

From Statistical Relational AI to Neural Symbolic Computation

Luc De Raedt, Sebastijan Dumancic, Robin Manhaeve, Giuseppe Marra
firstname.lastname@kuleuven.be

reusing some slides from previous tutorials with
Angelika Kimmig, Kristian Kersting, David Poole, and Sriraam Natarajan



LEUVEN.AI INSTITUTE



WASP | WALLENBERG AI AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM



You will find a version of this tutorial and additional content at

<https://dtai.cs.kuleuven.be/tutorials/nesytutorial>
(after the tutorial)

See also

De Raedt, Dumancic, Marra, Manhaeve
From Statistical Relational to Neuro-Symbolic Artificial Intelligence
IJCAI 20, and long version on AIJ 24



Ministry of Culture and Science
of the State of
North Rhine-Westphalia



 Funded by
the European Union



LEUVEN.AI INSTITUTE

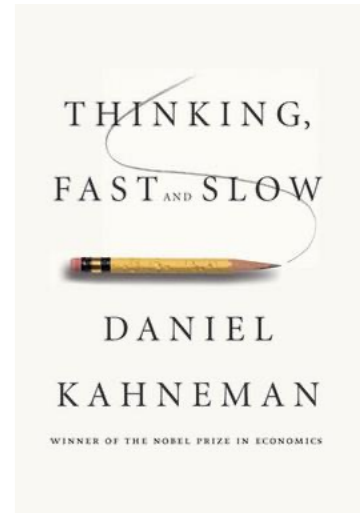


WASP | WALLENBERG AI
AUTONOMOUS SYSTEMS
AND SOFTWARE PROGRAM



Introduction

Learning and Reasoning both needed



- System 1 - thinking fast - can do things like $2+2 = ?$ and recognise objects in image
- System 2 - thinking slow - can reason about solving complex problems - planning a complex task
- alternative terms — data-driven vs knowledge-driven, symbolic vs subsymbolic, solvers and learners, neuro-symbolic...
- **A lot of work on integrating learning and reasoning, neural symbolic computation to integrate logic / symbols reasoning with neural networks**

- see also arguments by Marcus, Darwiche, Levesque, Tenenbaum, Geffner, Bengio, Le Cun, Kautz, ...
see also AI Debates

Real-life problems involve two important aspects.

 AUTO



<https://www.theorie-blokken.be/nl/gratis-proefexamen>

Who can go first ?

- A. The red car
- B. The blue van
- C. The white car

Real-life problems involve two important aspects.



Who can go first ?

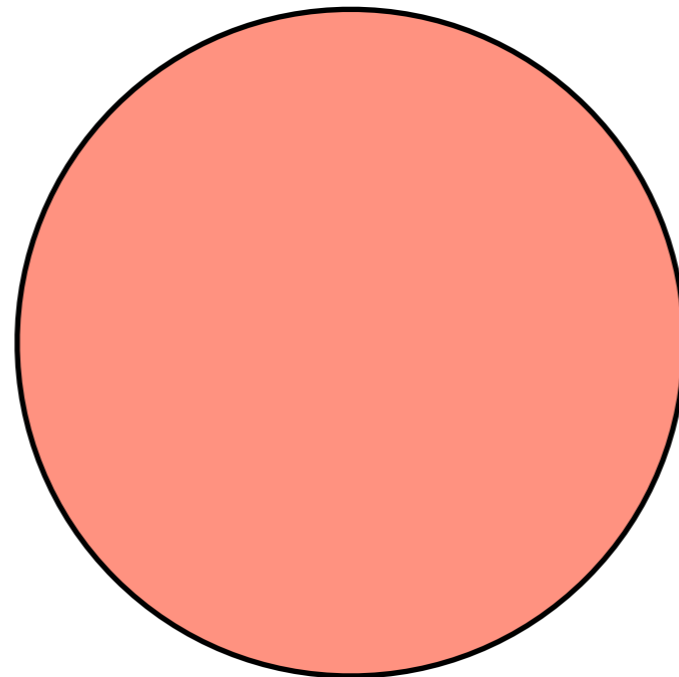
- A. The red car
- B. The blue van
- C. The white car

Reasoning

Sub-symbolic perception

Thinking fast

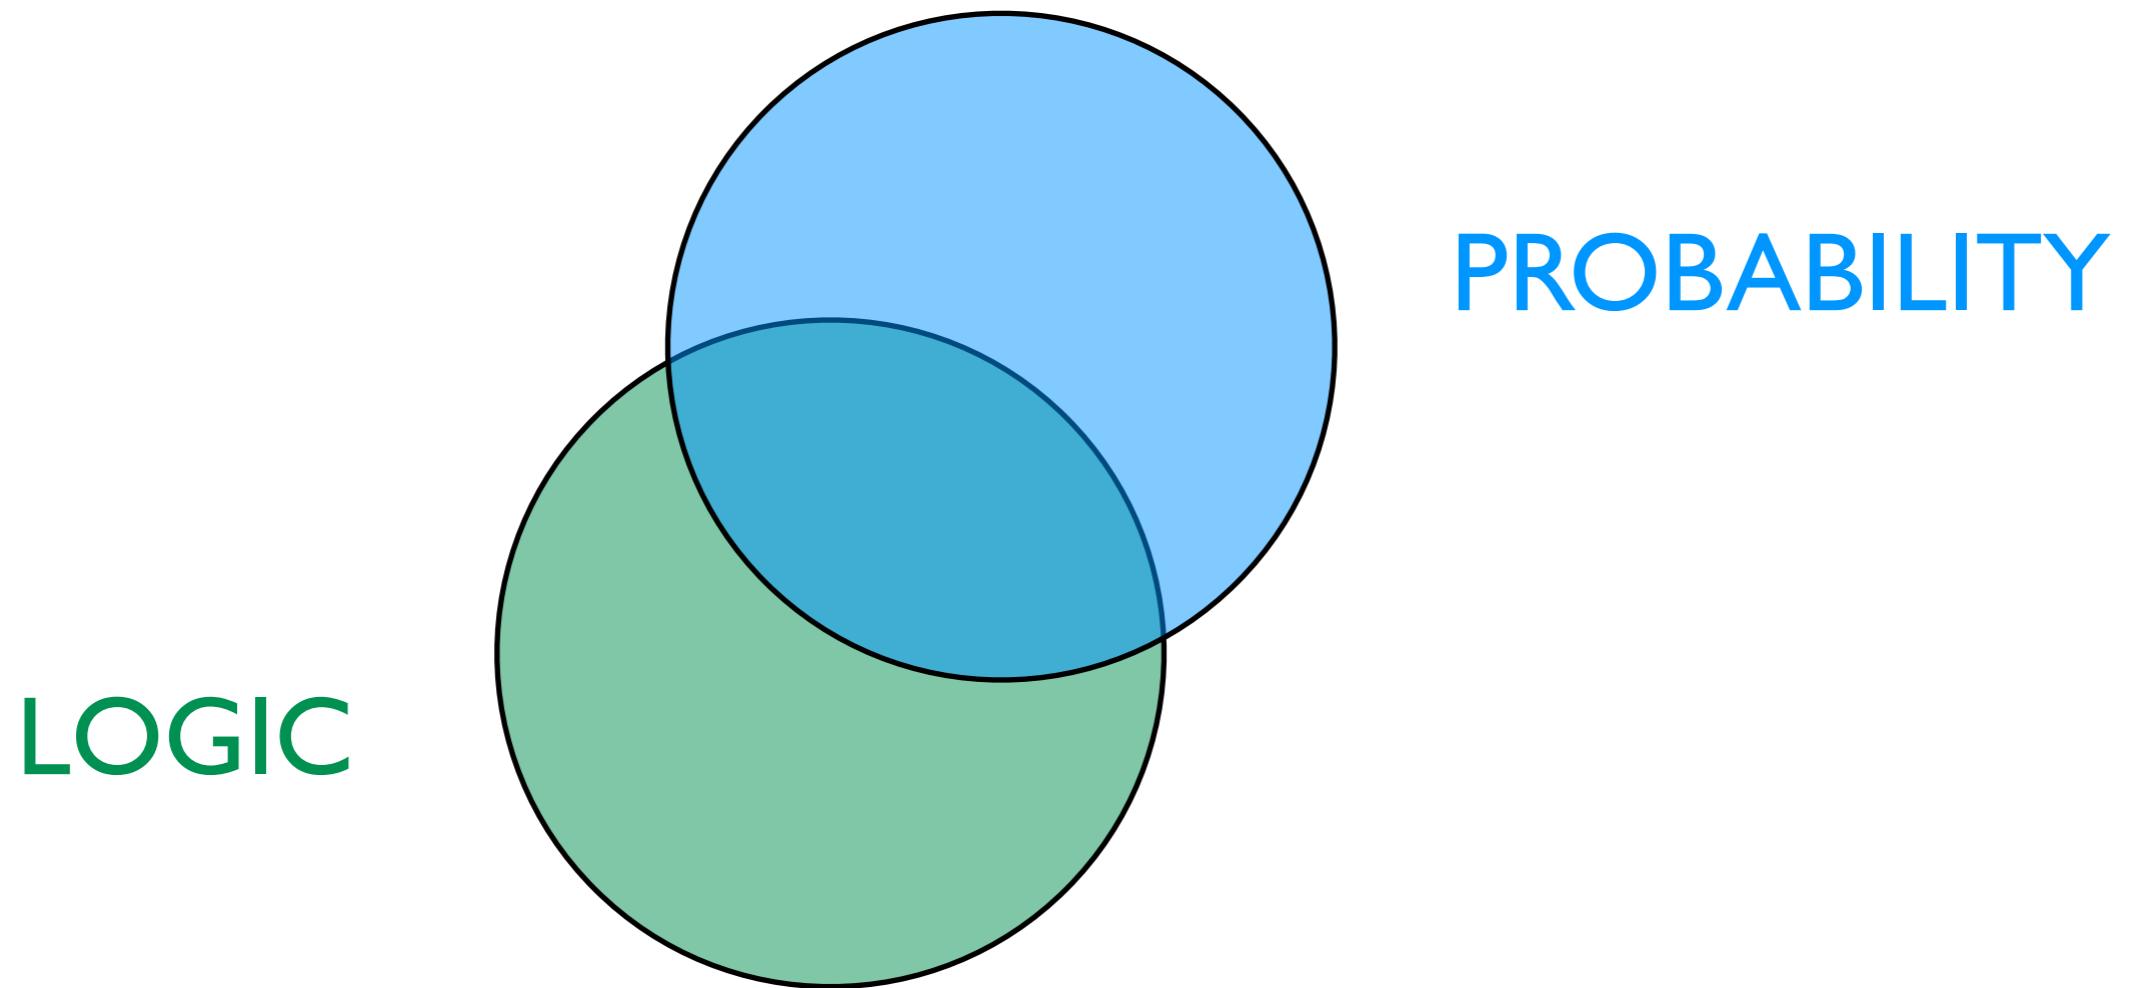
MAIN PARADIGM in AI
Focus on Learning



NEURAL

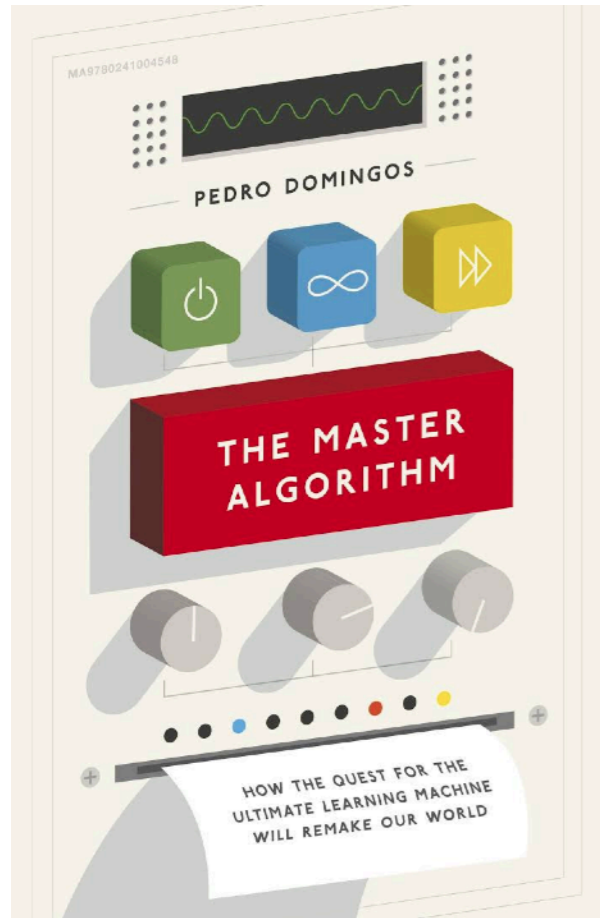
Thinking slow = reasoning

TWO MAIN PARADIGMS in AI

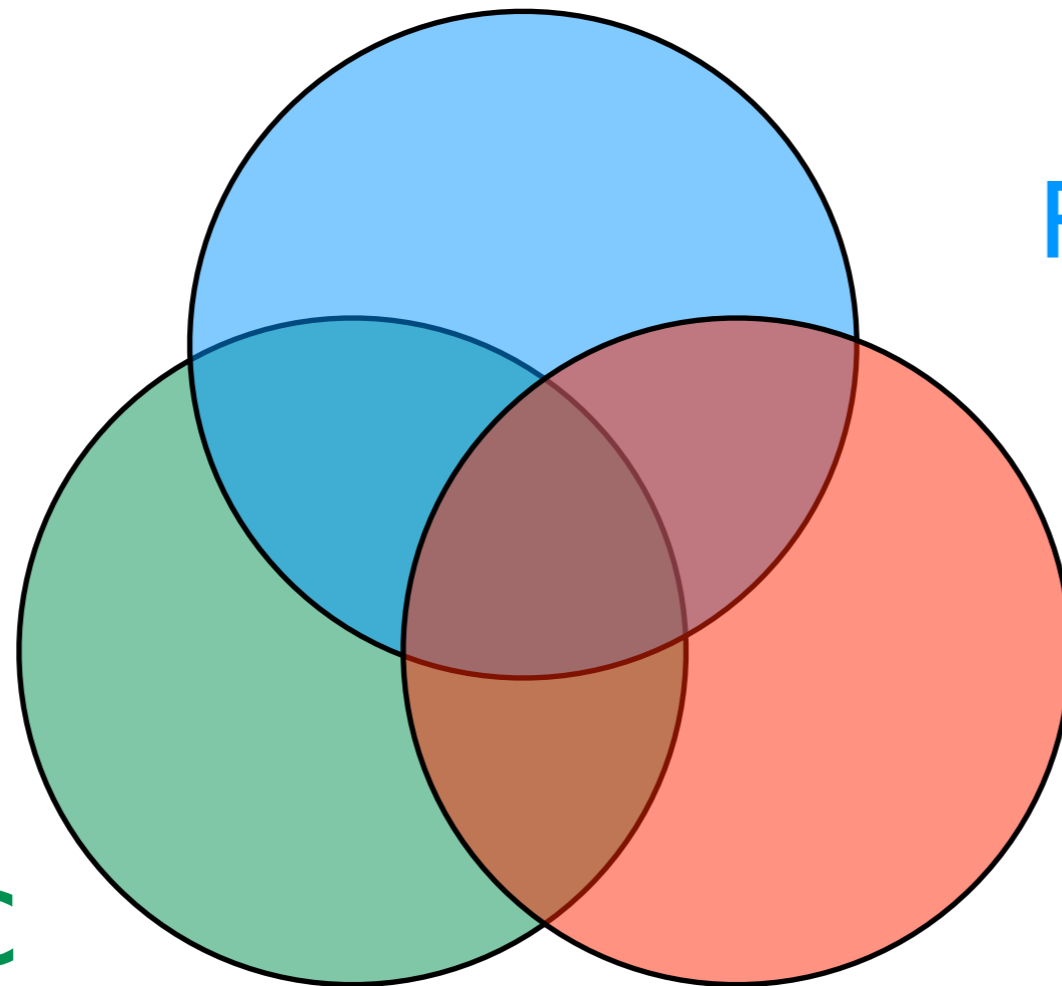


Their integration has been well studied in Probabilistic (Logic) Programming and Statistical Relational AI (StarAI)

Learning



LOGIC

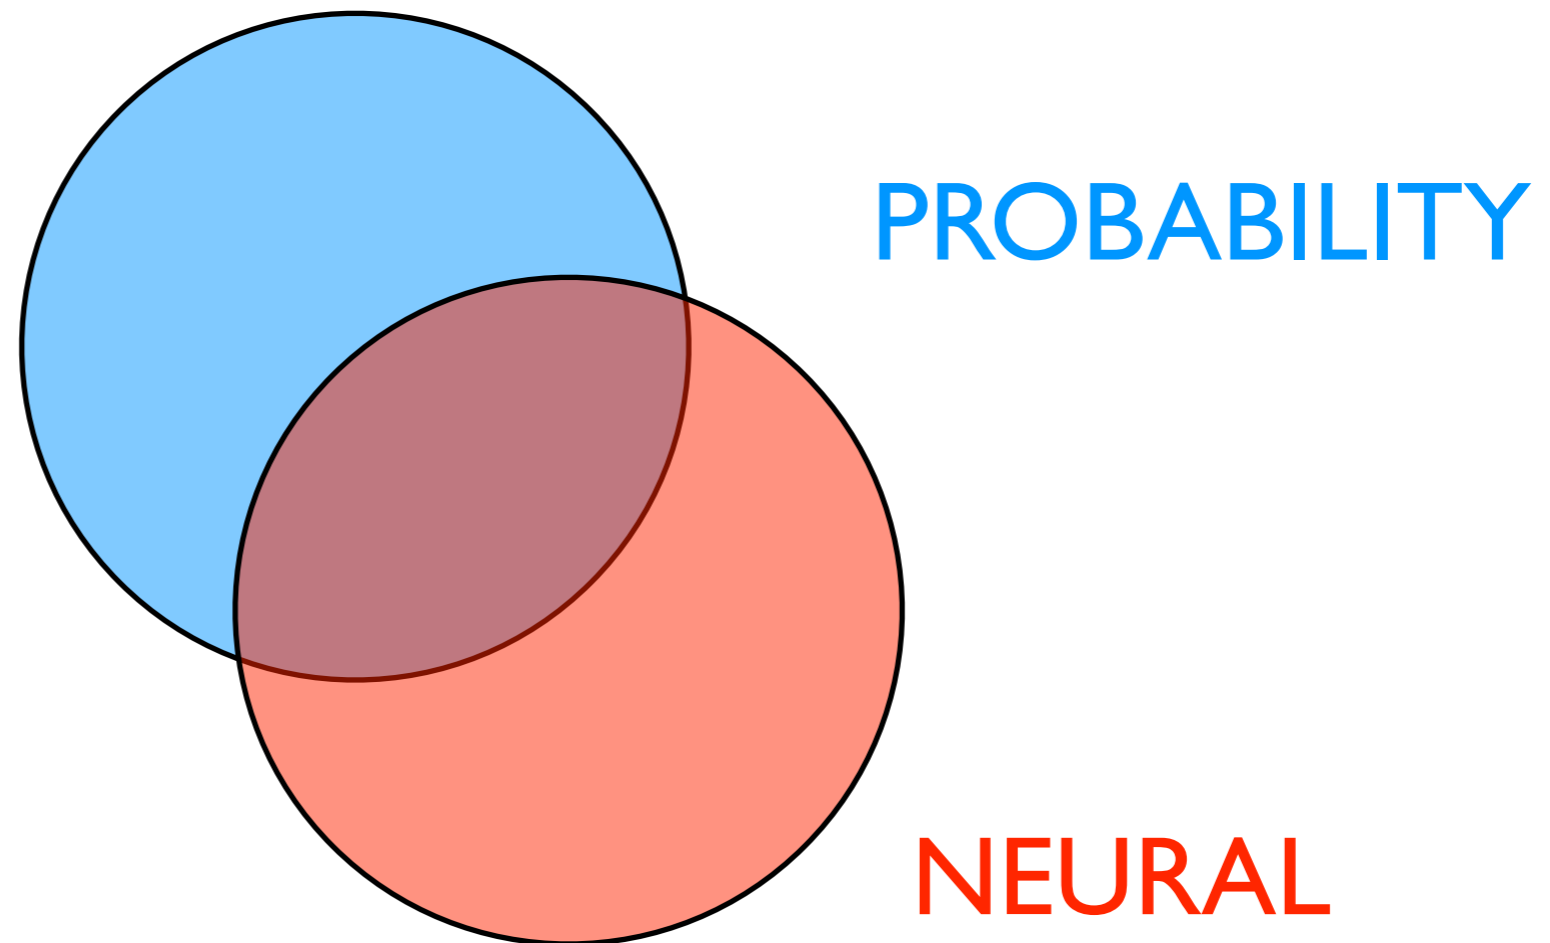


PROBABILITY

NEURAL

How to integrate these three paradigms in AI ?

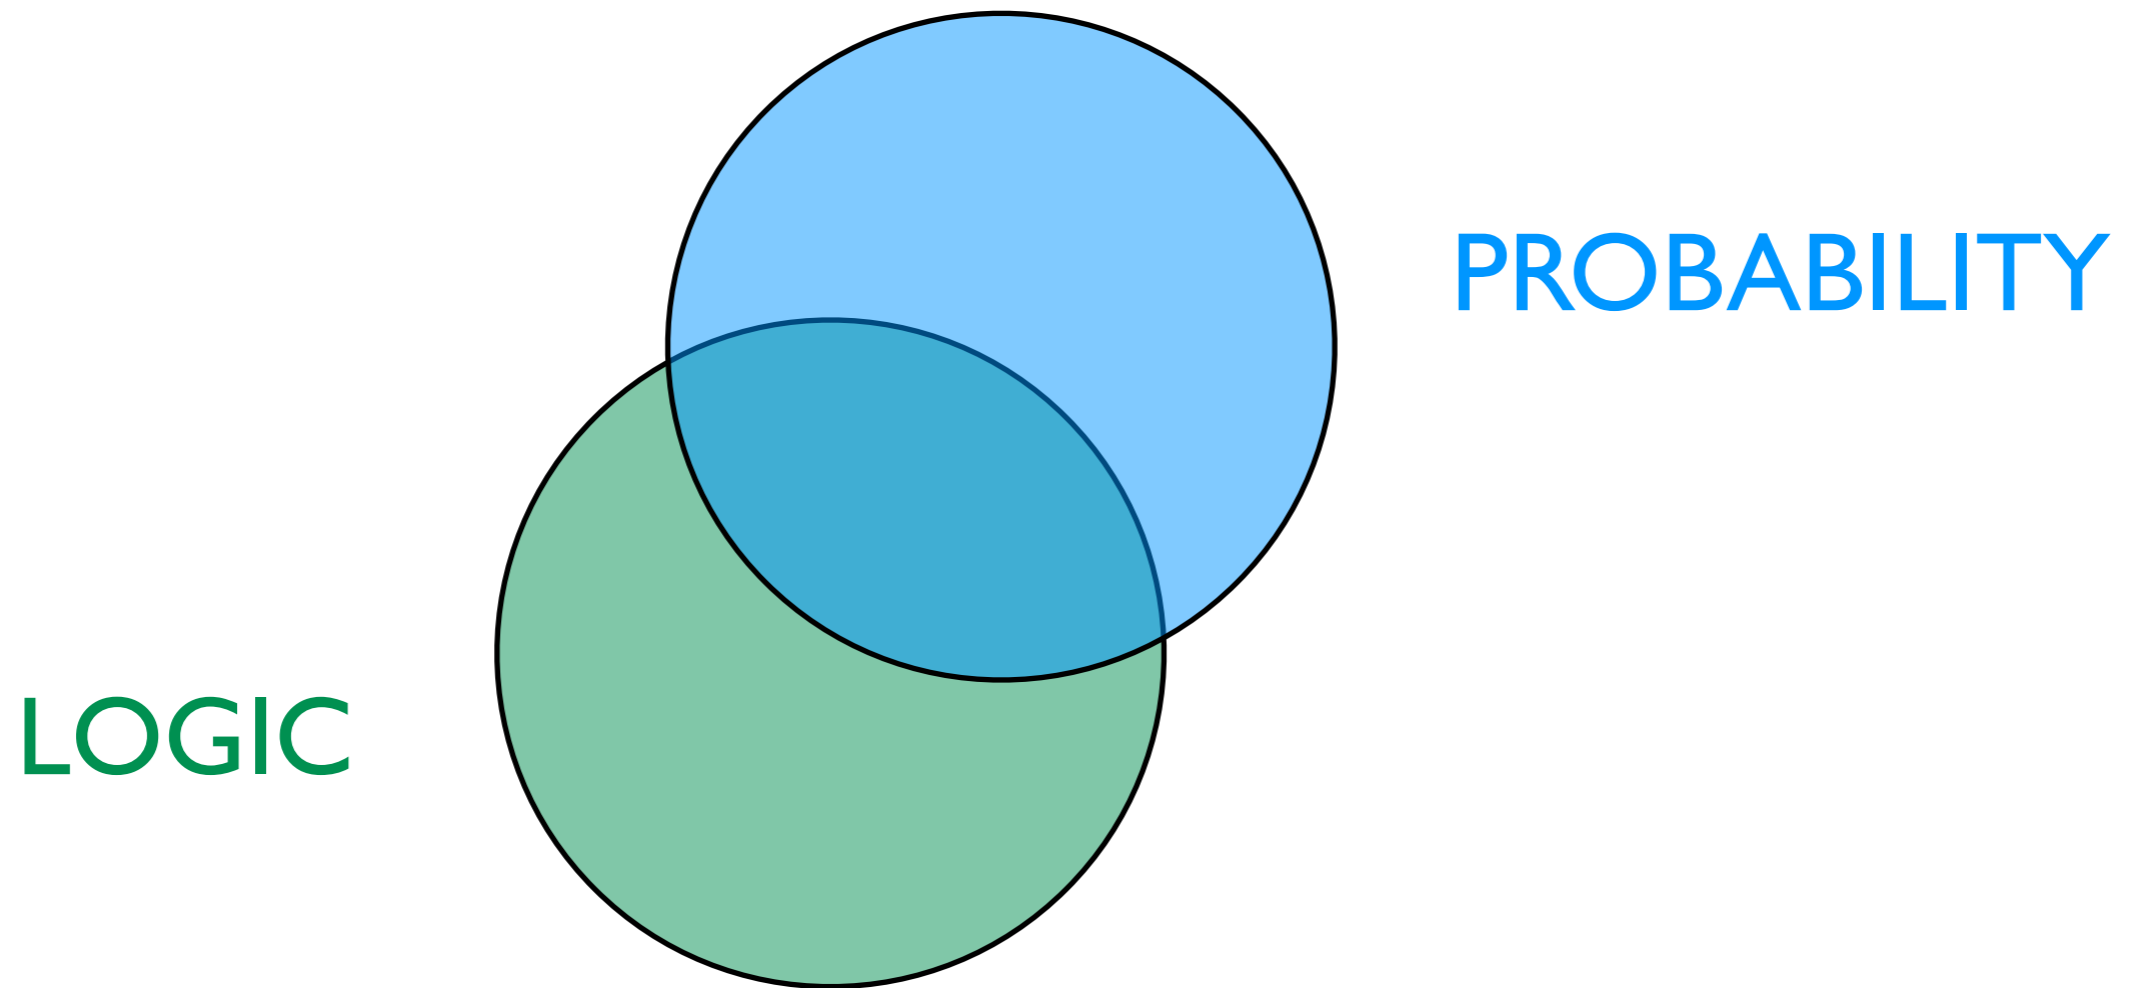
A lot of ML



**Well studied from a LEARNING perspective
in Deep Learning**

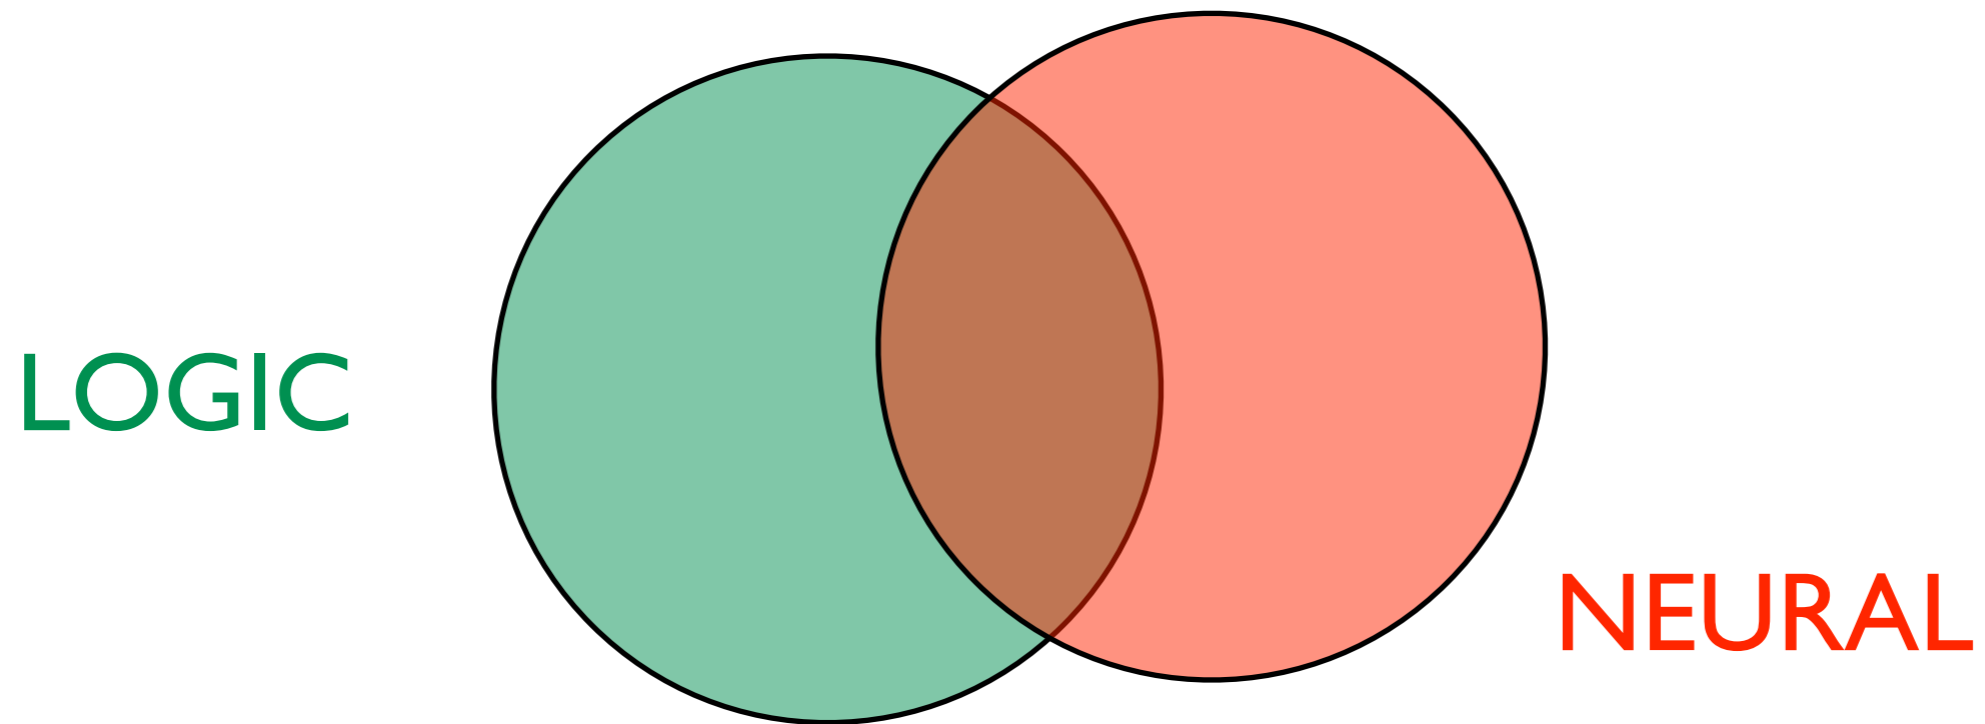
Thinking slow = reasoning

TWO MAIN PARADIGMS in AI



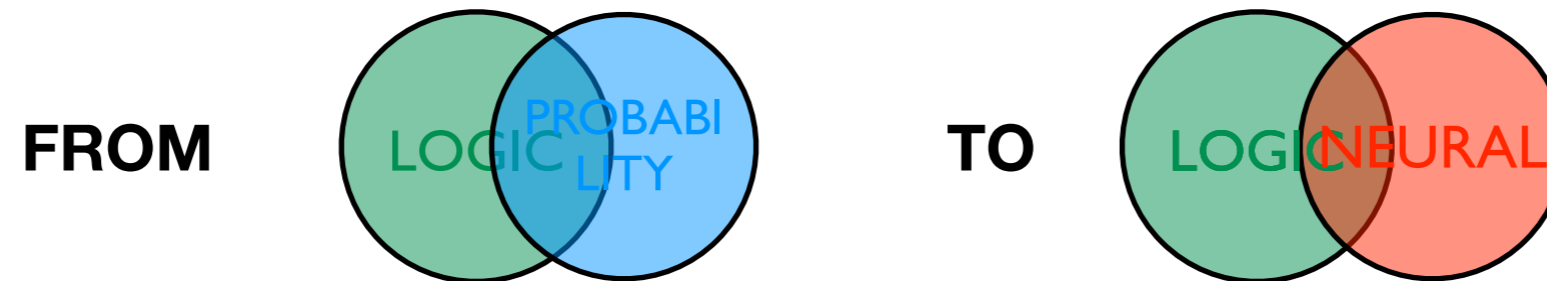
Their integration has been well studied in Probabilistic (Logic) Programming and Statistical Relational AI (StarAI)

State of the Art



**Being studied from a LEARNING perspective
in Neuro Symbolic Computation**

Key Message



**StarAI and NeSy share similar problems
and thus similar solutions apply**



WARNING

TALK MAY NOT COVER ALL of
NESY

See also

De Raedt, Dumancic, Marra, Manhaeve

From Statistical Relational to Neuro-Symbolic Artificial Intelligence

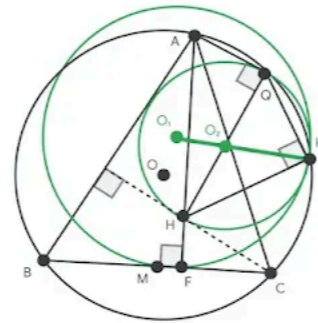
IJCAI 20, and long version on AIJ 24

Applications

Alpha Geometry

IMO 2015 P3

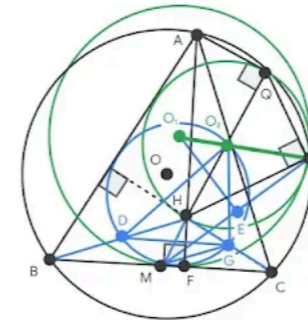
Let ABC be an acute triangle. Let (O) be its circumcircle, H its orthocenter, and F the foot of the altitude from A . Let M be the midpoint of BC . Let Q be the point on (O) such that $QH \perp QA$ and let K be the point on (O) such that $KH \perp KQ$. Prove that the circumcircles (O_1) and (O_2) of triangles FKM and KQH are tangent to each other.



AlphaGeometry

Solution

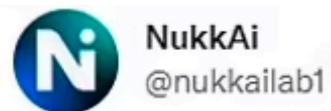
[...]
 Construct D: midpoint BH [a]
 [a], O_2 midpoint HQ \Rightarrow BQ \parallel O_2 D [20]
 [...]
 Construct G: midpoint HC [b]
 $\angle GMD = \angle GO_2D \Rightarrow$ M O_2 G D cyclic [26]
 [...]
 [a], [b] \Rightarrow BC \parallel DG [30]
 [...]
 Construct E: midpoint MK [c]
 [c] \Rightarrow $\angle KFC = \angle KO_1E$ [104]
 [...]
 $\angle FKO_1 = \angle FKO_2 \Rightarrow$ $KO_1 \parallel KO_2$ [109]
 [109] \Rightarrow O_1, O_2, K collinear \Rightarrow $(O_1)(O_2)$ tangent



AlphaGeometry solving an Olympiad problem: Problem 3 of the 2015 International Mathematics Olympiad (left) and a condensed version of AlphaGeometry's solution (right). The blue elements are added constructs. AlphaGeometry's solution has 109 logical steps.

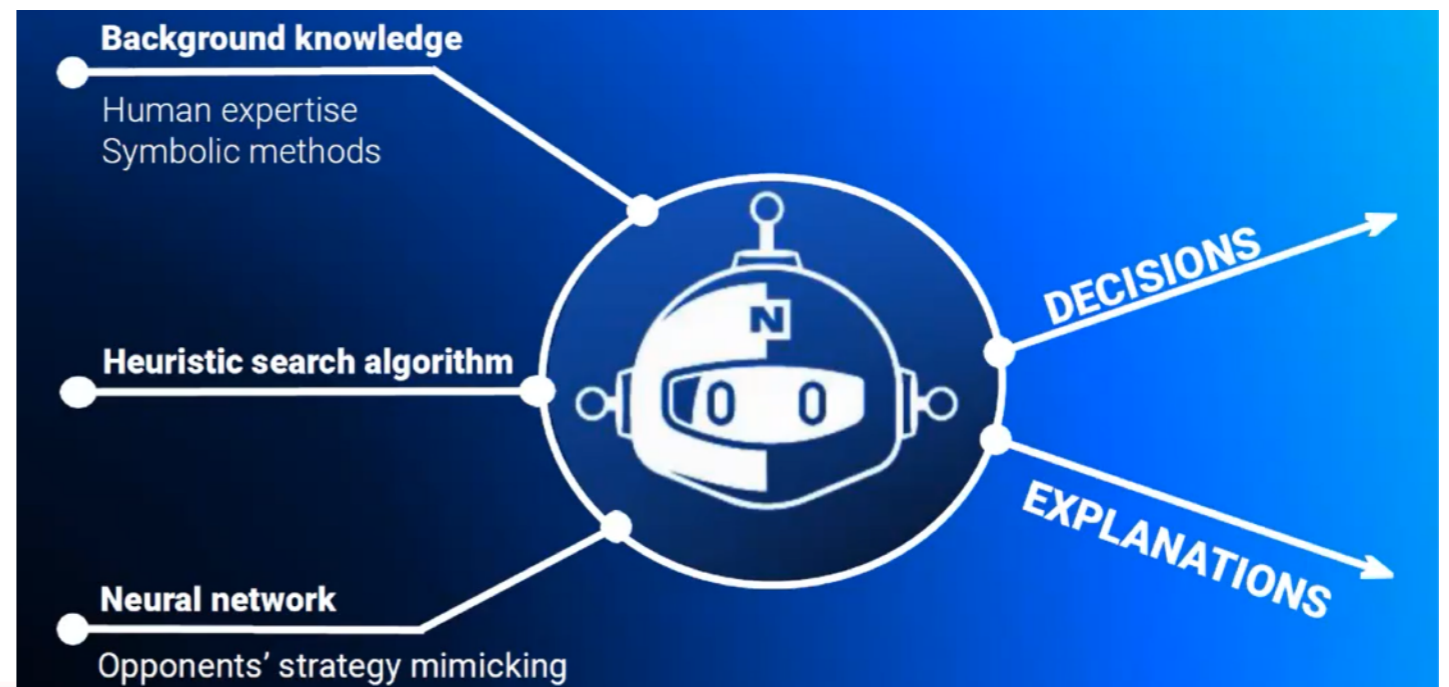
(New) Game Playing

The NeSy Nook system
defeats eight
world bridge champions
in Paris (2022)



🤖 : 6136
👤 : 5238

Nook won The Nukkai Challenge!



<https://challenge.nukk.ai/>

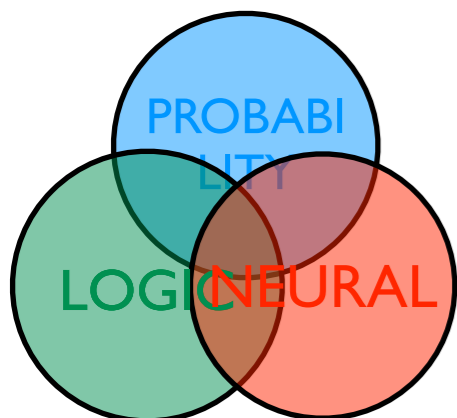
Addition

Learn to add the sum of lists of MNIST images


$$\mathbf{3} \mathbf{5} \mathbf{0} \mathbf{4} \mathbf{1} + \mathbf{9} \mathbf{2} \mathbf{1} = ? \quad \mathbf{35962}$$

example multi-addition predicate

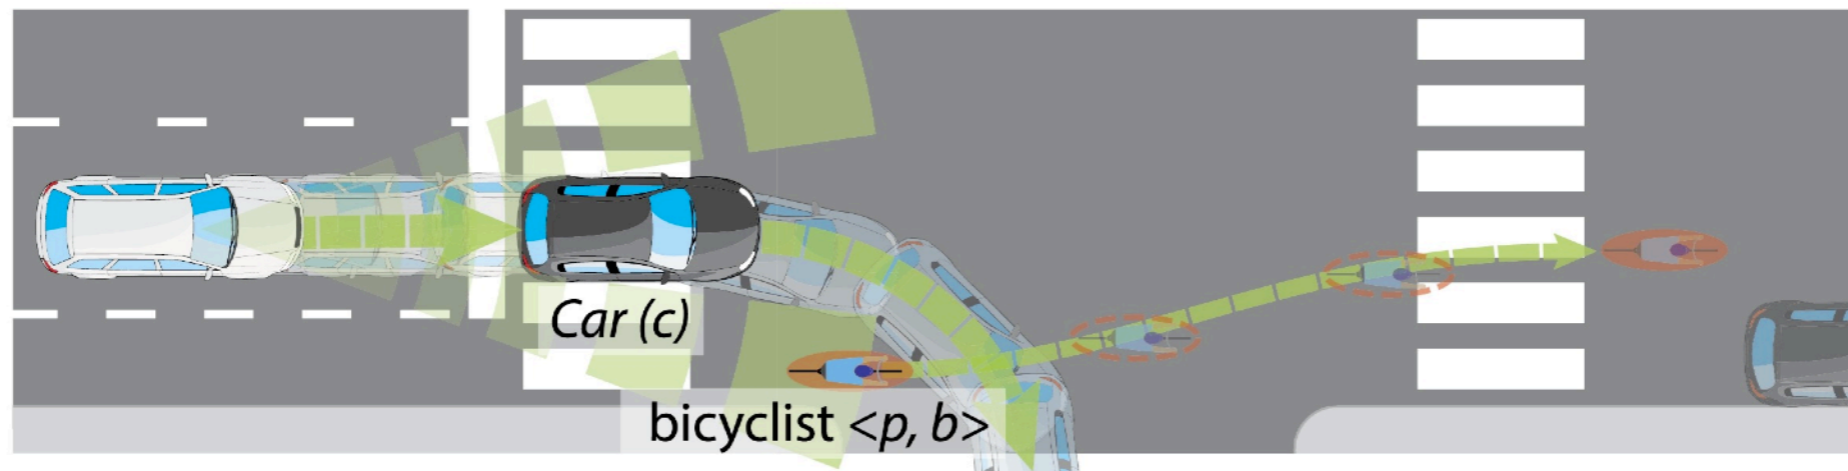
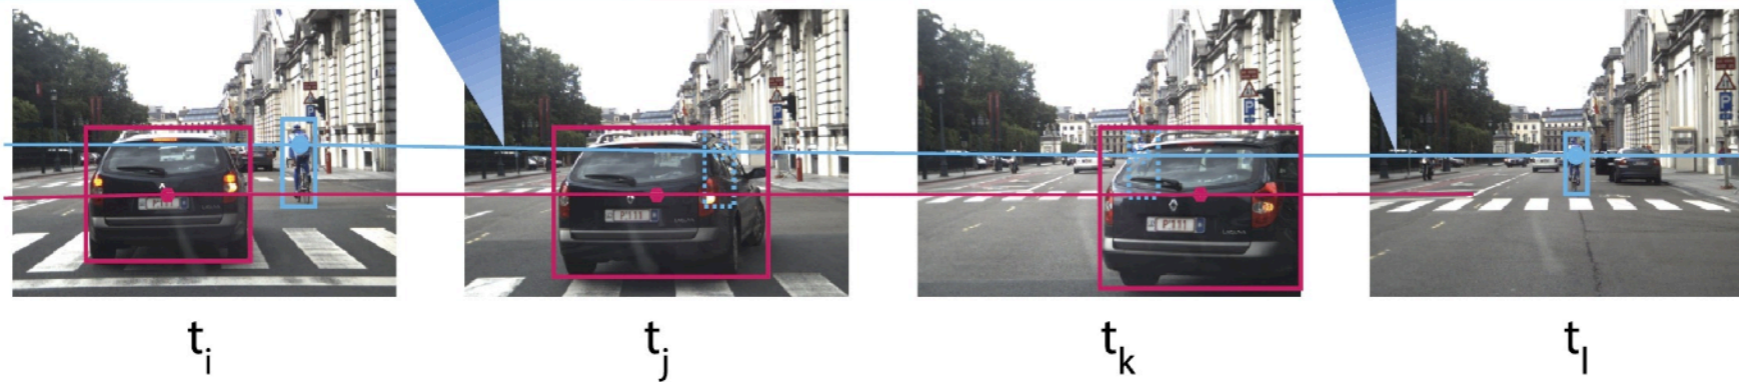
Assume you do not know how to map MNIST images to numbers, but do know the rules of addition. Can you learn from these examples how to map MNIST to numbers ?



Emerging applications

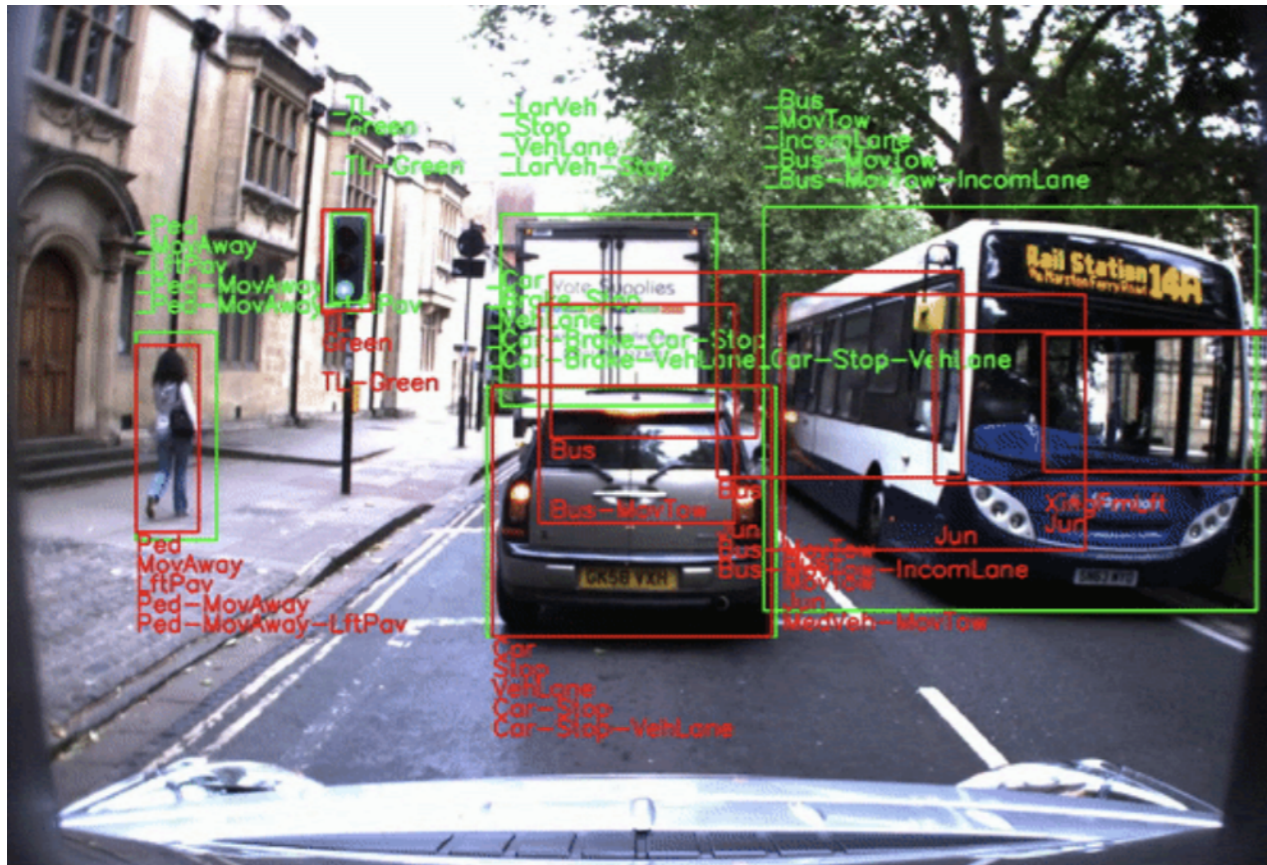


bicyclist $\langle p, b \rangle$ gets occluded by **Car (c)** bicyclist $\langle p, b \rangle$ reappears from behind **car (c)**



From Suchan, Bhatt and Varadarajan, AIJ 21

ROAD-R: The autonomous driving dataset with logical requirements



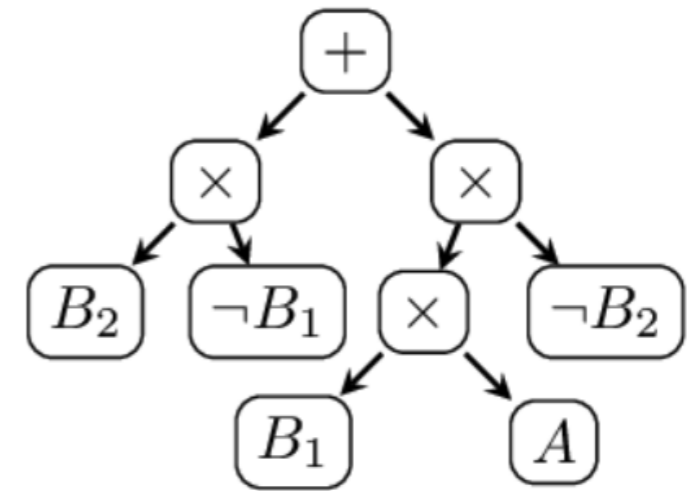
Natural Language Explanations

If an agent pushes an object then it is a pedestrian
 A pedestrian can only push objects, move away, etc.
 Only pedestriains, cars, cyclists, etc. can cross from left
 Only pedestrians and cyclists can wait to cross
 Only pedestrians, cars, cyclists, etc can stop
 Only pedestrians, cars, cyclists, etc can move
 Only pedestrians, cars, cyclists, etc can move towards
 Only pedestrians, cars, cyclists, etc can move away
 An emergency vehicle can only overtake, move away etc.
 Only emergency vehicles, cars etc. can have hazards lights on
 A bus can only overtake, move away move towards etc.
 A medium vehicle can only overtake, move away, move towards etc.

Giunchiglia, Eleonora, Mihaela Cătălina Stoian, Salman Khan, Fabio Cuzzolin, and Thomas Lukasiewicz. "ROAD-R: The autonomous driving dataset with logical requirements." *Machine Learning* (2023): 1-31.

ROAD-R: The autonomous driving dataset with logical requirements

- **Task:** road event-detection
multi-label classification
with constraints
- **Solution:** neuro-symbolic AI
Calculate most probable explanation
given constraints and neural outputs

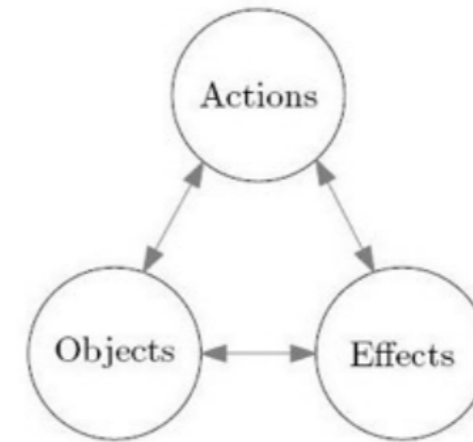


Giunchiglia, Eleonora, Mihaela Cătălina Stoian, Salman Khan, Fabio Cuzzolin, and Thomas Lukasiewicz. "ROAD-R: The autonomous driving dataset with logical requirements." *Machine Learning* (2023): 1-31.

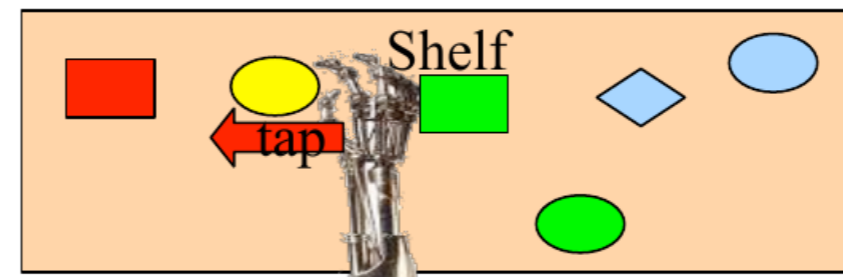
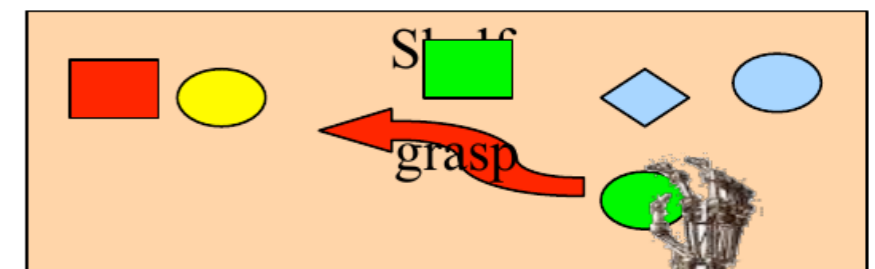
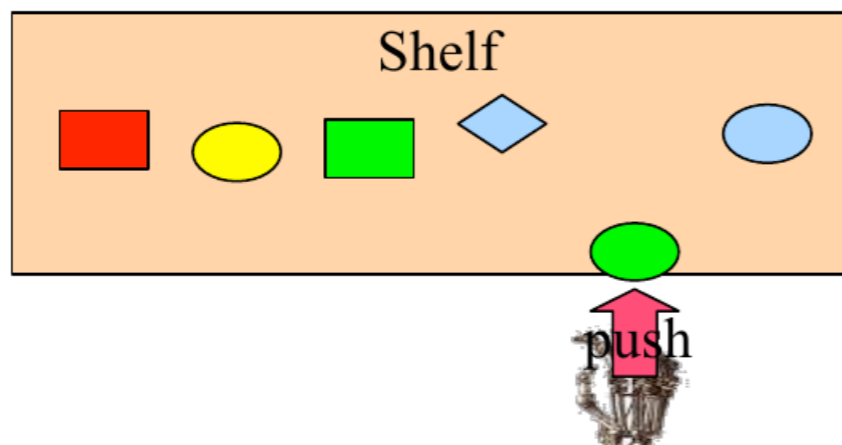
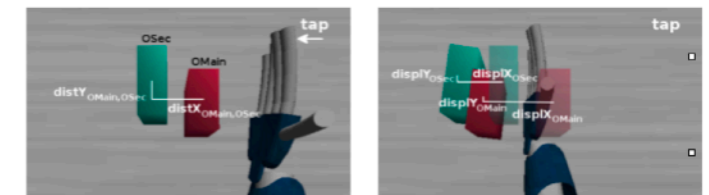
Relational Affordances

- **Object Affordance:**
What can one do with particular object?
- **Relational Affordance:**
in a particular context?

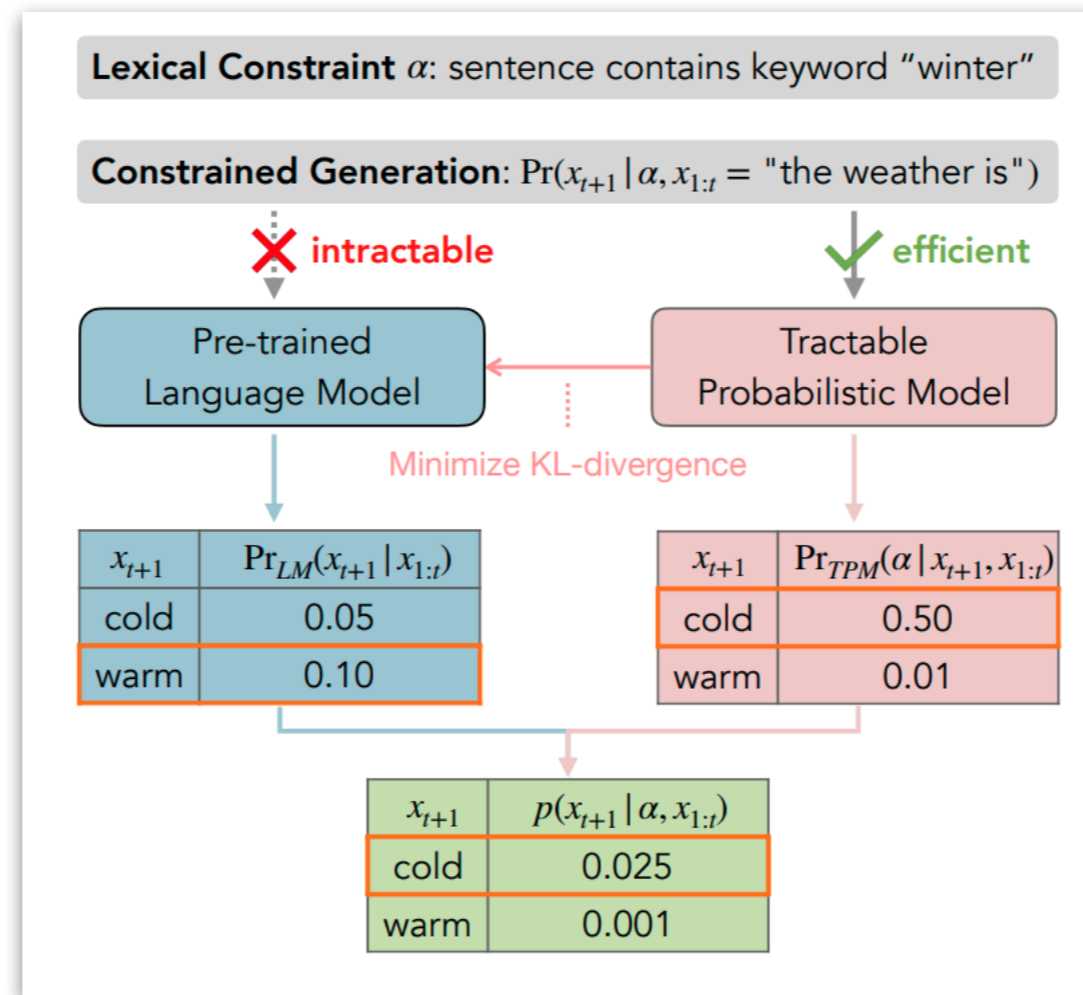
with multiple objects and relations among them
- Use of statistical relational learning, probabilistic programming for learning, reasoning and planning !



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



Constrained output of LLMs



Zhang, Honghua, Meihua Dang, Nanyun Peng, and Guy Van den Broeck. "Tractable control for autoregressive language generation." In International Conference on Machine Learning, pp. 40932-40945. PMLR, 2023.

Probabilistic Logic Shield for Reinforcement Learning

Wen-chi Yang et al, IJCAI 23 Distinguished paper award

Assuming noisy sensors



$$\begin{cases} \pi(\text{accelerate} | s) = 0.5 \\ \pi(\text{left} | s) = 0.3 \\ \pi(\text{right} | s) = 0.2 \end{cases}$$

```
0.8 :: obstc(front).
0.2 :: obstc(left).
0.5 :: obstc(right).
```

```
0.5 :: act(accel);
0.3 :: act(left);
0.2 :: act(right)
```

```
0.9 :: crash:- obstc(front), act(accel).
0.4 :: crash:- obstc(left), act(left).
0.4 :: crash:- obstc(right), act(right).
safe:- ¬crash.
```

Shield

Will stay undamaged?

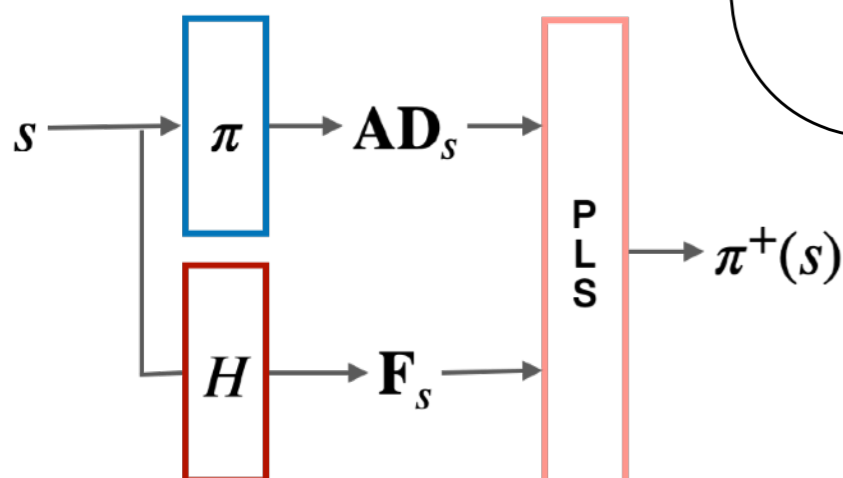
$$\mathbf{P}(\text{safe} | a, s) = \begin{cases} \text{accelerate} & \rightarrow 0.28 \\ \text{left} & \rightarrow 0.92 \\ \text{right} & \rightarrow 0.8 \end{cases}$$

Probability of staying safe if following π ?

$$\mathbf{P}_{\pi}(\text{safe} | s) = 0.576$$

What is a safer policy π^+ ?

$$\begin{cases} \pi^+(\text{accelerate} | s) = 0.24 \\ \pi^+(\text{left} | s) = 0.48 \\ \pi^+(\text{right} | s) = 0.28 \end{cases}$$



DeepProbLog Theory
(Manhaeve et al. AIJ)

Visual Reasoning and Question Answering

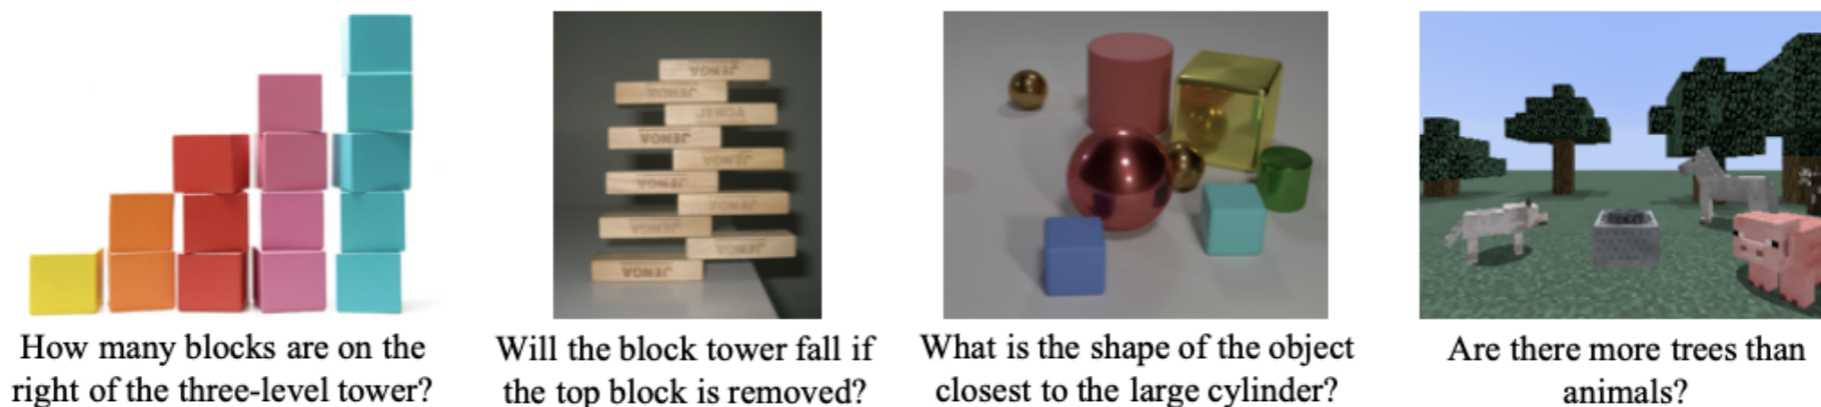


Figure 1: Human reasoning is interpretable and disentangled: we first draw abstract knowledge of the scene via visual perception and then perform logic reasoning on it. This enables compositional, accurate, and generalizable reasoning in rich visual contexts.

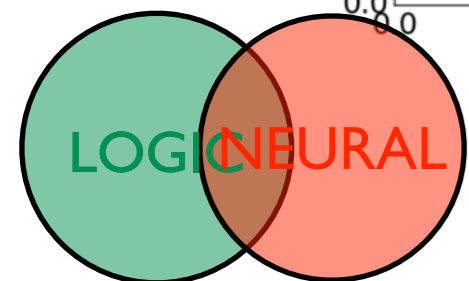
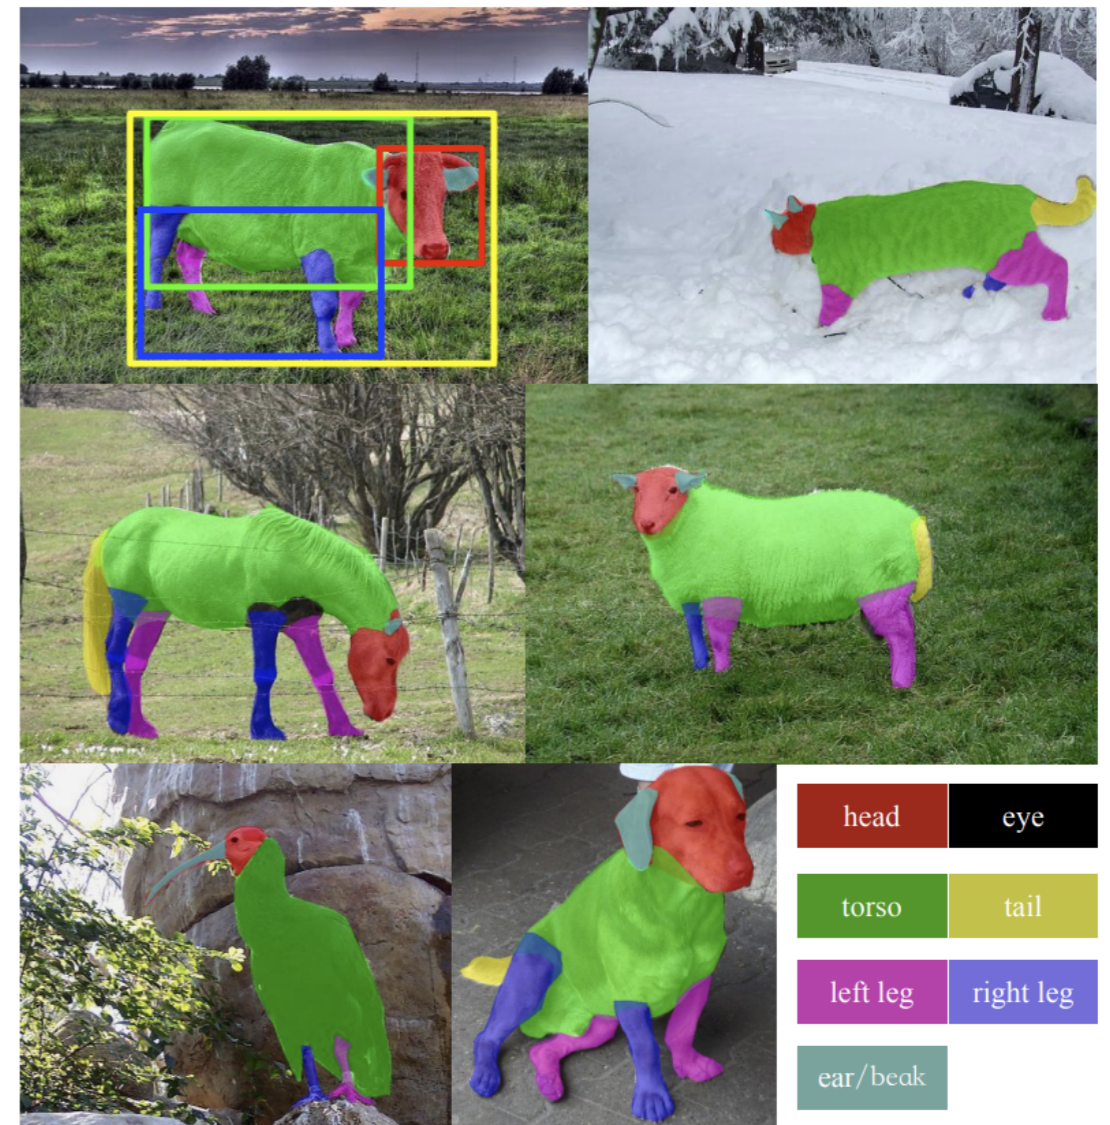
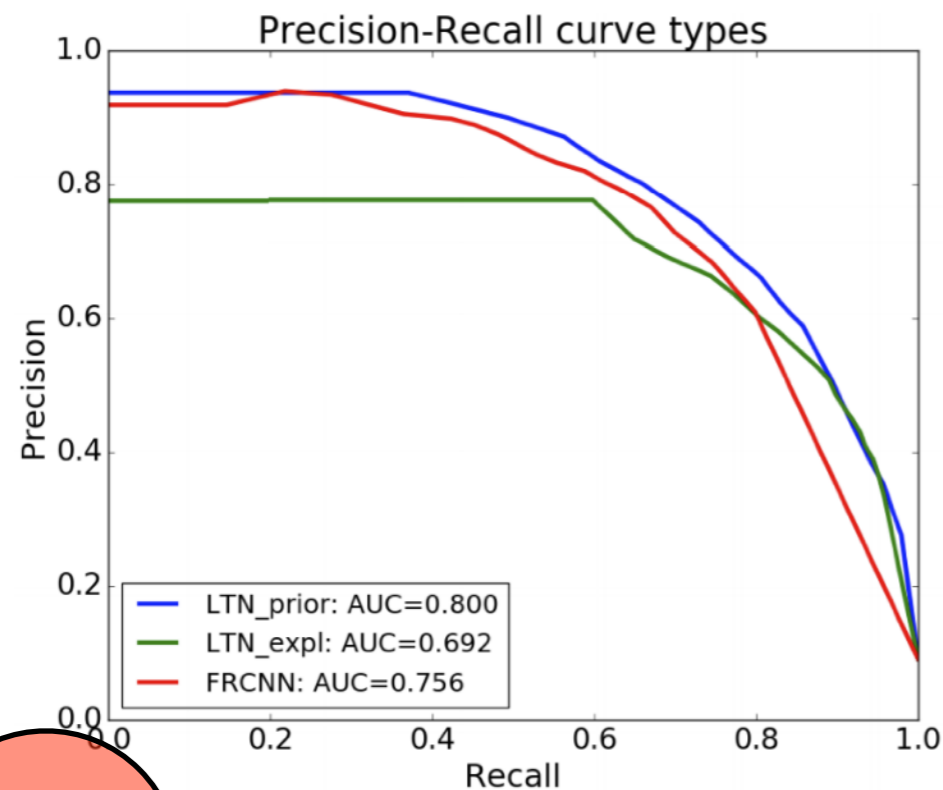
Adding a reasoning component on top of the perception can improve performance.

Semantic Image Interpretation

$$\forall xy(\text{partOf}(x, y) \rightarrow \neg \text{partOf}(y, x))$$

$$\forall xy(\text{Cat}(x) \wedge \text{partOf}(x, y) \rightarrow \text{Tail}(y) \vee \text{Muzzle}(y))$$

$$\forall xy(\text{Cat}(x) \rightarrow \neg \text{partOf}(x, y))$$

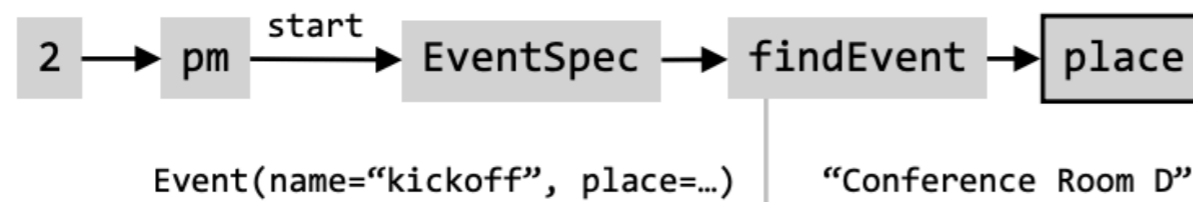


(New) Dialog Systems

User: *Where is my meeting at 2 this afternoon?*

```
place(findEvent(EventSpec(start=pm(2))))
```

(1)

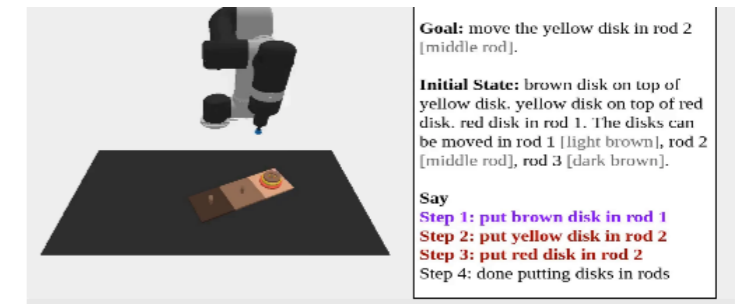
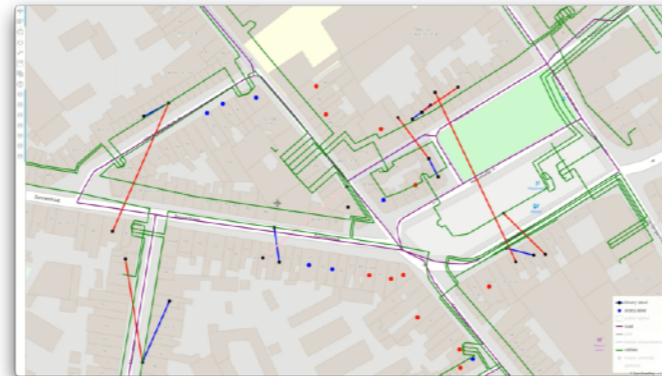
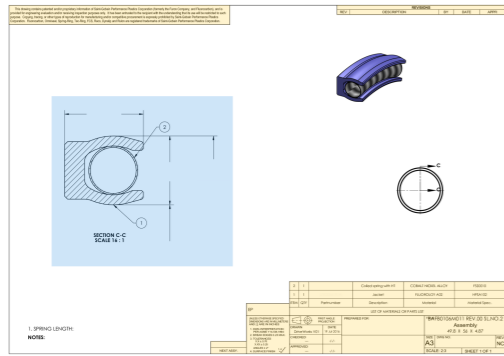


Agent: *It's in Conference Room D.*

Dialogues represented as symbolic programs (e.g. dataflow graphs)

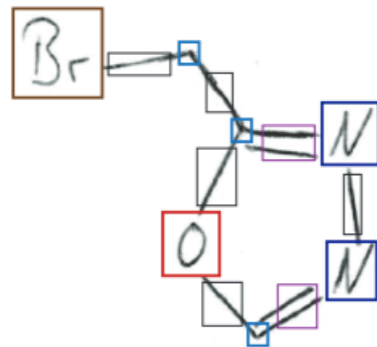
Andreas, Jacob, et al. ACL, 2020

Emerging applications



automated engineering assistant (IAAI 21)
interpret and correct designs and maps

planning, reinforcement learning and shielding (AAAI 24, IJCAI 23 distinguished paper award)



Intelligent OCR for chemical structures (ICLR 23)
and forms

Standard Prompting	Chain-of-Thought Prompting
<p>Model Input</p> <p>Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?</p> <p>A: The answer is 11.</p> <p>Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?</p>	<p>Model Input</p> <p>Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?</p> <p>A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.</p> <p>Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?</p>
<p>Model Output</p> <p>A: The answer is 27. ❌</p>	<p>Model Output</p> <p>A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅</p>

reasoning and mathematical problem solving JAIR 23, IJCAI 2017, EMNLP 21)

Both StarAI and NeSy

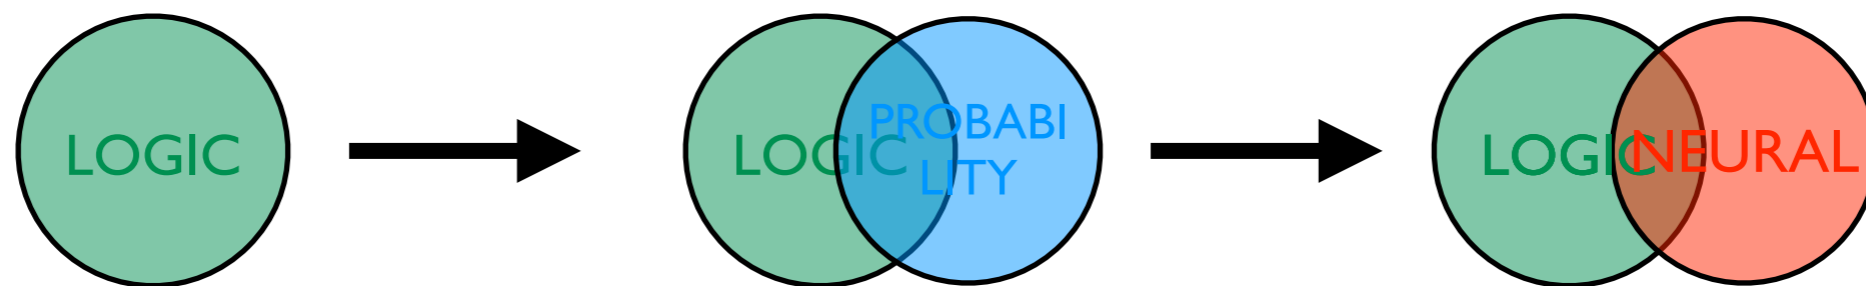
- Structured environments
 - objects, and
 - **relationships** amongst them
- and possibly
 - using **background knowledge**
- cope with **uncertainty and/or perception**
- **learn from data and reason with knowledge**

Power of Logic
Of Programs

The Seven Dimensions

1. Proof vs Model based
2. Directed vs Undirected
3. Type of Logic
4. Symbols vs Subsymbols
5. Parameter vs Structure Learning
6. Semantics
7. Logic vs Probability vs Neural

1. Proof vs Model based



1. Proof vs Model based

LOGIC

1. Proof vs Model based the logic dimension

- Model- vs proof-based
- First order / relational vs propositional
- Grounding
- Differences important for both StarAI and NeSY

Logic Programs

as in the programming language Prolog

Propositional logic program

```
burglary.  
hears_alarm_mary.
```

```
earthquake.  
hears_alarm_john.
```

facts :
burglary = true

```
alarm :- earthquake.
```

```
alarm :- burglary.
```

```
calls_mary :- alarm, hears_alarm_mary.
```

```
calls_john :- alarm, hears_alarm_john.
```

Logic Programs

as in the programming language Prolog

Propositional logic program

burglary.
hears_alarm_mary.

earthquake.
hears_alarm_john.

alarm :- earthquake.

alarm :- burglary. **rule: calls_mary = true IF alarm = true AND hears_alarm_mary = true**

calls_mary :- alarm, hears_alarm_mary.

calls_john :- alarm, hears_alarm_john.

Logic Programs

as in the programming language Prolog

Propositional logic program

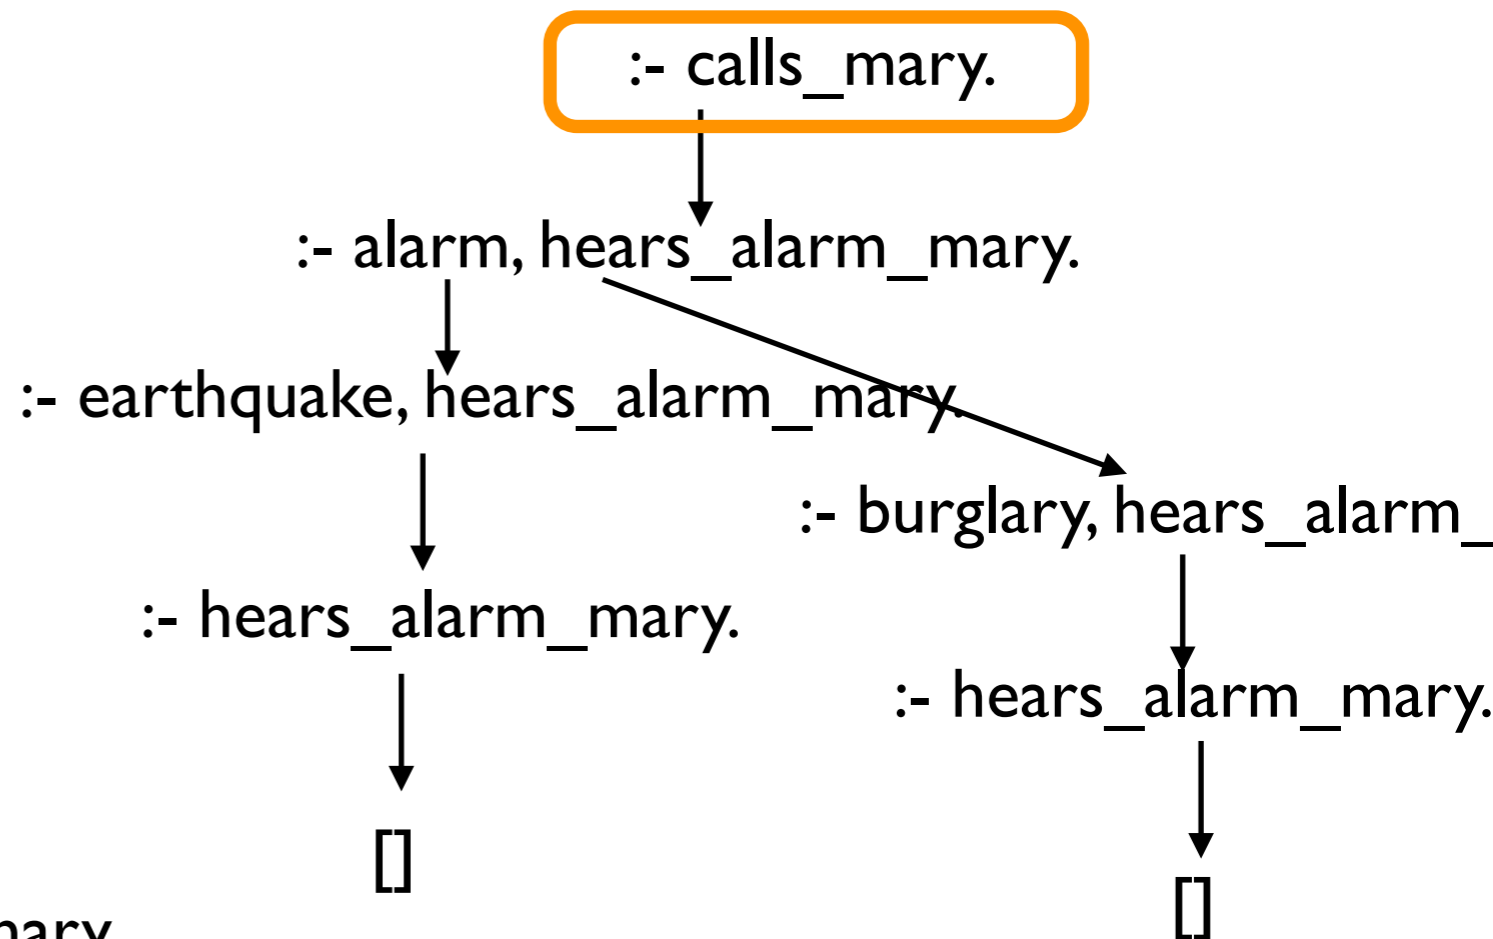
burglary.
hears_alarm_mary.

earthquake.
hears_alarm_john.

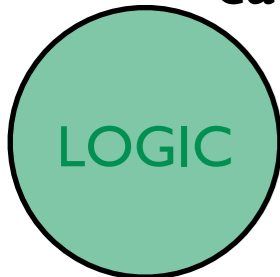
alarm :- earthquake.
alarm :- burglary.

calls_mary :- alarm, hears_alarm_mary.
calls_john :- alarm, hears_alarm_john.

Two proofs (by refutation)



A proof-theoretic view
backward chaining



Logic as constraints

as in SAT solvers

Propositional logic

Model / Possible World

IF

AND

$\text{calls}(\text{mary}) \leftarrow \text{hears_alarm}(\text{mary}) \wedge \text{alarm}$

$\text{calls}(\text{john}) \leftarrow \text{hears_alarm}(\text{john}) \wedge \text{alarm}$

OR

$\text{alarm} \leftarrow \text{earthquake} \vee \text{burglary}$

{ burglary,

hears_alarm(john),

alarm,

calls(john)}

**the facts that are true
in this model / possible world**

SAT: Find a model / possible world that satisfies all the constraints

SAT SOLVERS

A model-theoretic view

LOGIC

Relational/First Order Logic

Introduce Variables and Domains

The meaning of this is always the **GROUNDED** theory

allows to exploit symmetries / templates ...

burglary.
hears_alarm(**mary**).

earthquake.
hears_alarm(**john**).

alarm :- earthquake.

alarm :- burglary.
calls(**X**) :- alarm, hears_alarm(**X**).

Variable X

Domain = {mary, john}

burglary.
hears_alarm(mary).

earthquake.
hears_alarm(john).

alarm :- earthquake.

alarm :- burglary.
calls(mary) :- alarm, hears_alarm(mary).
calls(john) :- alarm, hears_alarm(john).

Grounded Theory

BOTH for model and proof-based approach

Logical Theory

GROUNDING OUT

```
stress (ann) .  
influences (ann , bob) .  
influences (bob , carl) .
```

```
smokes (ann) :- stress (ann) .  
smokes (bob) :- stress (bob) .  
smokes (carl) :- stress (carl) .
```

```
smokes (ann) :- influences (ann , ann) , smokes (ann) .  
smokes (ann) :- influences (bob , ann) , smokes (bob) .  
smokes (ann) :- influences (carl , ann) , smokes (carl) .
```

```
smokes (bob) :- influences (ann , bob) , smokes (ann) .  
smokes (bob) :- influences (bob , bob) , smokes (bob) .  
smokes (bob) :- influences (carl , bob) , smokes (carl) .
```

```
smokes (carl) :- influences (ann , carl) , smokes (ann) .  
smokes (carl) :- influences (bob , carl) , smokes (bob) .  
smokes (carl) :- influences (carl , carl) , smokes (carl) .
```

```
stress (ann) .  
influences (ann , bob) .  
influences (bob , carl) .
```

```
smokes (X) :- stress (X) .  
smokes (X) :-  
    influences (Y , X) ,  
    smokes (Y) .
```

**IF INTERESTED ONLY IN
CERTAIN QUERIES,
CLEVER TECHNIQUES EXIST
TO AVOID GROUNDING OUT
COMPLETELY**

Logical Reasoning: Model Theoretic

FINDING A MODEL

```
stress (ann) .  
influences (ann,bob) .  
influences (bob,carl) .
```

```
smokes (ann) :- stress (ann) .  
-> infer smokes (ann)
```

```
smokes (bob) :- influences (ann,bob) , smokes (ann)  
-> infer smokes (bob)
```

```
smokes (carl) :- influences (bob,carl) , smokes (bob) .  
-> infer smokes (carl) .
```

```
stress (ann) .  
influences (ann,bob) .  
influences (bob,carl) .
```

```
smokes (X) :- stress (X) .  
smokes (X) :-  
    influences (Y,X) ,  
    smokes (Y) .
```

FINDING A MODEL

here — the least Herbrand model as in Prolog using the Tp Operator (forward reasoning)

Logical Reasoning: Model Theoretic

Clark's completion AND call a SAT Solver

```
stress (ann) .  
influences (ann,bob) .  
influences (bob,carl) .
```

```
smokes (ann) <-> stress (ann)  
    v influences (ann,ann) , smokes (ann) here  
    v influences (bob,ann) , smokes (bob)  
    v influences (carl,ann) , smokes (carl)
```

```
smokes (bob) <-> stress (bob)  
    v influences (ann,bob) , smokes (ann)  
    v influences (bob,bob) , smokes (bob)  
    v influences (carl,bob) , smokes (carl)
```

```
smokes (carl) <-> stress (carl)  
    v influences (ann,carl) , smokes (ann)  
    v influences (bob,carl) , smokes (bob)  
    v influences (carl,carl) , smokes (carl)
```

```
stress (ann) .  
influences (ann,bob) .  
influences (bob,carl) .
```

```
smokes (X) :- stress (X) .  
smokes (X) :-  
    influences (Y,X) ,  
    smokes (Y) .
```

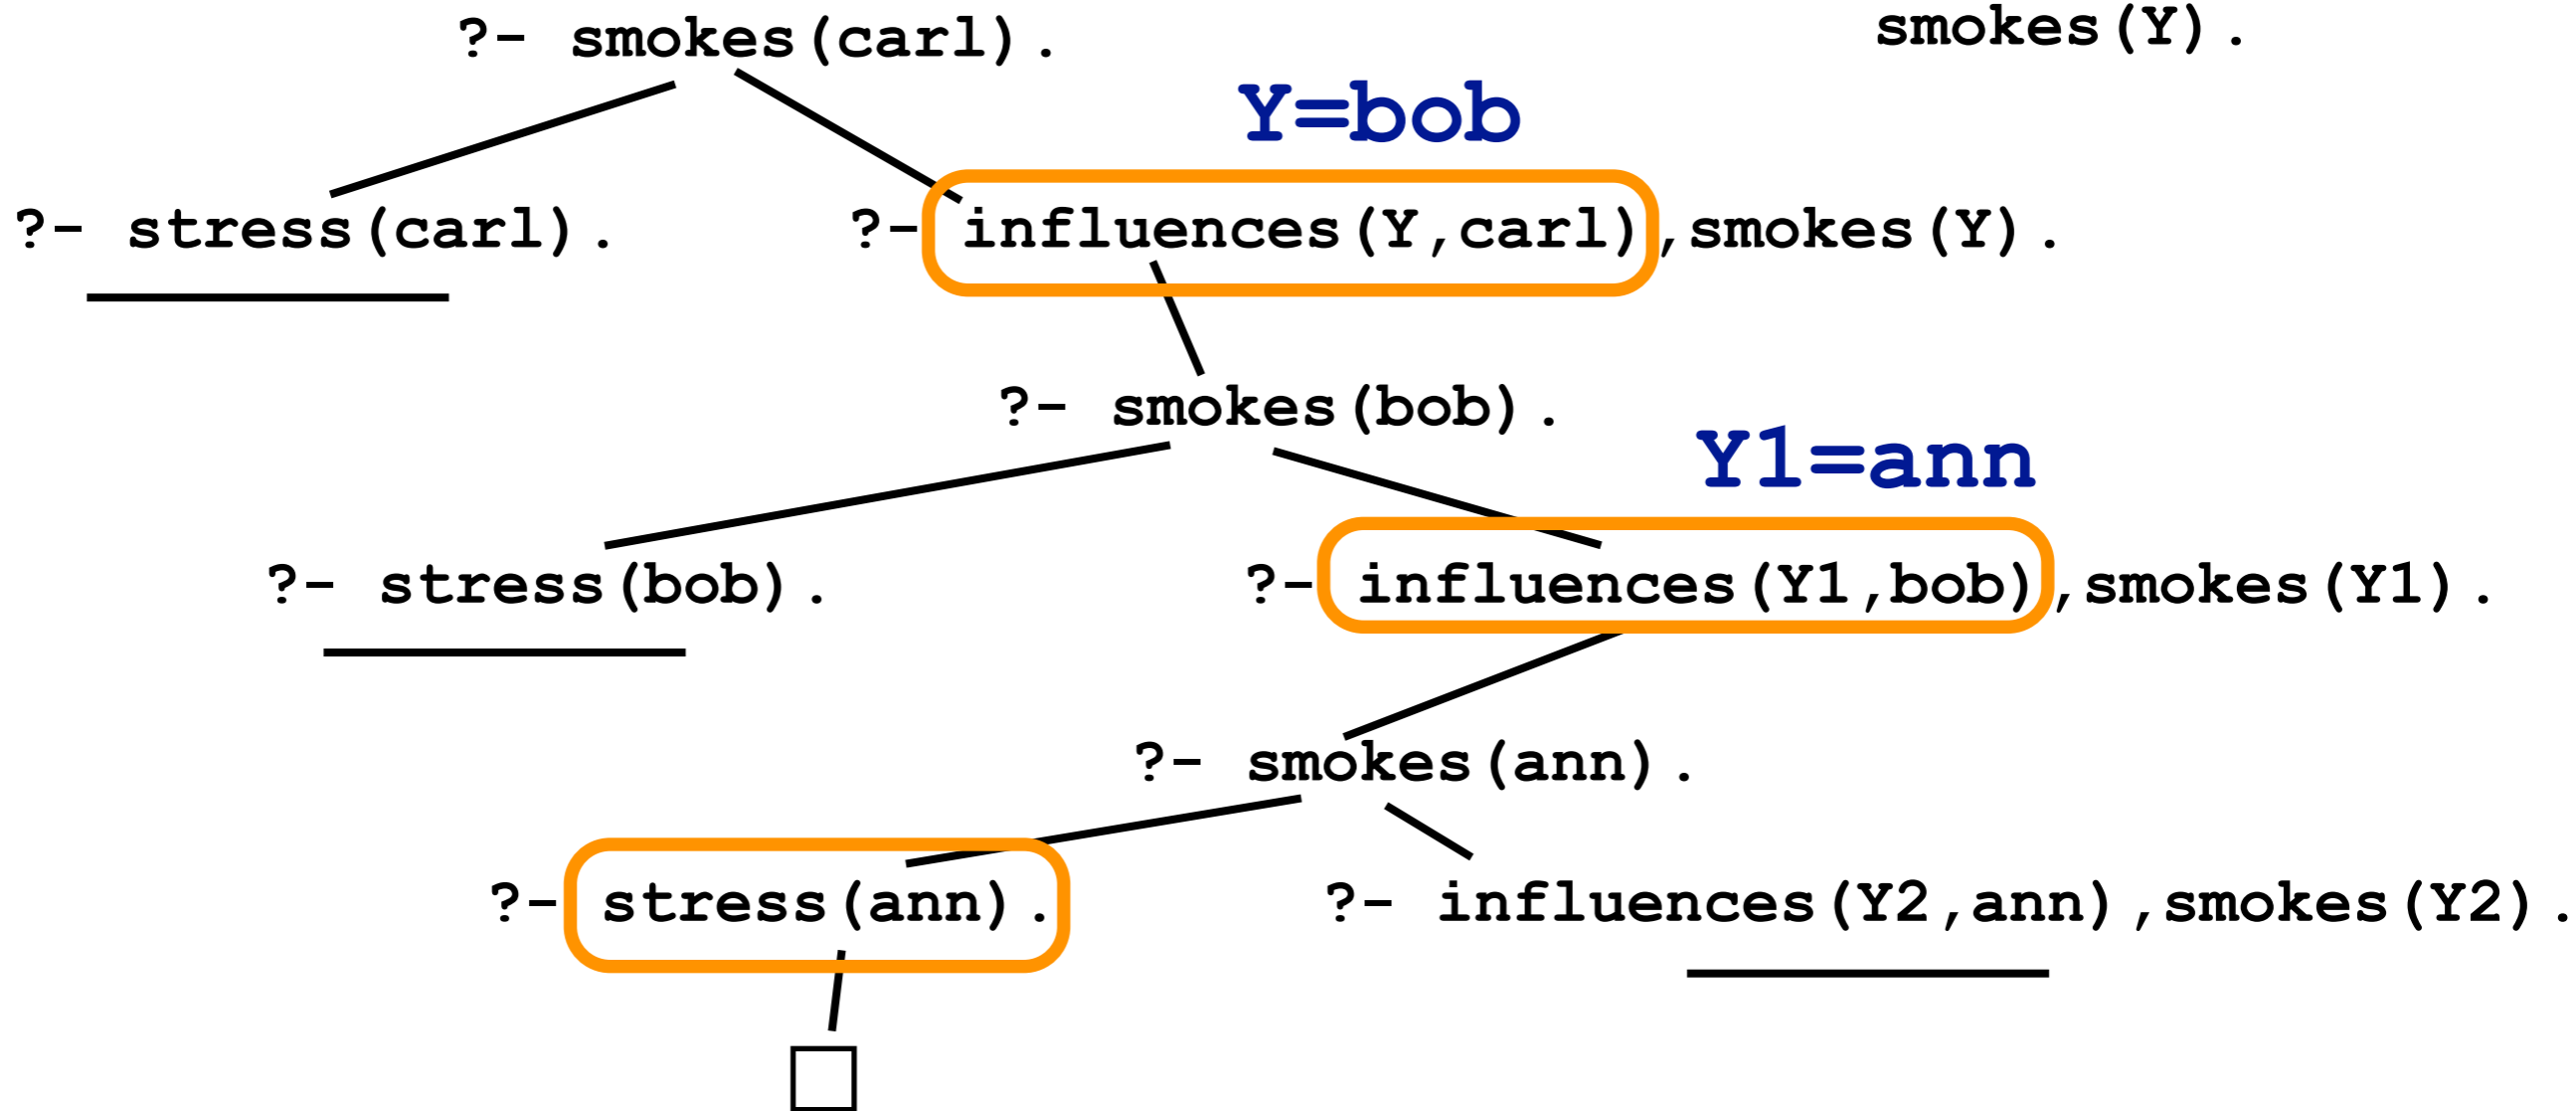
Clark's completion's as a
grounding is incorrect
for Prolog when there are cycles

but it is too hard to explain why

Logical Reasoning Proofs

```
stress(ann) .
influences(ann,bob) .
influences(bob,carl) .
```

```
smokes(X) :- stress(X) .
smokes(X) :-
    influences(Y,X) ,
    smokes(Y) .
```

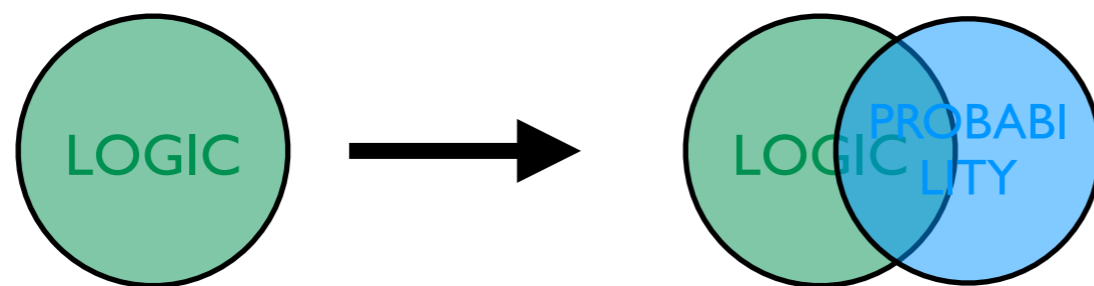


facts used in successful derivation:
`influences(bob,carl) & influences(ann,bob) & stress(ann)`

1. Proof vs Model based the logic dimension

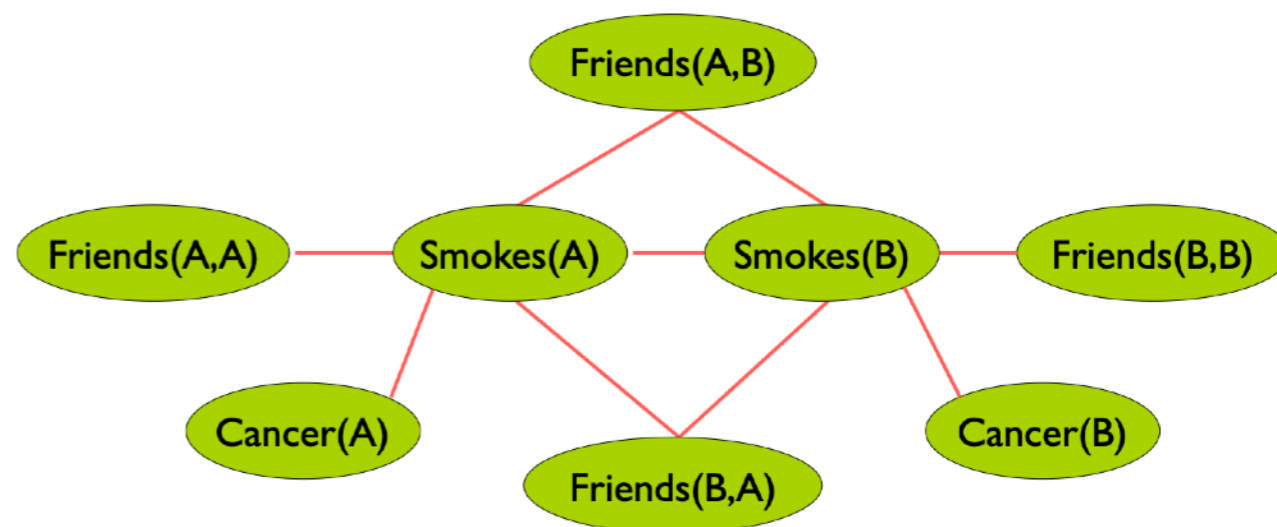
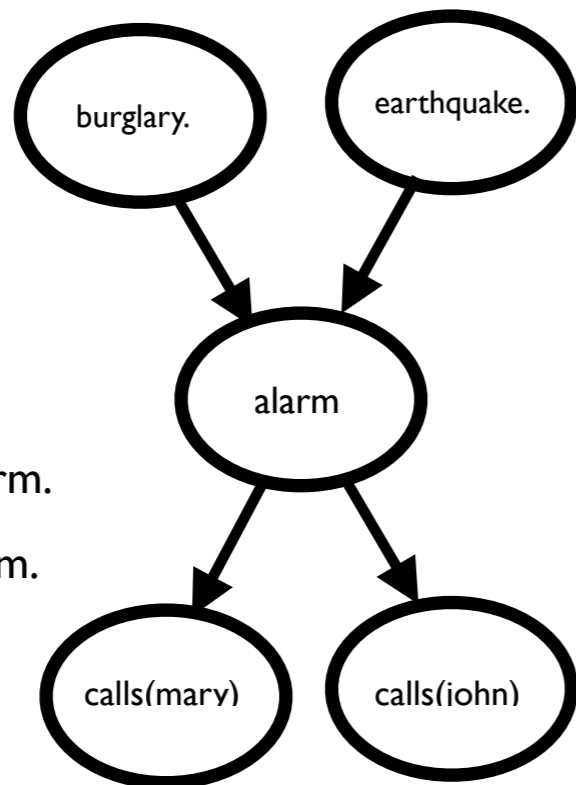
- Model- vs proof-based
- First order / relational vs propositional
- Grounding
- Differences important for both StarAI and NeSY

1. Proof vs Model based
2. Directed vs Undirected



2. Directed vs Undirected the PGM / StarAI dimension

0.1 :: burglary.
 0.05 :: earthquake.
 alarm :- earthquake.
 alarm :- burglary.
 0.7::calls(mary) :- alarm.
 0.6::calls(john) :- alarm.



1.5 $\forall x \text{ Smokes}(x) \Rightarrow \text{Cancer}(x)$

1.1 $\forall x, y \text{ Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$

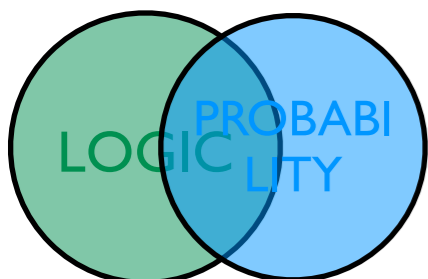
Probabilistic Logic Programs
ProbLog

directed
Bayesian Net

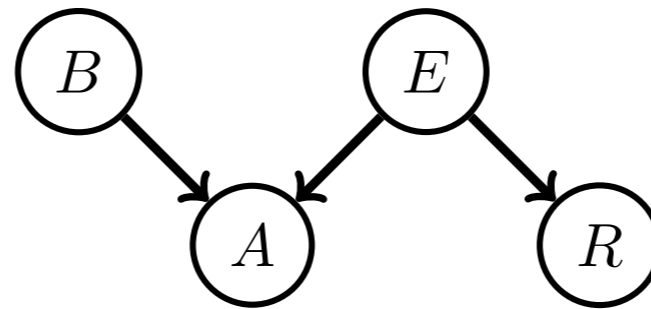
Markov Logic

undirected
Markov Net
model theoretic

key representatives



Bayesian Net



$$\mathbf{P}(A|B, E)$$

alarm (= true)	Burglar	Earthquake
0.9999	true	true
0.99	true	false
0.99	false	true
0.0001	false	true

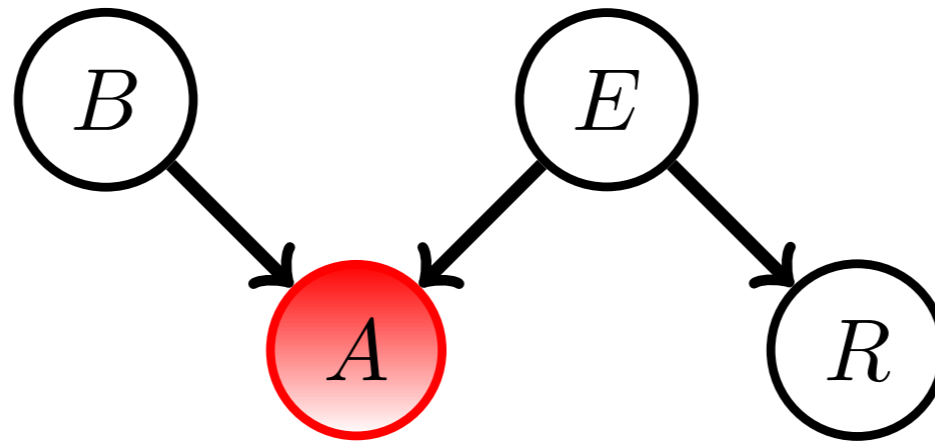
$$\mathbf{P}(R|E)$$

radio	Earthquake
1	true
0	false

The remaining tables are $P(b) = 0.01$ and $P(e) = 0.000001$. The tables and graphical structure fully specify the joint distribution $\mathbf{P}(A, R, E, B)$.

Queries

Initial evidence: The alarm is sounding



$$\begin{aligned} P(b|a) &= \frac{P(b, a)}{P(a)} = \frac{\sum_{e, r} P(b, e, a, e)}{\sum_{b, e, r} P(b, e, a, r)} \\ &= \frac{\sum_{e, r} P(r|b, e)P(b)P(e)P(r|e)}{\sum_{b, e, r} P(a|b, e)P(b)P(e)P(r|e)} \approx 0.99 \end{aligned}$$

Logic Programs

as in the programming language Prolog

Propositional logic program

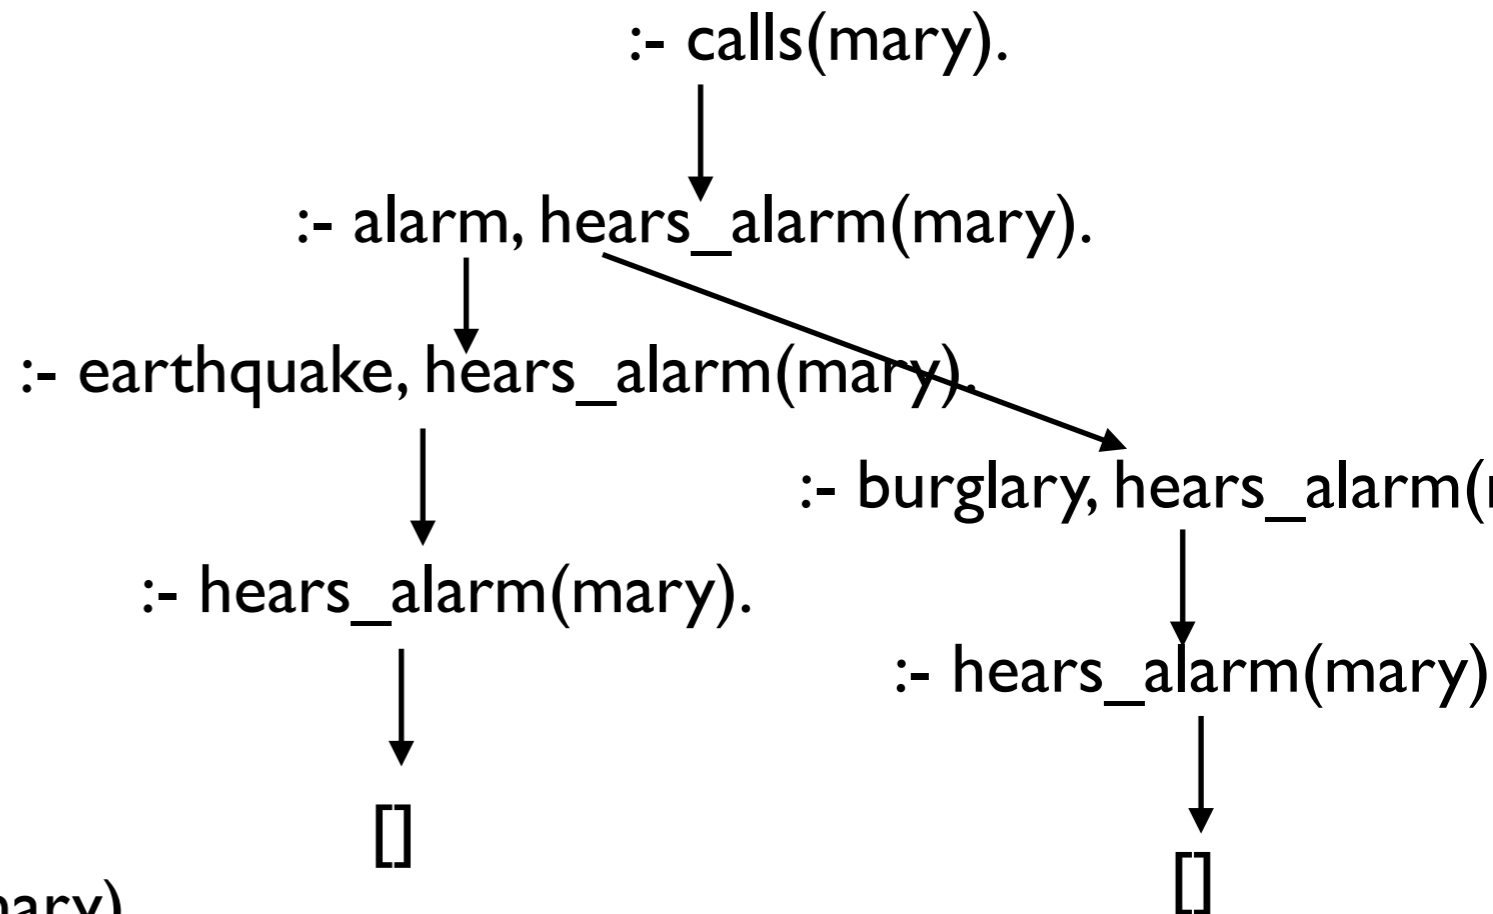
burglary.
hears_alarm(mary).

earthquake.
hears_alarm(john).

alarm :- earthquake.
alarm :- burglary.

calls(mary) :- alarm, hears_alarm(mary).
calls(john) :- alarm, hears_alarm(john).

Two proofs (by refutation)



A proof-theoretic view

Probabilistic Logic Programs

as in the probabilistic programming language ProbLog

Propositional logic program

0.1 :: burglary.

0.3 :: hears_alarm(mary).

Probabilistic facts

0.05 :: earthquake.

0.6 :: hears_alarm(john).

alarm :- earthquake.

alarm :- burglary.

calls(mary) :- alarm, hears_alarm(mary).

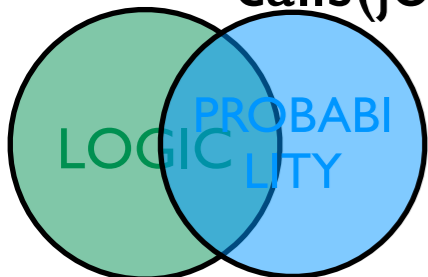
calls(john) :- alarm, hears_alarm(john).

Key Idea (Sato & Poole)
the distribution semantics:

**unify the basic concepts in logic
and probability:**

**random variable ~ propositional
variable**

**an interface between logic and
probability**



Probabilistic Logic Programs

as in the probabilistic programming language ProbLog

Propositional logic program

0.1 :: burglary.

0.3 :: hears_alarm(mary).

0.05 :: earthquake.

0.6 :: hears_alarm(john).

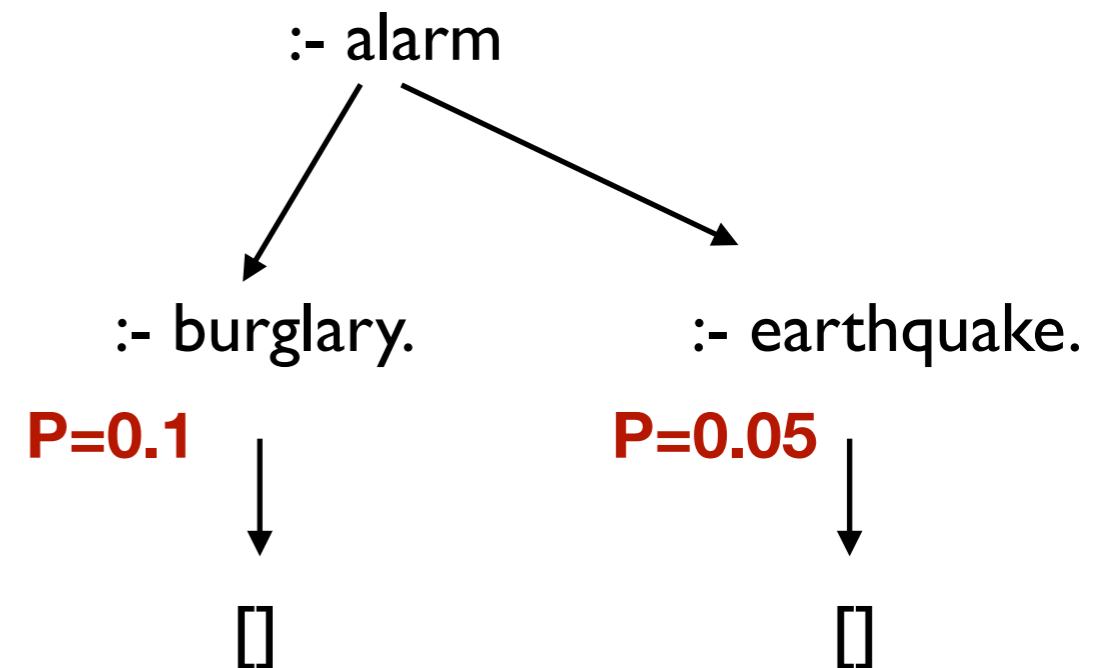
alarm :- earthquake.

alarm :- burglary.

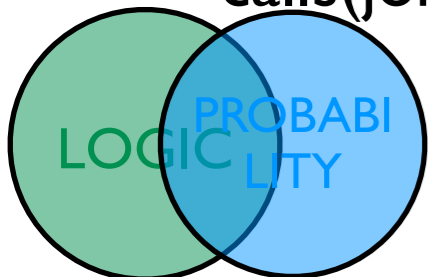
calls(mary) :- alarm, hears_alarm(mary).

calls(john) :- alarm, hears_alarm(john).

Two proofs (by refutation)



Probability of one proof : $\prod_{f: fact \in Proof} P_f$



Probabilistic Logic Programs

as in the probabilistic programming language ProbLog

Propositional logic program

0.1 :: burglary.
0.3 :: hears_alarm(mary).

0.05 :: earthquake.
0.6 :: hears_alarm(john).

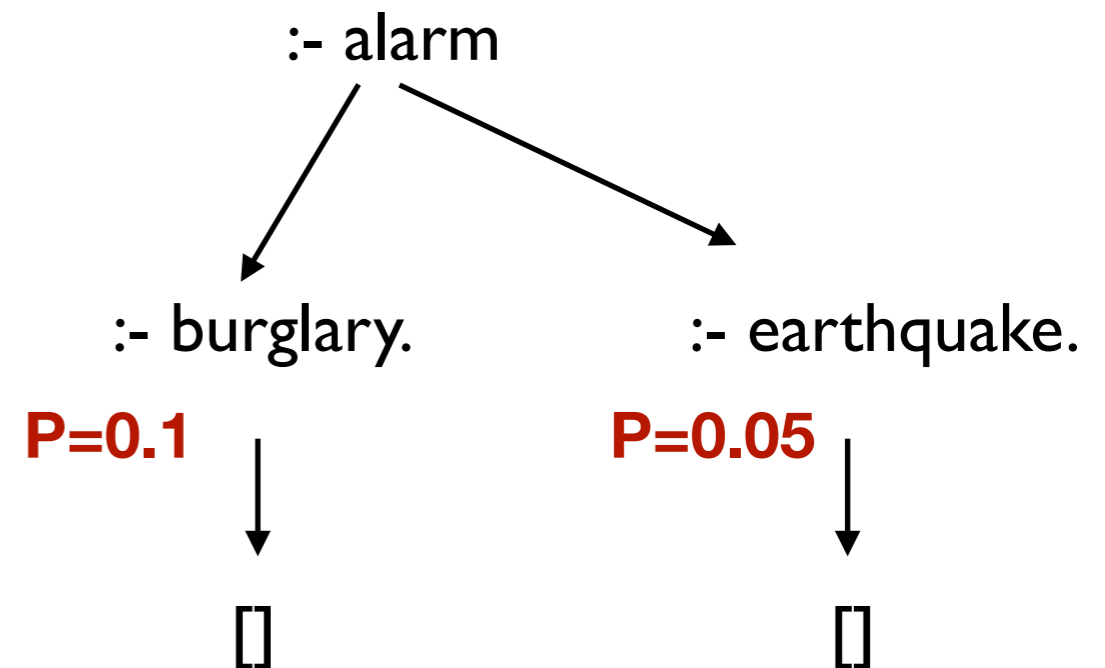
alarm :- earthquake.

alarm :- burglary.

calls(mary) :- alarm, hears_alarm(mary).

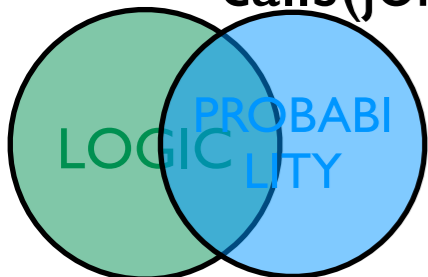
calls(john) :- alarm, hears_alarm(john).

Disjoint sum problem



Probability of one proof : $\prod_{f: fact \in Proof} P_f$

$$\begin{aligned} P(\text{alarm}) &= P(\text{burg OR earth}) \\ &= P(\text{burg}) + P(\text{earth}) - P(\text{burg AND earth}) \\ &\neq P(\text{burg}) + P(\text{earth}) \end{aligned}$$



Probabilistic Logic Program Semantics

earthquake.

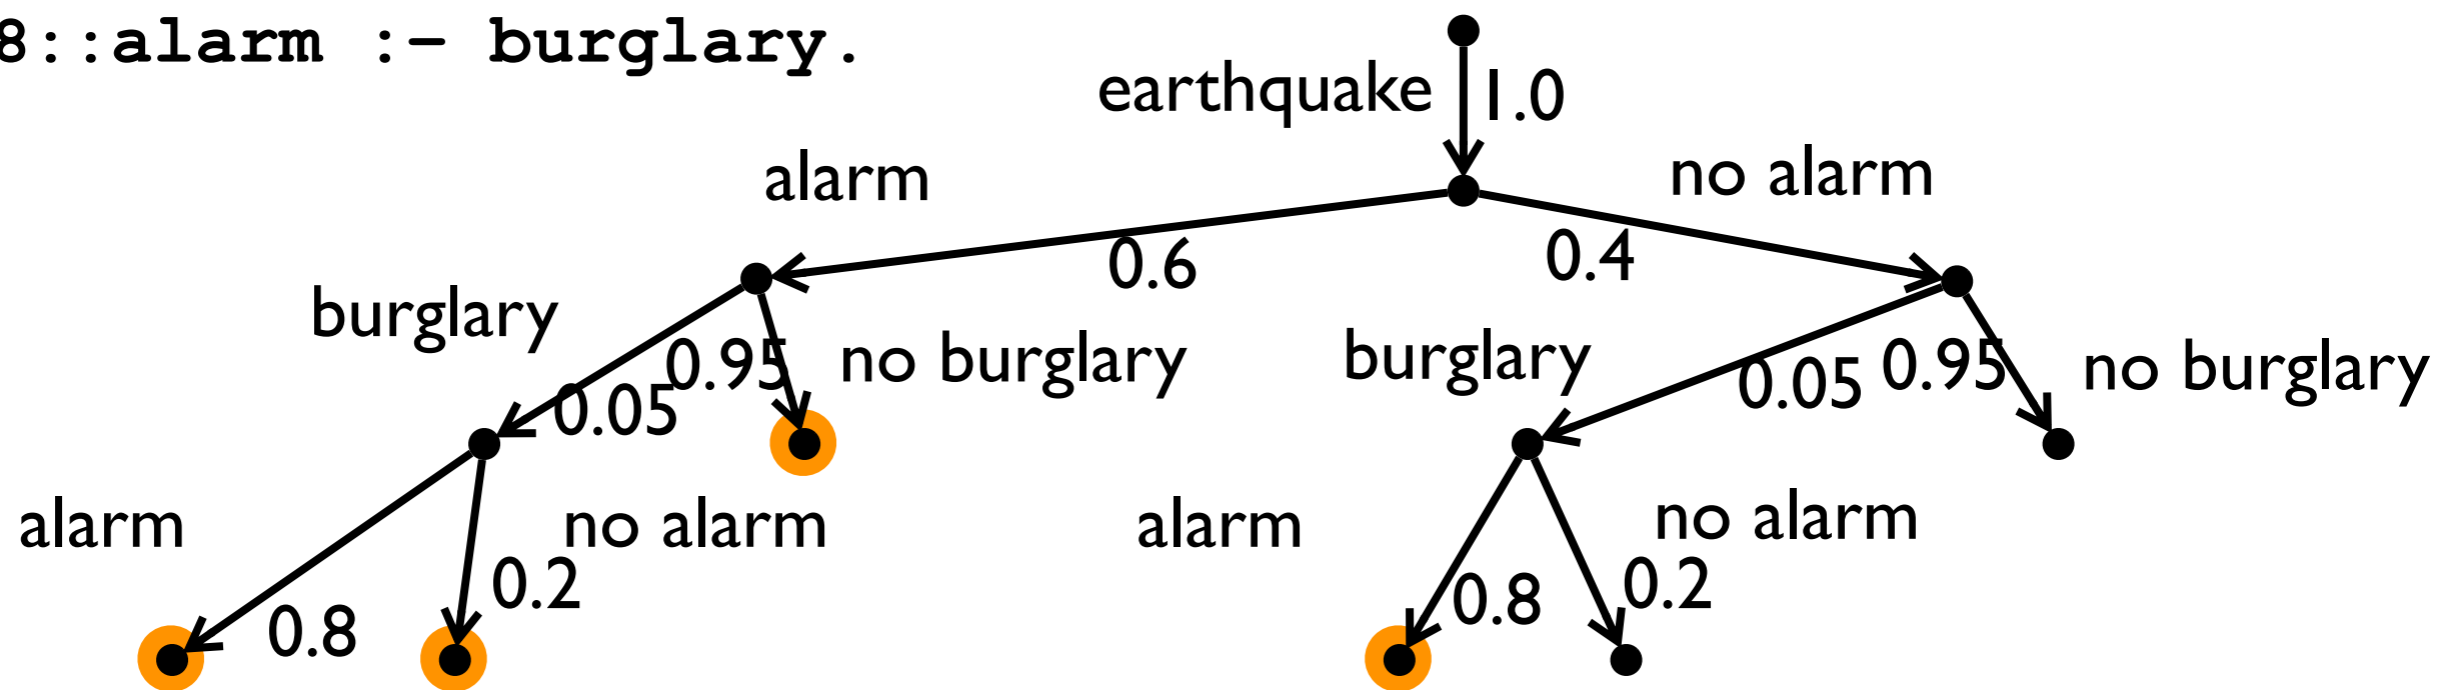
[Vennekens et al, ICLP 04]

0.05::burglary.

probabilistic causal laws

0.6::alarm :- earthquake.

0.8::alarm :- burglary.



$$P(\text{alarm}) = 0.6 \times 0.05 \times 0.8 + 0.6 \times 0.05 \times 0.2 + 0.6 \times 0.95 + 0.4 \times 0.05 \times 0.8$$

Probabilistic Logic Program Semantics

Propositional logic program

0.1 :: burglary.

0.05 :: earthquake.

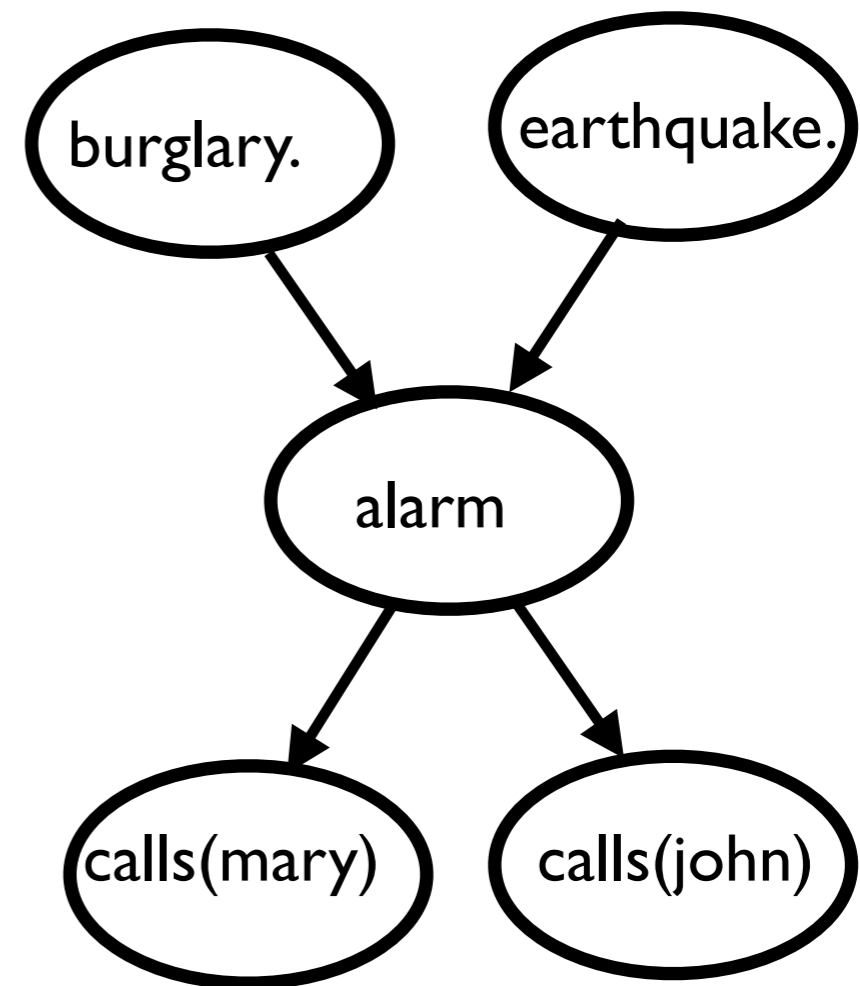
alarm :- earthquake.

alarm :- burglary.

0.7::calls(mary) :- alarm.

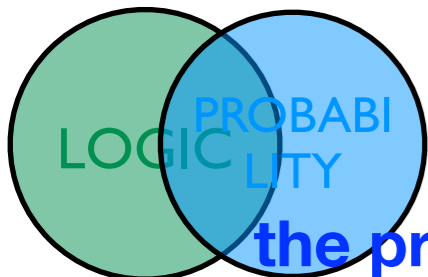
0.6::calls(john) :- alarm.

Bayesian Network

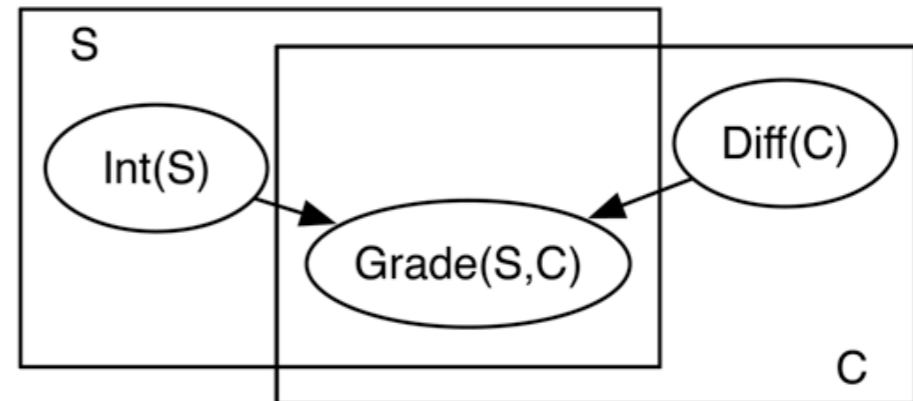


**Bayesian net encoded as Probabilistic Logic Program
PLPs correspond to directed graphical models**

**ProbLog has both (directed) probabilistic graphic models,
the programming language Prolog (and probabilistic databases) as special case**



Flexible and Compact Relational Model for Predicting Grades



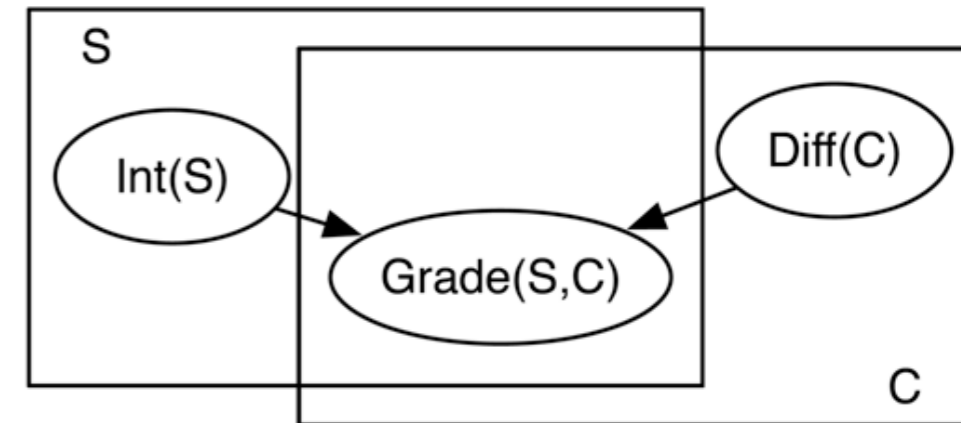
“Program” Abstraction:

- S, C **logical variable** representing students, courses
- the set of individuals of a type is called a **population**
- Int(S), Grade(S, C), D(C) are **parametrized random variables**

Grounding:

- for every student s , there is a random variable $\text{Int}(s)$
- for every course c , there is a random variable $\text{Di}(c)$
- for every s, c pair there is a random variable $\text{Grade}(s,c)$
- all instances share the same structure and parameters

ProbLog by example: Grading



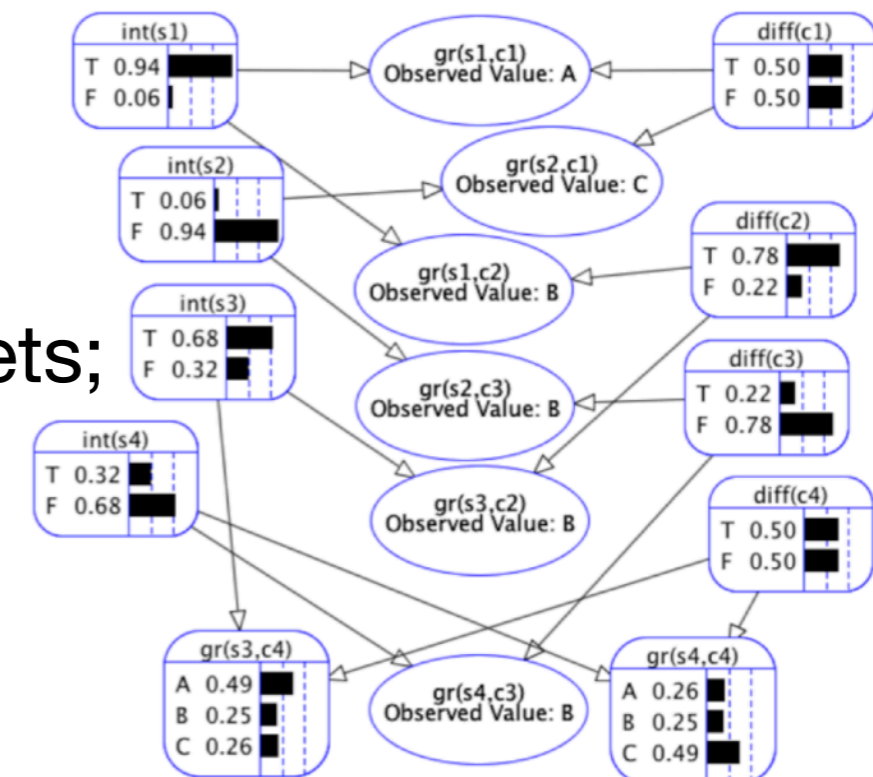
Shows relational structure

- grounded model: replace variables by constants

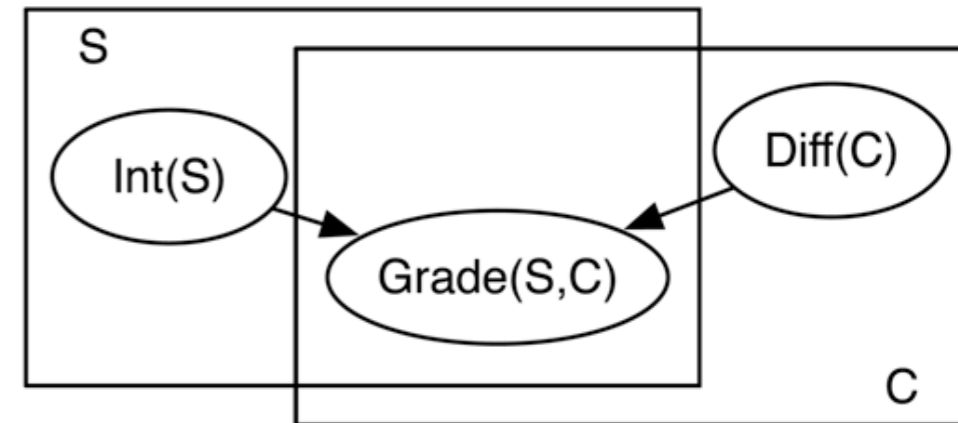
Works for any number of students / classes (for 1000 students and 100 classes, you get 101100 random variables); still only few parameters

With SRL / PP

- build and learn compact models,
- from one set of individuals - > other sets;
- reason also about exchangeability,
- build even more complex models,
- incorporate background knowledge



ProbLog by example: Grading



Shows relational structure

- grounded model: replace variables by constants

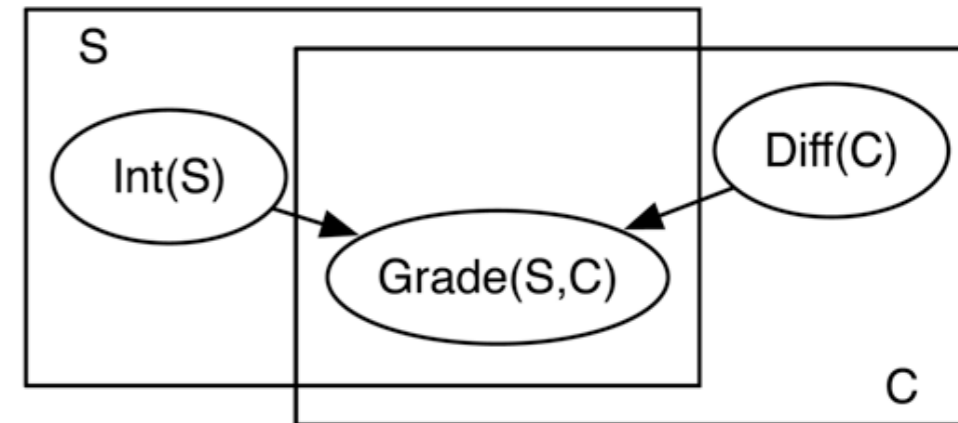
Works for any number of students / classes (for 1000 students and 100 classes, you get 101100 random variables); still only few parameters

With SRL / PP

- build and learn compact models,
- from one set of individuals - > other sets;
- reason also about exchangeability,
- build even more complex models,
- incorporate background knowledge

<i>Student</i>	<i>Course</i>	<i>Grade</i>
s_1	c_1	A
s_2	c_1	C
s_1	c_2	B
s_2	c_3	B
s_3	c_2	B
s_4	c_3	B
s_3	c_4	$?$
s_4	c_4	$?$

ProbLog by example: Grading



```
0.4 :: int(S) :- student(S).  
0.5 :: diff(C) :- course(C).
```

```
student(john). student(anna). student(bob).  
course(ai). course(ml). course(cs).
```

```
gr(S,C,a) :- int(S), not diff(C).
```

```
0.3 :: gr(S,C,a); 0.5 :: gr(S,C,b); 0.2 :: gr(S,C,c) :-  
int(S), diff(C).
```

```
0.1 :: gr(S,C,b); 0.2 :: gr(S,C,c); 0.2 :: gr(S,C,f) :-  
student(S), course(C),  
not int(S), not diff(C).
```

```
0.3 :: gr(S,C,c); 0.2 :: gr(S,C,f) :-  
not int(S), diff(C).
```


ProbLog by example: Grading

```
unsatisfactory(S) :- student(S), grade(S,C,f).
```

```
excellent(S) :- student(S), not(grade(S,C1,G),below(G,a)),  
                grade(S,C2,a).
```

```
0.4 :: int(S) :- student(S).
```

```
0.5 :: diff(C) :- course(C).
```

```
student(john). student(anna). student(bob).  
course(ai).    course(ml).    course(cs).
```

```
gr(S,C,a) :- int(S), not diff(C).
```

```
0.3 :: gr(S,C,a); 0.5 :: gr(S,C,b); 0.2 :: gr(S,C,c) :-  
                int(S), diff(C).
```

```
0.1 :: gr(S,C,b); 0.2 :: gr(S,C,c); 0.2 :: gr(S,C,f) :-  
                student(S), course(C),  
                not int(S), not diff(C).
```

```
0.3 :: gr(S,C,c); 0.2 :: gr(S,C,f) :-  
                not int(S), diff(C).
```

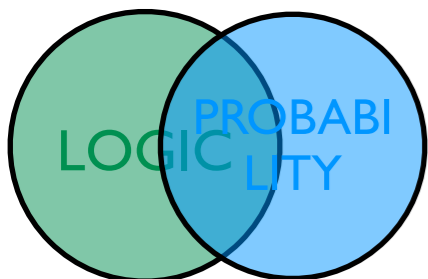
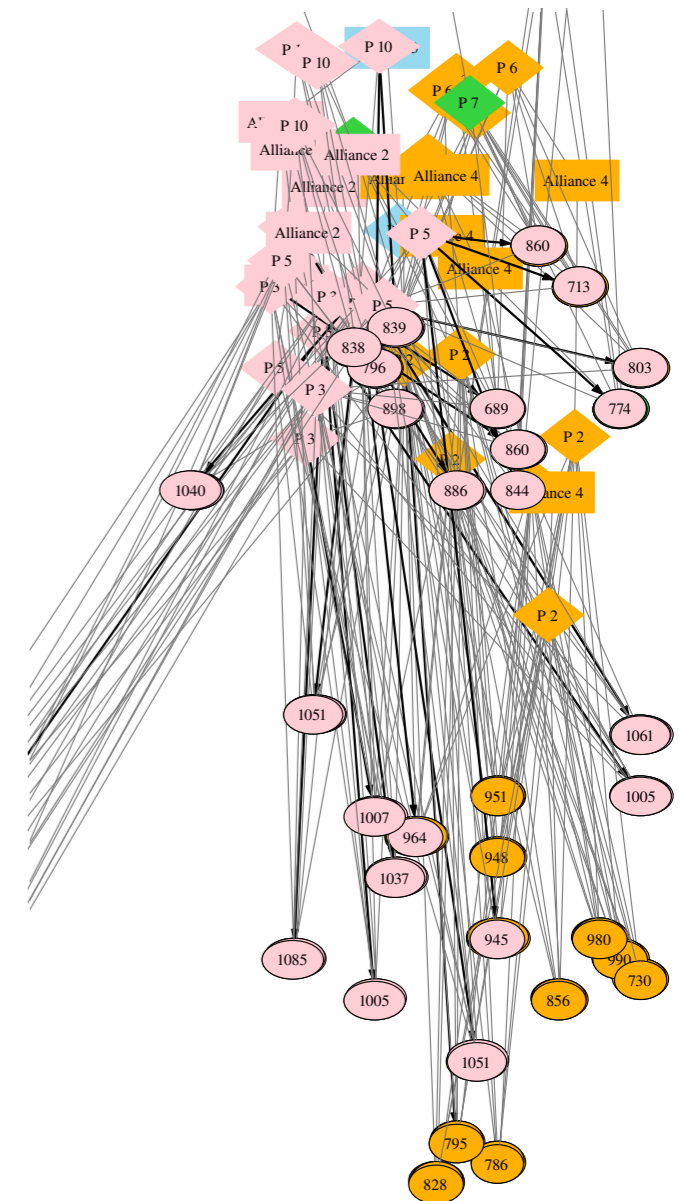
Dynamic networks



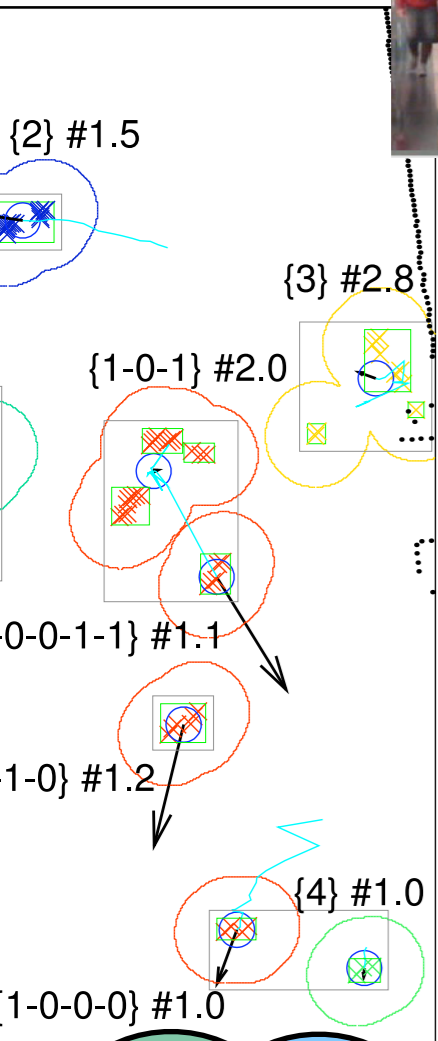
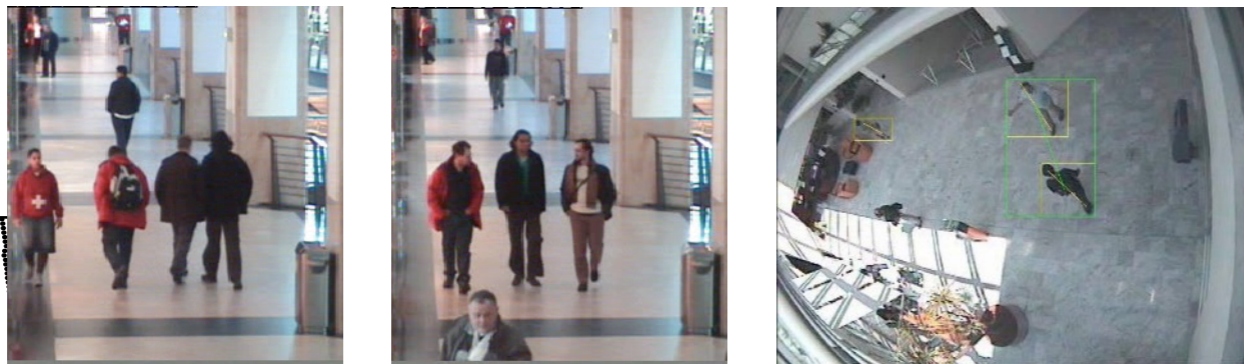
Travian: A massively multiplayer real-time strategy game

Can we build a model of this world ?

Can we use it for playing better ?

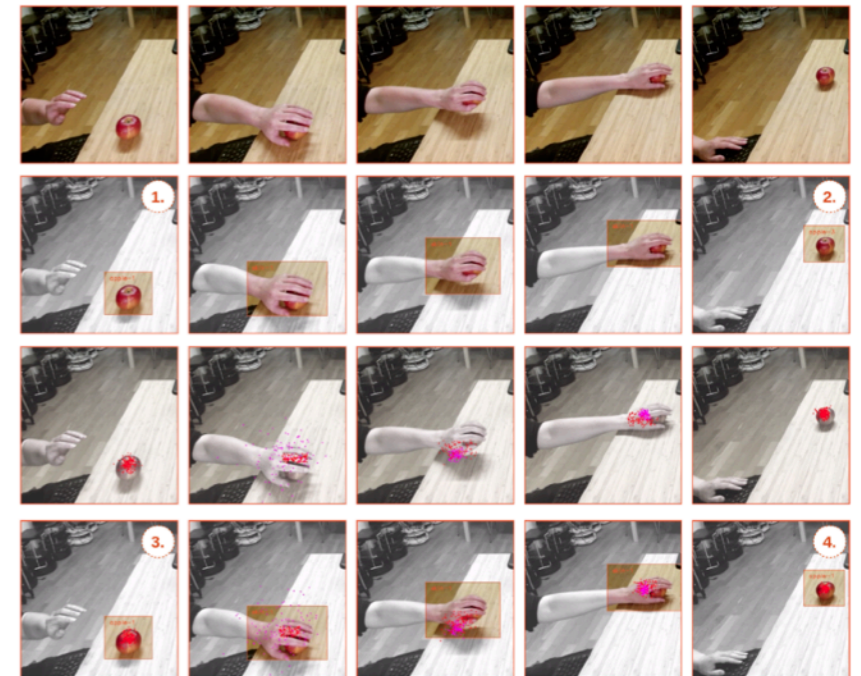


Activity analysis and tracking video analysis

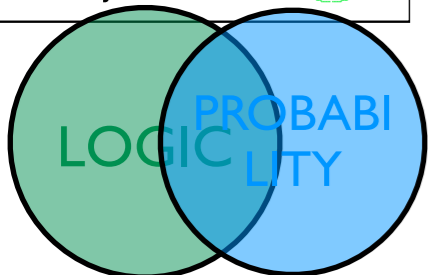


- Track people or objects over time? Even if temporarily hidden?
- Recognize activities?
- Infer object properties?

[Skarlatidis et al, TPLP 14;
Nitti et al, IROS 13, ICRA 14,
MLJ 16]



[Persson et al, IEEE Trans on
Cogn. & Dev. Sys. 19;
IJCAI 20]



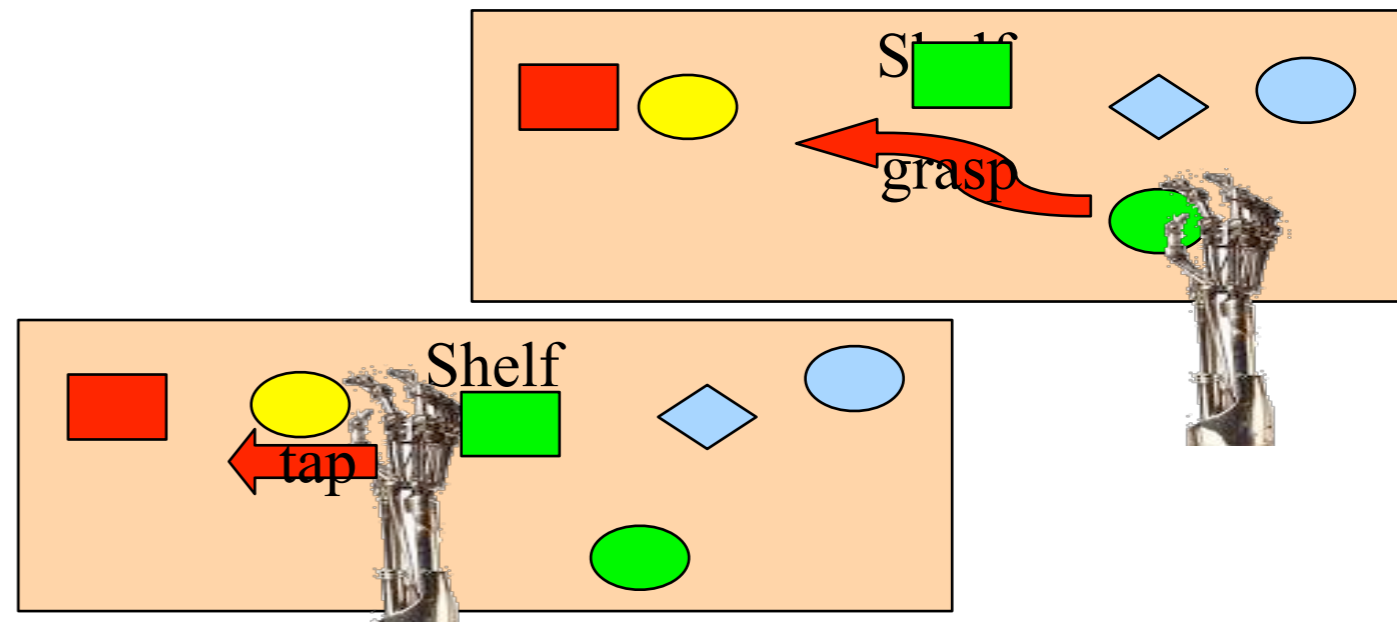
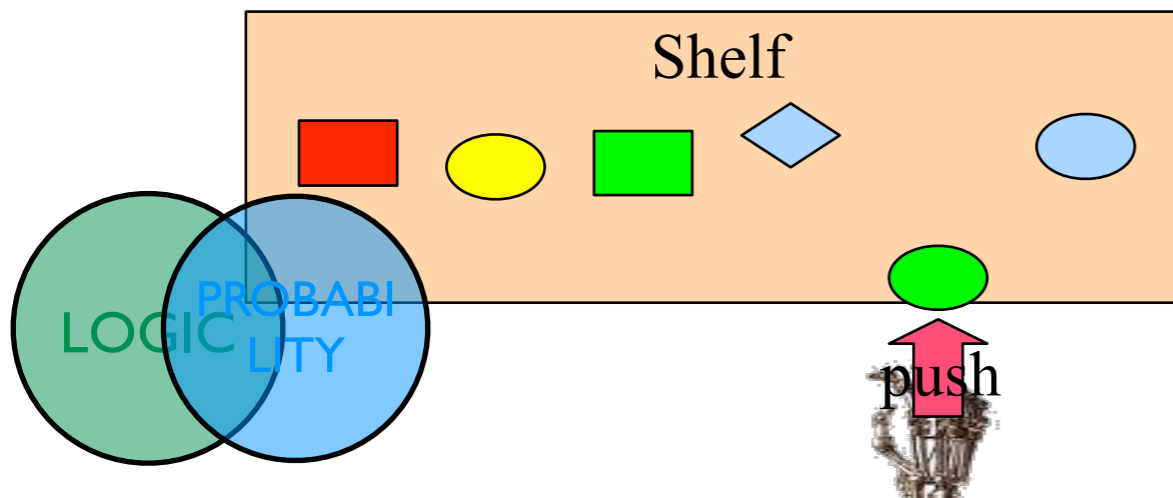
Learning relational affordances



Learning relational affordances between two objects (learnt by experience)

1), and similar to probabilistic Strips (with continuous distributions)

Moldovan et al. ICRA 12, 13, 14; Auton. Robots 18



Distributional Clauses (DC)

- Discrete- and continuous-valued random variables

random variable with Gaussian distribution

```
length(Obj) ~ gaussian(6.0,0.45) :- type(Obj,glass).
```

```
stackable(OBot,OTop) :-
```

```
    ≈length(OBot) ≥ ≈length(OTop),
```

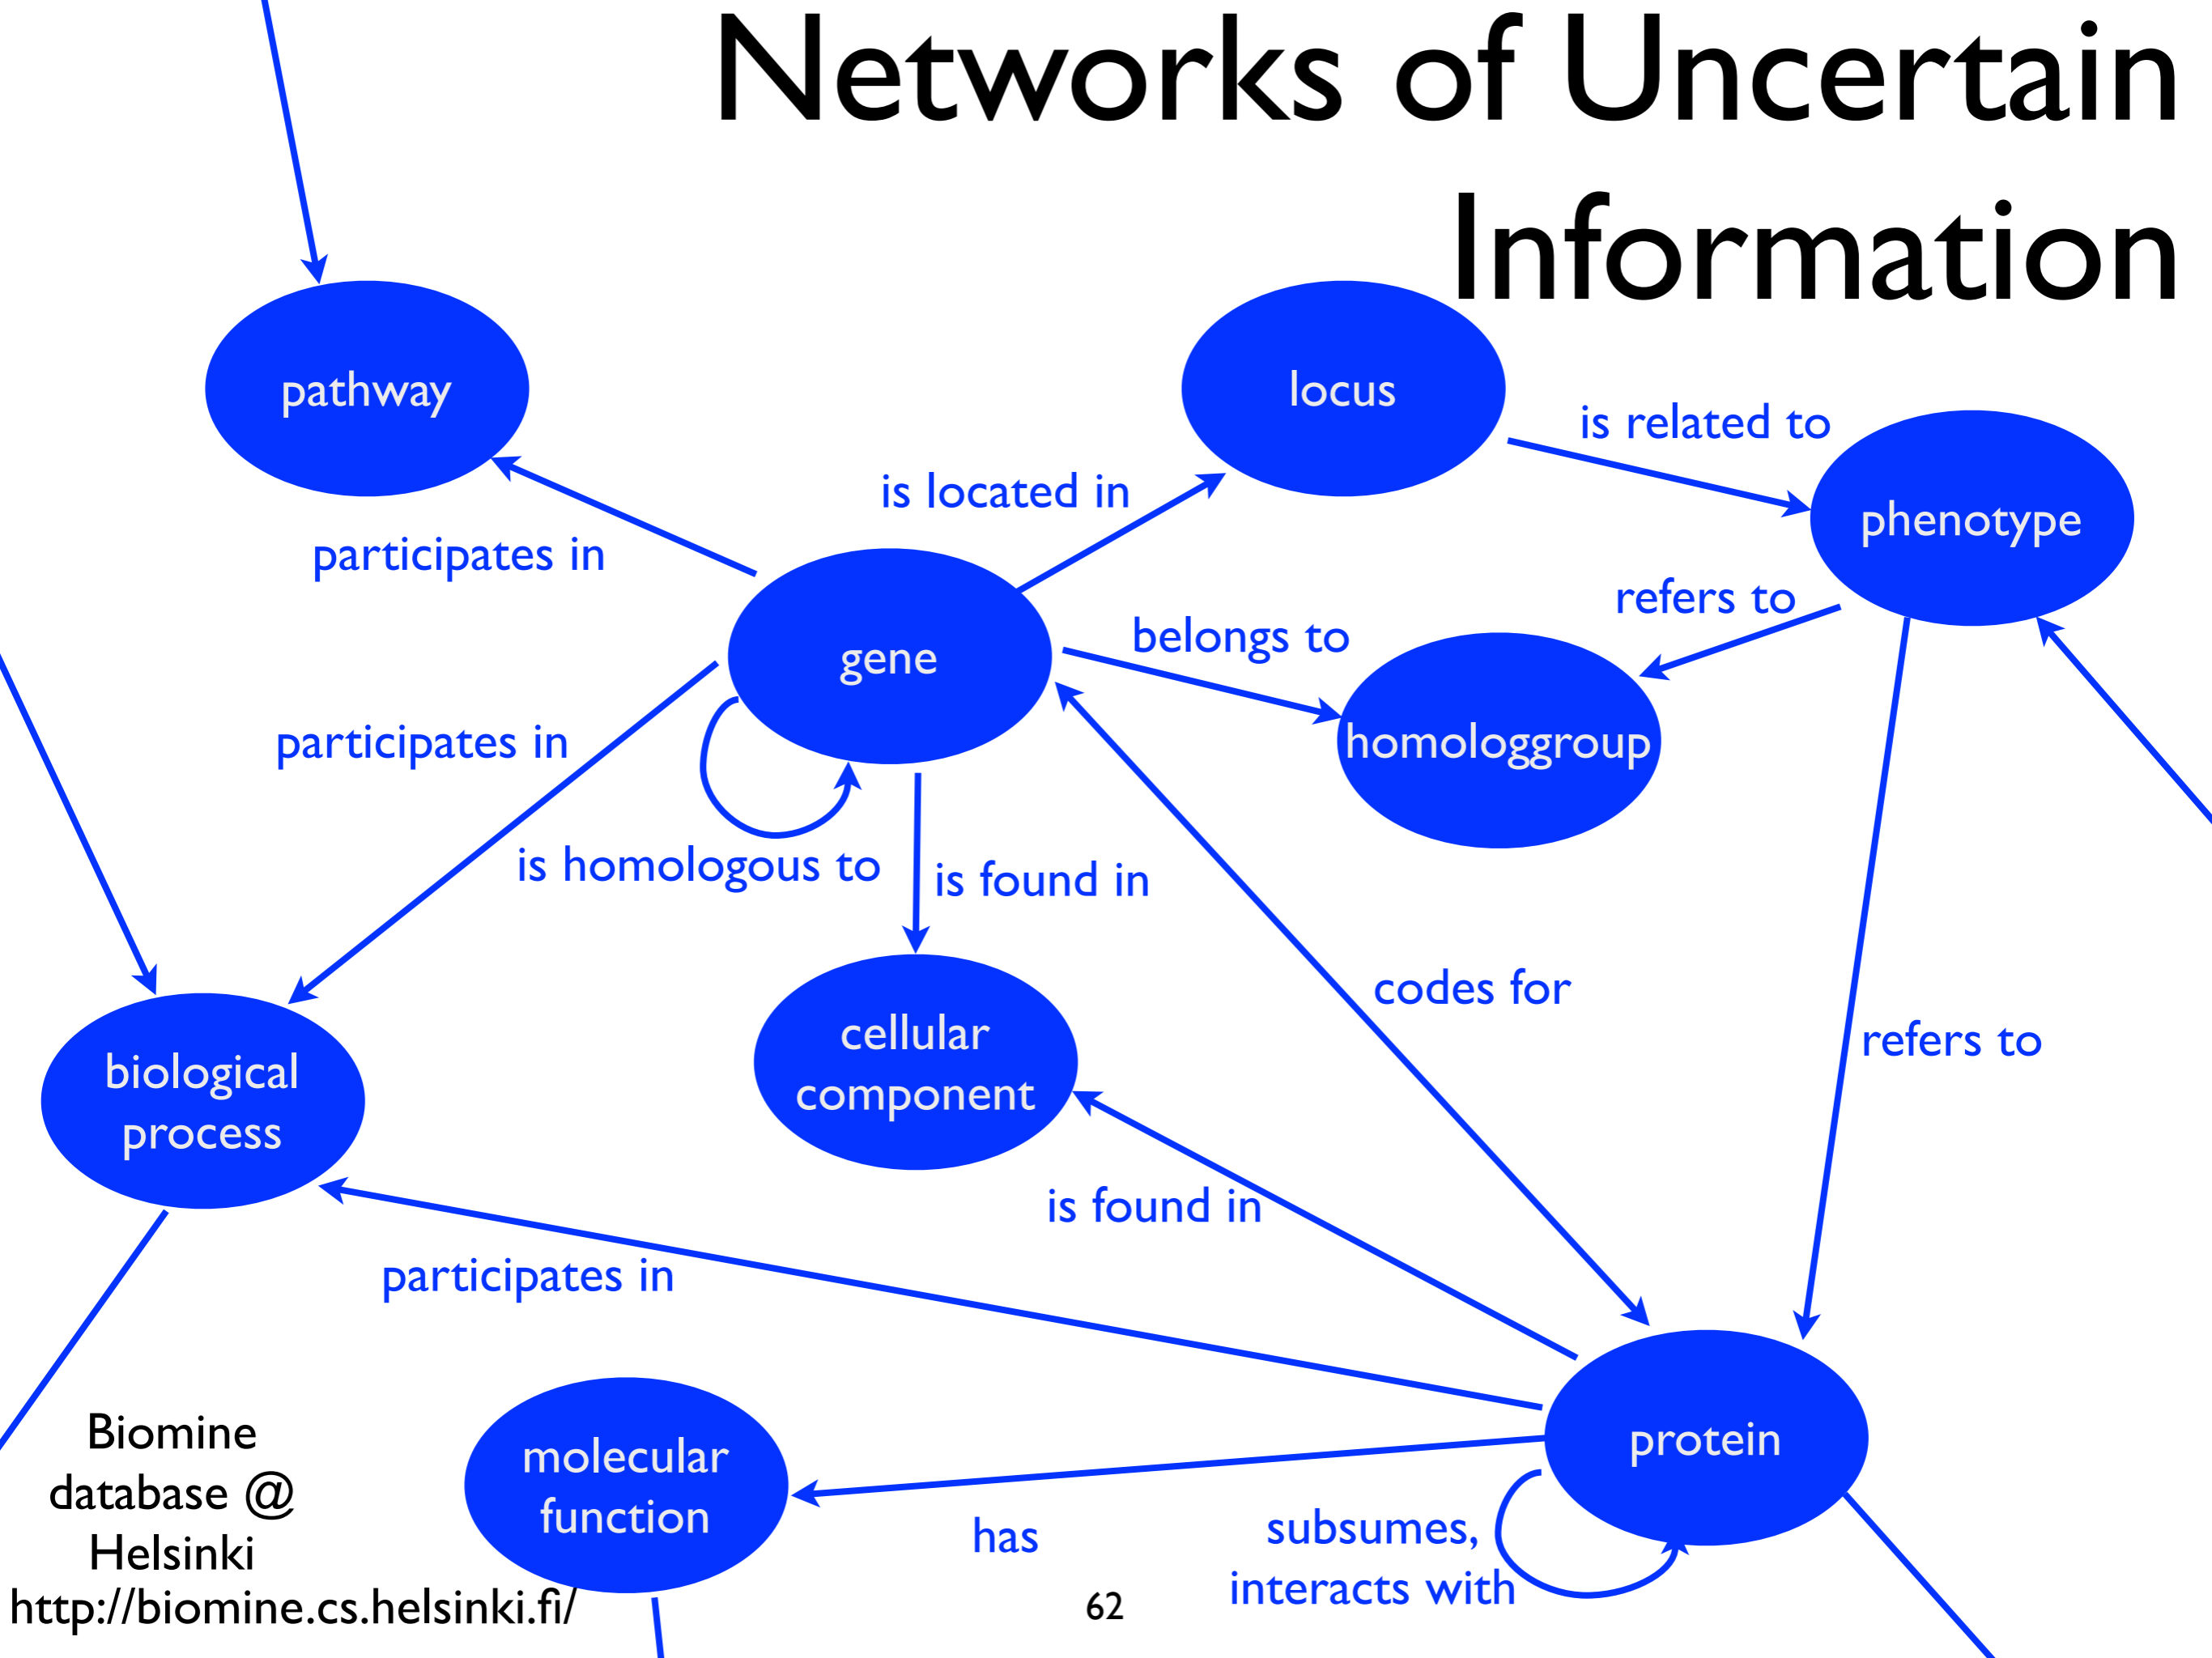
```
    ≈width(OBot) ≥ ≈width(OTop).
```

comparing values of
random variables



[Gutmann et al, TPLP 11; Nitti et al, IROS 13;
Nitti et al. MLJ]

Networks of Uncertain Information



Biology

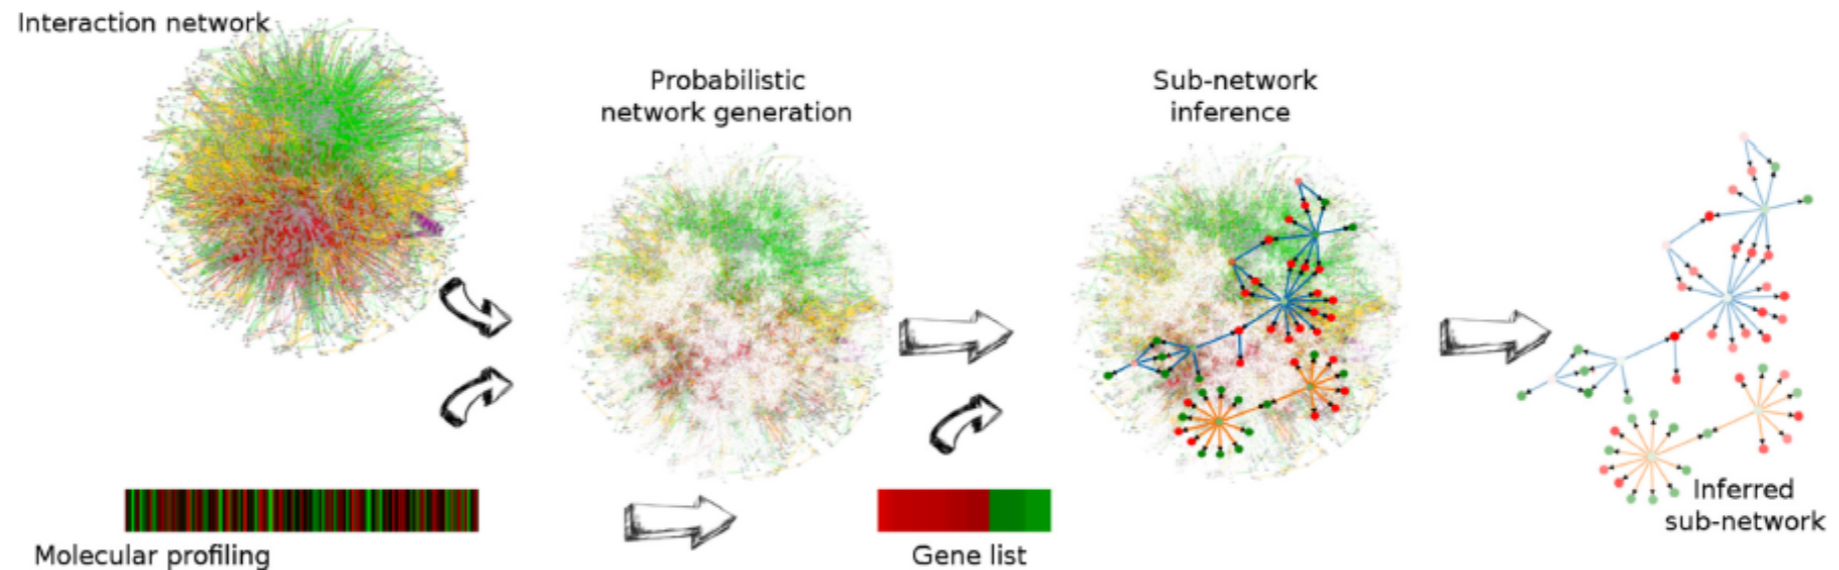
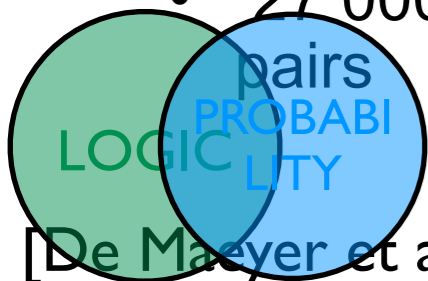
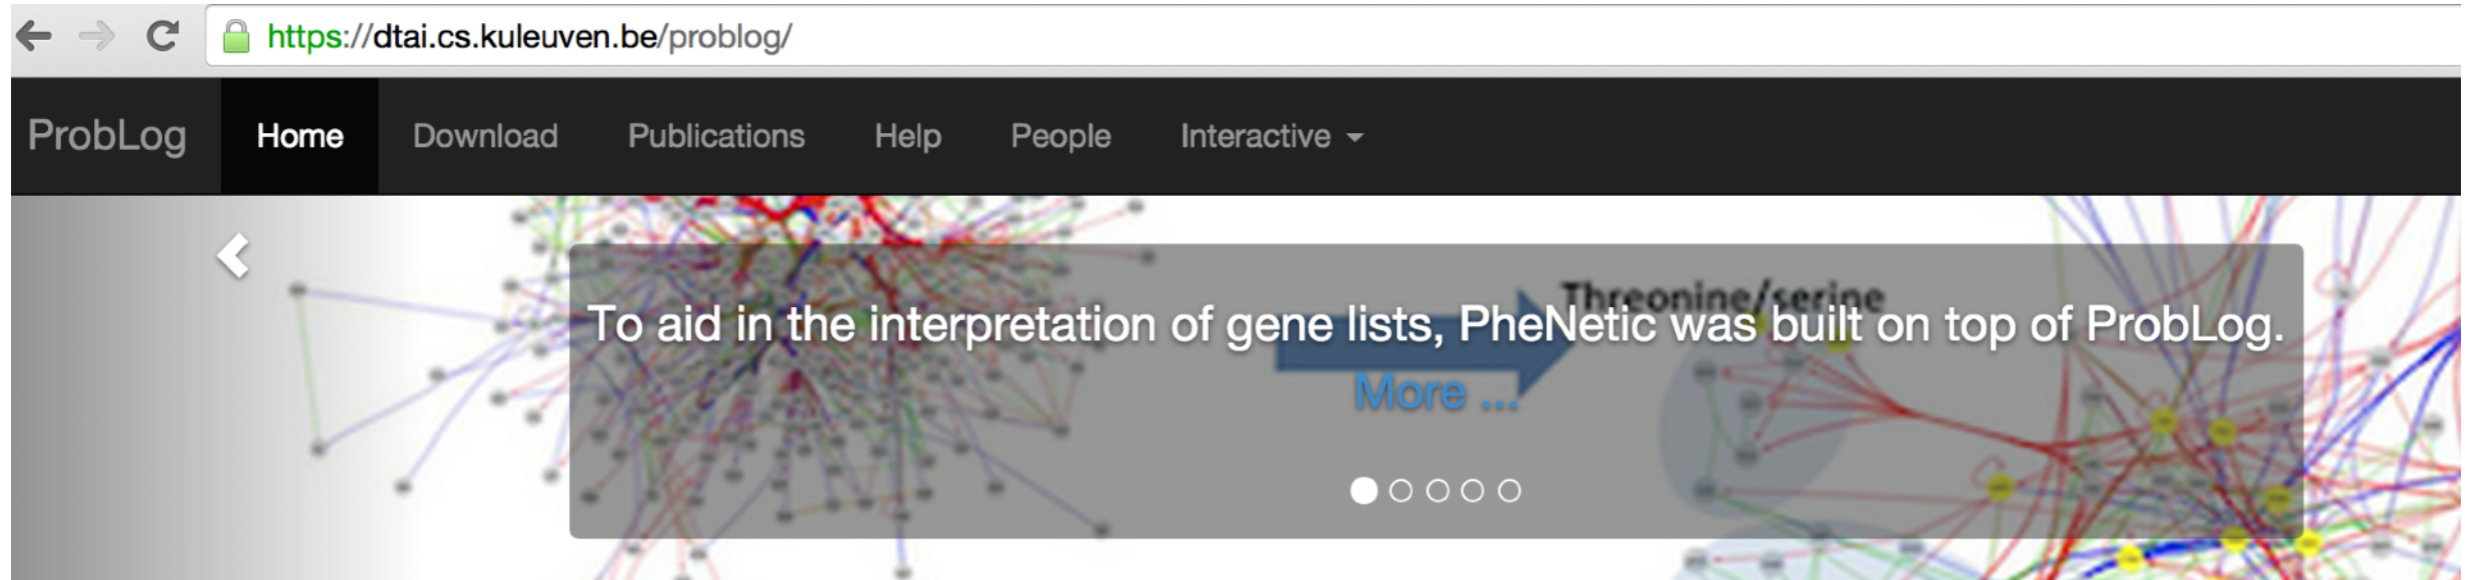


Figure 1. Overview of PheNetic, a web service for network-based interpretation of ‘omics’ data. The web service uses as input a genome wide interaction network for the organism of interest, a user generated molecular profiling data set and a gene list derived from these data. Interaction networks for a wide variety of organisms are readily available from the web server. Using the uploaded user-generated molecular data the interaction network is converted into a probabilistic network: edges receive a probability proportional to the levels measured for the terminal nodes in the molecular profiling data set. This probabilistic interaction network is used to infer the sub-network that best links the genes from the gene list. The inferred sub-network provides a trade-off between linking as many genes as possible from the gene list and selecting the least number of edges.

- Causes: Mutations
- All related to similar phenotype
- Effects: Differentially expressed genes
- 27 000 cause effect pairs
- Interaction network:
 - 3063 nodes
 - Genes
 - Proteins
 - 16794 edges
 - Molecular interactions
 - Uncertain
- Goal: connect causes to effects through common subnetwork
 - = Find mechanism
- Techniques:
 - DTPProbLog
 - Approximate inference





Introduction.

Probabilistic logic programs are logic programs in which some of the facts are annotated with probabilities.

ProbLog is a tool that allows you to intuitively build programs that do not only encode **complex interactions** between a large sets of **heterogenous components** but **uncertainties** that are present in real-life situations.

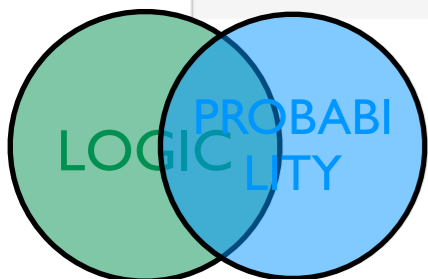
The engine tackles several tasks such as computing the marginals given evidence and learning from (partial) interpretations. ProbLog is a suite of efficient algorithms tasks. It is based on a conversion of the program and the queries and evidence to a weighted Boolean formula. This allows us to reduce the inference tasks to well-s weighted model counting, which can be solved using state-of-the-art methods known from the graphical model and knowledge compilation literature.

The Language. Probabilistic Logic Programming.

ProbLog makes it easy to express complex, probabilistic models.

```
0.3::stress(X) :- person(X).
0.2::influences(X,Y) :- person(X), person(Y).

smokes(X) :- stress(X).
smokes(X) :- friend(X,Y), influences(Y,X), smokes(Y).
```



Probabilistic Programming Languages outside LP

- IBAL [Pfeffer 01]
- Figaro [Pfeffer 09]
- Church [Goodman et al 08]
- BLOG [Milch et al 05]
- Stan & Edward & Anglican
- and many more appearing recently such

Church

probabilistic functional programming

[Goodman et al, UAI 08]

several possible executions

```
(define randplus5  
  (lambda (x) (if (flip 0.6)  
                  (+ x 5)  
                  x)))  
  
(map randplus5 '(1 2 3))
```

probabilistic primitives + functional program
→ distribution over possible executions

Dealing with
primitivity

Reasoning with
probabilistic data

one execution

```
(define plus5 (lambda (x) (+ x 5)))  
  
(map plus5 '(1 2 3))
```

Learning

Church vs ProbLog

```
(define randplus5 (lambda (x) (if (flip 0.6) (+ x 5) x)))
```

```
(map randplus5 '(1 2))
```

Church result: (1 2) with 0.4×0.4

(1 7) with 0.4×0.6

(6 2) with 0.6×0.4

(6 7) with 0.6×0.6

```
0.4::p5(N,N);0.6::p5(N,M) :- M is N+5.
```

```
lp5([],[]).
```

```
lp5([N|L],[M|K]) :-
```

```
  p5(N,M),
```

```
  lp5(L,K).
```

```
query(lp5([1,2],_)).
```

ProbLog result: (1 2) with 0.4×0.4

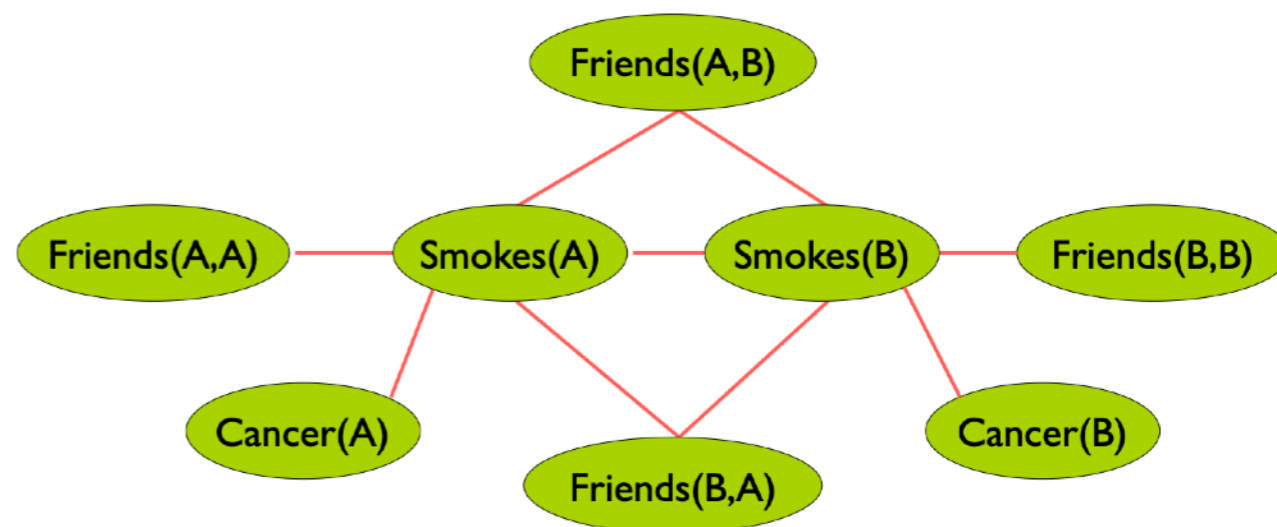
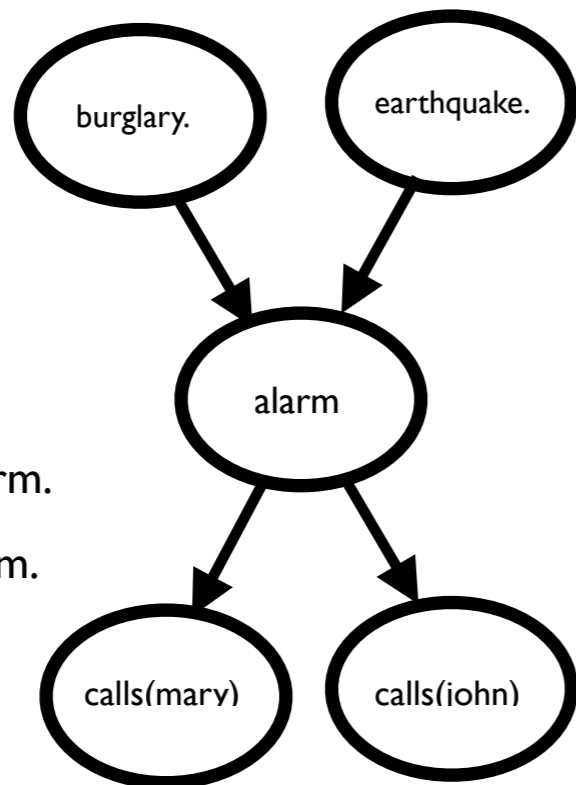
(1 7) with 0.4×0.6

(6 2) with 0.6×0.4

(6 7) with 0.6×0.6

2. Directed vs Undirected the PGM / StarAI dimension

0.1 :: burglary.
 0.05 :: earthquake.
 alarm :- earthquake.
 alarm :- burglary.
 0.7::calls(mary) :- alarm.
 0.6::calls(john) :- alarm.



$$1.5 \quad \forall x \text{ Smokes}(x) \Rightarrow \text{Cancer}(x)$$

$$1.1 \quad \forall x, y \text{ Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$$

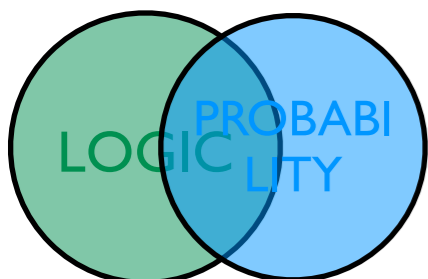
Probabilistic Logic Programs
ProbLog

directed
Bayesian Net

Markov Logic

undirected
Markov Net
model theoretic

key representatives



Markov Logic: Intuition

- *Undirected graphical model*
- A logical KB is a set of **hard constraints** on the set of possible worlds
- Let's make them **soft constraints**:
When a world violates a formula, it becomes less probable, not impossible
- Give each formula a **weight**
(Higher weight \Rightarrow Stronger constraint)

$$P(\text{world}) \propto \exp\left(\sum \text{weights of formulas it satisfies}\right)$$

A possible worlds view

Say we have two domain elements **Anna** and **Bob** as well as two predicates **Friends** and **Happy**

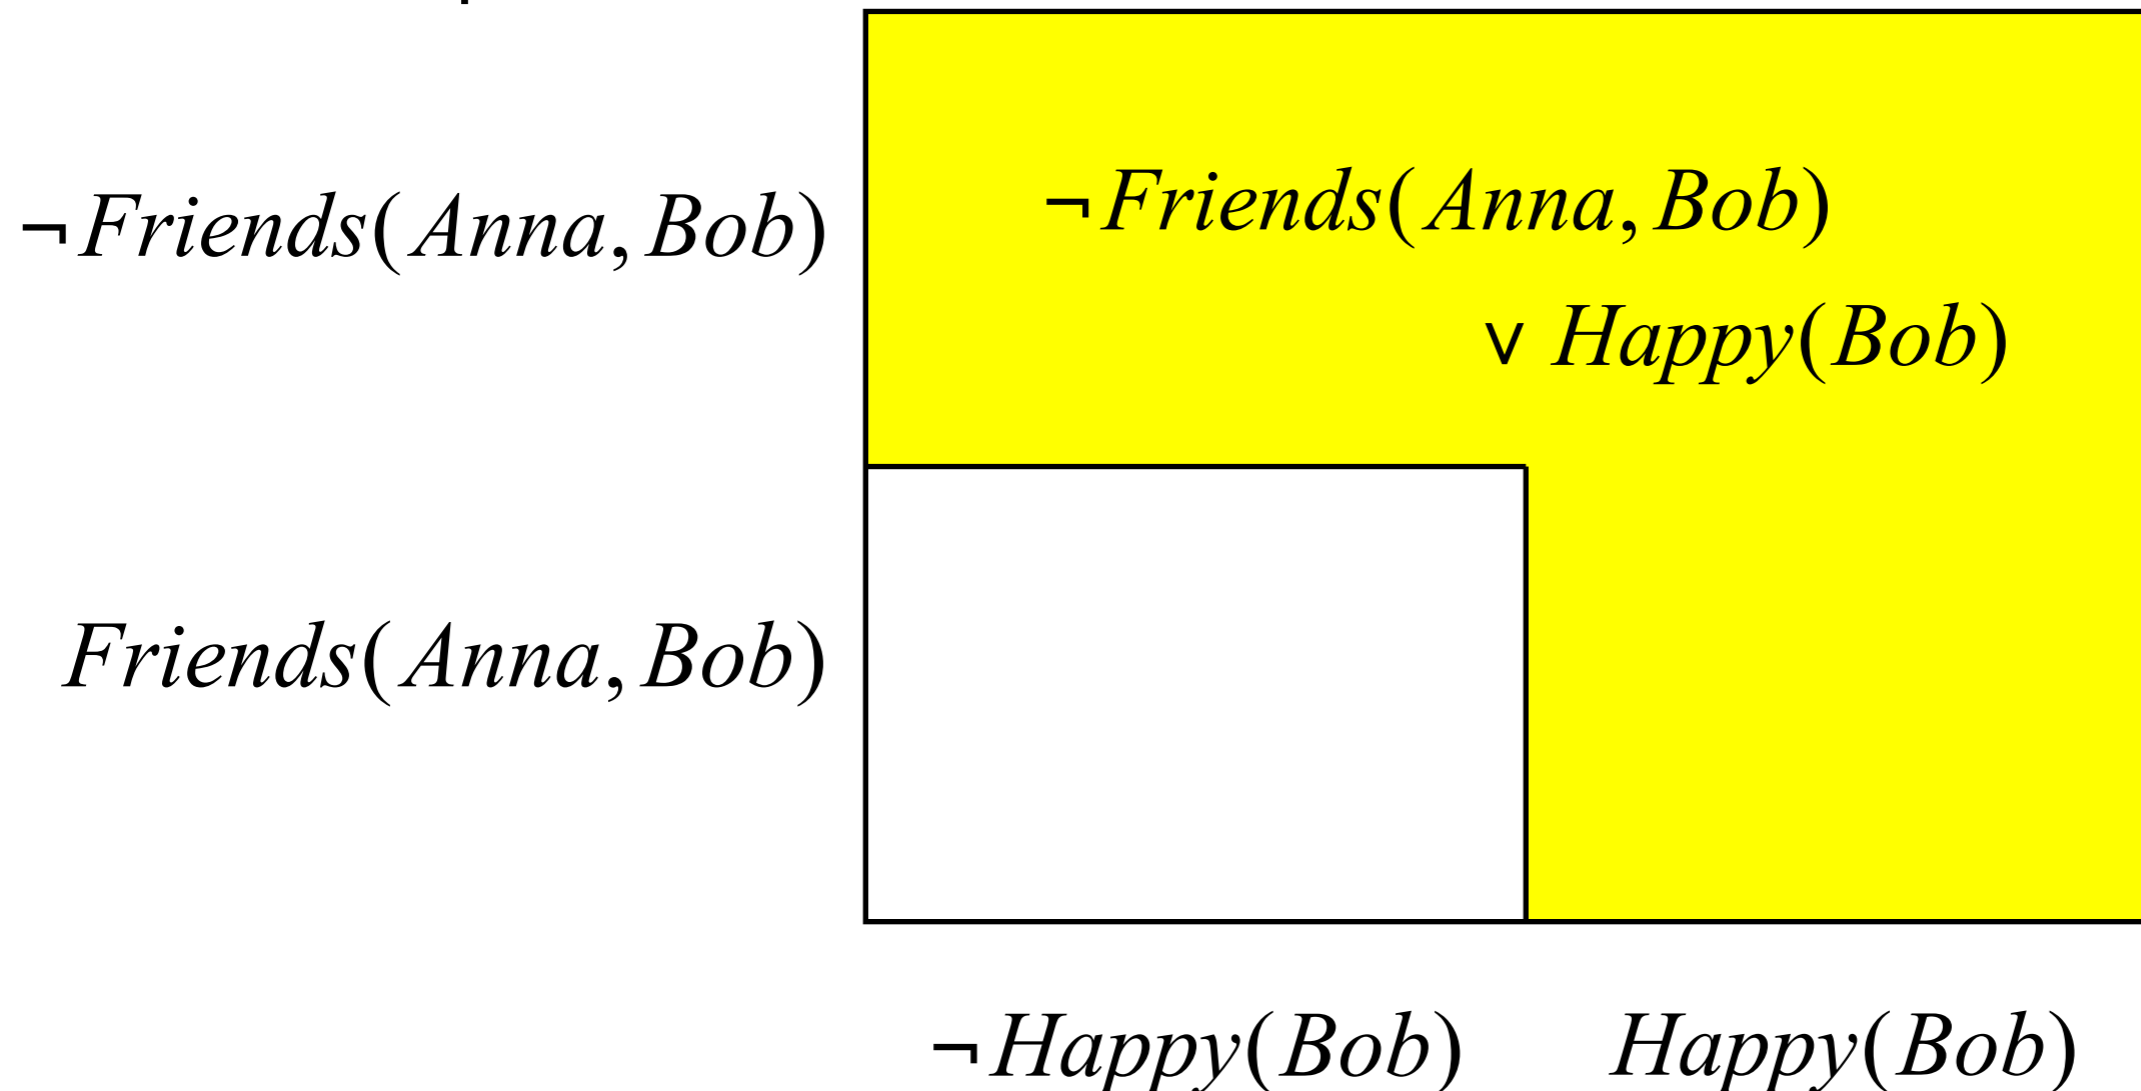
$\neg \text{Friends}(\text{Anna}, \text{Bob})$		
$\text{Friends}(\text{Anna}, \text{Bob})$		
	$\neg \text{Happy}(\text{Bob})$	$\text{Happy}(\text{Bob})$

A possible worlds view

Logical formulas such as

not Friends(Anna,Bob) or Happy(Bob)

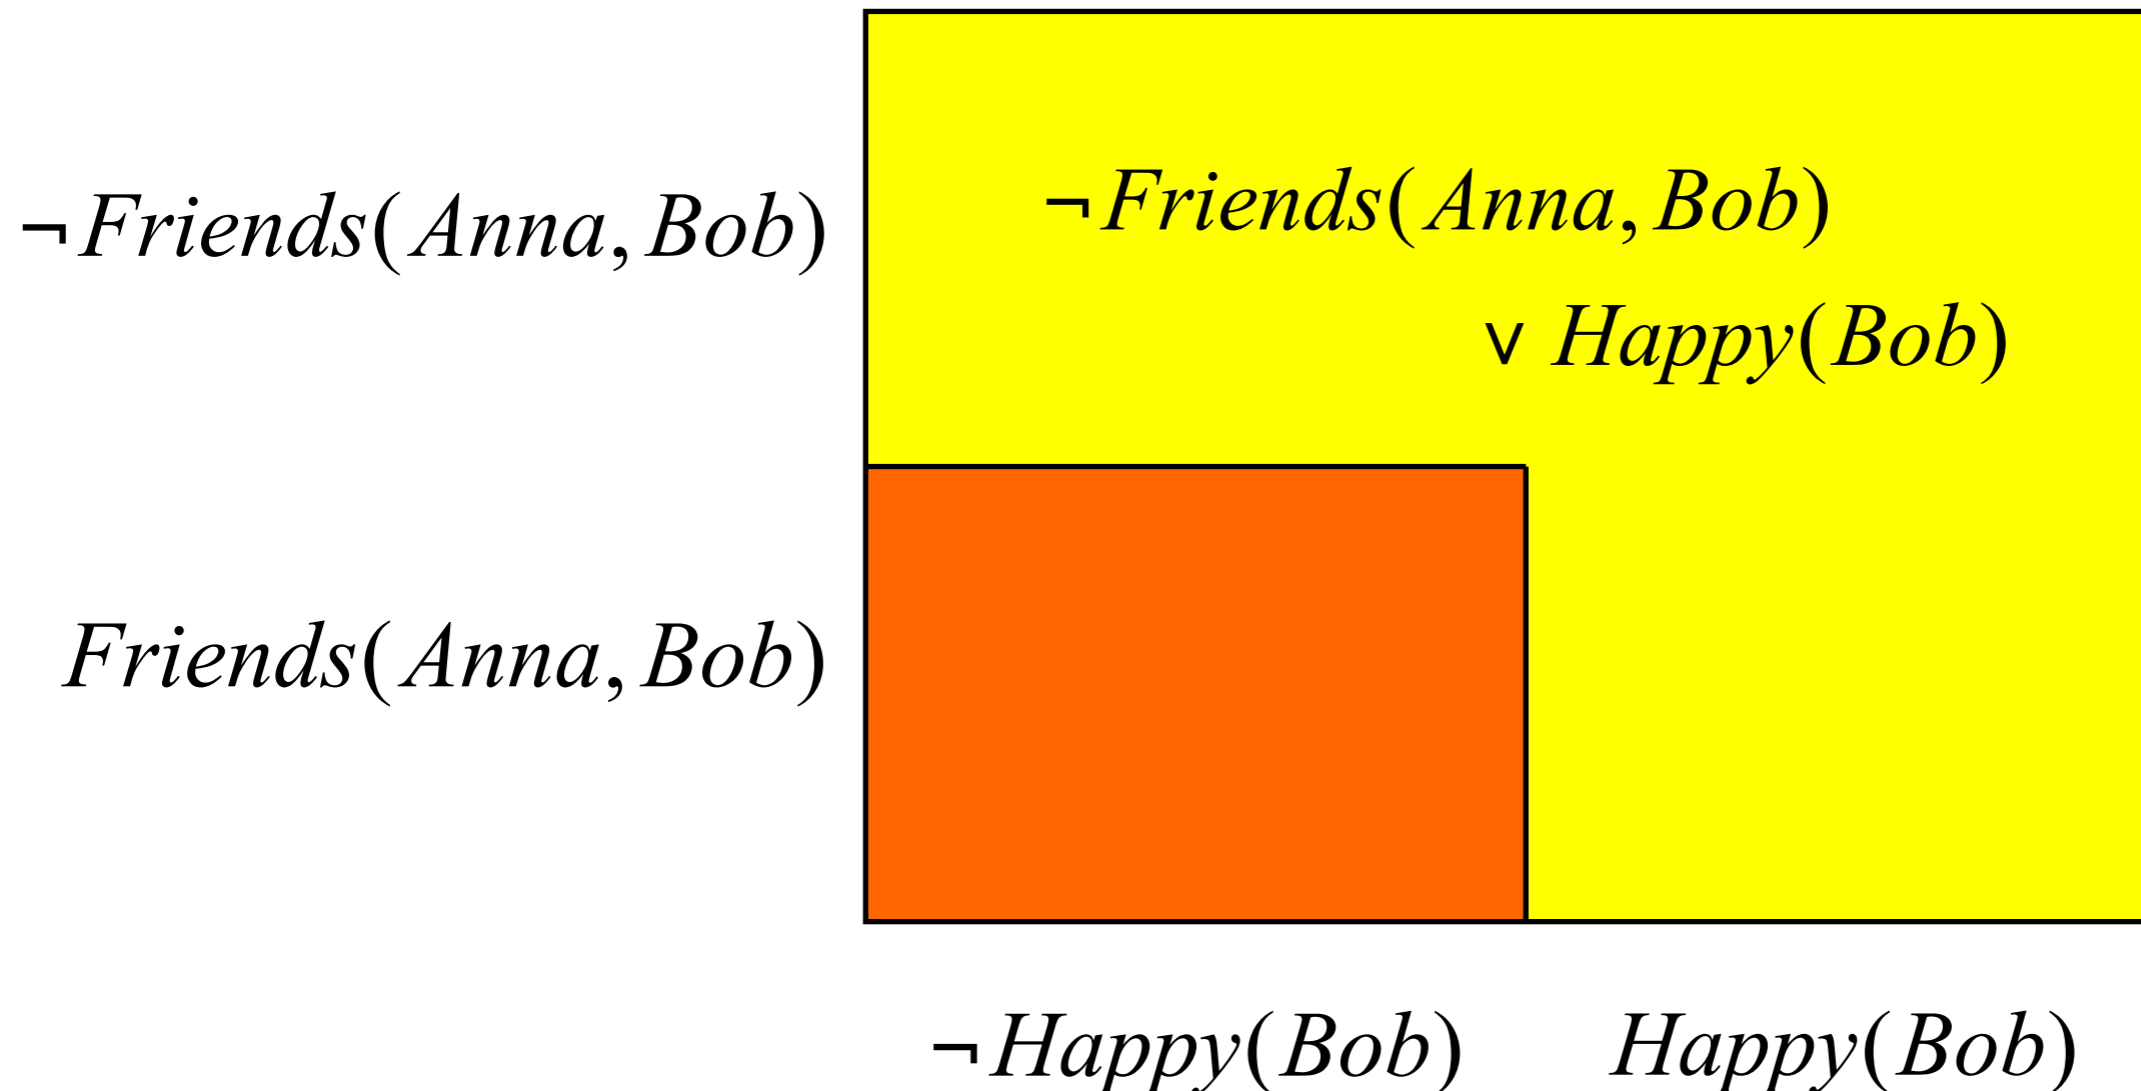
exclude possible worlds



A possible worlds view

Instead of excluding worlds, we want them to become less likely,
e.g.

$$P(\neg \text{Friends}(\text{Anna}, \text{Bob}) \vee \text{Happy}(\text{Bob})) = 0.8$$



A possible worlds view

four times as likely that rule holds

$$\Phi(\neg \text{Friends}(\text{Anna}, \text{Bob}) \vee \text{Happy}(\text{Bob})) = 1$$

$$\Phi(\text{Friends}(\text{Anna}, \text{Bob}) \wedge \neg \text{Happy}(\text{Bob})) = 0.75$$

$\neg \text{Friends}(\text{Anna}, \text{Bob})$	1	1
$\text{Friends}(\text{Anna}, \text{Bob})$	0.75	1
	$\neg \text{Happy}(\text{Bob})$	$\text{Happy}(\text{Bob})$

A possible worlds view

Or as log-linear model this is:

$$w(\Phi(\neg Friends(Anna, Bob) \vee Happy(Bob)))$$

$$= \log(1 / 0.75) = 0.29$$

$\neg Friends(Anna, Bob)$	1	1
$Friends(Anna, Bob)$	0.75	1
	$\neg Happy(Bob)$	$Happy(Bob)$

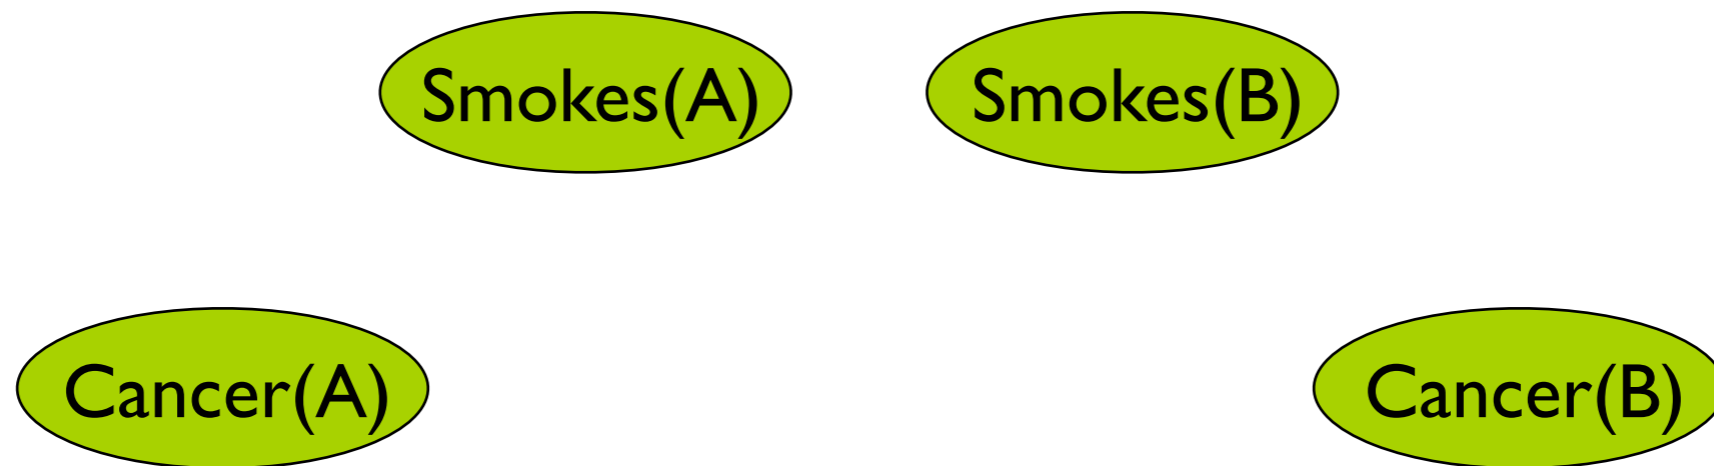
This can also be viewed as ⁷⁴ building a graphical model

Markov Logic

1.5 $\forall x \text{Smokes}(x) \Rightarrow \text{Cancer}(x)$

1.1 $\forall x, y \text{Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$

Suppose we have two constants: **Anna (A)** and **Bob (B)**

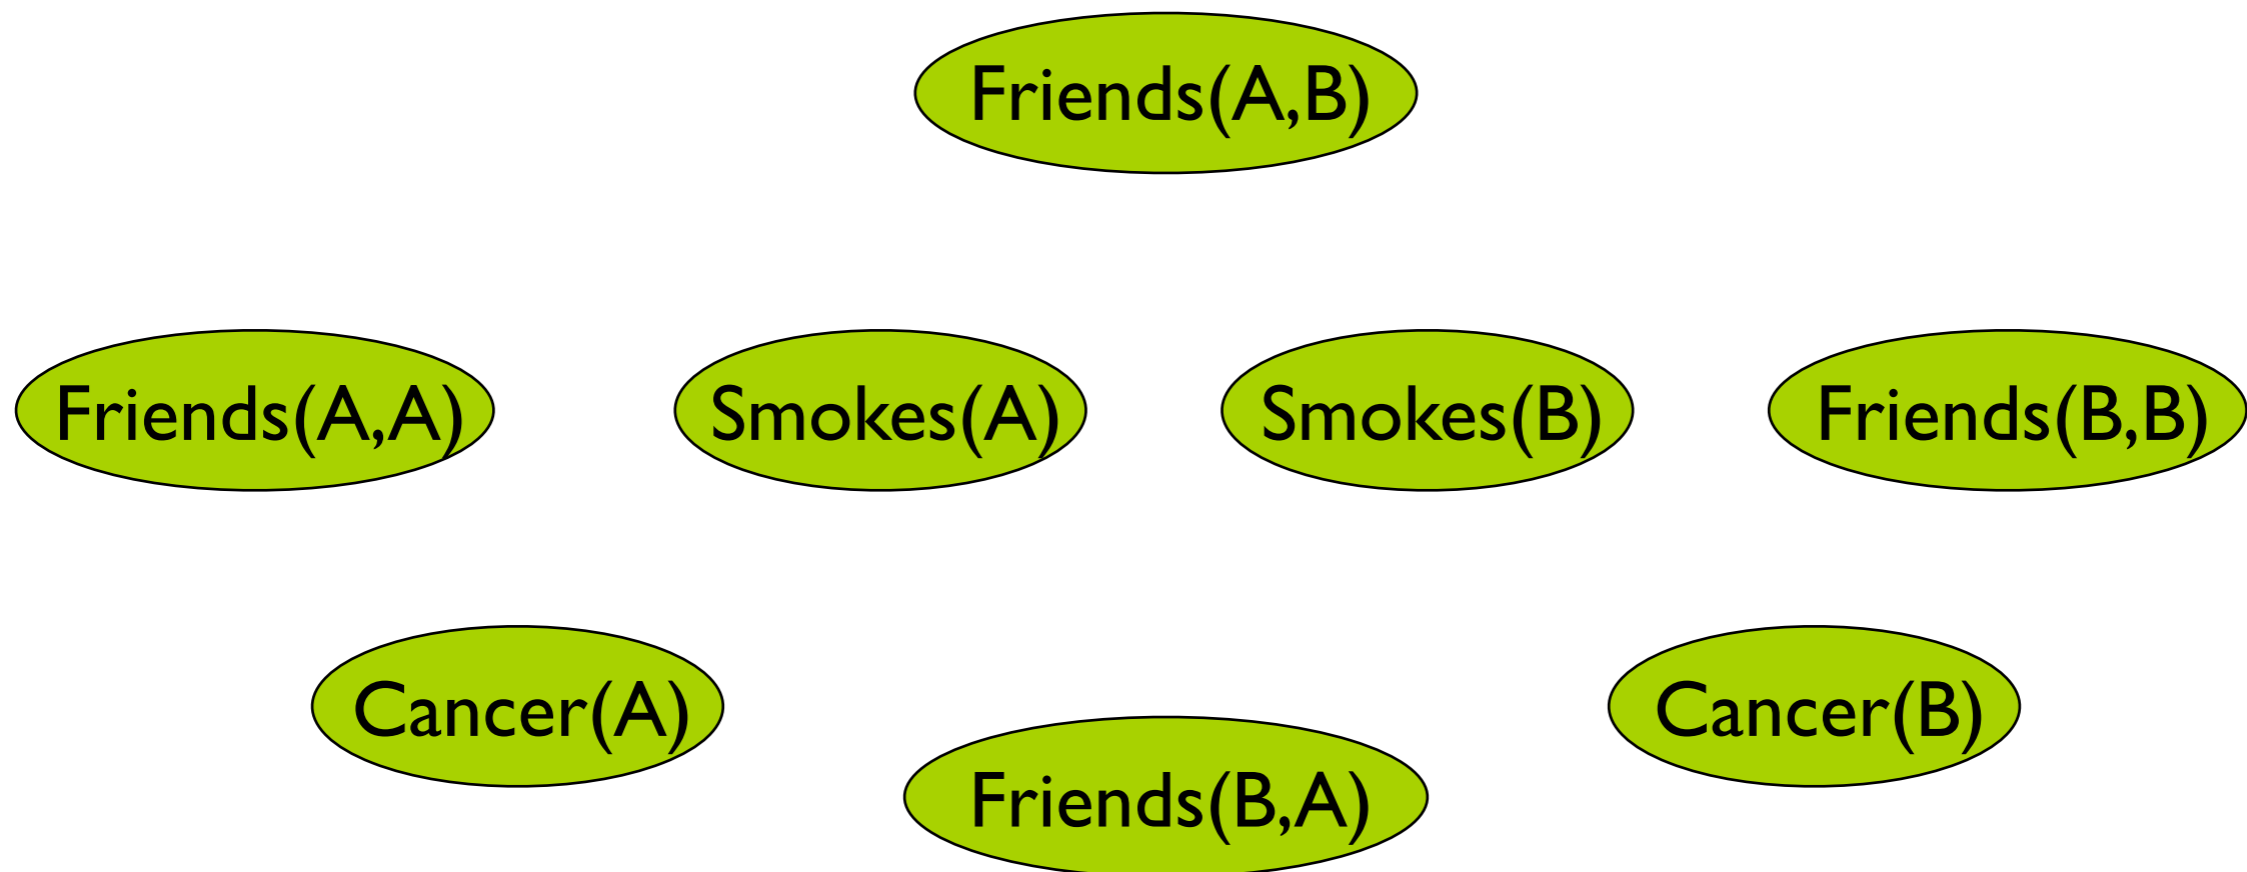


Markov Logic

1.5 $\forall x \text{Smokes}(x) \Rightarrow \text{Cancer}(x)$

1.1 $\forall x, y \text{Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$

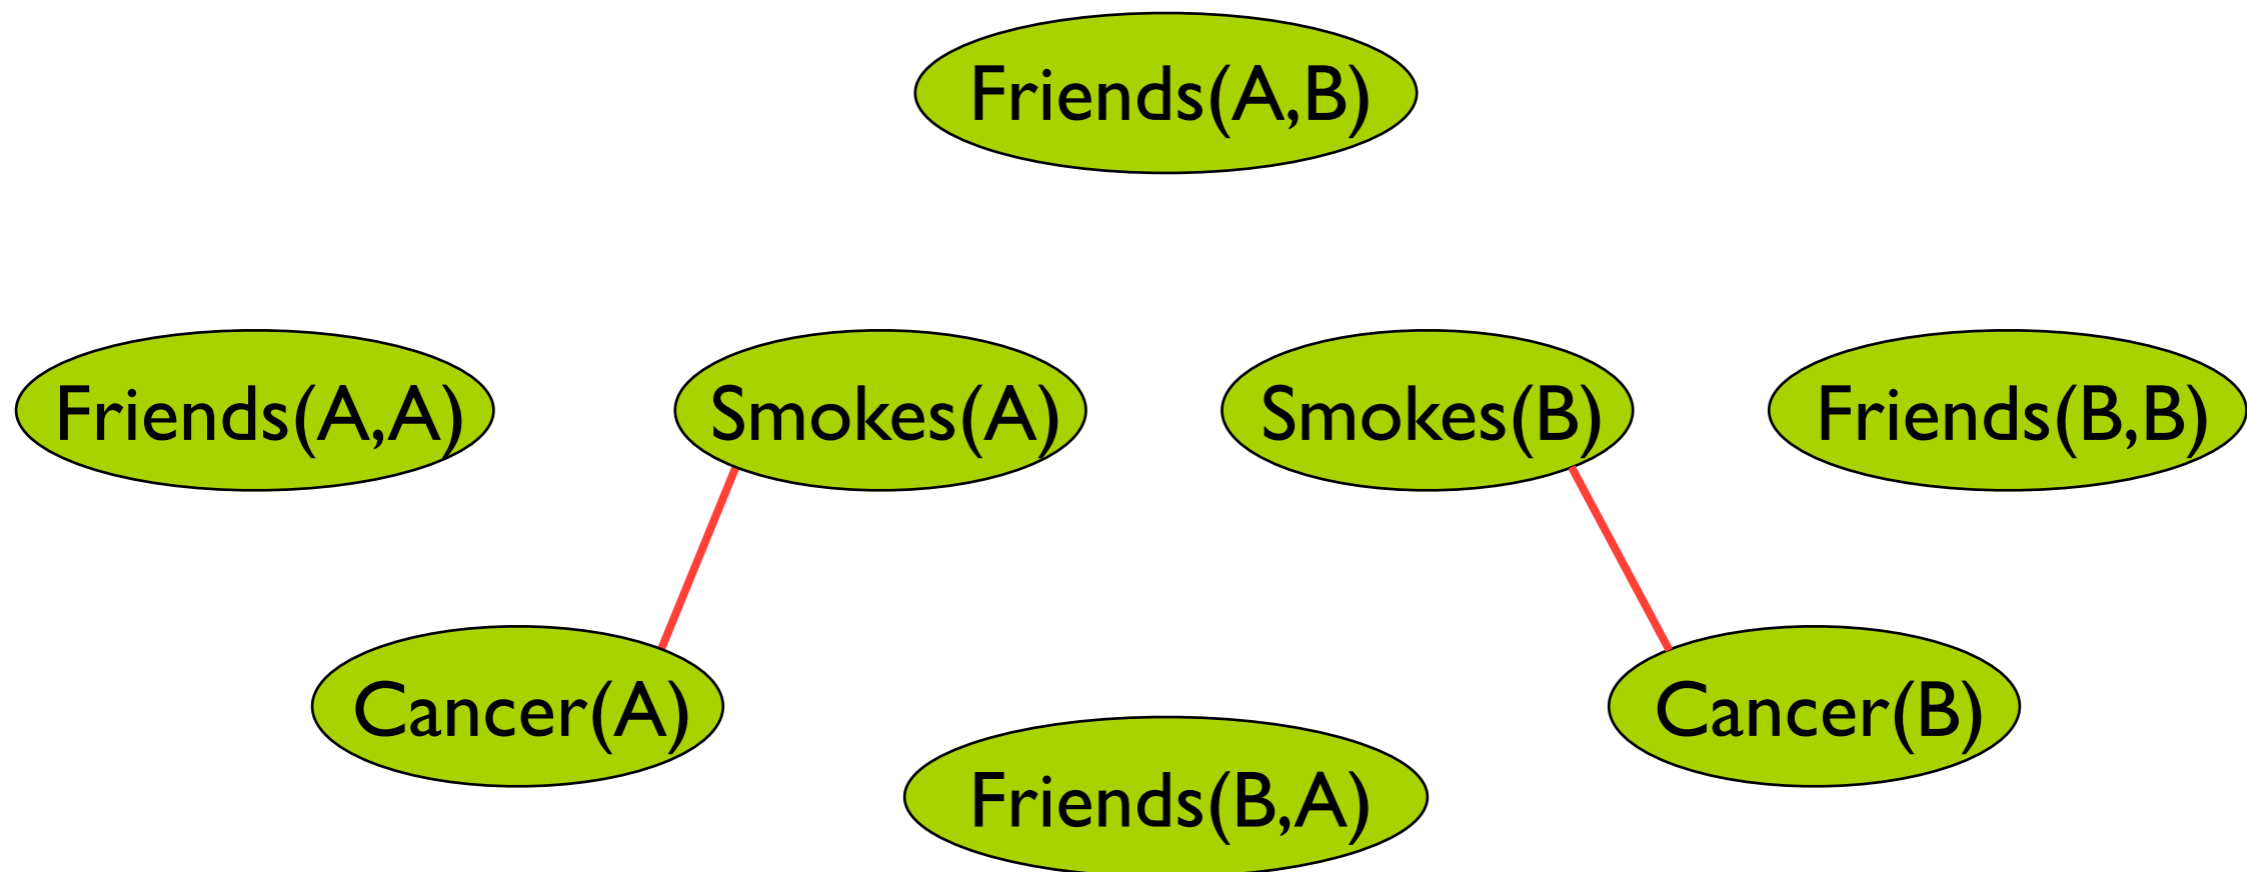
Suppose we have two constants: **Anna (A)** and **Bob (B)**



Markov Logic

1.5	$\forall x \text{ Smokes}(x) \Rightarrow \text{Cancer}(x)$
1.1	$\forall x, y \text{ Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$

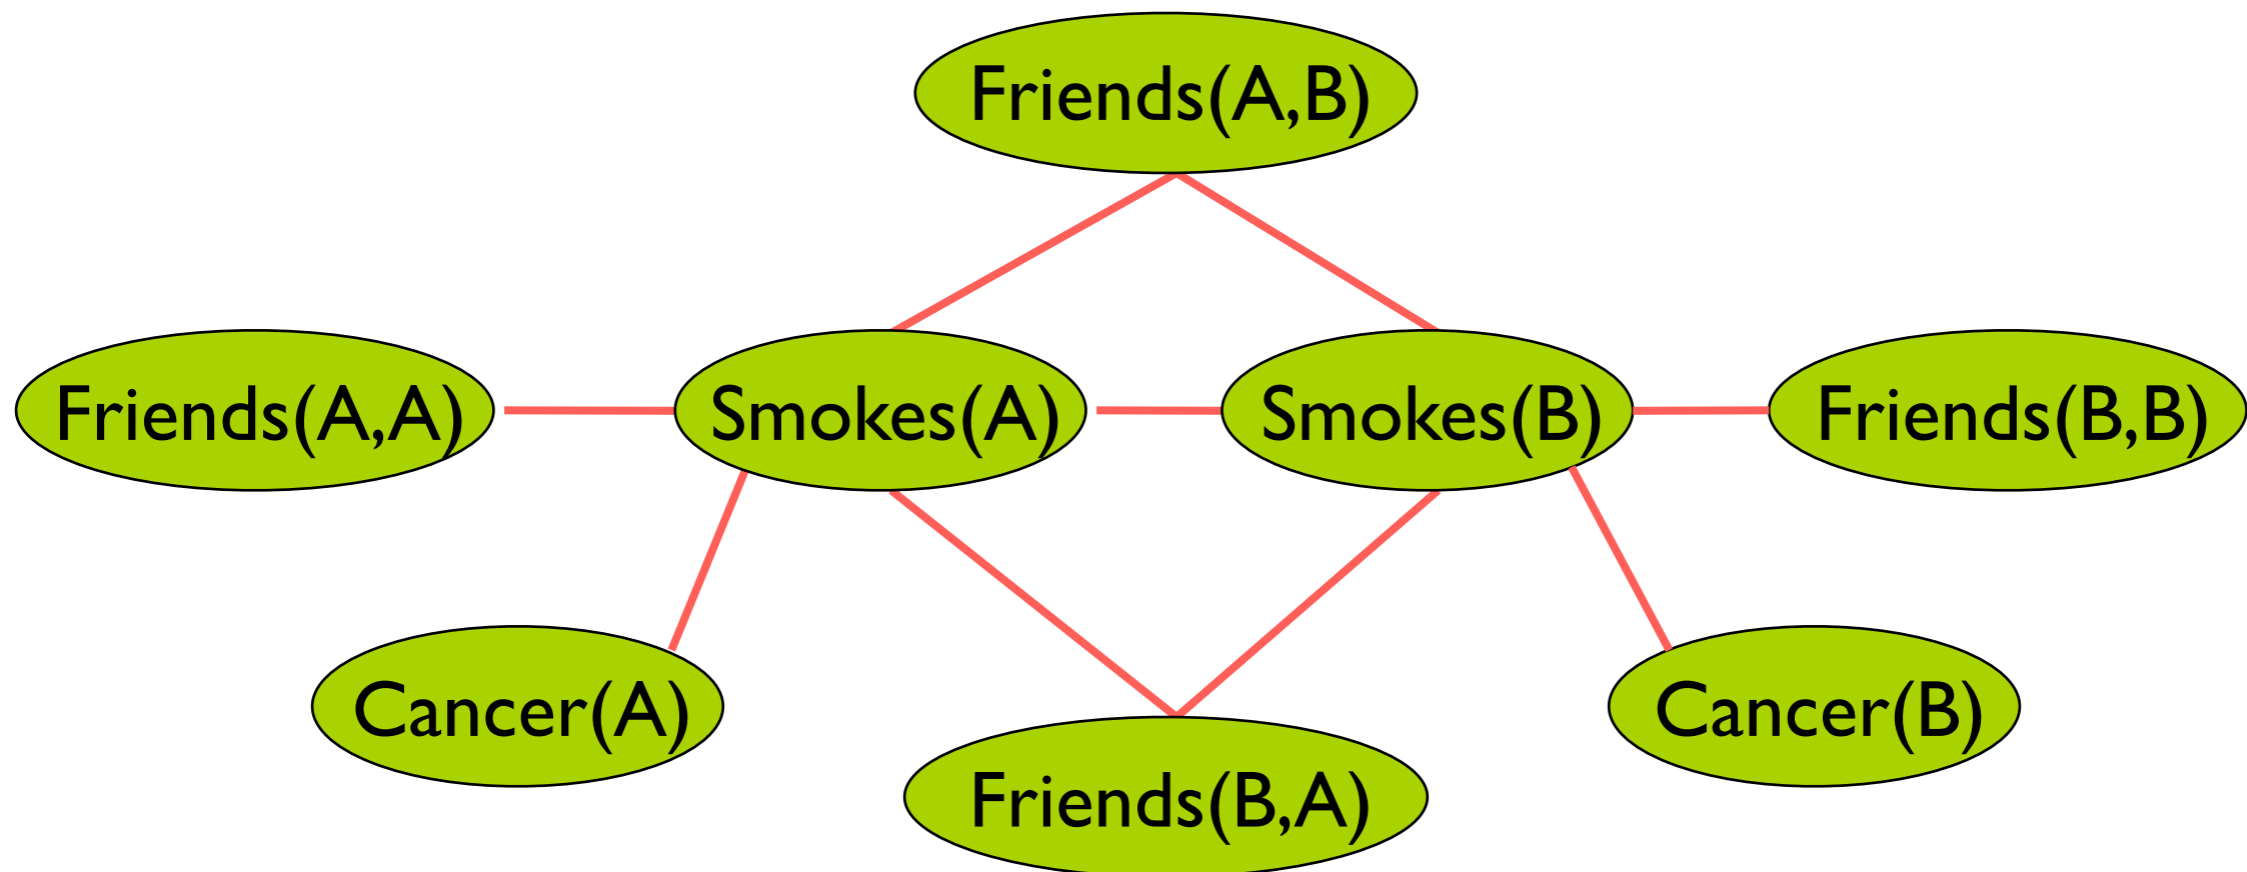
Suppose we have two constants: **Anna** (A) and **Bob** (B)



Markov Logic

1.5	$\forall x \text{ Smokes}(x) \Rightarrow \text{Cancer}(x)$
1.1	$\forall x, y \text{ Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$

Suppose we have two constants: **Anna (A)** and **Bob (B)**



Markov Logic

- A Markov Logic Network (MLN) is a set of pairs (F, w) where
 - F is a formula in first-order logic
 - w is a real number
- An MLN defines a Markov network with
 - One node for each grounding of each predicate in the MLN
 - One feature for each grounding of each formula F in the MLN, with the corresponding weight w
- Probability of a world

$$P(x) = \frac{1}{Z} \exp \left(\sum_i w_i n_i(x) \right)$$

Weight of formula i No. of true groundings of formula i in x

Possible Worlds

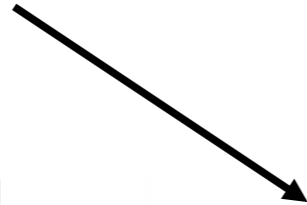
A vocabulary

Smokes(Alice)	Smokes(Bob)	Friends(Alice,Bob)	Friends(Bob,Alice)
0	0	0	0
⋮	⋮	⋮	⋮
1	0	1	0
⋮	⋮	⋮	⋮
1	1	1	1

Possible worlds
Logical interpretations

Possible Worlds

A logical theory



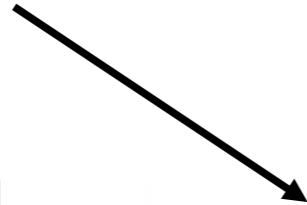
$$\forall x,y, \text{Smokes}(x) \wedge \text{Friends}(x,y) \Rightarrow \text{Smokes}(y)$$

Smokes(Alice)	Smokes(Bob)	Friends(Alice,Bob)	Friends(Bob,Alice)	theory
0	0	0	0	1
⋮	⋮	⋮	⋮	⋮
1	0	1	0	0
⋮	⋮	⋮	⋮	⋮
1	1	1	1	1

Interpretations that satisfy the theory
Models

First-Order Model Counting

A logical theory



$$\forall x,y, \text{Smokes}(x) \wedge \text{Friends}(x,y) \Rightarrow \text{Smokes}(y)$$

Smokes(Alice)	Smokes(Bob)	Friends(Alice,Bob)	Friends(Bob,Alice)	theory
0	0	0	0	1
⋮	⋮	⋮	⋮	⋮
1	0	1	0	0
⋮	⋮	⋮	⋮	⋮
1	1	1	1	1

A green bracket on the right side of the table groups the rows and is labeled with the summation symbol Σ .

First-order model count
 $\sim \#SAT$

Markov Logic

- MLNs are a template for ground Markov Networks
- Probability of a world/interpretation
- If $n_i = 0$ then $P(x) = \frac{1}{Z}$

$$P(x) = \frac{1}{Z} \exp \left(\sum_i w_i n_i(x) \right)$$

Weight of formula i

No. of true groundings of formula i in x

Markov Logic

A Markov Logic theory



Smokes(Alice)	Smokes(Bob)	Friends(Alice,Bob)	Friends(Bob,Alice)	theory
0	0	0	0	$\frac{1}{Z} \exp(1.5 * 2)$
1	0	1	0	$\frac{1}{Z} \exp(1.5 * 1)$
1	1	1	1	$\frac{1}{Z} \exp(1.5 * 2)$

$1.5 \forall x,y, \text{Smokes}(x) \wedge \text{Friends}(x,y) \Rightarrow \text{Smokes}(y)$

counting only substitutions for which $X \neq Y$
 $X=\text{Alice}, Y=\text{Bob}$
 $X=\text{Bob}, Y=\text{Alice}$

Markov Logic

A Markov Logic theory

Smokes(Alice)	Smokes(Bob)	Friends(Alice,Bob)	Friends(Bob,Alice)	theory
0	0	0	0	$\frac{1}{Z} \exp(1.5 * 2)$
1	0	1	0	$\frac{1}{Z} \exp(1.5 * 1)$
1	1	1	1	$\frac{1}{Z} \exp(1.5 * 2)$

$1.5 \forall x,y, \text{Smokes}(x) \wedge \text{Friends}(x,y) \Rightarrow \text{Smokes}(y)$

Z
 partition function

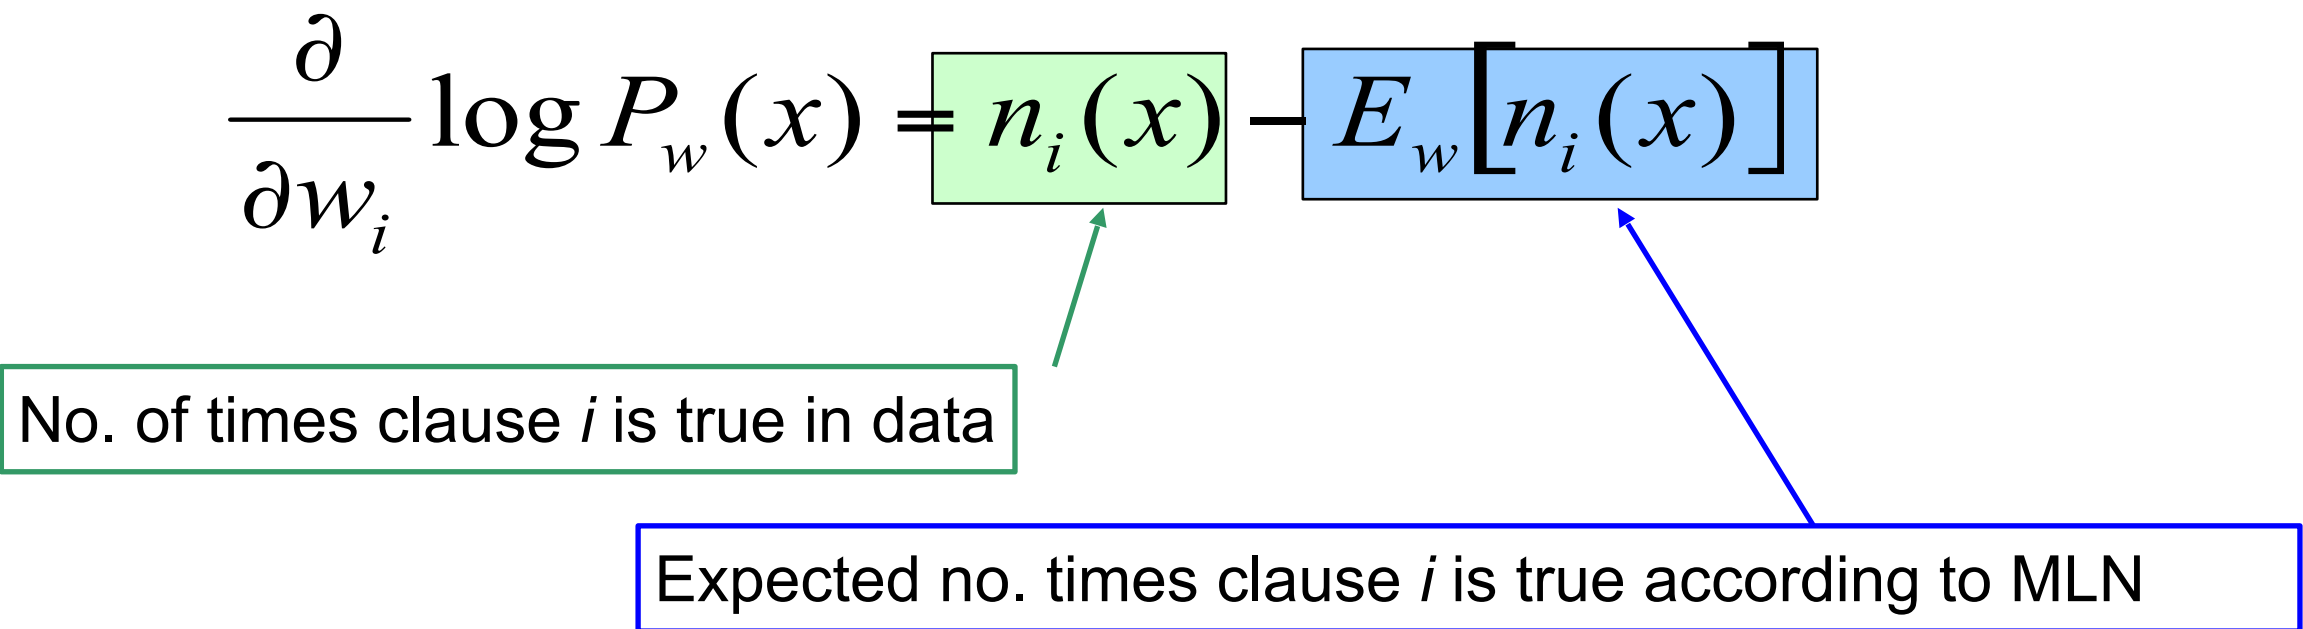
Markov Logic

- A Markov Logic Network (MLN) is a set of pairs (F, w) where
 - F is a formula in first-order logic
 - w is a real number
- An MLN defines a Markov network with
 - One node for each grounding of each predicate in the MLN
 - One feature for each grounding of each formula F in the MLN, with the corresponding weight w
- Probability of a world

$$P(x) = \frac{1}{Z} \exp \left(\sum_i w_i n_i(x) \right)$$

Weight of formula i No. of true groundings of formula i in x

Parameter Learning

$$\frac{\partial}{\partial w_i} \log P_w(x) = n_i(x) - E_w[n_i(x)]$$
The equation is presented with the term $n_i(x)$ enclosed in a light green box and $E_w[n_i(x)]$ enclosed in a light blue box. A green arrow points from the green box to a green-bordered text box below it. A blue arrow points from the blue box to a blue-bordered text box below it.

No. of times clause i is true in data

Expected no. times clause i is true according to MLN

Has been used for generative learning (Pseudolikelihood);
Many variations (also discriminative);
applications in networks, NLP, bioinformatics, ...

Applications

- Natural language processing, Collective Classification, Social Networks, Activity Recognition, ...

Alchemy: Open Source AI

Tutorial

Mailing Lists

[Alchemy](#)

[Alchemy-announce](#)

[Alchemy-update](#)

[Alchemy-discuss](#)

Repositories

[Code](#)

[Datasets](#)

[MLNs](#)

[Publications](#)

Related Links

Welcome to the Alchemy system! Alchemy is a software package providing a series of algorithms for statistical relational learning and probabilistic logic inference, based on the Markov logic representation. Alchemy allows you to easily develop a wide range of AI applications, including:

- Collective classification
- Link prediction
- Entity resolution
- Social network modeling
- Information extraction

Choose a version of Alchemy:

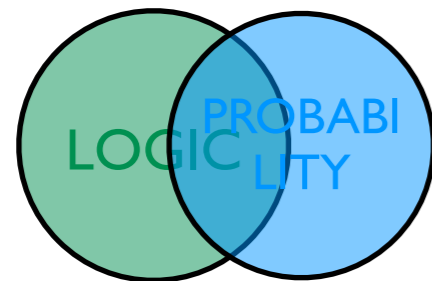
[Alchemy Lite](#)

Alchemy Lite is a software package for inference in Tractable Markov Logic (TML), the first tractable first-order probabilistic logic. Alchemy Lite allows for fast, exact inference for models formulated in TML. Alchemy Lite can be used in batch or interactive mode.

Why StarAI ?

- Reasoning (Probability + Logic) AND Learning
- SRL : Expressive Probabilistic Graphical Models
 - First order logic results supports entities + relationships + background knowledge — abstraction of multiple entities
 - Recursion (e.g. smokers cannot be represented by a plate model)
- PP : Power of a universal Turing machine = a prog. language
 - you can program in it and have builtin expressive prob. models
 - PP can learn -> so bring learning to programming languages
- ProbLog fits both paradigms

Inference



Inference / Reasoning

- Most of the work in PP and StarAI is on inference
 - It is hard (complexity wise)
 - Many inference methods
 - exact, approximate, sampling and lifted ...
- Inference is the key to learning

Two Steps

- **Logical inference** -
 - about a ground logical theory
 - proofs or model theoretic ...
 - *Result: Weighted Model Counting problem*
- **Probabilistic propositional inference** —
 - Knowledge Compilation
 - Backtracking search — DPLL, VE, RC based
- **Advanced** — lifted inference

ProbLog Inference

Answering a query in a ProbLog program happens in four steps

1. Grounding the program w.r.t. the query
- 2. Rewrite the ground logic program into a propositional logic formula**
3. Compile the formula into an arithmetic circuit
4. Evaluate the arithmetic circuit

0.1 :: burglary.

0.5 :: hears_alarm(mary).

0.2 :: earthquake.

0.4 :: hears_alarm(john).

alarm :- earthquake.

alarm :- burglary.

calls(mary) :- alarm, hears_alarm(mary).

calls(john) :- alarm, hears_alarm(john).

calls(mary)

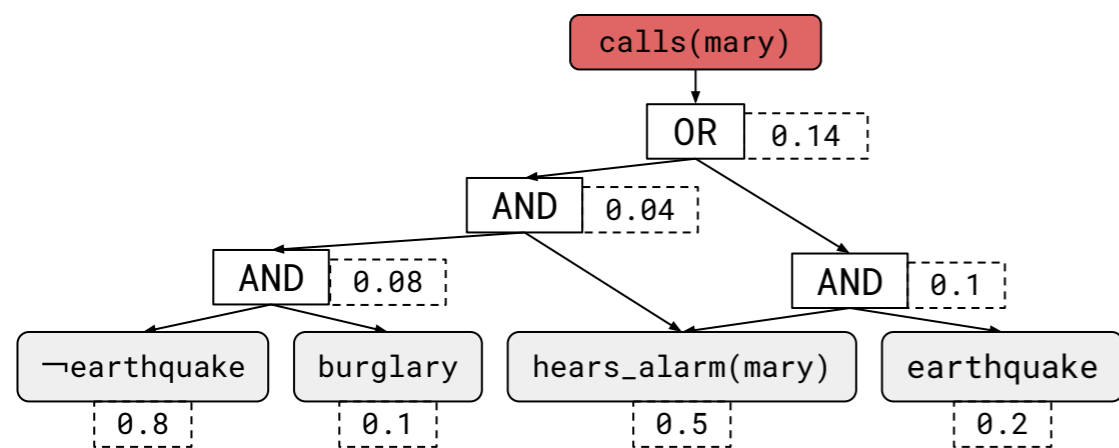
\Leftrightarrow

$\text{hears_alarm(mary)} \wedge (\text{burglary} \vee \text{earthquake})$

ProbLog Inference

Answering a query in a ProbLog program happens in four steps

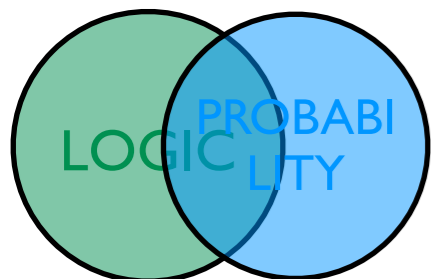
1. Grounding the program w.r.t. the query
2. Rewrite the ground logic program into a propositional logic formula
3. **Compile the formula into an arithmetic circuit (knowledge compilation)**
4. Evaluate the arithmetic circuit



`calls(mary)`

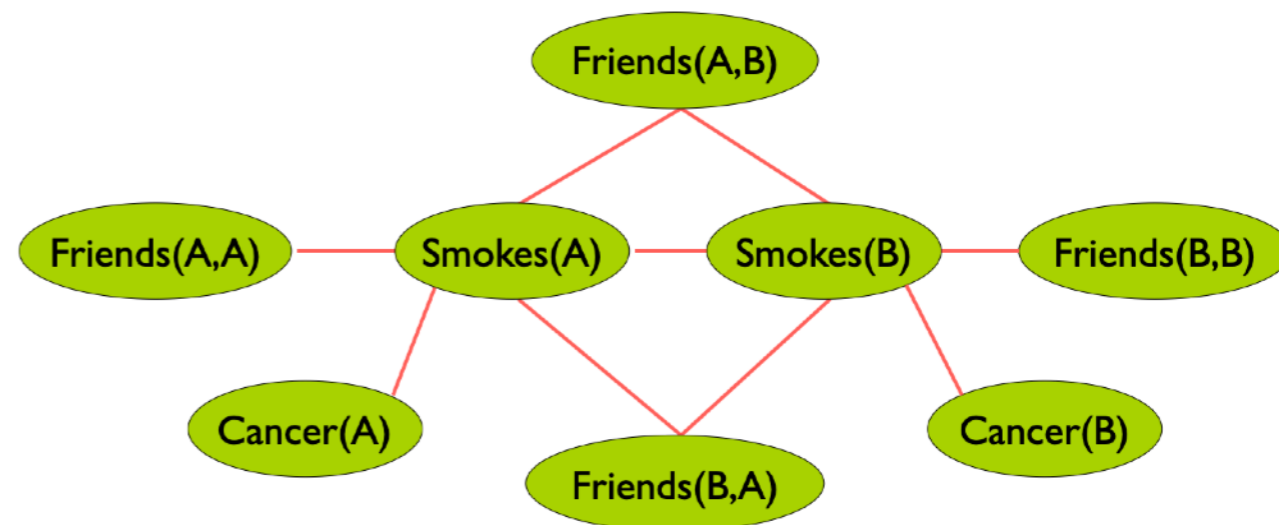
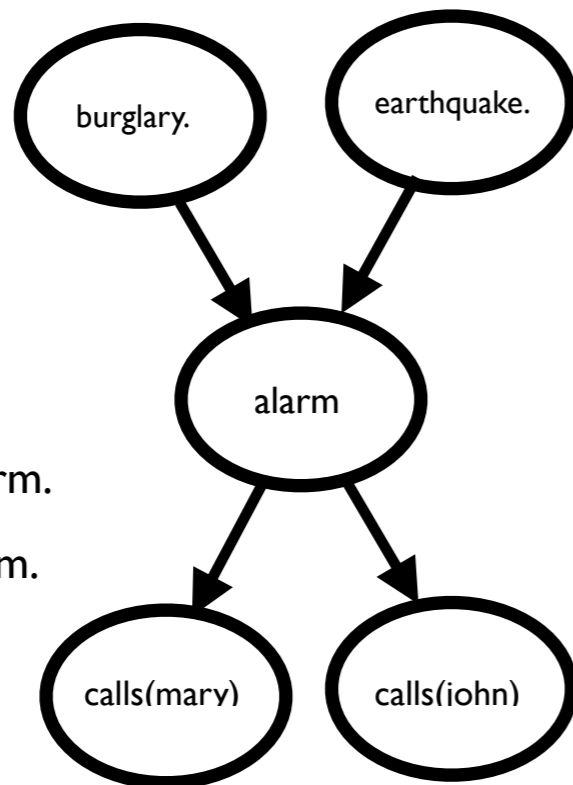
\Leftrightarrow

`hears_alarm(mary) \wedge (burglary \vee earthquake)`



2. Directed vs Undirected the PGM / StarAI dimension

0.1 :: burglary.
 0.05 :: earthquake.
 alarm :- earthquake.
 alarm :- burglary.
 0.7::calls(mary) :- alarm.
 0.6::calls(john) :- alarm.



$$1.5 \quad \forall x \text{ Smokes}(x) \Rightarrow \text{Cancer}(x)$$

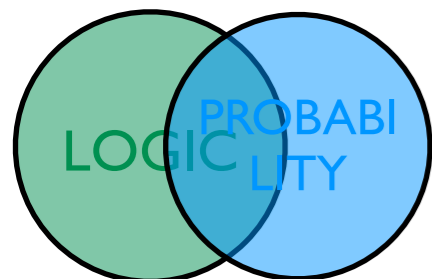
$$1.1 \quad \forall x, y \text{ Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$$

**Probabilistic Logic Programs
 ProbLog**

**directed
 Bayesian Net**

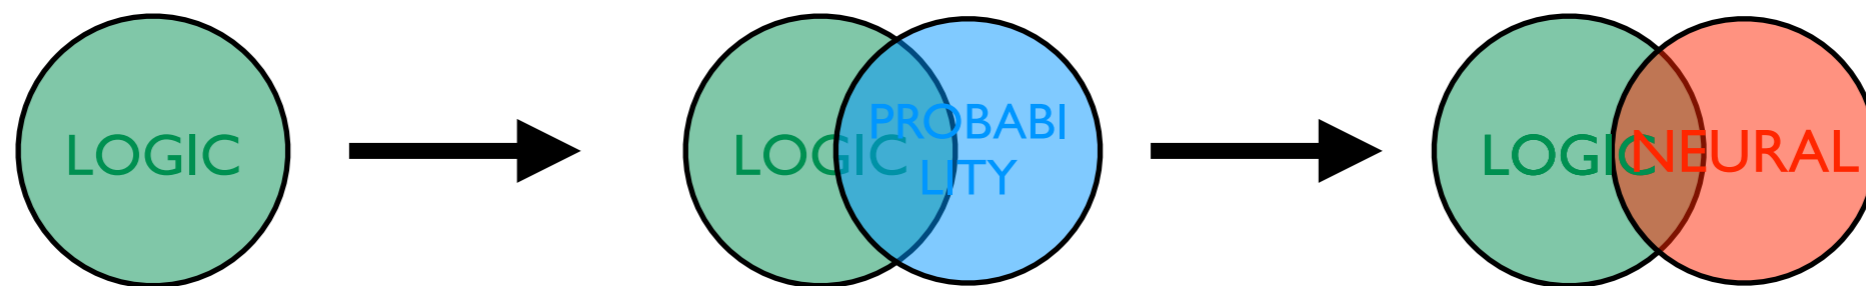
Markov Logic

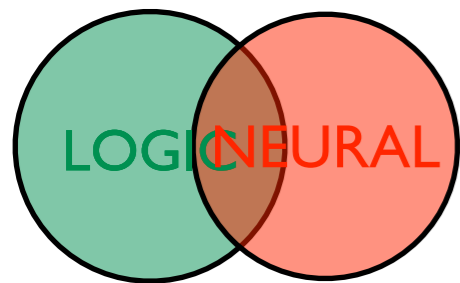
**undirected
 Markov Net
 model theoretic**



key representatives

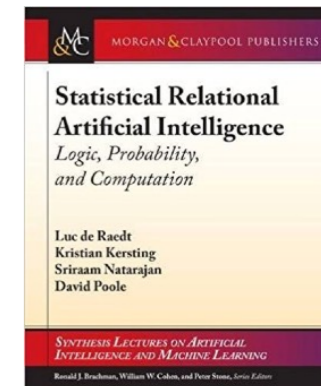
1. Proof vs Model based
2. Directed vs Undirected





2. Directed vs Undirected the NeSy dimension

Two types of Neural Symbolic Systems



Just like in StarAI

Logic as a kind of *neural program*

directed StarAI approach and logic programs

Logic as the *regularizer* (reminiscent of Markov Logic Networks)

undirected StarAI approach and (soft) constraints

Also, many NeSy systems are doing *knowledge based model construction KBMC* where logic is used as a template

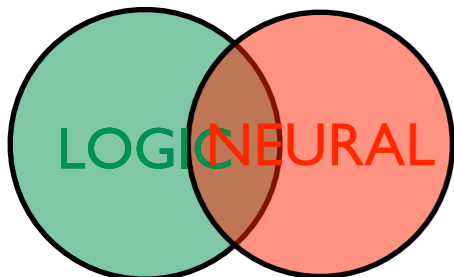
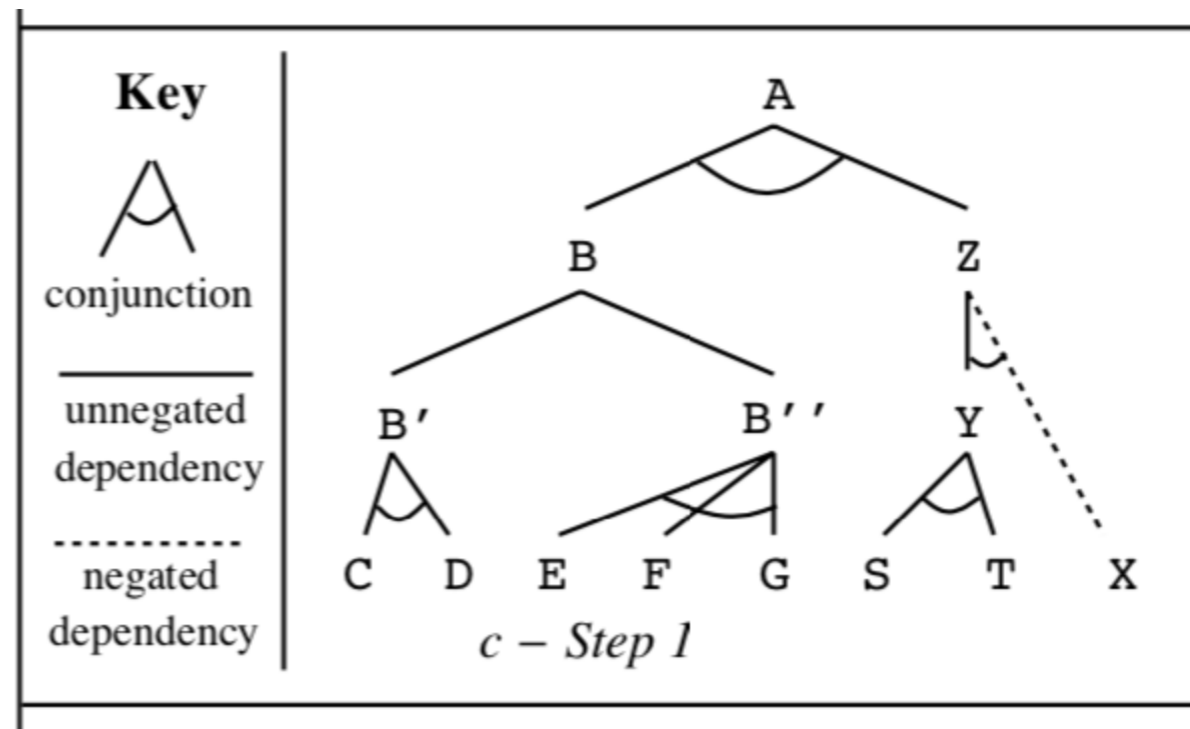
Just like in StarAI

Logic as a neural program

directed StarAI approach and logic programs

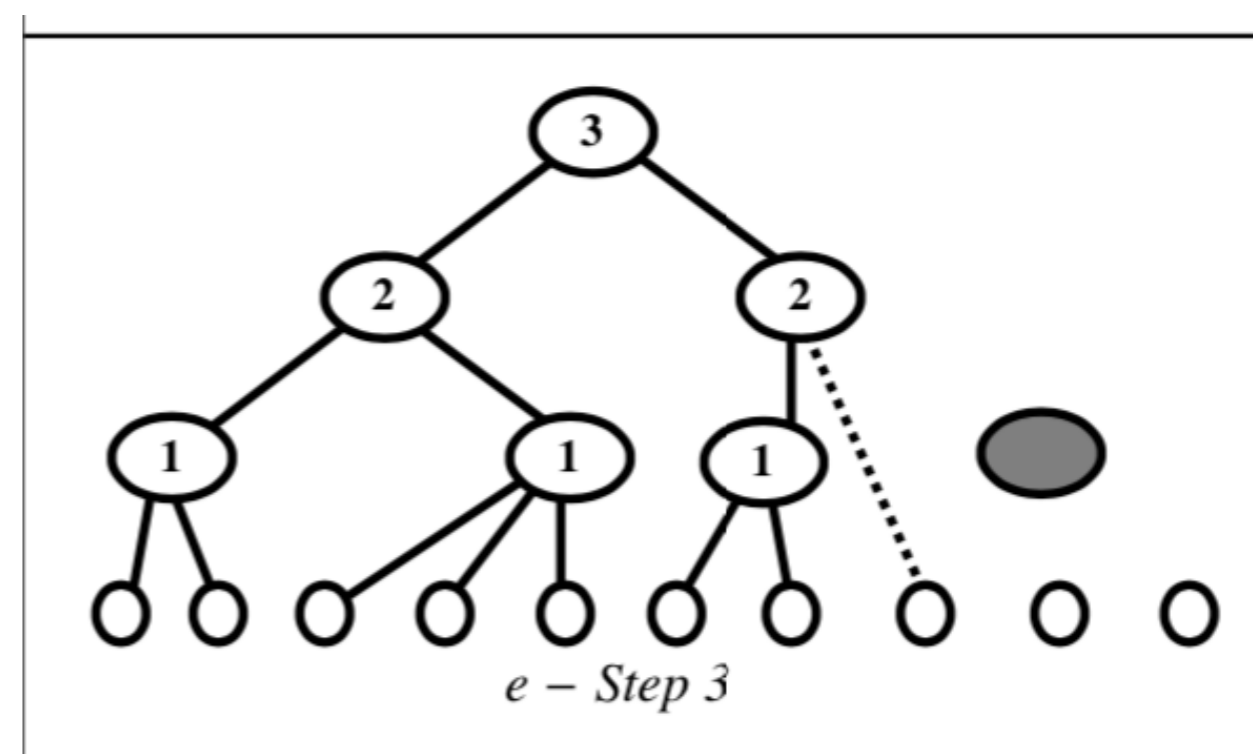
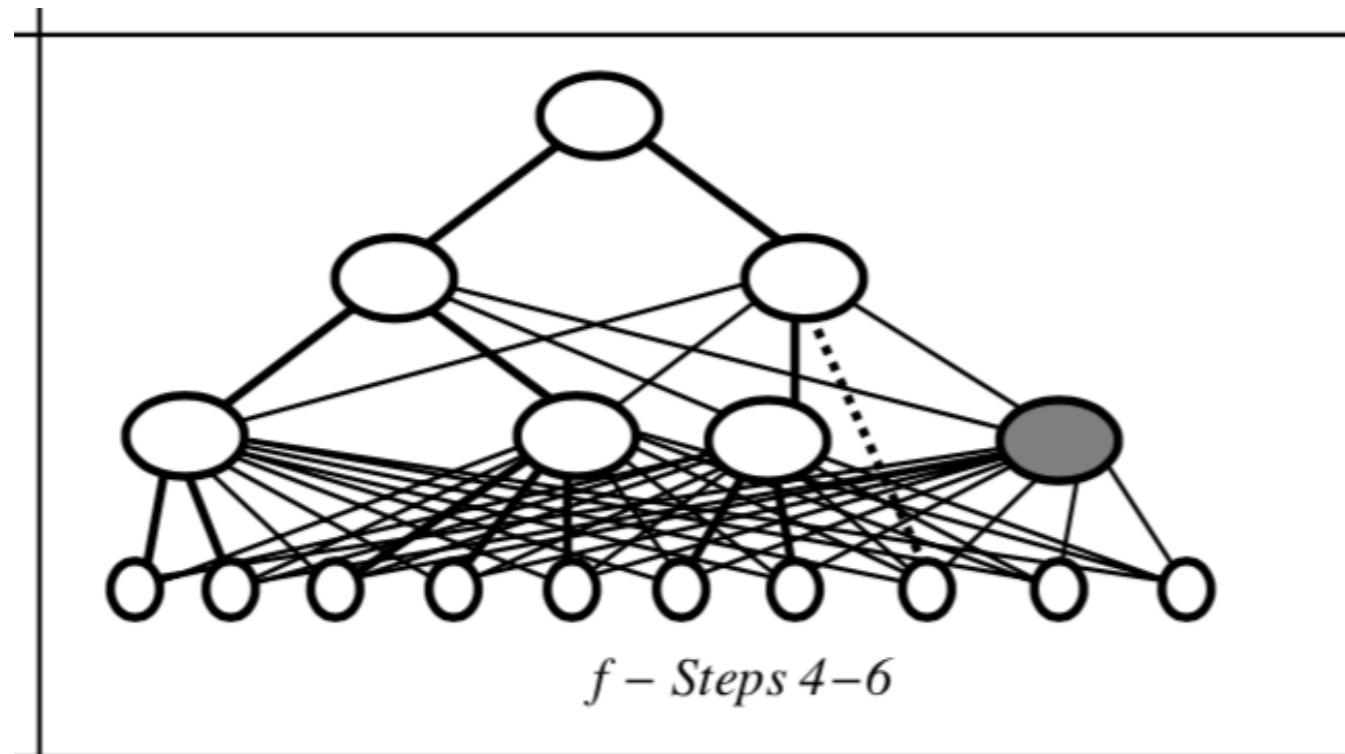
- KBANN (Towell and Shavlik AIJ 94)
- Turn a (propositional) Prolog program into a neural network and learn

A :- B, Z. REWRITE	A :- B, Z.
B :- C, D.	B :- B'.
B :- E, F, G.	B :- B''.
Z :- Y, not X.	B' :- C, D.
Y :- S, T.	B'' :- E, F, G.
	Z :- Y, not X.
	Y :- S, T.



Logic as a neural program

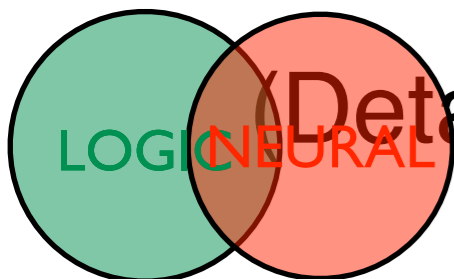
directed StarAI approach and logic programs



ADD LINKS — ALSO SPURIOUS ONES

HIDDEN UNIT

and then learn

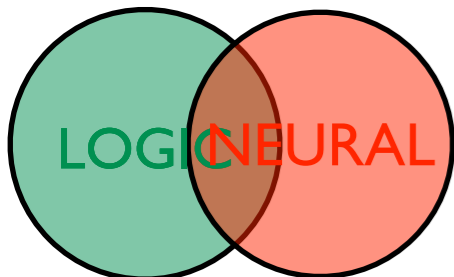
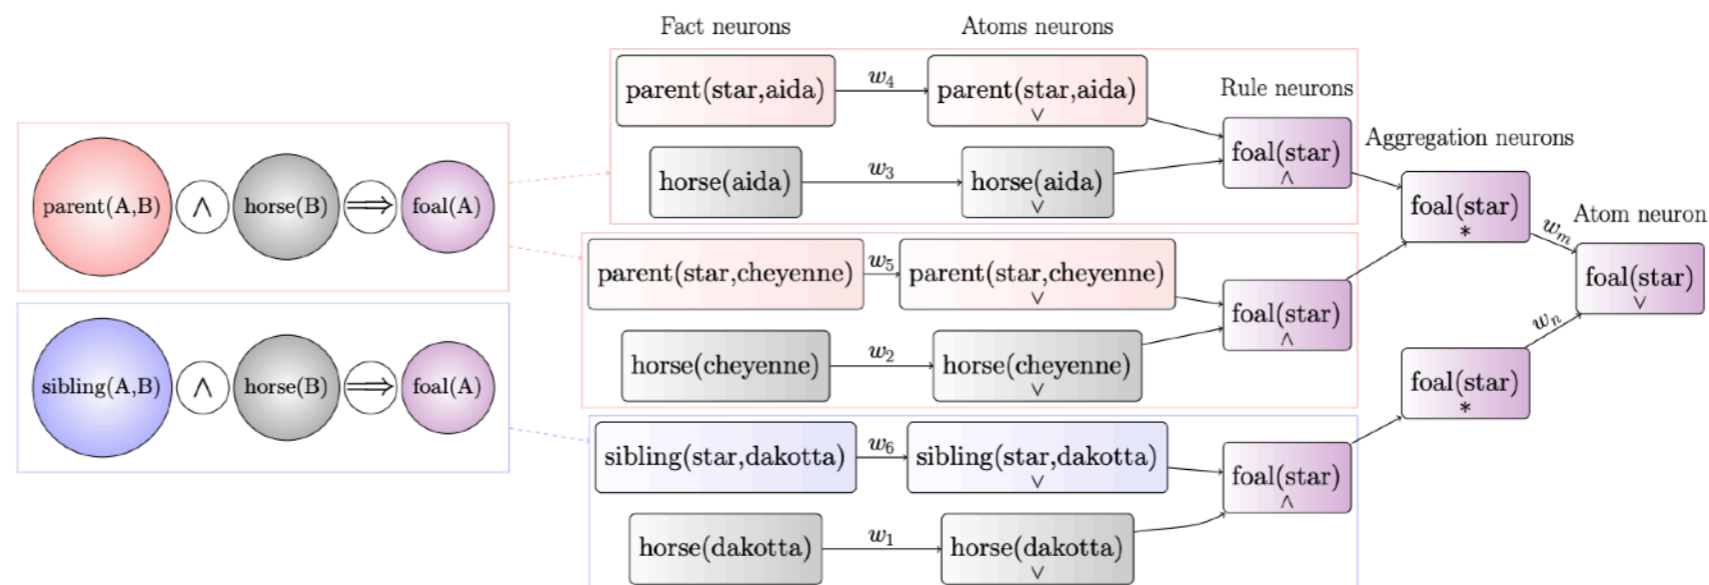


(Details of activation & loss functions not mentioned)

Lifted Relational Neural Networks

directed StarAI approach and logic programs

- Directed (fuzzy) NeSy
- similar in spirit to the Bayesian Logic Programs and Probabilistic Relational Models
- Of course, other kind of (fuzzy) operations for AND, OR and Aggregation (cf. later)



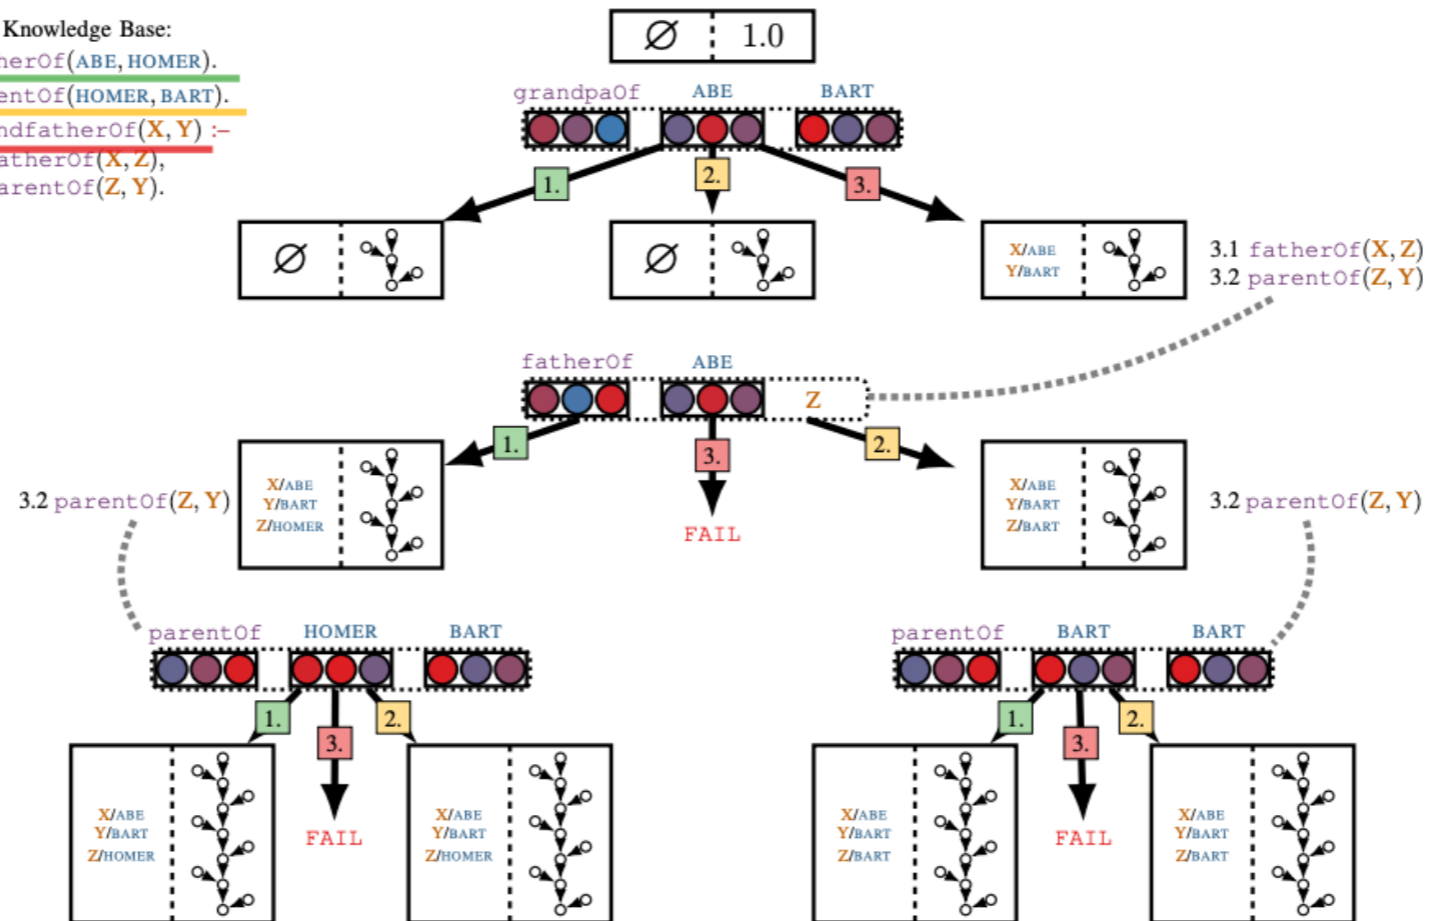
Neural Theorem Prover

directed StarAI approach and logic programs

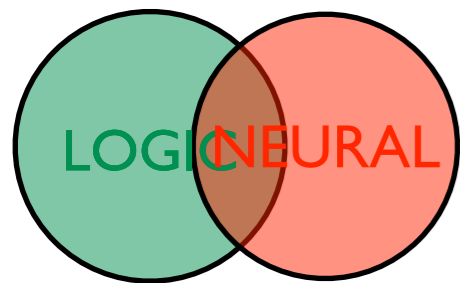
Towards Neural Theorem Proving at Scale

Example Knowledge Base:

1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
fatherOf(X, Z),
parentOf(Z, Y).`

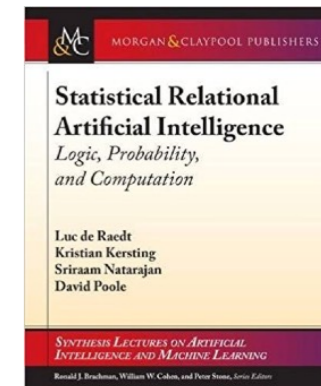


the logic is encoded in the network
how to reason logically ?



2. Directed vs Undirected the NeSy dimension

Two types of Neural Symbolic Systems



Just like in StarAI

Logic as a kind of *neural program*

directed StarAI approach and logic programs

Logic as the *regularizer* (reminiscent of Markov Logic Networks)

undirected StarAI approach and (soft) constraints

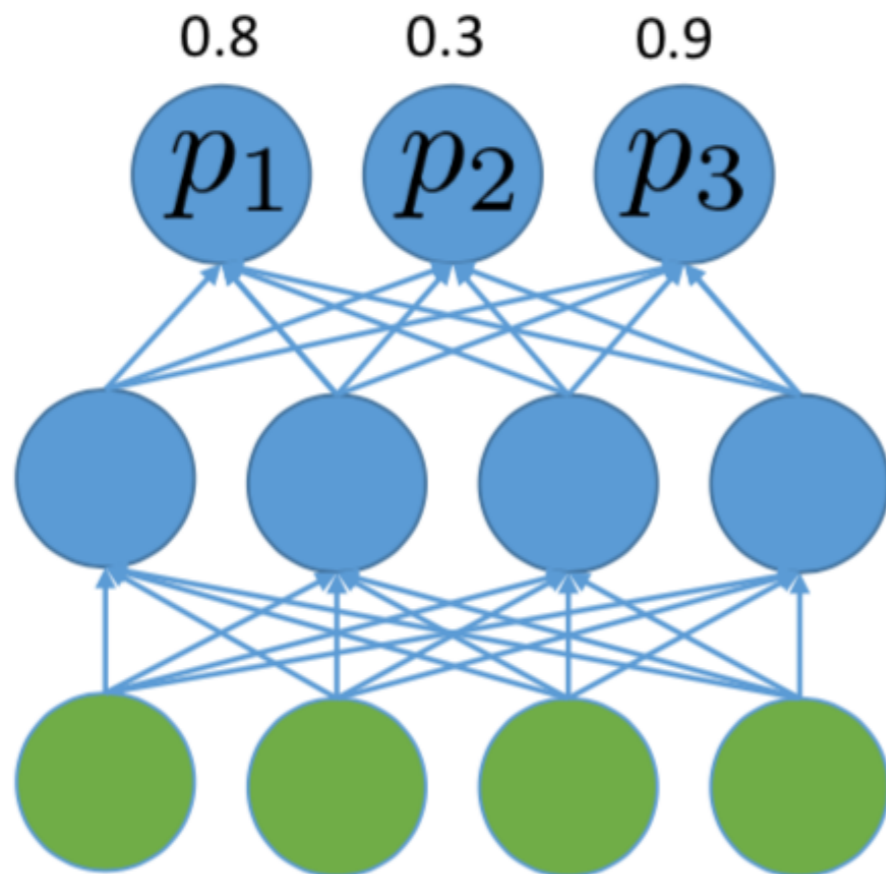
Also, many NeSy systems are doing *knowledge based model construction KBMC* where logic is used as a template

Just like in StarAI

Logic as constraints

undirected StarAI approach and (soft) constraints

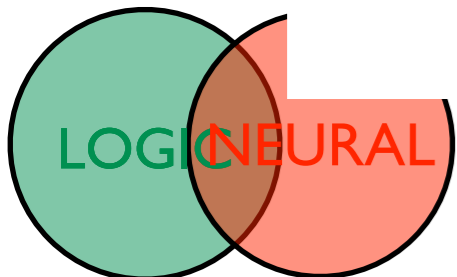
multi-class classification



This constraint should be satisfied

$$\begin{aligned} &(\neg x_1 \wedge \neg x_2 \wedge x_3) \vee \\ &(\neg x_1 \wedge x_2 \wedge \neg x_3) \vee \\ &(x_1 \wedge \neg x_2 \wedge \neg x_3) \end{aligned}$$

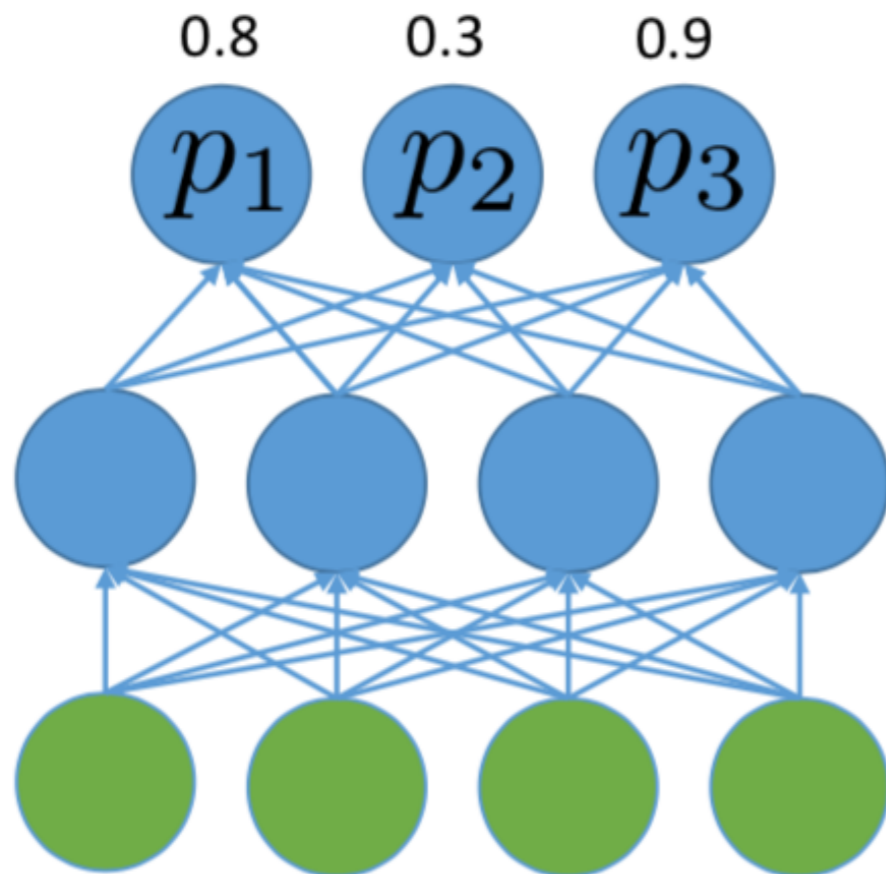
from Xu et al., ICML 2018



Logic as constraints

undirected StarAI approach and (soft) constraints

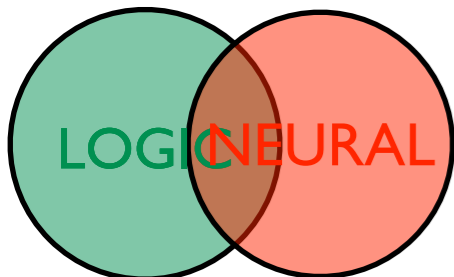
multi-class classification



Probability that constraint is satisfied

$$(1 - x_1)(1 - x_2)x_3 + (1 - x_1)x_2(1 - x_3) + x_1(1 - x_2)(1 - x_3)$$

basis for SEMANTIC LOSS
(weighted model counting)



Logic as a regularizer

undirected StarAI approach and (soft) constraints

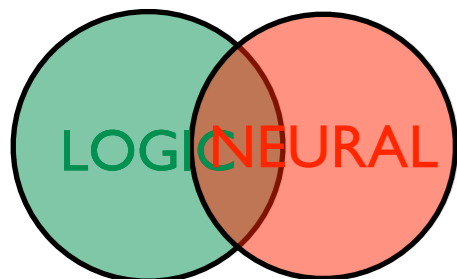
Semantic Loss:

- Use logic as constraints (very much like “propositional MLNs)

- Semantic loss $SLoss(T) \propto -\log \sum_{X \models T} \prod_{x \in X} p_i \prod_{\neg x \in X} (1 - p_i)$

- Used as regulariser $Loss = TraditionalLoss + w.SLoss$

- Use weighted model counting , close to StarAI



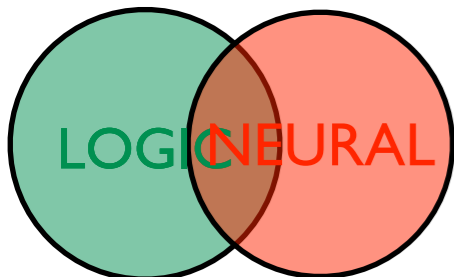
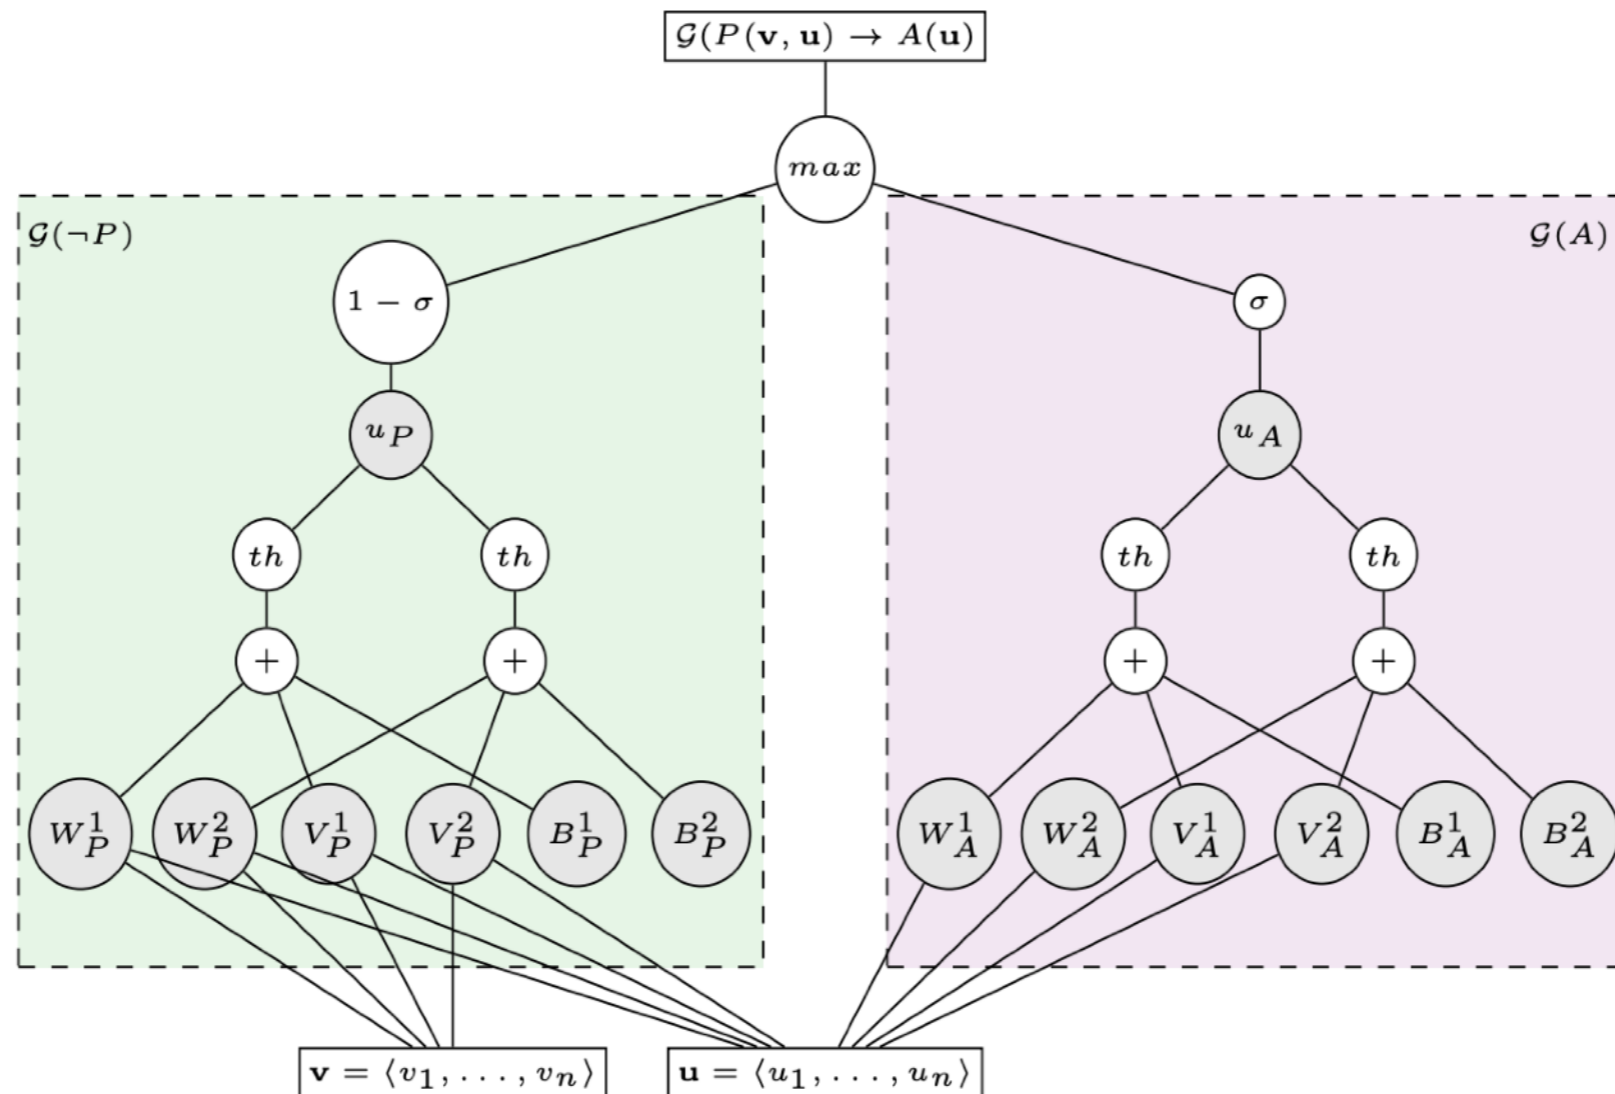
Logic as a regularizer

- Semantic Loss can be used with any logical constraint theory
- Examples with semi-supervised learning, where the constraint enforces that each example should have a class
- very nice properties :
 - differentiable, also monotonicity
 - if $\alpha \models \beta$ then $SLoss(\alpha) \geq SLoss(\beta)$

Logic Tensor Networks

undirected StarAI approach and (soft) constraints

$$P(x, y) \rightarrow A(y), \text{ with } \mathcal{G}(x) = \mathbf{v} \text{ and } \mathcal{G}(y) = \mathbf{u}$$



Semantic Based Regularization

undirected StarAI approach and (soft) constraints

$$F := \forall d P_A(d) \Rightarrow A(d)$$

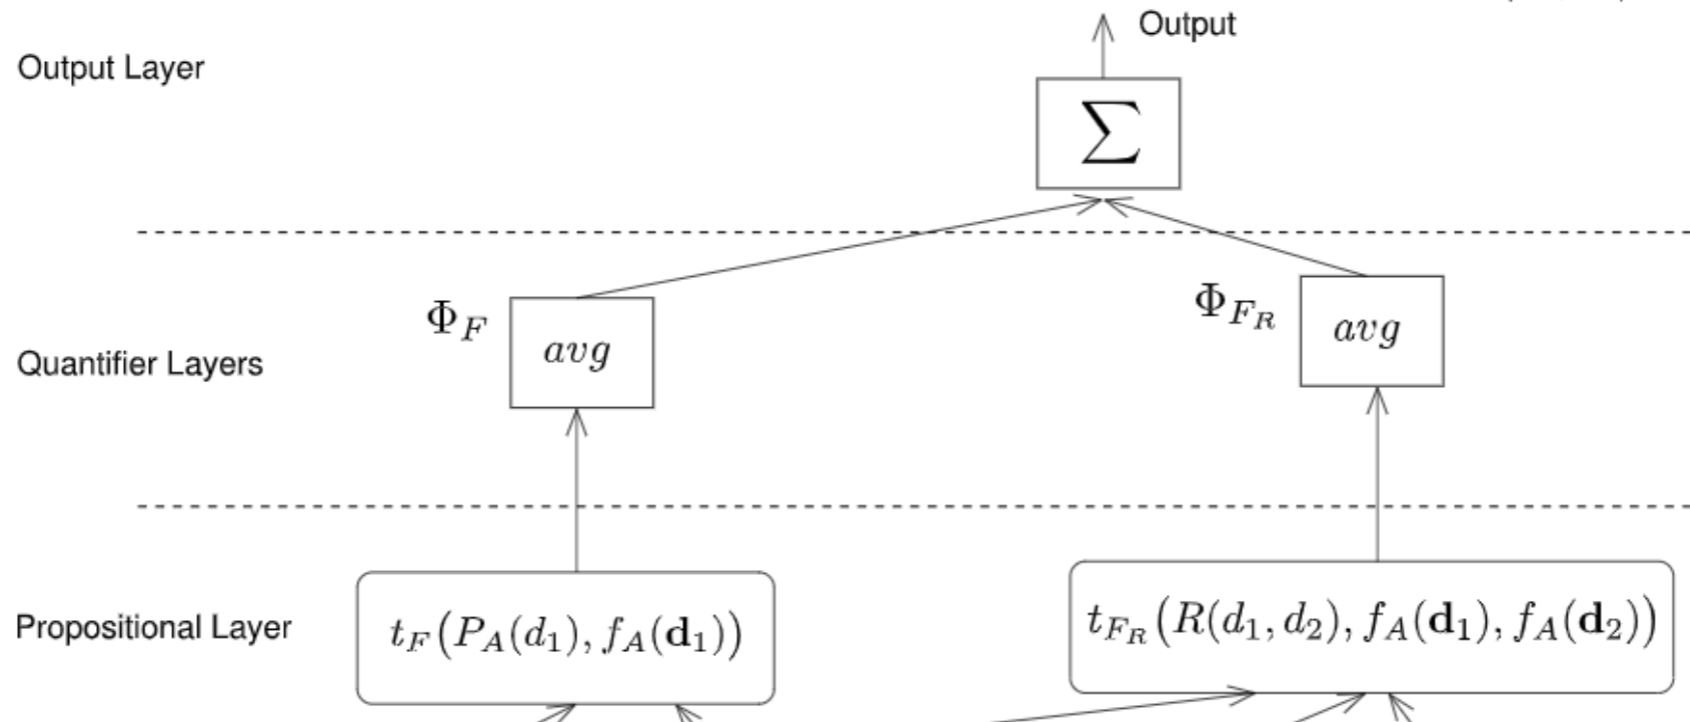
$$F_R := \forall d \forall d' R(d, d') \Rightarrow ((A(d) \wedge A(d')) \vee (\neg A(d) \wedge \neg A(d')))$$

$$C = \{d_1, d_2\}$$

Evidence Predicate
Groundings

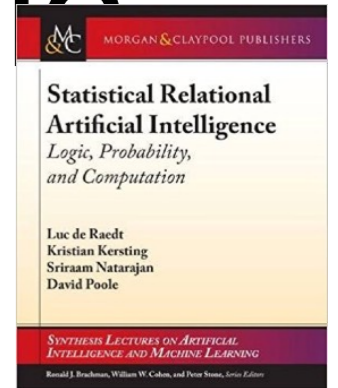
$$P_A(d_1) = 1$$

$$R(d_1, d_2) = 1$$



the logic is encoded in the network
how to reason logically ?

Two types of Neural Symbolic Systems



Just like in StarAI

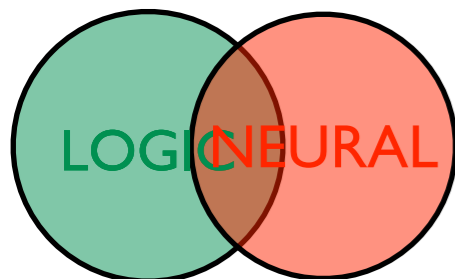
Logic as a kind of *neural program*

directed StarAI approach and logic programs

Logic as the *regularizer* (reminiscent of Markov Logic Networks)

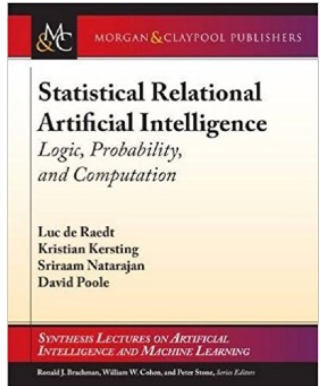
undirected StarAI approach and (soft) constraints

Consequence :
the logic is encoded in the network
the ability to logically reason is lost
logic is not a special case



2. Directed vs Undirected the NeSy dimension

Two types of Neural Symbolic Systems



Just like in StarAI

Logic as a kind of *neural program*

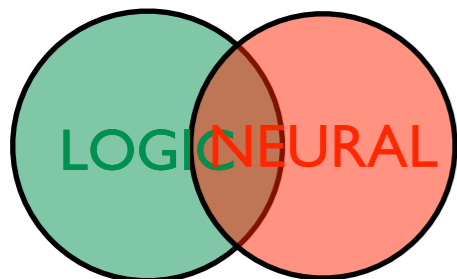
directed StarAI approach and
logic programs

Logic as the *regularizer*
(reminiscent of Markov Logic
Networks)

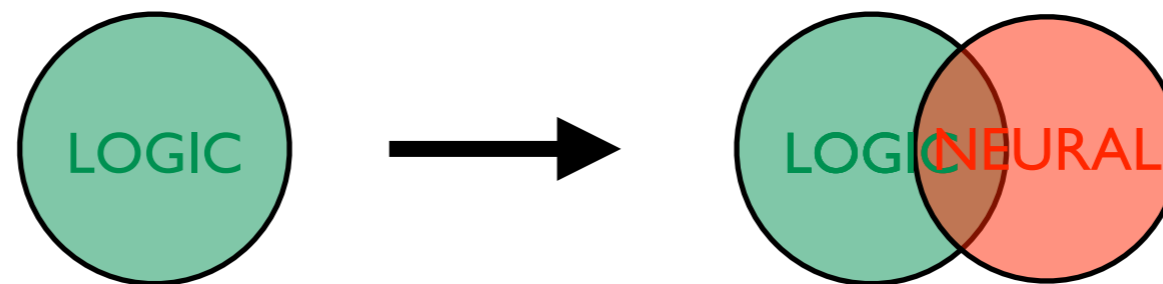
undirected StarAI approach and
(soft) constraints

Also, many NeSy systems are doing
knowledge based model construction KBMC
where logic is used as a template

Just like in StarAI



3. Types of Logic

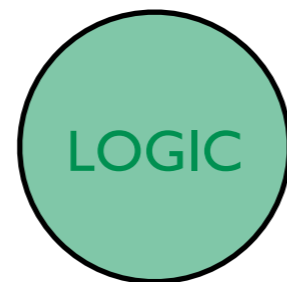


3. Types of Logic

Key Messages

- Different types of logic exist
- Different types of logic enable different functionalities

3. Types of Logic



Various flavours of logic

`alarm :- earthquake.`

`alarm :- burglary.`

`calls_mary :- alarm, hears_alarm_mary.`

`calls_john :- alarm, hears_alarm_john.`

`stress(ann).`

`influences(ann,bob).`

`influences(bob,carl).`

`smokes(X) :- stress(X).`

`smokes(X) :-`

`influences(Y,X),`

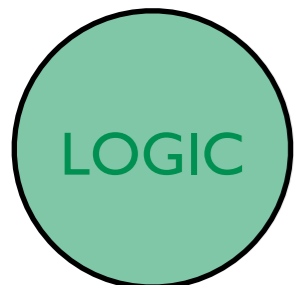
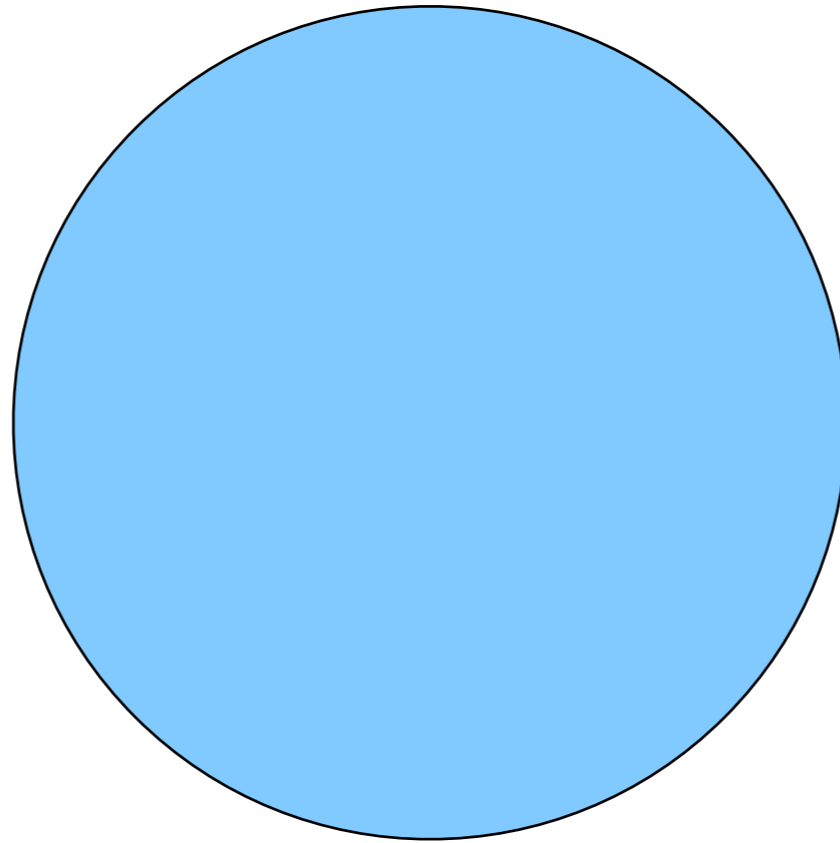
`smokes(Y).`

Propositional logic

First-order logic

Various flavours of first-order logic

Logic programs
= programming language



Logic programming and Prolog

Full-fledged programming language

structured terms

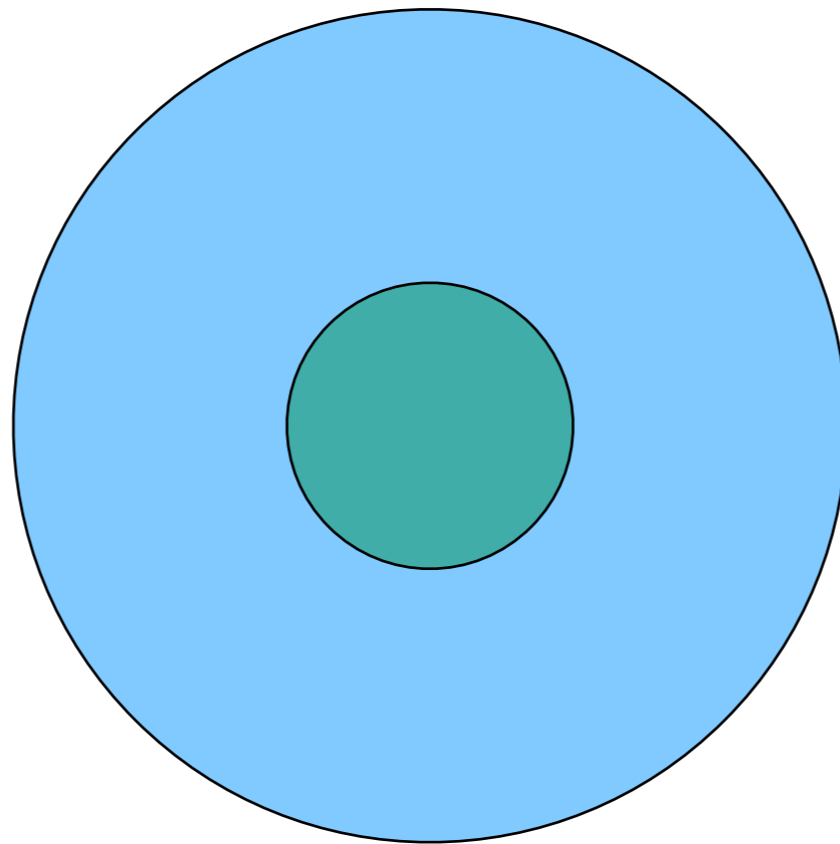
```
member(X, [X|_]).
```

```
member(X, [_|Tail]) :-  
    member(X, Tail).
```

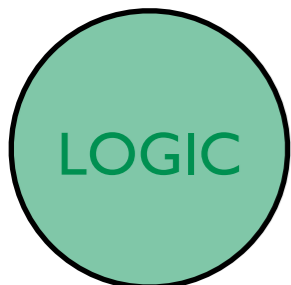
recursion

Various flavours of first-order logic

Logic programs
= programming language



Datalog
= Logic programs
that always terminate



Datalog

Query language for deductive databases

no structured terms

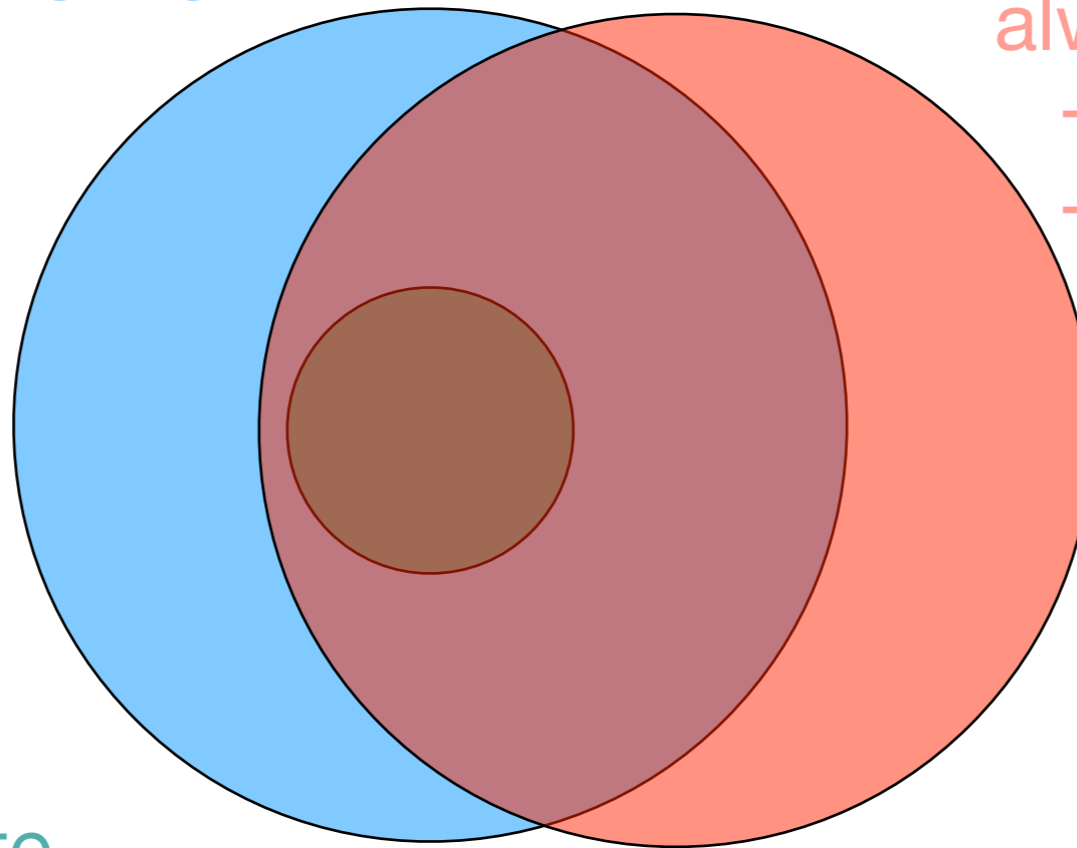
guaranteed to terminate

```
ancestor(X, Y) :- parent(X, Y).  
ancestor(X, Y) :- parent(X, Z), ancestor(Z, Y).
```

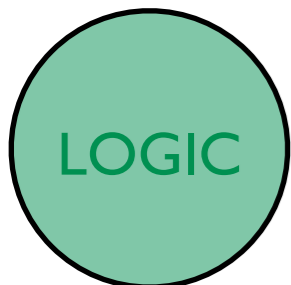
Various flavours of first-order logic

Logic programs
= programming language

Answer-set programs
= Logic programs with
multiple models that
always terminate
+ soft/hard constraints
+ preferences



Datalog
= Logic programs
that always terminate



Answer-set programming

Prolog with multiple models + interesting features

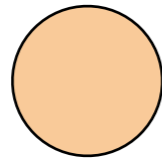
```
col(r). col(g). col(b).
```

```
1 {color(X,C) : col(C)} 1 :- node(X).  
:- edge(X,Y), color(X,C), color(Y,C).
```

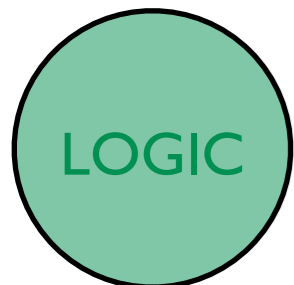
choice rules

constraint

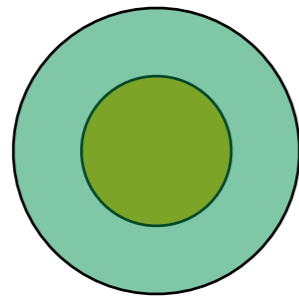
What can it do?



Propositional logic:
simple propositional reasoning

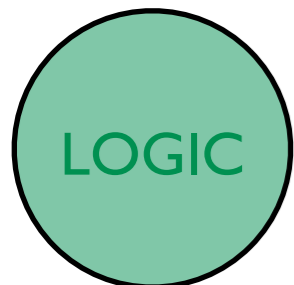


What can it do?

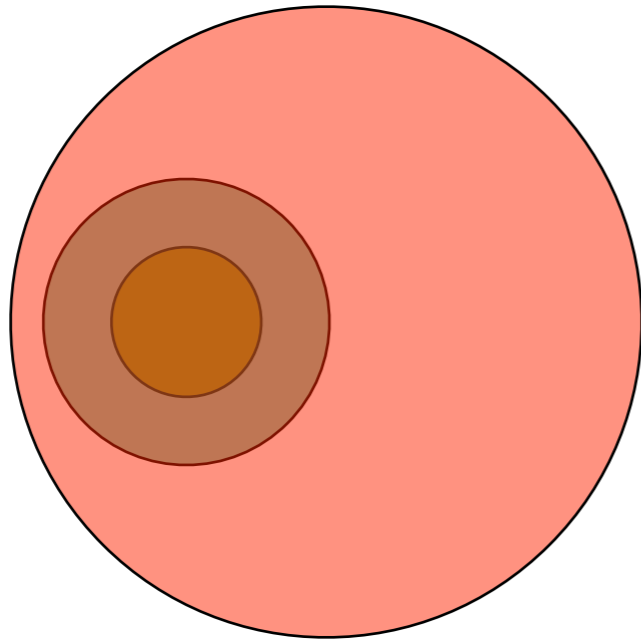


Datalog:
database queries

Propositional logic:
simple propositional reasoning



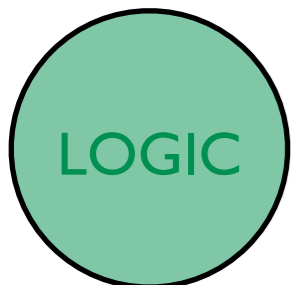
What can it do?



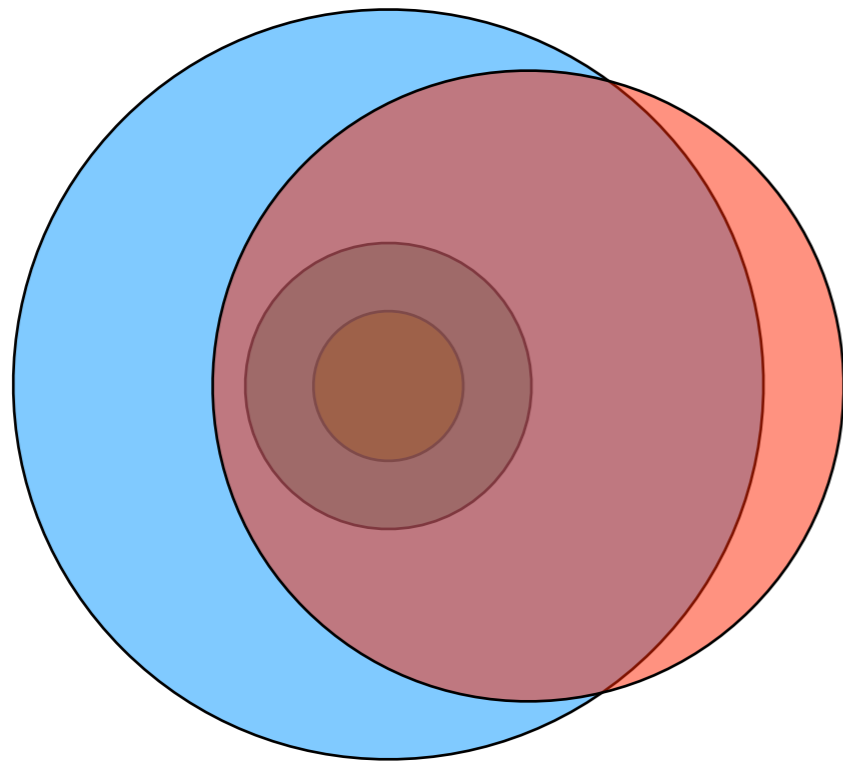
Answer-set programming:
database queries, common-sense
reasoning, preferences

Datalog:
database queries

Propositional logic:
simple propositional reasoning



What can it do?

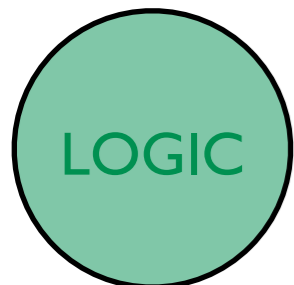


Logic programming:
programs manipulating structured
objects, infinite domains, ...

Answer-set programming:
database queries, common-sense
reasoning, preferences

Datalog:
database queries

Propositional logic:
simple propositional reasoning



Logic program vs First-order logic

Issues with transitive closure in first-order logic

$\text{edge}(1,2).$

$\text{path}(A,B) \leftarrow \text{edge}(A,B).$

$\text{path}(A,B) \leftarrow \text{edge}(A,C), \text{path}(C,B).$

Logic programs always
have one model

$\{\text{edge}(1,2), \text{path}(1,2)\}$

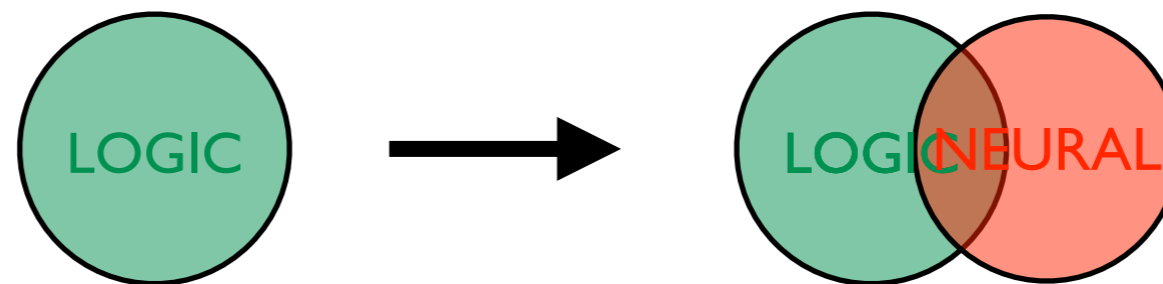
First-order logic can have
many models

$\{\text{edge}(1,2), \text{path}(1,2)\}$

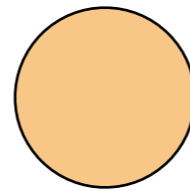
$\{\text{edge}(1,2), \text{path}(1,2), \text{path}(1,1)\}$

$\{\text{edge}(1,2), \text{path}(1,2), \text{path}(2,1)\}$

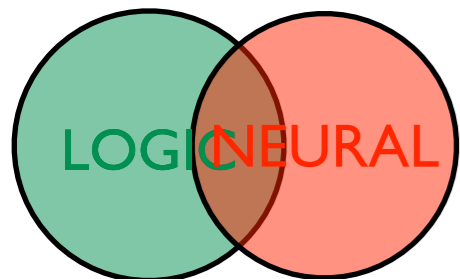
3. Types of Logic



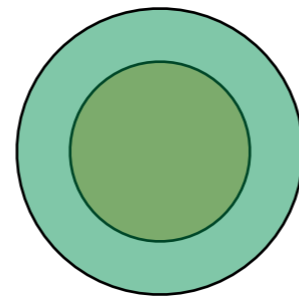
Logic in NeSy - Propositional logic



Semantic loss

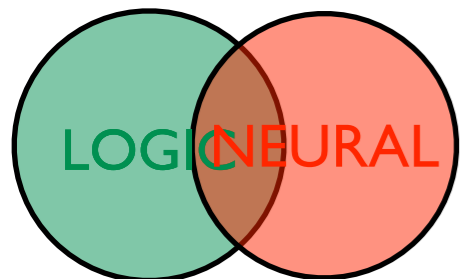


Logic in NeSy - Datalog

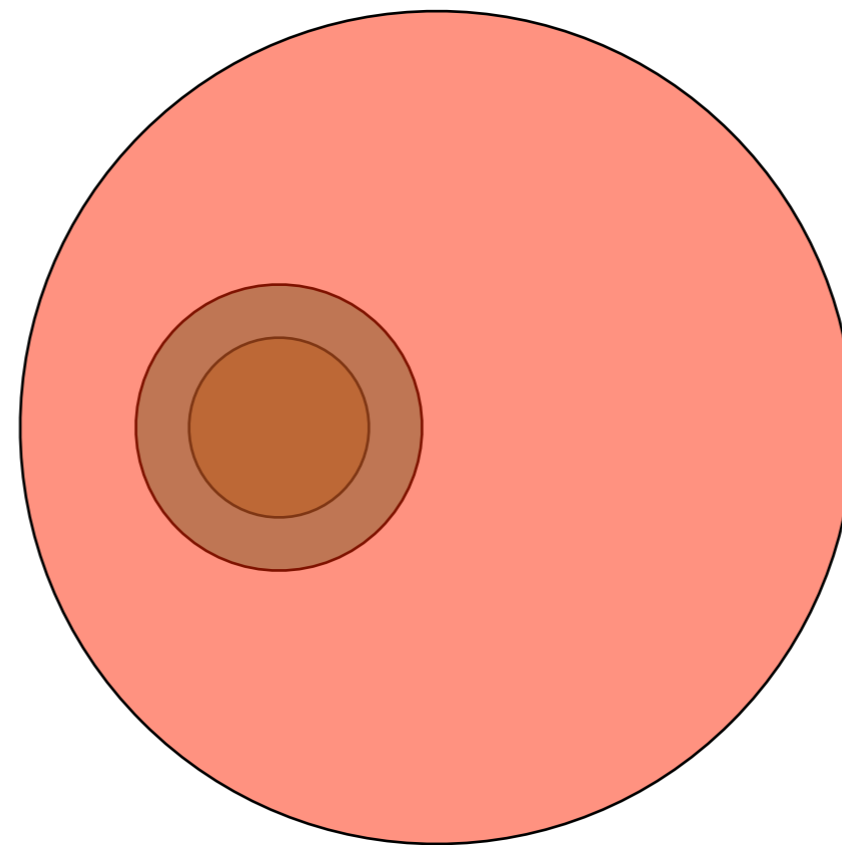


∂ ILP, Neural Theorem
Provers, LRNN, DiffLog, ...

Semantic loss



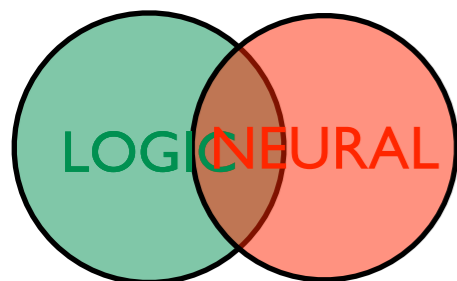
Logic in NeSy - Answer-set programming



NeurASP

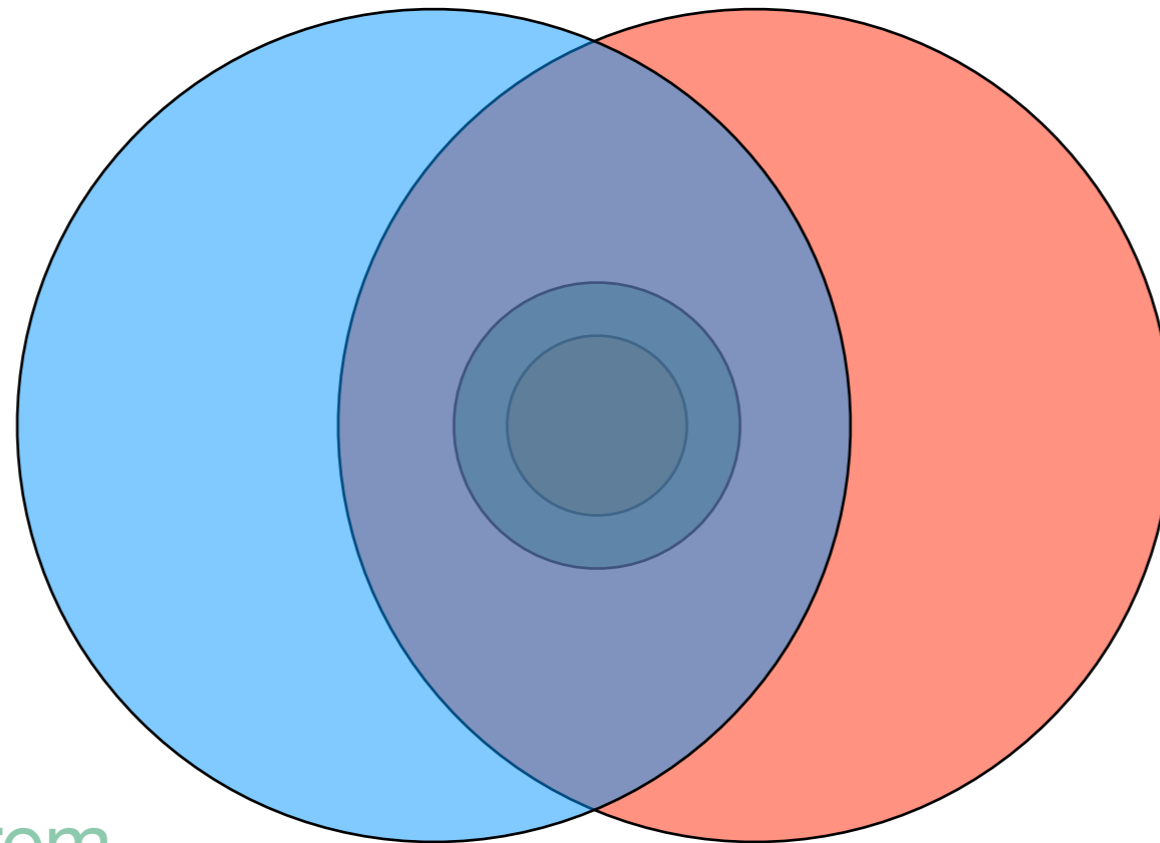
∂ ILP, Neural Theorem Provers, LRNN, DiffLog, ...

Semantic loss



Logic in NeSy - Logic programming

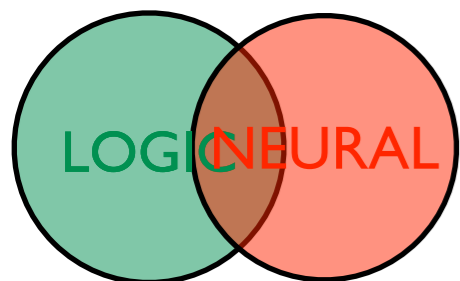
DeepProblog,
NLProlog



NeurASP

∂ ILP, Neural Theorem
Provers, LRNN, DiffLog, ...

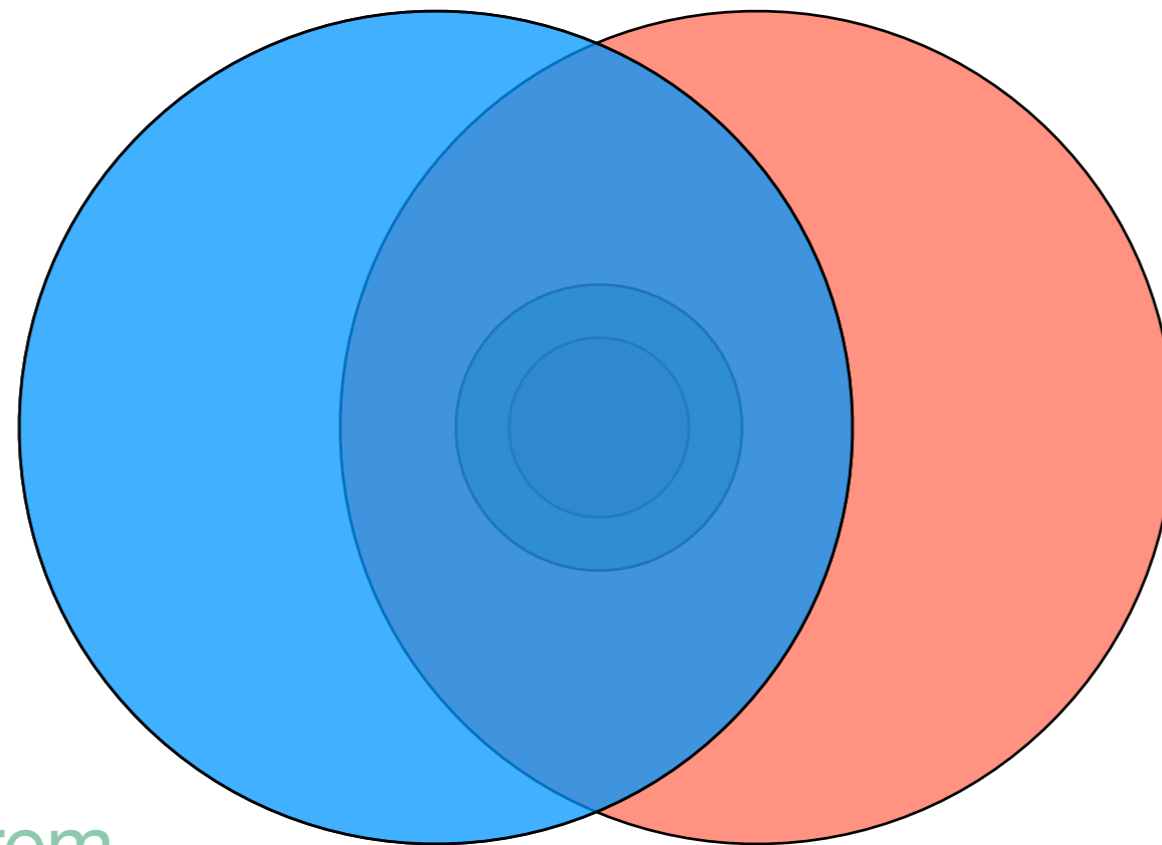
Semantic loss



Logic in NeSy - First-order logic

Logic tensor networks, NMLN,
SBT, RNM

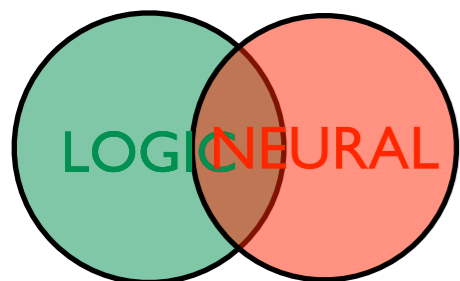
DeepProblog,
NLProlog



NeurASP

∂ ILP, Neural Theorem
Provers, LRNN, DiffLog, ...

Semantic loss



3. Types of Logic

Key Messages

- Different types of logic exist
- Different types of logic enable different functionalities

4. Symbolic vs sub-symbolic

4. Symbolic vs sub-symbolic

Key Messages

- Entities are represented very differently in symbolic and sub-symbolic systems, but they are complementary
- NeSy systems can be categorized by how they use symbolic and sub-symbolic intermediate representations

Symbolic representations

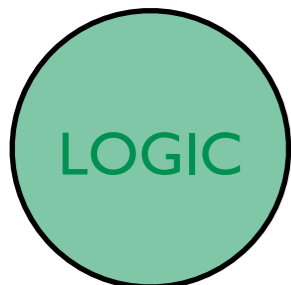
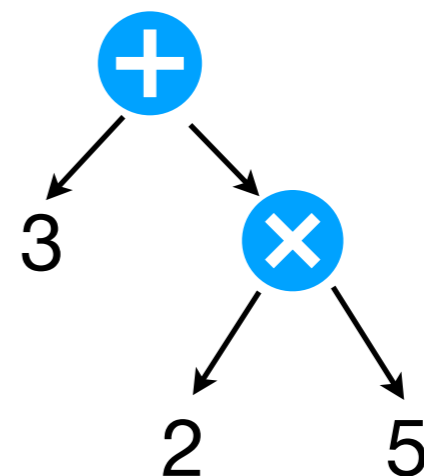
- Atoms: an, bob
- Numbers: 4, -3.5
- Variables: X,Y

- Structured terms: $f(t_1, \dots, t_n)$
 - motherOf(an)
 - [-0.1, 1.2, 0.5]
 - [[1, 2, 3], [4, 5, 6]]
 - plus(3, times(2, 5))
 - ...

an $\xrightarrow{\text{motherOf}}$ bob

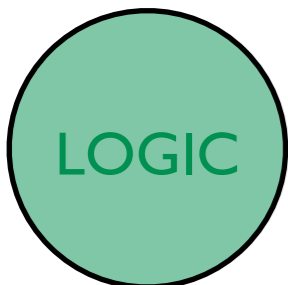
-0,1	1,2	0,5
------	-----	-----

1	2	3
4	5	6



Comparing symbols: unification

- Powerful mechanism for symbol matching
 - basis for many logic-based AI systems
- Finds substitution θ such that both symbols match
 - $\text{mother}(X, \text{bob}) = \text{mother}(\text{an}, Y)$
 - $\theta = \{X = \text{an}, Y = \text{bob}\}$
- Not useful to determine similarity
 - $\text{mother}(\text{an}, \text{bob}) \approx \text{mother}(\text{an}, \text{charlie})?$



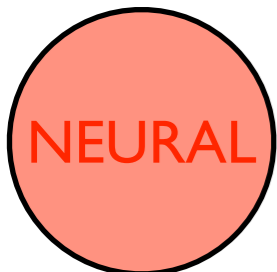
Sub-symbolic representations

- Sub-symbolic systems require numerical representation
- Often, entities are already numerical in nature



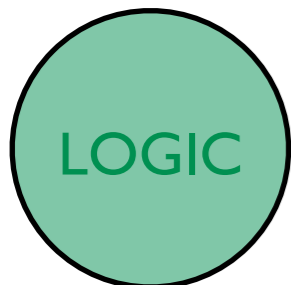
0,1	-0,3	...
-0,9	-0,2	...
...

- Generally, these representations are fixed in size and dimensionality
- Exceptions require special neural architectures, e.g.
 - Recurrent neural networks
 - Fully convolutional networks
 - ...



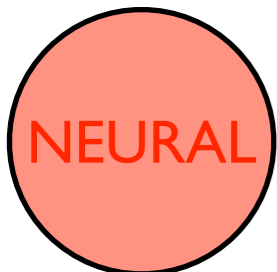
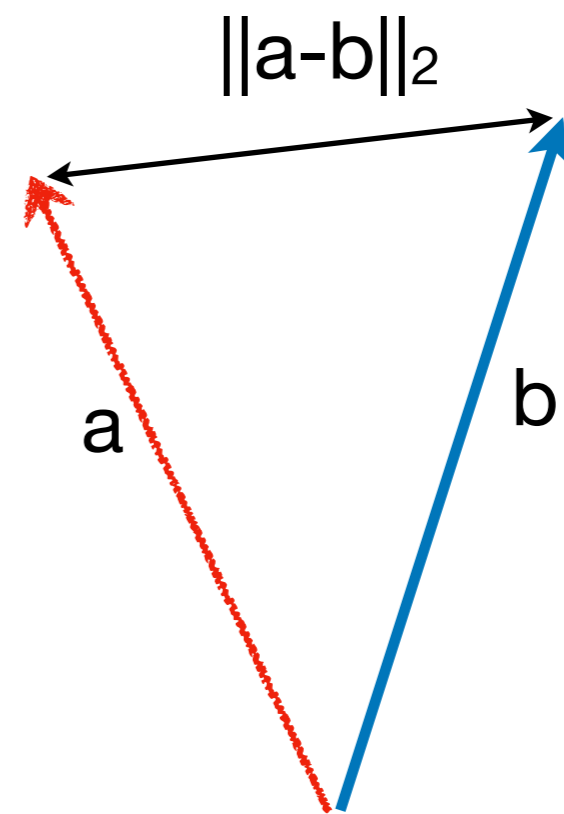
Sub-symbols in StarAI

- It is possible to represent these sub-symbols in logic
 - vectors: $[0.1, -0.5, 0.6]$
 - matrices: $\begin{bmatrix} [0.2, 0.4], \\ [0.3, 0.1] \end{bmatrix}$
 - ...
- However, they are not part of the computation mechanisms.
 - i.e. we cannot learn its parameters
- They are not first class citizens.

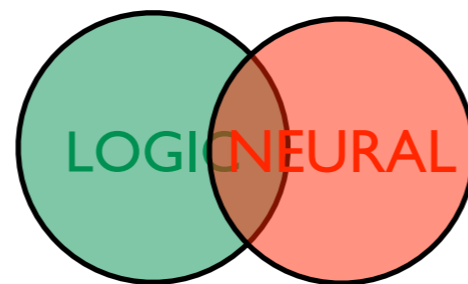


Comparing sub-symbols

- Similarity can be determined through various metrics
 - L1, L2, radial-basis function, ...
- Can only give a degree of similarity
- When is $a \neq b$? When is $a = b$?

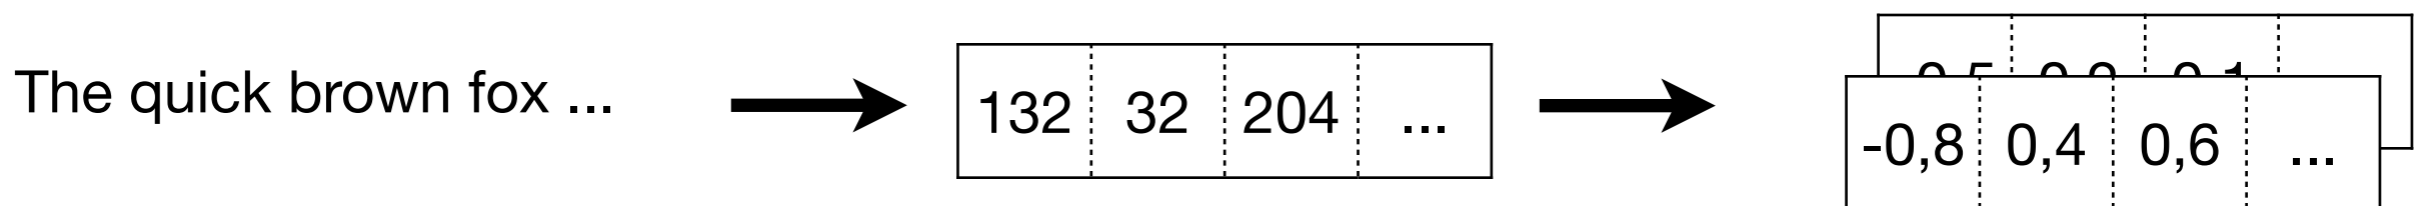


4. Symbolic vs sub-symbolic Translating between representations



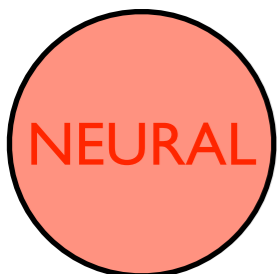
Symbols to sub-symbols

- A lot of deep learning research is on how to represent symbols



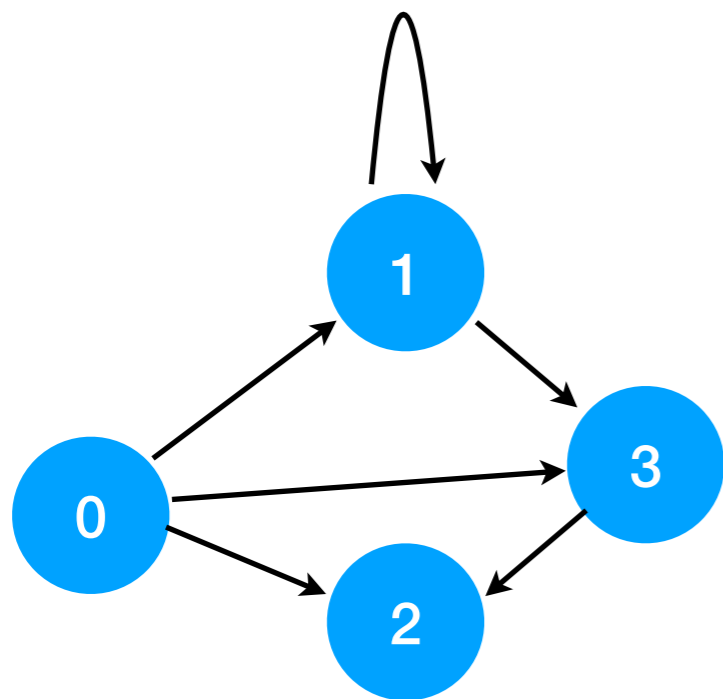
- Encoding relations $r(h,t)$
 - Many ways to structure embedding space

Models	score function $f(\mathbf{h}, \mathbf{r}, \mathbf{t})$
TransE [2]	$- \mathbf{h} + \mathbf{r} - \mathbf{t} _{1/2}$
TransR [10]	$- M_r \mathbf{h} + \mathbf{r} - M_r \mathbf{t} _2^2$
DistMult [20]	$\mathbf{h}^\top \text{diag}(\mathbf{r}) \mathbf{t}$
Complex [16]	$\text{Real}(\mathbf{h}^\top \text{diag}(\mathbf{r}) \bar{\mathbf{t}})$
RESCAL [12]	$\mathbf{h}^\top M_r \mathbf{t}$
RotatE [15]	$- \mathbf{h} \circ \mathbf{r} - \mathbf{t} ^2$

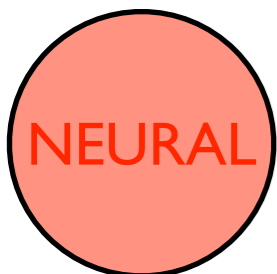
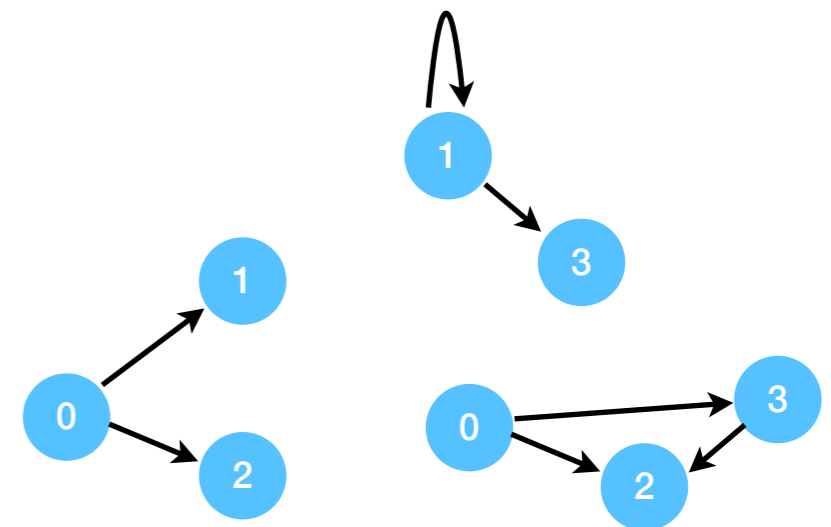
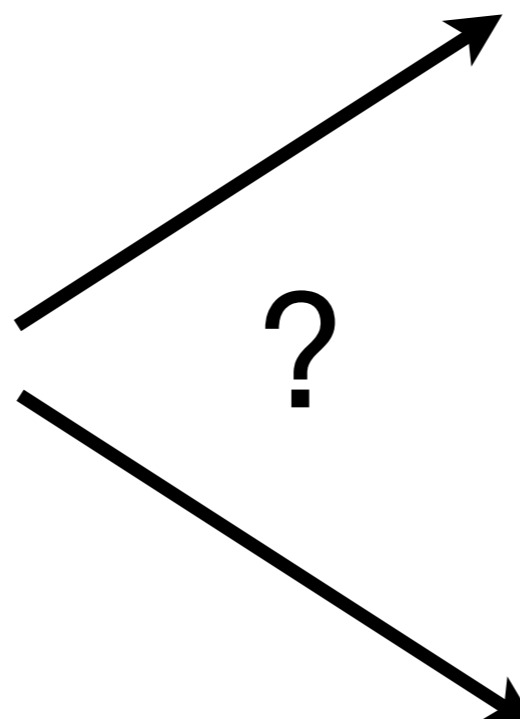


Symbols to sub-symbols

- What about graphs?

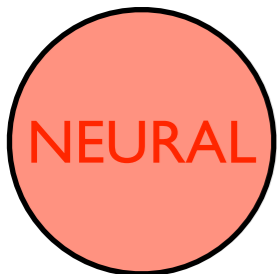
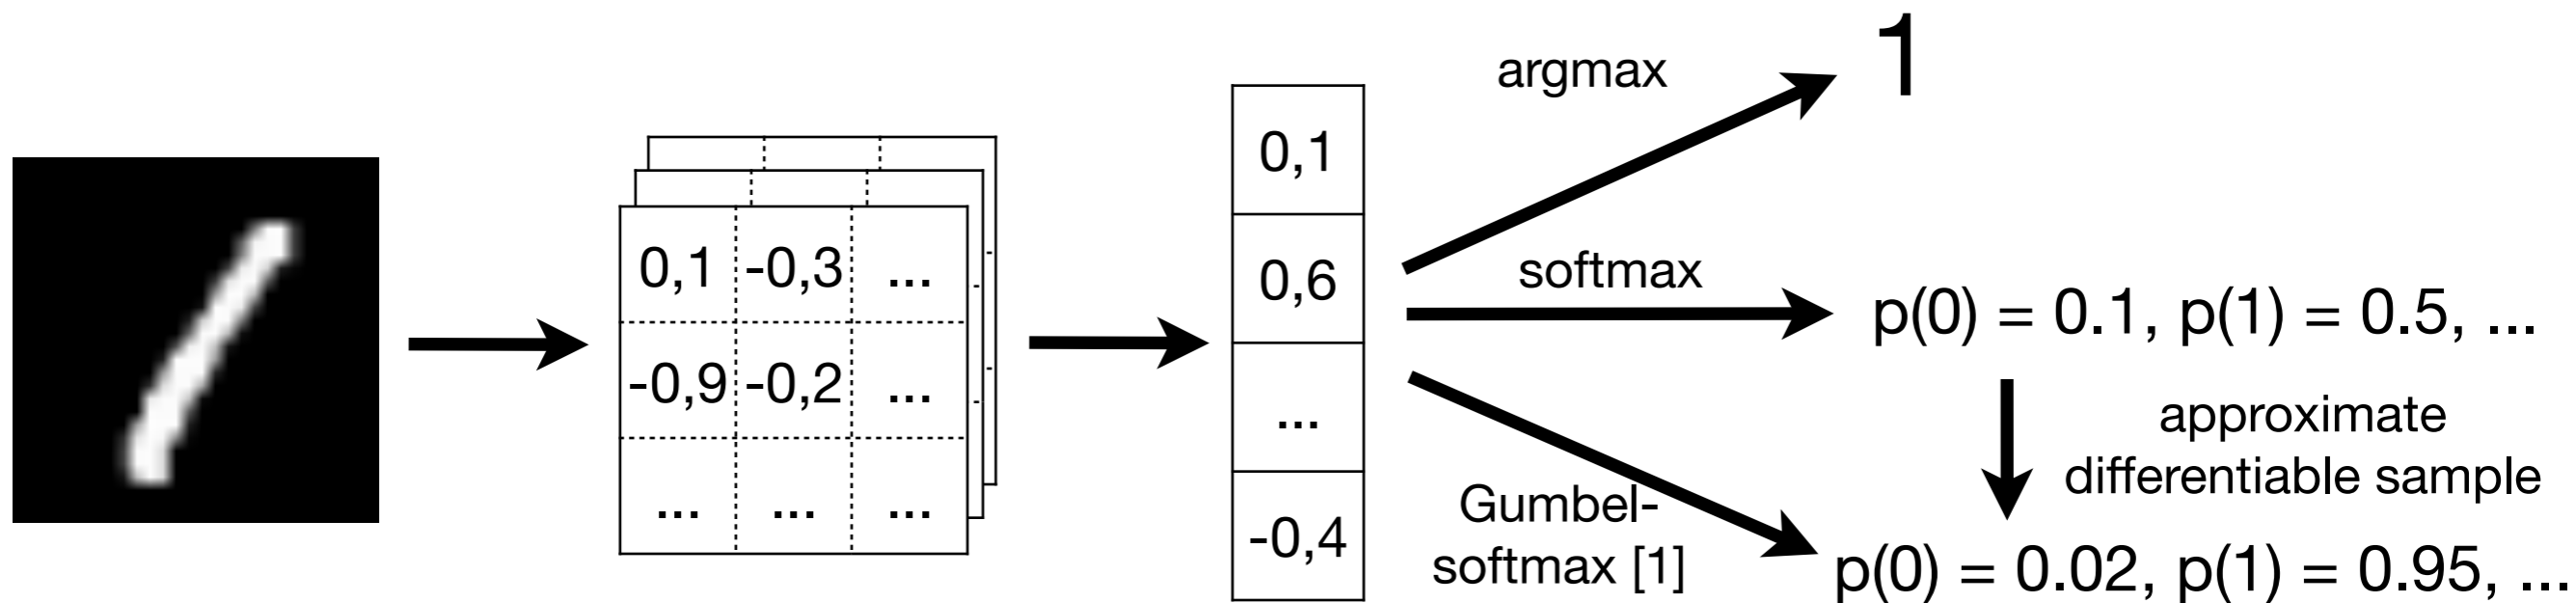


	0,3	-0,5	0,2	0,1	
0	0	0	0	0	0,6
1	1	1	0	0	0,2
2	1	0	0	1	0,4
3	1	1	0	0	

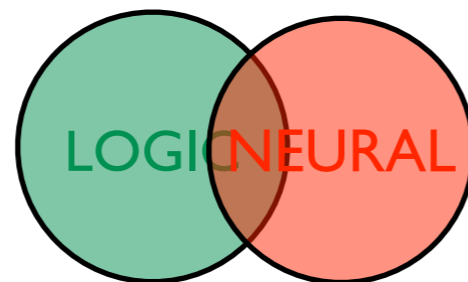


Sub-symbols to symbols

- E.g. in neural network classifiers
 - Turn real-valued vector into discrete classes
 - Final layer with specific activation function

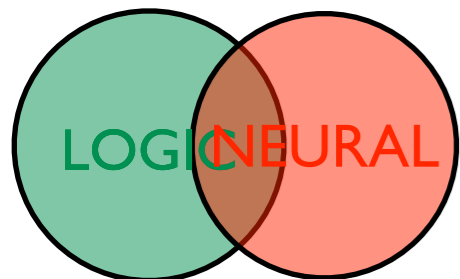


4. Symbolic vs sub-symbolic Representations in NeSy

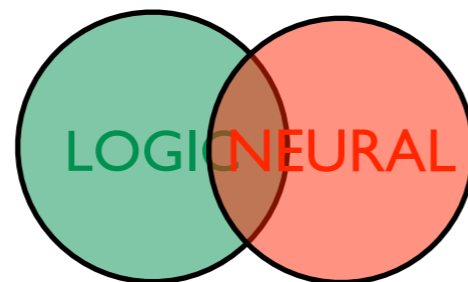


Representation in NeSy

- StarAI
 - Input = intermediate = output = symbolic representation
- Neural methods
 - Input = intermediate = sub-symbolic
 - Output =
 - Symbolic (classifier)
 - Or sub-symbolic (auto-encoder, GAN, regression, ...)
- NeSy
 - Intermediate representation = symbolic or sub-symbolic
 - We discern several approaches



4. Symbolic vs sub-symbolic Single translation step



Single translation step

- Symbolic input is mapped onto sub-symbols
 - One-hot encoding, relational embeddings, ...
- Afterwards, all reasoning happens in sub-symbolic space
- This approach is seen in most NeSy systems
- Examples include:
 - LTNs[1], SBR[2], NLMs[3], TensorLog[4]

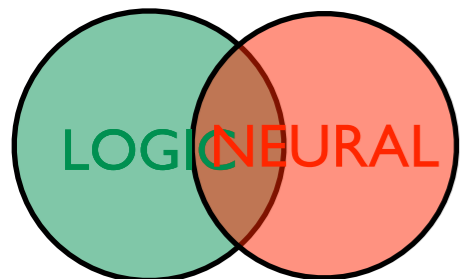
[1] Serafini, et al.: "Logic Tensor Networks:

Deep Learning and Logical Reasoning from Data and Knowledge", NeSy@HLAI 2016

[2] Diligenti et al.: "Semantic based regularization for learning and inference", Artificial Intelligence 2017

[3] Dong et al.: "Neural Logic Machines", ICLR 2019

[4] Cohen et al.: "Deep Learning meets Probabilistic DBs"



Logic Tensor Network

- This translations is made explicit in Logic Tensor Networks

Definition 1. A grounding \mathcal{G} for a first order language \mathcal{L} is a function from the signature of \mathcal{L} to the real numbers that satisfies the following conditions:

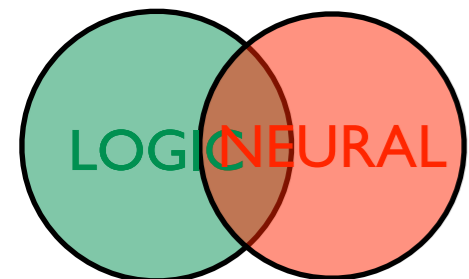
1. $\mathcal{G}(c) \in \mathbb{R}^n$ for every constant symbol $c \in \mathcal{C}$;
2. $\mathcal{G}(f) \in \mathbb{R}^{n \cdot \alpha(f)} \longrightarrow \mathbb{R}^n$ for every $f \in \mathcal{F}$;
3. $\mathcal{G}(P) \in \mathbb{R}^{n \cdot \alpha(P)} \longrightarrow [0, 1]$ for every $P \in \mathcal{P}$;

$$\mathcal{G}(f(t_1, \dots, t_m)) = \mathcal{G}(f)(\mathcal{G}(t_1), \dots, \mathcal{G}(t_m))$$

$$\mathcal{G}(P(t_1, \dots, t_m)) = \mathcal{G}(P)(\mathcal{G}(t_1), \dots, \mathcal{G}(t_m))$$

$$\mathcal{G}(\neg P(t_1, \dots, t_m)) = 1 - \mathcal{G}(P(t_1, \dots, t_m))$$

$$\mathcal{G}(\phi_1 \vee \dots \vee \phi_k) = \mu(\mathcal{G}(\phi_1), \dots, \mathcal{G}(\phi_k))$$



Logical Theory

GROUNDING OUT

```
stress (ann) .  
influences (ann , bob) .  
influences (bob , carl) .
```

```
smokes (ann) :- stress (ann) .  
smokes (bob) :- stress (bob) .  
smokes (carl) :- stress (carl) .
```

```
smokes (ann) :- influences (ann , ann) , smokes (ann) .  
smokes (ann) :- influences (bob , ann) , smokes (bob) .  
smokes (ann) :- influences (carl , ann) , smokes (carl) .
```

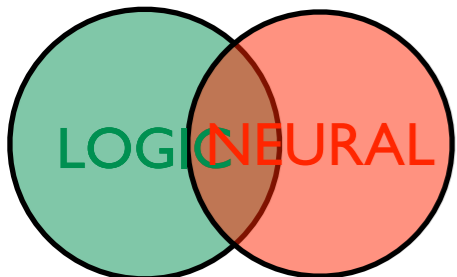
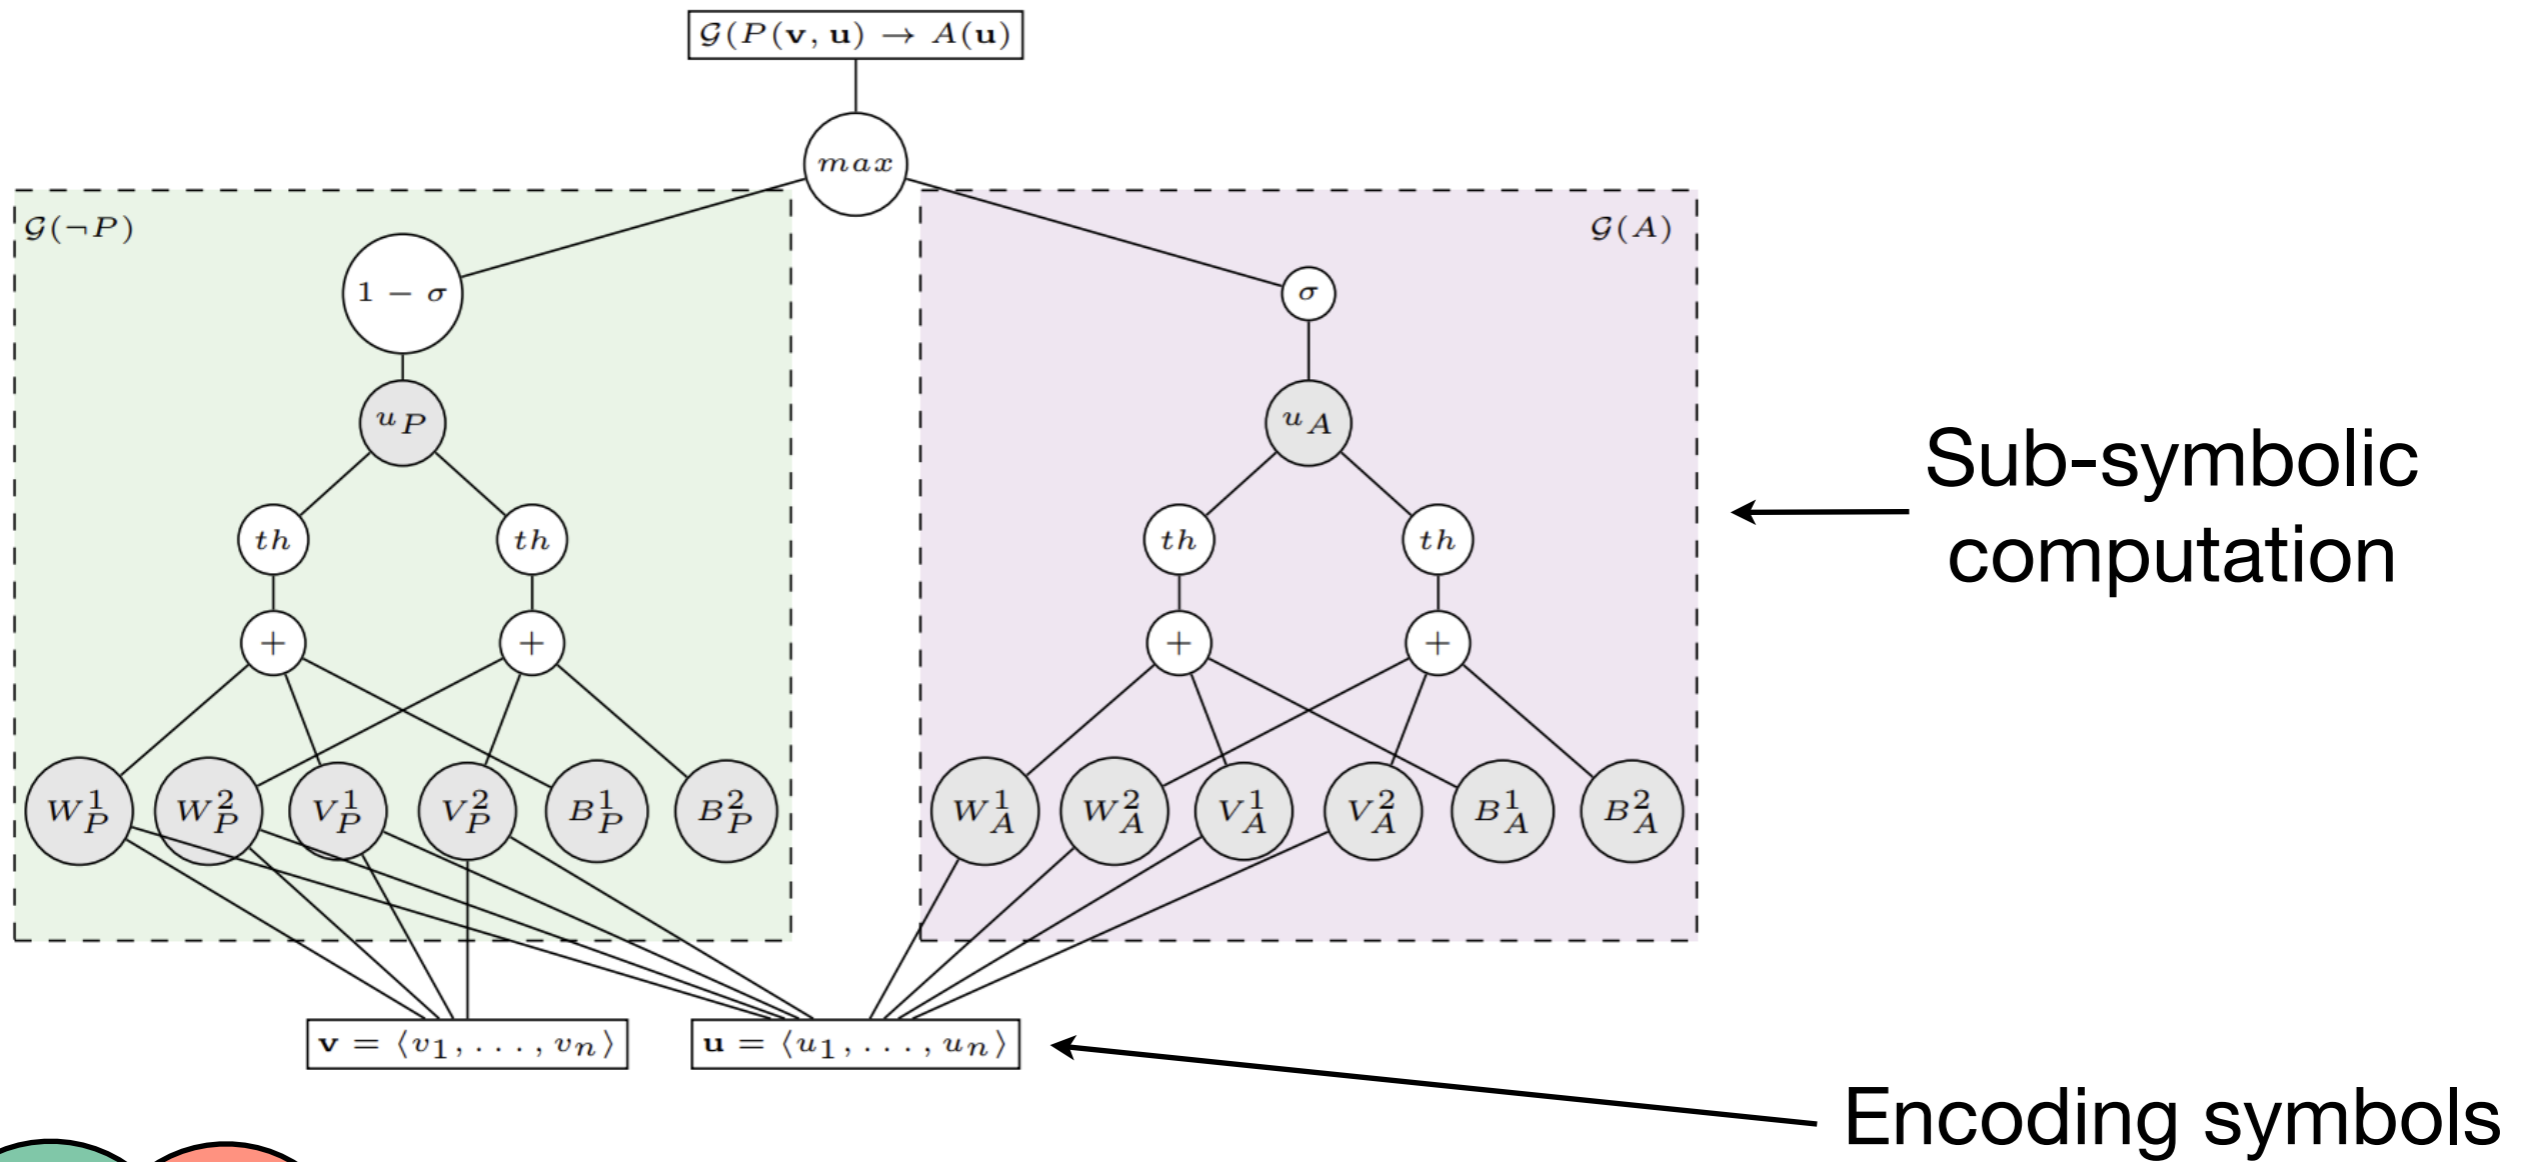
```
smokes (bob) :- influences (ann , bob) , smokes (ann) .  
smokes (bob) :- influences (bob , bob) , smokes (bob) .  
smokes (bob) :- influences (carl , bob) , smokes (carl) .
```

```
smokes (carl) :- influences (ann , carl) , smokes (ann) .  
smokes (carl) :- influences (bob , carl) , smokes (bob) .  
smokes (carl) :- influences (carl , carl) , smokes (carl) .
```

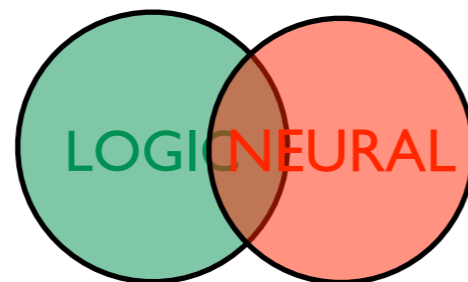
```
stress (ann) .  
influences (ann , bob) .  
influences (bob , carl) .  
  
smokes (X) :- stress (X) .  
smokes (X) :-  
    influences (Y , X) ,  
    smokes (Y) .
```

**IF INTERESTED ONLY IN
CERTAIN QUERIES,
CLEVER TECHNIQUES EXIST
TO AVOID GROUNDING OUT
COMPLETELY**

Logic Tensor Network

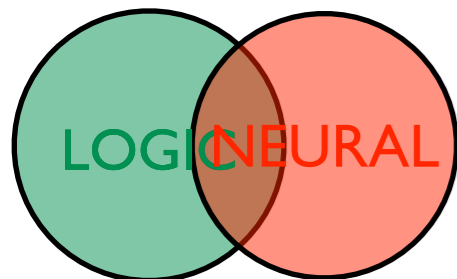


4. Symbolic vs sub-symbolic Alternating symbols and sub-symbols



Alternating symbols and sub-symbols

- Both symbolic and sub-symbolic representations are used
 - Not simultaneously by one component
 - Some components work on symbols, others on sub-symbols
- Indicative of systems that implement an interface
- Very natural for NeSy systems originating from a logical framework
- Examples include:
 - DeepProbLog[1], NeurASP[2], ...
 - ABL[3], NeuroLog[4], ..



[1] Manhaeve et al: "DeepProbLog: Neural Probabilistic Logic Programming", NeurIPS 2018

[2] Yang et al: "NeurASP: Embracing Neural Networks into Answer Set Programming", IJCAI 2020

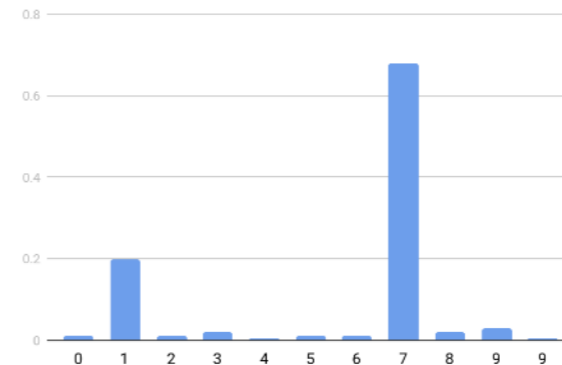
[3] Dai et al.: "Bridging Machine Learning and Logical Reasoning by Abductive Learning", NeurIPS 2019

[4] Tsamora et al. "Neural-symbolic integration: A compositional perspective"

Neural predicate

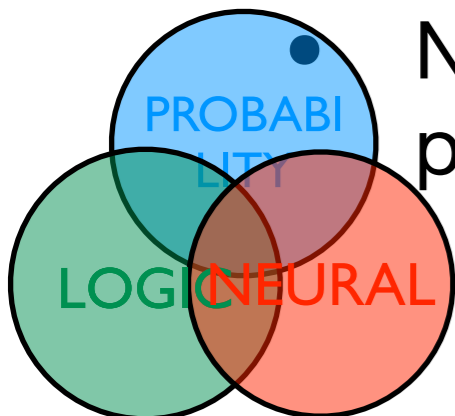


Output distribution



- Neural networks have uncertainty in their predictions
- A normalized output can be interpreted as a probability distribution
- Neural predicate models the output as probabilistic facts

No changes needed in the probabilistic host language



Key Idea DeepProbLog

unify the basic concepts in logic and neural networks:

neural predicate ~ neural net

an interface between logic and neural nets

DeepProbLog

- DeepProbLog: interface between PLP (ProbLog) and neural networks.
- This interface takes the form of the neural predicate
 - Output of neural networks represented as probabilistic facts

```
nn(mnist_net, [D], N, [0 ... 9] ) :: digit(D,N).  
addition(X,Y,Z) :- digit(X,N1), digit(Y,N2), Z is N1+N2.
```

- In the logic, the images are represented as constants
- Sub-symbolic properties are used in the neural network to make predictions
- This may seem as a limitation, but isn't

Examples:

```
addition( 3, 5, 8), addition( 0, 4, 4), addition( 9, 2, 11), ...
```

DeepProbLog exemplified: MNIST addition

Task: Classify pairs of MNIST digits with their sum

Benefit of DeepProbLog:

- Encode addition in logic
- Separate addition from digit classification



```
nn(mnist_net, [X], Y, [0 ... 9] ) :: digit(X,Y).
```

```
addition(X,Y,Z) :- digit(X,N1), digit(Y,N2), Z is N1+N2.
```

```
addition(3, 5, 8) :- digit(3, N1), digit(5, N2), 8 is N1 + N2.
```

Examples:

```
addition(3, , 8), addition(0, 4, 4), addition(9, 2, 11), ...
```

Example

Learn to classify the sum of pairs of MNIST digits

Individual digits are not labeled!

E.g. ( ,  , 8)

Could be done by a CNN: classify the concatenation of both images into 19 classes

However:      +    = ?

MNIST Addition

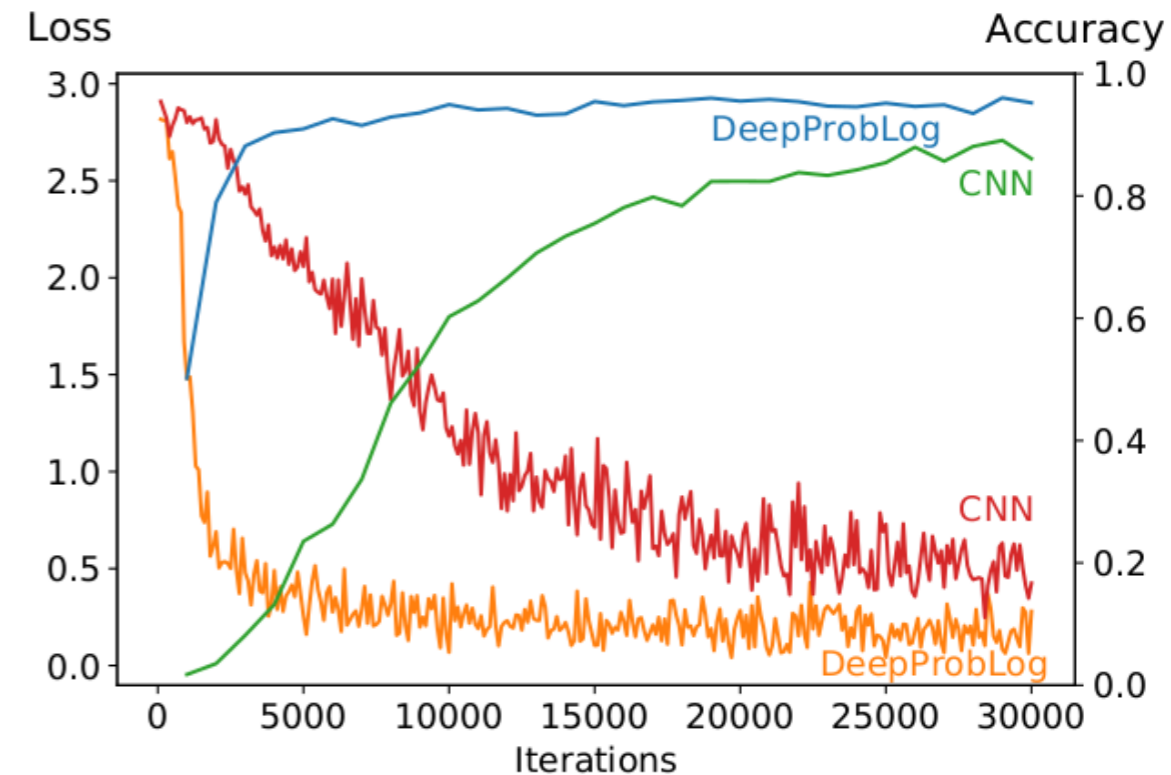
Pairs of MNIST images, labeled with sum

Baseline: CNN

- Classifies concatenation of both images into classes 0 ... 18

DeepProbLog:

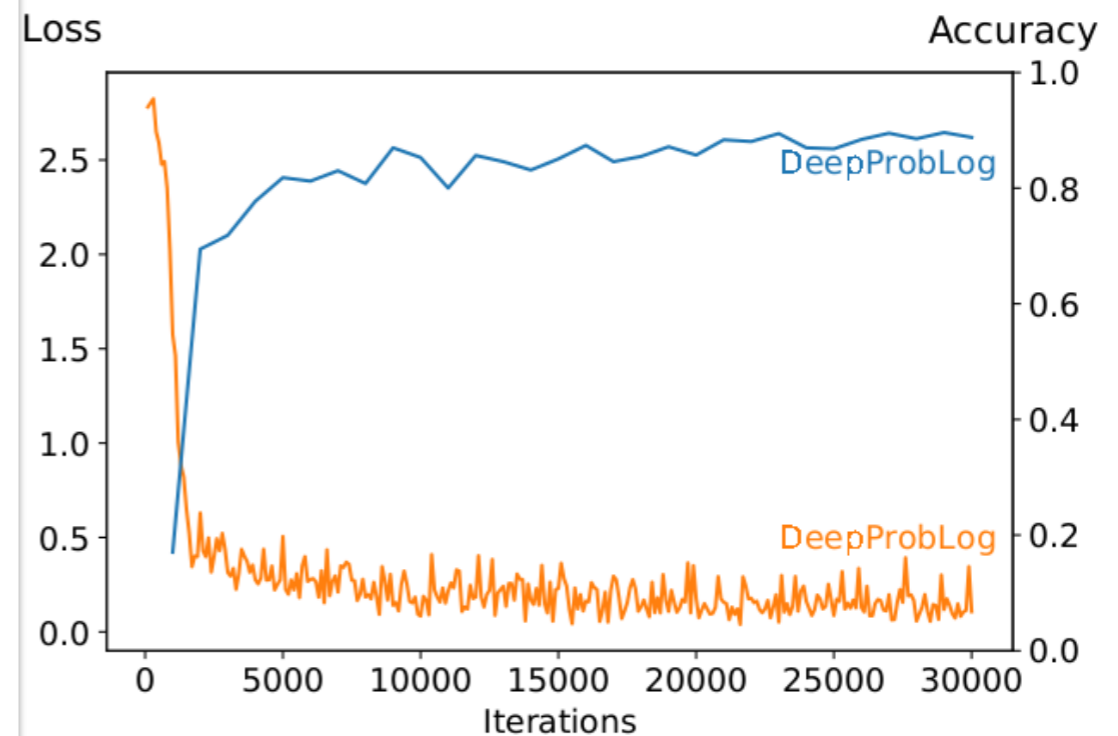
- CNN that classifies images into 0 ... 9
- Two lines of DeepProblog code



Multi-digit MNIST addition with MNIST

```
number ( [ ] , Result , Result ) .  
number ( [H | T ] , Acc , Result) :-  
    digit(H, Nr ) , Acc2 is Nr +10*Acc ,  
    number ( T , Acc2 , Result ) .  
number (X,Y) :- number (X, 0 ,Y) .
```

```
multiaddition(X, Y, Z) :-  
    number (X, X2 ) ,  
    number (Y, Y2 ) ,  
    Z is X2+Y2 .
```



(b) Multi-digit (T2)

Noisy Addition

```
mn(classifier, [X], Y, [0 .. 9]) :: digit(X,Y).  
t(0.2) :: noisy.
```

```
1/19 :: uniform(X,Y,0) ; ... ; 1/19 :: uniform(X,Y,18).
```

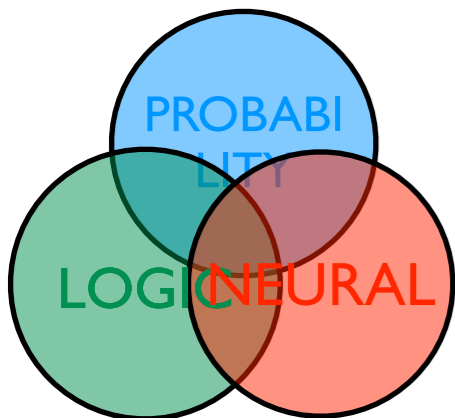
```
addition(X,Y,Z) :- noisy, uniform(X,Y,Z).
```

```
addition(X,Y,Z) :- \+noisy, digit(X,N1), digit(Y,N2), Z is N1+N2.
```

(a) The DeepProbLog program.

	Fraction of noise					
	0.0	0.2	0.4	0.6	0.8	1.0
Baseline	93.46	87.85	82.49	52.67	8.79	5.87
DeepProbLog	97.20	95.78	94.50	92.90	46.42	0.88
DeepProbLog w/ explicit noise	96.64	95.96	95.58	94.12	73.22	2.92
Learned fraction of noise	0.000	0.212	0.415	0.618	0.803	0.985

Table 3: The accuracy on the test set for **T4**.

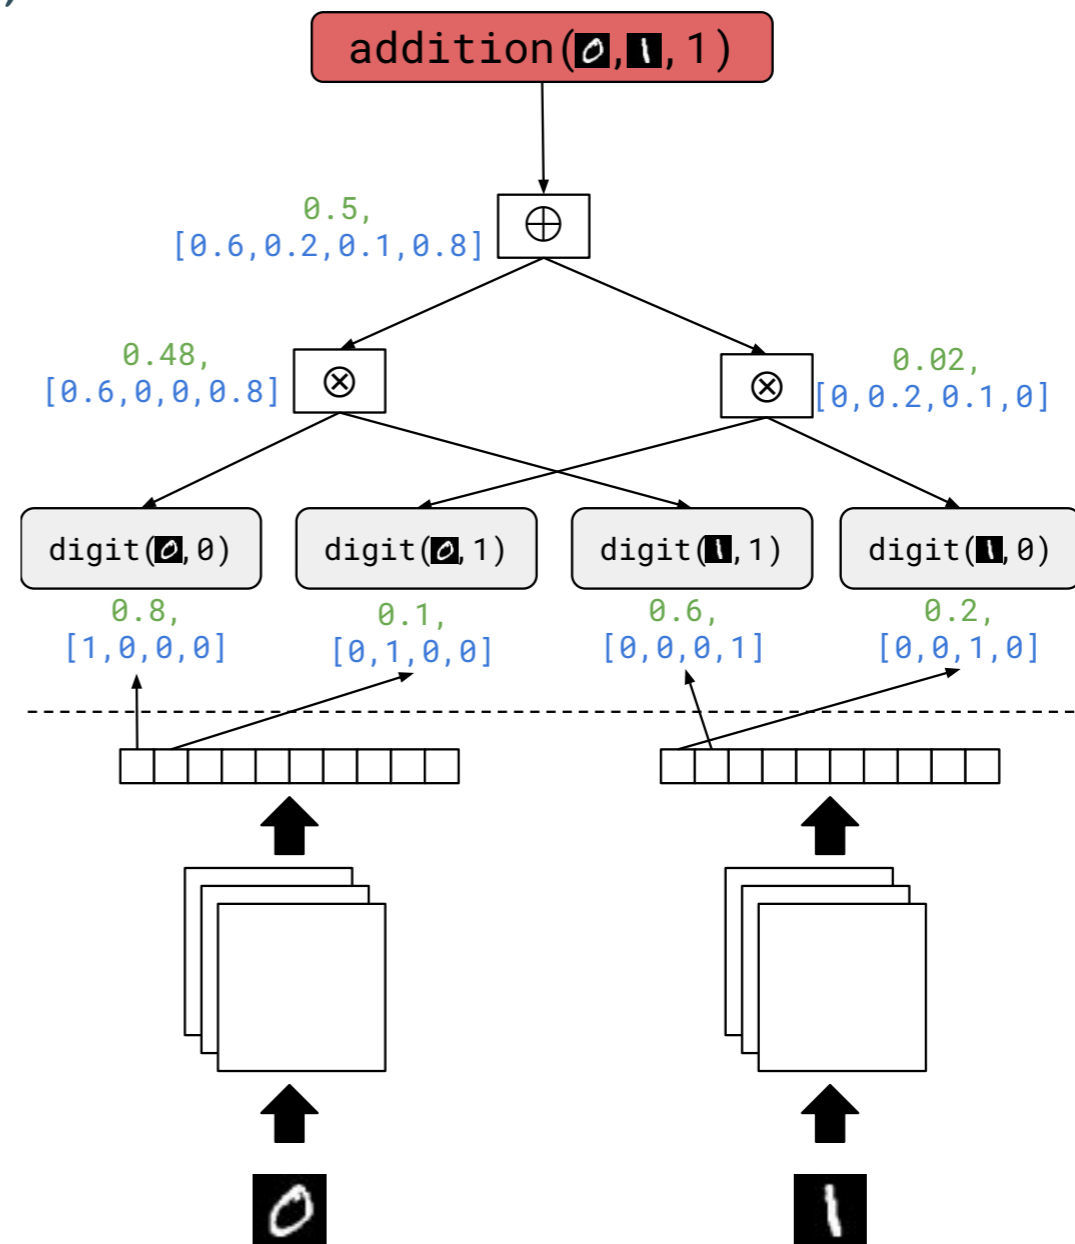
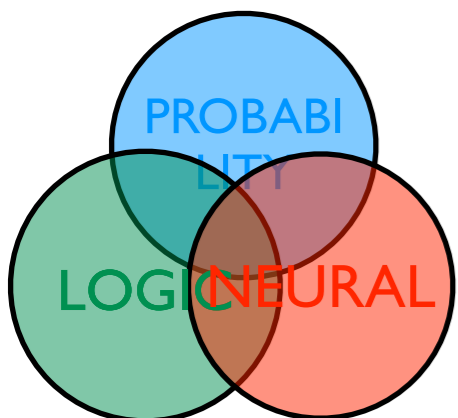


DeepProbLog

```
nn(mnist_net, [X], Y, [0 ... 9] ) ::
  digit(X,Y).
```

```
addition(X,Y,Z) :-
  digit(X,N1),
  digit(Y,N2),
  Z is N1+N2.
```

The ACs are differentiable and there is an interface with the neural nets



Useful Semirings

task	\mathcal{A}	e^\oplus	e^\otimes	\oplus	\otimes	$\alpha(v)$	$\alpha(\neg v)$	ref
SAT	$\{true, false\}$	<i>false</i>	<i>true</i>	\vee	\wedge	<i>true</i>	<i>true</i>	B, BT, G, GK, K, L, M
#SAT	\mathbb{N}	0	1	+	\cdot	1	1	B, G, GK, K, L
WMC	$\mathbb{R}_{\geq 0}$	0	1	+	\cdot	$\in \mathbb{R}_{\geq 0}$	$\in \mathbb{R}_{\geq 0}$	
PROB	$\mathbb{R}_{\geq 0}$	0	1	+	\cdot	$\in [0, 1]$	$1 - \alpha(v)$	B, BT, E, G, K
SENS	$\mathbb{R}[\mathcal{V}]$	0	1	+	\cdot	v or $\in [0, 1]$	$1 - \alpha(v)$	K
GRAD	$\mathbb{R}_{\geq 0} \times \mathbb{R}$	(0, 0)	(1, 0)	Eq. (4)	Eq. (5)	Eq. (2)	Eq. (3)	E, K
MPE	$\mathbb{R}_{\geq 0}$	0	1	max	\cdot	$\in [0, 1]$	$1 - \alpha(v)$	B, BT, G, K, L, M
S-PATH	\mathbb{N}^∞	∞	0	min	+	$\in \mathbb{N}$	0	BT, GK, K
W-PATH	\mathbb{N}^∞	0	∞	max	min	$\in \mathbb{N}$	∞	BT
FUZZY	[0, 1]	0	1	max	min	$\in [0, 1]$	1	GK, M
k WEIGHT	$\{0, \dots, k\}$	k	0	min	$+^k$	$\in \{0, \dots, k\}$	$\in \{0, \dots, k\}$	M
OBDD $_{<}$	OBDD $_{<}(\mathcal{V})$	OBDD $_{<}(0)$	OBDD $_{<}(1)$	\vee	\wedge	OBDD $_{<}(v)$	\neg OBDD $_{<}(v)$	K
WHY	$\mathcal{P}(\mathcal{V})$	\emptyset	\emptyset	\cup	\cup	$\{v\}$	n/a	GK
\mathcal{RA}^+	$\mathbb{N}[\mathcal{V}]$	0	1	+	\cdot	v	n/a	GK

Table 1: Examples of commutative semirings and labeling functions. The **WHY** and \mathcal{RA}^+ provenance semirings apply to positive literals only. Reference key: B (Bacchus et al., 2009), BT (Baras and Theodorakopoulos, 2010), E (Eisner, 2002), G (Goodman, 1999), GK (Green et al., 2007), K (Kimmig et al., 2011), L (Larrosa et al., 2010), M (Meseguer et al., 2006); more examples can be found in these references.

Program Induction/Sketching

In Neural Symbolic methods

- Rule Induction — work with templates

$$P(X) :- R(X,Y), Q(Y)$$

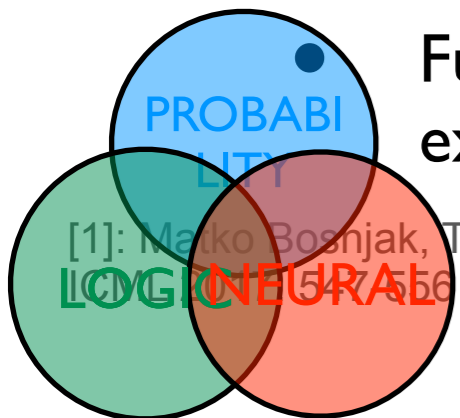
- and have the “predicate” variables / slots P,Q, R determined by the NN
- Simpler form, fill just a few slots / holes

Approach similar to ‘*Programming with a Differentiable Forth Interpreter*’ [1] 24

- Partially defined Forth program with slots / holes
- Slots are filled by neural network (encoder / decoder)

Fully differentiable interpreter: NNs are trained with input / output examples

[1]: Marko Boshnjak, Tim Rocktäschel, Jason Naradowsky, Sebastian Riedel: Programming with a Differentiable Forth Interpreter.



Example DeepProbLog

neural predicate

```
hole(X,Y,X,Y):-
  swap(X,Y,0).
```

```
hole(X,Y,Y,X):-
  swap(X,Y,1).
```

bubble sort

```
bubble([X],[],X).
bubble([H1,H2|T],[X1|T1],X):-
  hole(H1,H2,X1,X2),
  bubble([X2|T],T1,X).
```

```
bubblesort([],L,L).
```

```
bubblesort(L,L3,Sorted) :-
  bubble(L,L2,X),
  bubblesort(L2,[X|L3],Sorted).
```

```
sort(L,L2) :- bubblesort(L,[],L2).
```

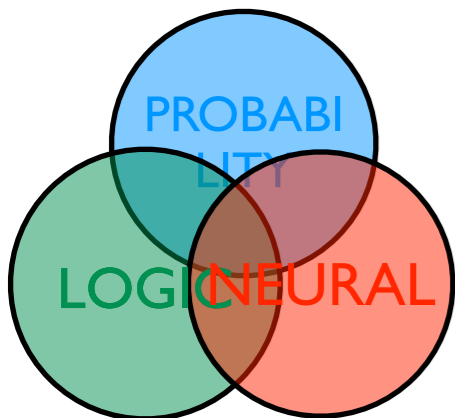
	Test Length	Sorting: Training length					Addition: training length		
		2	3	4	5	6	2	4	8
$\partial 4$ [Bošnjak et al., 2017]	8	100.0	100.0	49.22	-	-	100.0	100.0	100.0
	64	100.0	100.0	20.65	-	-	100.0	100.0	100.0
DeepProbLog	8	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
	64	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0

(a) Accuracy on the sorting and addition problems (results for $\partial 4$ reported by Bošnjak et al. [2017]).

Training length \longrightarrow	2	3	4	5	6
$\partial 4$ on GPU	42 s	160 s	-	-	-
$\partial 4$ on CPU	61 s	390 s	-	-	-
DeepProbLog on CPU	11 s	14 s	32 s	114 s	245 s

(b) Time until 100% accurate on test length 8 for the sorting problem.

Table 1: Results on the Differentiable Forth experiments



DeepSeaProbLog

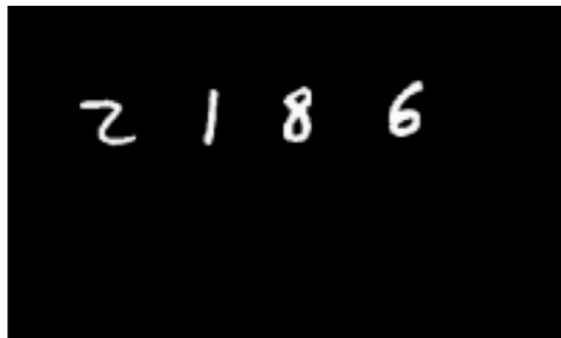
discrete and continuous distributions [De Smet UAI 23]

useful for robotics and perception

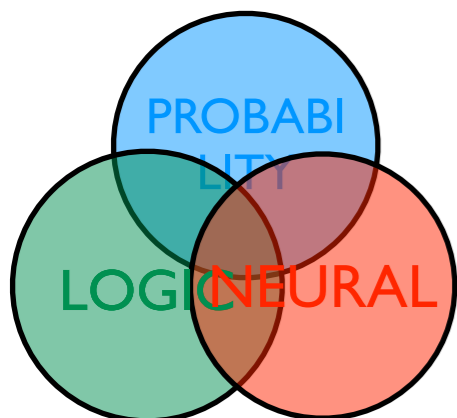
dim is neural net returning parameters of normal distribution.

`length (Obj) ~ normal (dim (Obj , Image)) .`

`large (Obj) :- length (Obj) > 100 .`







**determining order digits
to determine year**




DeepSeaProbLog

discrete and continuous distributions [De Smet UAI 23]

generative model with variational autoencoders (see also [Misoni et al NeurIPS 22])

So far from input   to output 11 so that **SUM**(  ,11) holds

In DeepSeaProblog, you can query **SUM**( , X, 5)

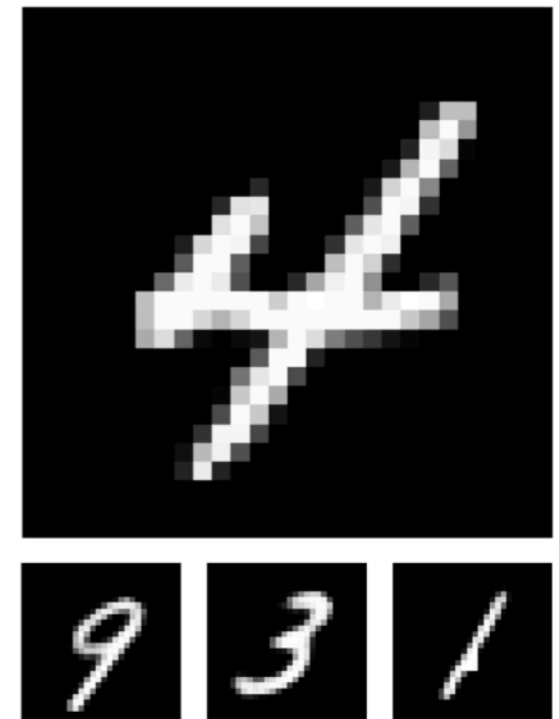
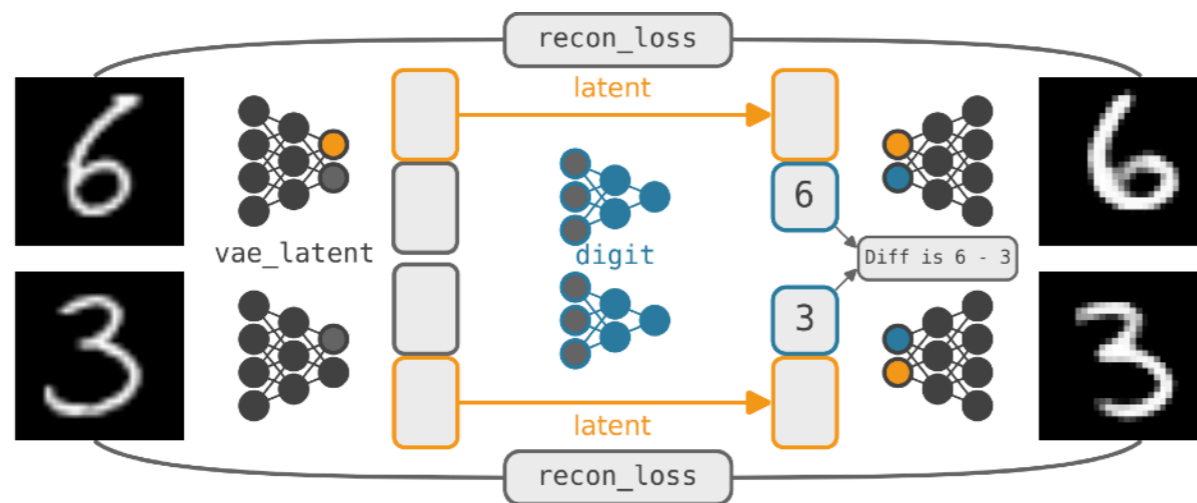

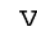
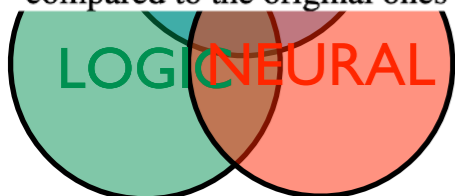
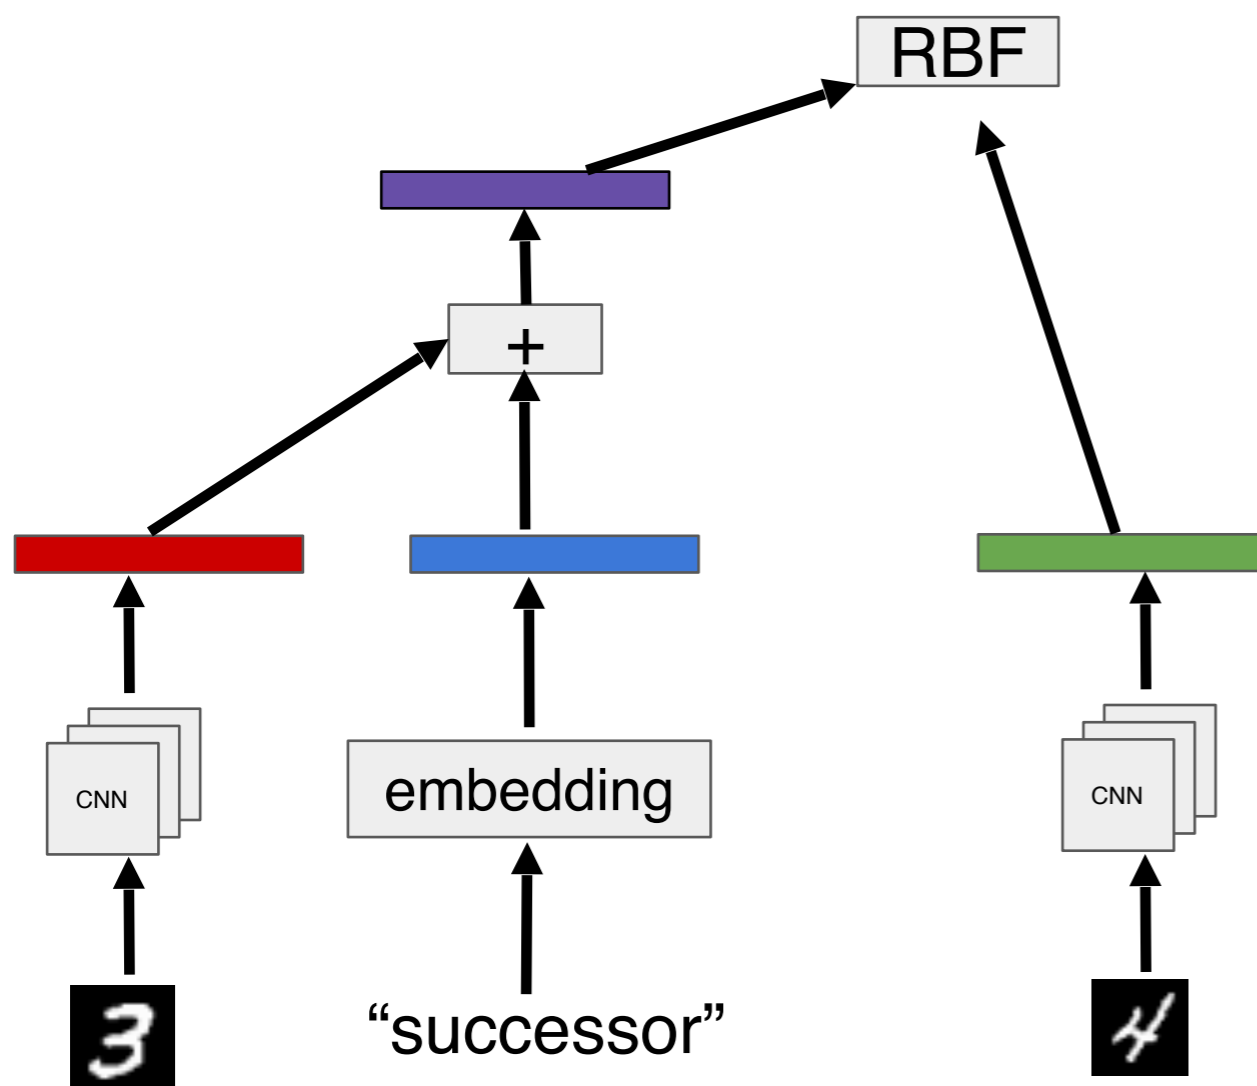


Figure 4: Given example pairs of images and the value of their subtraction, e.g., (, ) and 3, the CVAE encoder `vae_latent` first encodes each image into a multivariate normal NDF (`latent`) and a latent vector. The latter is the input of a categorical NDF `digit`, completing the CVAE latent space. Supervision is dual; generated images are compared to the original ones in a probabilistic reconstruction loss, while both digits need to subtract to the given value.



DeepProbLog: Embeddings as symbols

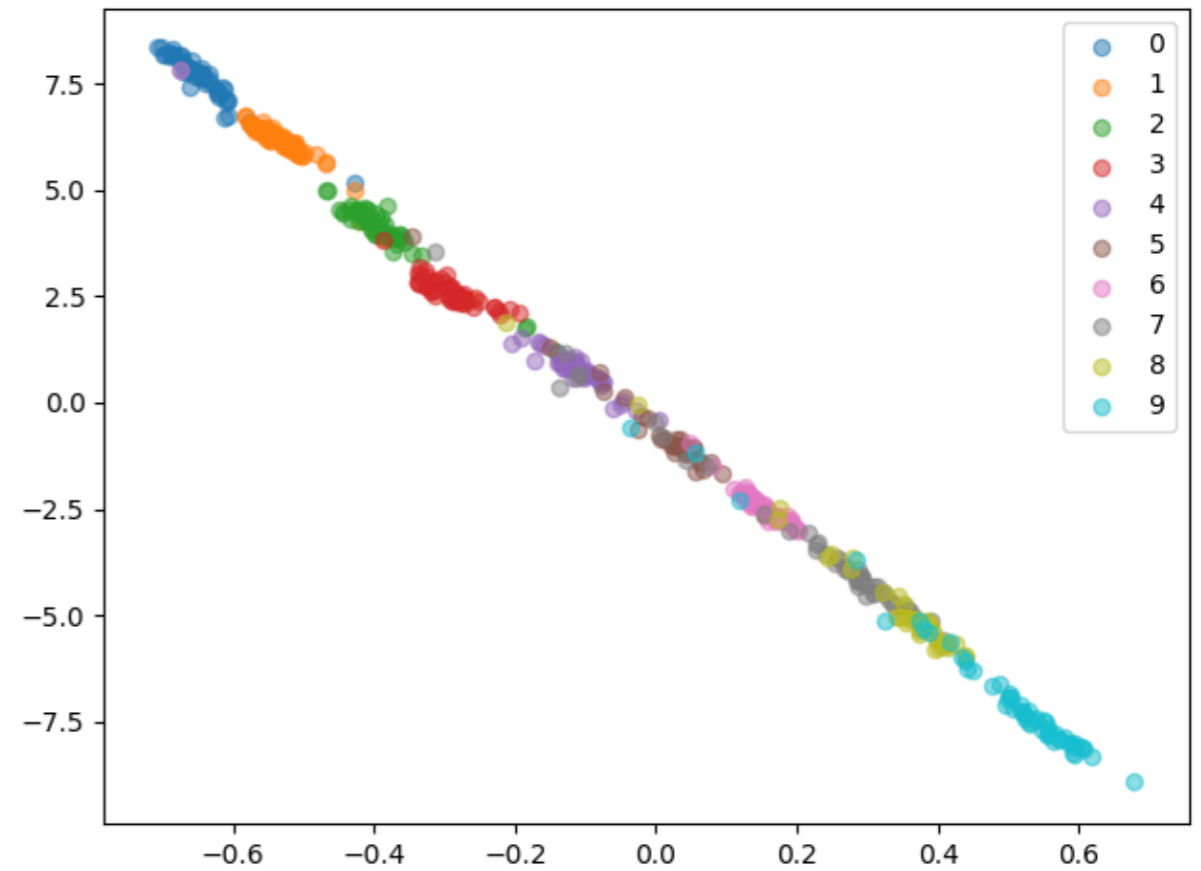
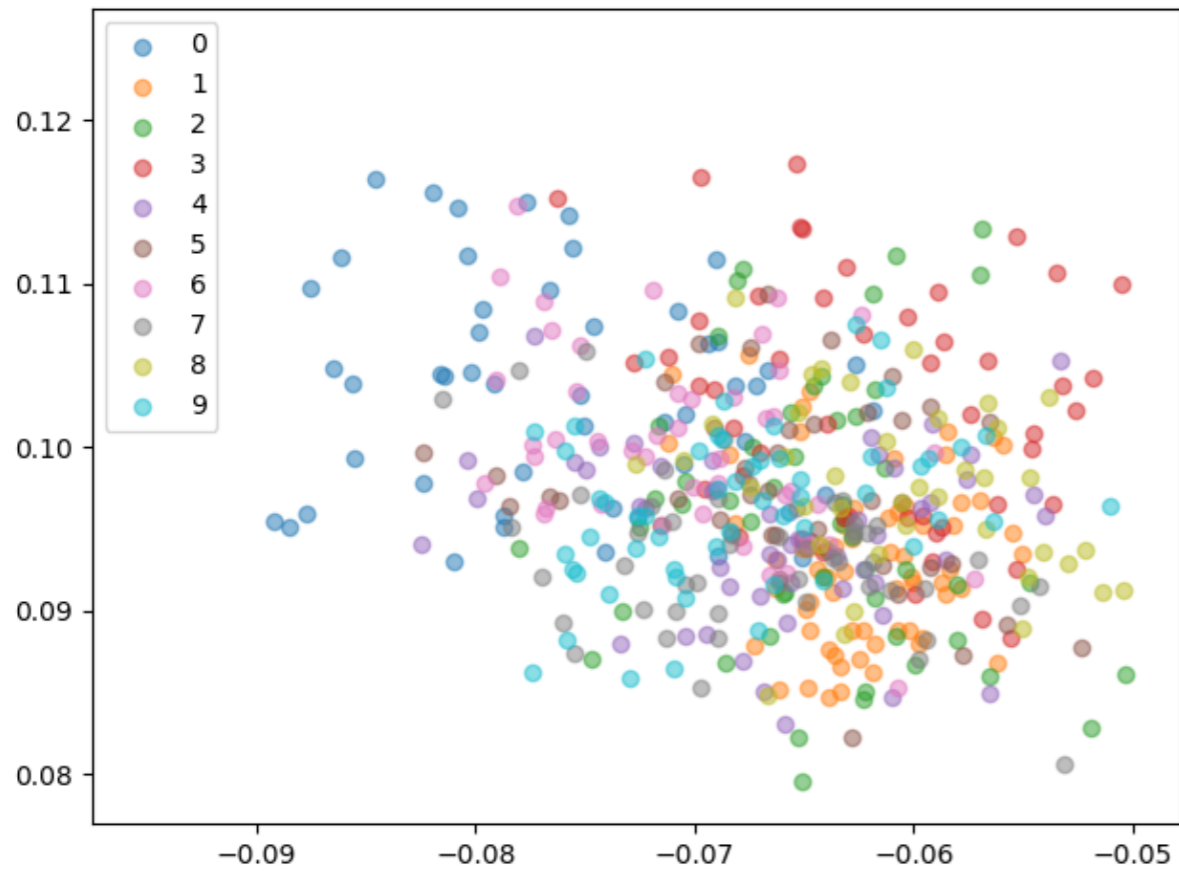
Computational Graph



```
successor(3, 4) :-  
  cnn_embed(3, e1),  
  cnn_embed(4, e2),  
  embed("successor", r),  
  add(r, e1, e3),  
  rbf(e2, e3).
```

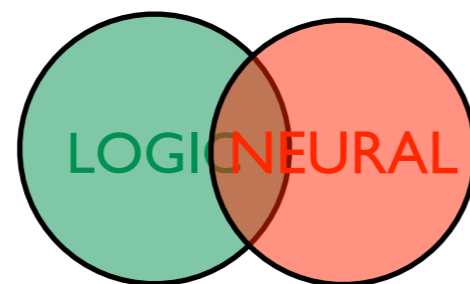
Idea of TransE [Bordes et al]

2D MNIST image embeddings



4. Symbolic vs sub-symbolic

Simultaneously symbolic and sub-symbolic



Neural Theorem Prover

Towards Neural Theorem Proving at Scale

Example Knowledge Base:

1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
`fatherOf(X, Z),`
`parentOf(Z, Y).`

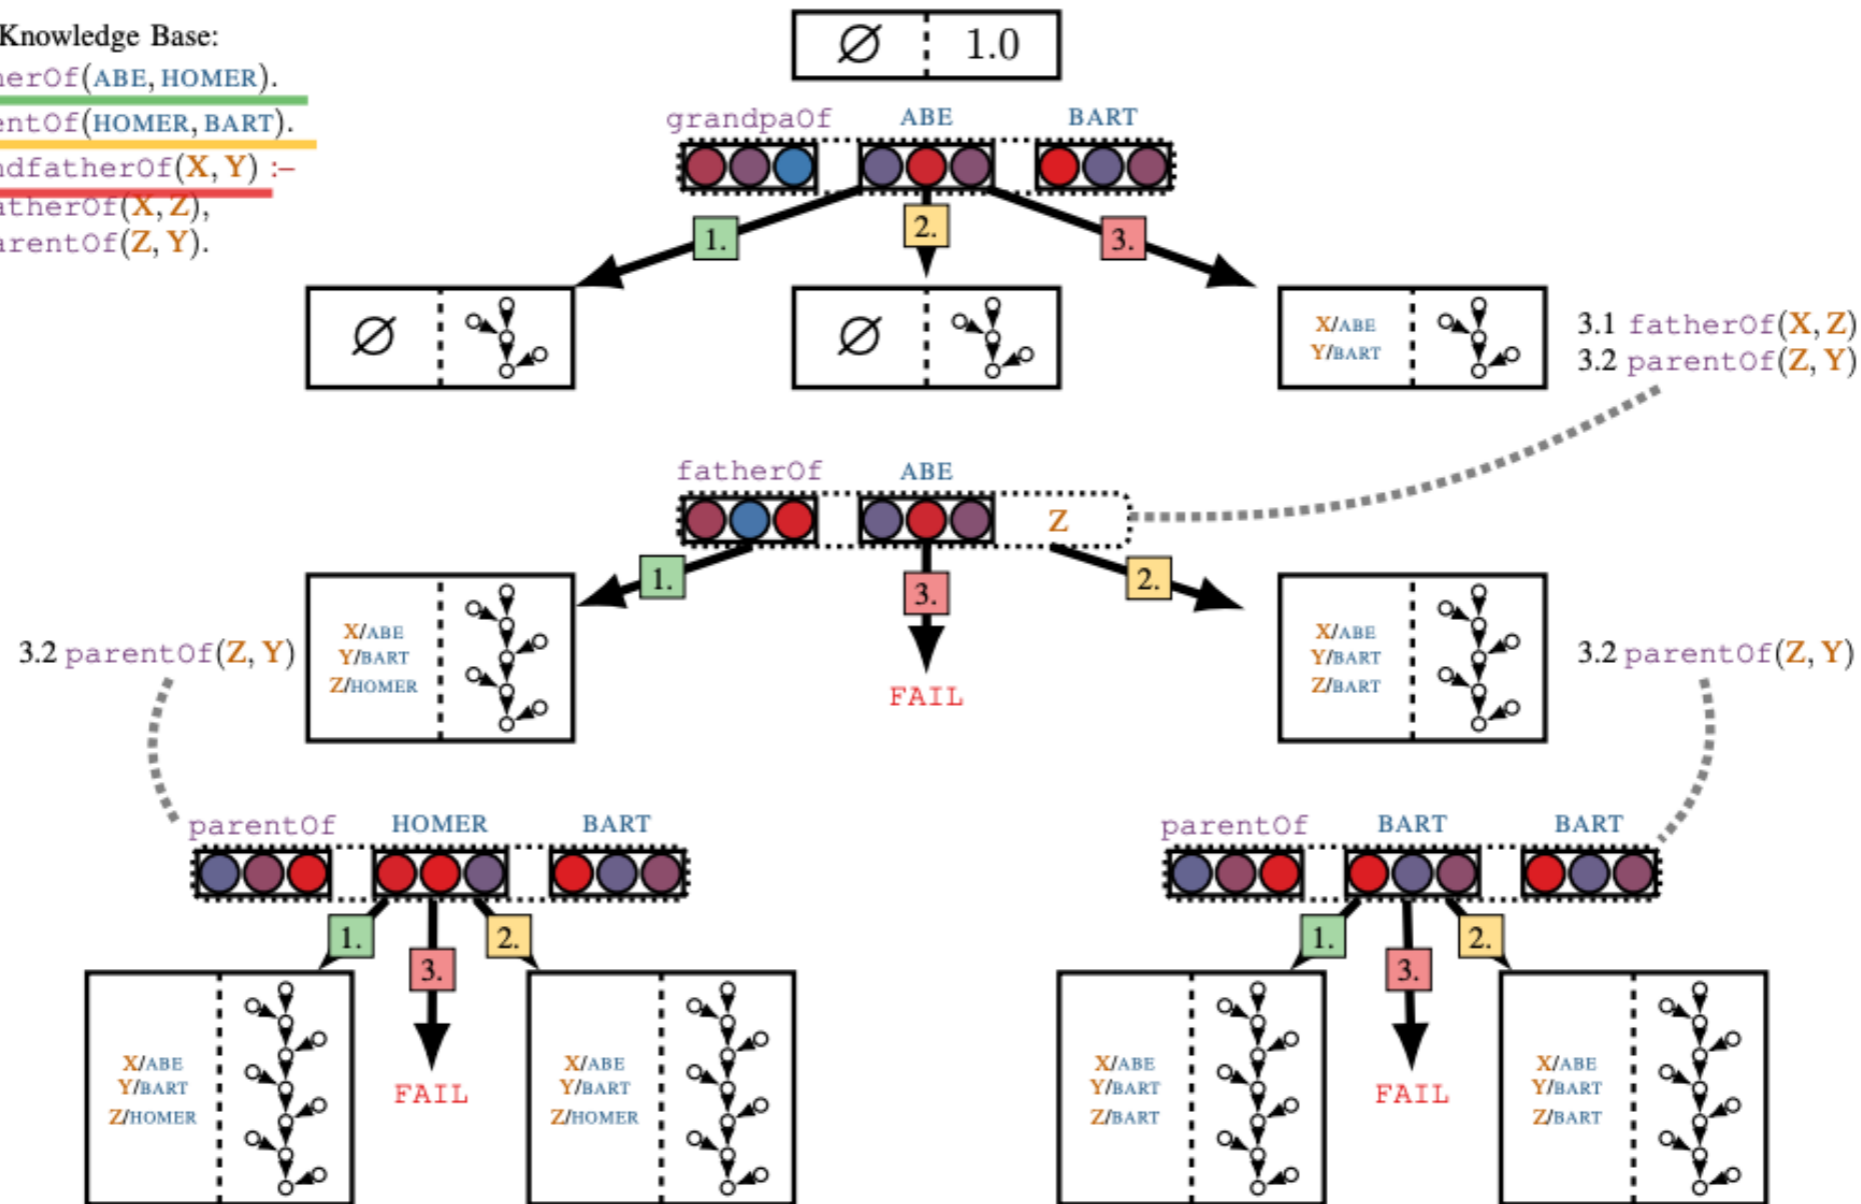
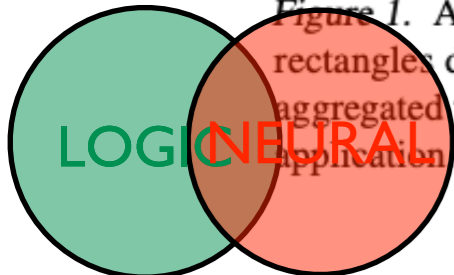
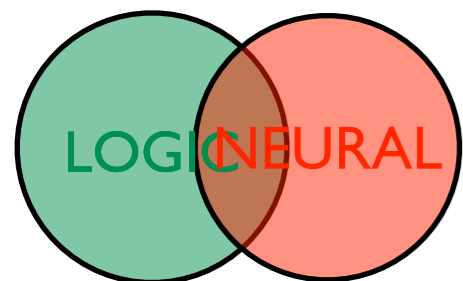


Figure 1. A visual depiction of the NTP's recursive computation graph construction, applied to a toy KB (top left). Dash-separated rectangles denote proof states (left: substitutions, right: proof score -generating neural network). All the non-FAIL proof states are aggregated to obtain the final proof success (depicted in Figure 2). Colours and indices on arrows correspond to the respective KB rule application.



Simultaneously symbolic and sub-symbolic

- Both symbolic and sub-symbolic representations are used
 - All entities have both representations
 - Reasoning uses both **simultaneously**
- Reasoning mechanism is extended
- Only used in a few systems
 - E.g. NTP[1], CTP[2]



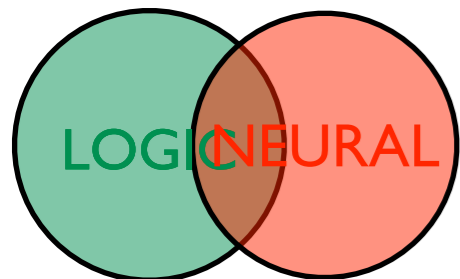
[1] Rocktäschel et al.: "End-to-end differentiable proving.", NeurIPS 2017.

[2] Minervini et al.: "Learning Reasoning Strategies in End-to-End Differentiable Proving", ICML 2020

Neural Theorem Prover

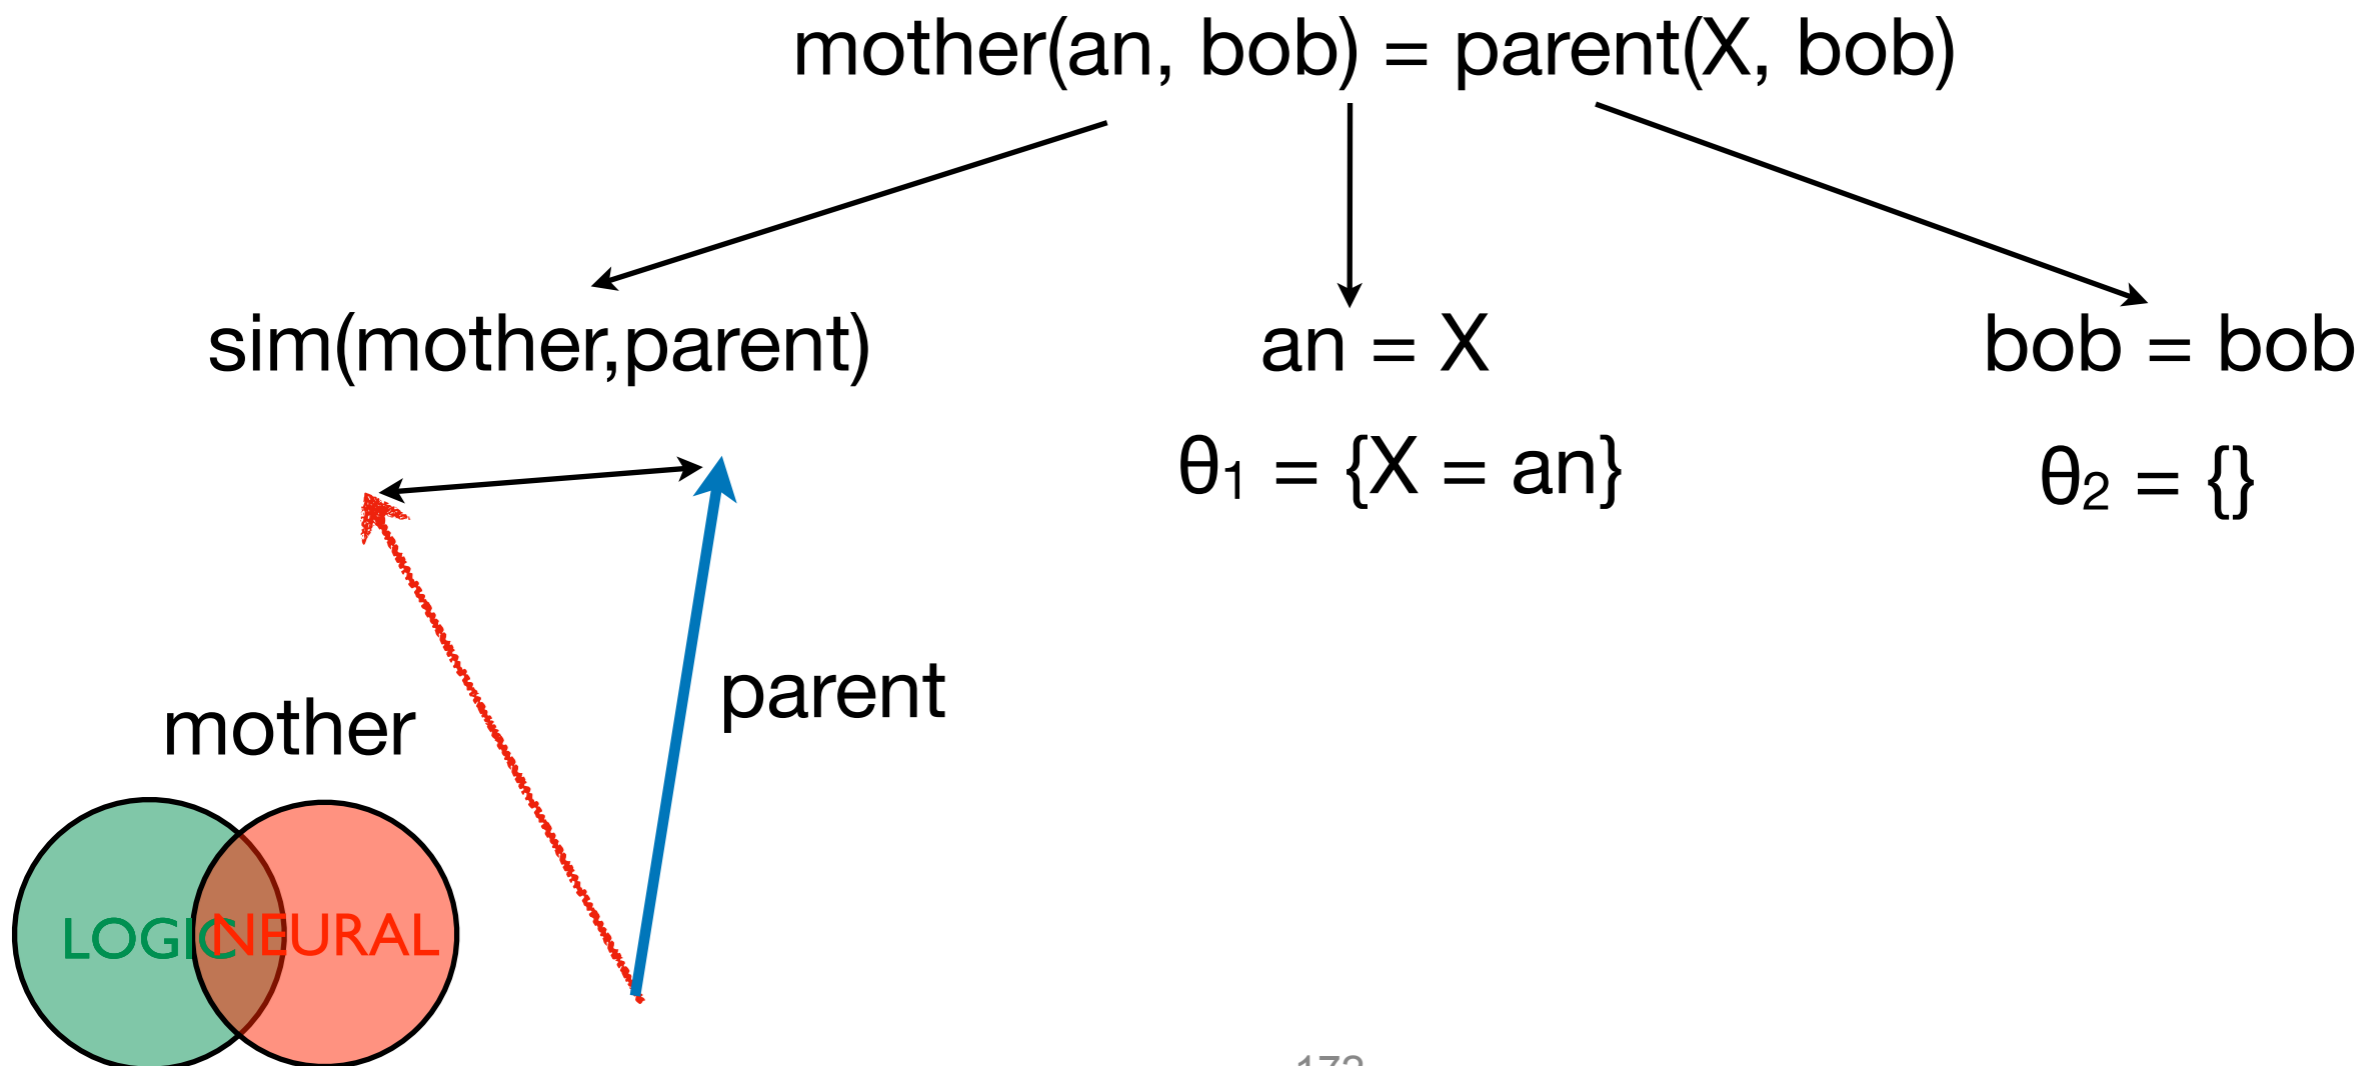
- The neural theorem prover uses both symbols and sub-symbols simultaneously
- Symbols retain their symbolic nature
- Each symbol has a learnable sub-symbol T

- Symbol comparison:
 - Normal unification
- Comparison of sub-symbols:
 - $\text{sim}(x,y) = \exp(-\|T_x - T_y\|_2)$



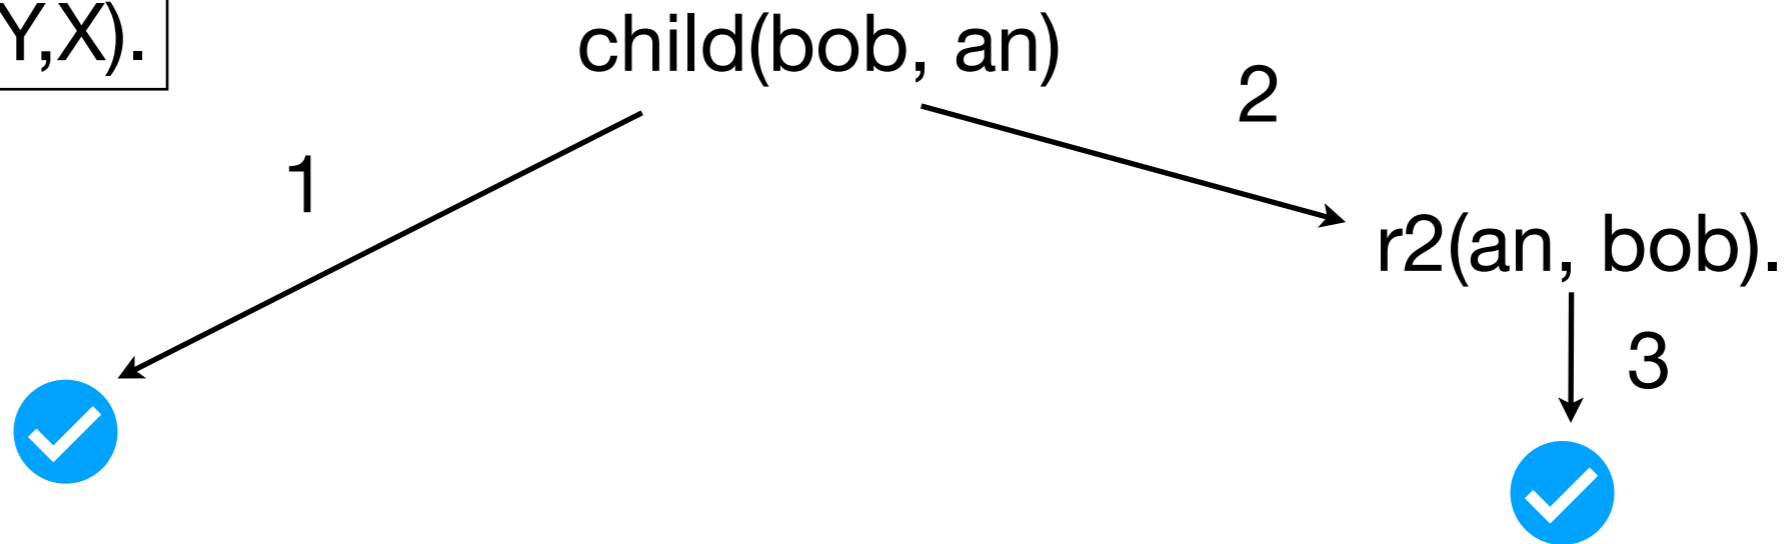
Soft unification

- Unify what can be unified
- Use similarity to compare other symbols and use it as a score



Example

mother(an, bob).
r1(X, Y) :- r2(Y, X).

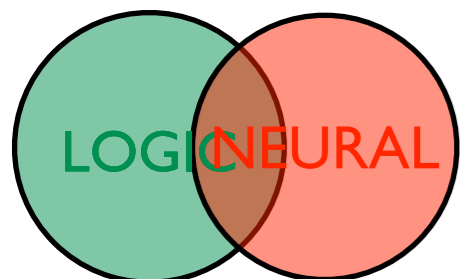


Unifications

1) mother(an,bob) = child(bob,an)
sim(mother,child)
sim(an,bob)

2) r1(X,Y) = child(bob,an)
sim(r1,child)
X = bob
Y = an

3) r2(an, bob) = mother(an, bob)
sim(r2,mother)

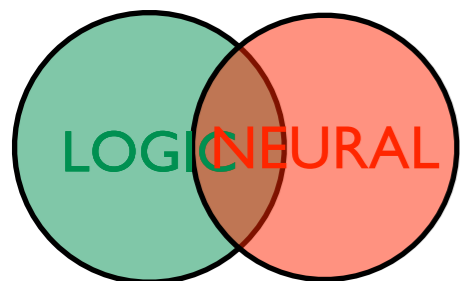


NTP

Knowledge base completion

Table 1: AUC-PR results on Countries and MRR and HITS@ m on Kinship, Nations, and UMLS.

Corpus	Metric	Model			Examples of induced rules and their confidence	
		Complex	NTP	NTP λ		
Countries	S1	AUC-PR	99.37 \pm 0.4	90.83 \pm 15.4	100.00 \pm 0.0	0.90 locatedIn(X,Y) :- locatedIn(X,Z), locatedIn(Z,Y).
	S2	AUC-PR	87.95 \pm 2.8	87.40 \pm 11.7	93.04 \pm 0.4	0.63 locatedIn(X,Y) :- neighborOf(X,Z), locatedIn(Z,Y).
	S3	AUC-PR	48.44 \pm 6.3	56.68 \pm 17.6	77.26 \pm 17.0	0.32 locatedIn(X,Y) :- neighborOf(X,Z), neighborOf(Z,W), locatedIn(W,Y).
Kinship	MRR		0.81	0.60	0.80	0.98 term15(X,Y) :- term5(Y,X)
	HITS@1		0.70	0.48	0.76	0.97 term18(X,Y) :- term18(Y,X)
	HITS@3		0.89	0.70	0.82	0.86 term4(X,Y) :- term4(Y,X)
	HITS@10		0.98	0.78	0.89	0.73 term12(X,Y) :- term10(X,Z), term12(Z,Y).
Nations	MRR		0.75	0.75	0.74	0.68 blockpositionindex(X,Y) :- blockpositionindex(Y,X).
	HITS@1		0.62	0.62	0.59	0.46 expeldiplomats(X,Y) :- negativebehavior(X,Y).
	HITS@3		0.84	0.86	0.89	0.38 negativecomm(X,Y) :- commonbloc0(X,Y).
	HITS@10		0.99	0.99	0.99	0.38 intergovorgs3(X,Y) :- intergovorgs(Y,X).
UMLS	MRR		0.89	0.88	0.93	0.88 interacts_with(X,Y) :- interacts_with(X,Z), interacts_with(Z,Y).
	HITS@1		0.82	0.82	0.87	
	HITS@3		0.96	0.92	0.98	0.77 isa(X,Y) :- isa(X,Z), isa(Z,Y).
	HITS@10		1.00	0.97	1.00	0.71 derivative_of(X,Y) :- derivative_of(X,Z), derivative_of(Z,Y).

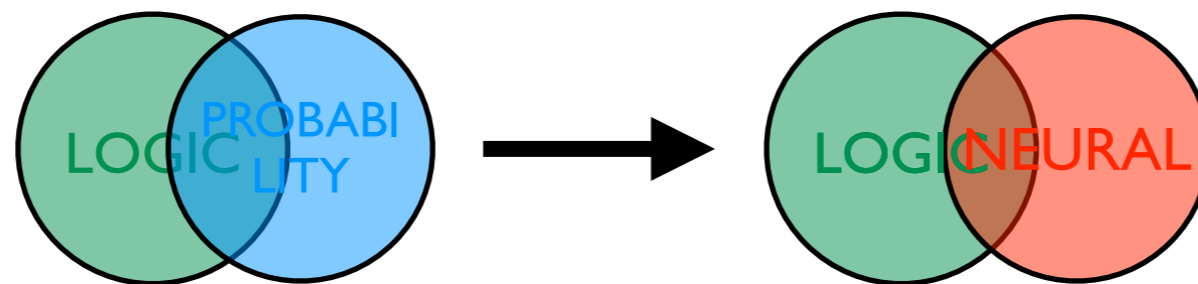


4. Symbolic vs sub-symbolic

Key Messages

- Entities are represented very differently in symbolic and sub-symbolic systems, but they are complementary
- NeSy systems can be categorized by how they use symbolic and sub-symbolic intermediate representations

5. Structure vs parameter learning

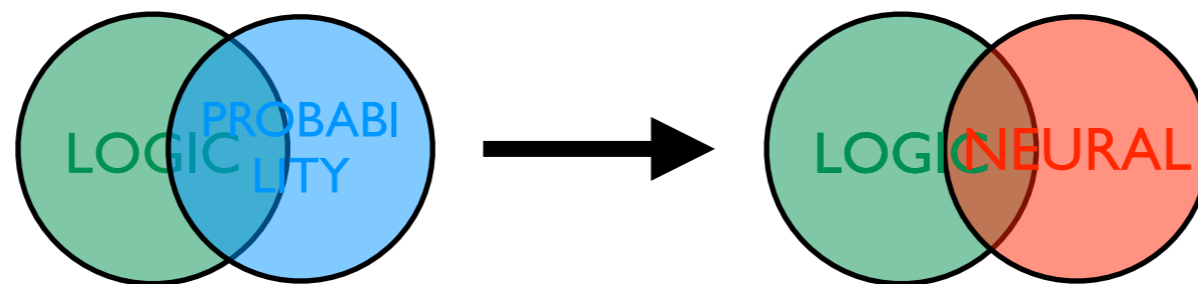


5. Learning

Key Messages

- Learning: finding logical formulas and estimating probabilities
- Structure learning: both formulas and probabilities
- Parameter learning: only probabilities
- Many flavours of learning in NeSy

5. Structure vs parameter learning



Spectrum of learning paradigms

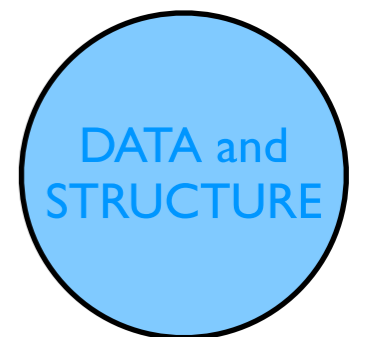
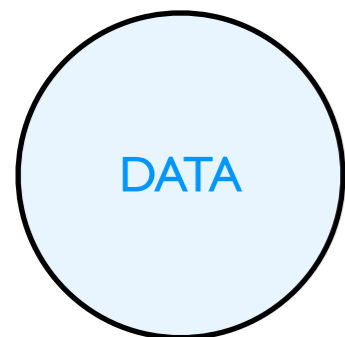
Soft patterns

Neural generation

Structure via
parameter learning

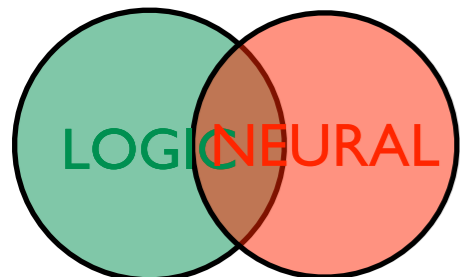
Neurally-guided
learning

Program sketching



Structure learning

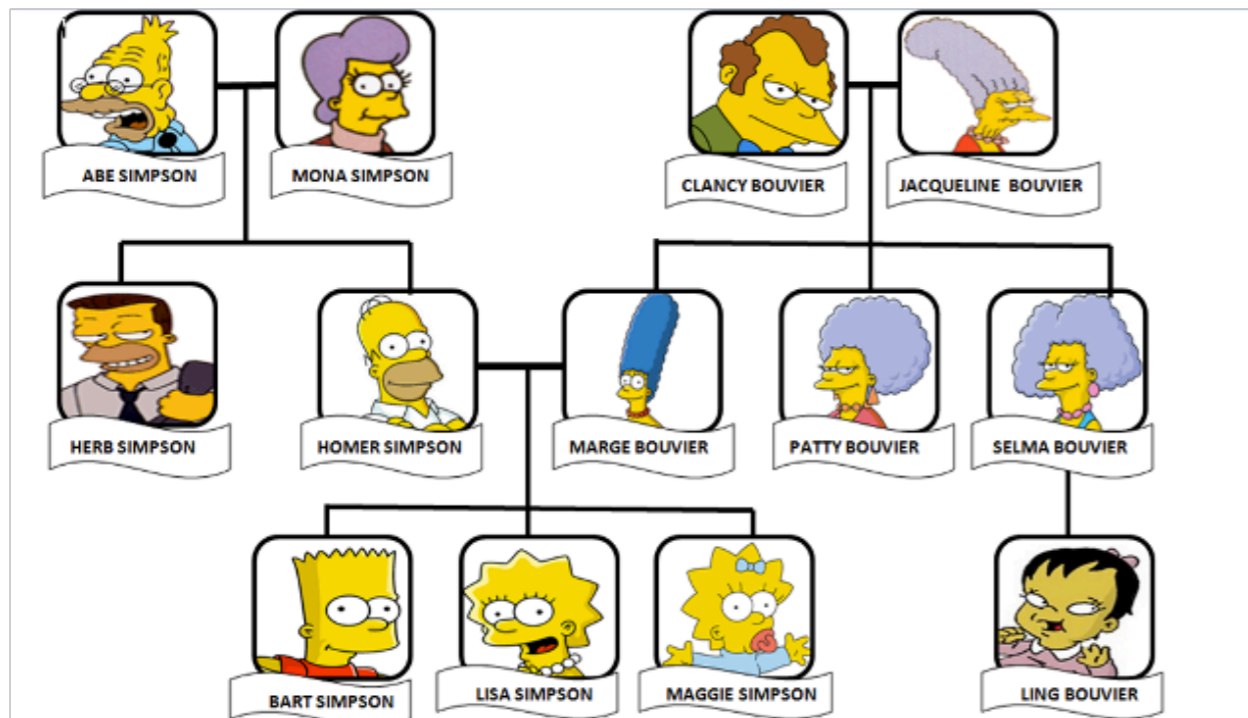
Parameter learning



Structure learning via parameter learning

[Su et al, 2019]

Enumerate (lots of) logical formulas from templates
and learn their probabilities/weights



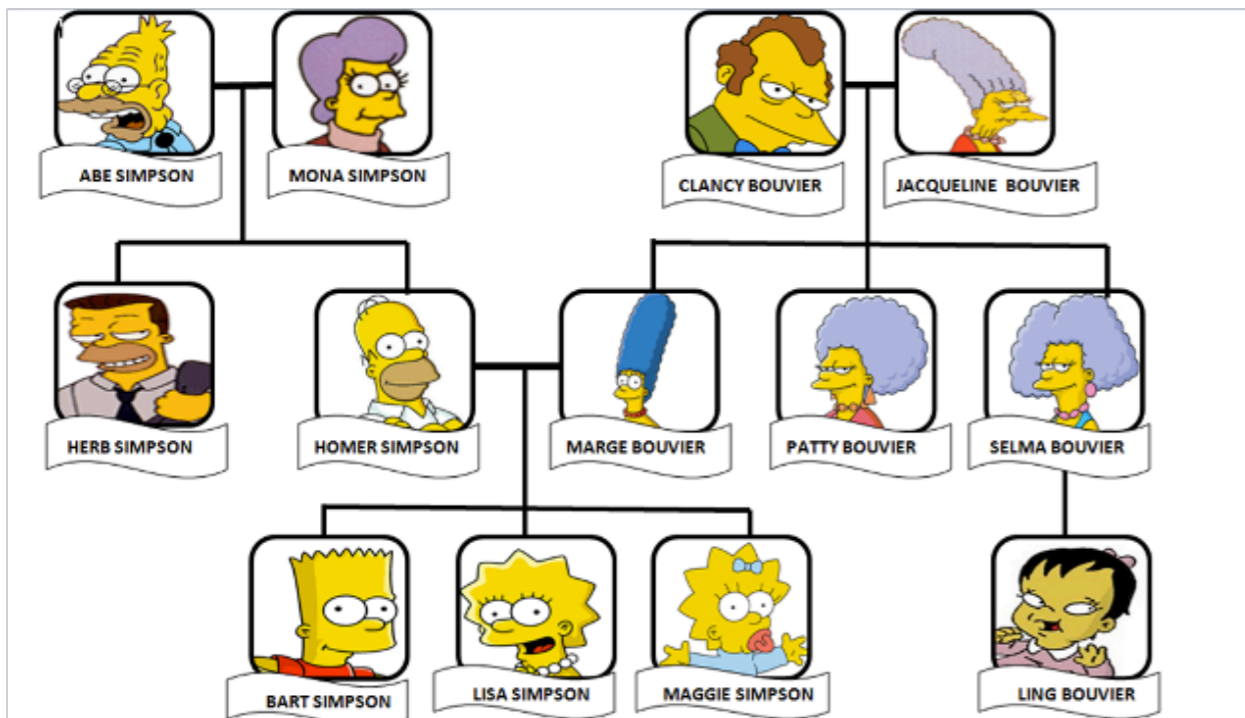
grandparent(abe,lisa).
grandparent(abe,bart).
grandparent(jacqueline,lisa).
grandparent(jacqueline,maggie.)



Program sketching

[Su et al, 2019]

Enumerate (lots of) logical formulas from templates
and learn their probabilities/weights



Program templates

$$T(X, Y) \leftarrow P(X, Y).$$

$$T(X, Y) \leftarrow P(Y, X).$$

$$T(X, Y) \leftarrow P(X, Z), Q(Z, Y).$$

Target: grandparent

Other predicates: father, mother



Program sketching

[Su et al, 2019]

Enumerate (lots of) logical formulas from templates
and learn their probabilities/weights

Program templates

$T(X,Y) \leftarrow P(X,Y).$

$T(X,Y) \leftarrow P(Y,X).$

$T(X,Y) \leftarrow P(X,Z), Q(Z,Y).$

$\text{grandparent}(X,Y) \leftarrow \text{father}(X,Y).$
 $\text{grandparent}(X,Y) \leftarrow \text{mother}(X,Y).$

$\text{grandparent}(X,Y) \leftarrow \text{father}(Y,X).$
 $\text{grandparent}(X,Y) \leftarrow \text{mother}(Y,X).$

$\text{grandparent}(X,Y) \leftarrow \text{mother}(X,Z), \text{mother}(Z,Y).$
 $\text{grandparent}(X,Y) \leftarrow \text{mother}(Y,X), \text{father}(Z,Y).$

.....

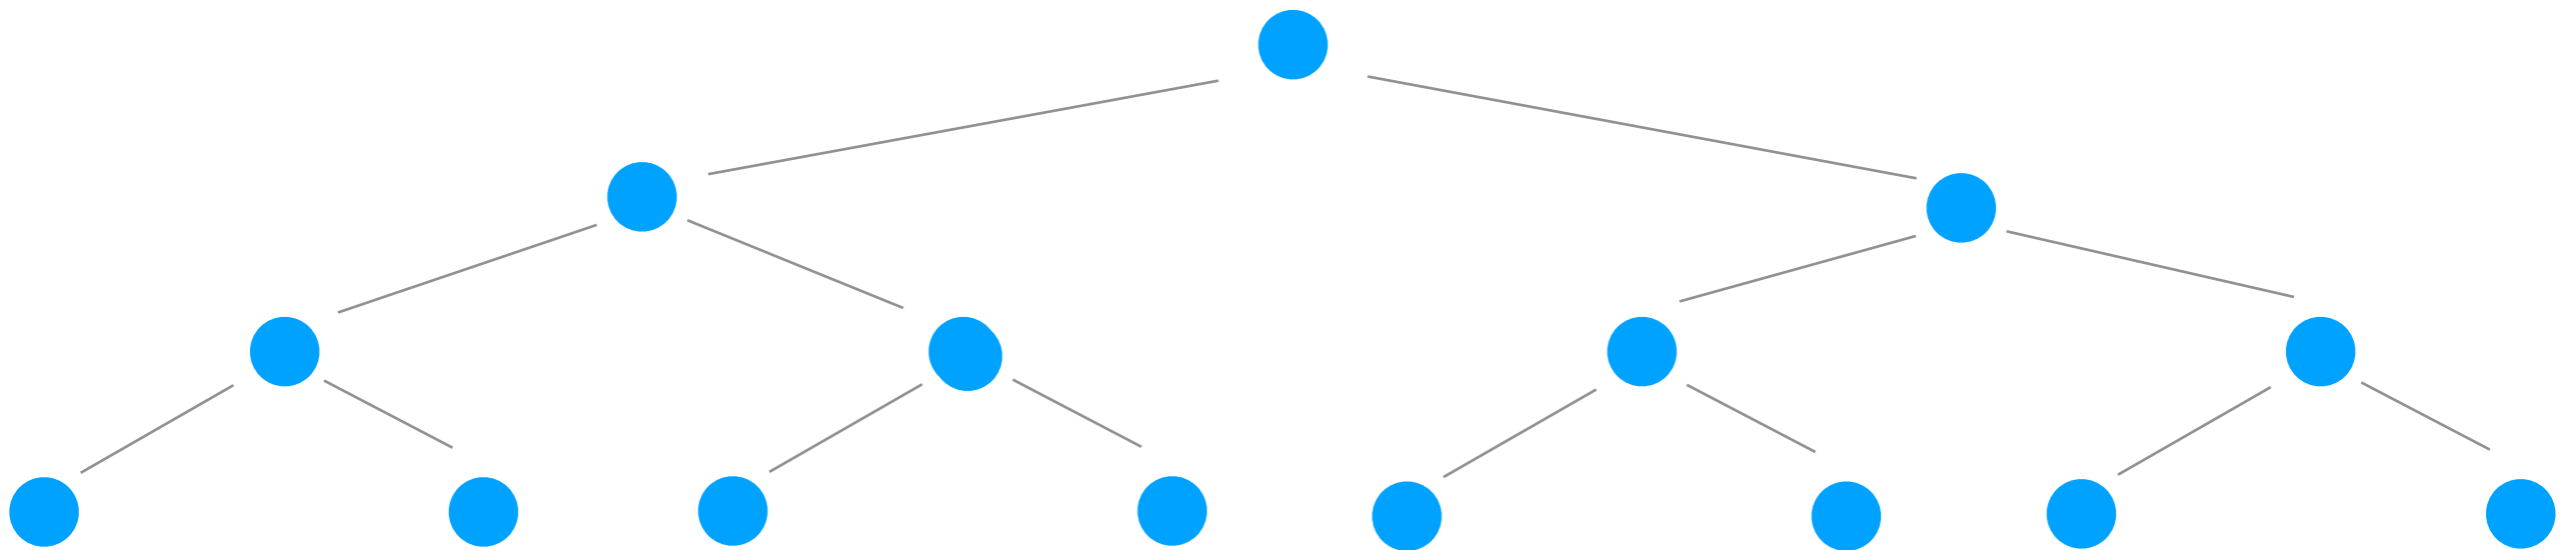
Target: grandparent

Other predicates: father, mother

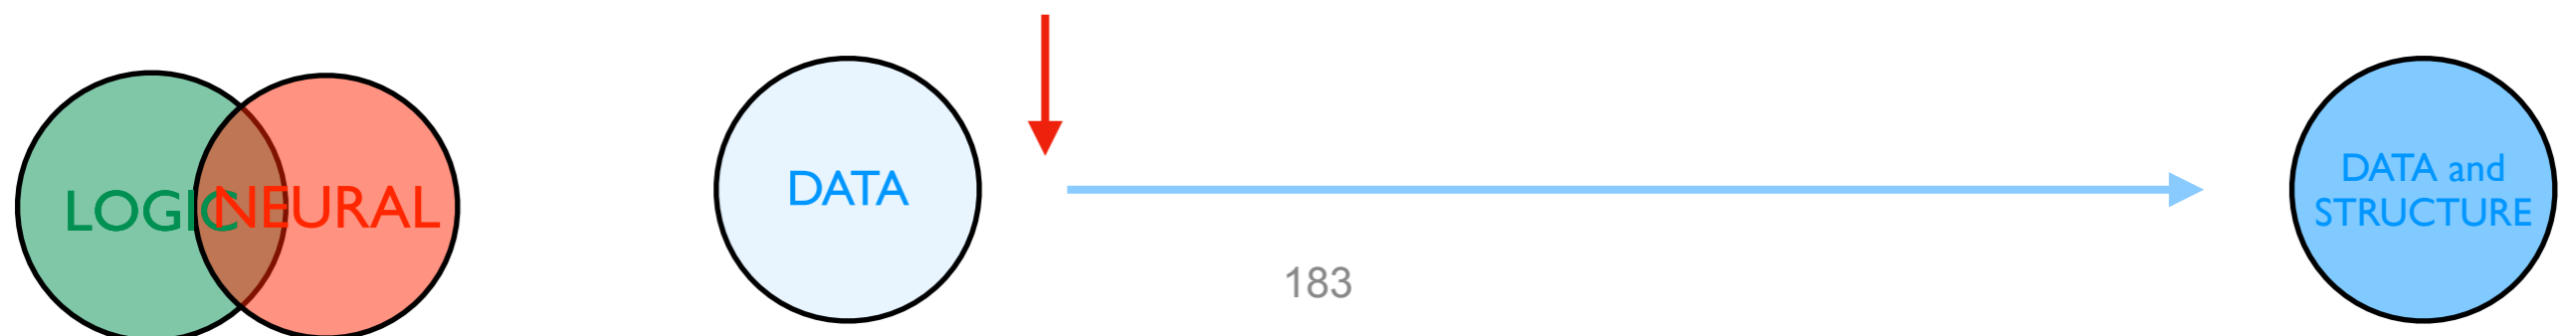


DeepCoder

[Balog et al, 2017]



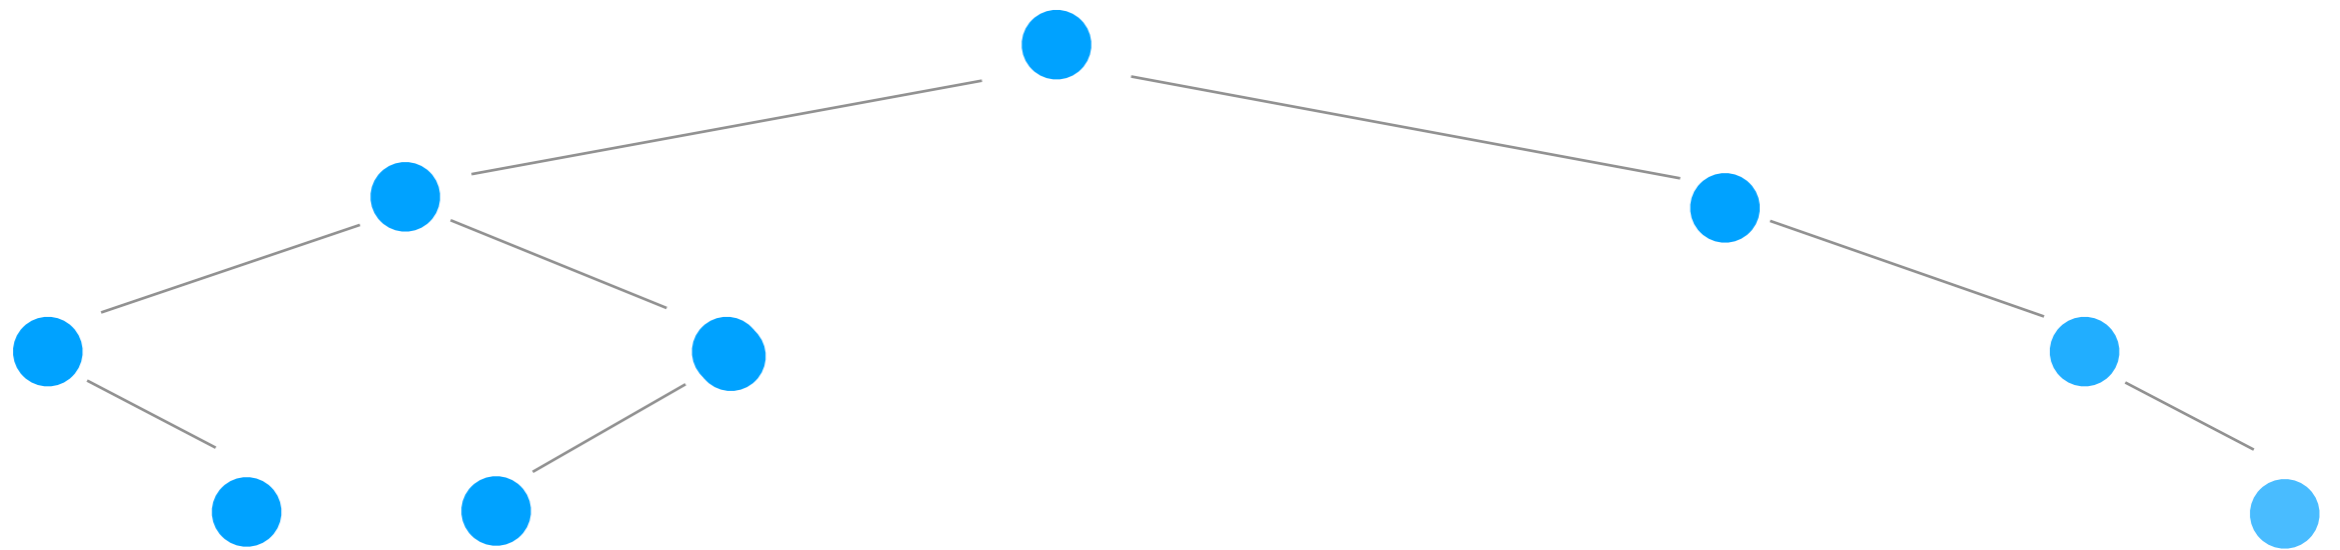
StarAI techniques search for clauses/rules systematically



DeepCoder

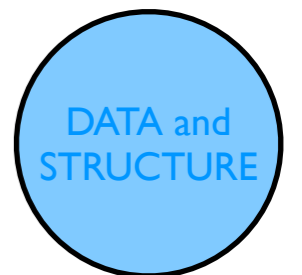
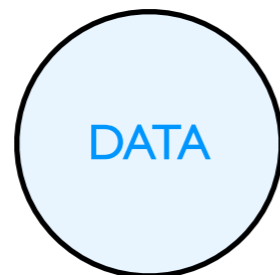
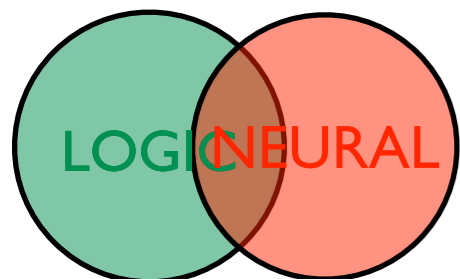
[Balog et al, 2017]

Preferences of learning 'primitives'



Explore the subpart of the space with primitives that are likely to solve the problem

likely to solve a problem = learned from data



DeepCoder

[Balog et al, 2017]

Preferences of learning ‘primitives’

Learn from pairs
(examples, program)

```
a ← [int]
b ← FILTER (<0) a
c ← MAP (*4) b
d ← SORT c
e ← REVERSE d
```

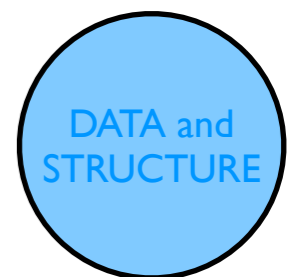
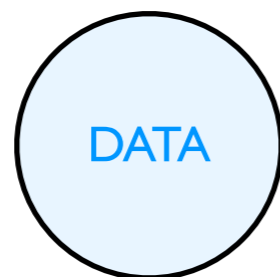
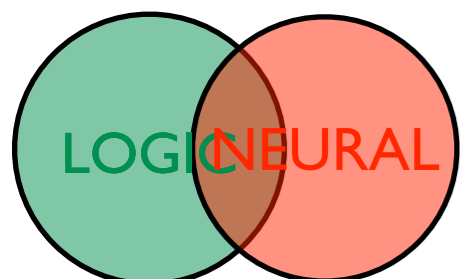
An input-output example:

Input:

[-17, -3, 4, 11, 0, -5, -9, 13, 6, 6, -8, 11]

Output:

[-12, -20, -32, -36, -68]



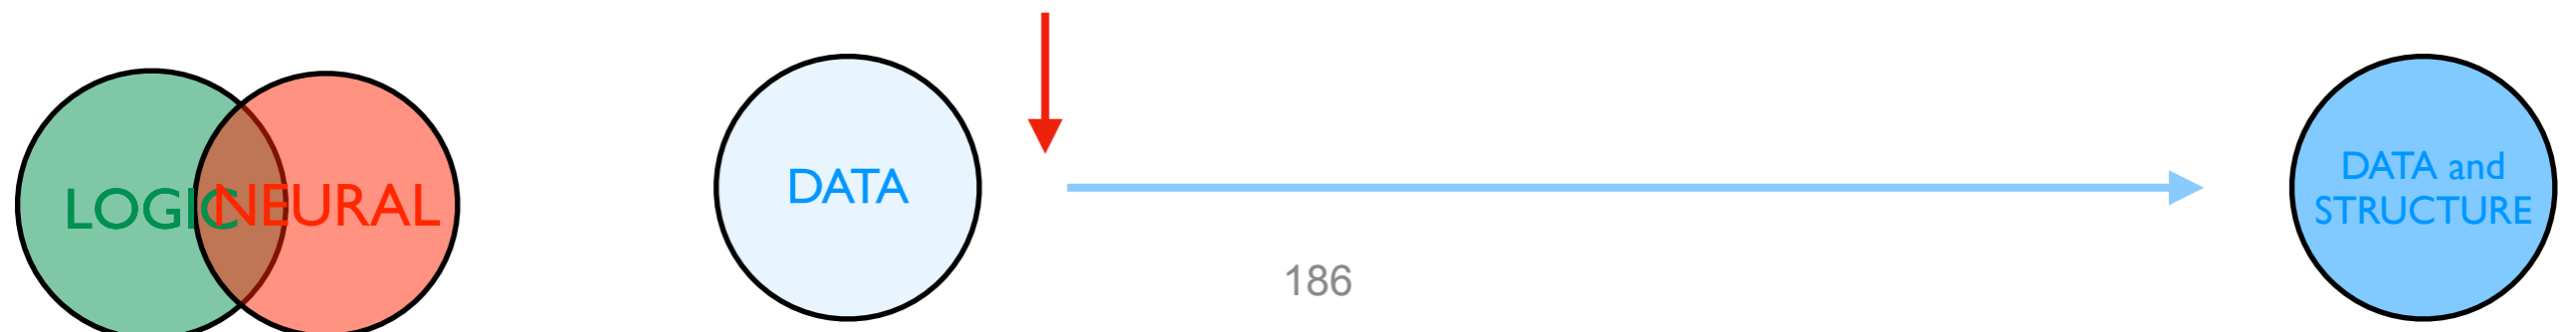
DreamCoder

[Ellis et al, 2018]

Distribution of primitives defines a generative model of programs

$$q(\text{programs} \mid \text{examples})$$

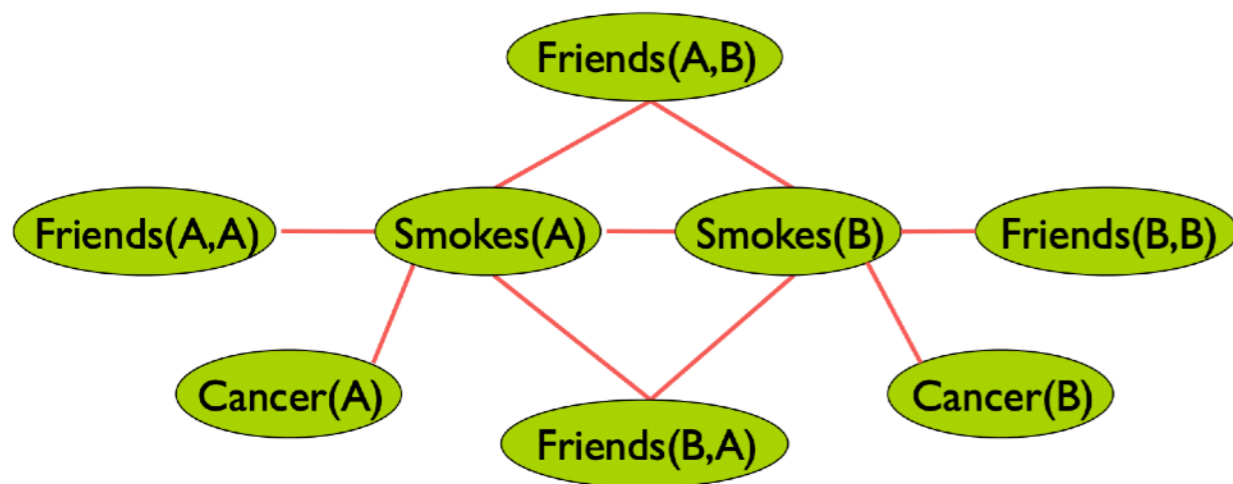
Neural network outputs the posterior distribution over programs likely to solve a specific task



Neural Markov Logic Networks

[Marra et al, 2020]

MLNs can be interpreted as log-linear models



$$P(X = x) = \frac{1}{Z} \prod_i \phi_i(x_{\{i\}})^{n_i(x)}$$

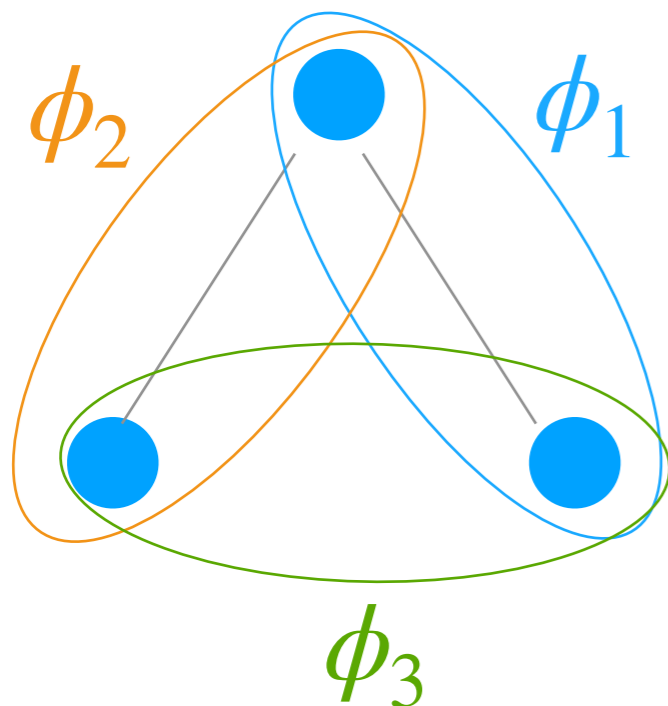
potentials come from formulas
provided by the expert
(cliques in Markov network)



Neural Markov Logic Networks

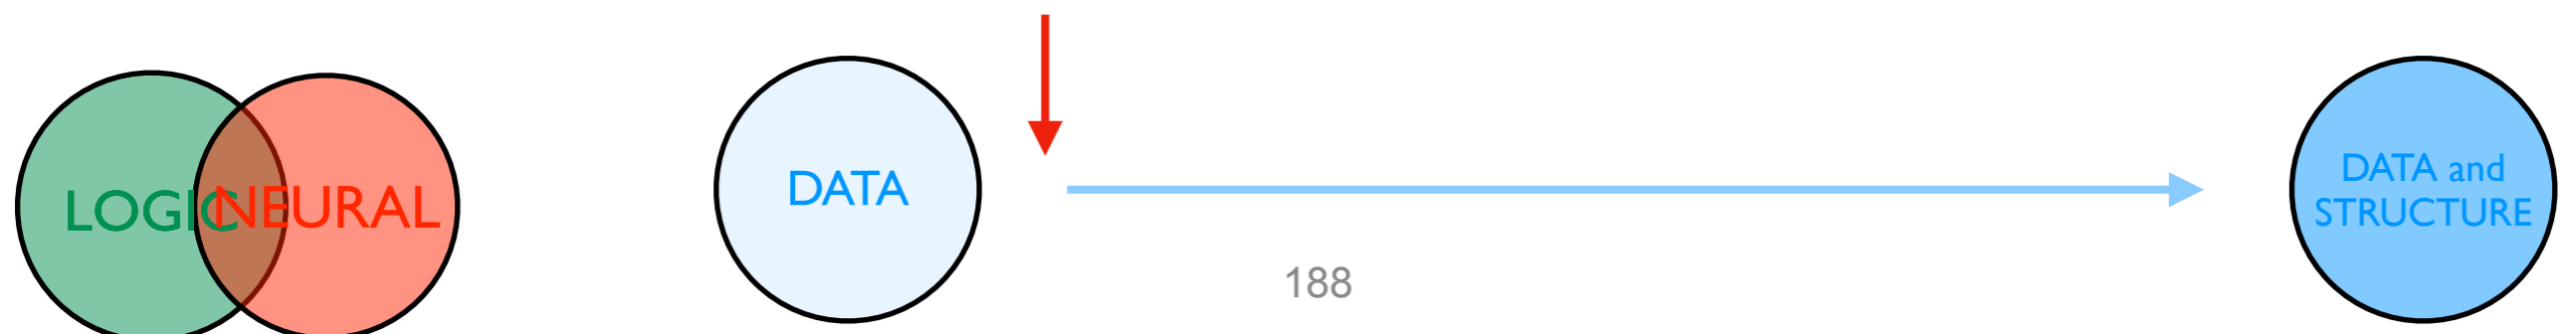
[Marra et al, 2020]

Learn neural potentials from fragments of data

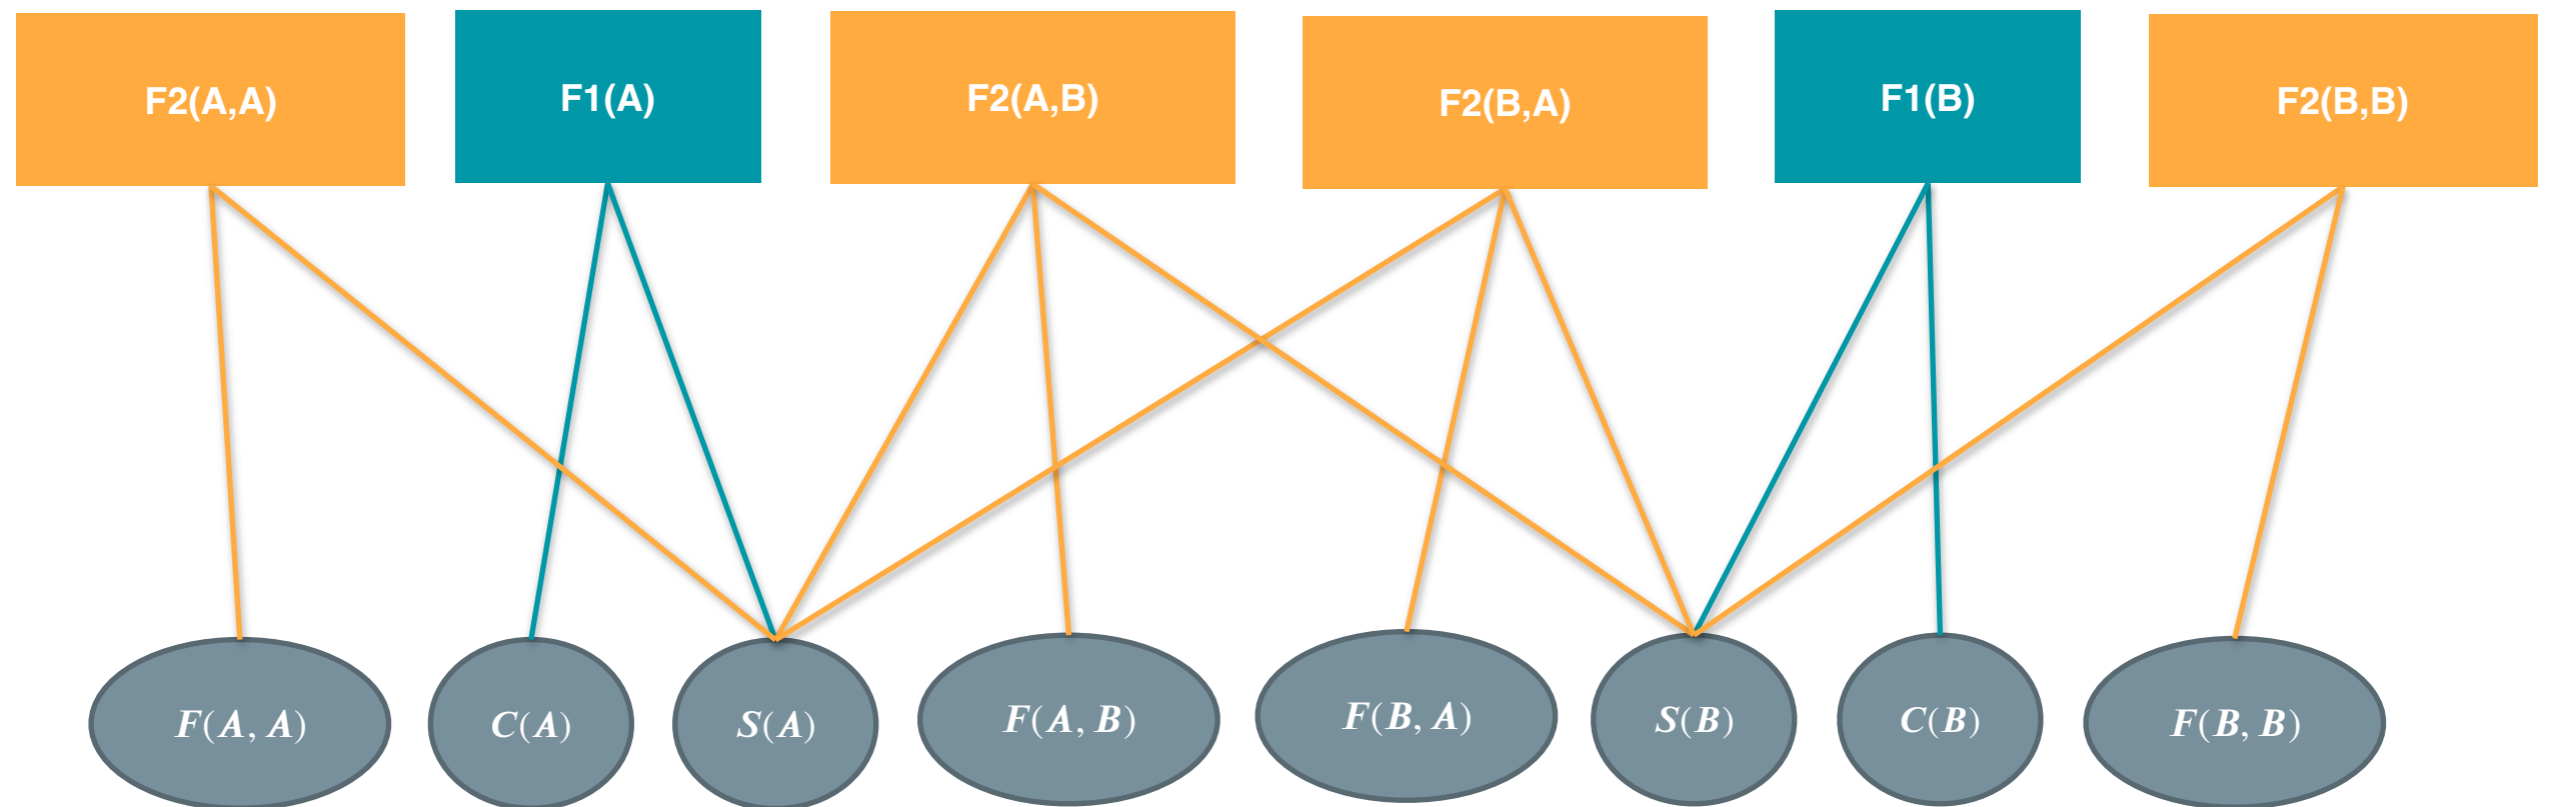


$$P(X = x) = \frac{1}{Z} \prod_i \phi_i(x_{\{i\}})^{n_i(x)}$$

potentials come from fragments of data (knowledge graph)



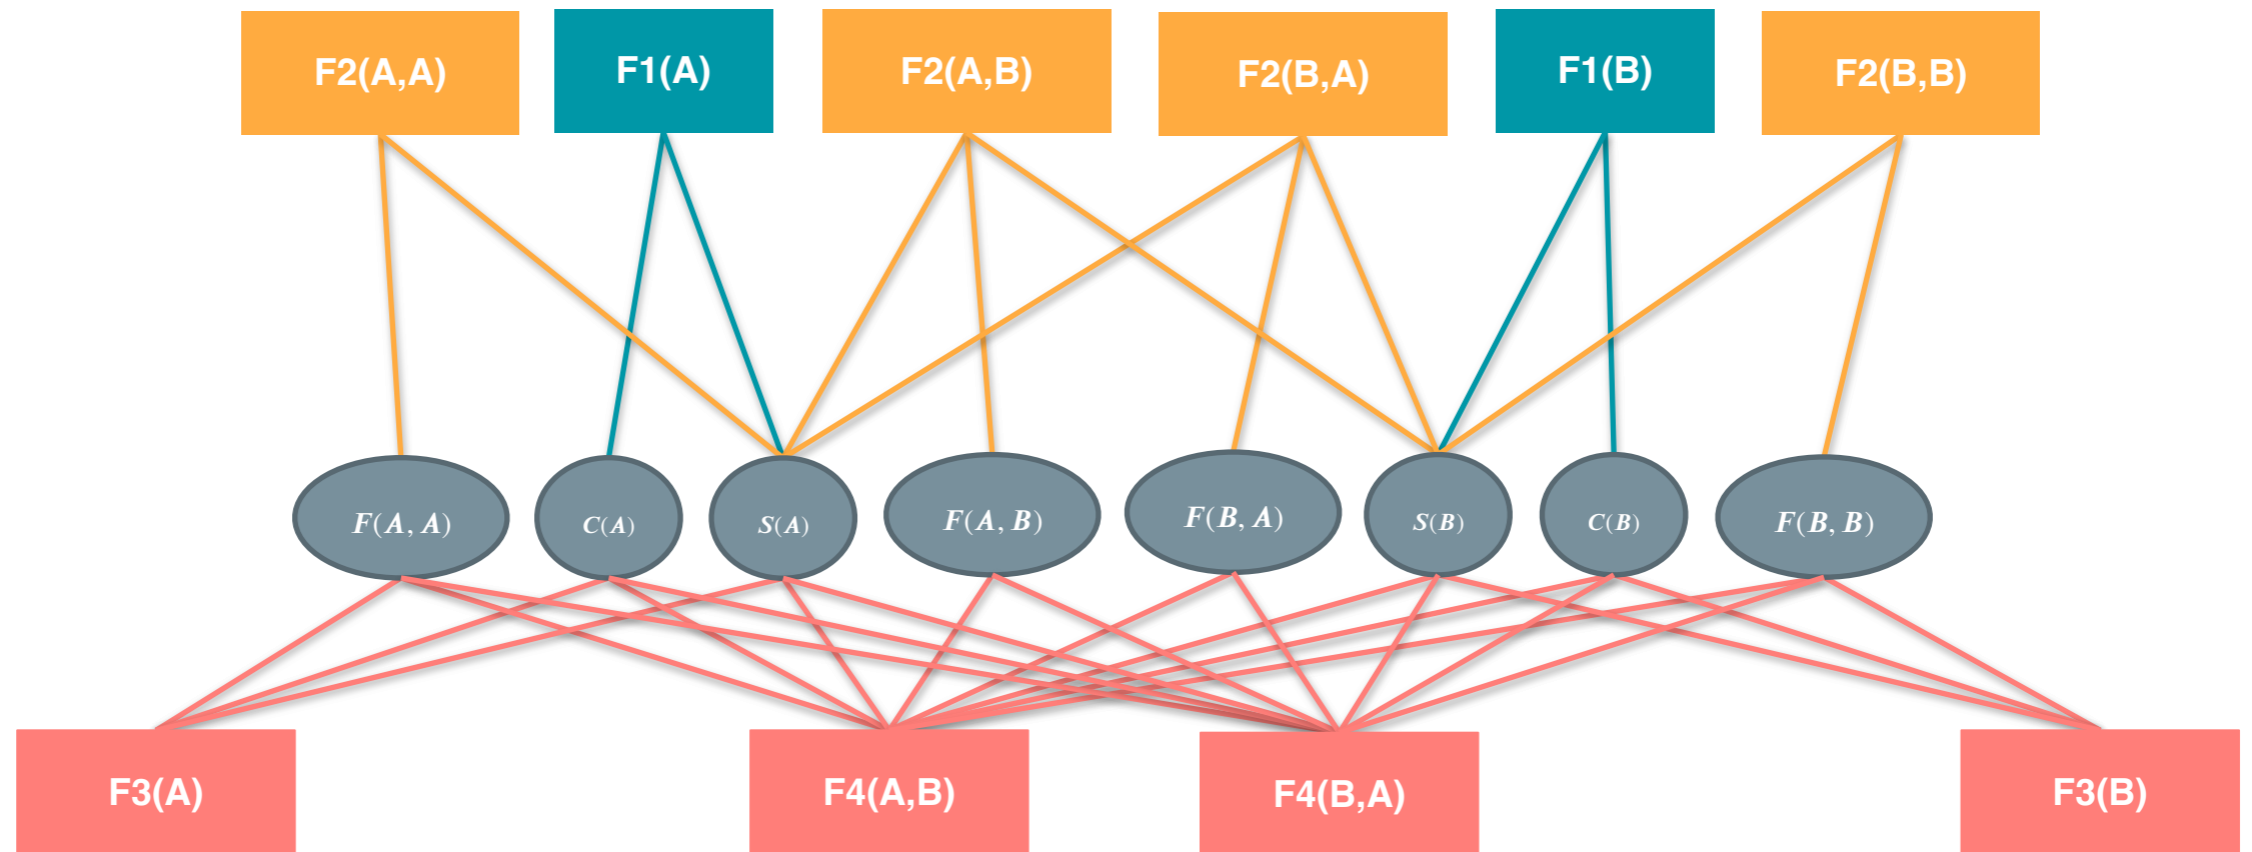
Markov Logic



represented as a factor graph

$$P(\text{Interpretation}) \propto \prod_i F_i(X, Y) = \prod_i \exp(w_i \mathbb{1}(\text{Interpretation} \models F_i))$$

Neural Markov Logic



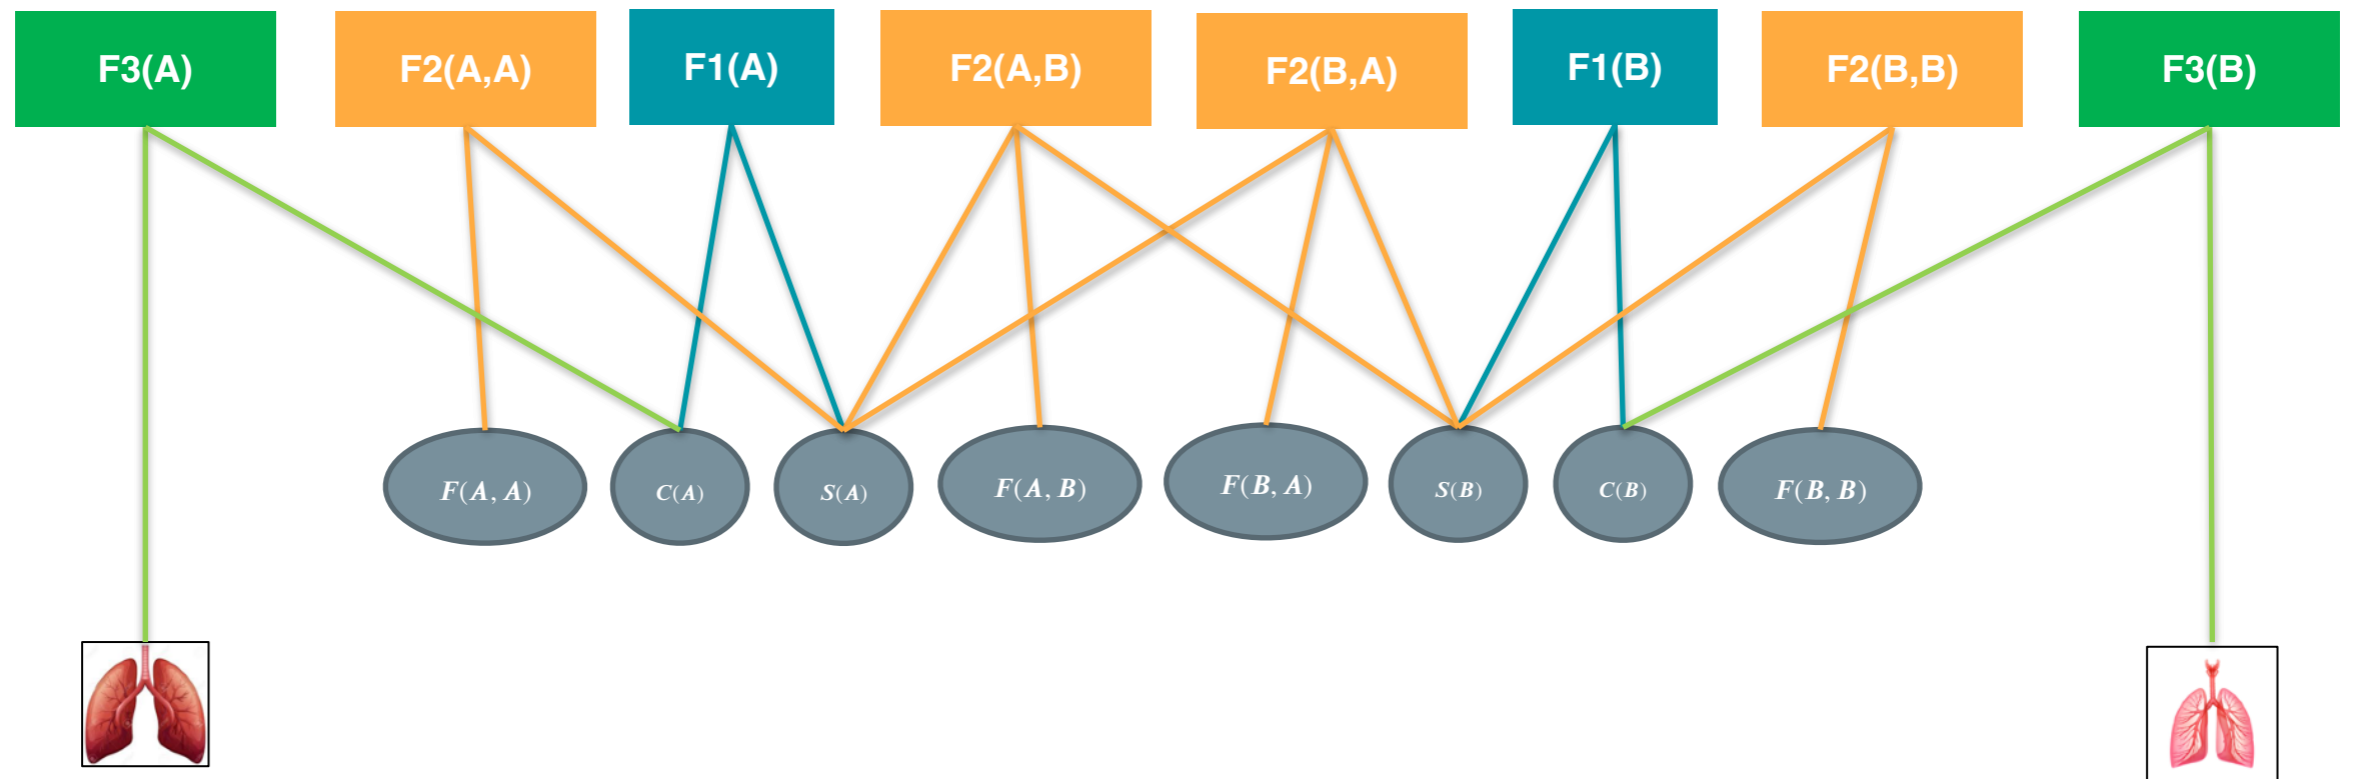
F3 and F4 are trainable factors

very much like in probabilistic graphical models and embeddings/hidden layers of a NN

F3 and F4 correspond in a sense to the logical rules in the other factors
this gives a kind of structure learning
F3 and F4 will not be “interpretable”

Relational Neural Machines

[Marra et al ECAI 20]



$$F3\left(\omega_{Cancer(Alice)}, \boxed{\text{Lung Image}}\right) = 1 - \left(CNN_{cancer}\left(\boxed{\text{Lung Image}}\right) - \omega_{Cancer(Alice)} \right)^2$$

The Neural Network is trained to become a FACTOR (or a part of it)

Pros

Cons

Neural guidance

makes discrete search tractable

lots of training data

Soft patterns

efficient learning

no explicit structure

Neural generation

focused combinatorial search

lots of training data

Sketching

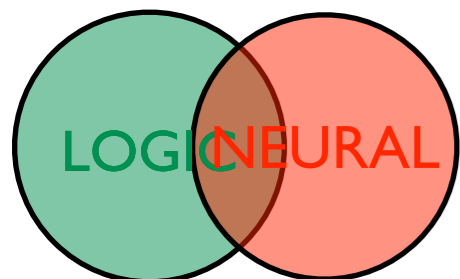
reduces combinatorial search

significant user effort

Structure via params

removes combinatorial search

spurious interactions



5. Learning

Key Messages

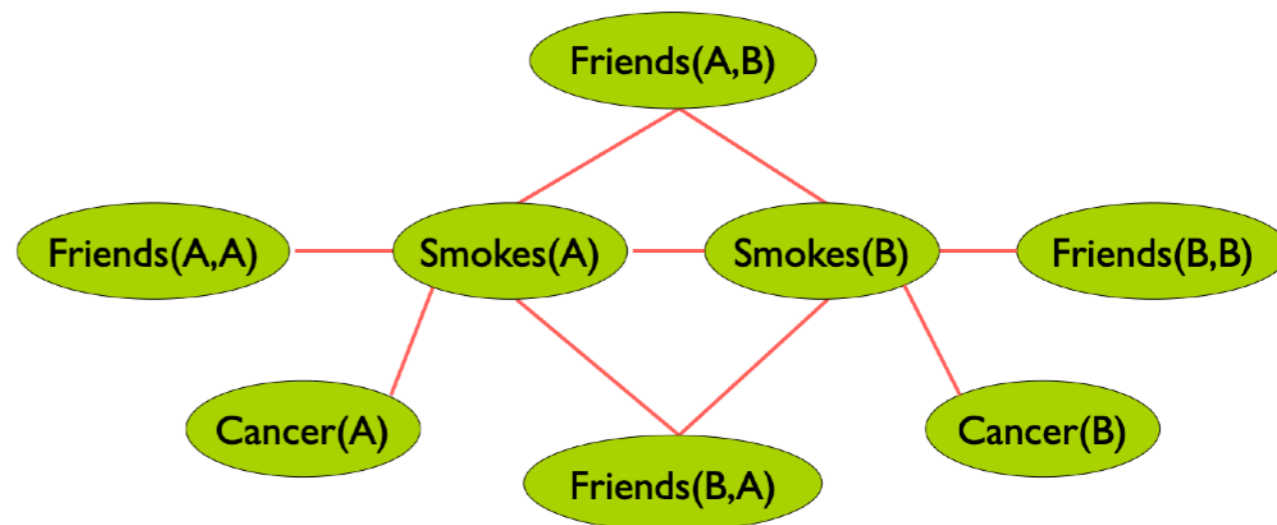
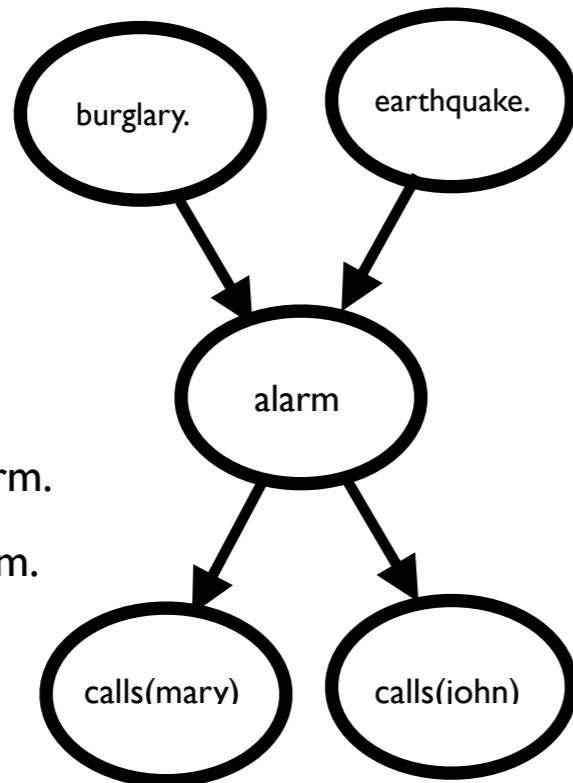
- Learning: finding logical formulas and estimating probabilities
- Structure learning: both formulas and probabilities
- Parameter learning: only probabilities
- Many flavours of learning in NeSy

The Seven Dimensions

1. Proof vs Model based
2. Directed vs Undirected
3. Type of Logic
4. Symbols vs Subsymbols
5. Parameter vs Structure Learning
6. Semantics
7. Logic vs Probability vs Neural

2. Directed vs Undirected the PGM / StarAI dimension

0.1 :: burglary.
 0.05 :: earthquake.
 alarm :- earthquake.
 alarm :- burglary.
 0.7::calls(mary) :- alarm.
 0.6::calls(john) :- alarm.



$$1.5 \quad \forall x \text{ Smokes}(x) \Rightarrow \text{Cancer}(x)$$

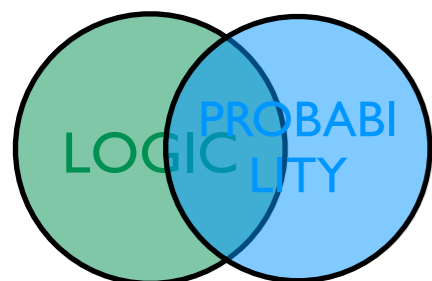
$$1.1 \quad \forall x, y \text{ Friends}(x, y) \Rightarrow (\text{Smokes}(x) \Leftrightarrow \text{Smokes}(y))$$

**Probabilistic Logic Programs
 ProbLog**

**directed
 Bayesian Net**

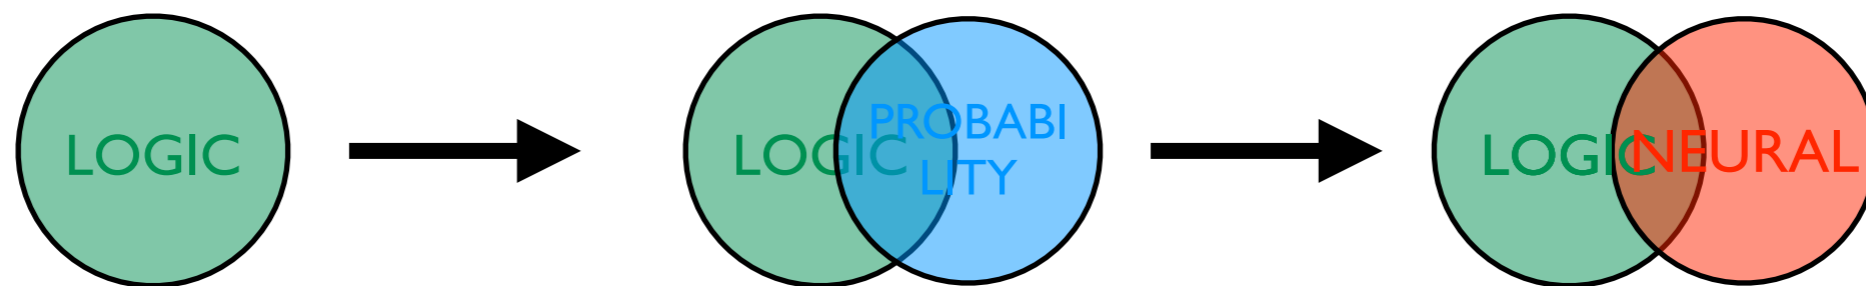
Markov Logic

**undirected
 Markov Net
 model theoretic**



key representatives

6. Semantics



6. Semantics

Key Messages

- StarAI and NeSy share the same underlying semantics
- Semantics can be described in terms of parametric circuits
- Differentiable semantics/circuits allows an easy integration
- NeSy models can be seen as neural reparameterization of StarAI models

Semantics

- In Logic, semantics is connected to the **interpretations** of logical sentences
- An interpretation assigns a **denotation** or a **value** to each symbol in that language.

“42(47)”

Semantics

- In Logic, semantics is connected to the **interpretations** of logical sentences
- An interpretation assigns a **denotation** or a **value** to each symbol in that language.

“42(47)”

42 is the property “being human” (or human/1)

47 is a constant referring to a particular human “Socrates”

human(Socrates) = True

Semantics

- We are interested in answering the following family of questions:

*Given a **sentence** of a propositional (or propositionalized through grounding) language, what is its **value**?*

The nature of what **value** is differs in the different semantics.

Semantics

For simplicity,

- **labelling function** is the function ℓ_S that assigns, to the **sentence Q**, the value **v** according to **semantics S**.

$$\ell_S(Q) = v$$

e.g.

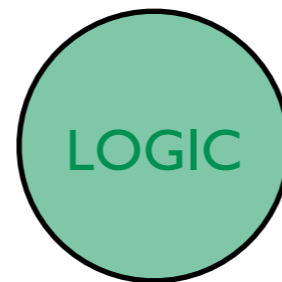
$$\ell_B(\text{human}(\text{socrates})) = \text{True}$$

$$\ell_F(\text{tall}(\text{john})) = 0.8$$

...

6. Semantics

Boolean logic



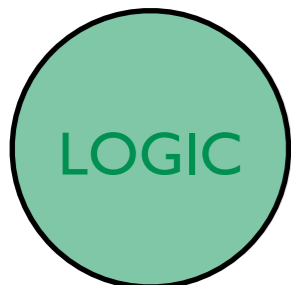
Semantics in Boolean Logic

- Defining a **semantics** for a propositional language L is about **assigning a truth value** to all the sentences of the logic
- Boolean truth values:

$\{True, False\}$

Three steps:

1. Truth values for propositions
2. Truth values for operators
3. Labelling formulas



Semantics in Boolean Logic

1. Providing the **labels** for propositions

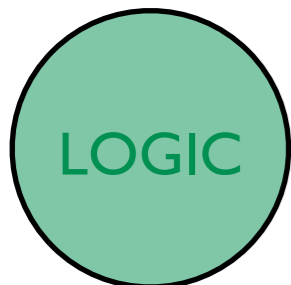
$L = \{burglary, earthquake, hears_alarm(john)\}$

$$\ell_B(burglary) = True$$

$$\ell_B(earthquake) = False$$

$$\ell_B(hears_alarm(john)) = True$$

*This is a **model** or a **possible world**, a “potential” assignment of truth values to all the propositional variables in the language.*



Semantics in Boolean Logic

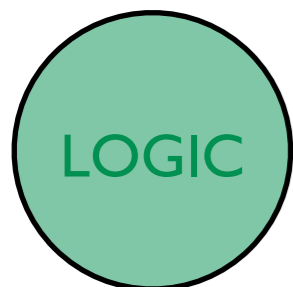
2. Providing the semantics for operators

p	q	$p \wedge q$
T	T	T
T	F	F
F	T	F
F	F	F

\mathcal{L}_B^\wedge

p	q	$p \rightarrow q$
T	T	T
T	F	F
F	T	T
F	F	T

$\mathcal{L}_B^\rightarrow$

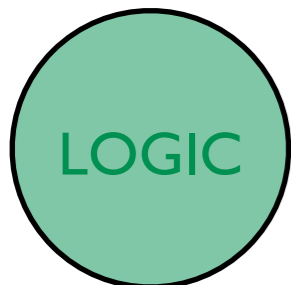


Semantics in Boolean Logic

3. The labels of **formulas** are defined **recursively** on the semantics of its components

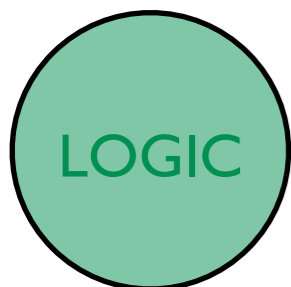
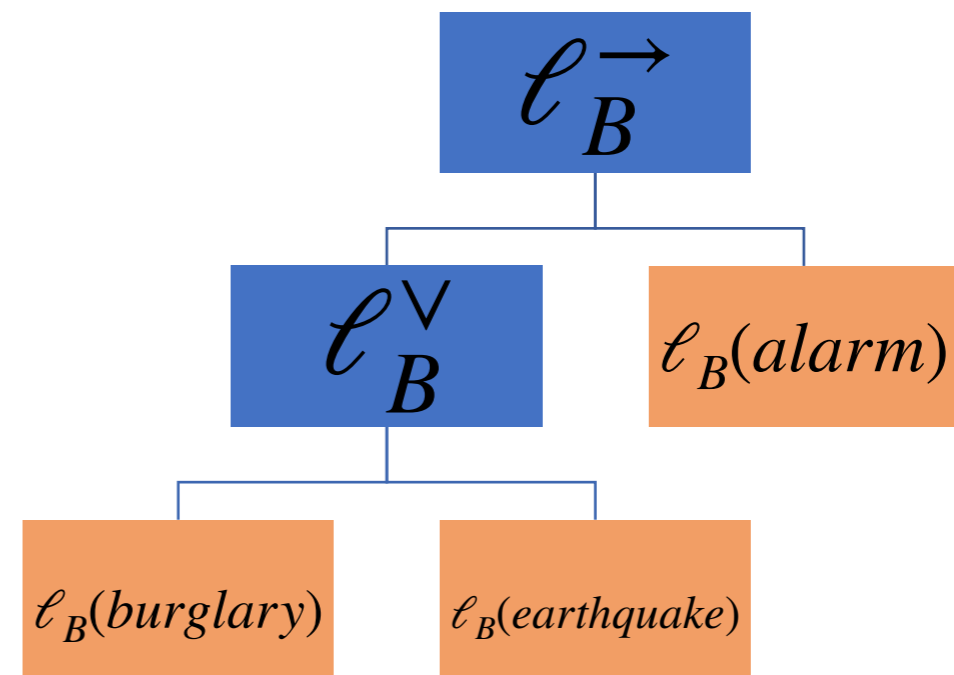
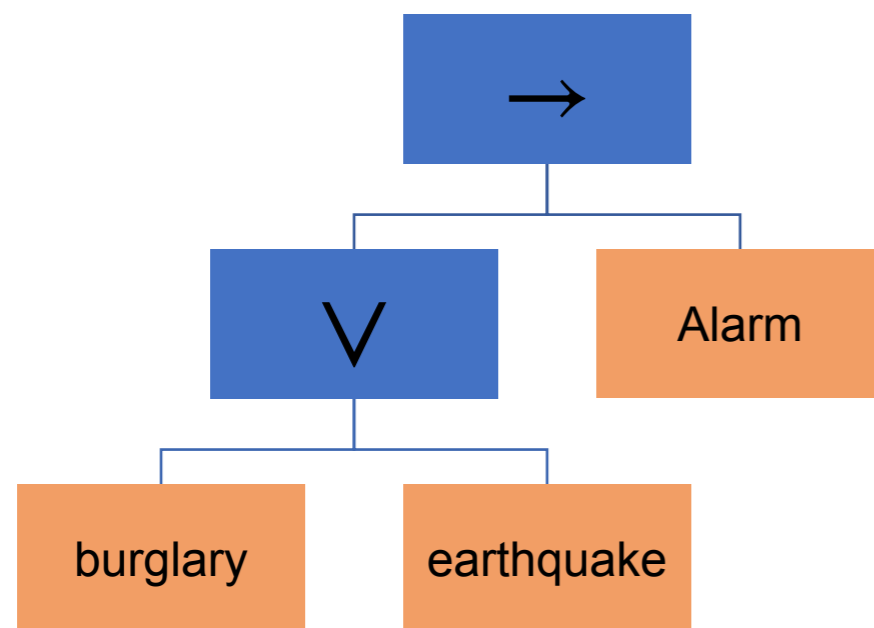
$$\ell_B(\text{earthquake} \wedge \text{burglary}) = \ell_B^\wedge(\ell_B(\text{earthquake}), \ell_B(\text{burglary}))$$

This recursive evaluation of formulas is said to be **extensional approach**.



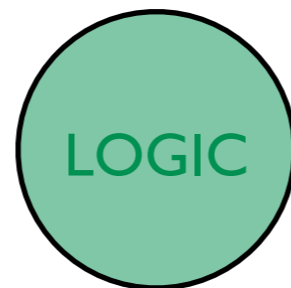
Semantics in Boolean Logic

- Consider: $(\text{burglary} \vee \text{earthquake}) \rightarrow \text{alarm}$

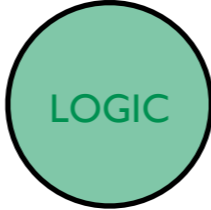


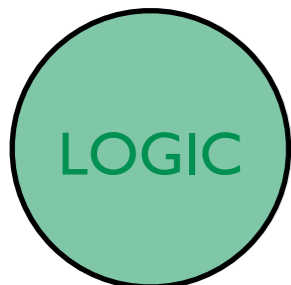
6. Semantics

Fuzzy logic



Semantics in Fuzzy Logic

- Still a pure **logic** semantics: 
- There are many fuzzy logics
- Here we are interested in a subclass, in particular *t-norm fuzzy logic*



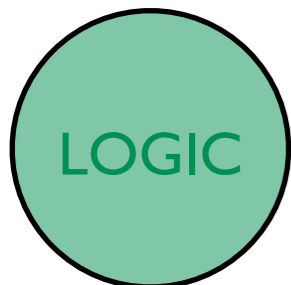
Semantics in Fuzzy Logic

- Defining a **semantics** for a propositional fuzzy language L is again about **assigning a membership degree** to all the sentences of the logic
- Fuzzy **truth/membership degrees**:

$$\ell_F: L \rightarrow [0,1]$$

Three steps:

1. Labels for propositions
2. Labels for operators
3. Labels for formulas



Semantics in Fuzzy Logic

1. Providing the **labels** for propositions

$L = \{burglary, earthquake, hears_alarm(john)\}$

$$\ell_F(burglary) = 0.9$$

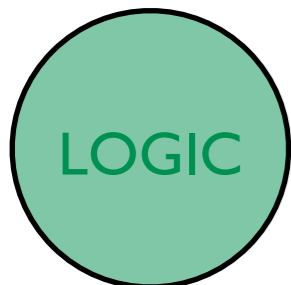
$$\ell_F(earthquake) = 0.1$$

$$\ell_F(hears_alarm(john)) = 0.8$$

Note: $\ell_F(earthquake) = 0.1$ -> very mild earthquake,

(\neq probability of earthquake = 0.1)

fuzzy is a measure of **intensity/vagueness** not of uncertainty



Semantics in Fuzzy Logic

2. Providing the labels for operators: t-norm theory

- A **t-norm** is a binary function that extends the **conjunction** to the continuous case

$$t : [0,1] \times [0,1] \rightarrow [0,1]$$

- There are **3 fundamental t-norms**:
 - **Lukasiewicz t-norm**: $t_L(x, y) = \max(0, x + y - 1)$
 - **Goedel t-norm**: $t_G(x, y) = \min(x, y)$
 - **Product t-norm**: $t_P(x, y) = x \cdot y$



LOGIC

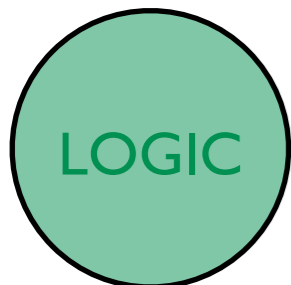
They are the continuous version of truth tables!!

Semantics in Fuzzy Logic

- All the other operators can be derived from the t-norm

	Product	Łukasiewicz	Gödel
$x \wedge y$	$x \cdot y$	$\max(0, x + y - 1)$	$\min(x, y)$
$x \vee y$	$x + y - x \cdot y$	$\min(1, x + y)$	$\max(x, y)$
$\neg x$	$1 - x$	$1 - x$	$1 - x$
$x \Rightarrow y$ ($x > y$)	y/x	$\min(1, 1 - x + y)$	y

They are the continuous version of truth tables!!



Semantics in Fuzzy Logic

3. The labels of **formulas** is defined **recursively** on the semantics of its components

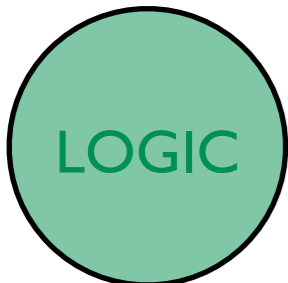
$$\ell_F(\textit{burglary} \rightarrow \textit{alarm}) = \ell_F^{\vec{}}(\ell_F(\textit{burglary}), \ell_F(\textit{alarm}))$$

This recursive evaluation of formulas is said to be **extensional approach**.

e.g.

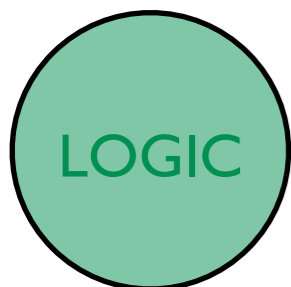
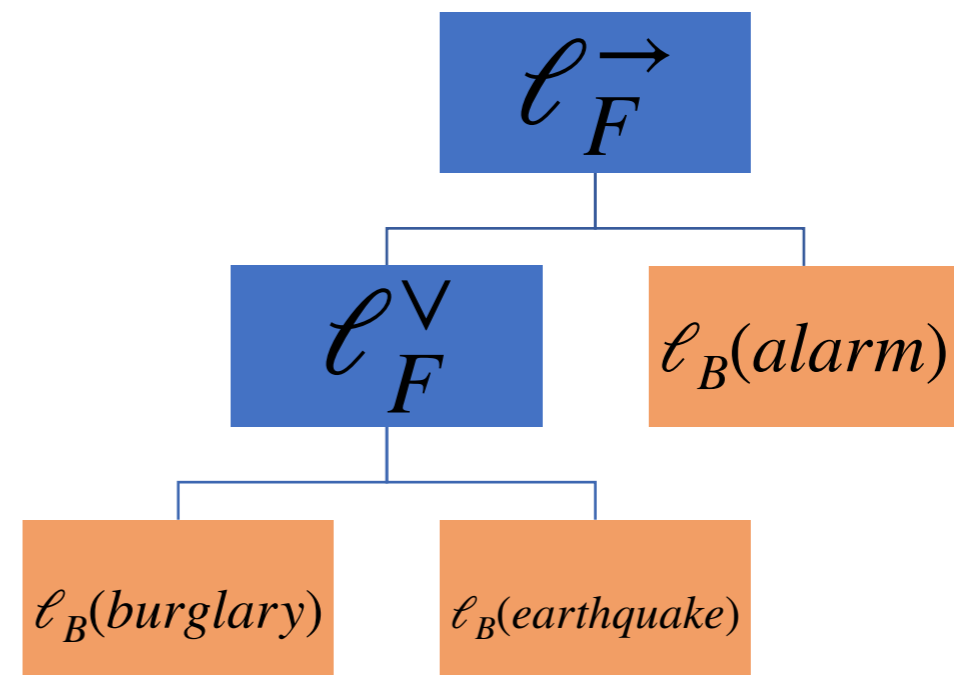
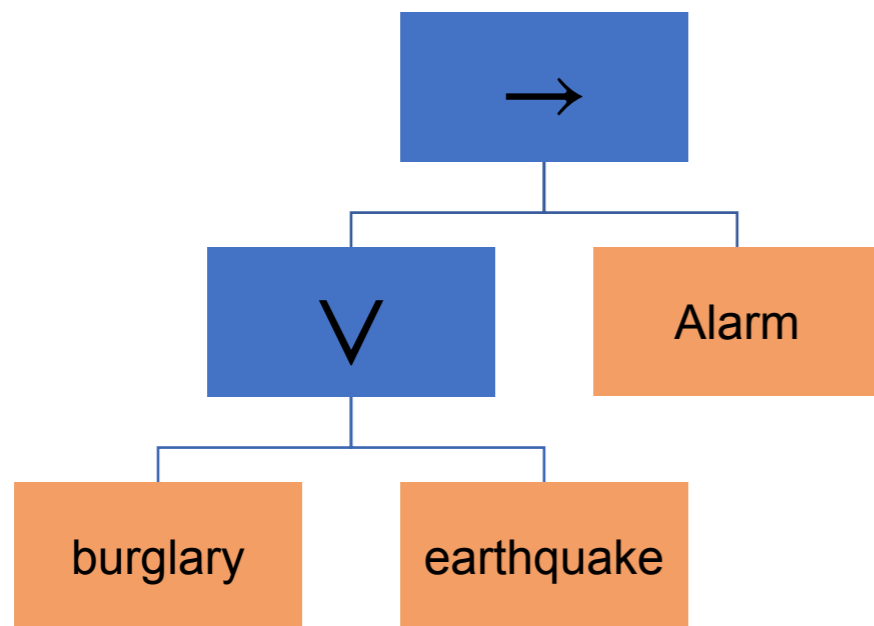
$$\ell_F(\textit{burglary}) = 0.9, \ell_F(\textit{alarm}) = 0.3,$$

$$\ell_F^{\vec{}} = \min(1, 1 - x + y) = \min(1, 1 - 0.9 + 0.3) = 0.4$$



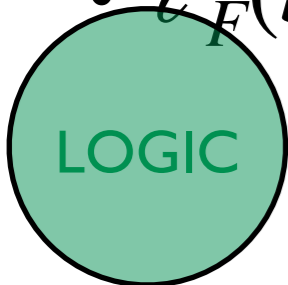
Semantics in Fuzzy Logic

- Consider: $(\text{burglary} \vee \text{earthquake}) \rightarrow \text{alarm}$



Fuzzy Logic Semantics

- Most common t-norms are:
 - **Continuous**
 - **Differentiable** -> This turns to be one of the reason of their adoption in NeSY
- Convex fragments of the logic can be defined (Giannini et al, 2019)
- But, $\ell_F(\text{human}(\text{Socrates})) = 0.5$??????
- $\ell_F(\text{bat}(\text{Socrates})) = 0.5$



Fuzzy vs Boolean

- Fuzzy and Boolean have different properties
- When fuzzy is used as a “relaxation” (**fuzzification**) of Boolean **undesired effects** can happen.

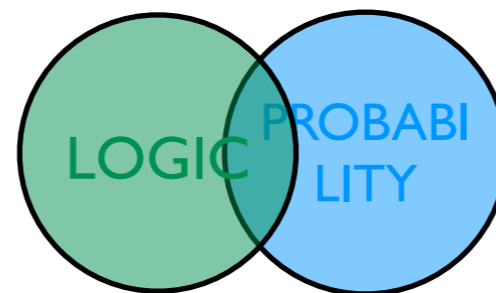
- Suppose: $A \vee B \vee C \vee D \vee E = 1$

- Satisfying assignments (Lukasiewicz)

- $A = B = C = D = E = 1$ (all true)
- $A = 1, B = C = D = E = 0$ (at least one true)
- $A = B = C = D = E = 0.2$

Semantics

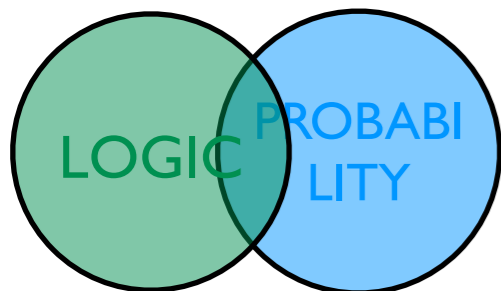
Probabilistic logic



Probabilistic Logic Semantics

Given a proposition language L , the basic idea is to introduce a **probability function** p :

$$p : L \rightarrow [0,1]$$



Probabilistic Logic Semantics

Two steps:

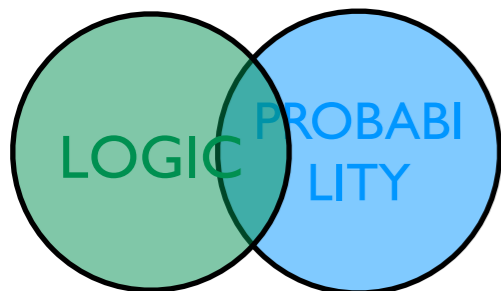
- Define a **probability distribution over interpretations / worlds (i.e. boolean semantics)**

$$p(\ell_B(x_1), \dots, \ell_B(x_n))$$

(E.g. $p(\ell_B(\text{burglary}) = \text{True}, \ell_B(\text{earthquake}) = \text{False}, \dots)$)

- Define a **the probability of sentence Q of L:**

$$p(Q) = \sum_{\ell_B(x_1), \dots, \ell_B(x_n) \models Q} p(\ell_B(x_1), \dots, \ell_B(x_n))$$



Probabilistic Logic Semantics

Problog

0.1 :: burglary. (B)

0.05 :: earthquake. (E)

0.6 :: hears_alarm(john). (H)

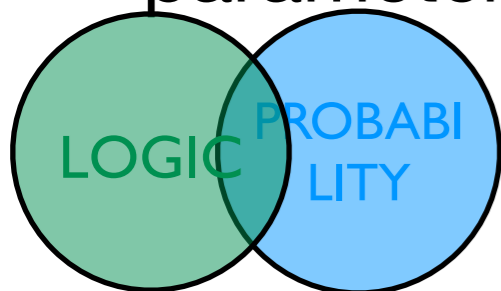
alarm :- earthquake.

alarm :- burglary.

calls(john) :- alarm, hears_alarm(john)

$$p(\ell_B(x_1), \dots, \ell_B(x_n)) = \prod_{i:\ell_B(x_i)=True} p(x_i) \prod_{i:\ell_B(x_i)=False} (1 - p(x_i))$$

parameters = the **labels for propositions** (i.e. probabilistic facts)



Probabilistic Logic Semantics

Problog

e.g. in Problog:

B	E	H	p(B,E,H)
F	F	F	0,342
F	F	T	0,513
F	T	F	0,018
F	T	T	0,027
T	F	F	0,038
T	F	T	0,057
T	T	F	0,002
T	T	T	0,003

0.1 :: burglary. (B)

0.05 :: earthquake. (E)

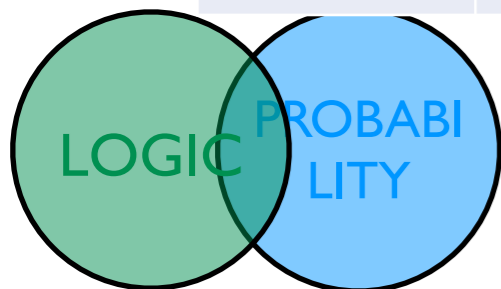
0.6 :: hears_alarm(john). (H)

alarm :- earthquake.

alarm :- burglary.

calls(john) :- alarm, hears_alarm(john)

$$0.1 \times 0.05 \times (1 - 0.6)$$



Probabilistic Logic Semantics

Markov Logic

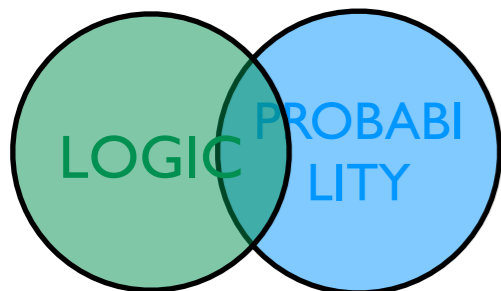
1.5 : calls(Mary) <- hears_alarm(Mary), alarm

2.0 : alarm <- earthquake

0.5 : alarm <- burglary

Weight formula 1 if α is True otherwise 0

$$p(\ell_B(x_1), \dots, \ell_B(x_n)) = \frac{1}{Z} \exp \left(\sum_{\alpha} w_{\alpha} \ell_B(\alpha) \right)$$



Probabilistic Logic Semantics

Markov Logic

1.5 : `calls(Mary) <- hears_alarm(Mary), alarm`

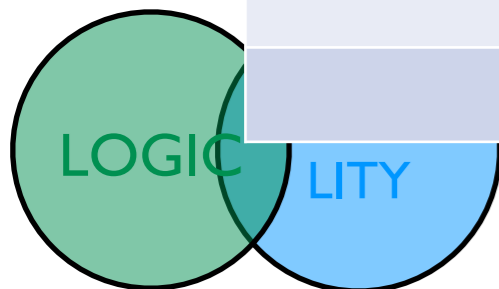
2.0 : `alarm <- earthquake`

0.5 : `alarm <- burglary`

B	E	A	H	C	p
T	F	T	T	T	0,05
T	F	T	T	F	0,01
...

$\propto \exp(1.5 + 2.0 + 0.5)$

$\propto \exp(0 + 2.0 + 0.5)$

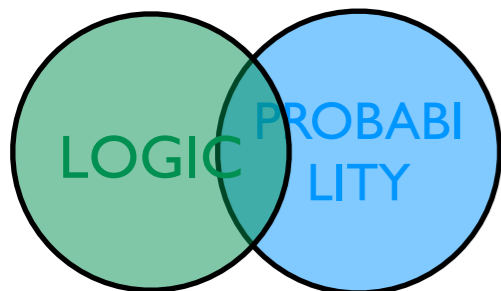


Probabilistic Logic Semantics

Given any **sentence** Q of the propositional language L , with variables x_1, \dots, x_n :

$$\ell_P(Q) = \sum_{\ell_B(x_1), \dots, \ell_B(x_n) \models Q} p(\ell_B(x_1), \dots, \ell_B(x_n))$$

WMC - Weighted Model Counting
(for both ProbLog and Markov Logic)



Probabilistic Logic Semantics

For example:

B	E	H	p(B,E,H)
F	F	F	0,342
F	F	T	0,513
F	T	F	0,018
F	T	T	0,027
T	F	F	0,038
T	F	T	0,057
T	T	F	0,002
T	T	T	0,003

0.1 :: burglary. (B)

0.05 :: earthquake. (E)

0.6 :: hears_alarm(john). (H)

alarm :- earthquake.

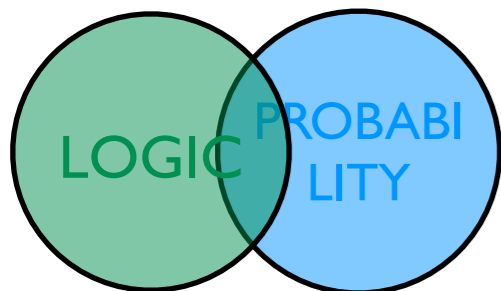
alarm :- burglary.

calls(john) :- alarm, hears_alarm(john)

Query = burglary ^ hears_alarm(john)

$$Q = B \wedge H$$

$$p(Q) = 0.06$$



Probabilistic Logic Semantics

For example:

B	E	H	p(B,E,H)
F	F	F	0,342
F	F	T	0,513
F	T	F	0,018
F	T	T	0,027
T	F	F	0,038
T	F	T	0,057
T	T	F	0,002
T	T	T	0,003

0.1 :: burglary. (B)

0.05 :: earthquake. (E)

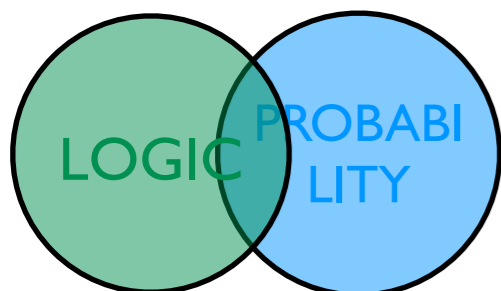
0.6 :: hears_alarm(john). (H)

alarm :- earthquake.

alarm :- burglary.

$$Q = (B \wedge H) \vee E$$

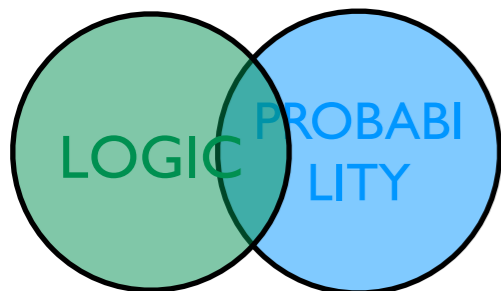
$$\ell_p(Q) = 0.105$$



Probabilistic Logic Semantics

Probabilistic Semantics is different from a pure logic semantics

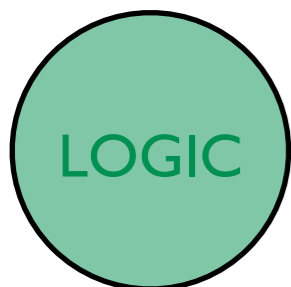
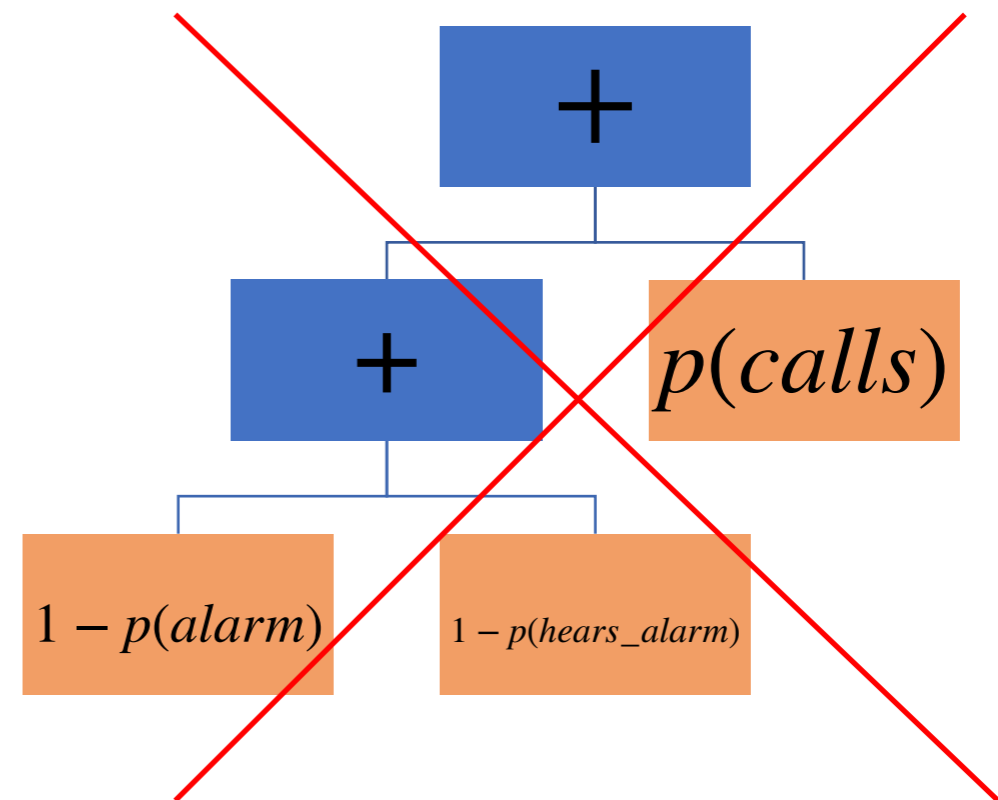
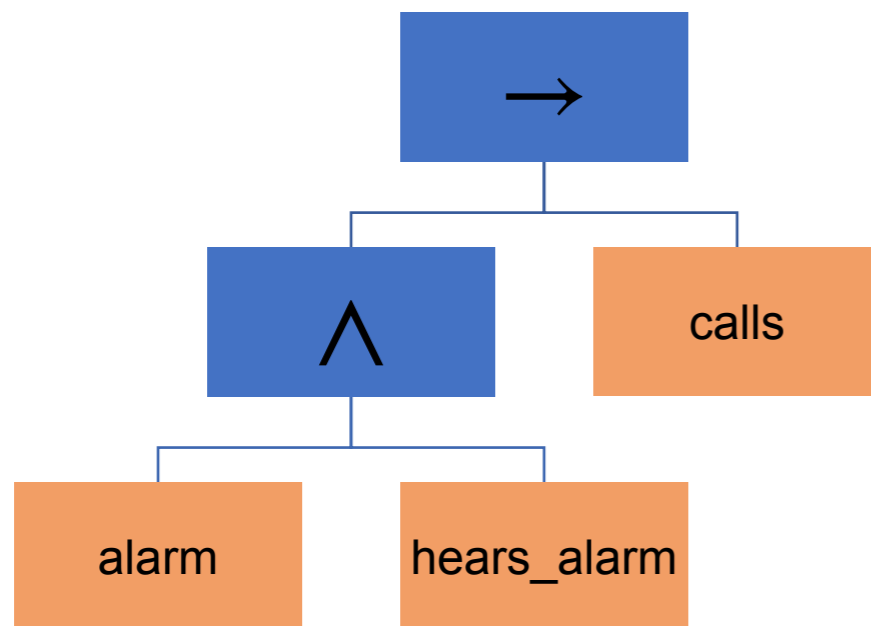
1. It is built on top of a logical semantics; $p(\ell_B(x_1), \dots, \ell_B(x_n))$.
2. Probability is **NOT extensional**, the probability of a formula
 - A. **cannot be defined recursively** by the probabilities of its arguments
 - B. requires **WMC**



Probabilistic Logic Semantics

$$(alarm \wedge hears_alarm) \rightarrow calls$$

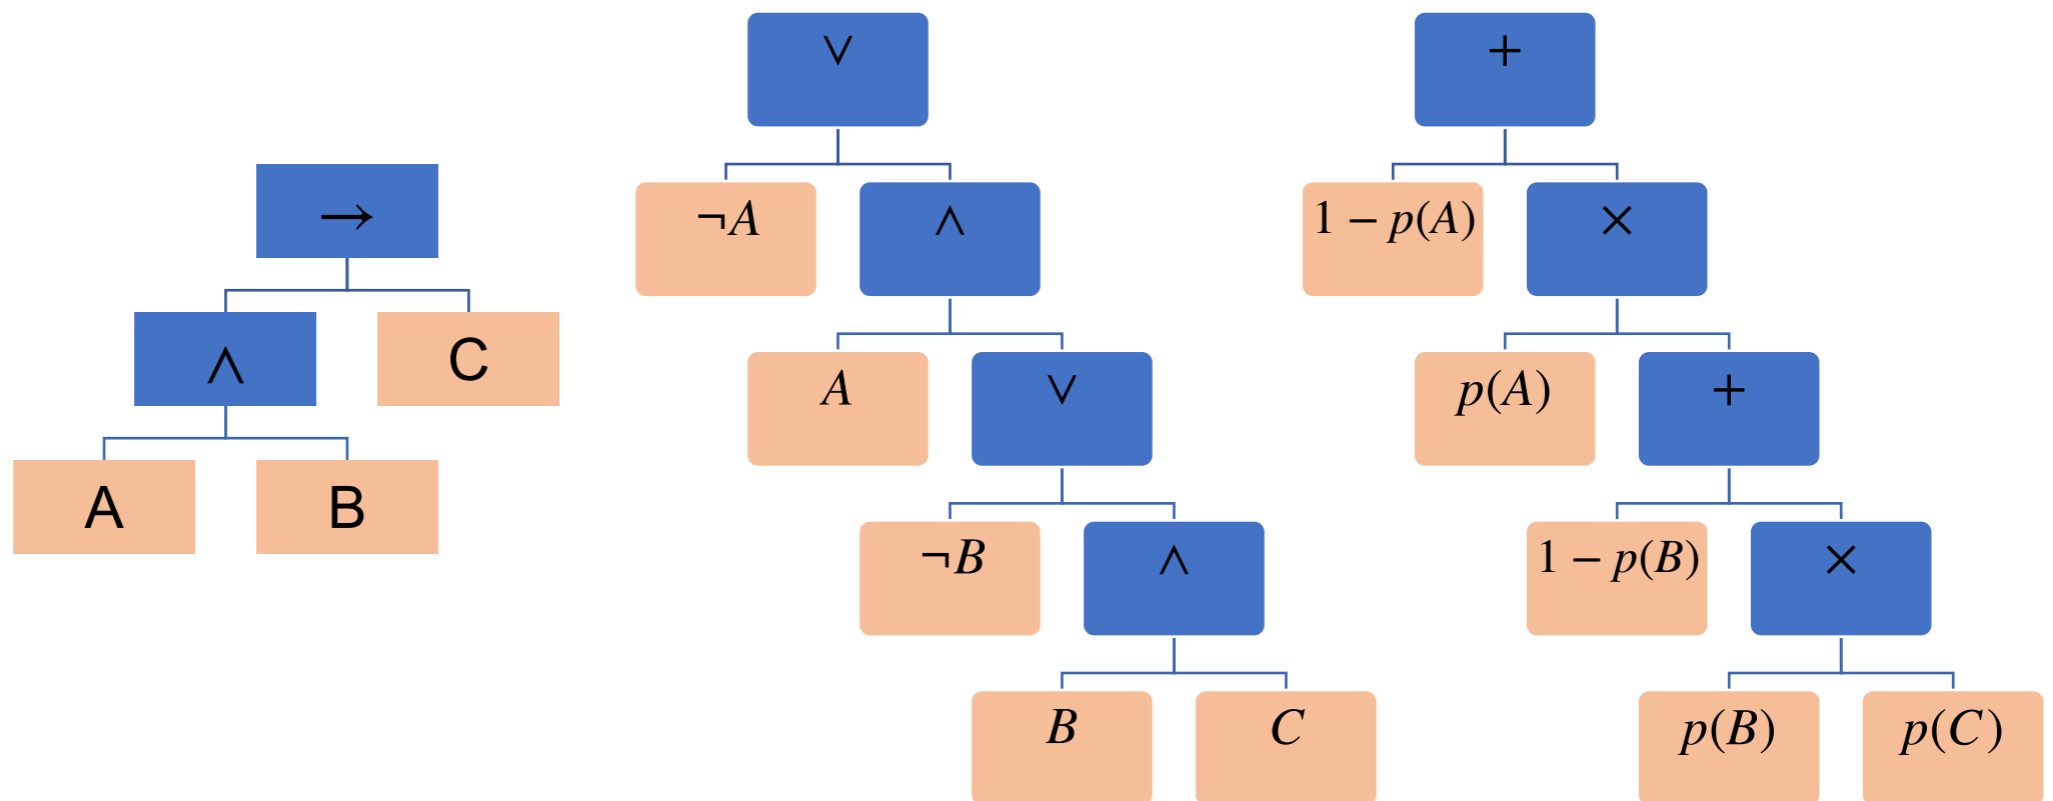
- Consider:



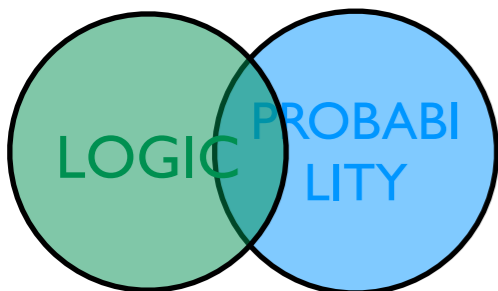
Probabilistic Logic Semantics

$$\ell_P(Q) = \sum_{\ell_B(x_1), \dots, \ell_B(x_n) \models Q} p(\ell_B(x_1), \dots, \ell_B(x_n))$$

$(A \wedge B) \rightarrow C$



Knowledge Compilation

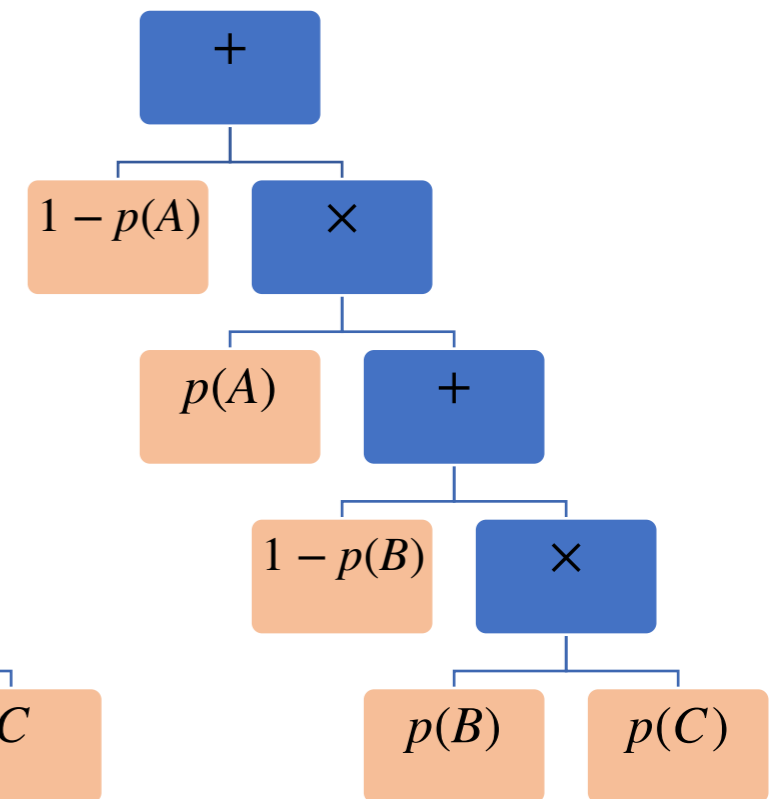
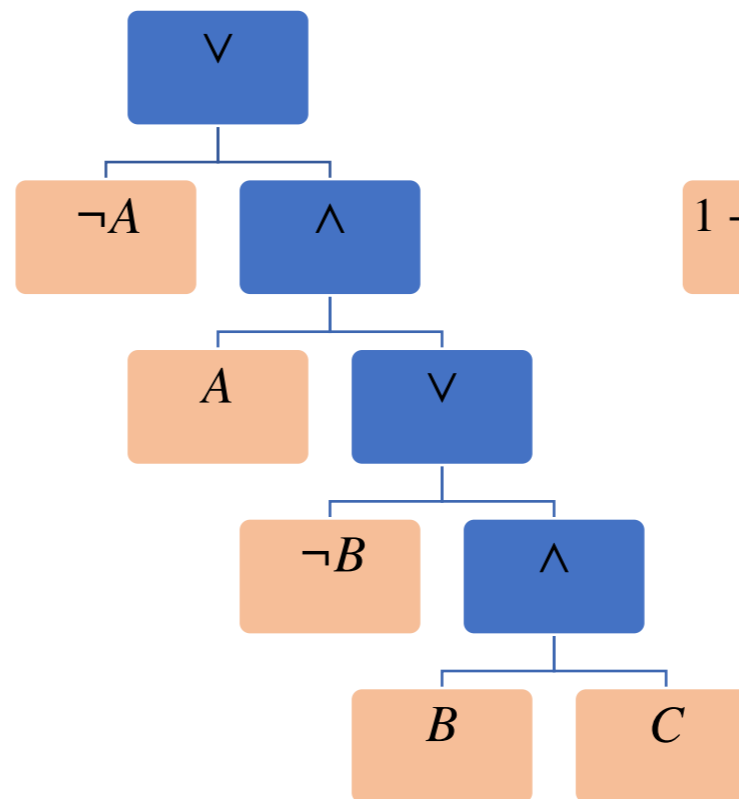
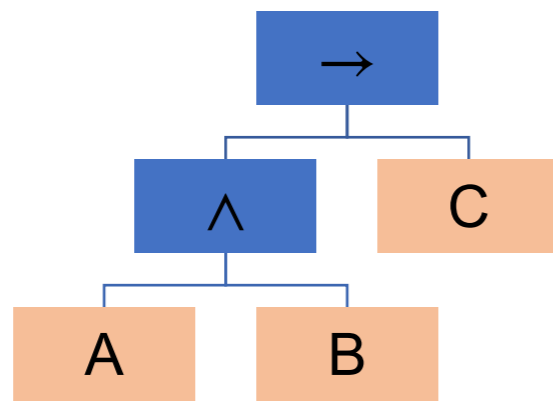


The probabilistic structure is now explicit in the compiled formula.

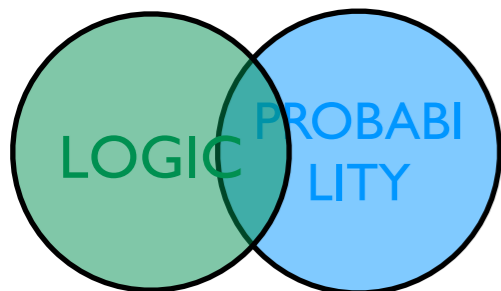
Probabilistic Logic Semantics

- Consider:

$$(A \wedge B) \rightarrow C$$



The circuit is differentiable!



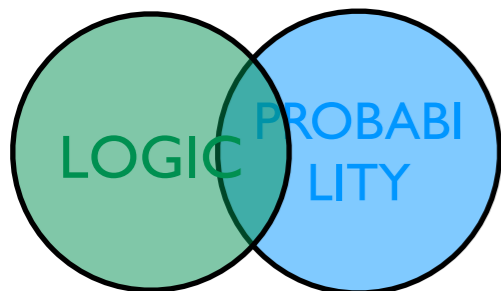
Probabilistic Logic Semantics

- WMC:

$$p(Q) = \sum_{\ell_B(x_1), \dots, \ell_B(x_n) \models Q} p(\ell_B(x_1), \dots, \ell_B(x_n))$$

- Another important inference task in MPE inference (connected to maxSAT)

$$\ell_B^\star(x_1), \dots, \ell_B^\star(x_n) = \max_{\ell_B(x_1), \dots, \ell_B(x_n) \models Q} p(\ell_B(x_1), \dots, \ell_B(x_n))$$



Boolean vs Fuzzy vs Probability

- Boolean and Fuzzy logic are two **alternative** logical semantics
- Probability is a semantics that is built on top of a logical one (i.e. “which is the **probability** of a given **truth assignments** / world?”)
- Can we have a probabilistic fuzzy logic as well?

Probabilistic Soft Logic (PSL)

Bach, Stephen H., et al. *JMLR* 2017

- Let's start by an example of a Markov Logic Network:

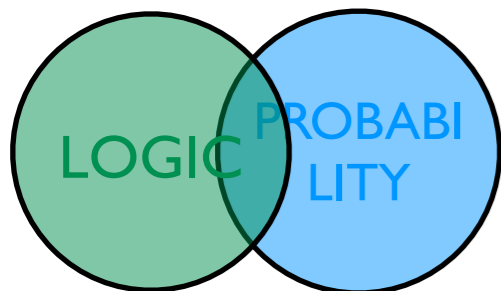
$$p(\ell_B(x_1), \dots, \ell_B(x_n)) = \frac{1}{Z} \exp\left(\sum_{\alpha} w_{\alpha} \ell_B(\alpha)\right)$$

- In PSL, we relax the **Boolean semantics** ℓ_B to a **fuzzy semantics** ℓ_F

$$p(\ell_F(x_1), \dots, \ell_F(x_n)) = \frac{1}{Z} \exp\left(\sum_{\alpha} \boxed{w_{\alpha}} \boxed{\ell_F(\alpha)}\right)$$

Weight formula

Each formula contributes with a value in $[0,1]$

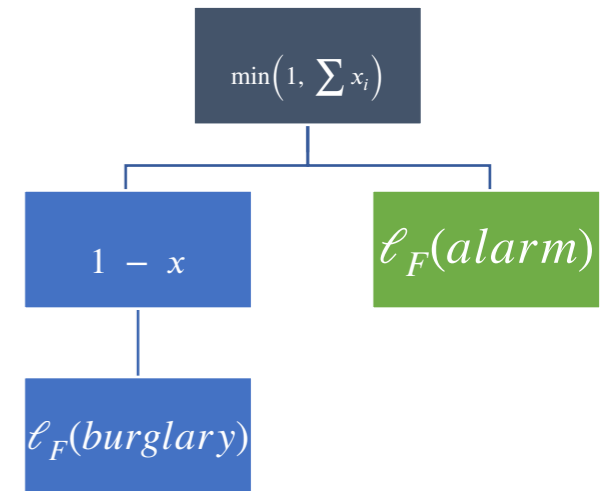


Probabilistic Soft Logic (PSL)

Using Lukasiewicz t-norm:

$\alpha : \textit{burglary} \rightarrow \textit{alarm}$

$$\ell_F(\alpha) = \min(1, 1 - \ell_F(\textit{burglary}) + \ell_F(\textit{alarm}))$$

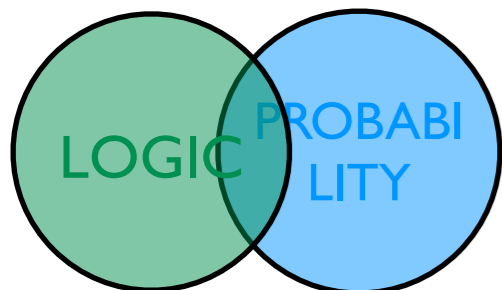


This is soft SAT
using fuzzy logic

MPE:

$$\max_{\ell_F(\textit{burglary}), \ell_F(\textit{alarm})} w_\alpha \ell_F(\alpha)$$

$$\ell_F(\textit{burglary}) = \ell_F(\textit{burglary}) + \lambda \frac{\partial w_\alpha \ell_F(\alpha)}{\partial \ell_F(\textit{burglary})}$$

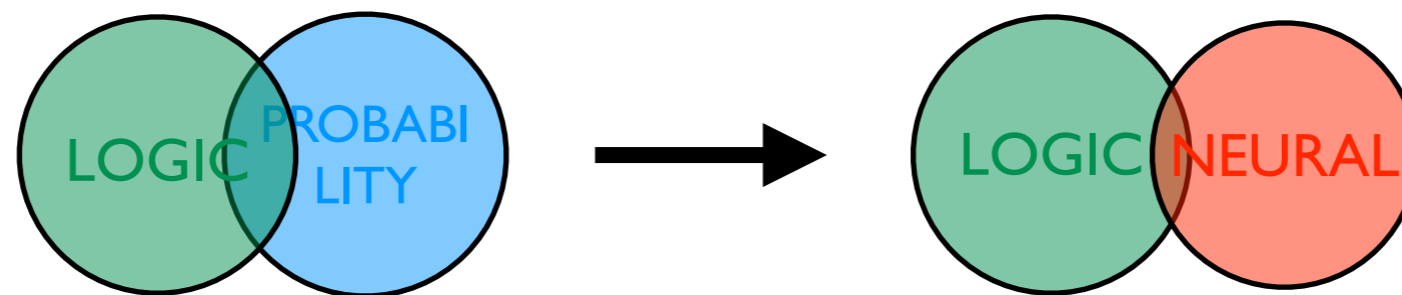


Probabilistic vs Fuzzy

- Fuzzy is an alternative logical semantics and it can still be coupled with the probabilistic ones
- Fuzzy logic is **sometimes** used as an approximation of MPE in probabilistic logic
- Fuzzy logic is **sometimes** used to solve **satisfiability** faster
 - **However**, it does not guarantee solutions coherent with the Boolean logic theory.
 - (Remember $A = B = C = D = E = 0.2$)

6. Semantics

Neural Symbolic



Neural Symbolic

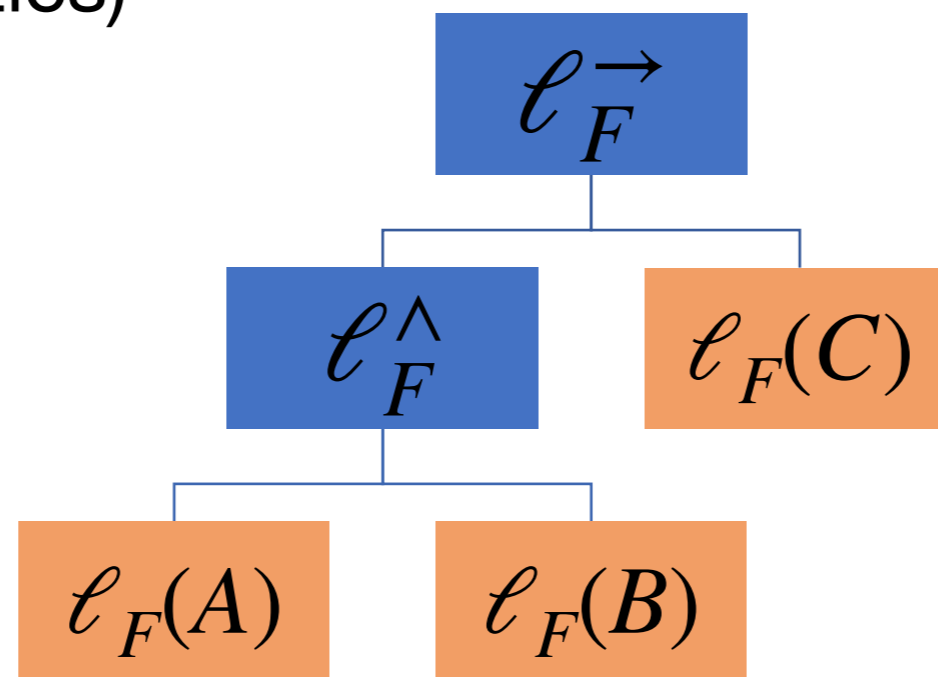
How to carry over concepts from the semantics of StarAI to neural symbolic?

$\ell(Q)$

Labelling functions
(semantics)

= Parametric circuit

$\ell_F((A \wedge B) \rightarrow C)$



The query Q determine the **structure** (potentially after knowledge compilation)

Neural Symbolic

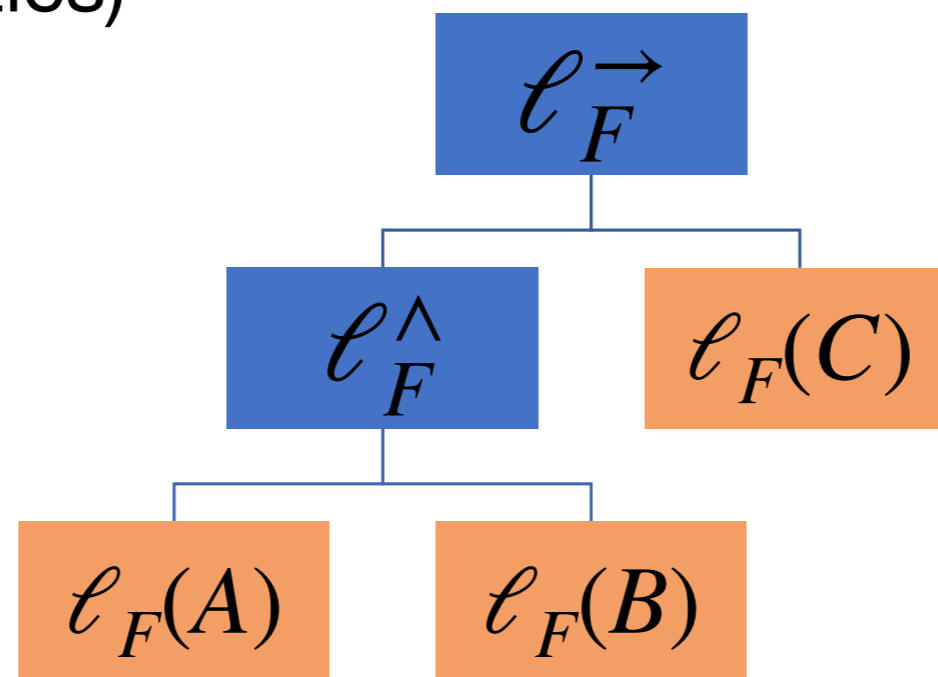
How to carry over concepts from the semantics of StarAI to neural symbolic?

$\ell(Q)$

Labelling functions
(semantics)

= Parametric circuit

$\ell_F((A \wedge B) \rightarrow C)$



The leaves
represent the
scalar parameters

Neural Symbolic

How to carry over concepts from the semantics of StarAI to neural symbolic?

- Atomic labels are just **scalar tables of parameters**



0.1 :: burglary. (B)
0.05 :: earthquake. (E)
0.6 :: hears_alarm(john). (H)
alarm :- earthquake.
alarm :- burglary.

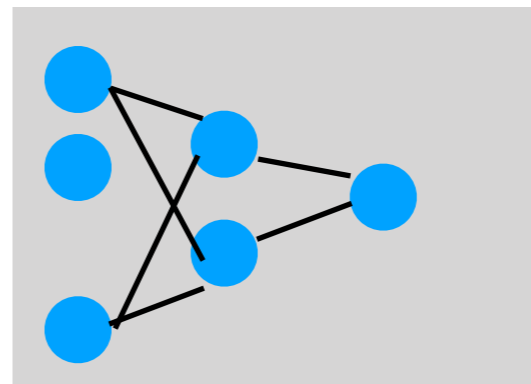
L	p
Burglary	0,1
Earthquake	0,05
...	

Neural Symbolic

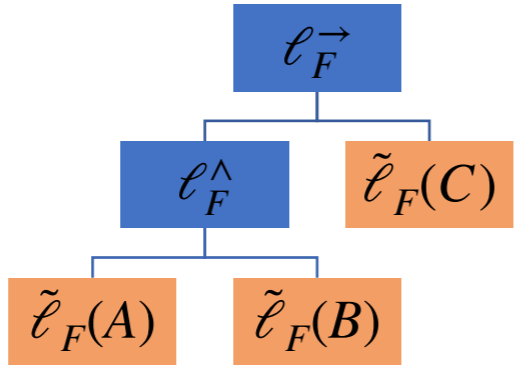
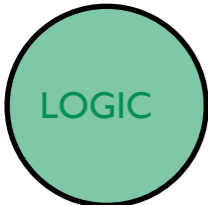
How to carry over concepts from the semantics of StarAI to neural symbolic?

- What if atomic labels are just **neural networks**?

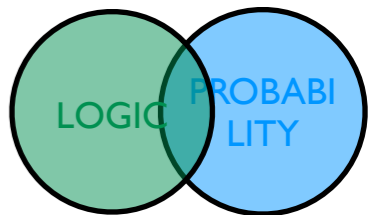
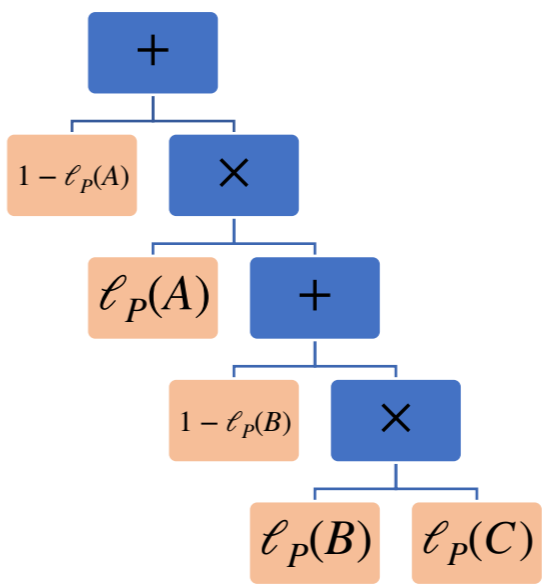
? :: burglary()
? :: earthquake. ()
? :: hears_alarm(john).
alarm :- earthquake.
alarm :- burglary.



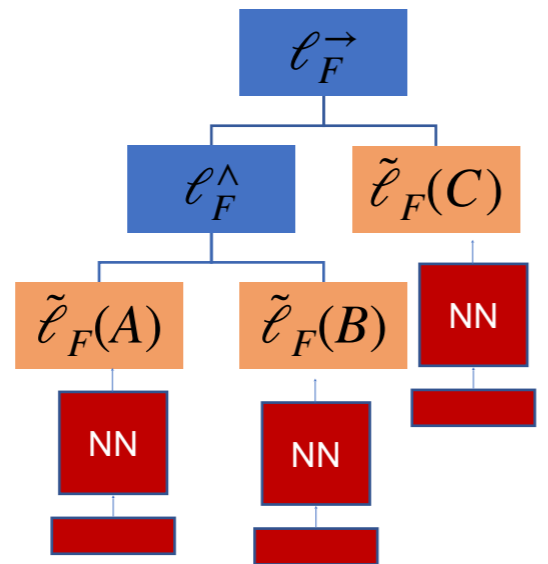
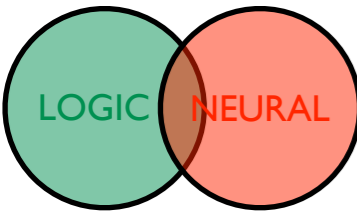
StarAI to Neural Symbolic



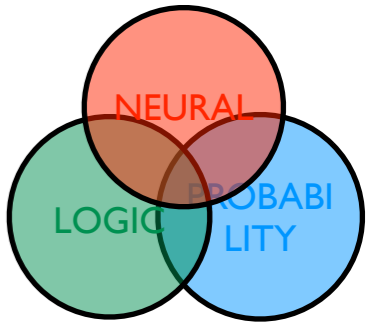
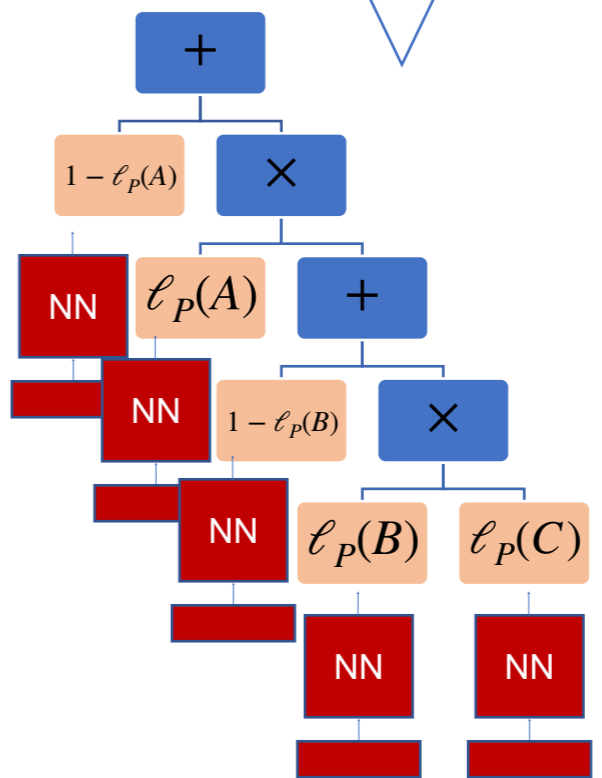
StarAI



REPARAMETERIZATION

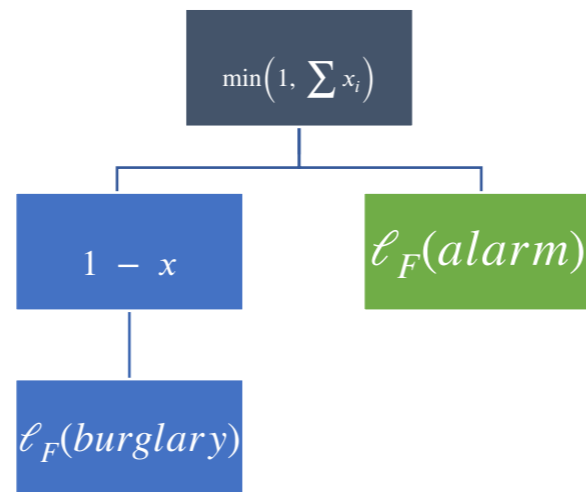


NeSy



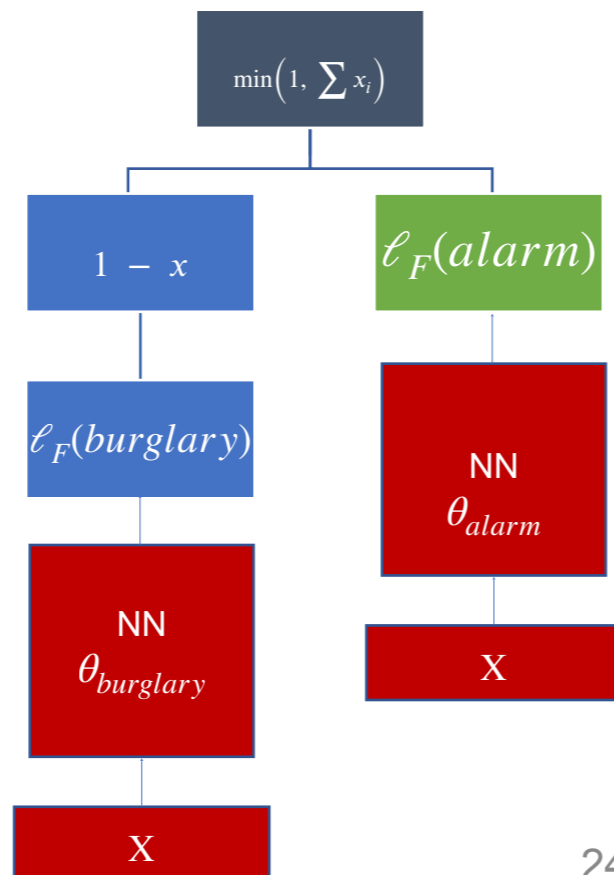
Fuzzy Reparameterization

$\alpha : burglary \rightarrow alarm$



Semantic Based
Regularization (Diligenti
et al, AI 2017)

Logic Tensor Network
(Donadello et al, IJCAI
2017)



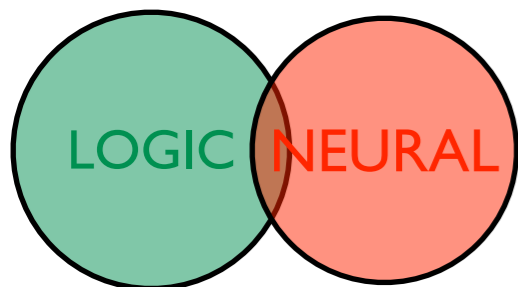
StarAI (PSL)

$$\max_{\ell_F(stress(X)), \ell_F(smokes(X))} w_\alpha \ell_F(\alpha)$$

NeSy (SBR, LTN)

$$\max_{\theta_{burglary}, \theta_{alarm}} w_\alpha \ell_F(\alpha)$$

Parameters of
the neural nets



Probabilistic Reparameterization

■ Probabilistic parameters

- ProbLog:

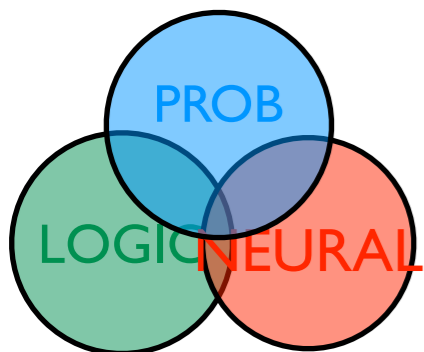
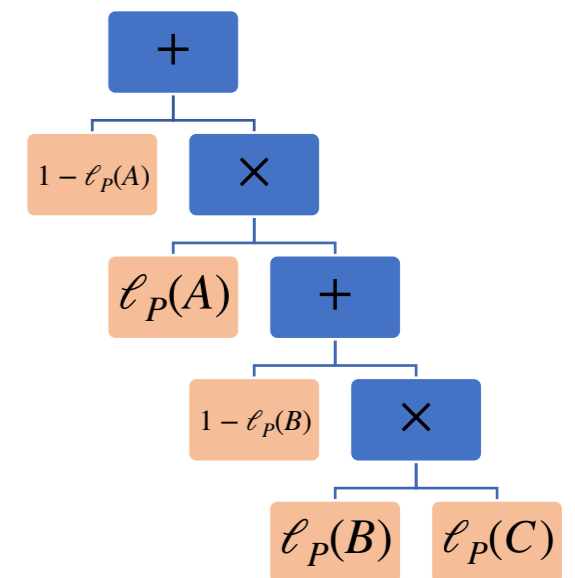
$$p(\ell_B(x_1), \dots, \ell_B(x_n)) = \prod_{i:\ell_B(x_i)=True} p(x_i) \prod_{i:\ell_B(x_i)=False} (1-p(x_i))$$

- Markov Logic:

$$p(\ell_B(x_1), \dots, \ell_B(x_n)) = \frac{1}{Z} \exp\left(\sum_{\alpha} w_{\alpha} \ell_B(\alpha)\right)$$

WMC

$$p(Q) = \sum_{\ell_B(x_1), \dots, \ell_B(x_n) \models Q} p(\ell_B(x_1), \dots, \ell_B(x_n))$$



Probabilistic Reparameterization

■ Neural parameters

- **DeepProbLog** (Manhaeve et al, NeurIPS (2018))

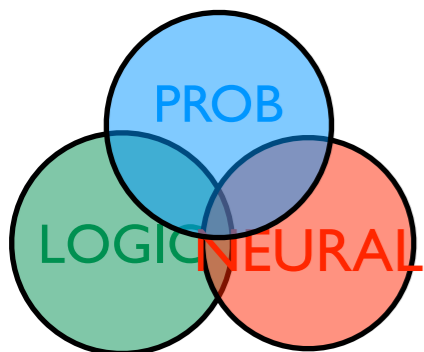
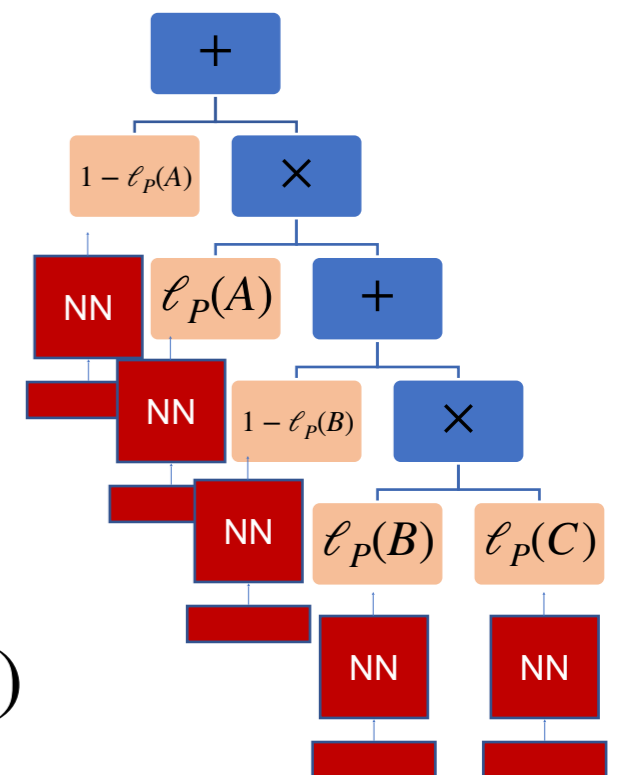
$$p(\ell_B(x_1), \dots, \ell_B(x_n)) = \prod_{i:\ell_B(x_i)=True} p(x_i) \prod_{i:\ell_B(x_i)=False} (1-p(x_i))$$

- **Relational Neural Machines** (Marra et al, ECAI 2020)

$$p(\ell_B(x_1), \dots, \ell_B(x_n)) = \frac{1}{Z} \exp\left(\sum_{\alpha} w_{\alpha} \ell_B(\alpha)\right)$$

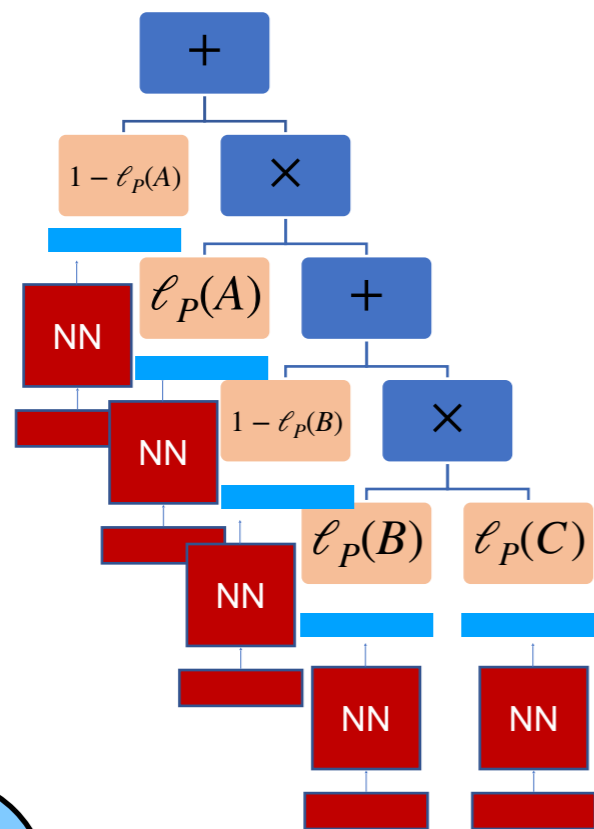
WMC

$$p(Q) = \sum_{\ell_B(x_1), \dots, \ell_B(x_n) \models Q} p(\ell_B(x_1), \dots, \ell_B(x_n))$$



Probabilistic Reparameterization

- **DeepProbLog** (Manhaeve et al, NeurIPS (2018))



Interface

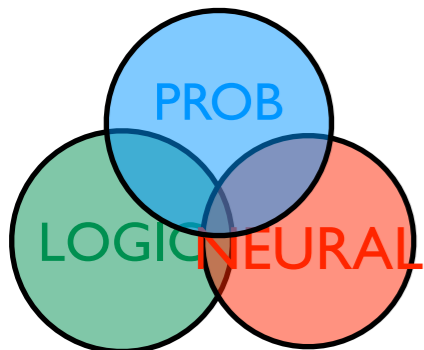
Probabilistic fact

0.01 :: burglary.



Neural Predicate

nn(mnist_net, [X], Y, [0 ... 9]) :: digit(X,Y).



6. Semantics

Key Messages

- StarAI and NeSy share the same underlying semantics
- Semantics can be described in terms of parametric circuits
- Differentiable semantics/circuits allow an easy integration
- NeSy models can be seen as neural reparameterization of StarAI models

A Recipe for NeSy

A recipe for NeSy

STEP 1

1. Take your favorite symbolic (logic / rule based) representation



(applied on DeepProbLog)

layout Pieter Robberechts

© Luc De Raedt

NeSy Data Point

`calls(image32, signal42, mary)`

`image32 =`  `signal42 =` 

NeSy Model

Logic Rules

```
alarm(B,E) IF burglary(B) OR earthquake(E).
calls(B,E,X) IF alarm(B,E) AND hears_alarm(X).
```

Logic Facts

```
hears_alarm(mary).
hears_alarm(john).
```

```
neural-net(image_perception(B))
neural-net(signal_analysis(E))
```

Neural Net Modules

`image_perception =`  `signal_analysis =` 

A recipe for NeSy



STEP 1

1. Take your favorite symbolic (logic / rule based) representation
2. Interpret neural networks as neural predicates

(applied on DeepProbLog)

NeSy Data Point

`calls(image32, signal42, mary)`

`image32 =`  `signal42 =` 

NeSy Model

Logic Rules

`alarm(B,E) IF burglary(B) OR earthquake(E).`
`calls(B,E,X) IF alarm(B,E) AND hears_alarm(X).`

Logic Facts

`hears_alarm(mary).`
`hears_alarm(john).`

Neural Predicates

`burglary(B) IF neural-net(image_perception(B))`
`earthquake(E) IF neural-net(signal_analysis(E))`

Neural Net Modules

`image_perception =`  `signal_analysis =` 

A recipe for NeSy

STEP 1

1. Take your favorite symbolic (logic / rule based) representation
2. Interpret neural networks as neural predicates
3. Turn the 0/1 or True/False into Probabilistic or Fuzzy Interpretation



(applied on DeepProbLog)

layout Pieter Robberechts

© Luc De Raedt

NeSy Data Point

`calls(image32, signal42, mary)`

`image32` =  `signal42` = 

NeSy Model

Logic Rules

`alarm(B,E) IF burglary(B) OR earthquake(E).`
`calls(B,E,X) IF alarm(B,E) AND hears_alarm(X).`

Logic Facts

`hears_alarm(mary).`
`hears_alarm(john).`

Probability

0.3
0.6

Neural Predicates

`burglary(B) IF neural-net(image_perception(B))`
`earthquake(E) IF neural-net(signal_analysis(E))`

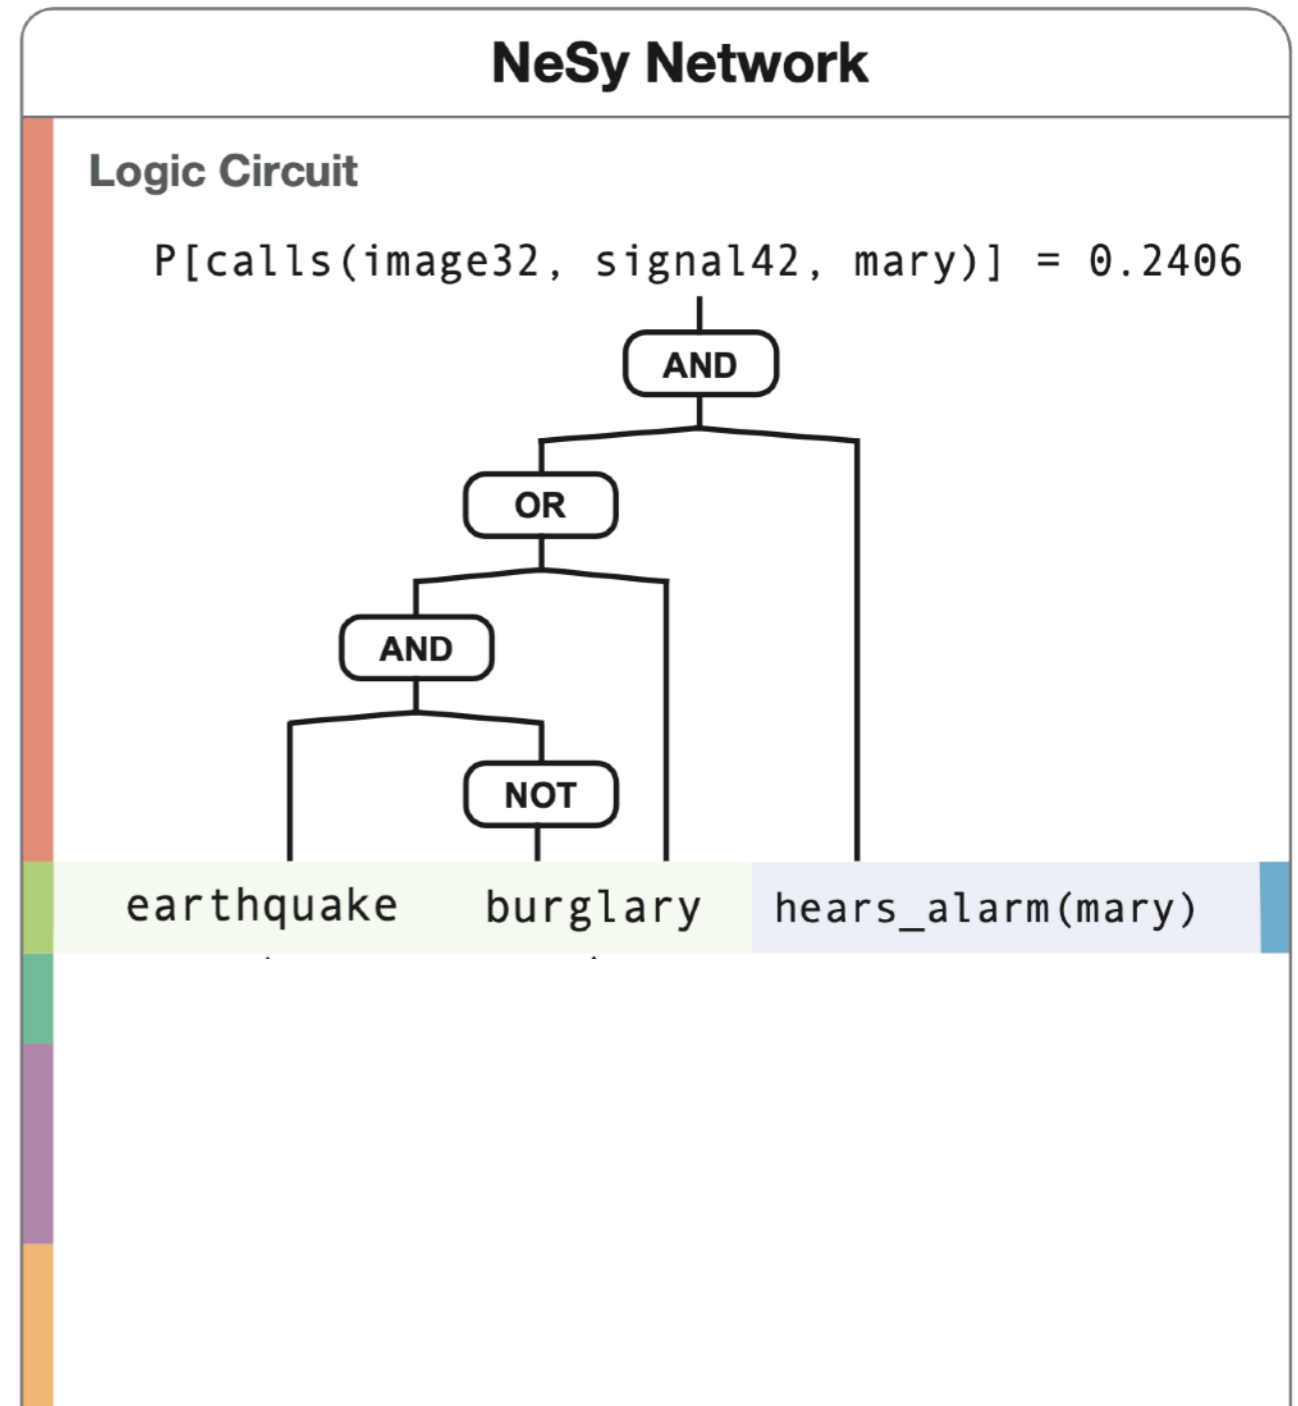
Neural Net Modules

`image_perception` =  `signal_analysis` = 

A recipe for NeSy

STEP 2

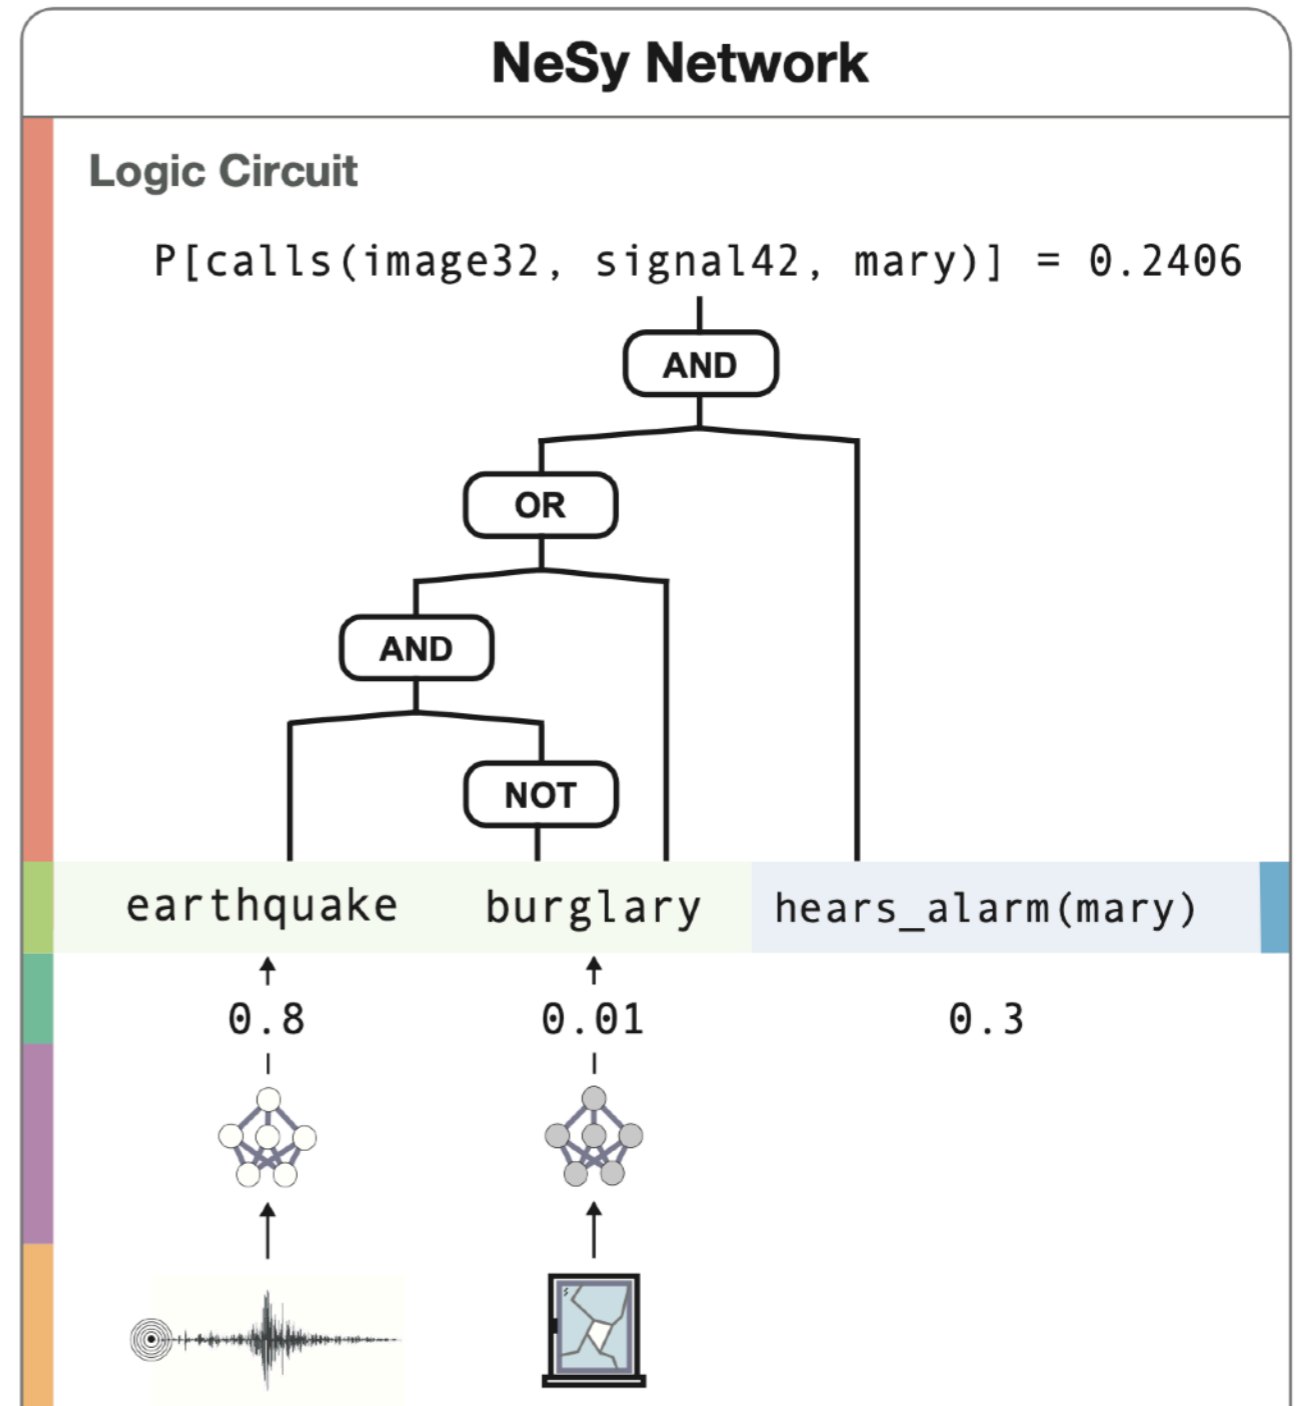
4. Construct logical proof / explanation for example



A recipe for NeSy

STEP 2

4. Construct logical proof / explanation for example
5. Add the neural networks to the corresponding predicates (*reparametrise*)



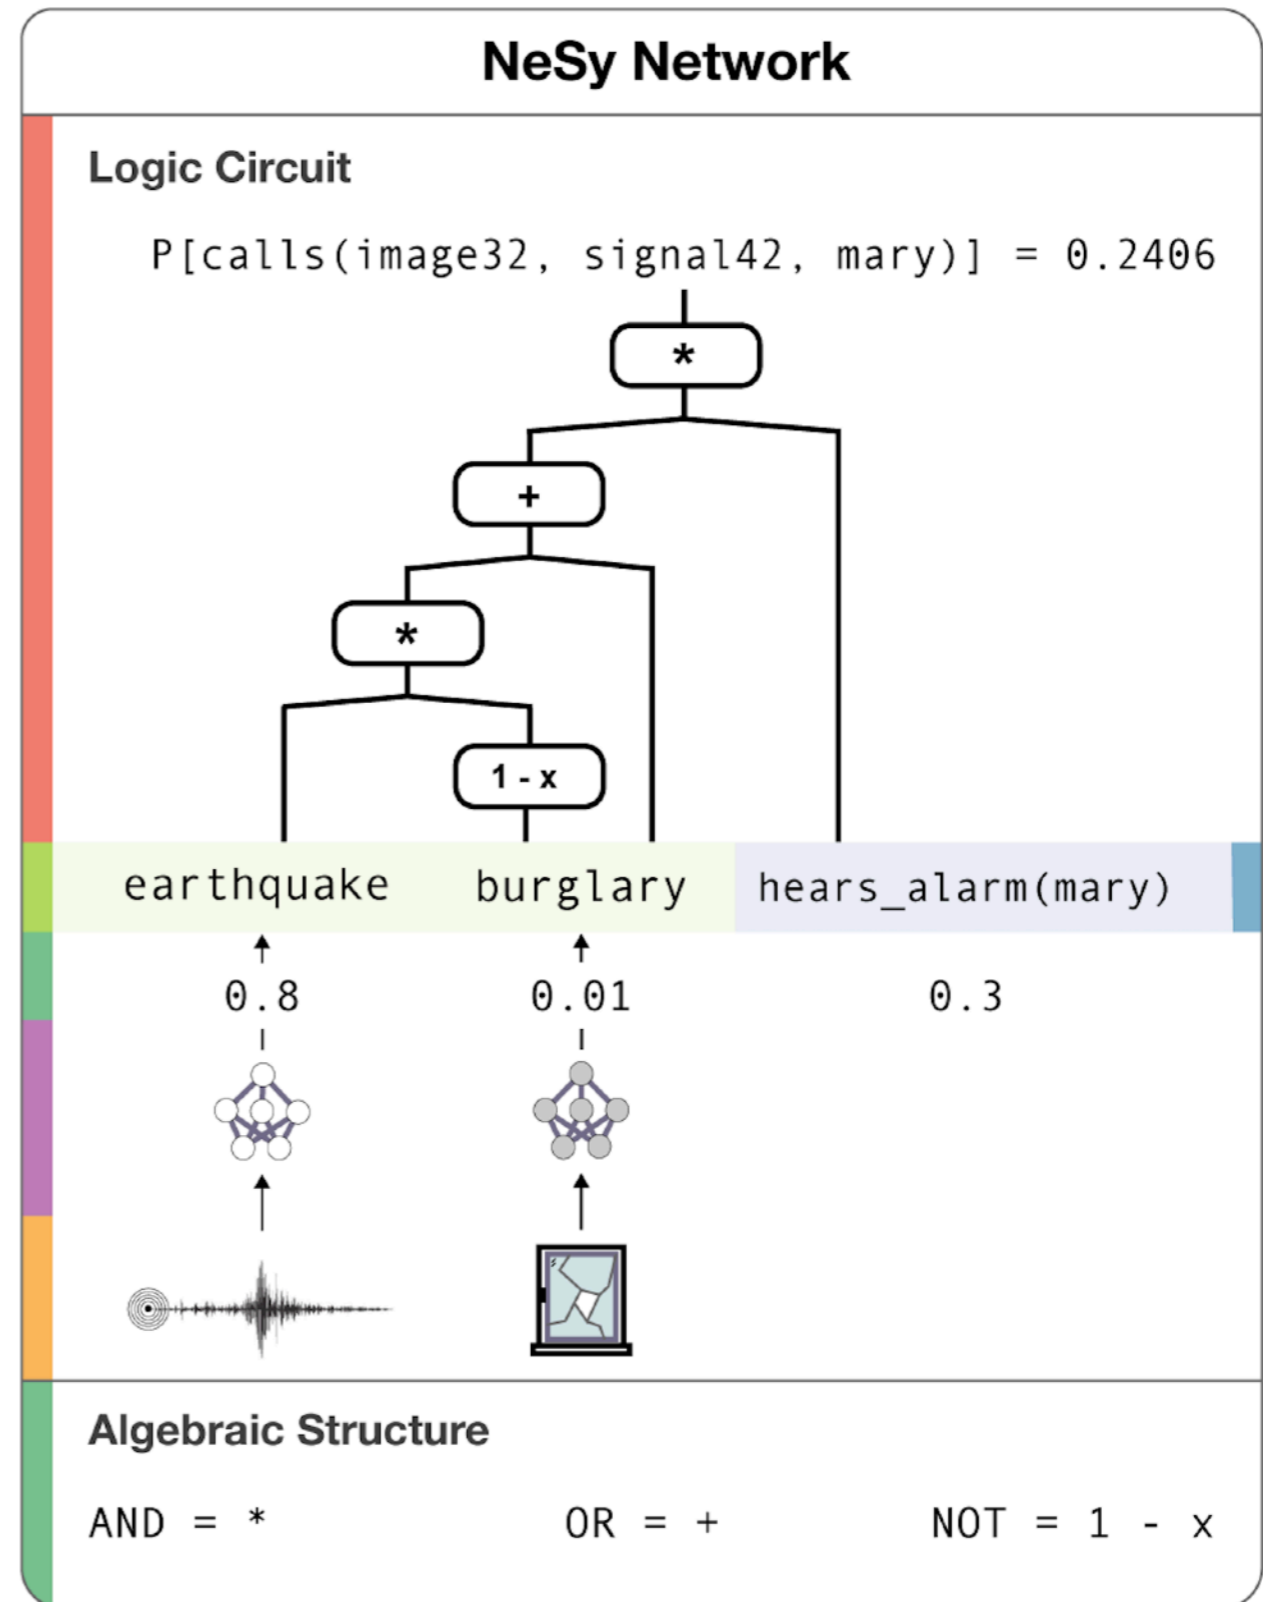
A recipe for NeSy

STEP 3

4. Construct logical proof / explanation for example
5. Add the neural networks to the corresponding predicates (*reparametrise*)
6. Replace OR and AND by \oplus and \otimes
7. Differentiate

layout Pieter Robberechts

© Luc De Raedt

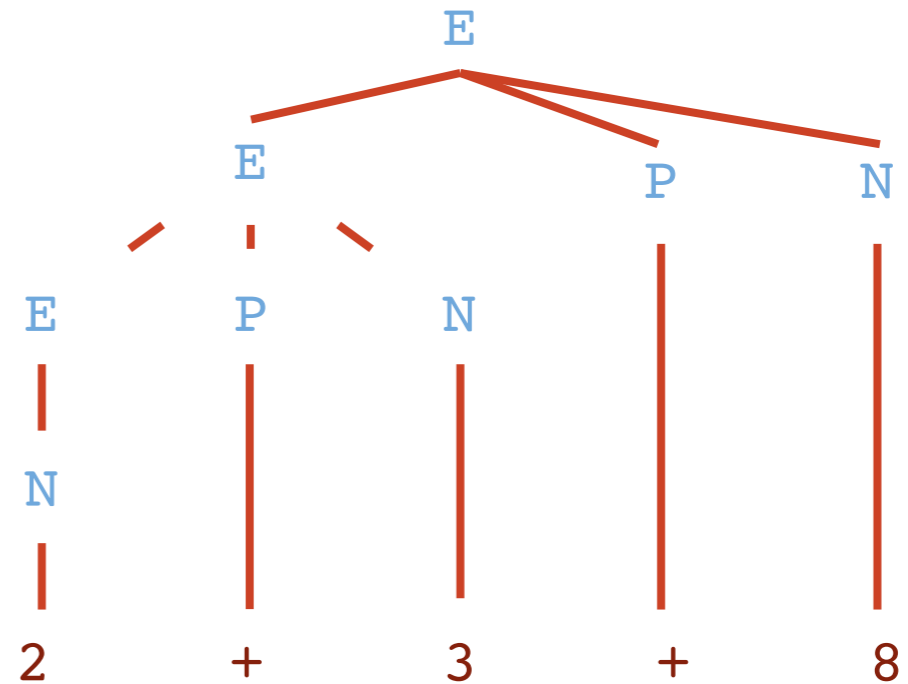


DeepStochLog

- Little sibling of DeepProbLog [Winters, Marra, et al AAAI 22]
- Based on a different semantics
 - **probabilistic graphical models vs grammars**
 - **random graphs vs random walks**
- Underlying StarAI representation is **Stochastic Logic Programs** (Muggleton, Cussens)
 - close to Probabilistic Definite Clause Grammars, aka probabilistic unification based grammar formalism
 - again the idea of neural predicates
- Scales better, is faster than DeepProbLog

CFG: Context-Free Grammar

$E \rightarrow N$
 $E \rightarrow E, P, N$
 $P \rightarrow ["+"]$
 $N \rightarrow ["0"]$
 $N \rightarrow ["1"]$
...
 $N \rightarrow ["9"]$



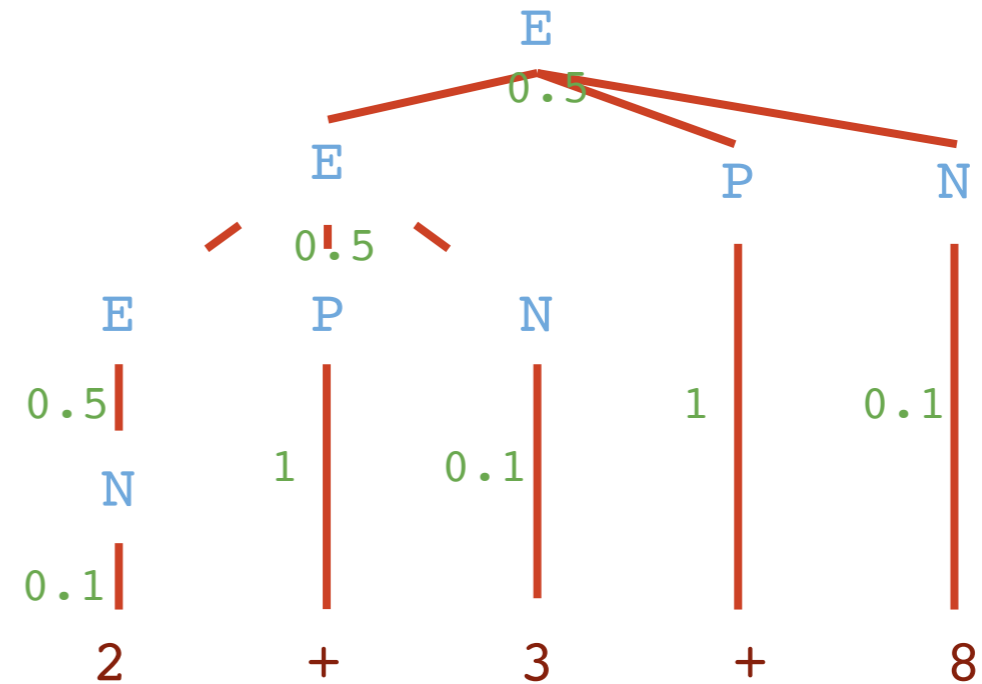
Useful for:

- Is sequence an **element of** the specified language?
- What is the *"part of speech"*-**tag** of a terminal
- **Generate** all elements of language

PCFG: Probabilistic Context-Free Grammar

0.5	:::	E	-->	N
0.5	:::	E	-->	E, P, N
1.0	:::	P	-->	["+"]
0.1	:::	N	-->	["0"]
0.1	:::	N	-->	["1"]
		...		
0.1	:::	N	-->	["9"]

Always sums to 1 per non-terminal



Probability of this parse = $0.5 * 0.5 * 0.5 * 0.1 * 1 * 0.1 * 1 * 0.1$
 = 0.000125

Useful for:

- What is the **most likely parse** for this sequence of terminals? *(useful for ambiguous grammars)*
- What is the **probability of generating** this string?

DCG: Definite Clause Grammar

$e(N) \text{ --> } n(N) \text{ .}$

$e(N) \text{ --> } e(N1), p, n(N2),$
 $\{N \text{ is } N1 + N2\} \text{ .}$

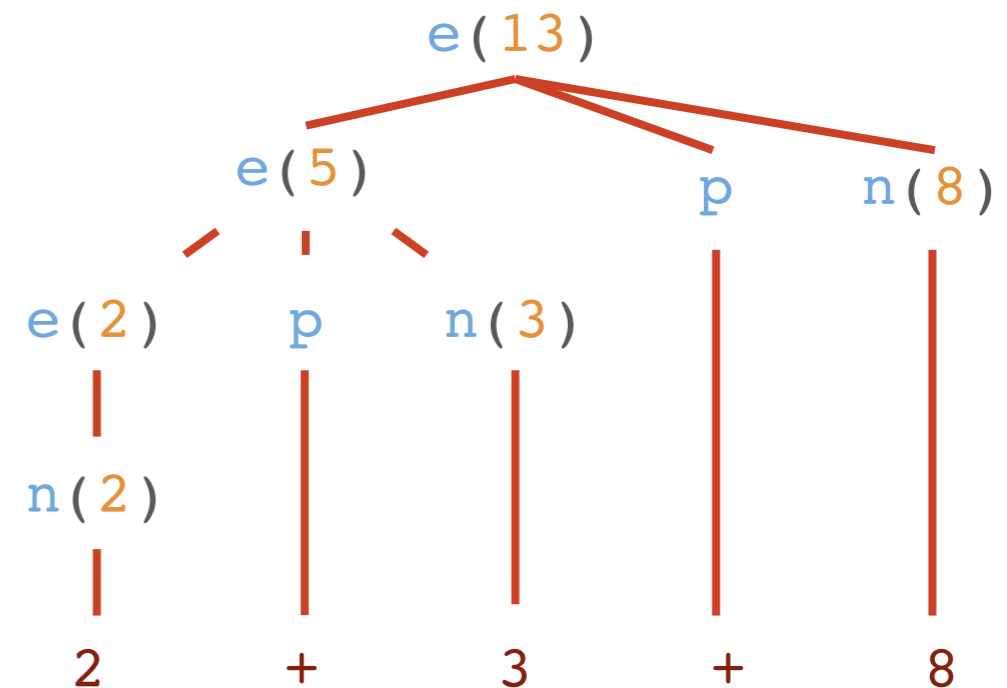
$p \text{ --> } ["+"] \text{ .}$

$n(0) \text{ --> } ["0"] \text{ .}$

$n(1) \text{ --> } ["1"] \text{ .}$

...

$n(9) \text{ --> } ["9"] \text{ .}$

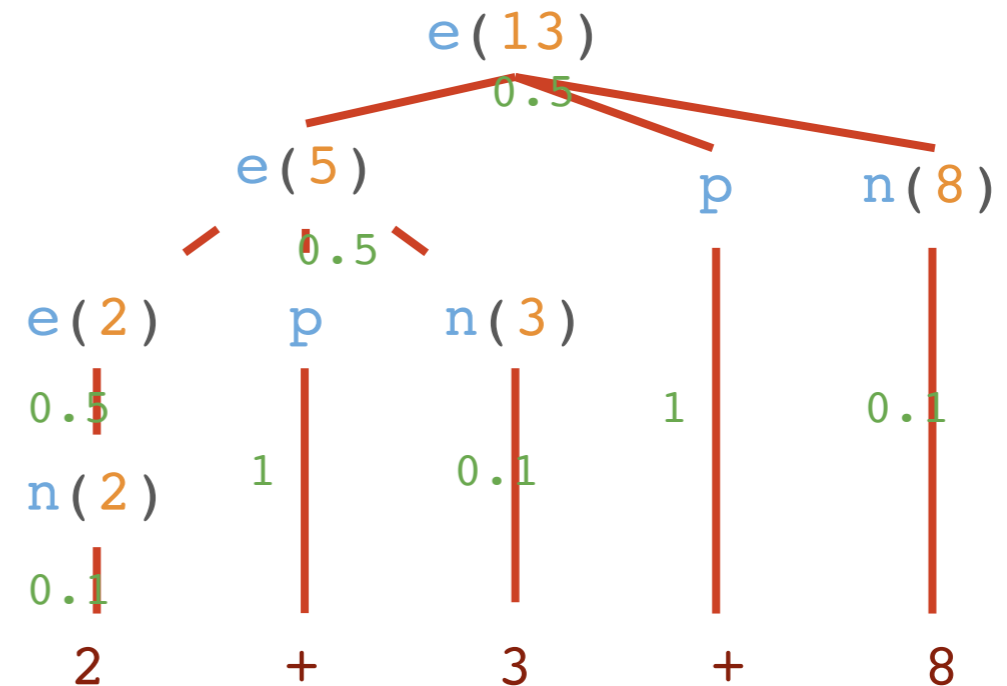


Useful for:

- Modelling **more complex** languages (e.g. context-sensitive)
- Adding constraints between non-terminals thanks to **Prolog** power (e.g. through unification)
- **Extra inputs & outputs** aside from terminal sequence (through unification of input variables)

SDCG: Stochastic Definite Clause Grammar

0.5 :: e(N) --> n(N) .
0.5 :: e(N) --> e(N1), p, n(N2),
 {N is N1 + N2} .
1.0 :: p --> ["+"] .
0.1 :: n(0) --> ["0"] .
0.1 :: n(1) --> ["1"] .
 ...
0.1 :: n(9) --> ["9"] .



Probability of this parse = $0.5 * 0.5 * 0.5 * 0.1 * 1 * 0.1 * 1 * 0.1$
= 0.000125

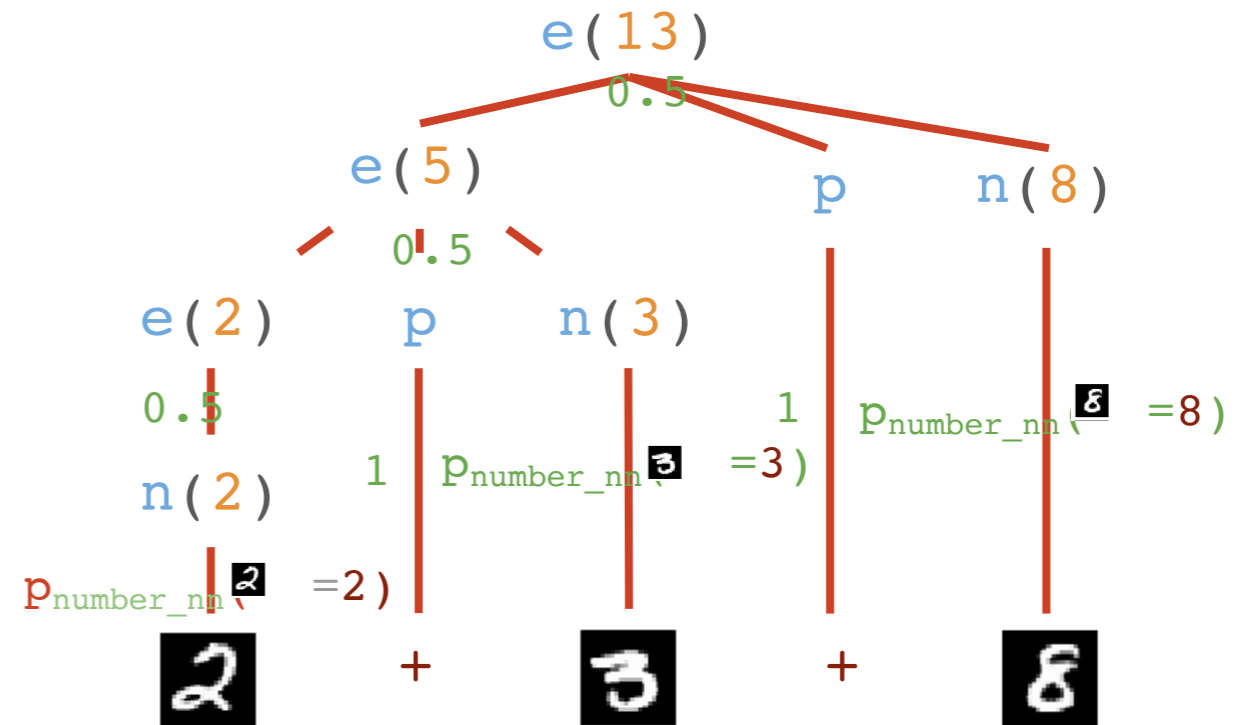
Useful for:

- Same benefits as PCFGs give to CFG (e.g. most likely parse)
- But: loss of probability mass possible due to failing derivations

NDCG: Neural Definite Clause Grammar (= DeepStochLog)

```

0.5 :: e(N) --> n(N).
0.5 :: e(N) --> e(N1), p, n(N2),
                {N is N1 + N2}.
1.0 :: p      --> ["+"].
nn(number_nn,[X],[Y],[digit])::
    n(Y) -> [X].
digit(Y) :-
    member(Y,[0,1,2,3,4,5,6,7,8,9]).
    
```



Probability of this parse =

$$0.5 * 0.5 * 0.5 * p_{\text{number_nn}}(2=2) * 1 * p_{\text{number_nn}}(3=3) * 1 * p_{\text{number_nn}}(8=8)$$

Useful for:

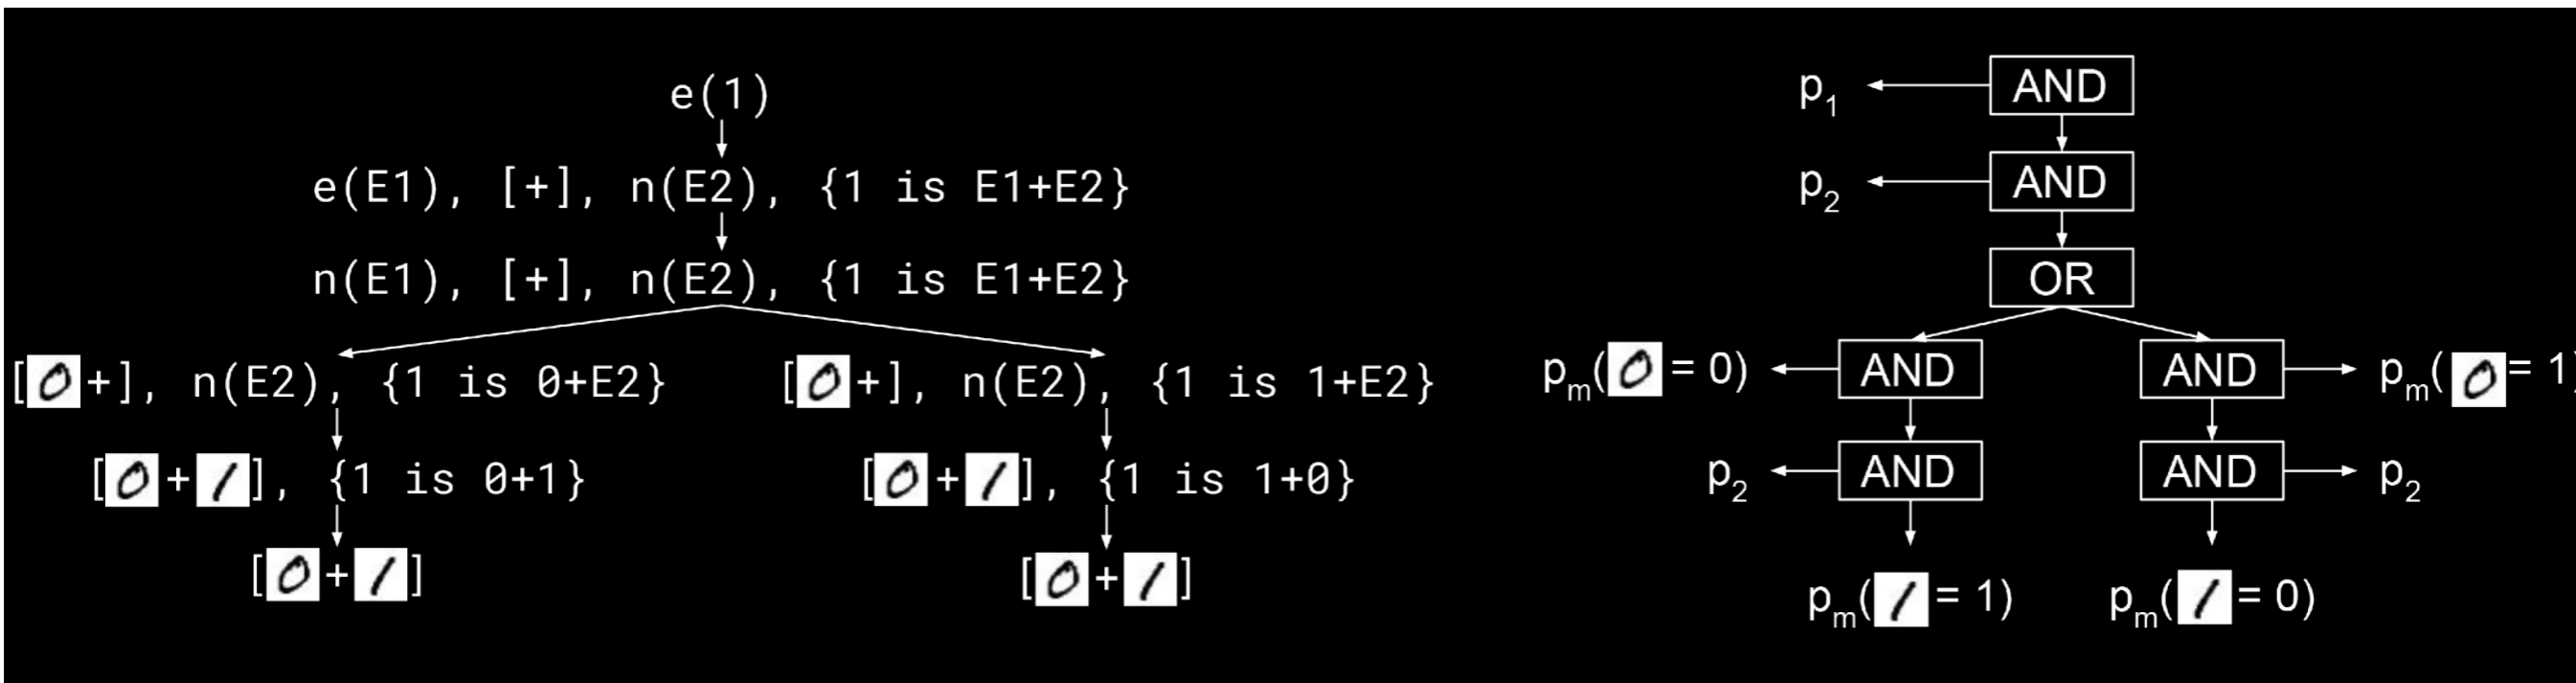
- **Subsymbolic** processing: e.g. tensors as terminals
- Learning rule probabilities using **neural networks**

DeepStochLog Inference

Deriving probability of goal for given terminals in NDCG

Proof derivations $d(e(1), [0+1])$

then turn it into and/or tree

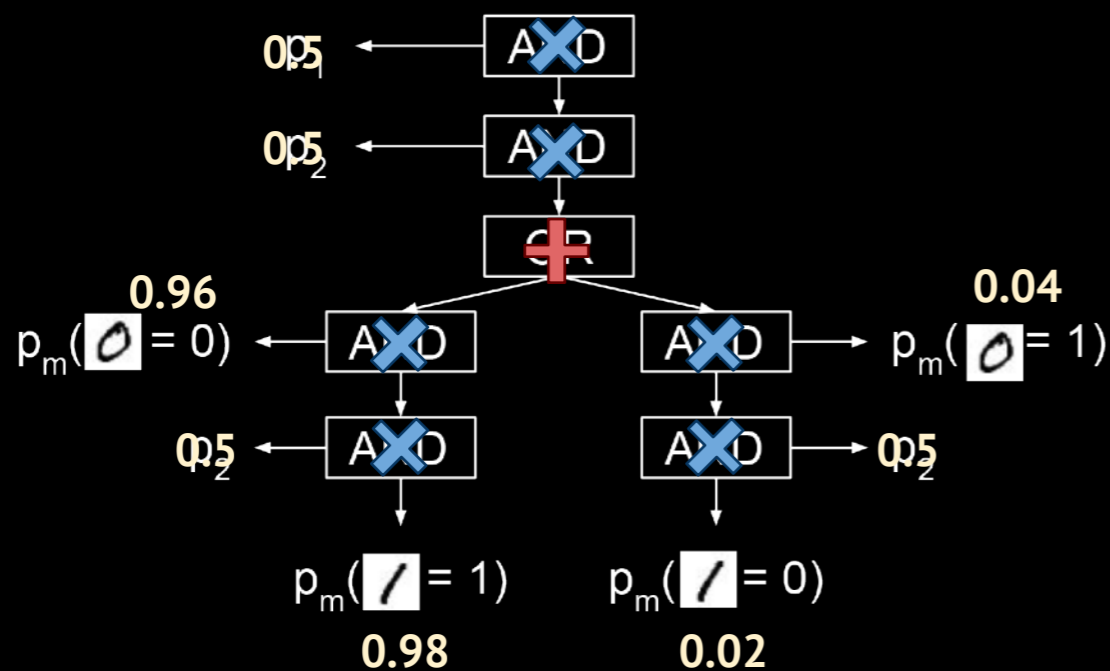


And/Or tree + semiring for different inference types

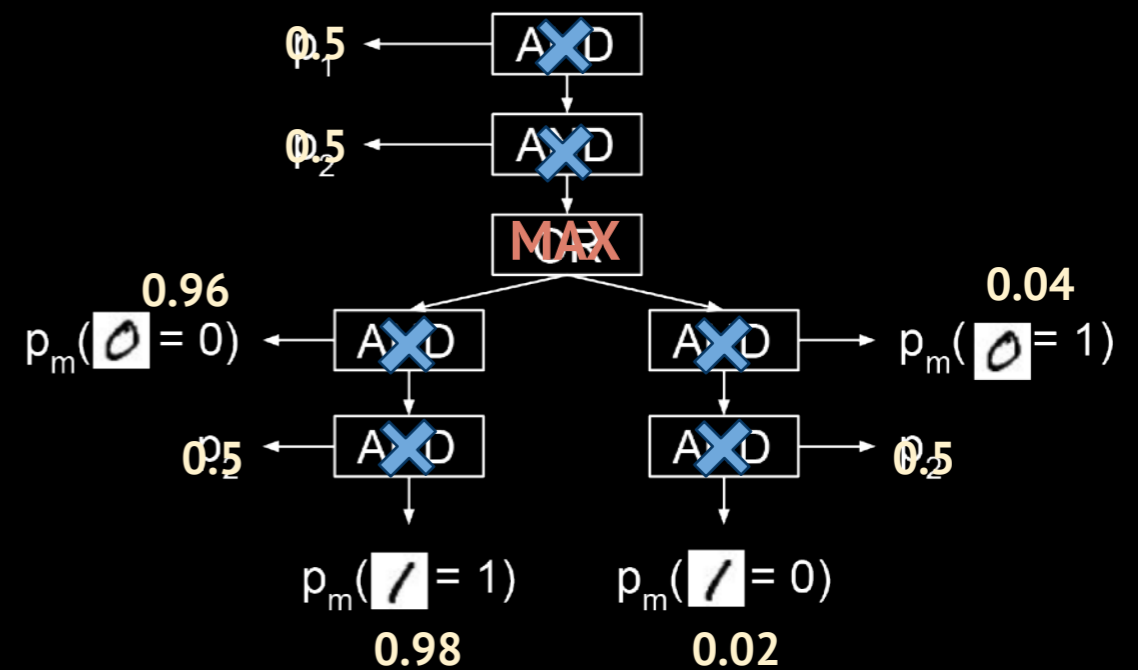
Probability of goal

Most likely derivation

$$P_G(\text{derives}(e(1), [0, +, 1]) = 0.1141$$



$$d_{\max}(e(1), [0, +, 1]) = \operatorname{argmax}_{d(e(t))=[0, +, 1]} P_G(d(e(1))) = [0, +, 1]$$



Inference optimisation

- Inference is **optimized** using
 - **SLG resolution**: Prolog tables the returned proof tree(s), and thus creates forest
 - Allows for reusing probability calculation results from intermediate nodes

Table 6: Q4 Parsing time in seconds (T2). Comparison of the DeepStochLog with and without tabling (SLD vs SLG resolution).

Lengths	# Answers	No Tabling	Tabling
1	10	0.067	0.060
3	95	0.081	0.096
5	1066	3.78	0.95
7	10386	30.42	10.95
9	68298	1494.23	132.26
11	416517	timeout	1996.09

- **Batched network calls**: Evaluate all the required neural network queries first
 - Very natural for neural networks to evaluate multiple instances at once using batching
 - & less overhead in logic & neural network communication

Mathematical expression outcome

T1: Summing MNIST numbers with pre-specified # digits


$$\boxed{5} \boxed{3} + \boxed{8} \boxed{4} = 137$$

T2: Expressions with images representing operator or single digit number.


$$7 + 4 \times 3 = 19$$

Table 1: The test accuracy (%) on the MNIST addition (T1).

Methods	Number of digits per number (N)			
	1	2	3	4
NeurASP	97.3 ± 0.3	93.9 ± 0.7	timeout	timeout
DeepProbLog	97.2 ± 0.5	95.2 ± 1.7	timeout	timeout
DeepStochLog	97.9 ± 0.1	96.4 ± 0.1	94.5 ± 1.1	92.7 ± 0.6

Table 2: The accuracy (%) on the HWF dataset (T2).

Method	Expression length			
	1	3	5	7
NGS	90.2 ± 1.6	85.7 ± 1.0	91.7 ± 1.3	20.4 ± 37.2
DeepProbLog	90.8 ± 1.3	85.6 ± 1.1	timeout	timeout
DeepStochLog	90.8 ± 1.0	86.3 ± 1.9	92.1 ± 1.4	94.8 ± 0.9

Classic grammars, but with MNIST images as terminals

T3: Well-formed brackets as input (without parse). Task: predict parse.



→ parse = () (() ())

T4: inputs are strings $a^k b^l c^m$ (or permutations of [a,b,c], and $(k+l+m) \bmod 3=0$). Predict 1 if $k=l=m$, otherwise 0.



Table 3: The parse accuracy (%) on the well-formed parentheses dataset (T3).

Method	Maximum expression length		
	10	14	18
DeepProbLog	100.0 ± 0.0	99.4 ± 0.5	99.2 ± 0.8
DeepStochLog	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0

Table 4: The accuracy (%) on the $a^n b^n c^n$ dataset (T4).

Method	Expression length		
	3-12	3-15	3-18
DeepProbLog	99.8 ± 0.3	timeout	timeout
DeepStochLog	99.4 ± 0.5	99.2 ± 0.4	98.8 ± 0.2

Citation networks

T5: Given scientific paper set with only few labels & citation network, find all labels

Table 5: **Q3** Accuracy (%) of the classification on the test nodes on task **T5**

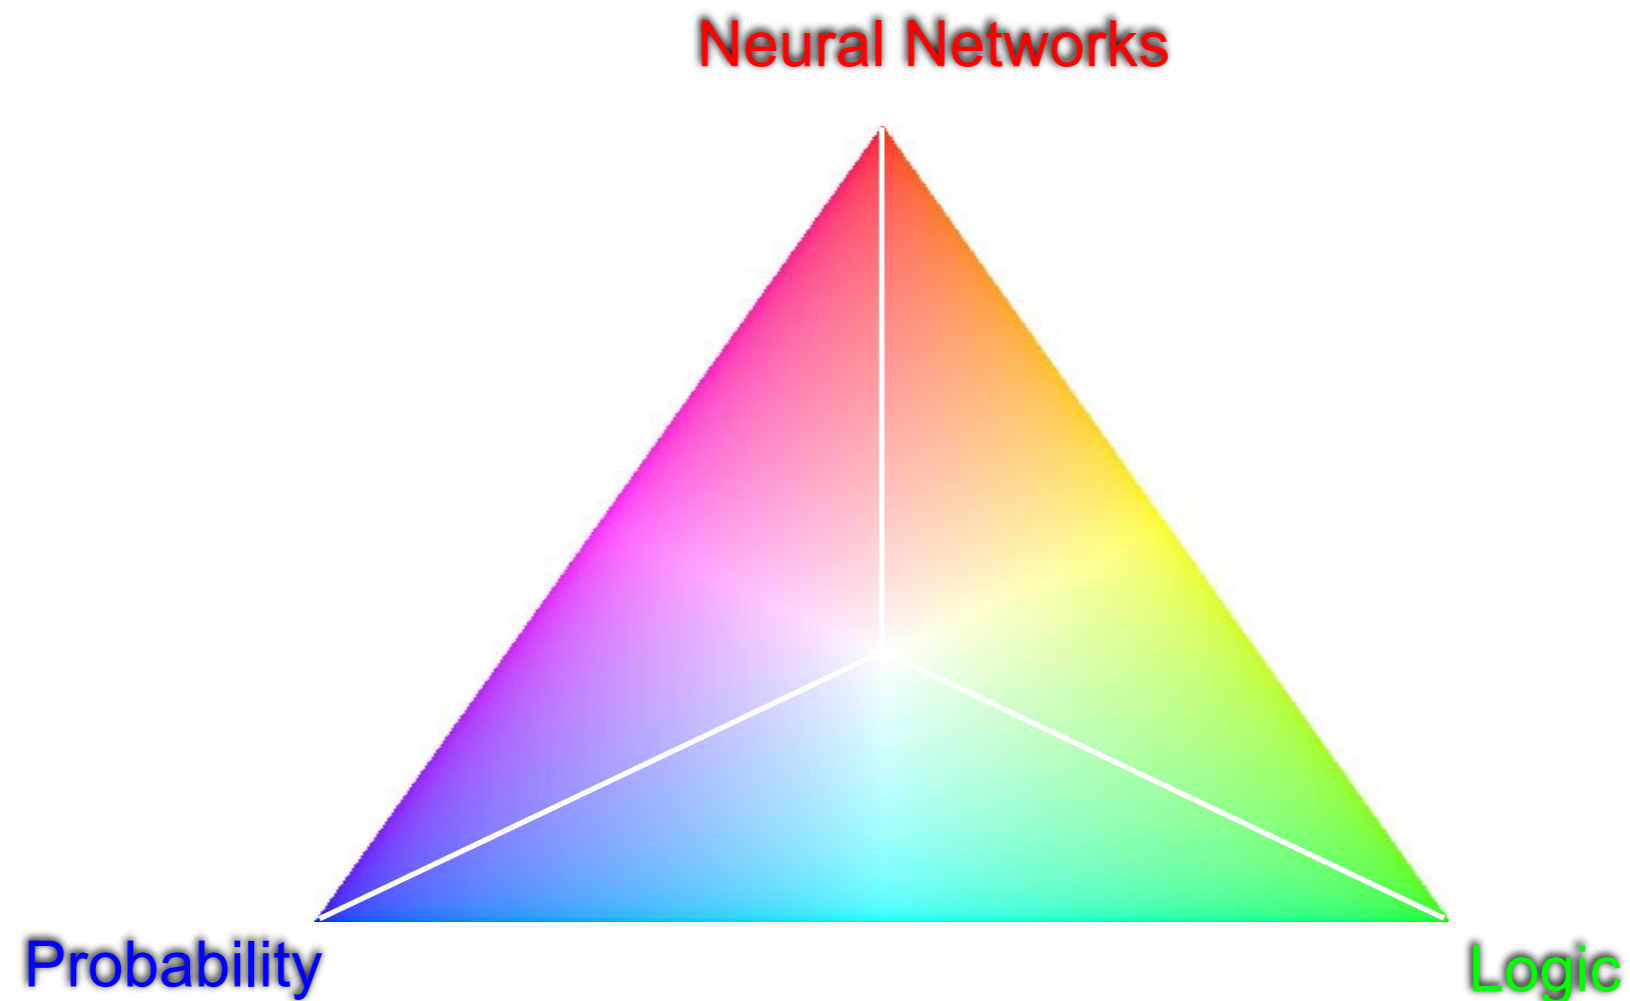
Method	Citeseer	Cora
ManiReg	60.1	59.5
SemiEmb	59.6	59.0
LP	45.3	68.0
DeepWalk	43.2	67.2
ICA	69.1	75.1
GCN	70.3	81.5
DeepProbLog	timeout	timeout
DeepStochLog	65.0	69.4

7. Logic vs Probability vs Neural

7. Logic vs Probability vs Neural Key Messages

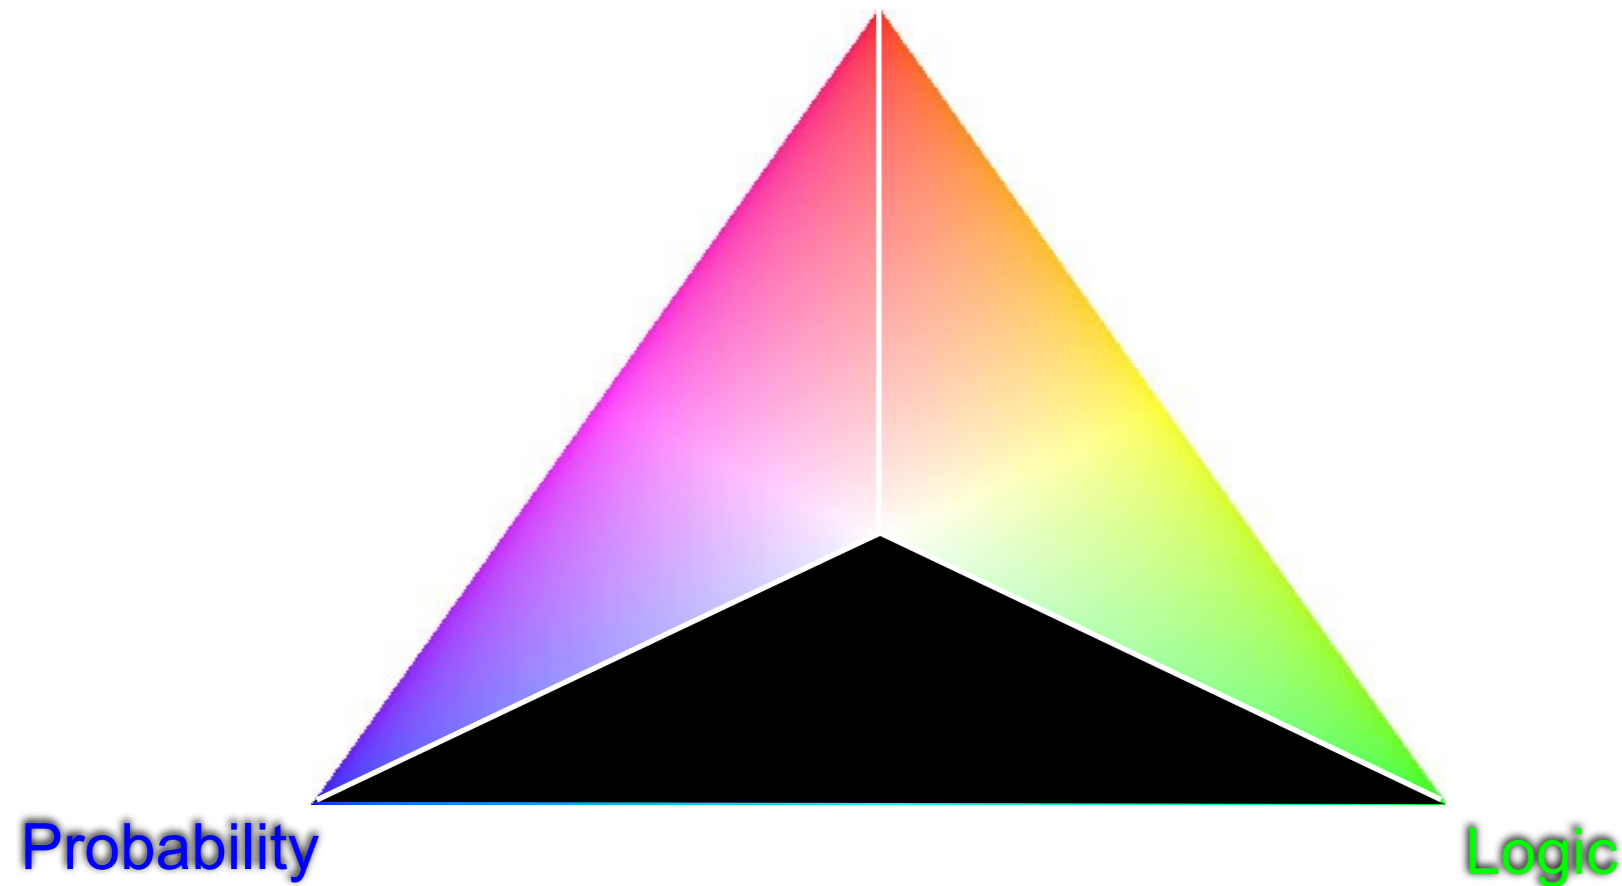
- We have three paradigms in the NeSy spectrum: **Logic**, **Probability** and **Neural Networks**
- An **integration** of the three should have the original paradigms as **special cases**
 - Computationally complex
- The integration is usually achieved by sacrificing the base paradigms
 - More scalable

About integration in neural symbolic



Statistical Relational AI

Neural Networks



They perfectly integrate probability theory (Probabilistic Graphical Models) and Logic.

Knowledge Graph Embeddings

Neural Networks

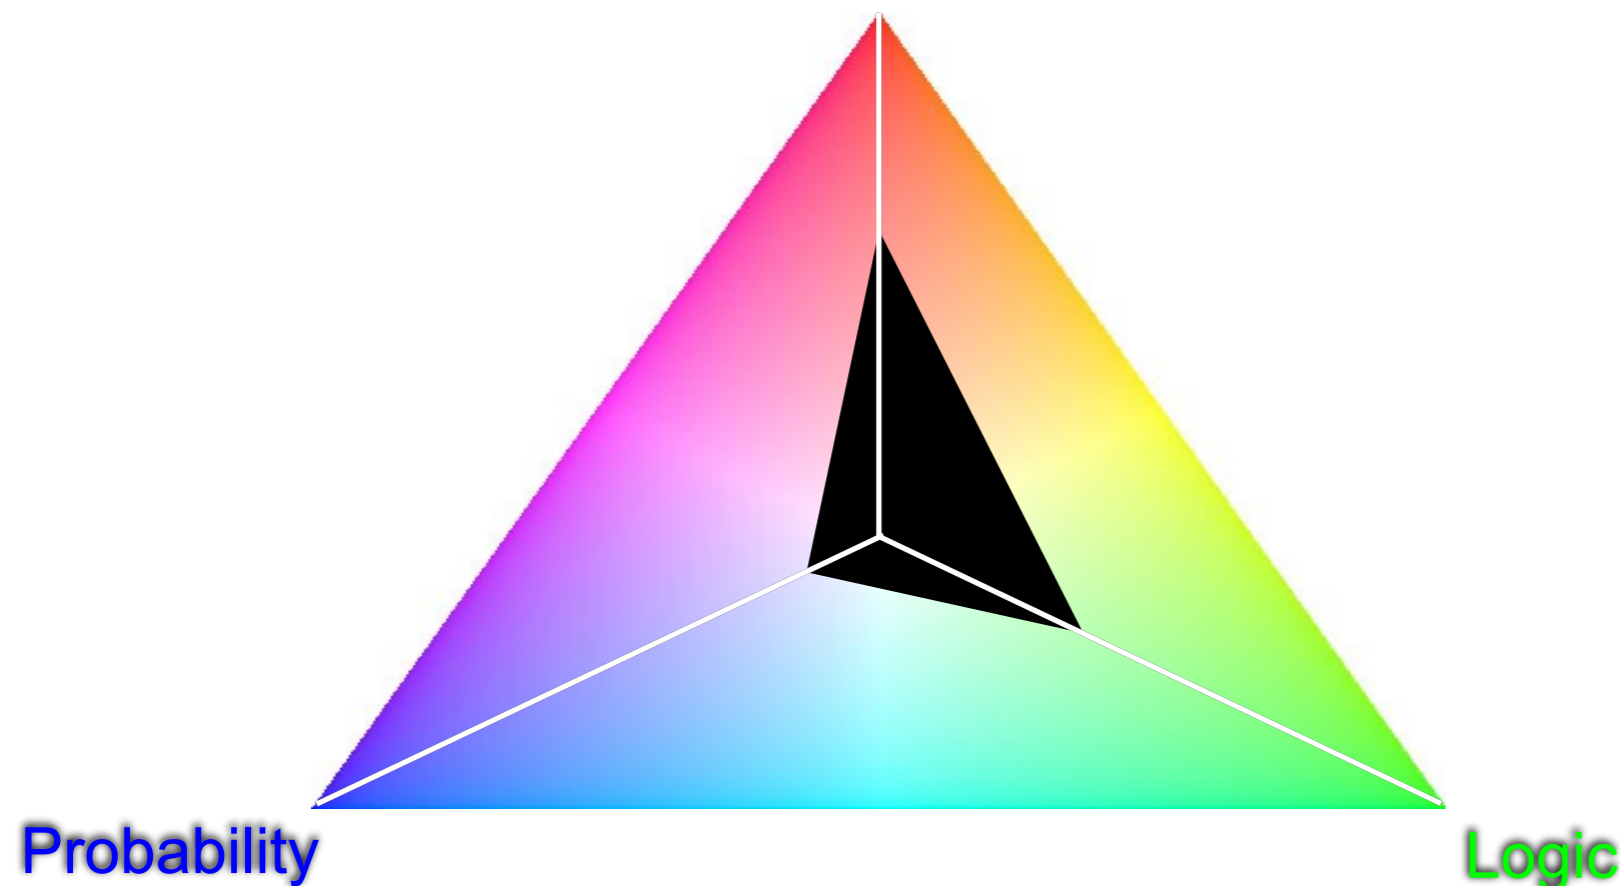
They use latent spaces, typical of neural computation to encode a relational structure of the data.

Neural networks cannot be recovered.

Logic is declined to encoding relations

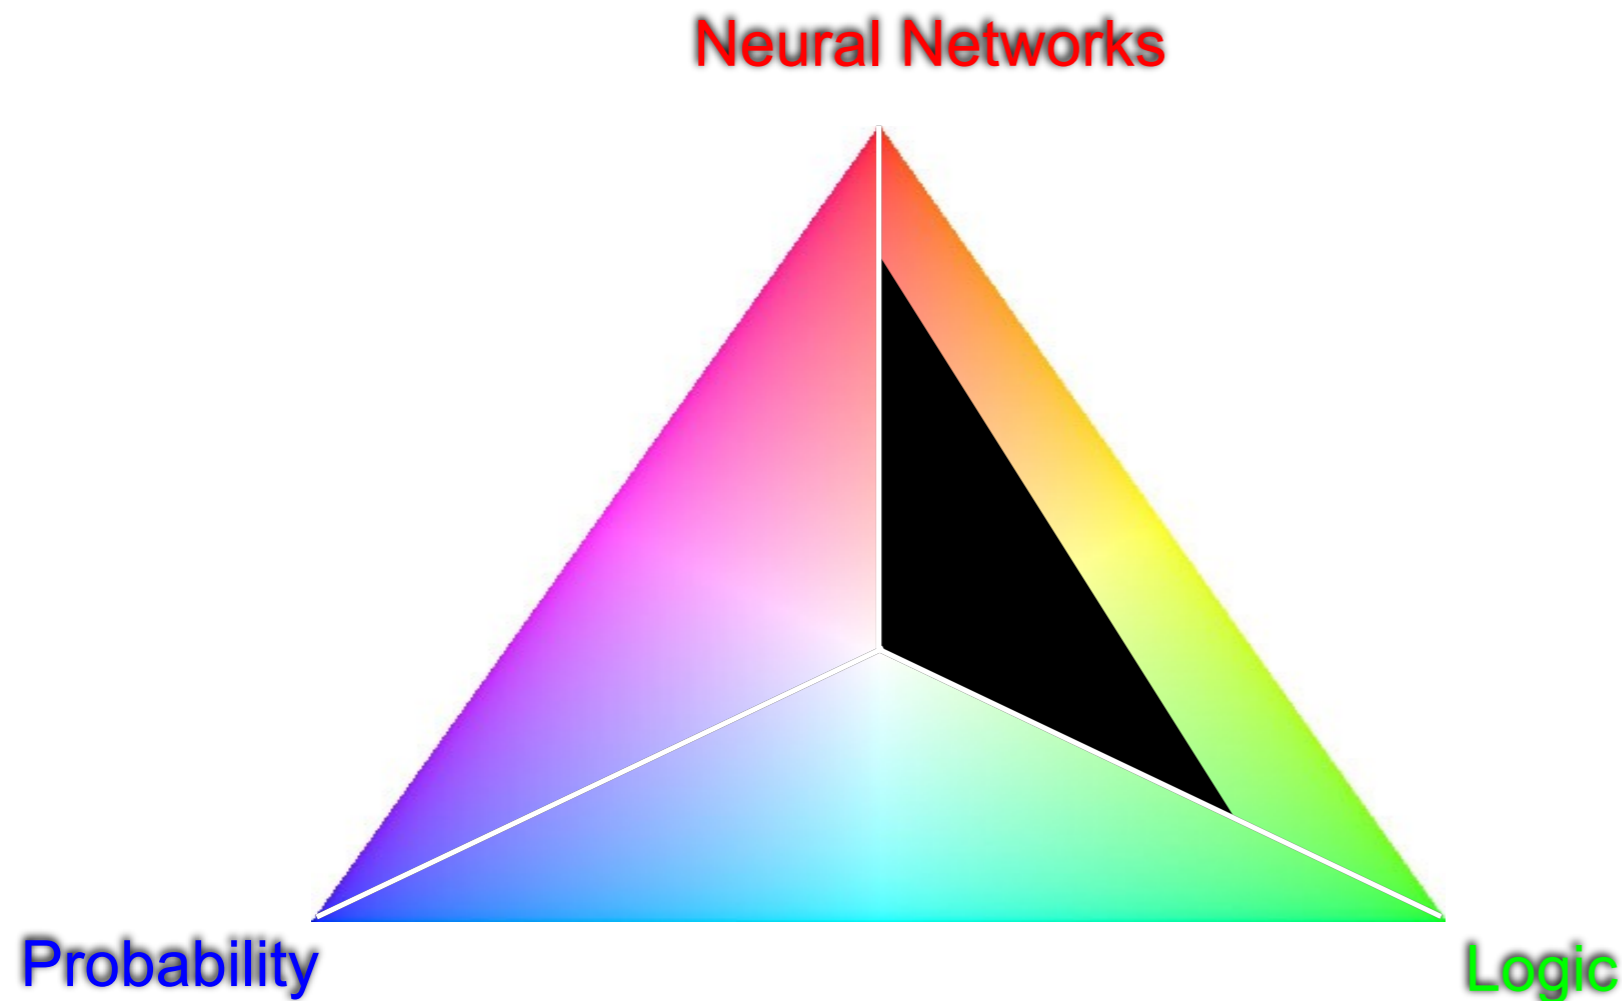
Probabilistic modelling is strongly approximated (e.g. atom mean field)

Most scalable solutions.



TransE (Bordes 2013)
DistMult (Yang, 2014)
Complex (Trouillon, 2016)
NTN (Socher, 2013)

Relaxed theorem provers

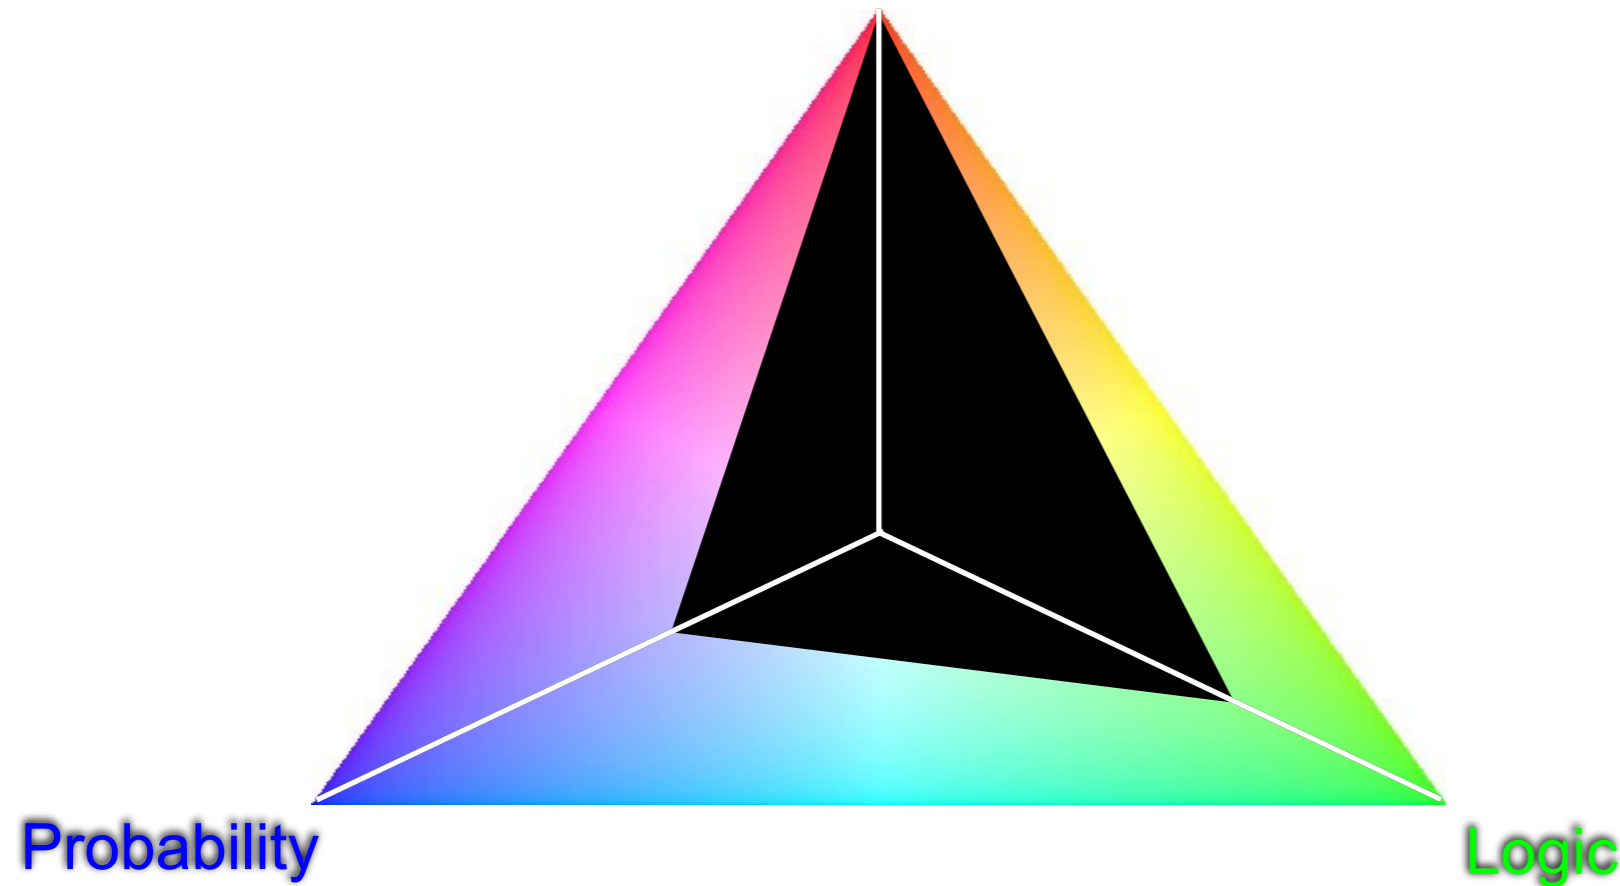


They sacrifice a bit the pure boolean semantics to obtain some soft neural capabilities (weighted reasoning, embeddings).

KBANN (Towell 1994)
LRNN (Sourek, 2017)
NTPs (Rocktäschel, 2017)
DiffLog (Si et al, 2018)
NN for Relational Data (2019)

Regularization methods

Neural Networks



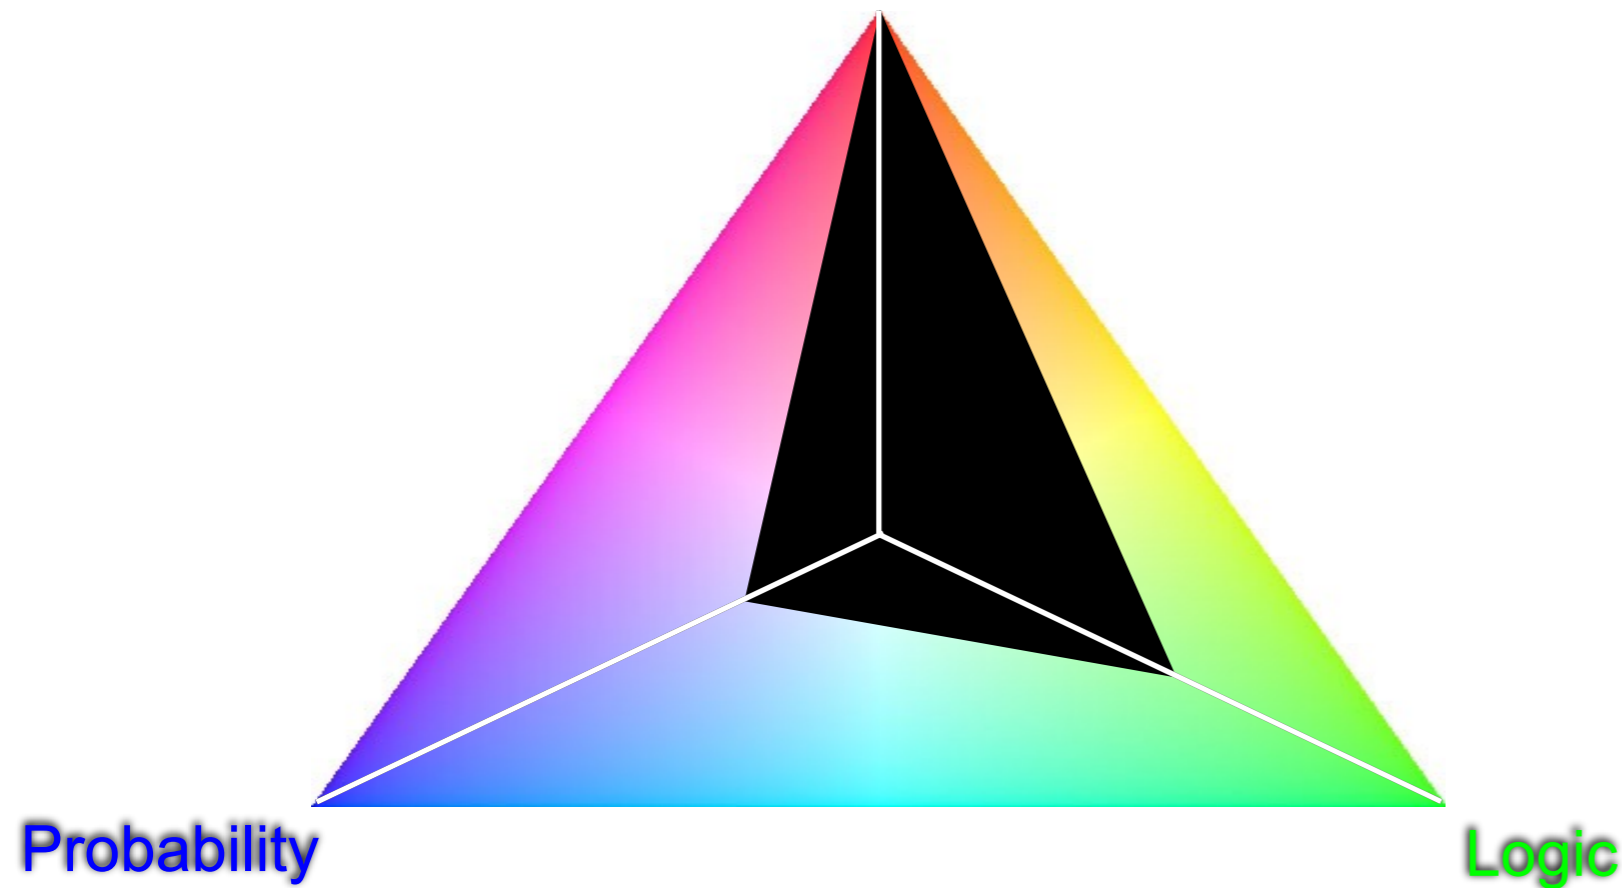
They sacrifice the logic and probability a lot by pushing everything inside the weights of the neural network.

Logic and probability are used only at training time. At inference time, only the neural net is used.

SBR (Diligenti et al, AI 2017)
LTN (Donatello et al, IJCAI 2017)
SL (Xu et al, ICML 2018)

Graph Neural Networks

Neural Networks



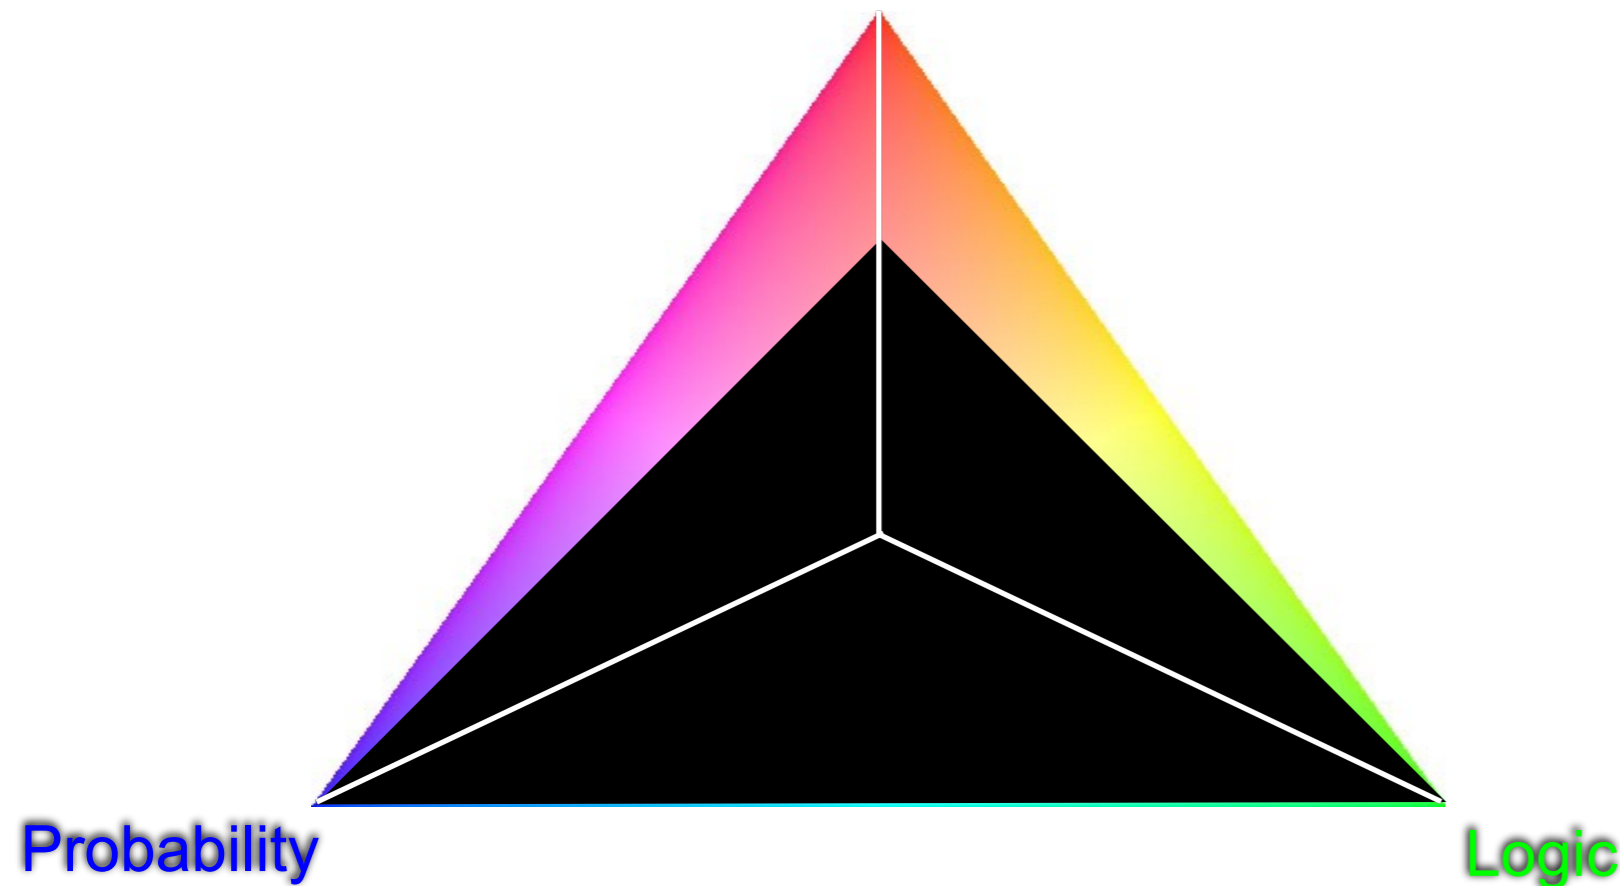
They extend neural network to provide some relational and multihop reasoning.

Logical semantics is not preserved.

R-GCN - Schlichtkrull et al, 2017

Probabilistic reparameterization

Neural Networks



They extend StarAI with perception capabilities.

Subsymbols at the level of the constants only

- Not at the level of the atoms (like KGE)
- Not at the level of the rules (like GNNs)

One of the most promising direction for NeSy.

Main problem is scalability.

DeepProbLog (Manhaeve, 2018)
RNM (Marra, 2020)

7. Logic vs Probability vs Neural Key Messages

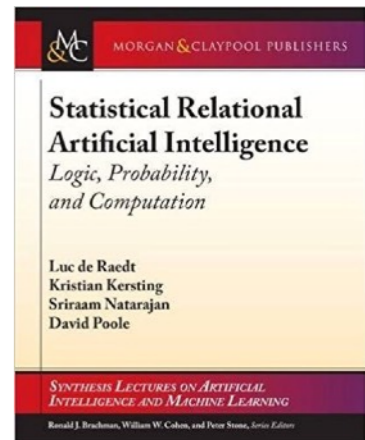
- We have three paradigms in the NeSy spectrum: **Logic**, **Probability** and **Neural Networks**
- An **integration** of the three should have the original paradigms as **special cases**
 - Computationally complex
- The integration is usually achieved by sacrificing the base paradigms
 - More scalable

Challenges

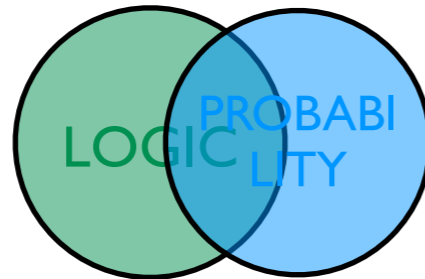
- For NeSy,
 - scaling up
 - which models and which knowledge to use
 - large scale life applications
 - peculiarities of neural nets & fuzzy logic
 - dynamics / continuous
- This is an excellent area for starting researchers / PhDs

Conclusions

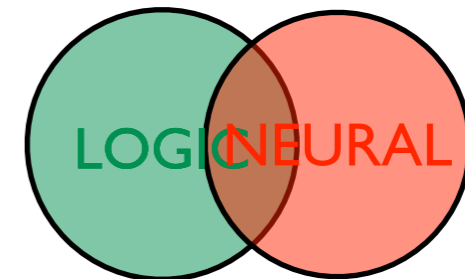
Key Message



FROM



TO



**StarAI and NeSy share similar problems
and thus similar solutions apply**

See also [De Raedt et al., IJCAI 20]

The Seven Dimensions

1. Proof vs Model based
2. Directed vs Undirected
3. Type of Logic
4. Symbols vs Subsymbols
5. Parameter vs Structure Learning
6. Semantics
7. Logic vs Probability vs Neural

Many questions to ask

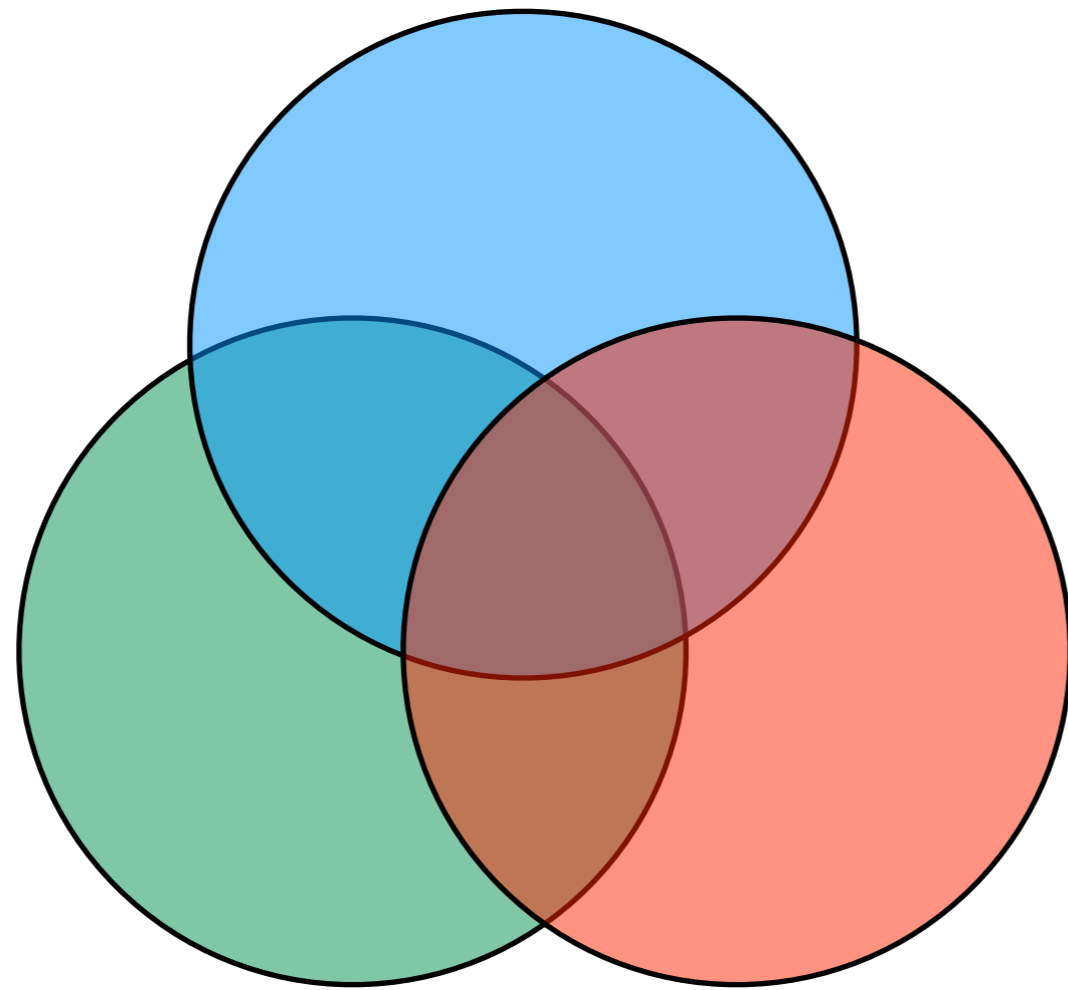
- What properties should integrated representations satisfy ?
 - Should one representation take over ? (As in most approaches to NeSy — push the logic inside and forget about it afterwards)
 - Should one build a pipeline or an interface between the integrated representations ?
 - Should one have the originals as a special case ?
 - (yes we believe you should be able to do all what you can do with the original representations)

Many questions to ask

- Which learning and reasoning techniques apply ?
 - Can you still reason logically / probabilistically ?
 - Can you still apply standard learning methods (like gradient descent) ?
 - Is everything explainable / trustworthy ?
- How to evaluate integrated representations ?
 - $1 + 1 = 3$?
 - Can they do what the originals can do, and can they do more ?
 - Can they do something different ?

Challenges

- For NeSy,
 - scaling up
 - which models to use
 - real life applications
 - peculiarities of neural nets
 - logical inference can be expensive
- **This is an excellent area for starting researchers / PhDs**



THANKS

References

- Tarek R. Besold, Artur S. d'Avila Garcez, Sebastian Bader, Howard Bowman, Pedro M. Domingos, Pascal Hitzler, Kai-Uwe Kühnberger, Luís C. Lamb, Daniel Lowd, Priscila Machado Vieira Lima, Leo de Penning, Gadi Pinkas, Hoifung Poon, and Gerson Zaverucha. Neural-symbolic learning and reasoning: A survey and interpretation. CoRR, abs/1711.03902, 2017.
- Matko Bošnjak, Tim Rocktäschel, Jason Naradowsky, and Sebastian Riedel. Programming with a differentiable forth interpreter. In ICML, 2017.
- William W. Cohen, Fan Yang, and Kathryn Mazaitis. Tensorlog: Deep learning meets probabilistic dbs. CoRR, abs/1707.05390, 2017.
- Andrew Cropper. Playgol: Learning programs through play. In IJCAI 2019, 2019.
- Andrew Cropper and Stephen H. Muggleton. Metagol system. <https://github.com/metagol/metagol>, 2016.
- Adnan Darwiche. Sdd: A new canonical representation of propositional knowledge bases. In IJCAI, 2011.
- Artur S. d'Avila Garcez, Marco Gori, Luís C. Lamb, Luciano Serafini, Michael Spranger, and Son N. Tran. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. FLAP, 6, 2019.
- Luc De Raedt, Sebastian Dumančić., Robin Manhaeve and Giuseppe Marra. From statistical relational to neuro-symbolic artificial intelligence. In IJCAI 2020.
- Luc De Raedt. Logical and relational learning. Springer, 2008.
- Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole. Statistical Relational Artificial Intelligence: Logic, Probability, and Computation. Morgan & Claypool Publishers, 2016.

References

- Luc De Raedt and Angelika Kimmig. Probabilistic (logic) programming concepts. *Machine Learning*, 100, 2015.
- Luc De Raedt, Robin Manhaeve, Sebastijan Dumančič, Thomas Demeester, and Angelika Kimmig. Neuro-symbolic= neural+ logical+probabilistic. In *NeSy @ IJCAI*, 2019.
- Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. Lifted rule injection for relation embeddings. In *EMNLP*, 2016.
- Michelangelo Diligenti, Marco Gori, and Claudio Saccà. Semantic-based regularization for learning and inference. *Artif. Intell.*, 244, 2017.
- Ivan Donadello, Luciano Serafini, and Artur S. d'Avila Garcez. Logic tensor networks for semantic image interpretation. In *IJCAI*, 2017.
- Honghua Dong, Jiayuan Mao, Tian Lin, Chong Wang, Lihong Li, and Denny Zhou. Neural logic machines. In *ICLR*, 2019.
- Sebastijan Dumančič, Tias Guns, Wannes Meert, and Hendrik Blockeel. Learning relational representations with auto-encoding logic programs. In *IJCAI*, 2019.
- Kevin Ellis, Lucas Morales, Mathias Sablé-Meyer, Armando Solar-Lezama, and Josh Tenenbaum. Learning libraries of subroutines for neurally-guided bayesian program induction. In *NeurIPS*, 2018.
- Kevin Ellis, Maxwell I. Nye, Yewen Pu, Felix Sosa, Josh Tenenbaum, and Armando Solar-Lezama. Write, execute, assess: Program synthesis with a REPL. *CoRR*, abs/1906.04604, 2019.
- Richard Evans and Edward Grefenstette. Learning explanatory rules from noisy data. *J. Artif. Intell. Res.*, 61, 2018.

References

- Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens, and Luc De Raedt. Inference and learning in probabilistic logic programs using weighted boolean formulas. *Theory and Practice of Logic Programming*, 15, 2015.
- Peter Flach. *Simply Logical: Intelligent Reasoning by Example*. John Wiley & Sons, Inc., 1994.
- Nir Friedman, Lise Getoor, Daphne Koller, and Avi Pfeffer. Learning probabilistic relational models. In *IJCAI*, 1999.
- Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. Answer set solving in practice. *Synthesis lectures on artificial intelligence and machine learning*, 6, 2012.
- L. Getoor and B. Taskar, editors. *An Introduction to Statistical Relational Learning*. MIT Press, 2007.
- Francesco Giannini, Michelangelo Diligenti, Marco Gori, and Marco Maggini. On a convex logic fragment for learning and reasoning. *IEEE TFS*, 27, 2018. CV Radhakrishnan et al.: Preprint submitted to Elsevier
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- Goldman, O., Laticinnik, V., Naveh, U., Globerson, A., & Berant, J.. Weakly-supervised semantic parsing with abstract examples. *ACL 2018*
- Bernd Gutmann, Angelika Kimmig, Kristian Kersting, and Luc De Raedt. Parameter learning in probabilistic databases: A least squares approach. In *ECML&PKDD*, 2008.
- Manfred Jaeger. Model-theoretic expressivity analysis. In Luc De Raedt, Paolo Frasconi, Kristian Kersting, and Stephen Muggleton, editors, *Probabilistic Inductive Logic Programming - Theory and Applications*, volume 4911 of LNCS. Springer, 2008.

References

- Ashwin Kalyan, Abhishek Mohta, Oleksandr Polozov, Dhruv Batra, Prateek Jain, and Sumit Gulwani. Neural-guided deductive search for real-time program synthesis from examples. In ICLR, 2018.
- Kristian Kersting and Luc De Raedt. Bayesian logic programming: Theory and tool. In L. Getoor and B. Taskar, editors, An introduction to Statistical Relational Learning. MIT Press, 2007.
- Stanley Kok and Pedro Domingos. Learning the structure of markov logic networks. In ICML, 2005.
- Daphne Koller and Nir Friedman. Probabilistic Graphical Models - Principles and Techniques. MIT Press, 2009.
- Marco Lippi and Paolo Frasconi. Prediction of protein beta-residue contacts by markov logic networks with grounding-specific weights. Bioinform., 25, 2009.
- John W Lloyd. Foundations of logic programming. Springer Science & Business Media, 2012.
- Daniel Lowd and Pedro Domingos. Efficient weight learning for markov logic networks. In ECML&PKDD, 2007.
- Robin Manhaeve, Sebastijan Dumanžić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Deepprolog: Neural probabilistic logic programming. In NeurIPS, 2018.
- Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In ICLR, 2019.
- Giuseppe Marra, Michelangelo Diligenti, Francesco Giannini, Marco Gori, and Marco Maggini. Relational neural machines. In ECAI, 2020.
- Giuseppe Marra and Ondrej Kuželka. Neural markov logic networks. CoRR, abs/1905.13462, 2019.

References

- Pasquale Minervini, Matko Bošnjak, Tim Rocktäschel, Sebastian Riedel, and Edward Grefenstette. Differentiable reasoning on large knowledgebases and natural language. In AAI, 2020.
- Pasquale Minervini, Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. Adversarial sets for regularising neural link predictors. In UAI, 2017.
- Stephen Muggleton. Stochastic logic programs. *Advances in inductive logic programming*, 32, 1996.
- Maxwell I. Nye, Armando Solar-Lezama, Josh Tenenbaum, and Brenden M. Lake. Learning compositional rules via neural program synthesis. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- David Poole. The independent choice logic and beyond. In *Probabilistic Inductive Logic Programming - Theory and Applications*, volume 4911 of LNCS. Springer, 2008.
- Matthew Richardson and Pedro M. Domingos. Markov logic networks. *Machine Learning*, 62, 2006.
- Tim Rocktäschel and Sebastian Riedel. End-to-end differentiable proving. In NIPS, 2017.
- Tim Rocktäschel, Sameer Singh, and Sebastian Riedel. Injecting logical background knowledge into embeddings for relation extraction. In NAACL HLT, 2015.
- Stuart Russell. Unifying logic and probability. *Communications of the ACM*, 58, 2015.

References

- Xujie Si, Mukund Raghothaman, Kihong Heo, and Mayur Naik. Synthesizing datalog programs using numerical relaxation. In IJCAI, 2019.
- Lazar Valkov, Dipak Chaudhari, Akash Srivastava, Charles A. Sutton, and Swarat Chaudhuri. Houdini: Lifelong learning as program synthesis. In NeurIPS, 2018.
- Guy Van den Broeck, Dan Suciu, et al. Query processing on probabilistic data: A survey. Foundations and Trends® in Databases, 7, 2017.
- Emile van Krieken, Erman Acar, and Frank van Harmelen. Analyzing differentiable fuzzy logic operators. CoRR, abs/2002.06100, 2020.
- Wenya Wang and Sinno Jialin Pan. Integrating deep learning with logic fusion for information extraction. CoRR, abs/1912.03041, 2019.
- Wang, P., Wu, Q., Shen, C., Hengel, A. V. D., & Dick, A. . Explicit knowledge-based reasoning for visual question answering. IJCAI 2017
- Leon Weber, Pasquale Minervini, Jannes Münchmeyer, Ulf Leser, and Tim Rocktäschel. Nlprolog: Reasoning with weak unification for question answering in natural language. In ACL, 2019.
- Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. A semantic loss function for deep learning with symbolicknowledge. In ICML, 2018.
- Fan Yang, Zhilin Yang, and William W Cohen. Differentiable learning of logical rules for knowledge base reasoning. In NIPS, 2017.
- Zhun Yang, Adam Ishay, and Joohyung Lee. Neurasp: Embracing neural networks into answer set programming. In Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI, pages 1755–1762,

References

- Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Josh Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In NeurIPS, 2018.
- Lotfi A Zadeh. Fuzzy logic and approximate reasoning. *Synthese*, 30(3-4):407–428, 1975.
- Pedro Zuidberg Dos Martires, Vincent Derkinderen, Robin Manhaeve, Wannes Meert, Angelika Kimmig, and Luc De Raedt. Transforming probabilistic programs into algebraic circuits for inference and learning. In Program Transformations for ML Workshop at NeurIPS, 2019.
- Gustav Šourek, Vojtech Aschenbrenner, Filip Zelezný, Steven Schockaert, and Ondrej Kuželka. Lifted relational neural networks: Efficient learning of latent relational structures. *J. Artif. Intell. Res.*, 62, 2018