



***subtractive mixture models***

***representation, learning & inference***

**antonio vergari** (he/him)



@nolovedeep learning

9th Sept 2025 - **Bristol AI Summer School**

***thanks to...***




Lorenzo Loconte  
*U of Edinburgh*



Lena Zellinger  
*U of Edinburgh*



Aleksanteri Sladek  
*Aalto U*



Gennaro Gala  
*TU Eindhoven*



Adrian Javaloy  
*U of Edinburgh*

***and moar...***

# *april*

`april-tools.github.io`

# *april*

***autonomous &  
provably  
reliable  
intelligent  
learners***

# *april*


*about*  
*probabilities*  
*integrals &*  
*logic*

# *april*

*april is  
probably a  
recursive  
identifier of a  
lab*

***today's topic...***


# *swiss-army knife of prob ML*




(hierarchical)  
mixtures




# *generalizing them as computational graphs*




# *a single formalism for many models*





***who knows mixture models?***




*who **loves** mixture models?*

## Hierarchical Gaussian Mixture Model Splatting for Efficient and Part Controllable 3D Generation

Qitong Yang, Mingtao Feng, Zijie Wu, Weisheng Dong, Fangfang Wu, Yaonan Wang, Ajmal Mian;  
Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR). 2025, pp.  
11104-11114

## Inversion of nitrogen and phosphorus contents in cotton leaves based on the Gaussian mixture model and differences in hyperspectral features of UAV

Lei Peng , Hui-Nan Xin , Cai-Xia Lv , Na Li , Yong-Fu Li , Qing-Long Geng ,  
Shu-Huang Chen , Ning Lai 


---

## Gaussian Mixture Flow Matching Models


---

Hansheng Chen<sup>1</sup> Kai Zhang<sup>2</sup> Hao Tan<sup>2</sup> Zexiang Xu<sup>3</sup> Fujun Luan<sup>2</sup>  
Leonidas Guibas<sup>1</sup> Gordon Wetzstein<sup>1</sup> Sai Bi<sup>2</sup>

***mixture models are everywhere***  
(still in 2025)



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



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

$$\int \sum_i w_i p_i(\mathbf{x}) d\mathbf{x} = \sum_i w_i \int p_i(\mathbf{x}) d\mathbf{x}$$

***mixture models can enable tractable inference***  
(if components are tractable, e.g., for marginals)



---

# Hierarchical Compositional Mixtures of Variational Autoencoders

---

Ping Liang Tan<sup>1,2</sup> Robert Peharz<sup>1</sup>

---

Mixtures of Laplace Approximations  
for Improved *Post-Hoc* Uncertainty in Deep Learning

---

Runa Eschenhagen<sup>\*1</sup> Erik Daxberger<sup>c,m</sup> Philipp Hennig<sup>l,m</sup> Agustinus Kristiadi<sup>‡</sup>

---


Efficient Mixture Learning in Black-Box Variational Inference

---

Alexandra Hotti<sup>\*1,2,3</sup> Oskar Kviman<sup>\*1,2</sup> Ricky Molén<sup>1,2</sup> Víctor Elvira<sup>4</sup> Jens Lagergren<sup>1,2</sup>


***mixture models can enable tractable inference***  
**(even in larger approximate inference pipelines)**

# *compile mixtures into circuits...*



# GMMs

*as computational graphs*




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

$\Rightarrow$  *translating inference to data structures...*

# GMMs

*as computational graphs*




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

$\Rightarrow$  ...e.g., as a weighted sum unit over Gaussian input distributions

# GMMs

as computational graphs




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

$\Rightarrow$  inference = feedforward evaluation

# GMMs

*as computational graphs*



A simplified notation:

$\Rightarrow$  **scopes** attached to inputs  
 $\Rightarrow$  edge directions omitted

***wait...!***

***how do we **learn** them?***

***wait...!***

***how do we **learn** them?***

***⇒ by maximizing the (log-)likelihood***



# ***which parameters?***

*how to reparameterize mixtures/circuits*

**Input distributions.**

**Sum unit parameters.**

# ***which parameters?***

*how to reparameterize mixtures/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 mixtures/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 **negative log likelihood** 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.,  $\sum_i w_i = 1$

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

***wait...!***

***how do we **learn** them?***

***⇒ by maximizing the (log-)likelihood***

***wait...!***

***how do we **learn** them?***

***⇒ by maximizing the (log-)likelihood***

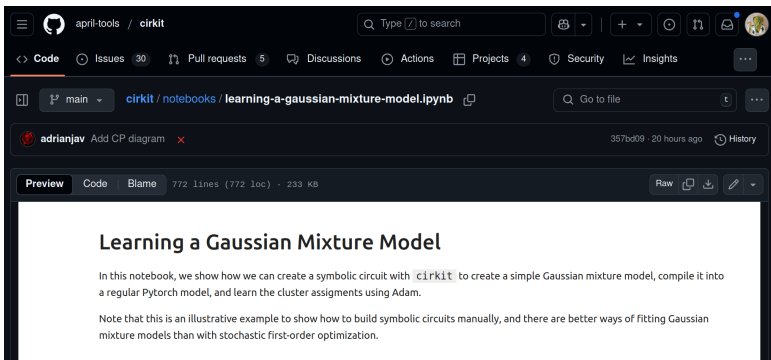
***just SGD your way as usual!***

***⇒ or any other gradient-based optimizer***



***learning & reasoning with circuits in pytorch***

`github.com/april-tools/circuit`



The screenshot shows a GitHub repository for 'april-tools / cirkit'. The main content is a Jupyter notebook titled 'learning-a-gaussian-mixture-model.ipynb'. The notebook's text is as follows:

### Learning a Gaussian Mixture Model


In this notebook, we show how we can create a symbolic circuit with `cirkit` to create a simple Gaussian mixture model, compile it into a regular Pytorch model, and learn the cluster assignments using Adam.

Note that this is an illustrative example to show how to build symbolic circuits manually, and there are better ways of fitting Gaussian mixture models than with stochastic first-order optimization.



## *a notebook on learning GMMs as circuits*

`https://github.com/april-tools/cirkit/blob/main/notebooks/learning-a-gaussian-mixture-model.ipynb`




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



# ***additive MMs***

*are so cool!*



easily represented as shallow PCs


these are *monotonic* PCs

if marginals/conditionals are tractable for the components, they are tractable for the MM

they are *universal approximators*...

# *additive MMs*

*are so cool!*



easily represented as shallow PCs


these are *monotonic* PCs

if marginals/conditionals are tractable for the components, they are tractable for the MM

they are *universal approximators*...

# ***additive MMs***

*are so cool!*



easily represented as shallow PCs


these are **monotonic** PCs

if marginals/conditionals are tractable for the components, they are tractable for the MM

they are **universal approximators**...

# ***additive MMs***

*are so cool!*




easily represented as shallow PCs

these are **monotonic** PCs


if marginals/conditionals are tractable for the components, they are tractable for the MM

they are **universal approximators**...

***however...***





*however...*




GMM ( $K = 2$ )

*however...*





GMM ( $K = 2$ )




GMM ( $K = 16$ )


*however...*



GMM ( $K = 2$ )



GMM ( $K = 16$ )




nGMM<sup>2</sup> ( $K = 2$ )



***spoiler***

**shallow mixtures  
with negative parameters  
can be *exponentially more compact* than  
deep ones with positive parameters**

## ***subtractive MMs***



also called negative/signed/**subtractive** MMs

⇒ or **non-monotonic** circuits,...


**issue:** how to preserve non-negative outputs?

well understood for simple parametric forms

e.g., Weibulls, Gaussians

⇒ *constraints on variance, mean*

## subtractive MMs



also called negative/signed/**subtractive** MMs

⇒ or **non-monotonic** circuits,...


**issue:** how to preserve non-negative outputs?

well understood for simple parametric forms

e.g., Weibulls, Gaussians

⇒ constraints on variance, mean

# ***subtractive MMs***



also called negative/signed/**subtractive** MMs

⇒ or **non-monotonic** circuits,...


**issue:** how to preserve non-negative outputs?

well understood for simple parametric forms

e.g., Weibulls, Gaussians

⇒ *constraints on variance, mean*

# subtractive MMs as circuits




a **non-monotonic** smooth and (structured)  
decomposable circuit

$\Rightarrow$  possibly with negative outputs

$$c(\mathbf{X}) = \sum_{i=1}^K w_i c_i(\mathbf{X}), \quad w_i \in \mathbb{R},$$


## ***squaring shallow MMs***



$$c^2(\mathbf{X}) = \left( \sum_{i=1}^K w_i c_i(\mathbf{X}) \right)^2$$


$\Rightarrow$  ensure non-negative output

## *squaring shallow MMs*



$$\begin{aligned} c^2(\mathbf{X}) &= \left( \sum_{i=1}^K w_i c_i(\mathbf{X}) \right)^2 \\ &= \sum_{i=1}^K \sum_{j=1}^K w_i w_j c_i(\mathbf{X}) c_j(\mathbf{X}) \end{aligned}$$

## squaring shallow MMs




$$\begin{aligned} c^2(\mathbf{X}) &= \left( \sum_{i=1}^K w_i c_i(\mathbf{X}) \right)^2 \\ &= \sum_{i=1}^K \sum_{j=1}^K w_i w_j c_i(\mathbf{X}) c_j(\mathbf{X}) \end{aligned}$$

still a smooth and (str) decomposable PC with  $\mathcal{O}(K^2)$  components!

$\Rightarrow$  but still  $\mathcal{O}(K)$  parameters




## squaring shallow MMs



$$\begin{aligned} c^2(\mathbf{X}) &= \left( \sum_{i=1}^K w_i c_i(\mathbf{X}) \right)^2 \\ &= \sum_{i=1}^K \sum_{j=1}^K w_i w_j c_i(\mathbf{X}) c_j(\mathbf{X}) \end{aligned}$$

how to **renormalize**?


## squaring shallow MMs



$$\begin{aligned} c^2(\mathbf{X}) &= \left( \sum_{i=1}^K w_i c_i(\mathbf{X}) \right)^2 \\ &= \sum_{i=1}^K \sum_{j=1}^K w_i w_j c_i(\mathbf{X}) c_j(\mathbf{X}) \end{aligned}$$

to **renormalize**, we have to compute  $\sum_i \sum_j w_i w_j \int c_i(\mathbf{x}) c_j(\mathbf{x}) d\mathbf{x}$

## squaring shallow MMs



$$\begin{aligned} c^2(\mathbf{X}) &= \left( \sum_{i=1}^K w_i c_i(\mathbf{X}) \right)^2 \\ &= \sum_{i=1}^K \sum_{j=1}^K w_i w_j c_i(\mathbf{X}) c_j(\mathbf{X}) \end{aligned}$$

to **renormalize**, we have to compute  $\sum_i \sum_j w_i w_j \int c_i(\mathbf{x}) c_j(\mathbf{x}) d\mathbf{x}$   
 $\Rightarrow$  closed-form for e.g., if  $c_i, c_j$  are **exponential families...**!

***wait...!***

***how do we **learn** them?***

***wait...!***

***how do we **learn** them?***

***⇒ by maximizing the (log-)likelihood***

## ***which parameters?***

*how to reparameterize non-monotonic mixtures/circuits*

**Input functions.**

**Sum unit parameters.**

# ***which parameters?***

*how to reparameterize non-monotonic mixtures/circuits*

**Input functions.** Each input can be a different parametric ***function***

⇒ *Bernoullis, Categoricals, Gaussians, **polynomials**, small NNs, ...*

**Sum unit parameters.**

# *which parameters?*

*how to reparameterize non-monotonic mixtures/circuits*

**Input functions.** Each input can be a different parametric *function*

**Sum unit parameters.** They can be negative, i.e.,  $w_i \in \mathbb{R}$  and we need to renormalize the *negative log likelihood* loss after squaring

$$\min_{\theta} - \left( \sum_{i=1}^N 2 \log c_{\theta}(\mathbf{x}^{(i)}) - \log \int c_{\theta}^2(\mathbf{x}^{(i)}) d\mathbf{X} \right)$$



***wait...!***

***how do we **learn** them?***

***⇒ by maximizing the (log-)likelihood***

***wait...!***

***how do we **learn** them?***

***⇒ by maximizing the (log-)likelihood***


***just SGD your way as usual!***

***⇒ or any other gradient-based optimizer***

*what about **deep** mixtures/circuits?*

# GMMs

*as computational graphs*




$$p(\mathbf{X}) = w_1 \cdot p_1(\mathbf{X}') \cdot p_1(\mathbf{X}'') + \\ w_2 \cdot p_2(\mathbf{X}''') \cdot p_2(\mathbf{X}''')$$

$\Rightarrow$  local factorizations...

# GMMs

as computational graphs



$$p(\mathbf{X}) = w_1 \cdot p_1(\mathbf{X}') \cdot p_1(\mathbf{X}'') + w_2 \cdot p_2(\mathbf{X}''') \cdot p_2(\mathbf{X}''')$$

$\Rightarrow$  ...are product units

# ***probabilistic circuits (PCs)***

*a grammar for tractable computational graphs*

I. A simple tractable function is a circuit


$\Rightarrow$  e.g., a multivariate Gaussian or small  
neural network

$\bigcirc$   
 $X_1$

# ***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 queries*** = ***feedforward*** evaluation

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




***probabilistic queries*** = ***feedforward*** evaluation

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



# *probabilistic queries* = *feedforward* evaluation

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



# ***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



***learning & reasoning with circuits in pytorch***

`github.com/april-tools/circuit`

A screenshot of a GitHub web interface. The repository is 'cirkit' by 'april-tools'. The file 'learning-a-circuit.ipynb' is selected under the 'notebooks' directory. The file is owned by 'adrianjav' and was last updated 4 months ago. The interface shows tabs for 'Preview', 'Code', and 'Blame'. The 'Preview' tab is active, displaying the title 'Learning and Evaluating a Probabilistic Circuit' and a paragraph of text describing the notebook's content: 'In this notebook, we instantiate, learn, and evaluate a probabilistic circuit using cirkit. The probabilistic circuit we build estimates the distribution of MNIST images, which is then evaluated on unseen images, compute marginal probabilities, and sample new images. Here, we focus on the simplest experimental setting, where we want to instantiate a probabilistic circuit for MNIST images using some hyperparameters of our own choice, such as the type of the layers, their size and how to parameterize them. Then, we learn the parameters of the circuit and perform inference using PyTorch.'

## *a notebook on learning a deep circuit on MNIST*

<https://github.com/april-tools/cirkit/blob/main/notebooks/learning-a-circuit.ipynb>





loreloc updated notebooks with respect to API changes ✖ e3e7e80 · 2 days ago 🕒

Preview Code Blame 1082 lines (1082 loc) · 793 KB Raw 📄 📥

## 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 images](#)<sup>1</sup>

***mix& match your structure and layers***


[https://github.com/april-tools/cirkit/blob/main/notebooks/  
region-graphs-and-parametrisation.ipynb](https://github.com/april-tools/cirkit/blob/main/notebooks/region-graphs-and-parametrisation.ipynb)

# 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***

***determinism***

***the combination of certain  
structural properties  
**guarantees**  
tractable computation of  
certain query classes***

# ***structural properties***

***property A***

***property B***

***property C***

***property D***

***the combination of certain  
structural properties  
guarantees  
tractable computation of  
certain query classes***

# structural properties

property A

property B

property C

property D

*tractable* computation of *arbitrary integrals*

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

$\Rightarrow$  *sufficient* and *necessary* conditions  
for a single feedforward evaluation

$\Rightarrow$  tractable partition function

$\Rightarrow$  also any *conditional* is tractable



# structural properties

**smoothness**

**tractable** computation of **arbitrary integrals**

**decomposability**

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

**property C**

$\Rightarrow$  **sufficient** and **necessary** conditions  
for a single feedforward evaluation

**property D**

$\Rightarrow$  tractable partition function  
 $\Rightarrow$  also any **conditional** is tractable

# ***structural properties***

***smoothness***

**smoothness**  $\wedge$  **decomposability**  $\implies$  **multilinearity**

***decomposability***

***property C***


***property D***

## ***tractable marginals on PCs***

Original



Missing




Conditional sample




***use tractable models  
inside intractable pipelines  
where it matters!***


# ***tractable + intractable***



***tractable conditioning over every missing mask***  
(under submission)



***better than (V)AEs for missing values***  
*(under submission)*



how to efficiently square (and **renormalize**) a deep PC?

# compositional inference I



```
1 from cirq.symbolic.functional import integrate, multiply
2
3 #
4 # create a deep circuit
5 c = build_symbolic_circuit('quad-tree-4')
6
7 #
8 # compute the partition function of  $c^2$ 
9 def renormalize(c):
10     c2 = multiply(c, c)
11     return integrate(c2)
```



***structural properties***

***smoothness***

***decomposability***

***property C***

***property D***

# *structural properties*

*smoothness*

*decomposability*


*compatibility*

*property D*

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

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


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



how to efficiently square (and **renormalize**) a deep PC?



# ***squaring deep PCs***

*the tensorized way*



# squaring deep PCs



the tensorized way



squaring a circuit = squaring layers


# squaring deep PCs

the tensorized way




exactly compute  $\int \mathbf{c}(\mathbf{x}) \mathbf{c}(\mathbf{x}) d\mathbf{X}$  in time  $O(LK^2)$

## ***theorem 1***



$\exists p'$  requiring exponentially large  
**monotonic circuits...**

# *theorem 1*




$$\left( \begin{matrix} \text{node} \\ \text{node} \\ \vdots \\ \text{node} \end{matrix} \right)^2$$

...but compact

**squared non-monotonic circuits**



*more expressive?*




data

monoPC

monoPC<sup>2</sup>



non - monoPC<sup>2</sup>

*more expressive?*




# how more expressive?

real-world data



## ***theorem II***


$\exists p''$  requiring exponentially large  
**squared non-mono circuits...**




## ***theorem II***




**...but compact  
monotonic circuits...!**







*what if we use more than one square?*

## theorem III



$\exists p'''$  requiring exponentially large **squared non-mono circuits**...


## theorem III




...exponentially large **monotonic circuits**...



# theorem III



...but compact **SOS circuits...**



***a hierarchy of subtractive mixtures***

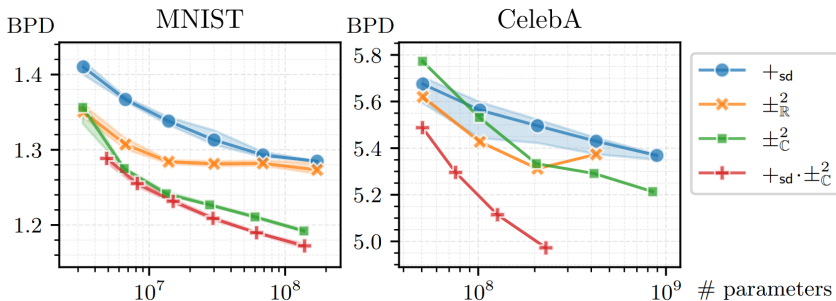
we can define circuits (and hence mixtures) over the Complex:

$$c^2(\mathbf{x}) = c(\mathbf{x})^\dagger c(\mathbf{x}), \quad c(\mathbf{x}) \in \mathbb{C}$$

and then we can note that they can be written as a SOS form

$$c^2(\mathbf{x}) = r(\mathbf{x})^2 + i(\mathbf{x})^2, \quad r(\mathbf{x}), i(\mathbf{x}) \in \mathbb{R}$$

***complex circuits are SOS (and scale better!)***



***complex circuits are SOS (and scale better!)***

## ***takeaway***

***“use squared mixtures  
over complex numbers  
(and you get a SOS for free)”***

## ***takeaway***

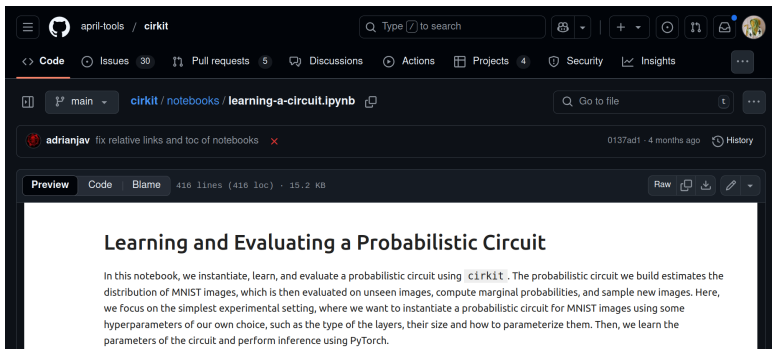
***“use squared mixtures  
over complex numbers  
(and you get a SOS for free)”***

***⇒ but how to **implement** them?***

# compositional inference I



```
1 from cirq.symbolic.functional import integrate, multiply,  
   ↪ conjugate  
2  
3 # create a deep circuit with complex parameters  
4 c = build_symbolic_complex_circuit('quad-tree-4')  
5  
6 # compute the partition function of  $c^2$   
7 def renormalize(c):  
8     c1 = conjugate(c)  
9     c2 = multiply(c, c1)  
10    return integrate(c2)
```




## *a notebook on learning SOS subtractive mixtures*


[https://github.com/april-tools/cirkit/blob/main/notebooks/  
sum-of-squares-circuits.ipynb](https://github.com/april-tools/cirkit/blob/main/notebooks/sum-of-squares-circuits.ipynb)




## *towards conclusions...*




**(deep)  
circuits**




**(hierarchical)  
tensor factorizations**



**decision trees**




**(hierarchical)  
mixtures**




**(deep)  
circuits**

$$\begin{aligned} & (d \rightarrow b) \wedge (e \rightarrow b) \\ & \vdash (\neg d \vee b) \wedge (\neg e \vee b) \\ & \vdash b \vee (\neg d \wedge \neg e) \end{aligned}$$


**logical  
formulas  
& constraints**




**(hierarchical)  
tensor factorizations**





***oh mixtures, you're so fine you blow my mind!***



$p(\mathbf{X})$






$$\sum_{i=1}^K w_i p_i(\mathbf{X}) \quad w_i > 0$$






$$p(\mathbf{X}) \quad \longrightarrow \quad \sum_{i=1}^K w_i p_i(\mathbf{X}) \quad w_i > 0$$

*“if someone publishes a paper on **model A**, there will be a paper about **mixtures of A** soon, with high probability”*

**A. Vergari**




$$p(\mathbf{X}) \xrightarrow{\quad} \sum_{i=1}^K w_i p_i(\mathbf{X}) \quad w_i > 0 \xrightarrow{\quad} \sum_{i=1}^{2^D} w_i p_i(\mathbf{X}) = PC(\mathbf{X})$$



$$p(\mathbf{X}) \longrightarrow \sum_{i=1}^K w_i p_i(\mathbf{X}) \quad w_i > 0 \longrightarrow \sum_{i=1}^{2^D} w_i p_i(\mathbf{X}) = \text{PC}(\mathbf{X})$$

$$\sum_j \left( \sum_{i=1}^K w_{ij} p_{ij}(\mathbf{X}) \right)^2 \quad w_{ij} \in \mathbb{R}$$






***learning & reasoning with circuits in pytorch***

`github.com/april-tools/circuit`





questions?

***structural properties***

***smoothness***


***decomposability***

***compatibility***


***determinism***

# determinism

the inputs of sum units are defined over disjoint supports



**deterministic circuit**



**non-deterministic circuit**