contents
The book contains a table of contents

Hier **in Otterlo** kun je mooie dingen zien.
Martijn **van Otterlo** woont in Leuven..... maar, **Otterlo** heeft een nieuw stadhuis, maar nu is Leuven jaloers. Mr. Mark **van Otterlo** zei dat dat gerechtvaardigd was.

Jeopardy and Watson

Dynamics: CPTL
(Thon,Landwehr,De Raedt 2010)

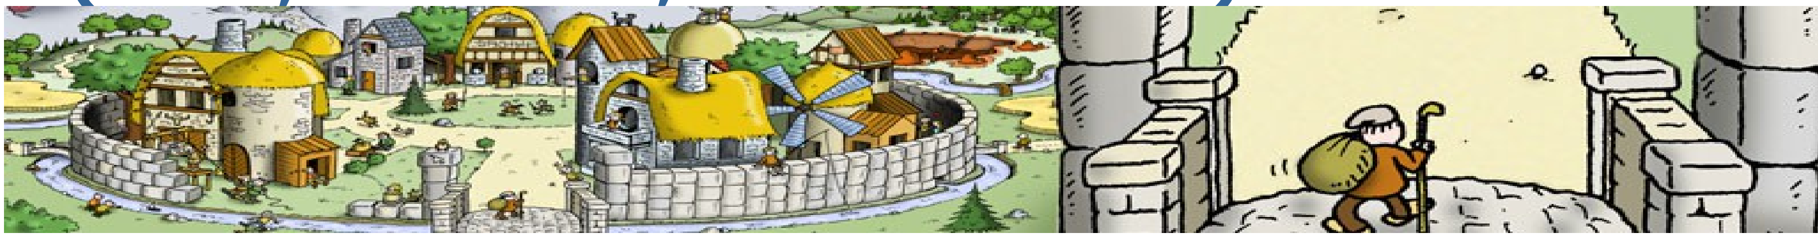

$$\underbrace{b_1, \ldots b_n}_{\text{cause (past)}} \rightarrow \underbrace{h_1 : p_1 \vee \ldots \vee h_m : p_m}_{\text{effect (future)}}$$

conquer a city which is close:

$$city(C, Owner), city(C2, Attacker), close(C, C2) \rightarrow$$
$$conquest(Attacker, C) \ : \ p \ \vee \ nil \ : \ (1-p)$$

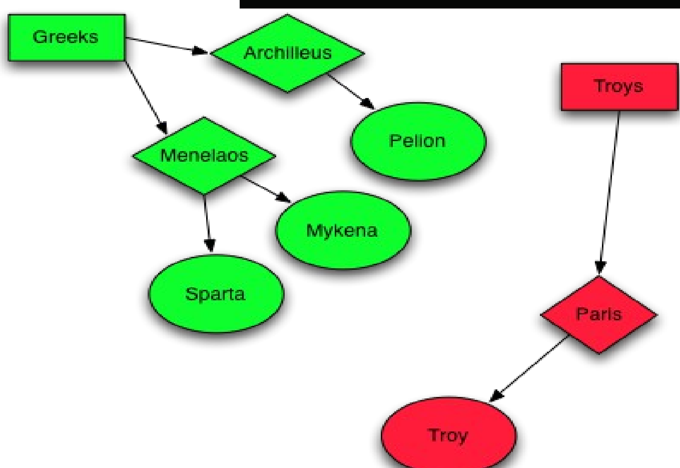

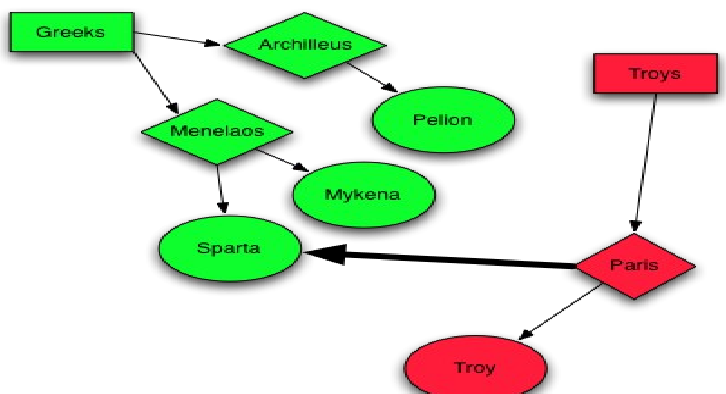

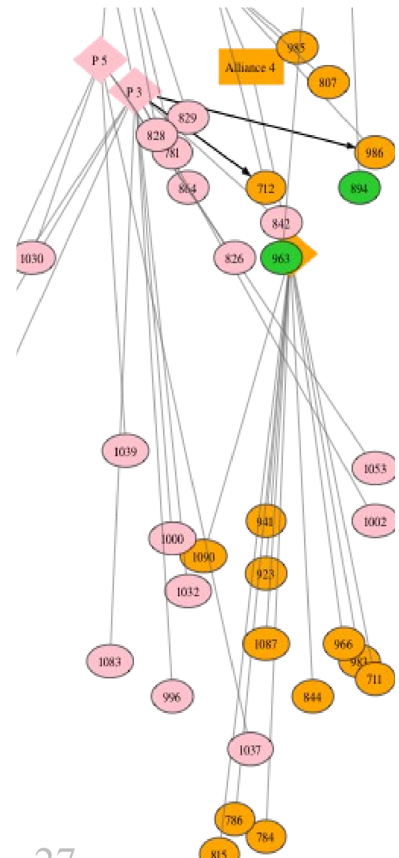

Again (AIPSML): Decision networks

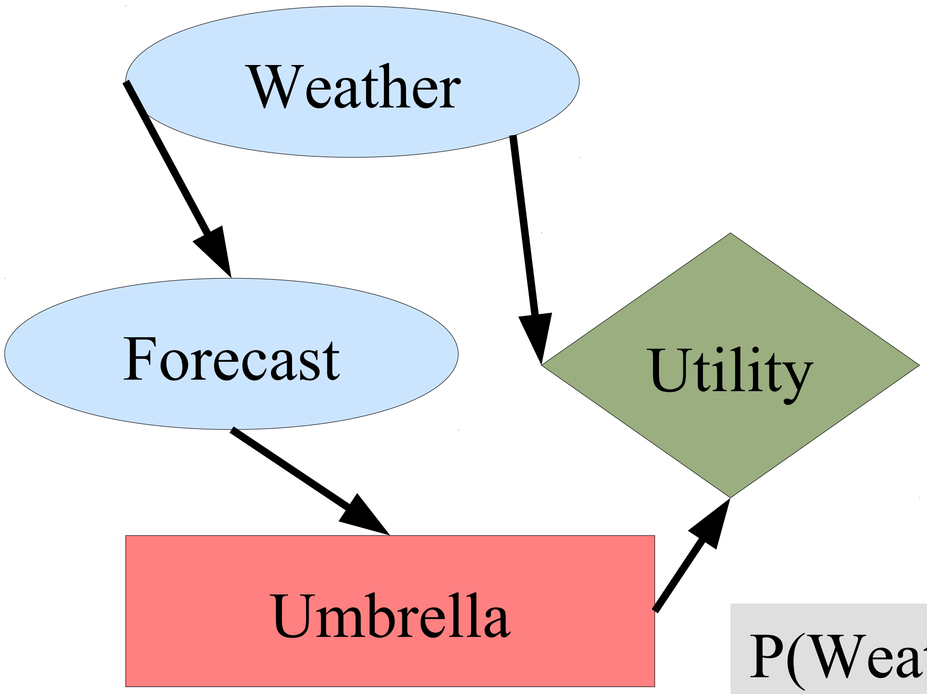

Weather	Forecast	Probability
norain	sunny	0.7
norain	cloudy	0.2
norain	rainy	0.1
rain	sunny	0.15
rain	cloudy	0.25
rain	rainy	0.6

$P(\text{Weather} = \text{rain}) = 0.3$

Weather	Umbrella	Utility
norain	takeIt	20
norain	leaveIt	100
rain	takeIt	70
rain	leaveIt	0

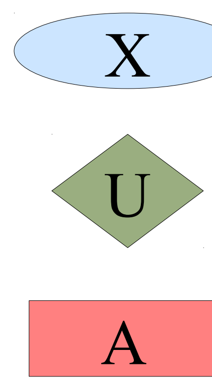

Chance node

Utility node

Decision node

Relational Reinforcement Learning

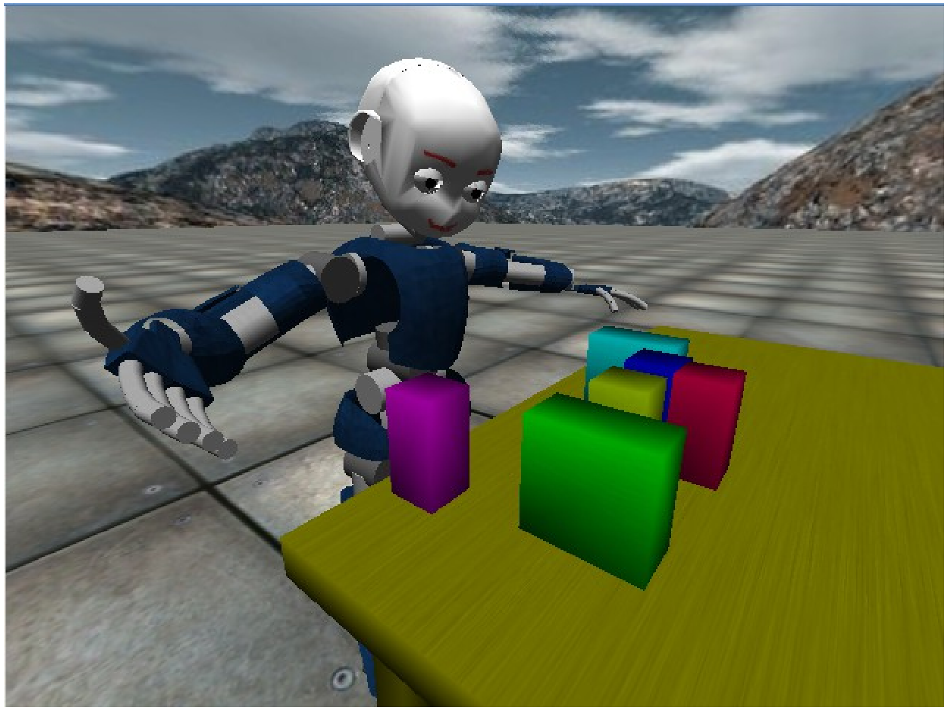

Representation and generalization in terms of **objects** and **relations**
Logical learning, reasoning, planning

Decision-theoretic planning
Reinforcement learning

Before state-of-the-art was **propositional**

Logic for abduction/deduction/induction
Probability for uncertainty
Utility for optimality and learning

→ decision-theoretic high-level cognition

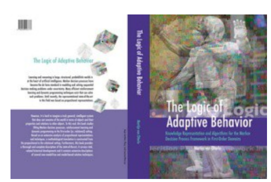

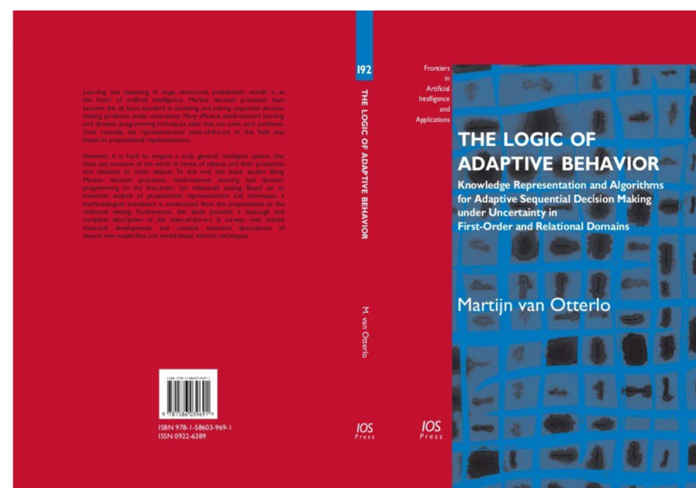

IOS Press (2009)

Model-free Value-based

Original RRL approach

by Dzeroski, De Raedt and Blockeel (ICML 98)

Episode based, tree induction (batch), Q-learning

Small, deterministic Blocks Worlds

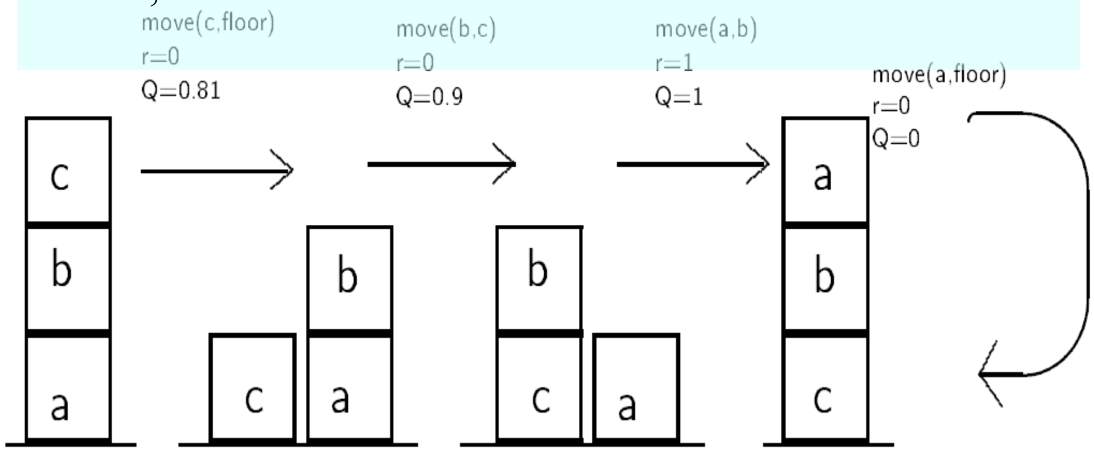


```
action(move(A,B)) , goal(on(C,D))
on(C,D) ?
+--yes: [0]
+--no:  action(move(C,D)) ?
        +--yes: [1]
        +--no:  action(move(D,B)) ?
                +--yes: [0.9]
                +--no:  [0.81]
```

```
qvalue(0) :-
  action(move(A,B)) , goal(on(C,D)) ,
  on(C,D), !.
qvalue(1) :-
  action(move(A,B)) , goal(on(C,D)) ,
  action(move(C,D)), !.
qvalue(0.9) :-
  action(move(A,B)) , goal(on(C,D)) ,
  action(move(D,B)), !.
qvalue(0.81).
```

qvalue(0.81). action(move(c,floor)). goal(on(a,b)). clear(c). on(c,b). on(b,a). on(a,floor).	qvalue(0.9). action(move(b,c)). goal(on(a,b)). clear(b). clear(c). on(b,a). on(a,floor). on(c,floor).	qvalue(1.0). action(move(a,b)). goal(on(a,b)). clear(a). clear(b). on(b,c). on(a,floor). on(c,floor).	qvalue(0.0). action(move(a,floor)). goal(on(a,b)). clear(a). on(a,b). on(b,c). on(c,floor).

T1a

DT-Problog (AAAI-2010)

Van den Broeck/Thon/van Otterlo/De Raedt
Decision-theoretic Prolog

Decision Facts

```
? :: umbrella.
? :: raincoat.
```

Probabilistic Facts

```
0.3 :: rainy.
0.5 :: windy.
```

Background Knowledge

```
dry :- rainy, umbrella, not(broken_umbrella).
dry :- rainy, raincoat.
dry :- not(rainy).

broken_umbrella :- umbrella, rainy, windy.
```

Utility Facts

```
umbrella => -2.
raincoat => -20.
```

```
dry => 60.
broken_umbrella => -40.
```

ProbLog

Probabilistic Facts

```
0.3 :: buy_trust(_,_).
0.2 :: buy_marketing(_).
```

Background Knowledge

```
buys(X) :-
    trusts(X,Y),
    buys(Y),
    buy_trust(X,Y).
```

```
buys(X) :-
    marketed(X),
    buy_marketing(X).
```

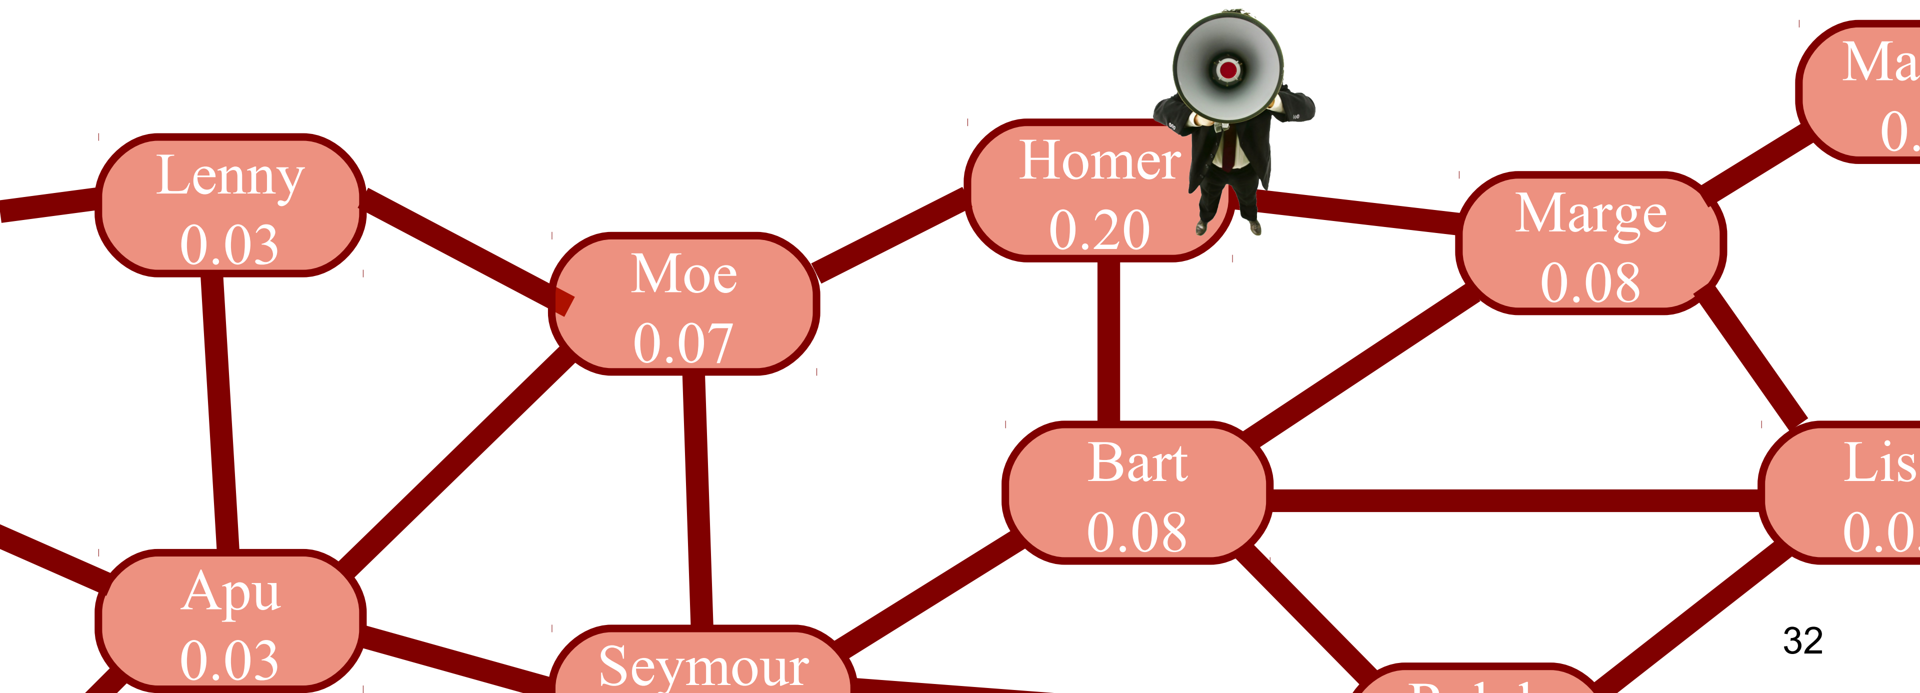

Probabilistic Facts

...

Background Knowledge

...

Decisions

```
? :: marketed(P) :- person(P).
```

Utility Facts

```
buys(P) => 5 :- person(P).
marketed(P) => -3 :- person(P).
```

DTProbLog

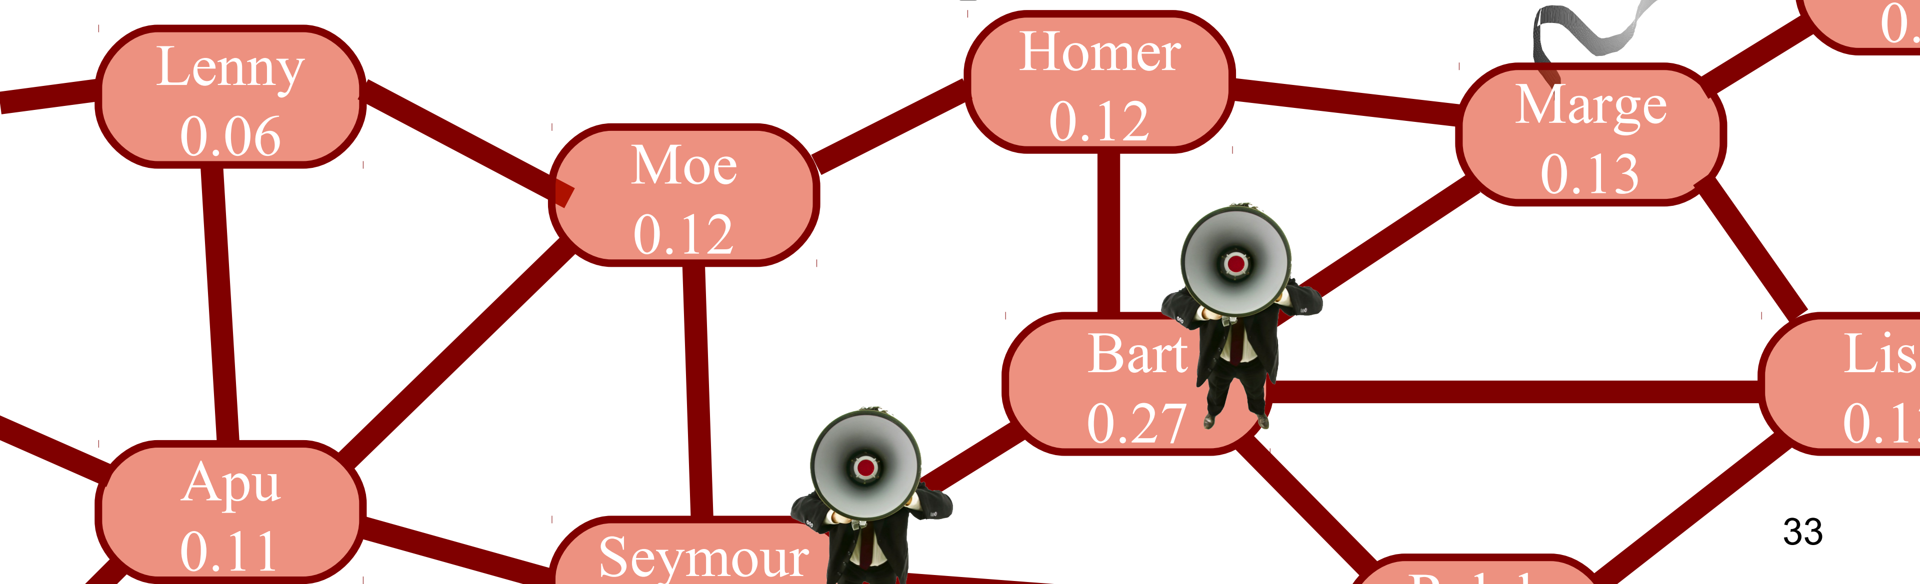

eindslide