A Probabilistic Neuro-symbolic Layer for Algebraic Constraint Satisfaction

Leander Kurscheidt

University of Edinburgh, UK

Paolo Morettin University of Trento, Italy **Roberto Sebastiani** University of Trento, Italy

Andrea Passerini University of Trento, Italy Antonio Vergari University of Edinburgh, UK

A Probabilistic Neuro-symbolic Layer for Algebraic Constraint Satisfaction

Leander Kurscheidt

University of Edinburgh, UK

Paolo Morettin University of Trento, Italy **Roberto Sebastiani** University of Trento, Italy

Andrea Passerini University of Trento, Italy Antonio Vergari University of Edinburgh, UK

A Probabilistic Neuro-symbolic Layer for Algebraic Constraint Satisfaction

Leander Kurscheidt

University of Edinburgh, UK

Paolo Morettin University of Trento, Italy **Roberto Sebastiani** University of Trento, Italy

Andrea Passerini University of Trento, Italy Antonio Vergari University of Edinburgh, UK



Leander Kurscheidt

University of Edinburgh, UK

Paolo Morettin University of Trento, Italy **Roberto Sebastiani** University of Trento, Italy

Andrea Passerini University of Trento, Italy Antonio Vergari University of Edinburgh, UK

Climate Science



$$\sum y_i = y_{obs}$$

[Harder et al. 2023]

Climate Science







[Harder et al. 2023]

5/32







Fairness

Trajectory Prediction



 $\sum y_i = y_{obs}$ [Harder et al. 2023]

[Dwork et al. 2012]

[Robicquet et al. 2016]











 $p(\mathbf{y}|\mathbf{x}) = ?$

How to model $p(\mathbf{y}|\mathbf{x})$ with a neural network?





$$p(\mathbf{y}|\mathbf{x}) = \text{GMM}(\mathbf{y}; \boldsymbol{\lambda} = f_{\boldsymbol{\theta}}(\mathbf{x}))$$

with
$$oldsymbol{\lambda} = (oldsymbol{\pi}, \{(\mu_i, \Sigma_i)\}_{i=1}^k)$$



issue: put densities/mass where constraints are not satisfied!

- \Rightarrow loosing probability mass
 - sampling invalid points
 - unreliable predictions

Bishop, "Mixture density networks", Tech Report, 1994



Can we do **better**?



Can we do better?

Yes! A plug&play layer that

puts mass inside constrains
guarantees valid predictions
end-to-end differentiable





Enter PAL



$$p_{\Theta}(\mathbf{Y} \mid \mathbf{x}) = \frac{q(\mathbf{Y}; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \cdot \mathbb{1}\{\mathbf{Y} \models \varphi\}}{\int q(\mathbf{y}'; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'}$$



$$p_{\Theta}(\mathbf{Y} \mid \mathbf{x}) = \frac{q(\mathbf{Y}; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \cdot \mathbf{1}\{\mathbf{Y} \models \varphi\}}{\int q(\mathbf{y}'; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \mathbf{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'}$$

10/32







Belle, Passerini, and Van den Broeck, "Probabilistic Inference in Hybrid Domains by Weighted Model Integration", IJCAI, 2015

How to compute the normalization constant?

How to compute the normalization constant?

\Rightarrow numerical approximation!?

Numerical Approximation



Smet et al., "Neural probabilistic logic programming in discrete-continuous domains", UAI, 2023 12/32

Numerical Approximation



Smet et al., "Neural probabilistic logic programming in discrete-continuous domains", UAI, 2023

How to compute the normalization constant?

\Rightarrow numerical approximation!? **no!**

Can we pick density q and constraints φ such that integration is efficient — and only needed once?



$$p_{\Theta}(\mathbf{Y} \mid \mathbf{x}) = \frac{q(\mathbf{Y}; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \cdot \mathbb{1}\{\mathbf{Y} \models \varphi\}}{\int q(\mathbf{y}'; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'}$$



$\mathbb{1}\{\mathbf{Y}\models\varphi\}$

 $\varphi \in \text{satisfiability modulo}$ linear real arithmetic (SMT(\mathcal{LRA}))

The Constraints



 $\varphi \in \mathsf{satisfiability} \mod \mathsf{loc}$ linear real arithmetic (SMT(\mathcal{LRA}))





Barrett and Tinelli, "Satisfiability modulo theories", Handbook of Model Checking, 2018





 $\varphi \in$ satisfiability modulo linear real arithmetic (SMT(\mathcal{LRA}))





$$p_{\Theta}(\mathbf{Y} \mid \mathbf{x}) = \frac{q(\mathbf{Y}; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \cdot \mathbb{1}\{\mathbf{Y} \models \varphi\}}{\int q(\mathbf{y}'; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \mathbb{1}\{\mathbf{y}' \models \varphi\} \, d\mathbf{Y}'}$$


Zeng et al., "Probabilistic Inference with Algebraic Constraints: Theoretical Limits and Practical Approximations", <u>NeurIPS</u>, 2020



Zeng et al., "Probabilistic Inference with Algebraic Constraints: Theoretical Limits and Practical Approximations", <u>NeurIPS</u>, 2020



Zeng et al., "Probabilistic Inference with Algebraic Constraints: Theoretical Limits and Practical Approximations", <u>NeurIPS</u>, 2020

 \Rightarrow

 \Rightarrow



Zeng et al., "Probabilistic Inference with Algebraic Constraints: Theoretical Limits and Practical Approximations", NeurIPS, 2020

$$q(\mathbf{Y}; \boldsymbol{\lambda} = f_{\boldsymbol{\psi}}(\mathbf{x}))$$

 \Rightarrow

piecewise polynomials, e.g., splines squared for non-negativity



Zeng et al., "Probabilistic Inference with Algebraic Constraints: Theoretical Limits and Practical Approximations", <u>NeurIPS</u>, 2020

$$q(\mathbf{Y}; \boldsymbol{\lambda} = f_{\boldsymbol{\psi}}(\mathbf{x}))$$

⇒ piecewise polynomials, e.g., splines ⇒ squared for non-negativity ⇒ exact integration on a simplex [Grundmann and Möller 1978]



Zeng et al., "Probabilistic Inference with Algebraic Constraints: Theoretical Limits and Practical Approximations", NeurIPS, 2020



$$p_{\Theta}(\mathbf{Y} \mid \mathbf{x}) = \frac{q(\mathbf{Y}; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \cdot \mathbb{1}\{\mathbf{Y} \models \varphi\}}{\int q(\mathbf{y}'; \boldsymbol{\lambda} = f_{\psi}(\mathbf{x})) \mathbb{1}\{\mathbf{y}' \models \varphi\} \, d\mathbf{Y}'}$$



- I. we introduce a new, GPU-accelerated integration backend called **GASP!**
- II. it calculates the symbolic integral in seconds
 - \Rightarrow **8 seconds** for a **81 piece** squared spline (**deg. 12**) on the SDD-Dataset
 - \Rightarrow up to $2~{\rm orders}~{\rm of}~{\rm magnitudes}$ faster than previous WMI-Solvers

III. the result is a function of λ

- \Rightarrow compile only once, evaluate cheaply many times
- ⇒ training independent on constraint complexity



I. we introduce a new, GPU-accelerated integration backend called **GASP!**

II. it calculates the symbolic integral in seconds

 \Rightarrow **8 seconds** for a **81 piece** squared spline (**deg. 12**) on the SDD-Dataset

 \Rightarrow up to $2~{\rm orders}~{\rm of}~{\rm magnitudes}$ faster than previous WMI-Solvers

III. the result is a function of λ

 \Rightarrow compile only once, evaluate cheaply many times

training independent on constraint complexity

$$\int q(\mathbf{y}'; \boldsymbol{\lambda}) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'$$
$$= I(\boldsymbol{\lambda})$$

- I. we introduce a new, GPU-accelerated integration backend called **GASP!**
- II. it calculates the symbolic integral in seconds
 - \Rightarrow **8 seconds** for a **81 piece** squared spline (**deg. 12**) on the SDD-Dataset
 - \Rightarrow up to $2~{\rm orders}~{\rm of}~{\rm magnitudes}$ faster than previous WMI-Solvers
- III. the result is a function of λ
 - \Rightarrow compile only once, evaluate cheaply many times
 - ⇒ training independent on constraint complexity

$$\int q(\mathbf{y}'; \boldsymbol{\lambda}) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'$$
$$= I(\boldsymbol{\lambda})$$

- I. we introduce a new, GPU-accelerated integration backend called **GASP!**
- II. it calculates the symbolic integral in seconds
 - \Rightarrow 8 seconds for a 81 piece squared spline (deg. 12) on the SDD-Dataset
 - \Rightarrow up to $2~{\rm orders}~{\rm of}~{\rm magnitudes}$ faster than previous WMI-Solvers
- lll. the result is a function of λ
 - \Rightarrow compile only once, evaluate cheaply many times
 - ⇒ training independent on constraint complexity

$$\int q(\mathbf{y}'; \boldsymbol{\lambda}) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'$$
$$= I(\boldsymbol{\lambda})$$

- I. we introduce a new, GPU-accelerated integration backend called **GASP!**
- II. it calculates the symbolic integral in seconds
 - \Rightarrow 8 seconds for a 81 piece squared spline (deg. 12) on the SDD-Dataset
 - \Rightarrow up to 2 orders of magnitudes faster than previous WMI-Solvers
- lll. the result is a function of $oldsymbol{\lambda}$
 - \Rightarrow compile only once, evaluate cheaply many times
 - ⇒ training independent on constraint complexity

$$\int q(\mathbf{y}'; \boldsymbol{\lambda}) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'$$
$$= I(\boldsymbol{\lambda})$$

- I. we introduce a new, GPU-accelerated integration backend called **GASP!**
- II. it calculates the symbolic integral in seconds
 - \Rightarrow 8 seconds for a 81 piece squared spline (deg. 12) on the SDD-Dataset
 - \Rightarrow up to 2 orders of magnitudes faster than previous WMI-Solvers

III. the result is a function of $oldsymbol{\lambda}$

- \Rightarrow compile only once, evaluate cheaply many times
- training independent on constraint complexity

$$\int q(\mathbf{y}'; \boldsymbol{\lambda}) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'$$
$$= I(\boldsymbol{\lambda})$$

- I. we introduce a new, GPU-accelerated integration backend called **GASP!**
- II. it calculates the symbolic integral in seconds
 - \Rightarrow 8 seconds for a 81 piece squared spline (deg. 12) on the SDD-Dataset
 - \Rightarrow up to 2 orders of magnitudes faster than previous WMI-Solvers
- III. the result is a function of $oldsymbol{\lambda}$
 - \Rightarrow compile only once, evaluate cheaply many times
 - ⇒ training independent on constraint complexity

$$\int q(\mathbf{y}'; \boldsymbol{\lambda}) \mathbb{1}\{\mathbf{y}' \models \varphi\} d\mathbf{Y}'$$
$$= I(\boldsymbol{\lambda})$$

- I. we introduce a new, GPU-accelerated integration backend called **GASP!**
- II. it calculates the symbolic integral in seconds
 - \Rightarrow 8 seconds for a 81 piece squared spline (deg. 12) on the SDD-Dataset
 - \Rightarrow up to 2 orders of magnitudes faster than previous WMI-Solvers
- III. the result is a function of λ
 - \Rightarrow compile only once, evaluate cheaply many times

⇒ training independent on constraint complexity



 $q:=\sum_i\lambda_i\prod_j y_j^{lpha_{ij}}$





$$q: \quad \sum_{i} \lambda_{i} \prod_{j} y_{j}^{\alpha_{ij}} \longrightarrow \left(\lambda_{1} \prod_{j} y_{j}^{\alpha_{1j}}, \lambda_{2} \prod_{j} y_{j}^{\alpha_{2j}}, ..., \lambda_{M} \prod_{j} y_{j}^{\alpha_{Mj}}\right)$$





$$q: \quad \sum_{i} \lambda_{i} \prod_{j} y_{j}^{\alpha_{ij}} \longrightarrow \left(\lambda_{1} \prod_{j} y_{j}^{\alpha_{1j}}, \lambda_{2} \prod_{j} y_{j}^{\alpha_{2j}}, ..., \lambda_{M} \prod_{j} y_{j}^{\alpha_{Mj}}\right)$$





































Smet et al., "Neural probabilistic logic programming in discrete-continuous domains", UAI, 2023 24/32





Results

		NN + PAL		NN + GMM		DeepSeaProblog
scene		10 knots	14 knots	K=50	K = 100	-
1	log-like $Pr(\neg arphi)$	$egin{array}{c} -2.08 \\ 0\% \end{array}$	$-2.27\\\mathbf{0\%}$	$-2.64 \approx 21\%$	$\begin{array}{c} -2.83 \\ \approx 20\% \end{array}$	$\begin{array}{c} -3.87 \\ \approx 49\% \end{array}$

Results

		NN + PAL		NN + GMM		DeepSeaProblog
scene		10 knots	14 knots	K=50	K = 100	-
1	log-like $Pr(\neg arphi)$	$egin{array}{c} -2.08 \ 0\% \end{array}$	-2.27 0%	$-2.64 \approx 21\%$	$-2.83 \approx 20\%$	$\begin{array}{c} -3.87 \\ \approx 49\% \end{array}$
2	log-like $Pr(eg arphi)$	-2.23 0%	$\begin{array}{c} -2.09\\0\%\end{array}$	$-2.39 \approx 15\%$	$-2.42 \approx 14\%$	$\begin{array}{c} -3.61 \\ \approx 36\% \end{array}$

Results

		NN + PAL		NN + GMM		DeepSeaProblog
scene		10 knots	14 knots	K=50	K = 100	-
1	log-like $Pr(eg arphi)$	$egin{array}{c} -2.08 \ 0\% \end{array}$	-2.27 0%	$-2.64 \approx 21\%$	$-2.83 \approx 20\%$	$\begin{array}{c} -3.87 \\ \approx 49\% \end{array}$
2	log-like $Pr(\neg arphi)$	-2.23 0%	$\begin{array}{c}-2.09\\0\%\end{array}$	$-2.39 \approx 15\%$	$-2.42 \approx 14\%$	$\begin{array}{c} -3.61 \\ \approx 36\% \end{array}$

are we done?

- I. Scaling to high-dimensional problems
 - \Rightarrow hybridize with approximations
- II. Parallelize sampling
 - \Rightarrow Exact integral can help with inverse-transform sampling
- III. Maximum a Posteriori (MAP)
 - \Rightarrow Picking the right piece to simplify the problem
- IV. Generalize the constraints to curves (e.g. NURBS)
 - \Rightarrow go beyond SMT(\mathcal{LRA})

Chin and Sukumar, "An efficient method to integrate polynomials over polytopes and curved solids", Computer Aided Geometric Design, 2020

I. Scaling to high-dimensional problems \Rightarrow hybridize with approximations

II. Parallelize sampling

- \Rightarrow Exact integral can help with inverse-transform sampling
- III. Maximum a Posteriori (MAP)
 - \Rightarrow Picking the right piece to simplify the problem
- IV. Generalize the constraints to curves (e.g. NURBS)
 - \Rightarrow go beyond $\mathsf{SMT}(\mathcal{LRA})$

Chin and Sukumar, "An efficient method to integrate polynomials over polytopes and curved solids", Computer Aided Geometric Design, 2020

- I. Scaling to high-dimensional problems \Rightarrow hybridize with approximations
- II. Parallelize sampling
 - \Rightarrow Exact integral can help with inverse-transform sampling
- II. Maximum a Posteriori (MAP)
 - \Rightarrow Picking the right piece to simplify the problem
- IV. Generalize the constraints to curves (e.g. NURBS)
 - \Rightarrow go beyond $\mathsf{SMT}(\mathcal{LRA})$

Chin and Sukumar, "An efficient method to integrate polynomials over polytopes and curved solids", Computer Aided Geometric Design, 2020

- I. Scaling to high-dimensional problems \Rightarrow hybridize with approximations
- II. Parallelize sampling
 - \Rightarrow Exact integral can help with inverse-transform sampling
- III. Maximum a Posteriori (MAP)
 - \Rightarrow Picking the right piece to simplify the problem
- IV. Generalize the constraints to curves (e.g. NURBS)
 - \Rightarrow go beyond SMT(\mathcal{LRA})

Chin and Sukumar, "An efficient method to integrate polynomials over polytopes and curved solids", Computer Aided Geometric Design, 2020

- I. Scaling to high-dimensional problems \Rightarrow hybridize with approximations
- II. Parallelize sampling
 - \Rightarrow Exact integral can help with inverse-transform sampling
- III. Maximum a Posteriori (MAP)
 - \Rightarrow Picking the right piece to simplify the problem
- IV. Generalize the constraints to curves (e.g. NURBS)
 - \Rightarrow go beyond SMT(\mathcal{LRA})

Chin and Sukumar, "An efficient method to integrate polynomials over polytopes and curved solids", <u>Computer Aided Geometric Design</u>, 2020

Takeaway

PAL combines:








Additional: Numerical Approximation



The Integral - Detail Decomposition

$$I(\boldsymbol{\lambda}) = \int q(\mathbf{y}'; \boldsymbol{\lambda}) \qquad \mathbb{1}\{\mathbf{y}' \models \varphi\} \ d\mathbf{Y}'$$
$$= \int \left(\sum_{\alpha} \lambda_{\alpha} \mathbf{y}'^{\alpha}\right) \qquad \mathbb{1}\{\mathbf{y}' \models \varphi\} \ d\mathbf{Y}'$$
$$= \sum_{\alpha} \left(\int \lambda_{\alpha} \mathbf{y}'^{\alpha} \qquad \mathbb{1}\{\mathbf{y}' \models \varphi\} \ d\mathbf{Y}'\right) \qquad (1)$$

Benchmark: GASP! vs LATTE



Benchmark: GASP! vs Rejection Sampling



Figure: Left and center: runtime of GASP! and rejection sampling (using $10^{\{2,4,6\}}$ samples) on integrals over 4 random simplices for increasing dimensions. We report the relative error (%) for rejection sampling. <u>Right</u>: Our amortization speed-up compared to rejection-sampling in 2-dimensions for the 12-degree polynomial.