



tractable probabilistic modeling with probabilistic circuits

antonio vergari (he/him)

 @tetraduzione

2nd Apr 2025 - Neuro-explicit retreat Saarbruecken

april

april-tools.github.io

april

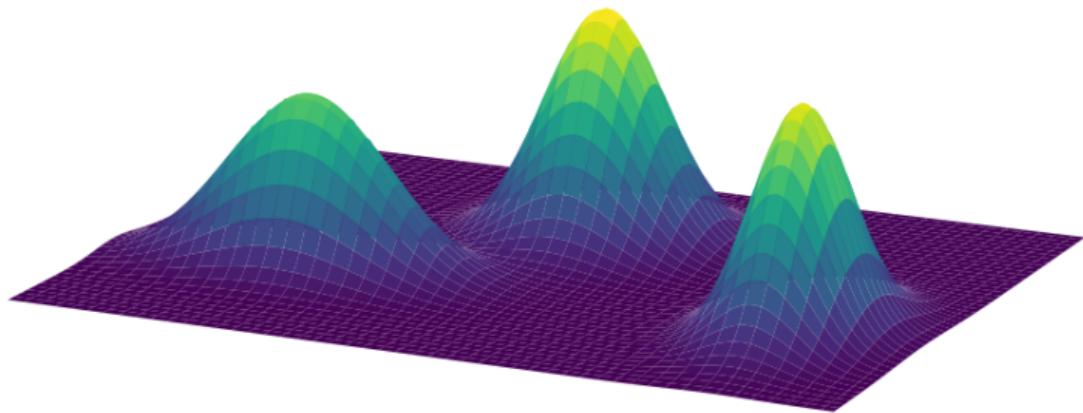
*about
probabilities
integrals &
logic*

deep generative models

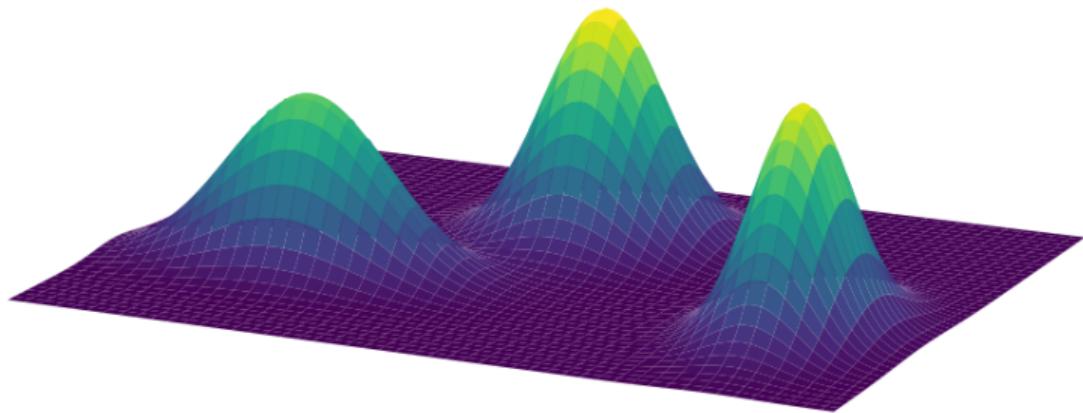
+

*flexible and reliable
(logic &) probabilistic reasoning?*

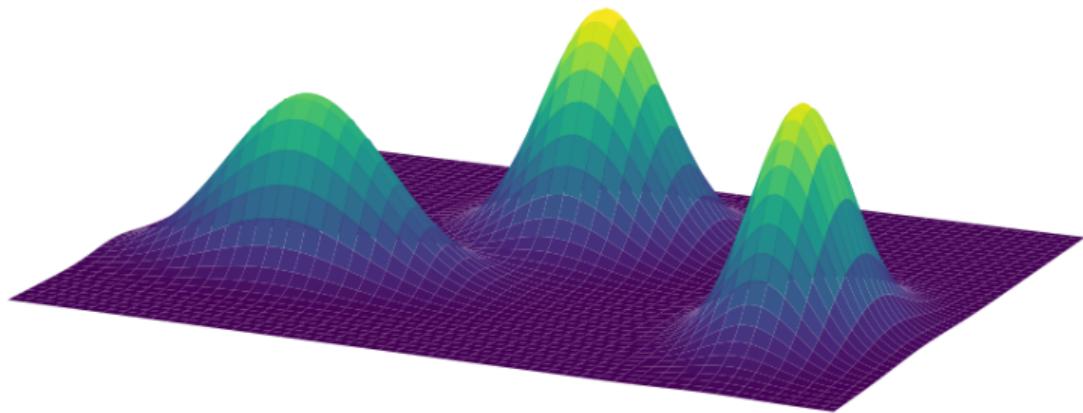
- i) probabilistic circuits: syntax and semantics*
- ii) reliable and efficient neuro-symbolic AI*
- iii) you talk: use cases from your research*



who knows mixture models?



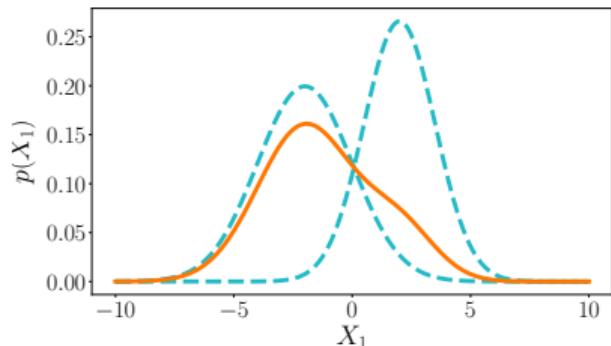
who loves mixture models?



$$c(\mathbf{X}) = \sum_{i=1}^K w_i c_i(\mathbf{X}), \quad \text{with} \quad w_i \geq 0, \quad \sum_{i=1}^K w_i = 1$$

GMMS

as computational graphs

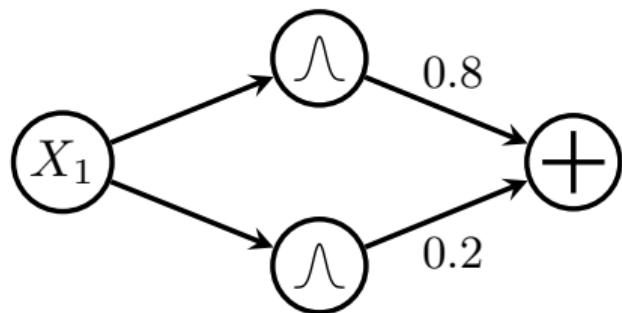


$$p(X) = w_1 \cdot p_1(X_1) + w_2 \cdot p_2(X_1)$$

⇒ translating inference to data structures...

GMMs

as computational graphs

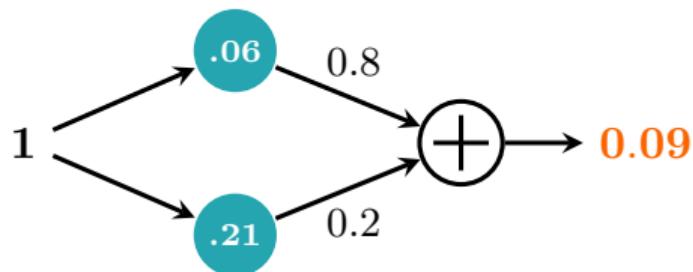


$$p(X_1) = 0.2 \cdot p_1(X_1) + 0.8 \cdot p_2(X_1)$$

⇒ ...e.g., as a weighted sum unit over Gaussian input distributions

GMMS

as computational graphs

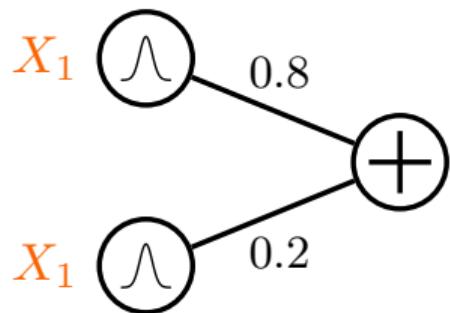


$$p(X = 1) = 0.2 \cdot p_1(X_1 = 1) + 0.8 \cdot p_2(X_1 = 1)$$

⇒ inference = feedforward evaluation

GMMs

as computational graphs

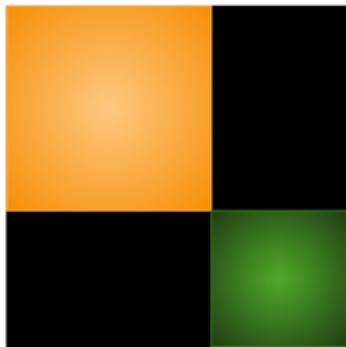
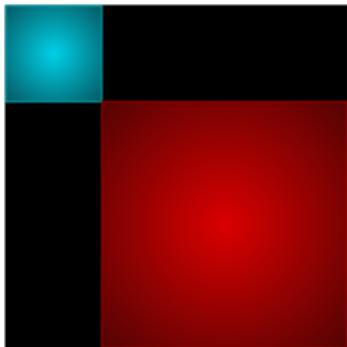


A simplified notation:

- ⇒ **scopes** attached to inputs
- ⇒ edge directions omitted

GMMS

as computational graphs

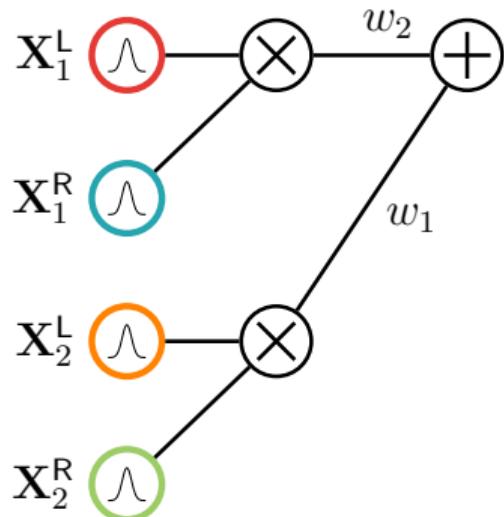


$$p(\mathbf{X}) = w_1 \cdot p_1(\mathbf{X}_1^L) \cdot p_1(\mathbf{X}_1^R) + \\ w_2 \cdot p_2(\mathbf{X}_2^L) \cdot p_2(\mathbf{X}_2^R)$$

⇒ local factorizations...

GMMs

as computational graphs



$$p(\mathbf{X}) = w_1 \cdot p_1(\mathbf{X}_1^L) \cdot p_1(\mathbf{X}_1^R) + \\ w_2 \cdot p_2(\mathbf{X}_2^L) \cdot p_2(\mathbf{X}_2^R)$$

⇒ ...are product units

probabilistic circuits (PCs)

a grammar for tractable computational graphs

I. A simple tractable function is a circuit

⇒ e.g., a multivariate Gaussian or
orthonormal polynomial



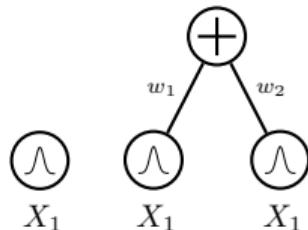
X_1

probabilistic circuits (PCs)

a grammar for tractable computational graphs

I. A simple tractable function is a circuit

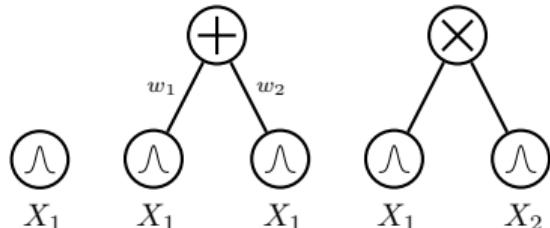
II. A weighted combination of circuits is a circuit



probabilistic circuits (PCs)

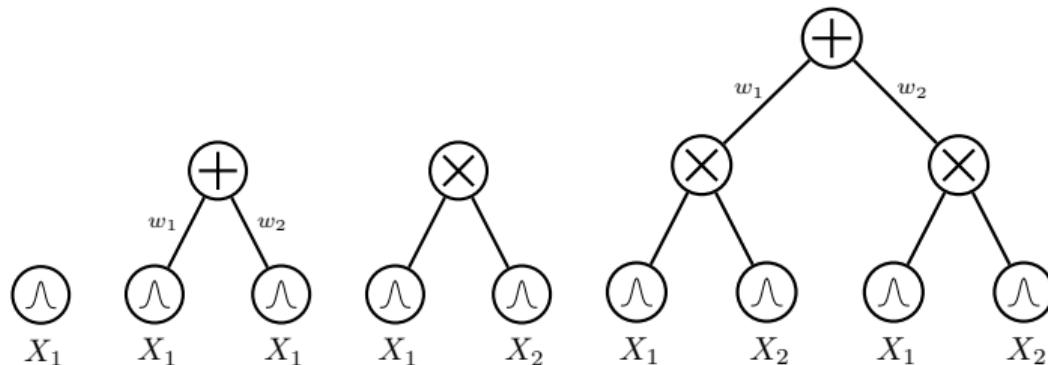
a grammar for tractable computational graphs

- 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



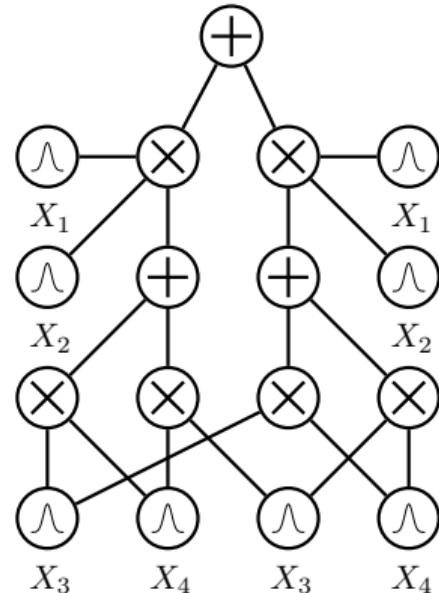
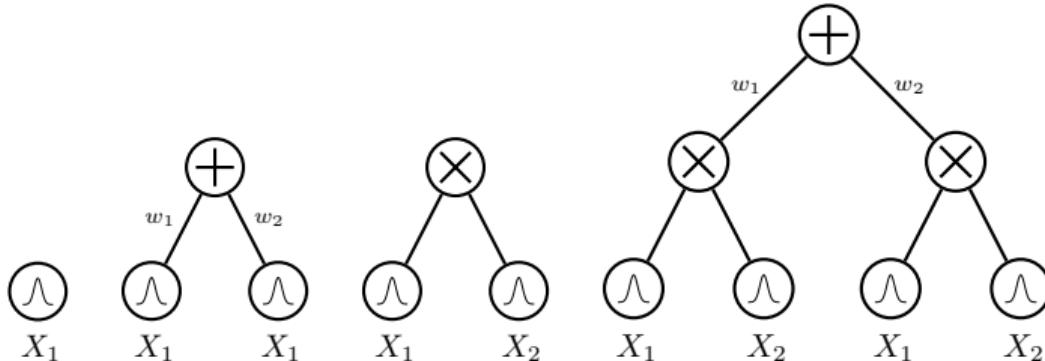
probabilistic circuits (PCs)

a grammar for tractable computational graphs



probabilistic circuits (PCs)

a grammar for tractable computational graphs



probabilistic circuits (PCs)

a tensorized definition

- I. A set of tractable functions is a circuit layer



probabilistic circuits (PCs)

a tensorized definition

- I. A set of tractable functions is a circuit layer
- II. A linear projection of a layer is a circuit layer

$$c(\mathbf{x}) = \mathbf{W}l(\mathbf{x})$$



probabilistic circuits (PCs)

a tensorized definition

- I. A set of tractable functions is a circuit layer
- II. A linear projection of a layer is a circuit layer

$$c(\mathbf{x}) = \mathbf{W}l(\mathbf{x})$$



probabilistic circuits (PCs)

a tensorized definition

- I. A set of tractable functions is a circuit layer
- II. A linear projection of a layer is a circuit layer
- III. The product of two layers is a circuit layer

$$c(\mathbf{x}) = \mathbf{l}(\mathbf{x}) \odot \mathbf{r}(\mathbf{x}) \quad // \text{ Hadamard}$$



probabilistic circuits (PCs)

a tensorized definition

- I. A set of tractable functions is a circuit layer
- II. A linear projection of a layer is a circuit layer
- III. The product of two layers is a circuit layer

$$c(\mathbf{x}) = \mathbf{l}(\mathbf{x}) \odot \mathbf{r}(\mathbf{x}) \quad // \text{ Hadamard}$$

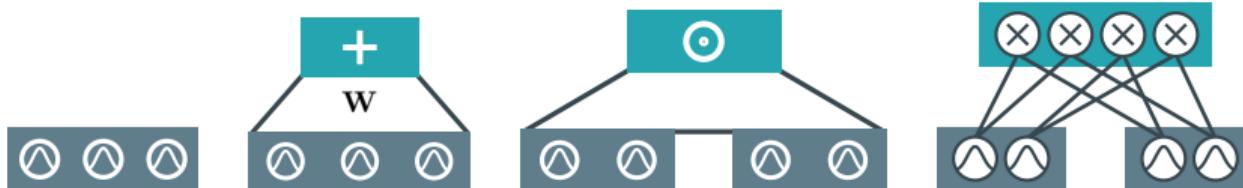


probabilistic circuits (PCs)

a tensorized definition

- I. A set of tractable functions is a circuit layer
- II. A linear projection of a layer is a circuit layer
- III. The product of two layers is a circuit layer

$$c(\mathbf{x}) = \text{vec}(\mathbf{l}(\mathbf{x})\mathbf{r}(\mathbf{x})^\top) \quad // \text{ Kronecker}$$

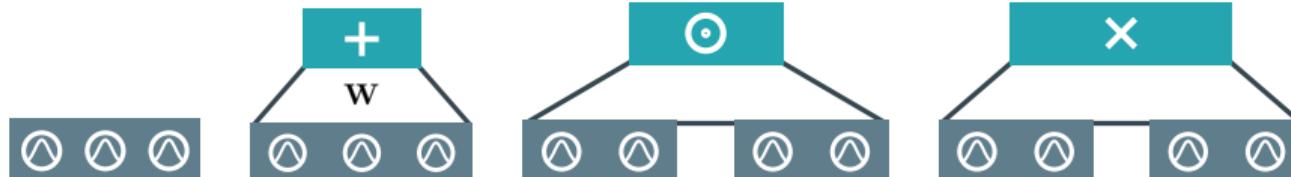


probabilistic circuits (PCs)

a tensorized definition

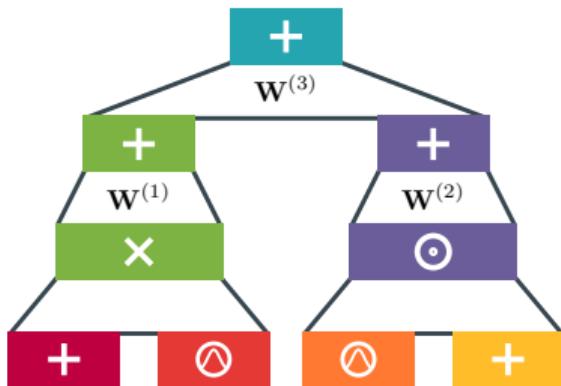
- I. A set of tractable functions is a circuit layer
- II. A linear projection of a layer is a circuit layer
- III. The product of two layers is a circuit layer

$$c(\mathbf{x}) = \text{vec}(\mathbf{l}(\mathbf{x})\mathbf{r}(\mathbf{x})^\top) \quad // \text{ Kronecker}$$



probabilistic circuits (PCs)

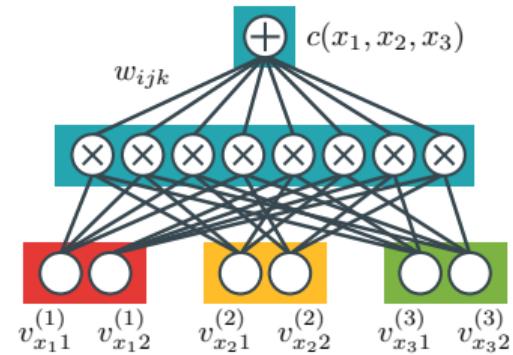
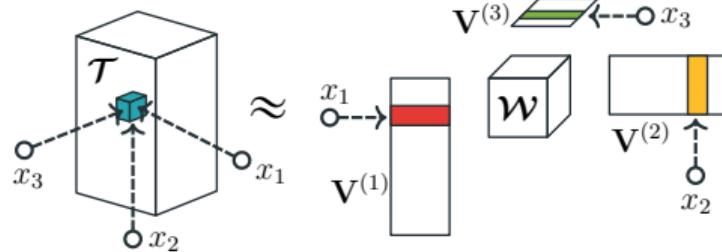
a tensorized definition



- I. A set of tractable functions is a circuit layer
 - II. A linear projection of a layer is a circuit layer
 - III. The product of two layers is a circuit layer
- stack layers to build a deep circuit!**

tensor factorizations

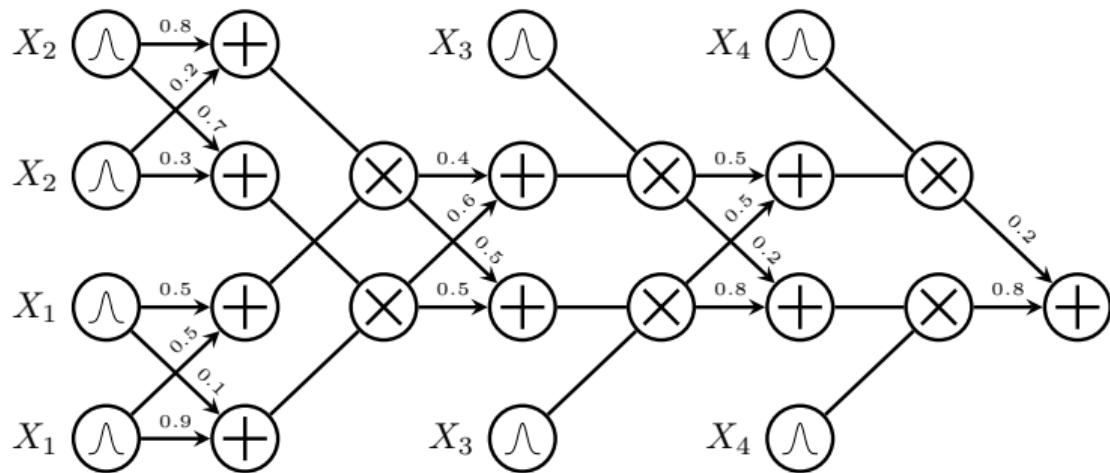
as circuits



Loconte et al., "What is the Relationship between Tensor Factorizations and Circuits (and How Can We Exploit it)?", TMLR, 2025

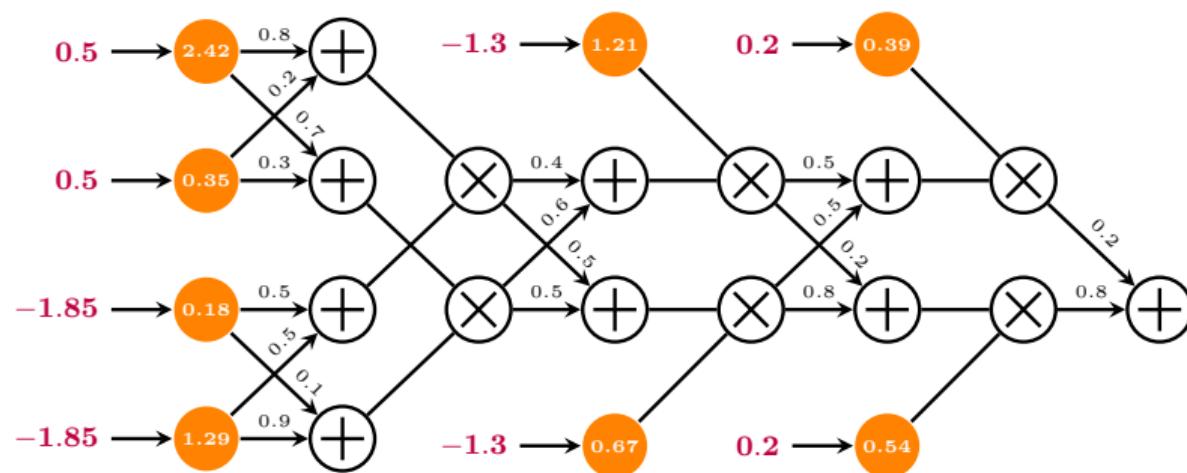
Probabilistic queries = feedforward evaluation

$$p(X_1 = -1.85, X_2 = 0.5, X_3 = -1.3, X_4 = 0.2)$$



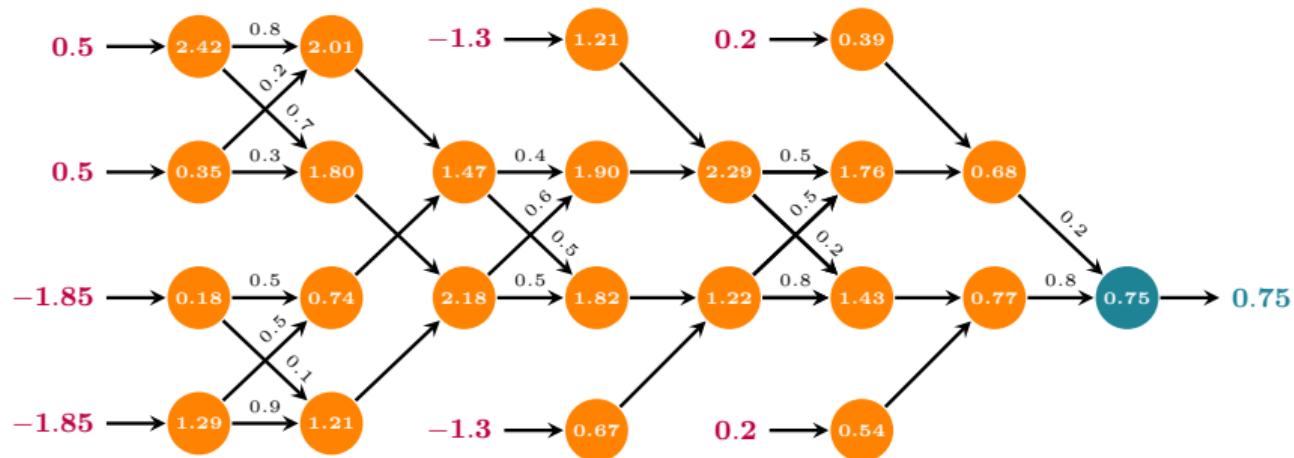
Probabilistic queries = *feedforward* evaluation

$$p(X_1 = \textcolor{red}{-1.85}, X_2 = \textcolor{red}{0.5}, X_3 = \textcolor{red}{-1.3}, X_4 = \textcolor{red}{0.2})$$



Probabilistic queries = feedforward evaluation

$$p(X_1 = -1.85, X_2 = 0.5, X_3 = -1.3, X_4 = 0.2) = 0.75$$





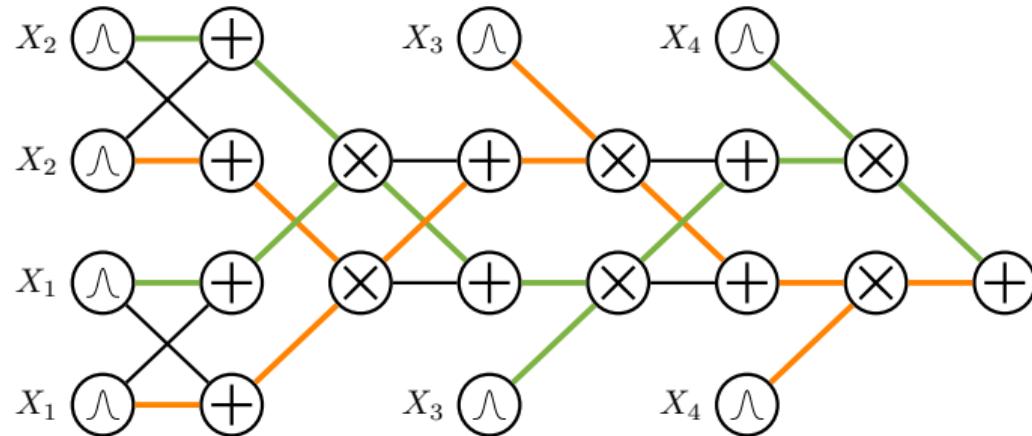
learning & reasoning with circuits in pytorch

github.com/april-tools/cirkit

```
1 from cirkit.templates import circuit_templates  
2  
3 symbolic_circuit = circuit_templates.image_data(  
4     (1, 28, 28),                      # The shape of MNIST  
5     region_graph='quad-graph',  
6     input_layer='categorical',          # input distributions  
7     sum_product_layer='cp',            # CP, Tucker, CP-T  
8     num_input_units=64,                # overparameterizing  
9     num_sum_units=64,  
10    sum_weight_param=circuit_templates.Parameterization(  
11        activation='softmax',  
12        initialization='normal'  
13    )  
14 )
```

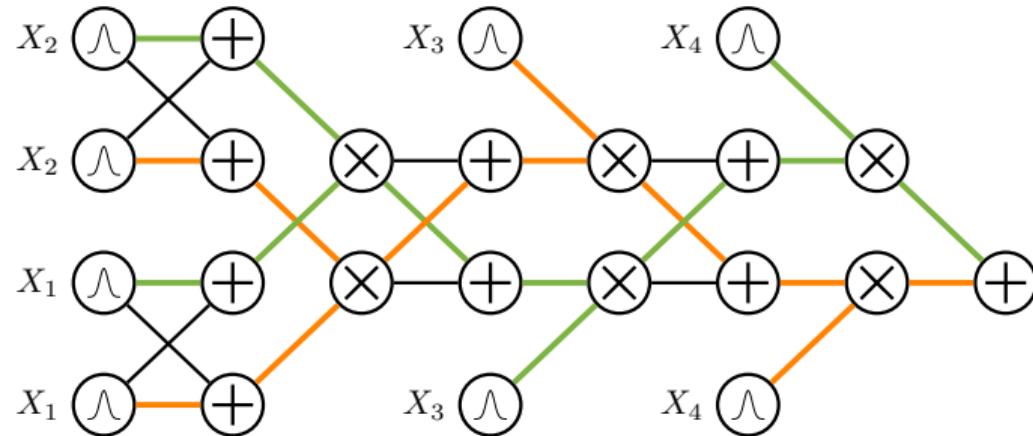
```
1 from cirkit.pipeline import compile
2 circuit = compile(symbolic_circuit)
3
4 with torch.no_grad():
5     test_lls = 0.0
6     for batch, _ in test_dataloader:
7         batch = batch.to(device).unsqueeze(dim=1)
8         log_likelihoods = circuit(batch)
9         test_lls += log_likelihoods.sum().item()
10 average_ll = test_lls / len(data_test)
11 bpd = -average_ll / (28 * 28 * np.log(2.0))
12 print(f"Average LL: {average_ll:.3f}") # Average LL:
13     → -682.916
14 print(f"Bits per dim: {bpds:.3f}") # Bits per dim: 1.257
```

deep mixtures



$$p(\mathbf{x}) = \sum_{\mathcal{T}} \left(\prod_{w_j \in \mathbf{w}_{\mathcal{T}}} w_j \right) \prod_{l \in \text{leaves}(\mathcal{T})} p_l(\mathbf{x})$$

deep mixtures



an exponential number of mixture components!

...why PCs?

1. A grammar for tractable models

One formalism to represent many probabilistic models

⇒ #HMMs #Trees #XGBoost, Tensor Networks, ...

...why PCs?

1. A grammar for tractable models

One formalism to represent many probabilistic models

⇒ #HMMs #Trees #XGBoost, Tensor Networks, ...

2. Tractability == structural properties!!!

Exact computations of reasoning tasks are certified by guaranteeing certain structural properties. #marginals #expectations #MAP, #product ...

structural properties

smoothness

decomposability

compatibility

structural properties

property A

property B

property C

structural properties

smoothness

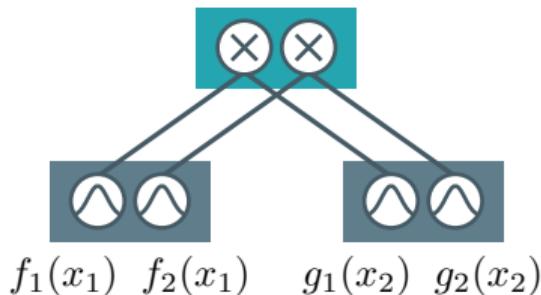
decomposability

property C

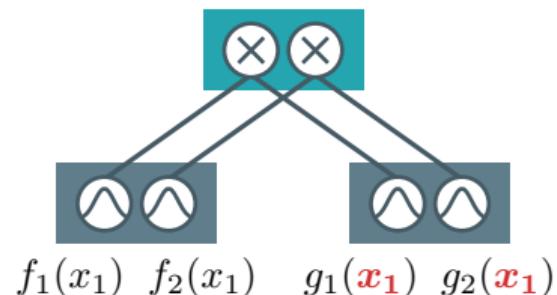
smoothness \wedge decomposability
 \implies multilinearity

Multilinearity in circuits

the inputs of product units are defined over disjoint sets of variables



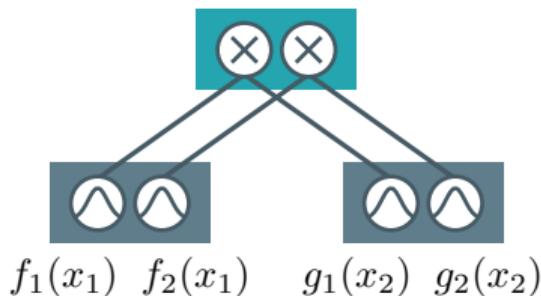
✓ **multilinear**



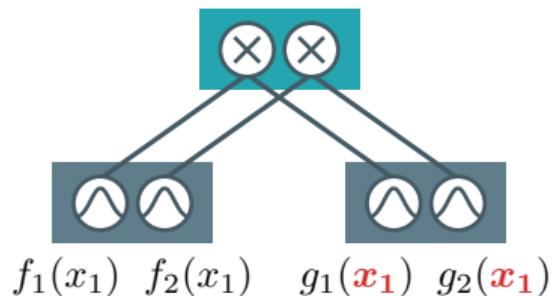
✗ **not multilinear**

Multilinearity in circuits

the inputs of product units are defined over disjoint sets of variables



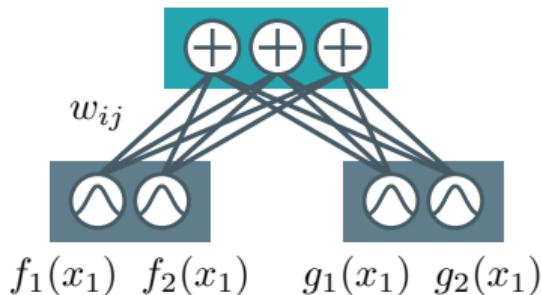
decomposable circuit



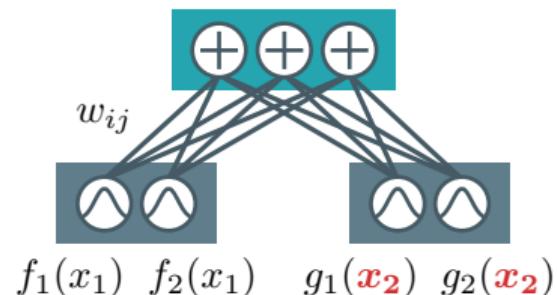
non-decomposable circuit

Multilinearity in circuits

the inputs of sum units are defined over the same variables



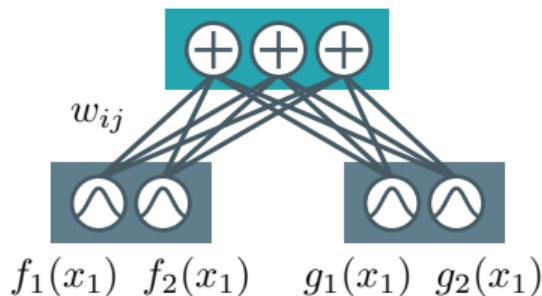
✓ **multilinear**



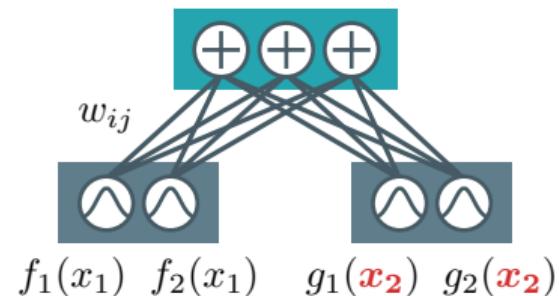
✗ **not multilinear**

Multilinearity in circuits

the inputs of sum units are defined over the same variables



smooth circuit



non-smooth circuit

structural properties

smoothness

decomposability

property C

smoothness \wedge decomposability
 \implies multilinearity

structural properties

smoothness

tractable computation of arbitrary integrals
in probabilistic circuits

decomposability

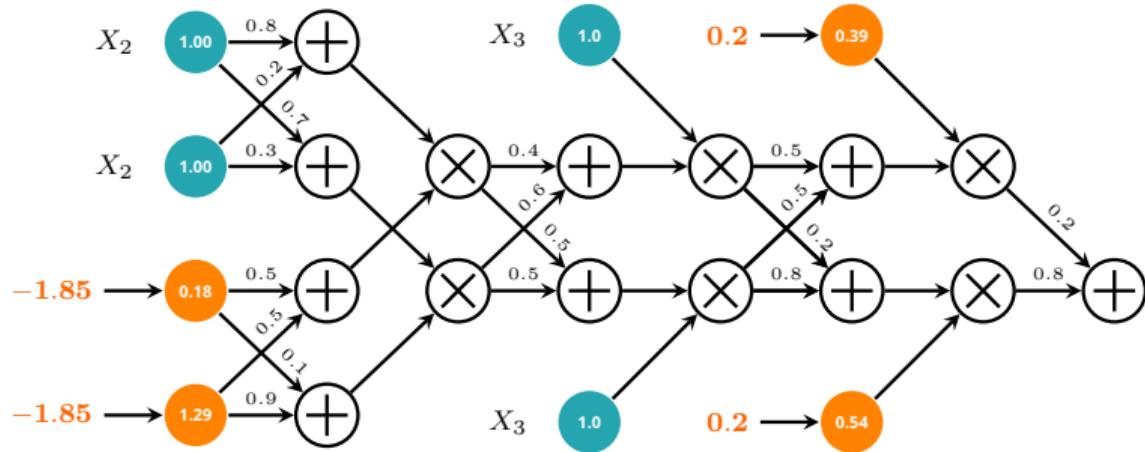
$$p(\mathbf{y}) = \int p(\mathbf{y}, \mathbf{z}) d\mathbf{z}, \quad \forall \mathbf{Y} \subseteq \mathbf{X}, \quad \mathbf{Z} = \mathbf{X} \setminus \mathbf{Y}$$

property C

⇒ tractable partition function
⇒ also any conditional is tractable

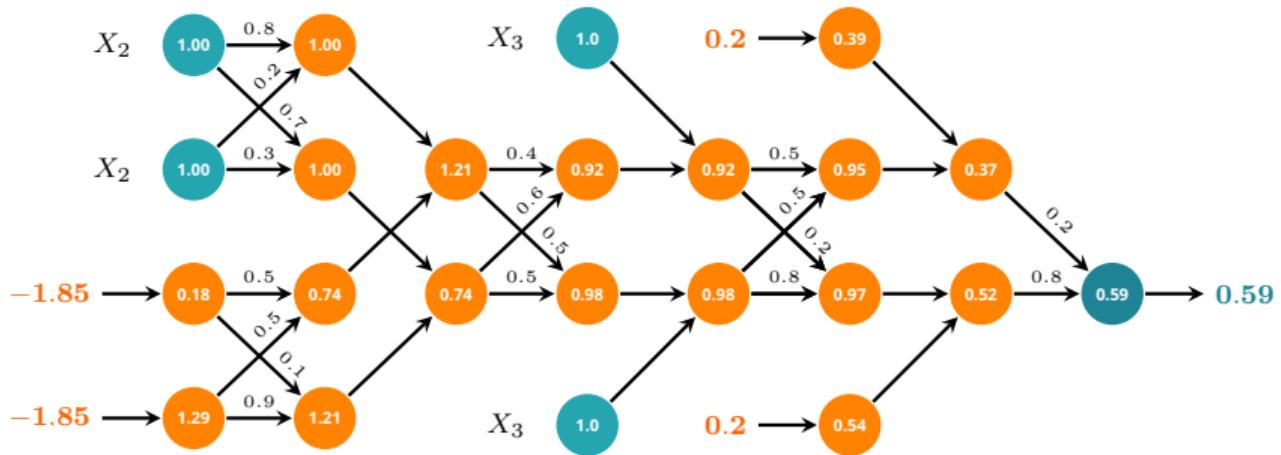
Probabilistic queries = feedforward evaluation

$$p(X_1 = -1.85, X_4 = 0.2)$$



Probabilistic queries = feedforward evaluation

$$p(X_1 = -1.85, X_4 = 0.2)$$



***smooth* + *decomposable* circuits = ...**

Computing arbitrary integrations (or summations)

⇒ *linear in circuit size!*

E.g., suppose we want to compute Z:

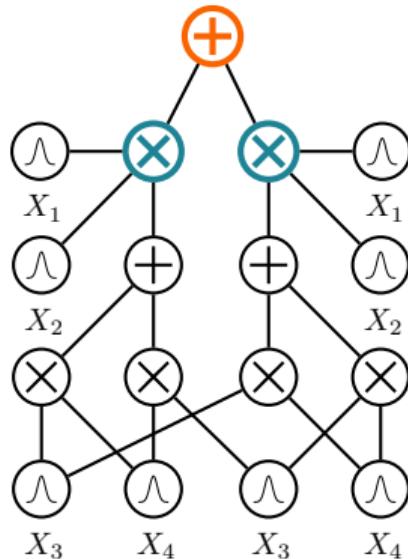
$$\int p(\mathbf{x}) d\mathbf{x}$$

***smooth* + *decomposable* circuits = ...**

If $\mathbf{p}(\mathbf{x}) = \sum_i w_i \mathbf{p}_i(\mathbf{x})$, (*smoothness*):

$$\begin{aligned}\int \mathbf{p}(\mathbf{x}) d\mathbf{x} &= \int \sum_i w_i \mathbf{p}_i(\mathbf{x}) d\mathbf{x} = \\ &= \sum_i w_i \int \mathbf{p}_i(\mathbf{x}) d\mathbf{x}\end{aligned}$$

\Rightarrow integrals are “pushed down” to inputs

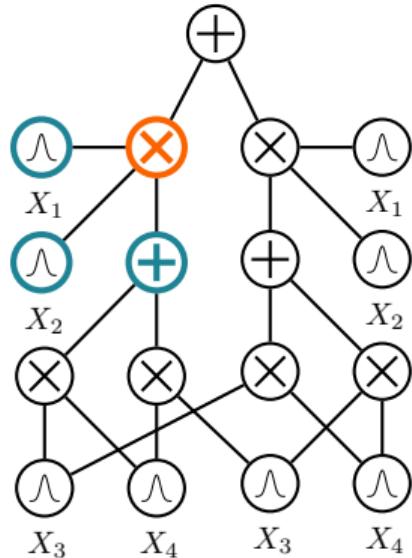


smooth + decomposable circuits = ...

If $\mathbf{p}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \mathbf{p}(\mathbf{x})\mathbf{p}(\mathbf{y})\mathbf{p}(\mathbf{z})$, (*decomposability*):

$$\begin{aligned}& \int \int \int \mathbf{p}(\mathbf{x}, \mathbf{y}, \mathbf{z}) d\mathbf{x}d\mathbf{y}d\mathbf{z} = \\&= \int \int \int \mathbf{p}(\mathbf{x})\mathbf{p}(\mathbf{y})\mathbf{p}(\mathbf{z}) d\mathbf{x}d\mathbf{y}d\mathbf{z} = \\&= \int \mathbf{p}(\mathbf{x}) d\mathbf{x} \int \mathbf{p}(\mathbf{y}) d\mathbf{y} \int \mathbf{p}(\mathbf{z}) d\mathbf{z}\end{aligned}$$

⇒ integrals decompose into easier ones



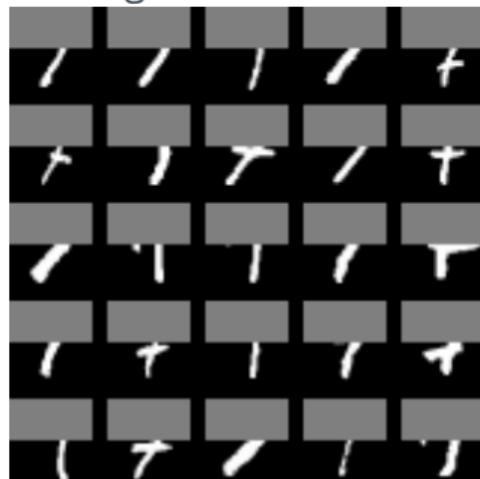
```
1 from cirkit.backend.torch.queries import IntegrateQuery
2 marginal_query = IntegrateQuery(circuit)
3
4 with torch.no_grad():
5     test_marginal_lls = 0.0
6
7     for batch, _ in test_dataloader:
8         batch = batch.to(device).unsqueeze(dim=1)
9         marginal_log_likelihoods = marginal_query(batch,
10             ↪ integrate_vars=vars_to_marginalize)
11         test_marginal_lls +=
12             ↪ marginal_log_likelihoods.sum().item()
13
14     marg_ll = test_marginal_lls / len(data_test)
15     print(f"marg LL: {marg_ll:.3f}") # marg LL:: -378.417
```

tractable marginals on PCs

Original

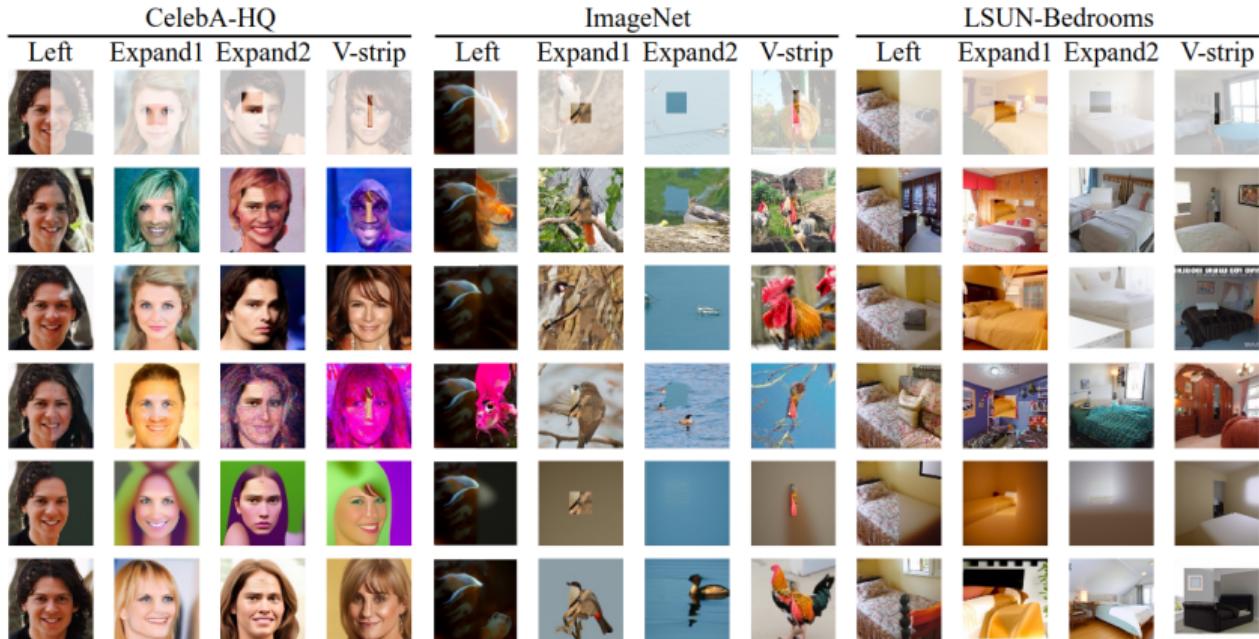


Missing



Conditional sample





structural properties

smoothness

Integrals involving two or more functions:
e.g., expectations

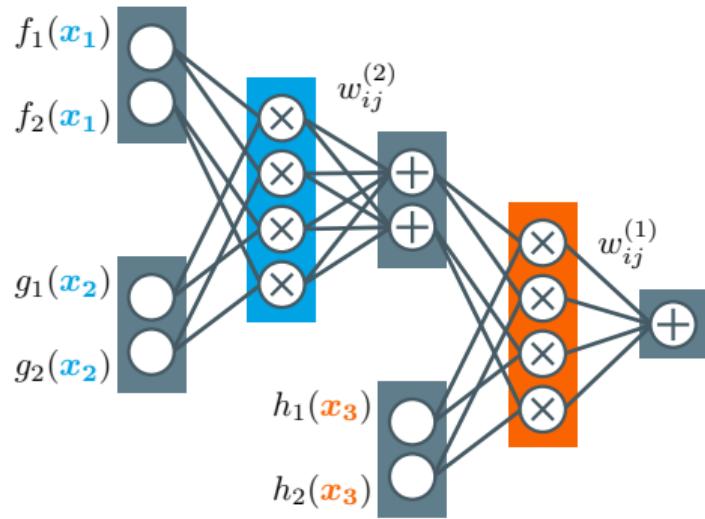
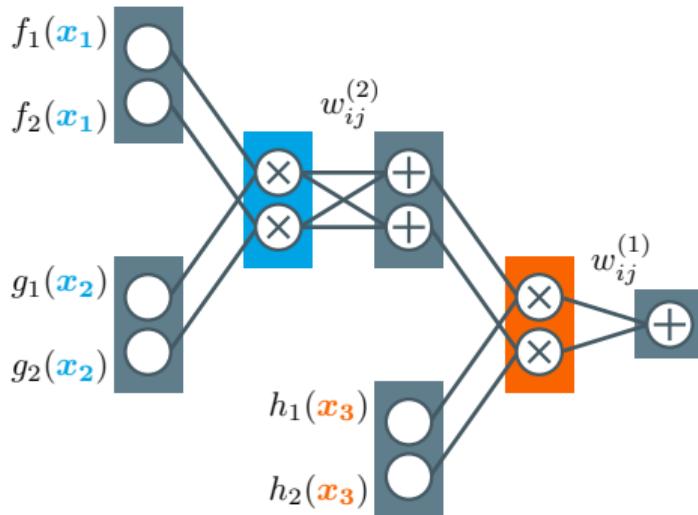
decomposability

$$\mathbb{E}_{\mathbf{x} \sim p} [f(\mathbf{x})] = \int p(\mathbf{x}) f(\mathbf{x}) d\mathbf{x}$$

compatibility

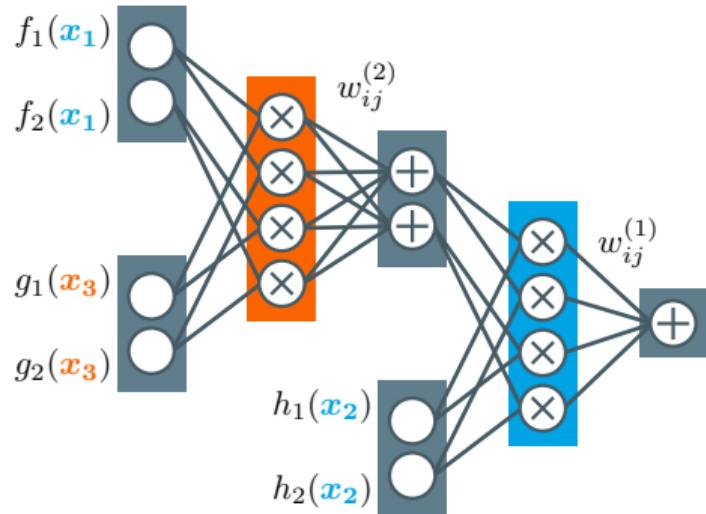
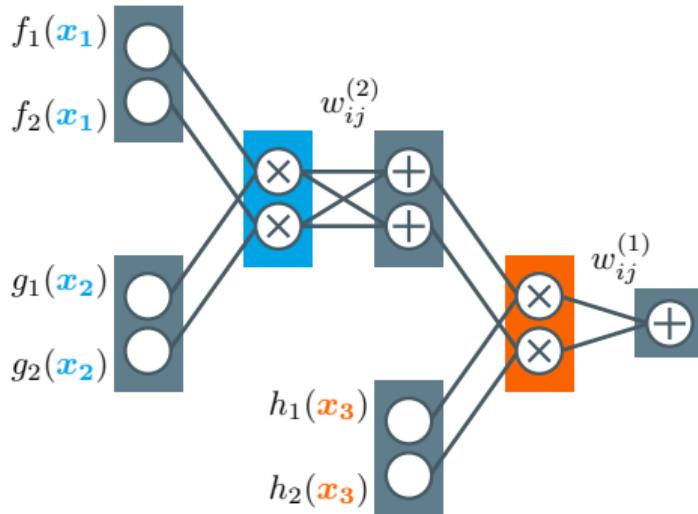
when both $p(\mathbf{x})$ and $f(\mathbf{x})$ are circuits

compatibility



compatible circuits

compatibility



non-compatible circuits

structural properties

smoothness

compatibility

decomposability

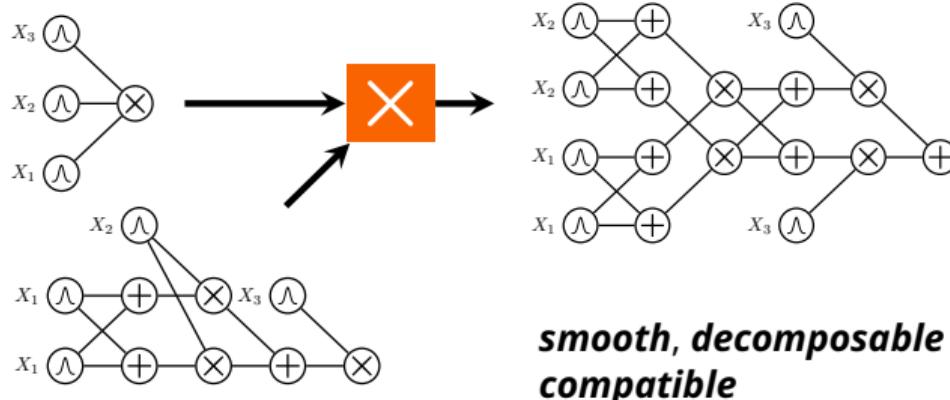


smoothness \wedge decomposability

compatibility

compatibility \Rightarrow tractable expectations

Tractable products



compute $\mathbb{E}_{\mathbf{x} \sim p} f(\mathbf{x}) = \int p(\mathbf{x}) f(\mathbf{x}) \mathrm{d}\mathbf{x}$ in $O(|p| |f|)$

```
1 from cirkit.symbolic.circuit import Circuit
2 from cirkit.symbolic.functional import (
3     integrate, multiply)
4
5 # Circuits expectation  $\int [p(x) f(x)] dx$ 
6 def expectation(p: Circuit, f: Circuit) -> Circuit:
7     i = multiply(p, f)
8     return integrate(i)
9
10 # Squared loss  $\int [p(x)-q(x)]^2 dx = E_p[p] + E_q[q] - 2E_p[q]$ 
11 #           =  $\int p^2(x) dx + \int q^2(x) dx - 2\int p(x)q(x) dx$ 
12 def squared_loss(p: Circuit, q: Circuit) -> Circuit:
13     p2 = multiply(p, p)
14     q2 = multiply(q, q)
15     pq = multiply(p, q)
16     return integrate(p2) + integrate(q2) - 2 * integrate(pq)
```

...why PCs?

1. A grammar for tractable models

One formalism to represent many probabilistic models

⇒ #HMMs #Trees #XGBoost, Tensor Networks, ...

2. Tractability == structural properties!!!

Exact computations of reasoning tasks are certified by guaranteeing certain structural properties. #marginals #expectations #MAP, #product ...

...why PCs?

1. A grammar for tractable models

One formalism to represent many probabilistic models

→ #HMMs #Trees #XGBoost, Tensor Networks, ...

2. Tractability == structural properties!!!

Exact computations of reasoning tasks are certified by guaranteeing certain structural properties. #marginals #expectations #MAP, #product ...

3. Reliable neuro-symbolic AI

logical constraints as circuits, multiplied to probabilistic circuits

```
1 from cirkit.templates import circuit_templates  
2  
3 symbolic_circuit = circuit_templates.image_data(  
4     (1, 28, 28),                      # The shape of MNIST  
5     region_graph='quad-graph',  
6     input_layer='categorical',          # input distributions  
7     sum_product_layer='cp',            # CP, Tucker, CP-T  
8     num_input_units=64,                # overparameterizing  
9     num_sum_units=64,  
10    sum_weight_param=circuit_templates.Parameterization(  
11        activation='softmax',  
12        initialization='normal'  
13    )  
14 )
```

learning probabilistic circuits

learning probabilistic circuits

Probabilistic circuits are (peculiar) neural networks...*just backprop with SGD!*

learning probabilistic circuits

Probabilistic circuits are (peculiar) neural networks...*just backprop with SGD!*

...end of Learning section!

learning probabilistic circuits

Probabilistic circuits are (peculiar) neural networks...*just backprop with SGD!*

wait but...

which loss?

how to learn normalized weights?

how to exploit structural properties?

maximum likelihood

the go-to objective in ProbML

Given a dataset $\mathcal{D} = \{\mathbf{x}^{(i)}\}_{i=1}^N$ and your parametric model $p_\theta(\mathbf{X})$ solve

$$\hat{\theta}_{\text{ML}} = \max_{\theta} \prod_{i=1}^N p_\theta(\mathbf{x}^{(i)}) = \min_{\theta} - \sum_{i=1}^N \log p_\theta(\mathbf{x}^{(i)})$$

\Rightarrow minimize the negative log-likelihood (NLL)

```
1 from cirkit.templates import circuit_templates  
2  
3 symbolic_circuit = circuit_templates.image_data(  
4     (1, 28, 28),                      # The shape of MNIST  
5     region_graph='quad-graph',  
6     input_layer='categorical',          # input distributions  
7     sum_product_layer='cp',            # CP, Tucker, CP-T  
8     num_input_units=64,                # overparameterizing  
9     num_sum_units=64,  
10    sum_weight_param=circuit_templates.Parameterization(  
11        activation='softmax',  
12        initialization='normal'  
13    )  
14 )
```

which parameters?

how to reparameterize circuits

Input distributions.

Sum unit parameters.

which parameters?

how to reparameterize circuits

Input distributions. Each input can be a different parametric distribution

⇒ *Bernoullis, Categoricals, Gaussians, exponential families, small NNs, ...*

Sum unit parameters.

which parameters?

how to reparameterize circuits

Input distributions. Each input can be a different parametric distribution

Sum unit parameters. Enforce them to be non-negative, i.e., $w_i \geq 0$ but unnormalized

$$w_i = \exp(\alpha_i), \quad \alpha_i \in \mathbb{R}, \quad i = 1, \dots, K$$

and renormalize the loss

$$\min_{\theta} - \left(\sum_{i=1}^N \log \tilde{p}_{\theta}(\mathbf{x}^{(i)}) - \log \int \tilde{p}_{\theta}(\mathbf{x}^{(i)}) d\mathbf{X} \right)$$

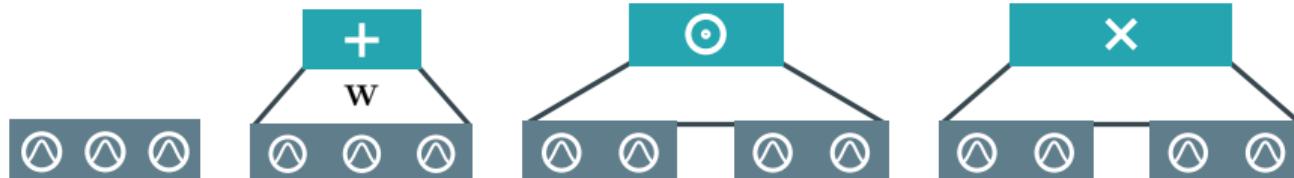
or just renormalize the weights, i.e., $\sum_i w_i = 1$

$$\mathbf{w} = \text{softmax}(\boldsymbol{\alpha}), \quad \boldsymbol{\alpha} \in \mathbb{R}^K$$

```
1 from cirkit.templates import circuit_templates  
2  
3 symbolic_circuit = circuit_templates.image_data(  
4     (1, 28, 28),                      # The shape of MNIST  
5     region_graph='quad-graph',  
6     input_layer='categorical',          # input distributions  
7     sum_product_layer='cp',            # CP, Tucker, CP-T  
8     num_input_units=64,                # overparameterizing  
9     num_sum_units=64,  
10    sum_weight_param=circuit_templates.Parameterization(  
11        activation='softmax',  
12        initialization='normal'  
13    )  
14 )
```

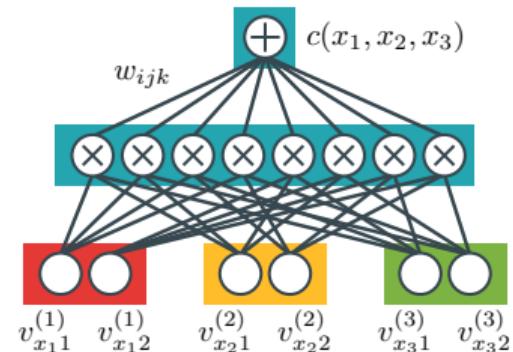
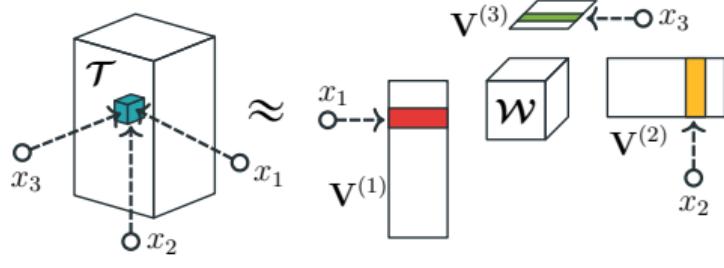
Probabilistic Circuits (PCs)

the layer-wise definition



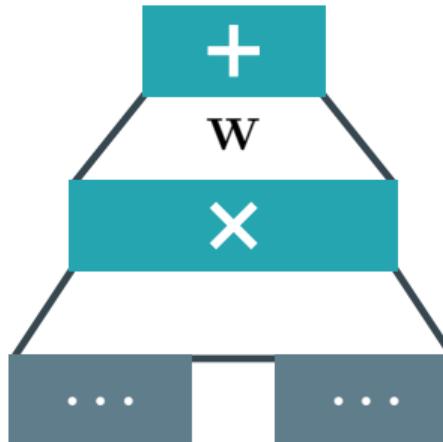
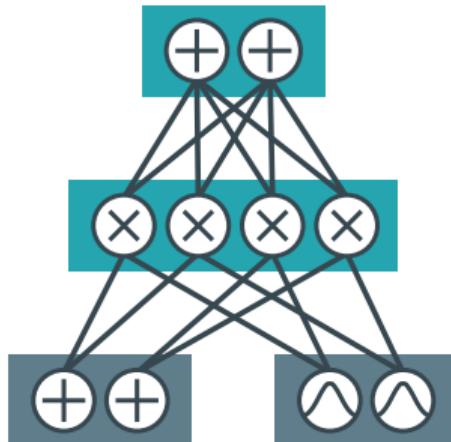
circuits layers

as tensor factorizations



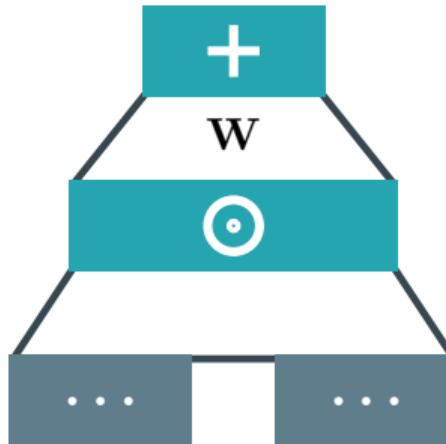
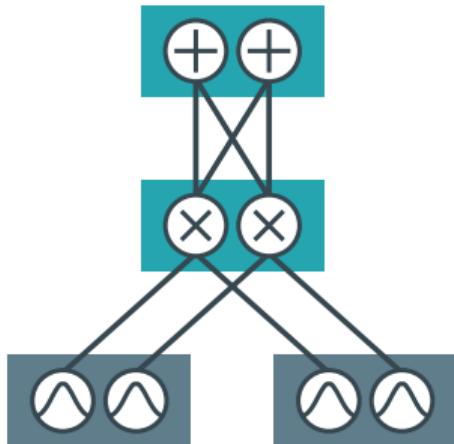
more layers

Tucker decomposition

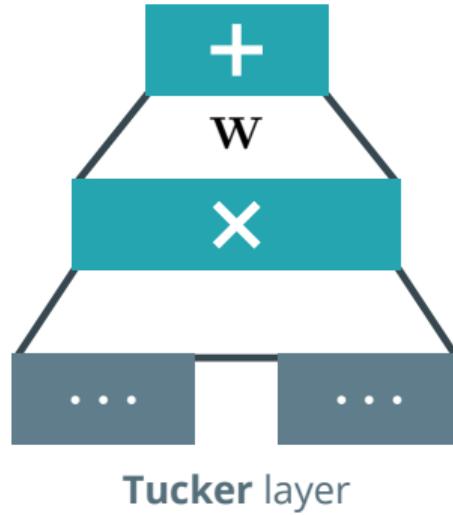
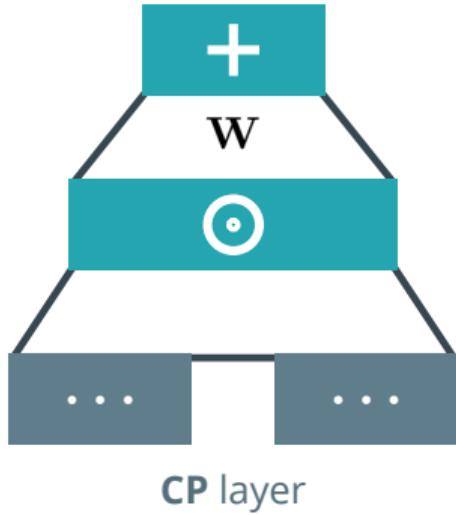


more layers

Candecomp Parafac (CP) decomposition



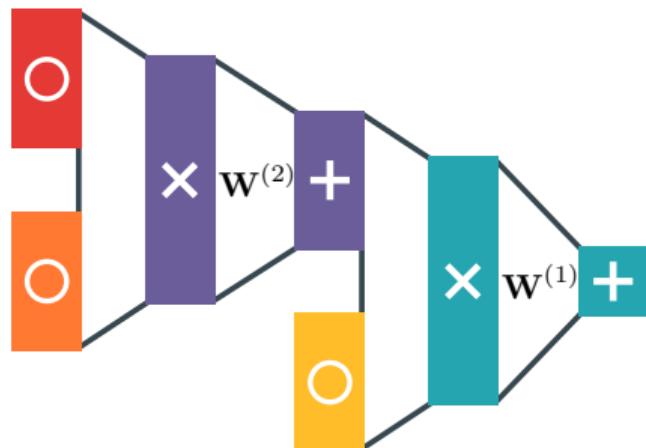
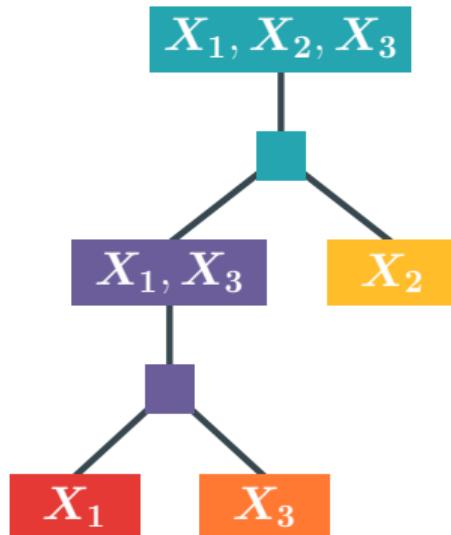
more layers



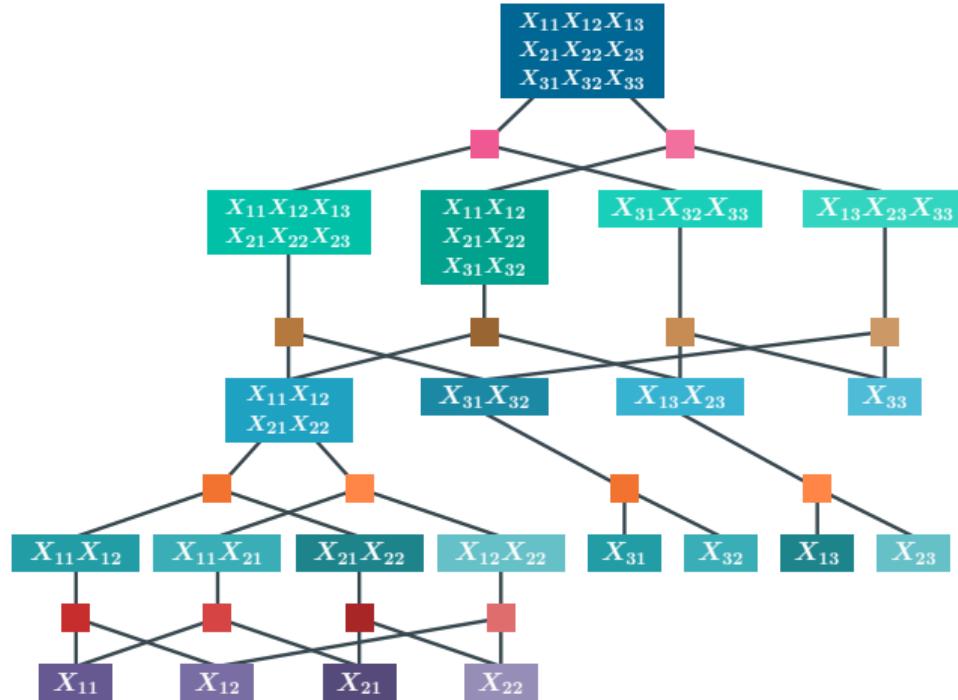
```
1 from cirkit.templates import circuit_templates  
2  
3 symbolic_circuit = circuit_templates.image_data(  
4     (1, 28, 28),  
5     region_graph='quad-graph',  
6     input_layer='categorical',    # input distributions  
7     sum_product_layer='cp',      # CP, Tucker, CP-T  
8     num_input_units=64,          # overparameterizing  
9     num_sum_units=64,  
10    sum_weight_param=circuit_templates.Parameterization(  
11        activation='softmax',  
12        initialization='normal'  
13    )  
14 )
```

region graphs

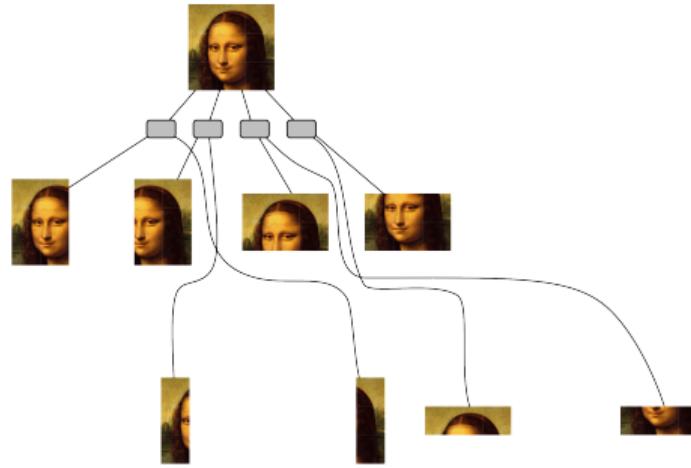
a template for smooth&decomposable PCs

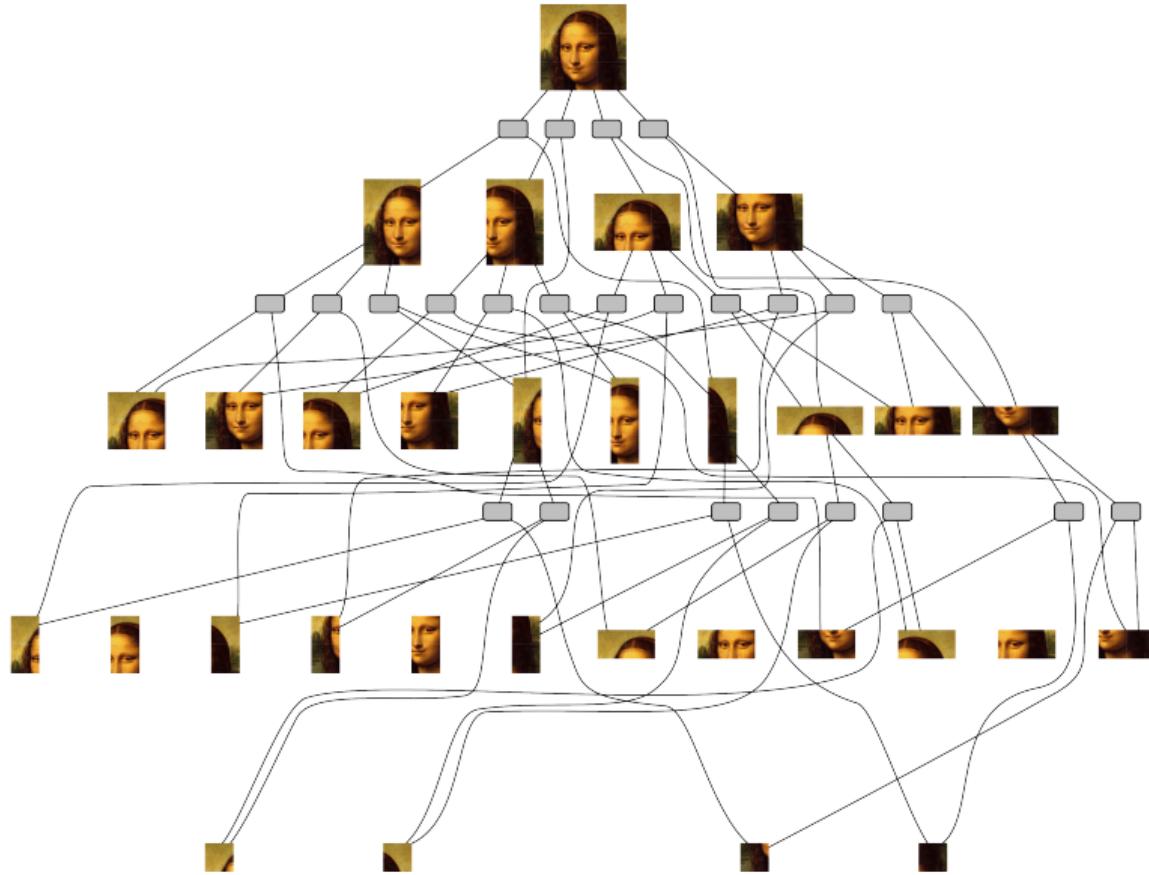


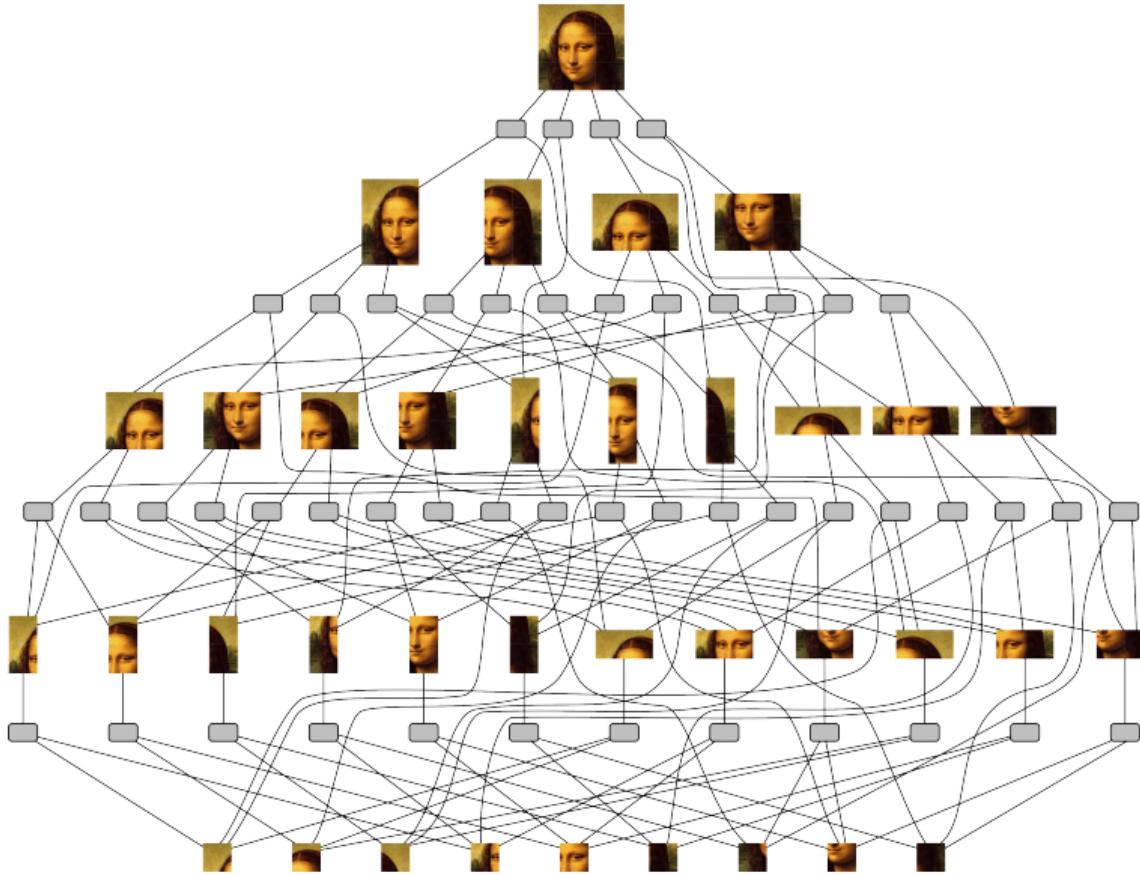
which region graph?







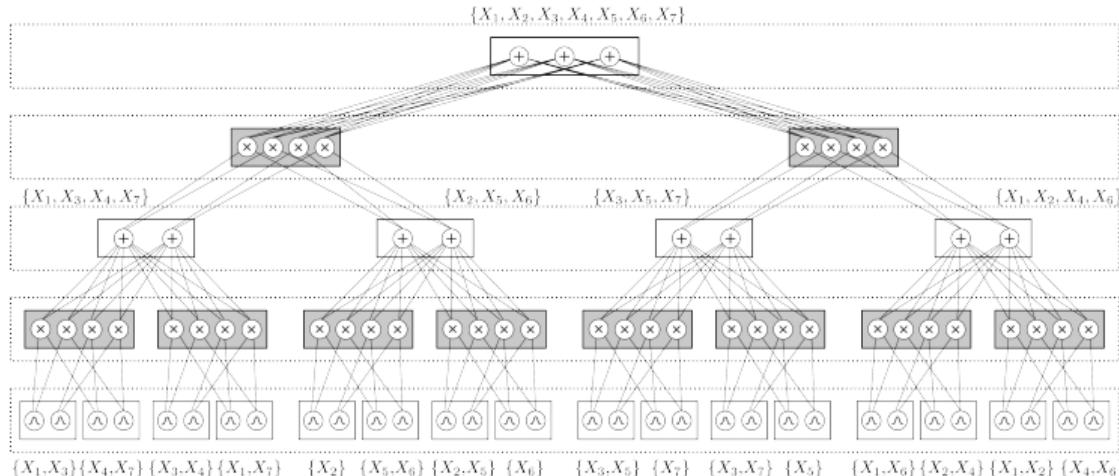




random regions graphs

The “no-learning” option

Generating a random region graph, by recursively splitting \mathbf{X} into two random parts:



The screenshot shows a Jupyter Notebook interface. At the top, there's a header bar with a file icon, a dropdown menu labeled "main", and the path "cirkit / notebooks / region-graphs-and-parametrisation.ipynb". To the right of the path is a search bar with the placeholder "Go to file" and a refresh icon. Below the header, a commit message from "loreloc" is displayed: "updated notebooks with respect to API changes" with a timestamp "e3e7e80 · 2 days ago" and a clock icon. Underneath the commit message, there are buttons for "Preview", "Code", and "Blame", followed by the statistics "1082 lines (1082 loc) · 793 KB". On the far right of the header, there are buttons for "Raw", "Copy", and "Download".

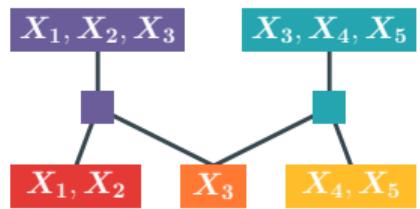
Notebook on Region Graphs and Sum Product Layers

Goals

By the end of this tutorial you will:

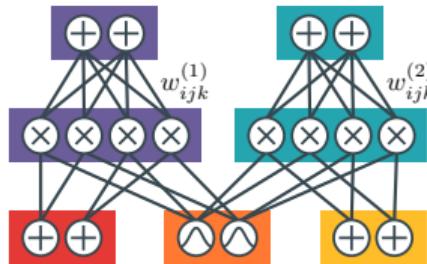
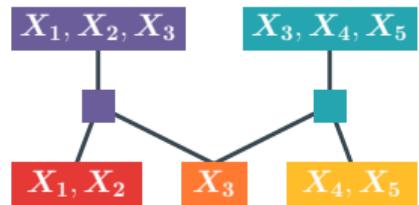
- know what a region graph is
- know how to choose between region graphs for your circuit
- understand how to parametrize a circuit by choosing a sum product layer
- build circuits to tractably estimate a probability distribution over images¹

learning recipe



1) Build a *region graph*

learning recipe

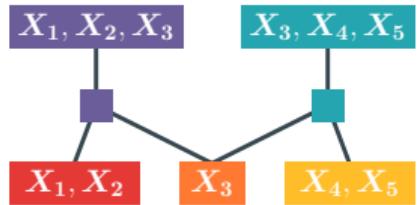


1) Build a *region graph*

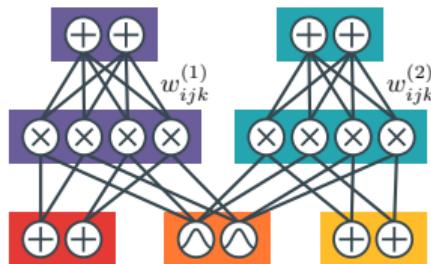
2) Overparameterize

- 2.1) pick a (composite) layer type**
- 2.2) choose how many units per layer**

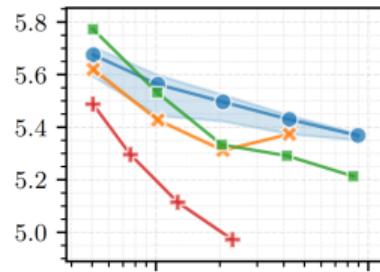
learning recipe



1) Build a *region graph*



2) Overparameterize



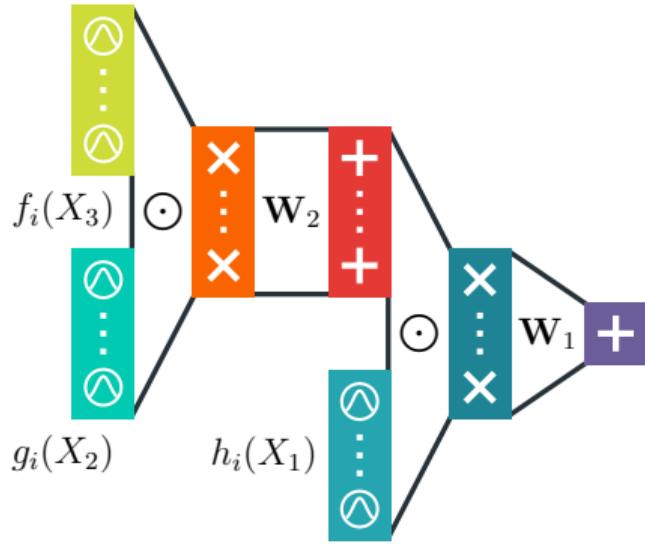
3) Learn parameters

use any optimizer in pytorch



learning & reasoning with circuits in pytorch

<https://github.com/april-tools/cirkit>



questions?