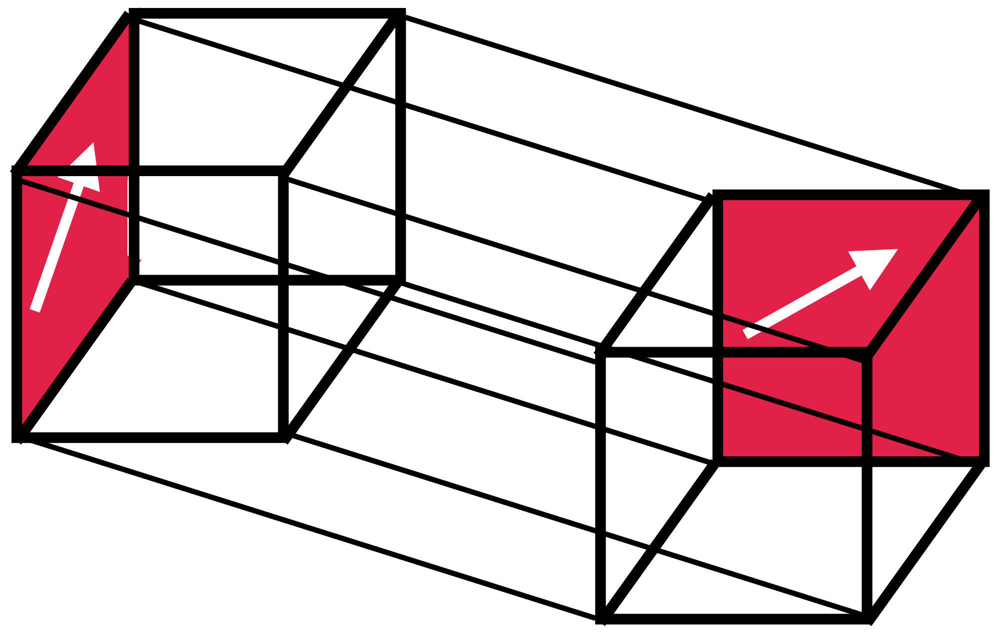

Neural+Symbolic Learning

A Probabilistic Journey

Jaron Maene - DTAI Seminar - 7/5/2024

Outline

1. What is “neural-symbolic”?
   => Introduce the *probabilistic view* from first-principles
2. Contribution 1: embeddings & neural-symbolic
3. Contribution 2: complexity & gradients

Part 1:
What is neural-symbolic?

$$p(X)$$

$p(\text{Vehicle} = \text{ambulance} \mid \text{Image} =$  $)$

Supervised Learning

$$p_\theta(Y \mid X)$$

Find θ that maximises $p_\theta(Y \mid X)$ over dataset

X	Y
	ambulance
...	...

$$p(\text{Cross} \mid \text{Vehicle} = \text{car}, \text{Lights} = \text{red})$$

Logical Reasoning

$$p_\phi(X = x) = \mathbb{1}(x \models \phi)$$

To keep it simple: ϕ is propositional Boolean formula

$$\begin{aligned} \phi = (&\text{cross} \leftrightarrow \text{ambulance} \vee (\text{car} \wedge \text{green\_lights})) \\ &\wedge (\text{car} \leftrightarrow \neg\text{ambulance}) \\ &\wedge (\text{lights\_green} \leftrightarrow \neg\text{lights\_red}) \end{aligned}$$

$p(\text{Cross} \mid \text{Image} =$  $, \text{Lights} = \text{red}) = ?$

- Not just logic…
- Not just a neural network…

Marginalise

$p(\text{Cross} \mid \text{Image} =$  $, \text{Lights} = \text{red}) =$

$p(\text{Cross}, \text{Vehicle} = \text{car} \mid \text{Image} =$  $, \text{Lights} = \text{red})$

$+p(\text{Cross}, \text{Vehicle} = \text{ambulance} \mid \text{Image} =$  $, \text{Lights} = \text{red})$

Decompose

$p(\text{Cross} \mid \text{Image} =$  $, \text{Lights} = \text{red}) =$

$p(\text{V} = \text{car} \mid \text{I} =$  $) \cdot p(\text{C} \mid \text{V} = \text{car}, \text{L} = \text{red})$

$+p(\text{V} = \text{ambulance} \mid \text{I} =$  $) \cdot p(\text{C} \mid \text{V} = \text{ambulance}, \text{L} = \text{red})$

We re-invented DeepProbLog*!

Logic

Neural networks

$$p(Y \mid X) = \sum_{z} p(Y \mid Z = z) \prod_{z_i} p(Z_i = z_i \mid X)$$

possible worlds

Weighted Model Counting

$$p(Y \mid X) = \sum_{z:z \models Y} \prod_{i} p(Z_i = z_i \mid X)$$

- Linear combination of models.
- Generalisation of SAT solving.
- #P-hard in general, but decades of research have created strong solvers.

Neural-Symbolic Learning

Find θ such that

$p_\theta(\text{Cross} \mid \text{Vehicle} =$  $, \text{Lights} =$ 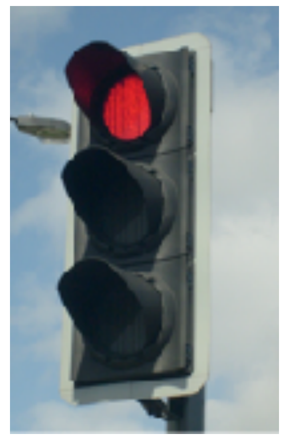 $) = 1$

=> Weakly-supervised

=> Find "right" solution (c.f. reasoning shortcuts)

Neural-Symbolic Learning

$$\nabla p(Y \mid X) = \sum_{z} p(Y \mid Z = z) \ \nabla p(Z = z \mid X)$$

End-to-end differentiable!

=> Parameter learning (vs. Structure Learning)

Why would you do this?

1. **ML view**: bias-variance trade-off.
   => Craft bias for your problem (c.f. probabilistic programming)
   => Better performance / less data required
2. **Safety view**: get guarantees, more trustworthy.
3. **Causality view**: out-of-distribution generalisation.
4. **KR view**: many more possible questions you can ask.
5. **Interpretability view**: why did my model do this?
6. **Psychology view**: Compositionality, System 1 vs. System 2

Part 2: Embeddings & Neural-Symbolic

$p(\text{Cross} \mid \text{Image} =$  $, \text{Lights} = \text{orange}) = ?$

$p(\text{Cross} \mid \text{Image} =$  $, \text{Lights} = \text{orange}) =$

$p(\text{orange} = \text{red})p(\text{Cross} \mid \text{Image} =$  $, \text{Lights} = \text{red})$
$+p(\text{orange} = \text{green})p(\text{Cross} \mid \text{Image} =$  $, \text{Lights} = \text{green})$

If you want a well-defined equality relation, you need a metric space.

In other words, probabilistic equality = **embeddings**

$$p(x = y) = e^{-d(e(x), e(y))}$$

1. $p(x = x)$
2. $p(x = y) = p(y = x)$
3. $p(x = y) \geq p(x = z) \cdot p(z = y)$

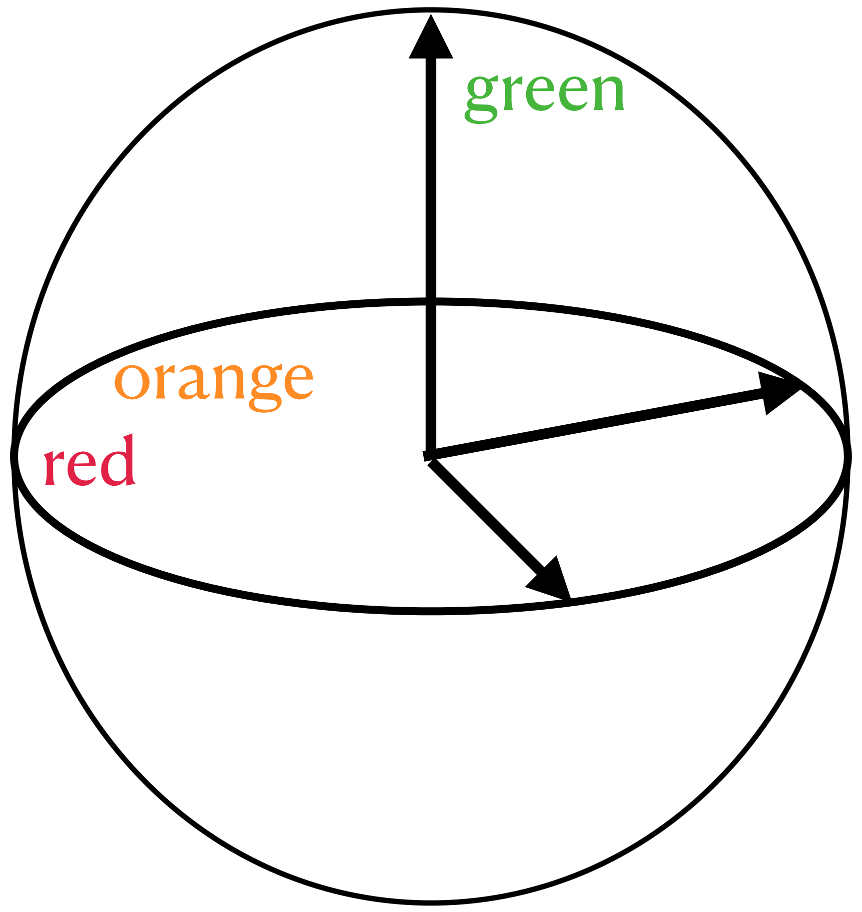

Soft-Unification

- Corresponds with Soft-Unification
- Corresponds with DeepProbLog: neural facts <=> embeddings.
- End-to-end differentiable learning of embeddings from proofs.
- Generalises knowledge graph embeddings

```
cross :- vehicle(car), light(green).
light(orange).
vehicle(car).
query(cross).
```

light(orange) == light(green)
=> Unification: No
=> Soft-unification: p(green == orange)

Part 3:
Complexity & Gradients

Reduction of ∇WMC to WMC

If you have a WMC $p(\phi) = \sum_{I \models \phi} \prod_{x \in I} p(x)$, then $\frac{\partial p(\phi)}{\partial x} = p(\phi \mid x) - p(\phi \mid \neg x)$.

Proof: Using linearity of WMC. (Theorem 3.1, Maene et al., 2024)

Corollaries

1. Computing $\nabla p(\phi)$ is #P-complete.
2. (ϵ, δ)-approximating $\nabla p(\phi)$ is NP-hard.

Scalable gradients for NeSy

- Gradients can be approximate. (No-one uses exact gradients!)
- We do want unbiased and (ϵ, δ)-approximate (i.e. high probability of being close to true gradient).
- Existing gradient estimators unbiased, or require exponential number of samples.

Score function estimator (aka REINFORCE): $\dfrac{\partial p(\phi)}{\partial x} = \mathbb{E}_{I \sim p(\cdot)} \left[\mathbb{1}(I \models \phi) \dfrac{\partial \log p(x)}{\partial x} \right]$

=> Does become provably tractable for low-entropy distributions.

Weighted Model Estimator (WeightME)

=> *Lot's of progress in approximate WMC solvers, can we leverage this for* ∇ *WMC?*

=> *Sample models instead of interpretations!*

$$\frac{\partial \log p(\phi)}{\partial x} = \mathbb{E}_{I \sim p(\cdot|\phi)} \left[\frac{\partial \log p(x)}{\partial x} \mathbb{1}(I \models x) + \frac{\partial \log p(\neg x)}{\partial x} \mathbb{1}(I \models \neg x) \right]$$

=> Under mild assumptions, (ϵ, δ)-approximation using constant number of samples.

=> Requires a logarithmic number of SAT calls.

References

- J. Maene & L. De Raedt, "**Soft-Unification in Deep Probabilistic Logic**", *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- J. Maene, V. Derkinderen & L. De Raedt, "**On the Hardness of Probabilistic Neurosymbolic Learning**", *International Conference on Machine Learning (ICML)*, 2024.

More general overview of Probabilistic NeSy:

- V. Derkinderen et al., "**Semirings for Probabilistic & Neuro-probabilistic Programming**", *International Journal of Approximate Reasoning*, 2024.
- G. Marra et al., "**From Statistical Relational to Neurosymbolic AI**", *Artificial Intelligence Journal*, 2024.