



logically-consistent deep learning via probabilistic circuits

antonio vergari (he/him)

 @tetraduazione

2nd Apr 2025 - **Neuro-explicit retreat** Saarbruecken

probabilistic circuits (PCs)

A grammar for tractable computational graphs

1. *A simple tractable function is a circuit*

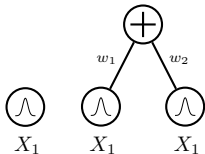
\Rightarrow *e.g., a multivariate Gaussian, or a logical literal*

\bigwedge
 X_1

probabilistic circuits (PCs)

A grammar for tractable computational graphs


- I. *A simple tractable function is a circuit*
- II. *A weighted combination of circuits is a circuit*



probabilistic circuits (PCs)

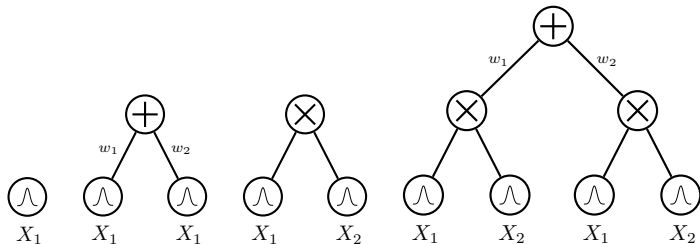
A grammar for tractable computational graphs

- I. A simple tractable function is a circuit
- II. A weighted combination of circuits is a circuit
- III. A product of circuits is a circuit



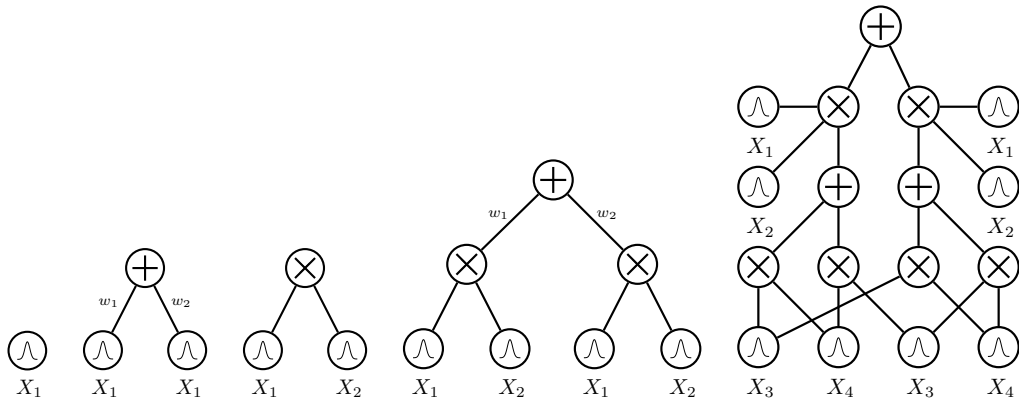
probabilistic circuits (PCs)

A grammar for tractable computational graphs



probabilistic circuits (PCs)

A grammar for tractable computational graphs



structural properties

smoothness

decomposability

compatibility

determinism

structural properties

smoothness

decomposability

compatibility

determinism

structural properties

smoothness

decomposability

compatibility

determinism

structural properties

smoothness


decomposability

compatibility


determinism

determinism



the inputs of sum units are defined over disjoint supports



deterministic circuit




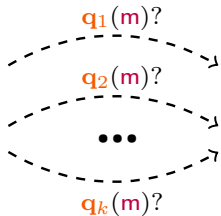
non-deterministic circuit



\approx

	X^1	X^2	X^3	X^4	X^5
x_8					
x_7					
x_6					
x_5					
x_4					
x_3					
x_2					
x_1					

generative models that can reason probabilistically



\approx

	X^1	X^2	X^3	X^4	X^5
x_8					
x_7					
x_6					
x_5					
x_4					
x_3					
x_2					
x_1					

but some events should have zero probabilities!

Goal

***“How can neural nets
reason and learn with
symbolic constraints
reliably and efficiently?”***

the issues!

- I) Logical constraints can be hard to represent in a unified way
⇒ ***a single framework** for implications, negation, paths, hierarchies, ...*
- II) How to integrate logic and probabilities in a single architecture
⇒ ***combining soft and hard constraints***
- III) Logical constraints are piecewise constant functions!
⇒ *differentiable almost everywhere but **gradient is zero!***

the issues!

- I) Logical constraints can be hard to represent in a unified way
⇒ ***a single framework*** for implications, negation, paths, hierarchies, ...
- II) How to integrate logic and probabilities in a single architecture
⇒ ***combining soft and hard constraints***
- III) Logical constraints are piecewise constant functions!
⇒ differentiable almost everywhere but ***gradient is zero!***

the issues!

- I) Logical constraints can be hard to represent in a unified way
⇒ ***a single framework*** for implications, negation, paths, hierarchies, ...
- II) How to integrate logic and probabilities in a single architecture
⇒ ***combining soft and hard constraints***
- III) Logical constraints are piecewise constant functions!
⇒ differentiable almost everywhere but ***gradient is zero!***

hard vs soft constraints

logic vs probabilities

logic

“If X is a bird, X flies”

$$A(X) \implies B(X)$$

prob logic

“How likely is that if X is a bird, X flies?”

$$p(A(X) \implies B(X))$$

which logic?

or which kind of constraints to represent?

propositional logic (zeroth-order)

$$(a \wedge b) \vee d \implies c$$

first-order logic (FOL)

$$\forall a \exists b : R(a, b) \vee Q(d) \implies C(x)$$

satisfiability modulo theory (SMT)

$$(\alpha X_i - \beta X_j \leq 100) \vee (X_j + X_k \geq 0) \implies (X_j X_k \leq X_i)$$

which logic?

or which kind of constraints to represent?

propositional logic (zeroth-order)

$$(a \wedge b) \vee d \implies c$$

first-order logic (FOL)


$$\forall a \exists b : R(a, b) \vee Q(d) \implies C(x)$$

satisfiability modulo theory (SMT)

$$(\alpha X_i - \beta X_j \leq 100) \vee (X_j + X_k \geq 0) \implies (X_j X_k \leq X_i)$$


Goal

***“How can we enforce
symbolic constraints
in a probability distribution?”***




$q(\mathbf{x})$

start from a distribution $q(\mathbf{x})$...




$q(\mathbf{x})$




$c(\mathbf{x})$


...and cut its support by a constraint $c(\mathbf{x})$



$q(\mathbf{x})$




$c(\mathbf{x})$




$q(\mathbf{x}) \cdot c(\mathbf{x})$


by multiplying them $q(\mathbf{x})c(\mathbf{x})...$



$q(\mathbf{x})$






$c(\mathbf{x})$



$\frac{q(\mathbf{x}) \cdot c(\mathbf{x})}{\sum_{\mathbf{x}} q(\mathbf{x}) \cdot c(\mathbf{x})}$

and then renormalizing them!


 $q(\mathbf{x})$

 $c(\mathbf{x})$

 $\frac{q(\mathbf{x}) \cdot c(\mathbf{x})}{\sum_{\mathbf{x}} q(\mathbf{x}) \cdot c(\mathbf{x})}$

**states with zero probability will never be predicted
(nor sampled)**

Goal

*Can we design q and c
to be **expressive models**
yet yielding a tractable product?*

Goal


*Can we design q and c
to be **deep computational graphs**
yet yielding a tractable product?*

Goal

*Can we design q and c
to be **deep computational graphs**
yet yielding a tractable product?*


*yes! as **circuits!***

Tractable products



exactly compute \mathcal{Z} in time $O(|q||c|)$

knowledge compilation



compiling logical formulas into circuits

(smooth, structured-decomposable, deterministic)

knowledge compilation

$$K : (Y_1 = 1 \implies Y_3 = 1)$$

$$\wedge (Y_2 = 1 \implies Y_3 = 1)$$

$$\mathbb{1}\{Y_1 = 0\} \odot$$

$$\mathbb{1}\{Y_1 = 1\} \odot$$

$$\mathbb{1}\{Y_2 = 0\} \odot$$

$$\mathbb{1}\{Y_2 = 1\} \odot$$


$$\mathbb{1}\{Y_3 = 0\} \odot$$

$$\mathbb{1}\{Y_3 = 1\} \odot$$

knowledge compilation

$$K : (Y_1 = 1 \implies Y_3 = 1)$$


$$\wedge (Y_2 = 1 \implies Y_3 = 1)$$



knowledge compilation

$$K : (Y_1 = 1 \implies Y_3 = 1)$$


$$\wedge (Y_2 = 1 \implies Y_3 = 1)$$



knowledge compilation

$$K : (Y_1 = 1 \implies Y_3 = 1)$$


$$\wedge (Y_2 = 1 \implies Y_3 = 1)$$



knowledge compilation

$$K : (Y_1 = 1 \implies Y_3 = 1)$$

$$\wedge (Y_2 = 1 \implies Y_3 = 1)$$




 main

cirkit / notebooks / logic-circuits.ipynb

PreviewCodeBlame1183 lines (1183 loc) · 81.5 KBRaw

Compiling a propositional formula using cirkit

As we said before, a propositional formula can be easily represented as a circuit with a tree-structure. For instance, α presented before can be represented as the following tree:



check the notebook

how

to enforce constraints?

$$\max p_{\theta}(\mathbf{K}_i)$$

maximise the probability of the constraint to hold!

*Xu et al., "A Semantic Loss Function for Deep Learning with Symbolic Knowledge",
Proceedings of the 35th International Conference on Machine Learning (ICML), 2018*

how

to enforce constraints?

$$\min \mathcal{L}(\mathbf{K}_i, p_\theta) = \min -\log \sum_{\mathbf{z} \models \mathbf{K}_i} \prod_{j: \mathbf{z} \models z_{f_j}} p_\theta(z_{f_j}) \prod_{j: \mathbf{z} \models \neg z_{f_j}} (1 - p_\theta(z_{f_j}))$$

minimize the semantic loss

*Xu et al., "A Semantic Loss Function for Deep Learning with Symbolic Knowledge",
Proceedings of the 35th International Conference on Machine Learning (ICML), 2018*

computing the probability of logical formulas

$$p_{\theta}(\mathbf{K}(\mathbf{z})) = \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [\mathbb{1}\{\mathbf{z} \models \mathbf{K}\}]$$

computing the probability of \mathbf{K}

WMC

computing the probability of logical formulas

$$\mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [\mathbb{1}\{\mathbf{z} \models K\}] = \sum_{\mathbf{z}} p(\mathbf{z}) \mathbb{1}\{\mathbf{z} \models K\} = \sum_{\mathbf{z} \models K} p(\mathbf{z})$$

computing the **weighted model count** (WMC) of K

computing the probability of logical formulas

$$\mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [\mathbb{1}\{\mathbf{z} \models \mathbf{K}\}] = \sum_{\mathbf{z} \models \mathbf{K}} \prod_{i: \mathbf{z} \models z_i} p(z_i) \prod_{i: \mathbf{z} \models \neg z_i} (1 - p(z_i))$$

*assuming independence of \mathbf{z} (but be careful!)*¹

¹van Krieken et al., “On the Independence Assumption in Neurosymbolic Learning”, 2024
Xu et al., “A Semantic Loss Function for Deep Learning with Symbolic Knowledge”,
Proceedings of the 35th International Conference on Machine Learning (ICML), 2018

WMC

computing the probability of logical formulas

$$\mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [\mathbb{1}\{\mathbf{z} \models \mathbf{K}\}] = \sum_{\mathbf{z} \models \mathbf{K}} \prod_{i: \mathbf{z} \models z_i} p(z_i) \prod_{i: \mathbf{z} \models \neg z_i} (1 - p(z_i))$$

computing WMC is #P-hard in general : (

*Xu et al., "A Semantic Loss Function for Deep Learning with Symbolic Knowledge",
Proceedings of the 35th International Conference on Machine Learning (ICML), 2018*


Goal

***Can we encode K
to yield a tractable WMC?***

Goal

***Can we encode K
to yield a tractable WMC?
yes, as a **circuit**!***

tractable WMC



exactly compute **WMC** in time $O(|c|)$

SL recipe


$$K : (Y_1 = 1 \implies Y_3 = 1)$$

$$\wedge (Y_2 = 1 \implies Y_3 = 1)$$

1) Take a
logical constraint

SL recipe

$$K : (Y_1 = 1 \implies Y_3 = 1) \\ \wedge (Y_2 = 1 \implies Y_3 = 1)$$




1) Take a
logical constraint

2) Compile it into
a constraint circuit

SL recipe

$$K : (Y_1 = 1 \implies Y_3 = 1) \\ \wedge (Y_2 = 1 \implies Y_3 = 1)$$



1) Take a
logical constraint


2) Compile it into
a constraint circuit

$$-\log \text{WMC}(K_i, p_\theta)$$

3) minimize the semantic loss

SL recipe

$$K : (Y_1 = 1 \implies Y_3 = 1) \\ \wedge (Y_2 = 1 \implies Y_3 = 1)$$




$$-\log \text{WMC}(K_i, p_\theta)$$

1) Take a
logical constraint

2) Compile it into
a constraint circuit

3) minimize the semantic loss

4) train end-to-end by sgd!

 main

circuit / notebooks / semantic-loss.ipynb

PreviewCodeBlame625 lines (625 loc) · 62.2 KBRawCopyDownload

Semantic loss using circuit

So far, we used our logic circuit in a *traditional* setting, by running queries related to its logical semantics. However, logic circuits also enable neuro-symbolic methods! In this notebook, we implement the **semantic loss** (Xu et al, 2018), which relies on a logic circuit to implement constraints on the predictions of a neural network, using an SDD.

Imagine we want to enforce constraints on the output \hat{y} of a neural network and we want to model those constraints as a logic formula F by considering each element of \hat{y} as a literal in F . We can compile the formula F into a logic circuit c and we can efficiently perform a bunch of operations on it (see [logic circuits](#)), but can we actually use the values of \hat{y} as inputs to c ? Or in other words, can we quantify if the input y is a model for c ($y \models c$)?

<https://github.com/n28div/circuit/blob/main/notebooks/semantic-loss.ipynb>

nice...

...but!

assuming facts to be independent...

nice...

...but!

assuming facts to be independent...

***no guarantees to satisfy
constraints at test time...***

nice...

...but!

assuming facts to be independent...

***no guarantees to satisfy
constraints at test time...***

WMC

on the independence assumption

$$K : \neg \textcolor{red}{r} \vee \neg \textcolor{green}{g}$$

*a neural net should not output that a traffic light is both **red** and **green***

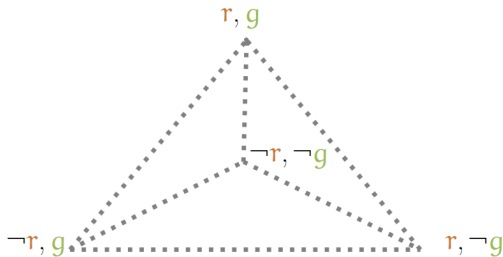


WMC

on the independence assumption

$$K : \neg \mathbf{r} \vee \neg \mathbf{g}$$

a neural net should not output that a traffic light is both **red** and **green**



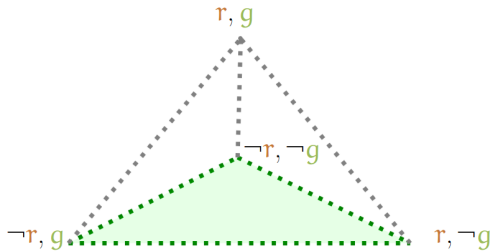
WMC

on the independence assumption

$$K : \neg \mathbf{r} \vee \neg \mathbf{g}$$

a neural net should not output that a traffic light is both **red** and **green**

only some probability assignments should be non-zero (lower triangle)




WMC

on the independence assumption

$$K : \neg \mathbf{r} \vee \neg \mathbf{g}$$

a neural net should not output that a traffic light is both **red** and **green**

but assuming $p(\mathbf{r}, \mathbf{g}) = p(\mathbf{r})p(\mathbf{g})$ restricts this even further (only blue lines)



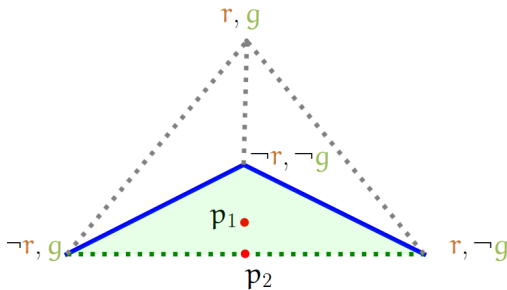
WMC

on the independence assumption

$$K : \neg \mathbf{r} \vee \neg \mathbf{g}$$

a neural net should not output that a traffic light is both **red** and **green**

but assuming $p(\mathbf{r}, \mathbf{g}) = p(\mathbf{r})p(\mathbf{g})$ restricts this even further (only blue lines)



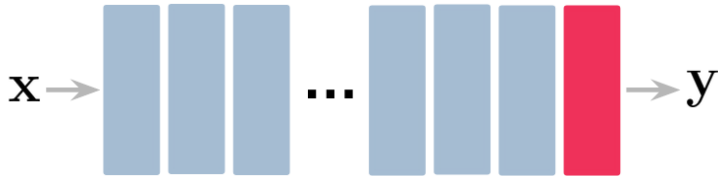
nice...

...but!

assuming facts to be independent...

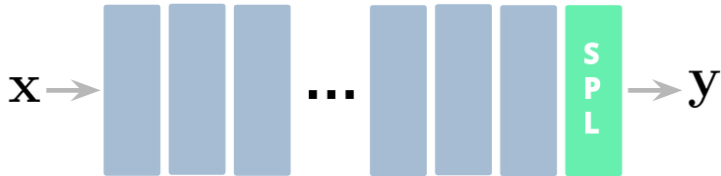
***no guarantees to satisfy
constraints at test time...***

how to



make any neural network architecture...

how to



...guarantee all predictions to conform to constraints?

Semantic Probabilistic Layers for Neuro-Symbolic Learning

Kareem Ahmed
CS Department
UCLA
ahmedk@cs.ucla.edu


Stefano Teso
CIMEC and DISI
University of Trento
stefano.teso@unitn.it

Kai-Wei Chang
CS Department
UCLA
kwchang@cs.ucla.edu

Guy Van den Broeck
CS Department
UCLA
guyvdb@cs.ucla.edu

Antonio Vergari
School of Informatics
University of Edinburgh
avergari@ed.ac.uk

enforce constraints in neural networks at NeurIPS 2022



$$p(\mathbf{y} | \mathbf{x}) = \mathbf{q}_{\Theta}(\mathbf{y} | g(\mathbf{z})) \cdot \mathbf{c}_{\mathbf{K}}(\mathbf{x}, \mathbf{y}) / \mathbf{Z}(\mathbf{x})$$

$$\mathbf{Z}(\mathbf{x}) = \sum_{\mathbf{y}} \mathbf{q}_{\Theta}(\mathbf{y} | \mathbf{x}) \cdot \mathbf{c}_{\mathbf{K}}(\mathbf{x}, \mathbf{y})$$

SPL



Ground Truth



ResNet-18



Semantic Loss



circuits

predictions guarantee a logical constraint 100% of the time!


SPL

(and variants)

everywhere

Tractable Control for Autoregressive Language Generation


Honghua Zhang^{*1} Meihua Dang^{*1} Nanyun Peng¹ Guy Van den Broeck¹



constrained text generation with LLMs (ICML 2023)

Safe Reinforcement Learning via Probabilistic Logic Shields

Wen-Chi Yang¹, Giuseppe Marra¹, Gavin Rens and Luc De Raedt^{1,2}



reliable reinforcement learning (AAAI 23)

How to Turn Your Knowledge Graph Embeddings into Generative Models

Lorenzo Loconte
University of Edinburgh, UK
l.loconte@sms.ed.ac.uk





























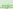

Nicola Di Mauro
University of Bari, Italy
nicola.dimauro@uniba.it

Robert Peharz
TU Graz, Austria
robert.peharz@tugraz.at

Antonio Vergari
University of Edinburgh, UK
avergari@ed.ac.uk

enforce constraints in knowledge graph embeddings
oral at NeurIPS 2023

Logically Consistent Language Models via Neuro-Symbolic Integration

 = LLaMa 2  = LLaMa 2 	Forward Implication	Reverse Implication	Negation
	$A \rightarrow B$ A: (albatross, isA, bird) B: (albatross, isA, fish)	$\neg B \rightarrow \neg A$ B: (albatross, isNotA, organism) A: (albatross, isNotA, living thing)	$A \leftrightarrow A$ A: (computer, isA, airplane) A: (computer, isNotA, airplane)
	<div>Is an albatross a bird? </div> <div> Yes.</div> <div>Is an albatross a fish? </div> <div> Yes. Logical:  Factual: </div> <div> No. Logical:  Factual: </div>	<div>Is it true that an albatross is not an organism? </div> <div> No.</div> <div>Is it true that an albatross is not a living thing? </div> <div> Yes. Logical:  Factual: </div> <div> No. Logical:  Factual: </div>	<div>Is a computer a airplane? </div> <div> No.</div> <div>Is it true that a computer is not a airplane? </div> <div> No. Logical:  Factual: </div> <div> Yes. Logical:  Factual: </div>

improving logical (self-)consistency in LLMs at ICLR 2025

open problems


I constraints over continuous variables

II scaling to H U G E constraints


III learn (partial) constraints

IV revise constraints (continual learning)


Ground Truth




GMM



Flow



PAL (*ours*)



extending it to SMT constraints

A Probabilistic Neuro-symbolic Layer for Algebraic Constraint Satisfaction

Leander Kurscheidt¹

Paolo Morettin²

Roberto Sebastiani²


Andrea Passerini²

Antonio Vergari¹


¹School of Informatics, University of Edinburgh, UK

²DiSI, University of Trento, Italy


extending it to SMT constraints



SPL & LTN & DPL




CBMs




NN + ex-post


NeSy models are concept bottlenecks



SPL & LTN & DPL



CBMs




NN + ex-post

NeSy models can suffer from reasoning shortcuts!

Task	Example Data	Knowledge K	Example RS	Impact
MNIST math	$\begin{cases} 2 \cdot \mathbf{2} + \mathbf{2} = 6 \\ \mathbf{3} + \mathbf{4} = 7 \end{cases}$	Equations must hold.	$\begin{cases} \mathbf{2} \rightarrow 2 \\ \mathbf{3} \rightarrow 4 \\ \mathbf{4} \rightarrow 3 \end{cases}$	$\mathbf{2} + \mathbf{4} = \mathbf{5}$

NeSy models can suffer from reasoning shortcuts!



how to **detect** and **mitigate** them


Marconato et al., "Not all neuro-symbolic concepts are created equal: Analysis and mitigation of reasoning shortcuts", NeurIPS, 2023

Bortolotti et al., "A Benchmark Suite for Systematically Evaluating Reasoning Shortcuts", NeurIPS Benchmark track, 2024



learning & reasoning with circuits in pytorch

`github.com/april-tools/circuit`



questions?