

Subtractive mixture models representation and learning

antonio vergari (he/him)

 @tetraduzione

joint work with Lorenzo Loconte, Aleksanteri M. Sladek, Stefan Mangel, Martin Trapp, Arno Solin, Nicolas Gillis

19th Oct 2023 - **Simons Institute**

april

*april is
probably a
recursive,
identifier of a
lab*

april

*about
probabilities
reasoning,
integrals &
logic*

why subtractions in mixture models

how to represent them as deep squared circuits?

what inference and model classes they support?

when are they more expressive

open problems

why subtractions in mixture models

how to represent them as deep squared circuits?

what inference and model classes they support?

when are they more expressive

open problems

why subtractions in mixture models

how to represent them as deep squared circuits?

what inference and model classes they support?

when are they more expressive

open problems

why subtractions in mixture models

how to represent them as deep squared circuits?

what inference and model classes they support?

when are they more expressive

open problems

why subtractions in mixture models

how to represent them as deep squared circuits?

what inference and model classes they support?

when are they more expressive

open problems


TCS crowd

**a circuit lowerbound to play with
connections with mixtures/PGMs/learning**


ML crowd

**(some) new tractable model(s) to play with
a tensorized way to represent circuits**


why subtractions in mixture models



mixtures are a staple in probML




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



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

additive MMs

are so cool!



easily represented as shallow **probabilistic circuits** (PCs)

\Rightarrow *smooth, (structured) decomposable*


these are **monotonic** PCs

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

universal approximators...

additive MMs

are so cool!



easily represented as shallow **probabilistic circuits** (PCs)

\Rightarrow *smooth, (structured) decomposable*


these are **monotonic** PCs

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

universal approximators...

additive MMs

are so cool!



easily represented as shallow **probabilistic circuits** (PCs)

\Rightarrow *smooth, (structured) decomposable*


these are **monotonic** PCs

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

universal approximators...

additive MMs

are so cool!



easily represented as shallow **probabilistic circuits** (PCs)


\Rightarrow *smooth, (structured) decomposable*

these are **monotonic** PCs


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

universal approximators...


additive MMs




additive MMs




additive MMs





additive MMs



however...





however...




GMM ($K = 2$)

however...





GMM ($K = 2$)




GMM ($K = 16$)


however...



GMM ($K = 2$)



GMM ($K = 16$)




nGMM² ($K = 2$)

SPOILER ALERT

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

subtractive MMs



sometimes called **negative** 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



sometimes called **negative** 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



sometimes called **negative** 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*

tl;dr

“Understand when and how we can use negative parameters in deep **subtractive mixture models**”

tl;dr


***“Understand when and how
we can use negative parameters
in deep **non-monotonic squared circuits**”***

tl;dr

***“Understand when and how
we can use negative parameters
in deep **non-monotonic squared circuits**”***

⇒ later PSD kernel models, tensor networks, ...

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



e.g., a squared GMM with negative parameters w_i


but we do not require non-negative inputs!

⇒ e.g. use *splines*

to renormalize we need to compute

$$\int c^2(\mathbf{x}) d\mathbf{x} = \sum_{i=1}^K \sum_{j=1}^K w_i w_j \int c_i(\mathbf{x}) c_j(\mathbf{x}) d\mathbf{x}$$

squaring shallow MMs



e.g., a squared GMM with negative parameters w_i


but we do not require non-negative inputs!

⇒ e.g. use **splines**

to renormalize we need to compute

$$\int c^2(\mathbf{x}) d\mathbf{x} = \sum_{i=1}^K \sum_{j=1}^K w_i w_j \int c_i(\mathbf{x}) c_j(\mathbf{x}) d\mathbf{x}$$

squaring shallow MMs



e.g., a squared GMM with negative parameters w_i

but we do not require non-negative inputs!

⇒ e.g. use **splines**

to renormalize we need to compute


$$\int c^2(\mathbf{x}) d\mathbf{x} = \sum_{i=1}^K \sum_{j=1}^K w_i w_j \int c_i(\mathbf{x}) c_j(\mathbf{x}) d\mathbf{x}$$


why subtractions in mixture models

how to represent them as deep squared circuits?

Circuits


A grammar for tractable computational graphs





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


Tractable square




smooth

structured decomposable

Tractable square





smooth
structured decomposable



smooth
structured decomposable

Tractable square



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

Tractable square

Algorithm 3 MULTIPLY(p, q, cache)

```
1: Input: two circuits  $p(\mathbf{Z})$  and  $q(\mathbf{Y})$  that are compatible over  $\mathbf{X} = \mathbf{Z} \cap \mathbf{Y}$  and a cache for
   memoization
2: Output: their product circuit  $m(\mathbf{Z} \cup \mathbf{Y}) = p(\mathbf{Z})q(\mathbf{Y})$ 
3: if  $(p, q) \in \text{cache}$  then return  $\text{cache}(p, q)$ 
4: if  $\phi(p) \cap \phi(q) = \emptyset$  then
5:    $m \leftarrow \text{PRODUCT}(\{p, q\}); s \leftarrow \text{True}$ 
6: else if  $p, q$  are input units then
7:    $m \leftarrow \text{INPUT}(p(\mathbf{Z}) \cdot q(\mathbf{Y}), \mathbf{Z} \cup \mathbf{Y})$ 
8:    $s \leftarrow [\text{supp}(p(\mathbf{X})) \cap \text{supp}(q(\mathbf{X})) \neq \emptyset]$ 
9: else if  $p$  is an input unit then
10:   $n \leftarrow \{\}; s \leftarrow \text{False} // q(\mathbf{Y}) = \sum_j \theta_j^q(\mathbf{Y})$ 
11:  for  $j = 1$  to  $|\text{in}(q)|$  do
12:     $n', s' \leftarrow \text{MULTIPLY}(p, q_j, \text{cache})$ 
13:     $n \leftarrow n \cup \{n'\}; s \leftarrow s \vee s'$ 
14:  if  $s$  then  $m \leftarrow \text{SUM}(n, \{\theta_j^q\}_{j=1}^{|\text{in}(q)|})$  else  $m \leftarrow \text{null}$ 
15: else if  $q$  is an input unit then
16:   $n \leftarrow \{\}; s \leftarrow \text{False} // p(\mathbf{Z}) = \sum_i \theta_i^p(\mathbf{Z})$ 
17:  for  $i = 1$  to  $|\text{in}(p)|$  do
18:     $n', s' \leftarrow \text{MULTIPLY}(p_i, q, \text{cache})$ 
19:     $n \leftarrow n \cup \{n'\}; s \leftarrow s \vee s'$ 
20:  if  $s$  then  $m \leftarrow \text{SUM}(n, \{\theta_i^p\}_{i=1}^{|\text{in}(p)|})$  else  $m \leftarrow \text{null}$ 
21: else if  $p, q$  are product units then
22:   $n \leftarrow \{\}; s \leftarrow \text{True}$ 
23:   $\{p_i, q_i\}_{i=1}^k \leftarrow \text{sortPairsByScope}(p, q, \mathbf{X})$ 
24:  for  $i = 1$  to  $k$  do
25:     $n', s' \leftarrow \text{MULTIPLY}(p_i, q_i, \text{cache})$ 
26:     $n \leftarrow n \cup \{n'\}; s \leftarrow s \wedge s'$ 
27:  if  $s$  then  $m \leftarrow \text{PRODUCT}(n)$  else  $m \leftarrow \text{null}$ 
28: else if  $p, q$  are sum units then
29:   $n \leftarrow \{\}; w \leftarrow \{\}; s \leftarrow \text{False}$ 
30:  for  $i = 1$  to  $|\text{in}(p)|, j = 1$  to  $|\text{in}(q)|$  do
31:     $n', s' \leftarrow \text{MULTIPLY}(p_i, q_j, \text{cache})$ 
32:     $n \leftarrow n \cup n'; w \leftarrow w \cup \{\theta_i^p, \theta_j^q\}; s \leftarrow s \vee s'$ 
33:  if  $s$  then  $m \leftarrow \text{SUM}(n, w)$  else  $m \leftarrow \text{null}$ 
34:   $\text{cache}(p, q) \leftarrow (m, s)$ 
35: return  $m, s$ 
```


on tensorized PCs


Algorithm 1 squareTensorizedCircuit(ℓ, \mathcal{R})

Input: A tensorized circuit having output layer ℓ and defined on a tree RG rooted by \mathcal{R} .


Output: The tensorized squared circuit defined on the same tree RG having ℓ^2 as output layer computing $\ell \otimes \ell$.

1: if ℓ is an input layer then	9: return $\ell_i^2 \odot \ell_{ii}^2$
2: ℓ computes K functions $f_i(\mathcal{R})$	10: else $\triangleright \ell$ is a sum layer
3: return An input layer ℓ^2 computing all K^2	11: $\{(\ell_i, \mathcal{R}_i)\} \leftarrow \text{getInputs}(\ell, \mathcal{R})$
4: product combinations $f_i(\mathcal{R})f_j(\mathcal{R})$	12: $\ell_i^2 \leftarrow \text{squareTensorizedCircuit}(\ell_i, \mathcal{R}_i)$
5: else if ℓ is a product layer then	13: $\mathbf{W} \in \mathbb{R}^{S \times K} \leftarrow \text{getParameters}(\ell)$
6: $\{(\ell_i, \mathcal{R}_i), (\ell_{ii}, \mathcal{R}_{ii})\} \leftarrow \text{getInputs}(\ell, \mathcal{R})$	14: $\mathbf{W}' \in \mathbb{R}^{S^2 \times K^2} \leftarrow \mathbf{W} \otimes \mathbf{W}$
7: $\ell_i^2 \leftarrow \text{squareTensorizedCircuit}(\ell_i, \mathcal{R}_i)$	15: return $\mathbf{W}' \ell_i^2$
8: $\ell_{ii}^2 \leftarrow \text{squareTensorizedCircuit}(\ell_{ii}, \mathcal{R}_{ii})$	





Tensorizing str-dec PCs





abstract computations into *layers*

group units with the same scope


parameterize connections by matrix/vector operations

Tensorizing str-dec PCs

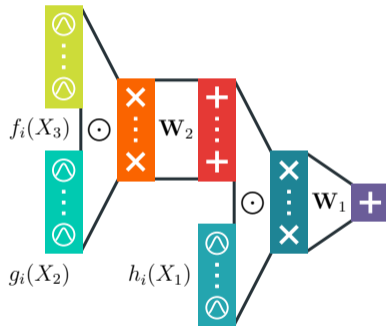


Mari, Vessio, and Vergari, "Unifying and Understanding Overparameterized Circuit Representations via Low-Rank Tensor Decompositions", 2023

Tensorizing str-dec PCs




region graph / vtrees / pseudotree





squaring deep PCs

the tensorized way




squaring deep PCs

the tensorized way





squaring deep PCs

the tensorized way




squaring deep PCs

the tensorized way




squaring reduces to square layers

Tractable squares




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

Tractable squares



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

Tractable squares

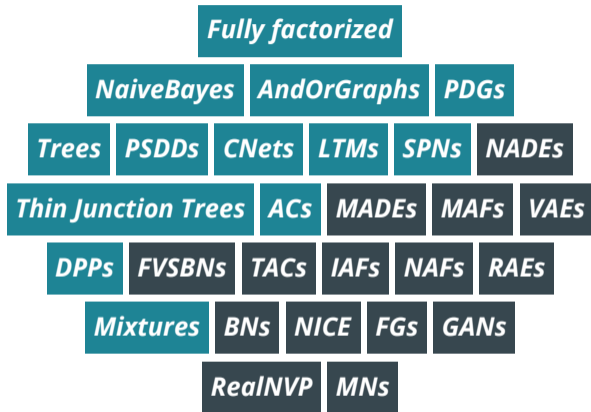


exactly compute $\int \mathbf{c}(\mathbf{x})\mathbf{c}(\mathbf{x})d\mathbf{X}$ in time $O(LK^2)$

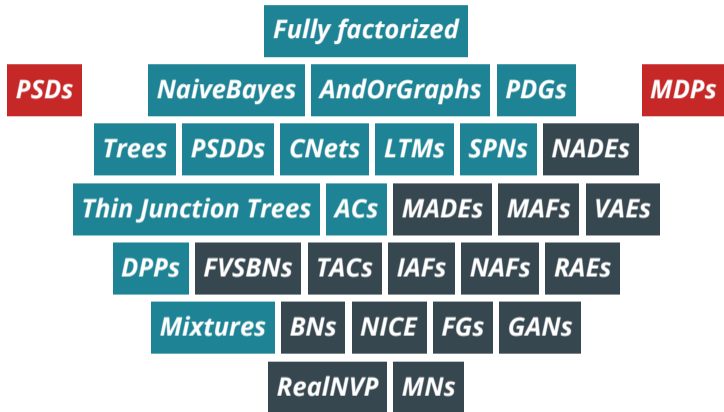
why subtractions in mixture models

how to represent them as deep squared circuits?

what inference and model classes they support?



the alphabet soup of *tractable* models



new entries in the family!

PSD kernels

Given a **kernel** κ and a set of d data points $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(d)}$ with $\boldsymbol{\kappa}(\mathbf{x}) = [\kappa(\mathbf{x}, \mathbf{x}^{(1)}), \dots, \kappa(\mathbf{x}, \mathbf{x}^{(d)})]^\top \in \mathbb{R}^d$, define the non-negative function

$$f(\mathbf{x}; \mathbf{A}, \boldsymbol{\kappa}) = \boldsymbol{\kappa}(\mathbf{x})^\top \mathbf{A} \boldsymbol{\kappa}(\mathbf{x})$$

where \mathbf{A} is a real $d \times d$ **positive semi-definite** matrix.

PSD kernels

Given a kernel κ and a set of d data points $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(d)}$ with $\boldsymbol{\kappa}(\mathbf{x}) = [\kappa(\mathbf{x}, \mathbf{x}^{(1)}), \dots, \kappa(\mathbf{x}, \mathbf{x}^{(d)})]^\top \in \mathbb{R}^d$, define the non-negative function

$$f(\mathbf{x}; \mathbf{A}, \boldsymbol{\kappa}) = \boldsymbol{\kappa}(\mathbf{x})^\top \mathbf{A} \boldsymbol{\kappa}(\mathbf{x})$$

where \mathbf{A} is a real $d \times d$ positive semi-definite matrix. Just a *mixture of squared PCs*

$$f(\mathbf{x}; \mathbf{A}, \boldsymbol{\kappa}) = \boldsymbol{\kappa}(\mathbf{x})^\top \left(\sum_{i=1}^r \lambda_i \mathbf{u}_i \mathbf{u}_i^\top \right) \boldsymbol{\kappa}(\mathbf{x}) = \sum_{i=1}^r \lambda_i \left(\mathbf{u}_i^\top \boldsymbol{\kappa}(\mathbf{x}) \right)^2,$$

tensor networks

matrix-product state, tensor trains, Born machines,...

A **matrix-product state** or **tensor-train** factorizes a D -dimensional tensor \mathcal{T} as

$$\mathcal{T}[x_1, \dots, x_D] = \sum_{i_1=1}^r \sum_{i_2=1}^r \cdots \sum_{i_{D-1}=1}^r \mathbf{A}_1[x_1, i_1] \mathbf{A}_2[x_2, i_1, i_2] \cdots \mathbf{A}_D[x_D, i_{D-1}]$$

and a **Born machine** squares it

$$\mathcal{B}[x_1, \dots, x_D] = \left(\sum_{i_1=1}^r \sum_{i_2=1}^r \cdots \sum_{i_{D-1}=1}^r \mathbf{A}_1[x_1, i_1] \mathbf{A}_2[x_2, i_1, i_2] \cdots \mathbf{A}_D[x_D, i_{D-1}] \right)^2$$

tensor networks

matrix-product state, tensor trains, Born machines,...

A matrix-product state or tensor train factorizes a D -dimensional tensor \mathcal{T} as

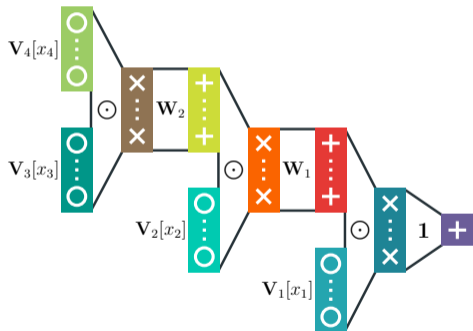

$$\mathcal{T}[x_1, \dots, x_D] = \sum_{i_1=1}^r \sum_{i_2=1}^r \cdots \sum_{i_{D-1}=1}^r \mathbf{A}_1[x_1, i_1] \mathbf{A}_2[x_2, i_1, i_2] \cdots \mathbf{A}_D[x_D, i_{D-1}]$$

and a Born machine squares it

$$\mathcal{B}[x_1, \dots, x_D] = \left(\sum_{i_1=1}^r \sum_{i_2=1}^r \cdots \sum_{i_{D-1}=1}^r \mathbf{A}_1[x_1, i_1] \mathbf{A}_2[x_2, i_1, i_2] \cdots \mathbf{A}_D[x_D, i_{D-1}] \right)^2$$

tensor networks


matrix-product state, tensor trains, Born machines,...



Marginals

and conditionals

monotonic




$$\mathcal{O}(|c|)$$

Marginals


and conditionals

monotonic



$$\mathcal{O}(|c|)$$


non-monotonic



$$\mathcal{O}(|c|^2)$$

Sampling


monotonic



$$\mathcal{O}(|c|)$$


Sampling

monotonic



$$\mathcal{O}(|c|)$$

non-monotonic




$$\mathcal{O}(|\mathbf{X}||c|^2)$$

by autoregressive sampling


Sampling

monotonic



$$\mathcal{O}(|c|)$$

non-monotonic




$$\mathcal{O}(\log |\mathbf{X}| |c|^2)$$

for balanced vtrees

MAP

monotonic




$$\mathcal{O}(|c|)$$

if deterministic

MAP


monotonic



$$\mathcal{O}(|c|)$$

if deterministic

non-monotonic



$$\mathcal{O}(|c|)$$

if deterministic

why subtractions in mixture models


how to represent them as deep squared circuits?

what inference and model classes they support?

when are they more expressive

more expressive?

squared probabilistic (o)BDDs, SDDs, str-d-DNNF?




smooth, **deterministic**
structured decomposable

?

more expressive?

squared probabilistic (o)BDDs, SDDs, str-d-DNNF?



no increase in size


no increase in expressiveness

no negative weights

smooth, **deterministic**
structured decomposable

more expressive?

squared probabilistic (o)BDDs, SDDs, str-d-DNNF?



no increase in size


no increase in expressiveness

no negative weights

smooth, **deterministic**
structured decomposable

more expressive?

squared probabilistic (o)BDDs, SDDs, str-d-DNNF?



no increase in size


no increase in expressiveness

no negative weights

smooth, **deterministic**
structured decomposable

more expressive?

squared probabilistic (o)BDDs, SDDs, str-d-DNNF?



no increase in size

no increase in expressiveness

no negative weights

smooth, **deterministic**
structured decomposable

more expressive?

exponential separation

Theorem: there is a class of non-negative functions \mathcal{F} over variables \mathbf{X} that can be represented by **compact a squared non-monotonic str-dec PC** but for which the smallest monotonic str-dec PC computing $F \in \mathcal{F}$ has size $2^{\Omega(|\mathbf{X}|)}$ **smallest monotonic str-dec PC computing $F \in \mathcal{F}$ has exponential size $2^{\Omega(|\mathbf{X}|)}$.**

more expressive?

exponential separation

Theorem: there is a class of non-negative functions \mathcal{F} over variables \mathbf{X} that can be represented by compact a squared non-monotonic str-dec PC but for which the ***smallest monotonic str-dec PC computing $F \in \mathcal{F}$ has exponential size $2^{\Omega(|\mathbf{X}|)}$.***

more expressive?


exponential separation

Theorem: there is a class of non-negative functions \mathcal{F} over variables \mathbf{X} that can be represented by compact a squared non-monotonic str-dec PC but for which the smallest monotonic str-dec PC computing $F \in \mathcal{F}$ has size $2^{\Omega(|\mathbf{X}|)}$ **smallest monotonic str-dec PC computing $F \in \mathcal{F}$ has exponential size $2^{\Omega(|\mathbf{X}|)}$.**

$$\text{UDISJ}_G(\mathbf{X}_v) := \left(1 - \sum_{uv \in E} X_u X_v \right)^2$$

Y\Z	000	100	010	001	110	101	011	111
000	1	1	1	1	1	1	1	1
100	1	0	1	1	0	0	1	0
010	1	1	0	1	0	1	0	0
001	1	1	1	0	1	0	0	0
110	1	0	0	1	1	0	0	1
101	1	0	1	0	0	1	0	1
011	1	1	0	0	0	0	1	1
111	1	0	0	0	1	1	1	4

more expressive?




data

monoPC

monoPC²

non – monoPC²

more expressive?




data

monoPC

monoPC²


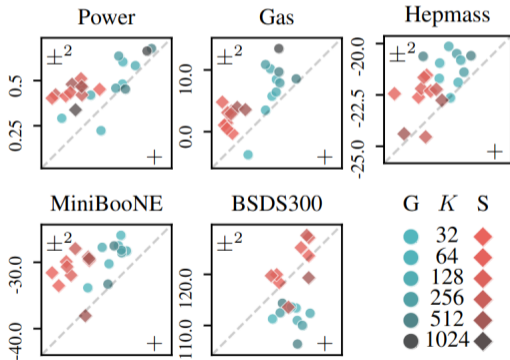
non - monoPC²

more expressive?



how more expressive?

for the ML crowd



why subtractions in mixture models

how to represent them as deep squared circuits?

what inference and model classes they support?

when are they more expressive

open problems

???

I how to retrieve a latent variable semantics?

II how to perform structure learning?

III more expressive than other circuit classes?

IV use logic-as-circuits for physics

???

I how to retrieve a latent variable semantics?

II how to perform structure learning?

III more expressive than other circuit classes?

IV use logic-as-circuits for physics

???

I how to retrieve a latent variable semantics?

II how to perform structure learning?

III more expressive than other circuit classes?

IV use logic-as-circuits for physics

???

I how to retrieve a latent variable semantics?

II how to perform structure learning?

III more expressive than other circuit classes?

IV use logic-as-circuits for physics

TCS crowd

ML crowd


a circuit lowerbound to play with
connections with mixtures/PGMs/learning

ML crowd

TCS crowd


(some) new tractable model(s) to play with
a tensorized way to represent circuits

$$\int p(\mathbf{x}) \times \log \left(p(\mathbf{x}) / q(\mathbf{x}) \right) d\mathbf{X}$$



build a LEGO-like query calculus...

+, **×**, **pow**, **log**, **exp**, **/**



property A, property B
property C

property A, property B

automating probabilistic reasoning

The TCS perspective

	Query	Tract. Conditions
CROSS ENTROPY	$-\int p(\mathbf{x}) \log q(\mathbf{x}) d\mathbf{X}$	Cmp, q Det
SHANNON ENTROPY	$-\sum p(\mathbf{x}) \log p(\mathbf{x})$	Sm, Dec, Det
RÉNYI ENTROPY	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{N}$	SD
MUTUAL INFORMATION	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{R}_+$	Sm, Dec, Det
KULLBACK-LEIBLER DIV.	$\int p(\mathbf{x}, \mathbf{y}) \log(p(\mathbf{x}, \mathbf{y}) / (p(\mathbf{x})p(\mathbf{y})))$	Sm, SD, Det*
RÉNYI'S ALPHA DIV.	$\int p(\mathbf{x}) \log(p(\mathbf{x}) / q(\mathbf{x})) d\mathbf{X}$	Cmp, Det
ITAKURA-SAITO DIV.	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) q^{1-\alpha}(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{N}$	Cmp, q Det
CAUCHY-SCHWARZ DIV.	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) q^{1-\alpha}(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{R}$	Cmp, Det
SQUARED LOSS	$\int [p(\mathbf{x}) / q(\mathbf{x}) - \log(p(\mathbf{x}) / q(\mathbf{x})) - 1] d\mathbf{X}$	Cmp, Det
	$-\log \frac{\int p(\mathbf{x}) q(\mathbf{x}) d\mathbf{X}}{\sqrt{\int p^2(\mathbf{x}) d\mathbf{X} \int q^2(\mathbf{x}) d\mathbf{X}}}$	Cmp
	$\int (p(\mathbf{x}) - q(\mathbf{x}))^2 d\mathbf{X}$	Cmp

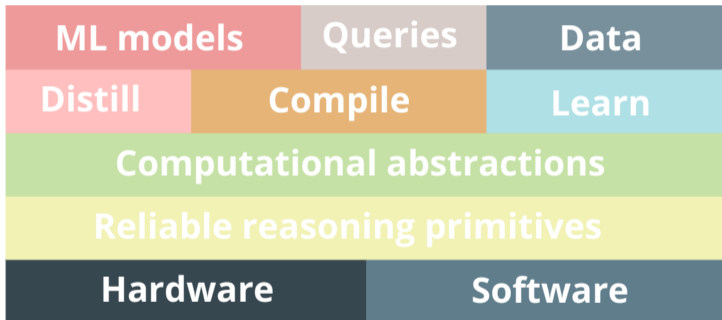
compositionally derive the tractability of many more queries

The TCS perspective

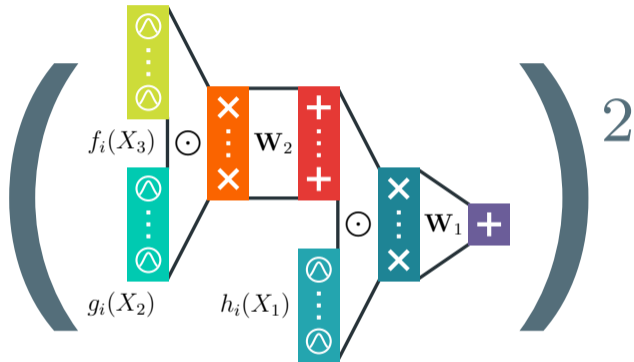
	Query	Tract. Conditions	Hardness
CROSS ENTROPY	$-\int p(\mathbf{x}) \log q(\mathbf{x}) d\mathbf{X}$	Cmp, q Det	#P-hard w/o Det
SHANNON ENTROPY	$-\sum p(\mathbf{x}) \log p(\mathbf{x})$	Sm, Dec, Det	coNP-hard w/o Det
RÉNYI ENTROPY	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{N}$	SD	#P-hard w/o SD
	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{R}_+$	Sm, Dec, Det	#P-hard w/o Det
MUTUAL INFORMATION	$\int p(\mathbf{x}, \mathbf{y}) \log(p(\mathbf{x}, \mathbf{y}) / (p(\mathbf{x})p(\mathbf{y})))$	Sm, SD, Det*	coNP-hard w/o SD
KULLBACK-LEIBLER DIV.	$\int p(\mathbf{x}) \log(p(\mathbf{x}) / q(\mathbf{x})) d\mathbf{X}$	Cmp, Det	#P-hard w/o Det
RÉNYI'S ALPHA DIV.	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) q^{1-\alpha}(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{N}$	Cmp, q Det	#P-hard w/o Det
	$(1 - \alpha)^{-1} \log \int p^\alpha(\mathbf{x}) q^{1-\alpha}(\mathbf{x}) d\mathbf{X}, \alpha \in \mathbb{R}$	Cmp, Det	#P-hard w/o Det
ITAKURA-SAITO DIV.	$\int [p(\mathbf{x}) / q(\mathbf{x}) - \log(p(\mathbf{x}) / q(\mathbf{x})) - 1] d\mathbf{X}$	Cmp, Det	#P-hard w/o Det
CAUCHY-SCHWARZ DIV.	$-\log \frac{\int p(\mathbf{x}) q(\mathbf{x}) d\mathbf{X}}{\sqrt{\int p^2(\mathbf{x}) d\mathbf{X} \int q^2(\mathbf{x}) d\mathbf{X}}}$	Cmp	#P-hard w/o Cmp
SQUARED LOSS	$\int (p(\mathbf{x}) - q(\mathbf{x}))^2 d\mathbf{X}$	Cmp	#P-hard w/o Cmp

proving hardness for when some input properties are not satisfied

UNREAL



realizing a full “virtual machine” for reasoning



questions?