

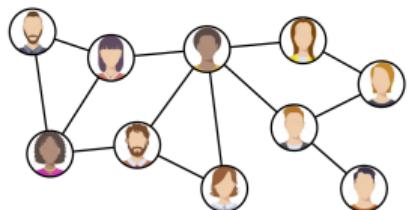
Graph Neural Networks for Data-driven Optimization

Christopher Morris @chrsmrrs

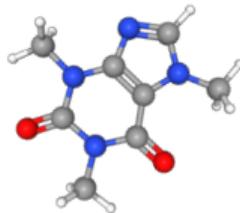
November 1, 2022

RWTH Aachen University

Graph and relational data



Social Networks



Molecules



Knowledge Bases

$$\operatorname{argmin}_x c^T x$$

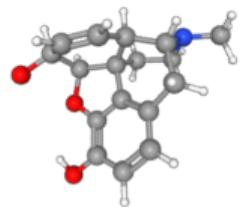
$$Ax \geq 0$$

$$x \in \{0, 1\}^n$$

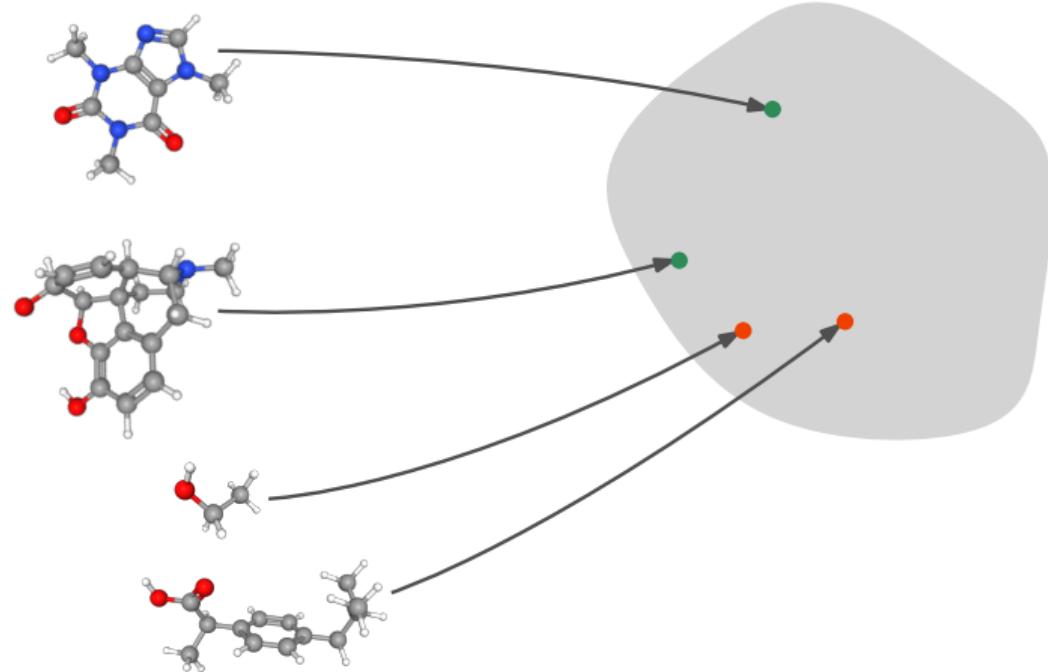
Discrete Optimization

Learning with graphs

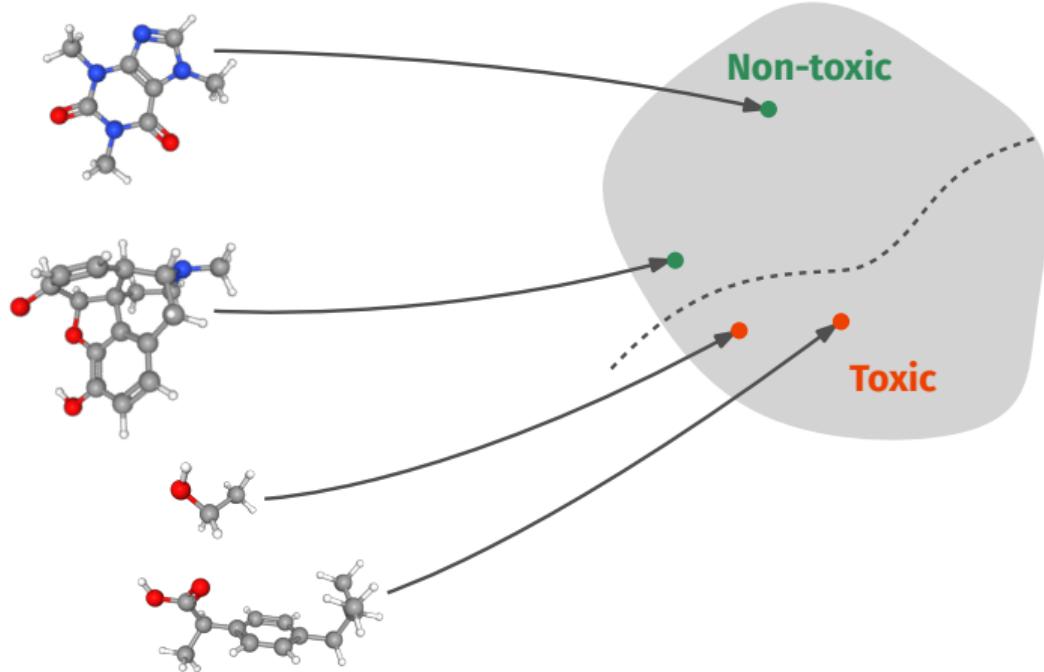
Learning with graphs



Learning with graphs

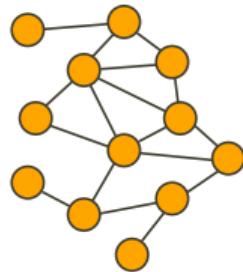


Learning with graphs



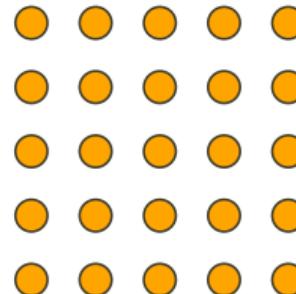
Why is this hard?

Networks are complex!



Network

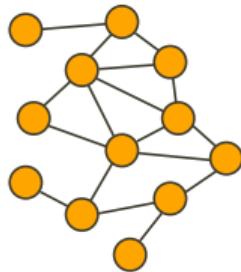
vs.



Image

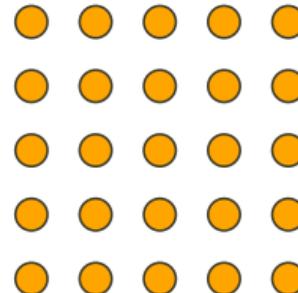
Why is this hard?

Networks are complex!



Network

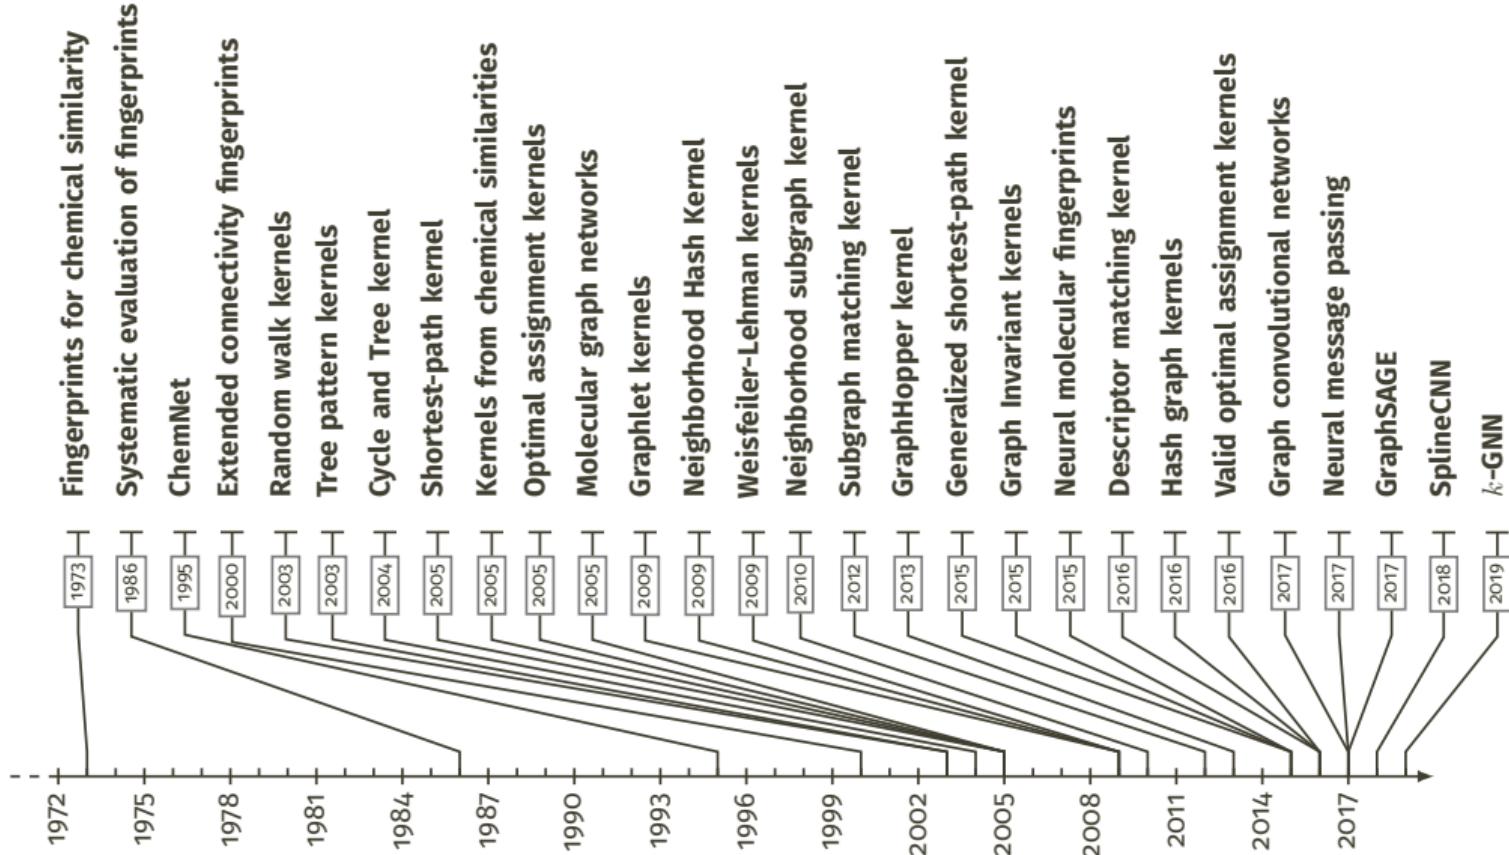
vs.



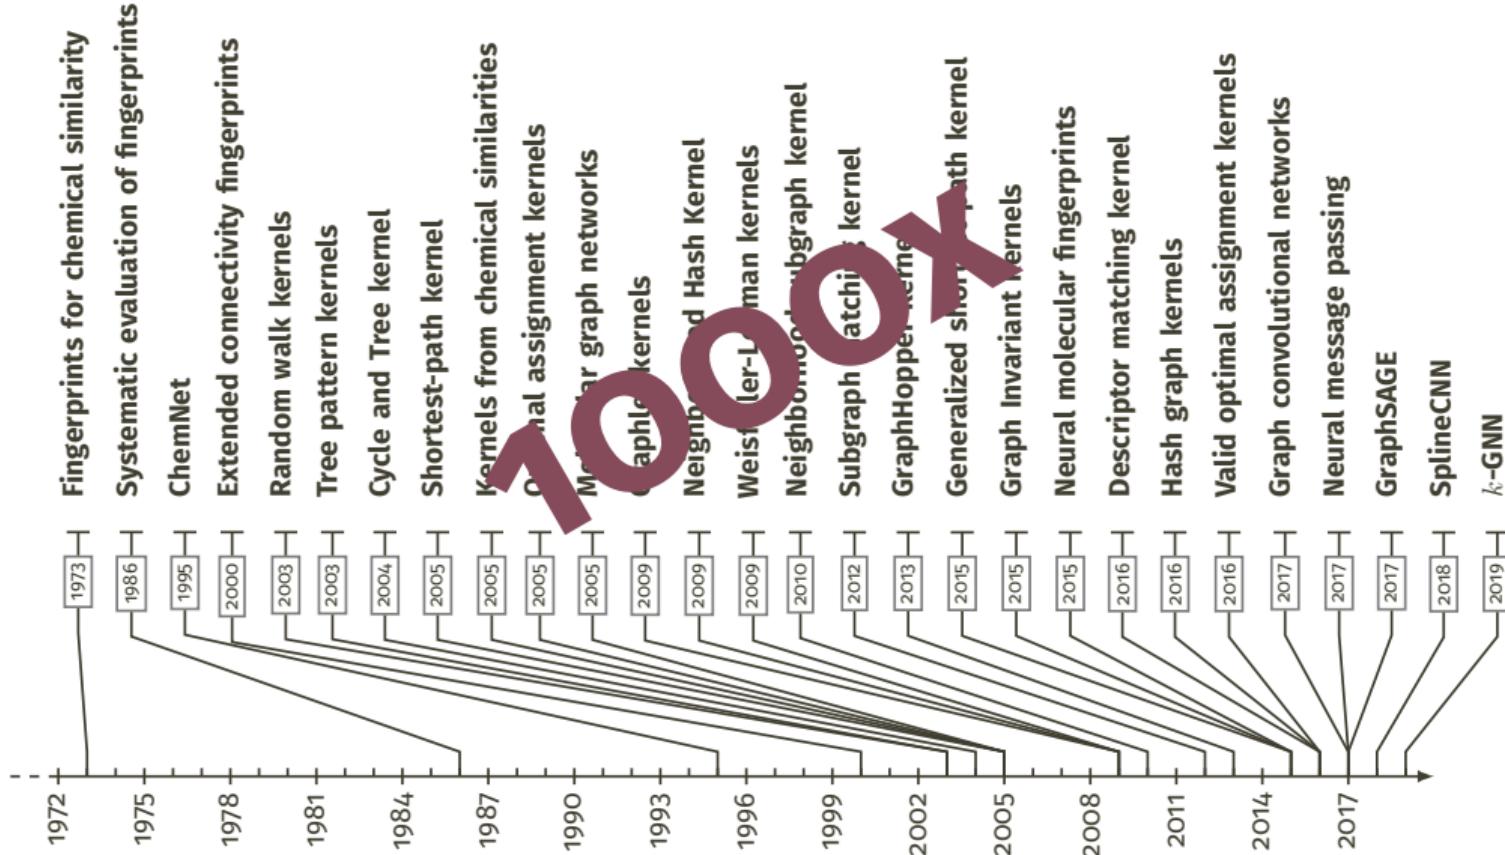
Image

- **Arbitrary size**
- **No fixed node ordering** (invariance to permutation)
- **Hierarchical structure**
- ...

The state of the art?



The state of the art?



The state of the art!



Le bien, la pauvreté l'âge mür la jeunesse

Qui fait, ou l'infortune, ou la felicité

Vérité

A practical theory of machine learning with graphs

Mathematical Foundations

A practical theory of machine learning with graphs

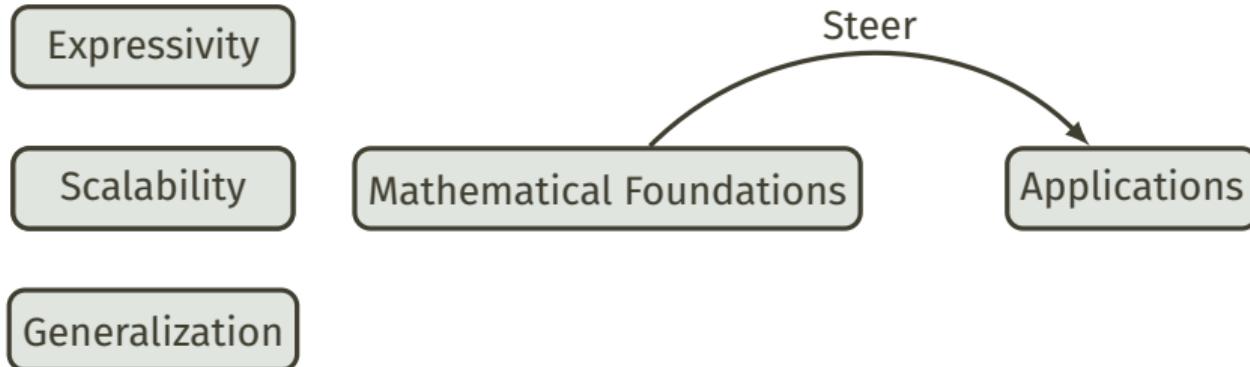
Expressivity

Scalability

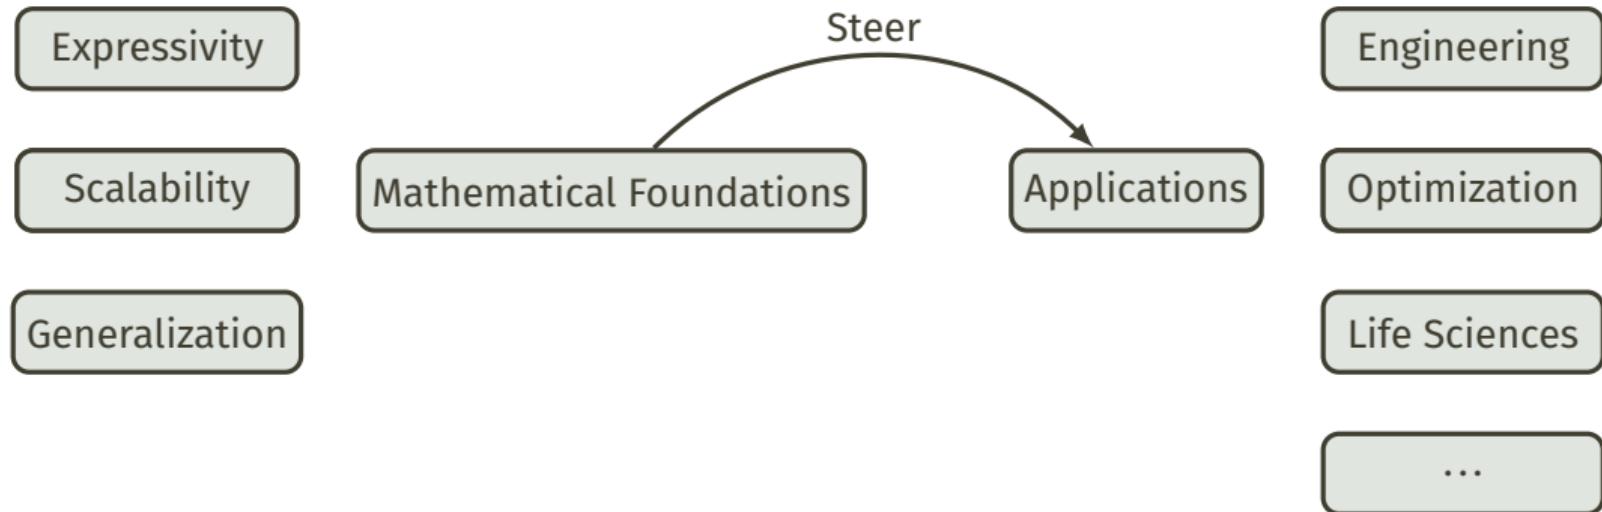
Mathematical Foundations

Generalization

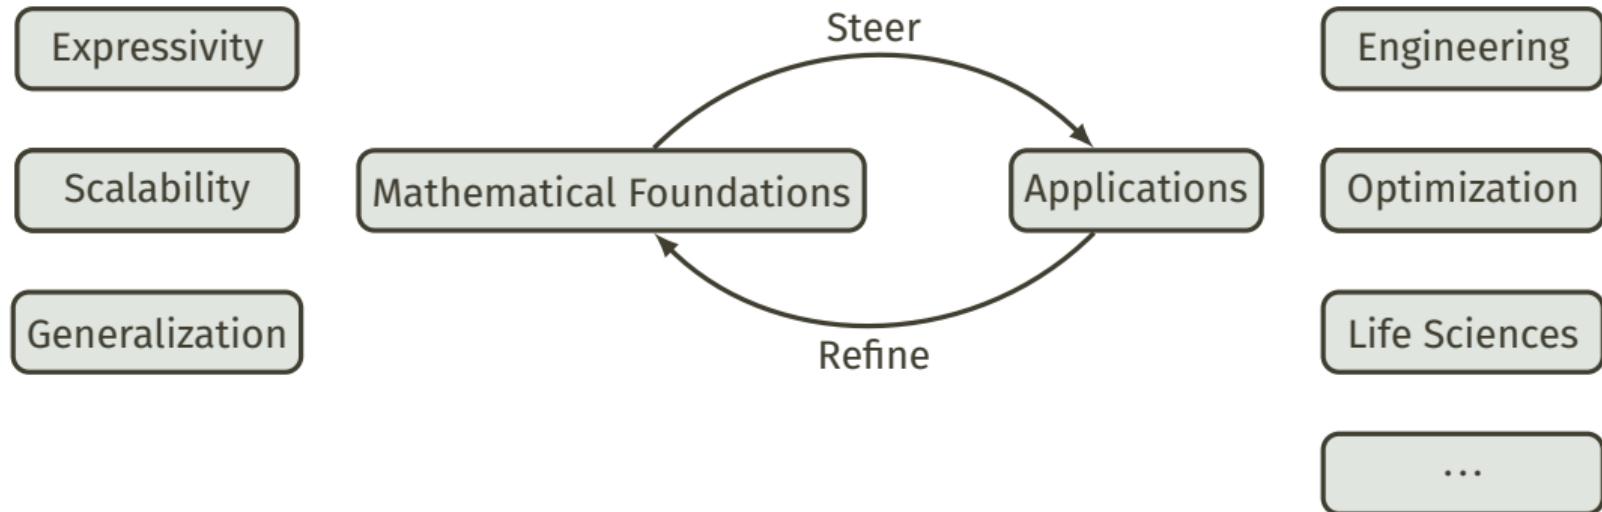
A practical theory of machine learning with graphs



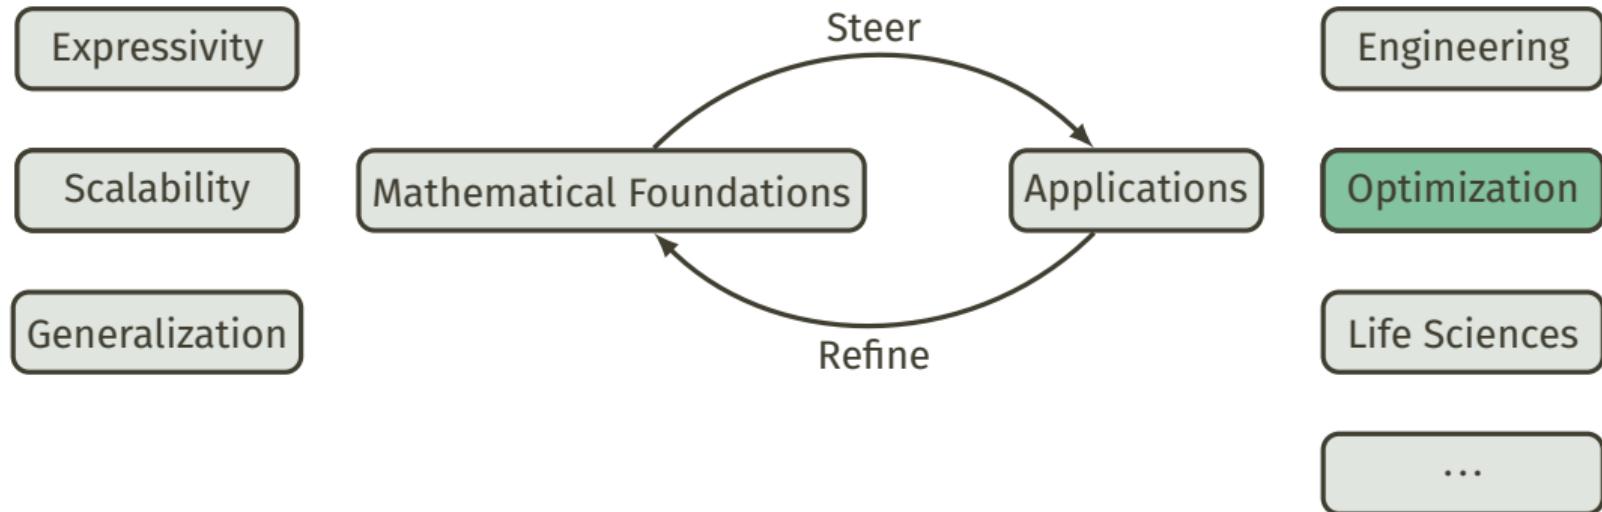
A practical theory of machine learning with graphs



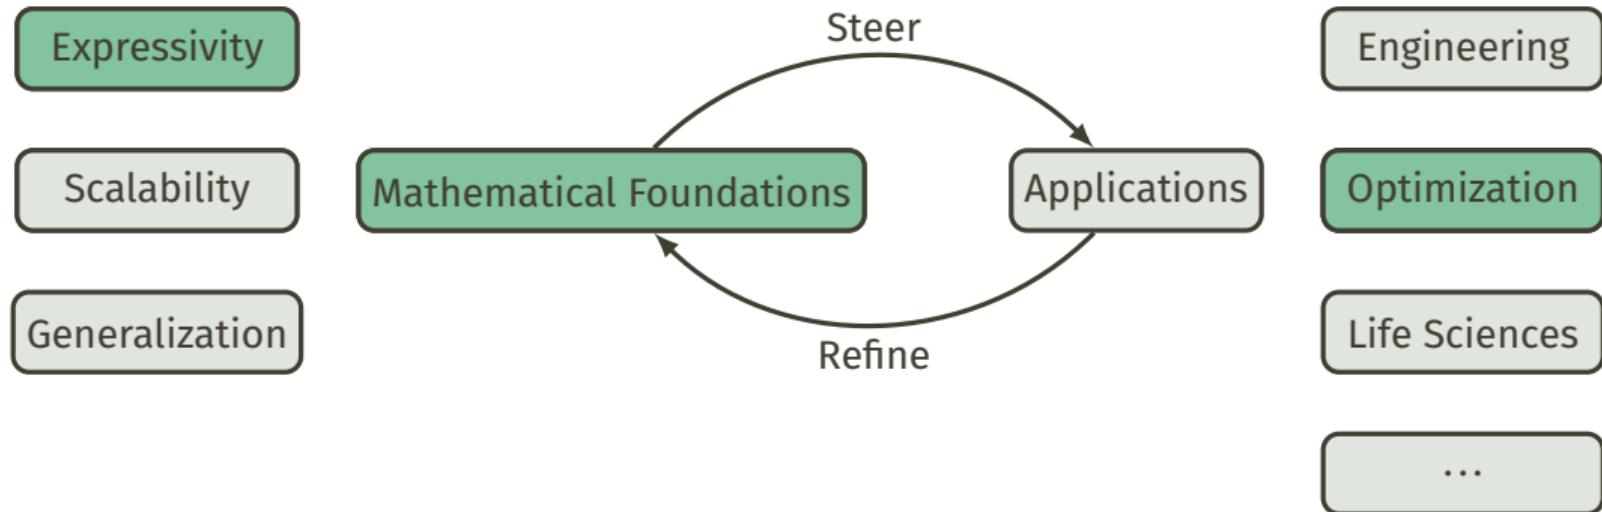
A practical theory of machine learning with graphs



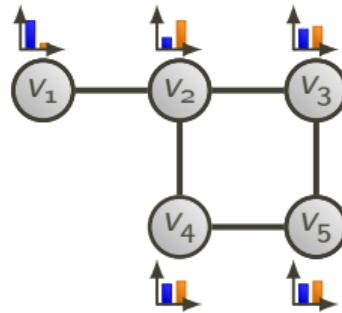
A practical theory of machine learning with graphs



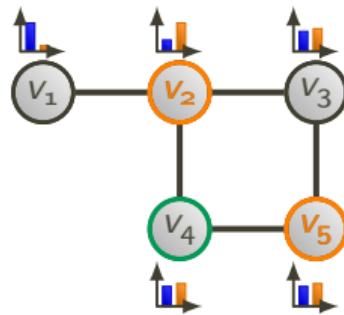
A practical theory of machine learning with graphs



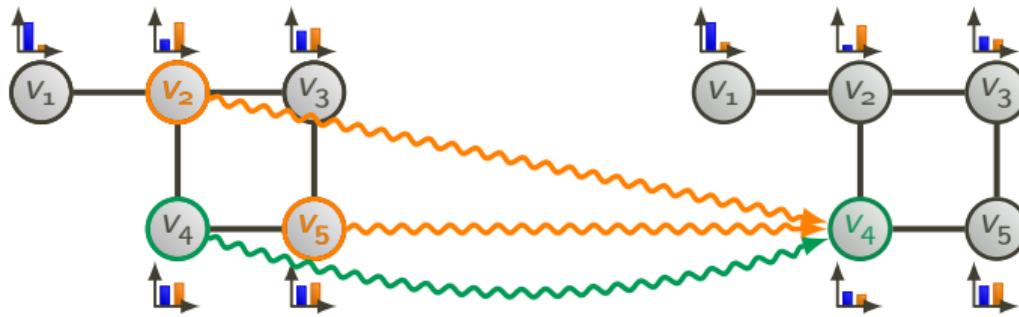
Graph Neural Networks



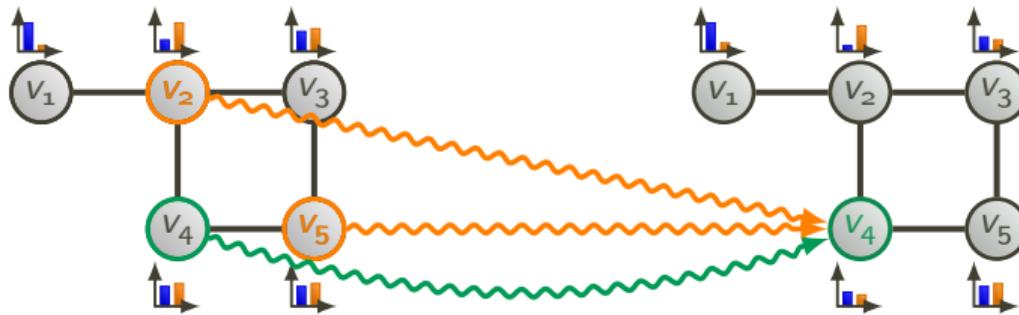
Graph Neural Networks



Graph Neural Networks

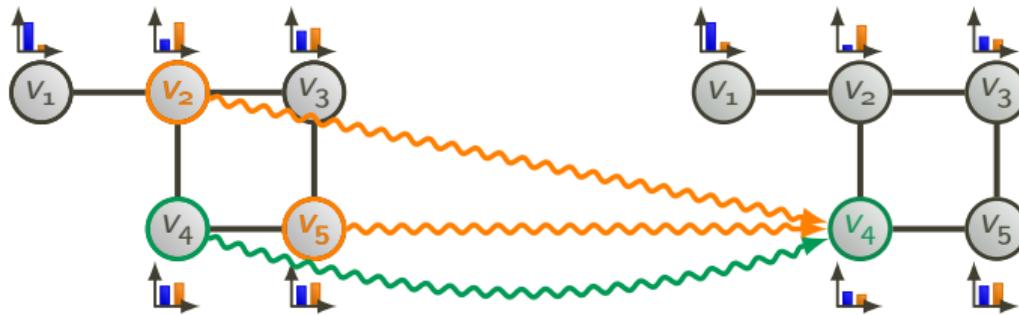


Graph Neural Networks



$$f^{(t)}(v) = \sigma(W_1 f^{(t-1)}(v) + W_2 \sum_{w \in N(v)} f^{(t-1)}(w))$$

Graph Neural Networks



$$f^{(t)}(v) = f_{\text{merge}}^{W_1} \left(f^{(t-1)}(v), f_{\text{aggr}}^{W_2} \left(\{ f^{(t-1)}(w) \mid w \in N(v) \} \right) \right)$$

*How can we use **machine learning** to enhance **combinatorial solvers**?*

Data-driven optimization: Binary Integer Optimization

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

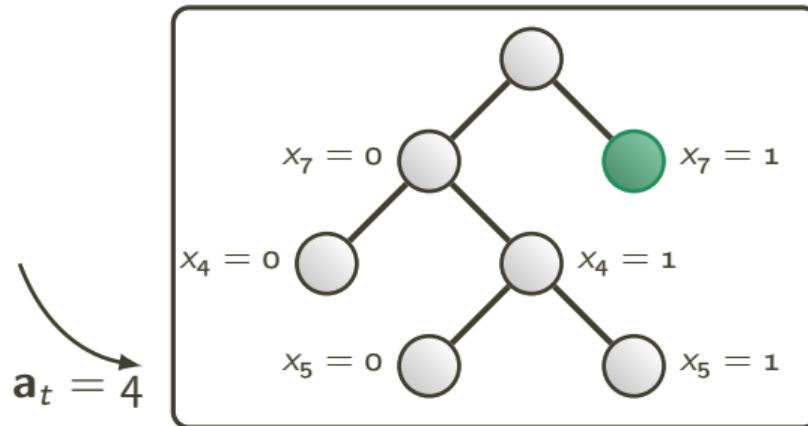
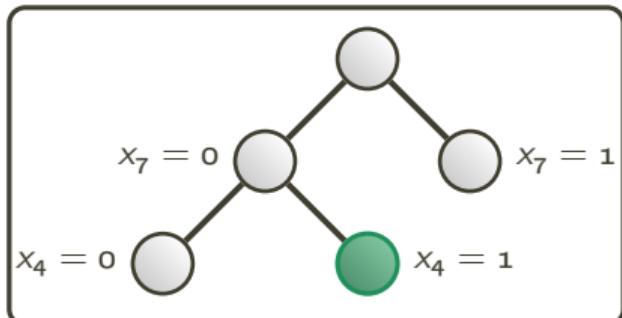
$$x_i \in \{0, 1\}$$

Data-driven optimization: Binary Integer Optimization

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



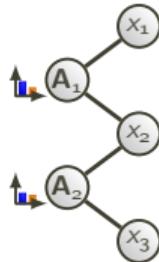
$$\mathbf{a}_t = 4$$

Data-driven optimization: Variable selection via GNNs

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



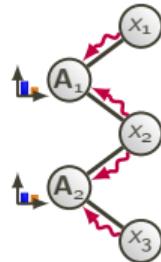
M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. "Exact Combinatorial Optimization with Graph Convolutional Neural Networks". In: *NeurIPS*. 2019

Data-driven optimization: Variable selection via GNNs

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



Update constraints

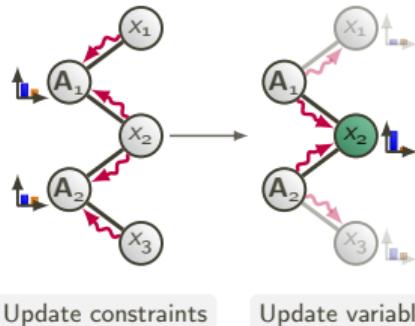
M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. "Exact Combinatorial Optimization with Graph Convolutional Neural Networks". In: *NeurIPS*. 2019

Data-driven optimization: Variable selection via GNNs

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



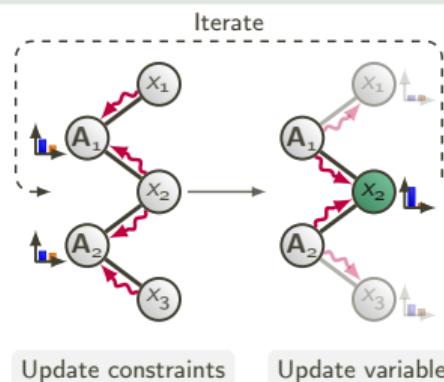
M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. "Exact Combinatorial Optimization with Graph Convolutional Neural Networks". In: *NeurIPS*. 2019

Data-driven optimization: Variable selection via GNNs

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



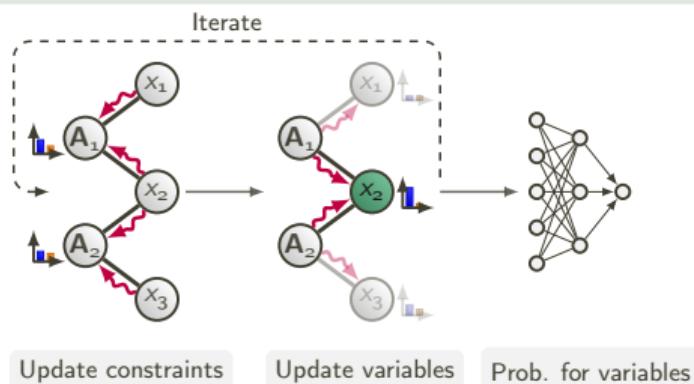
M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. "Exact Combinatorial Optimization with Graph Convolutional Neural Networks". In: *NeurIPS*. 2019

Data-driven optimization: Variable selection via GNNs

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



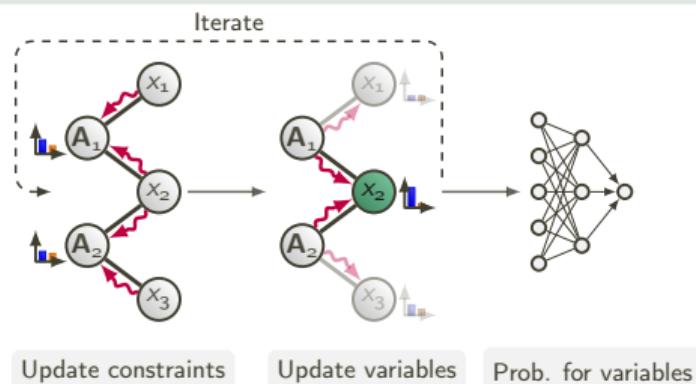
M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. "Exact Combinatorial Optimization with Graph Convolutional Neural Networks". In: *NeurIPS*. 2019

Data-driven optimization: Variable selection via GNNs

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



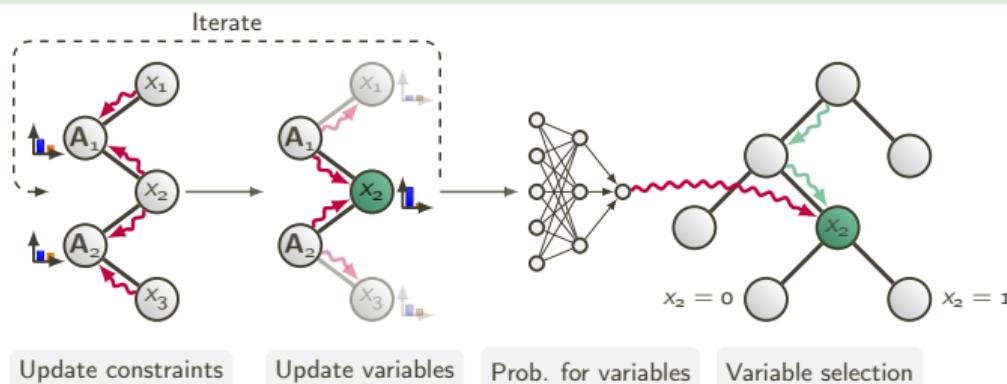
M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. "Exact Combinatorial Optimization with Graph Convolutional Neural Networks". In: *NeurIPS*. 2019

Data-driven optimization: Variable selection via GNNs

$$\arg \min_{\mathbf{x}} \mathbf{c}^T \mathbf{x}$$

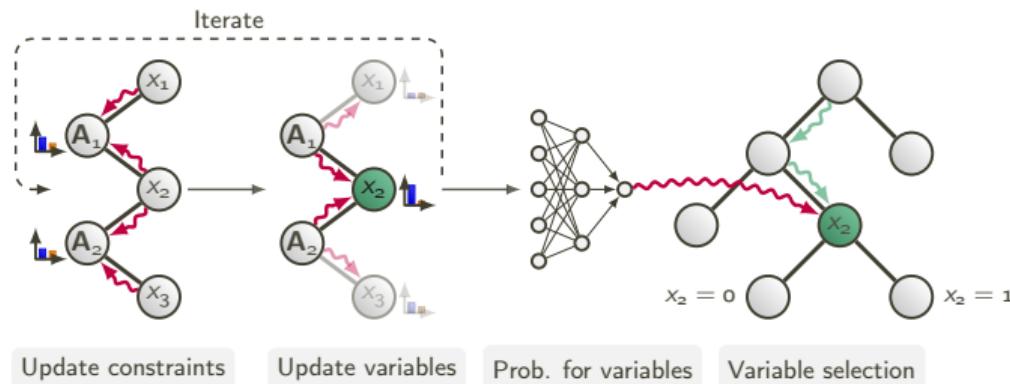
$$\mathbf{A}\mathbf{x} \geq \mathbf{0}$$

$$x_i \in \{0, 1\}$$



M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. "Exact Combinatorial Optimization with Graph Convolutional Neural Networks". In: *NeurIPS*. 2019

Data-driven optimization: Variable selection via GNNs



Data-driven optimization: Variable selection via GNNs

- Aim at simulate **strong branching heuristic**
- Trained in **supervised-fashion** against cross-entropy loss
- Beats default **variable-selection heuristics** of SCIP

Combinatorial Optimization and Reasoning with Graph Neural Networks

Quentin Cappart

Department of Computer Engineering and Software Engineering
Polytechnique Montréal
Montréal, Canada

QUENTIN.CAPPART@POLYMTL.CA

Didier Chételat

CERC in Data Science for Real-Time Decision-Making
Polytechnique Montréal
Montréal, Canada

DIDIER.CHETELAT@POLYMTL.CA

Elias B. Khalil

Department of Mechanical & Industrial Engineering,
University of Toronto
Toronto, Canada

KHALIL@MIE.UTORONTO.CA

Andrea Lodi

Jacobs Technion-Cornell Institute
Cornell Tech and Technion - IIT
New York, USA

ANDREA.LODI@CORNELL.EDU

Christopher Morris

Department of Computer Science
RWTH Aachen University
Aachen, Germany

MORRIS@CS.RWTH-AACHEN.DE

Petar Veličković

DeepMind
London, UK

Data-driven optimization: Variable selection via GNNs

Problem

- Only works for a **single heuristic**
- **Not generic**, e.g., **different models** needed for variable and node selection

*Design **generic approach** such that **once-trained model** can
replace many heuristics!*

Variable Biases I

- Let I be a problem instance with a corresponding **BIP formulation** (A, b, c) .
- Consider **feasible solutions** that are **close** to a **optimal solution** x^* for the instance I :

$$F_\varepsilon^*(I) = \{x \in F_{\text{Int}}(I) : |c^T x^* - c^T x| \leq \varepsilon\}$$

- Near-optimal solutions can be enumerated with MIP solvers' "solution pool" feature.

MIP-GNN: Variable biases

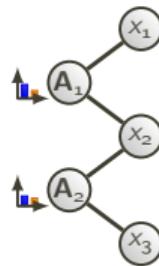
Variable Biases II

The vector of **variable biases** $\bar{\mathbf{b}}(I) \in \mathbb{R}^n$ of I w.r.t. to $F_\varepsilon^*(I)$ is the **component-wise average** over all elements in $F_\varepsilon^*(I)$, namely

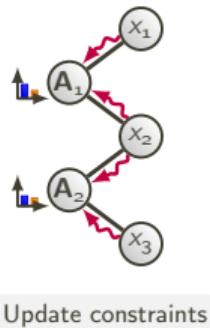
$$\bar{\mathbf{b}}(I) = 1/|F_\varepsilon^*| \sum_{\mathbf{x} \in F_\varepsilon^*(I)} \mathbf{x}.$$

Probability of a variable being 0 or 1 in an optimal solution

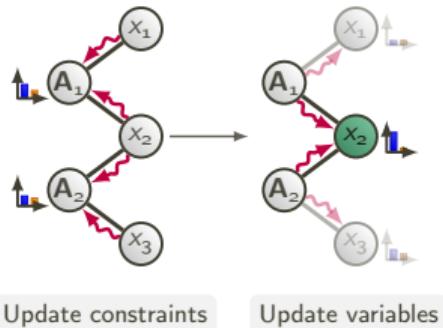
MIP-GNN: Node selection



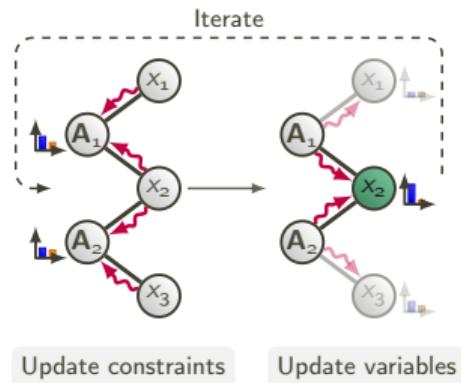
MIP-GNN: Node selection



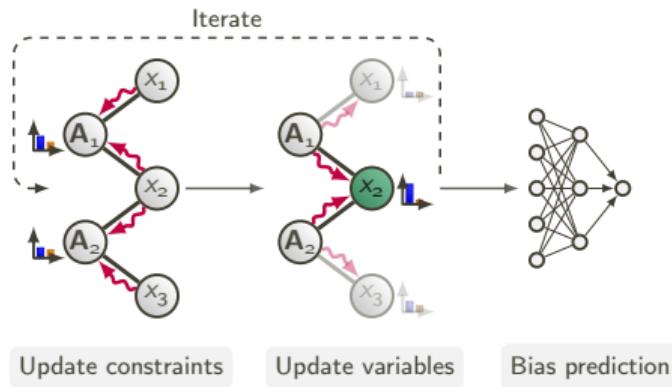
MIP-GNN: Node selection



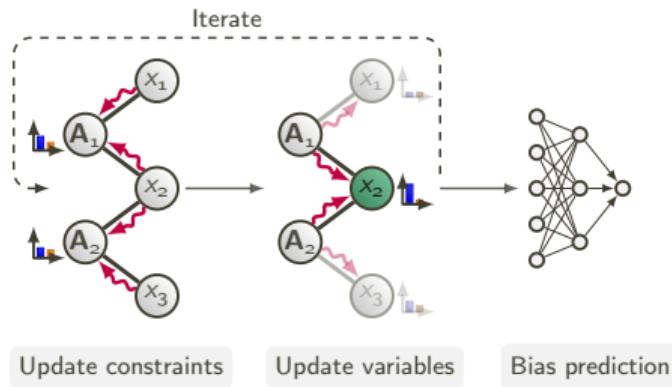
MIP-GNN: Node selection



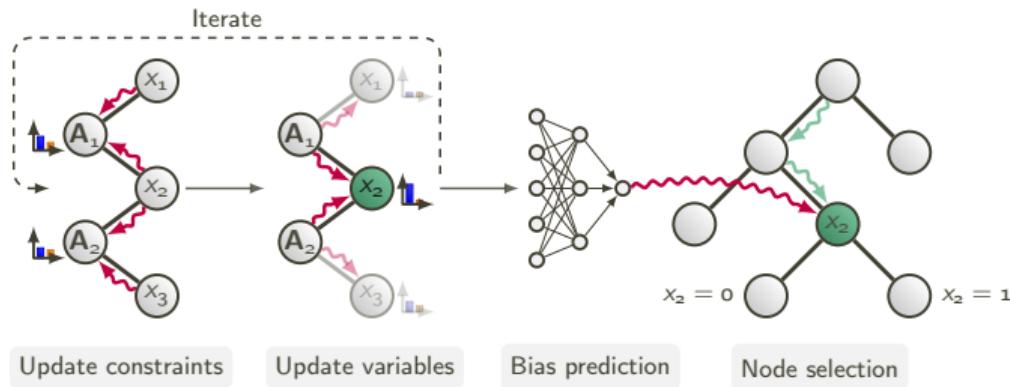
MIP-GNN: Node selection



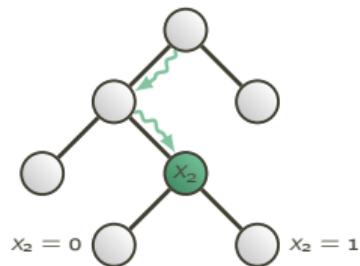
MIP-GNN: Node selection



MIP-GNN: Node selection



MIP-GNN: Node selection



Node selection via bias prediction

- **Confidence score** $\text{score}(\hat{\mathbf{p}}_i) = 1 - |\hat{\mathbf{p}}_i - \lfloor \hat{\mathbf{p}}_i \rfloor|$
- **Node score**

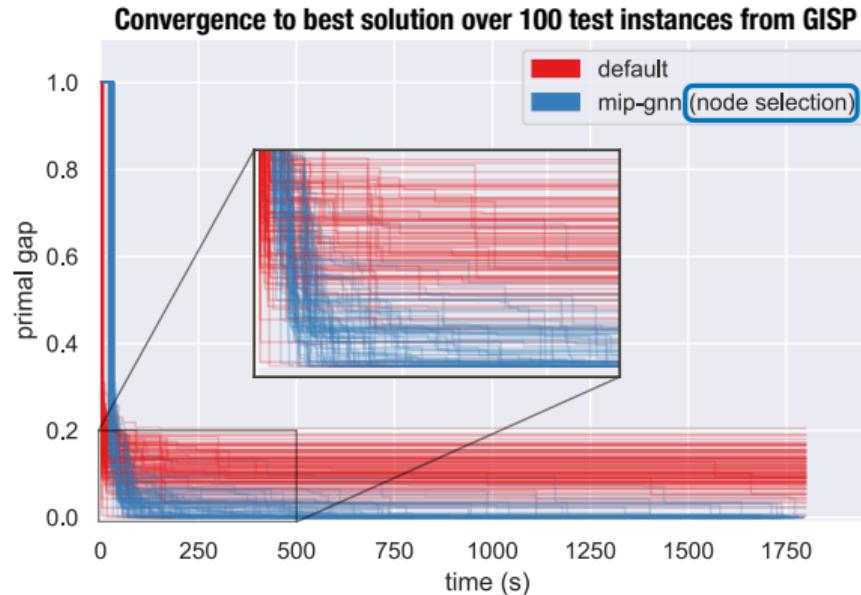
$$\text{node-score}(N; \hat{\mathbf{p}}) = \sum_{i \in \text{fixed-vars}(N)} \begin{cases} \text{score}(\hat{\mathbf{p}}_i), & \text{if } x_i^N = \lfloor \hat{\mathbf{p}}_i \rfloor, \\ 1 - \text{score}(\hat{\mathbf{p}}_i), & \text{otherwise,} \end{cases}$$

MIP-GNN: Experimental results

Experimental setup

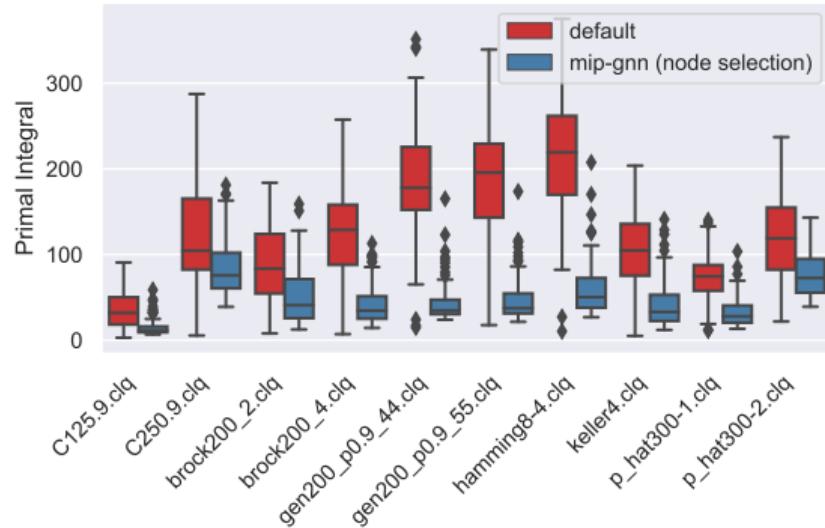
- **Benchmark problems:**
 - Generalized Independent Set Problem (GISP)
- **Solver details:**
 - IBM CPLEX 12.10.0
 - *Full-fledged setup:* Presolve, cuts, and primal heuristics enabled

MIP-GNN: Experimental results



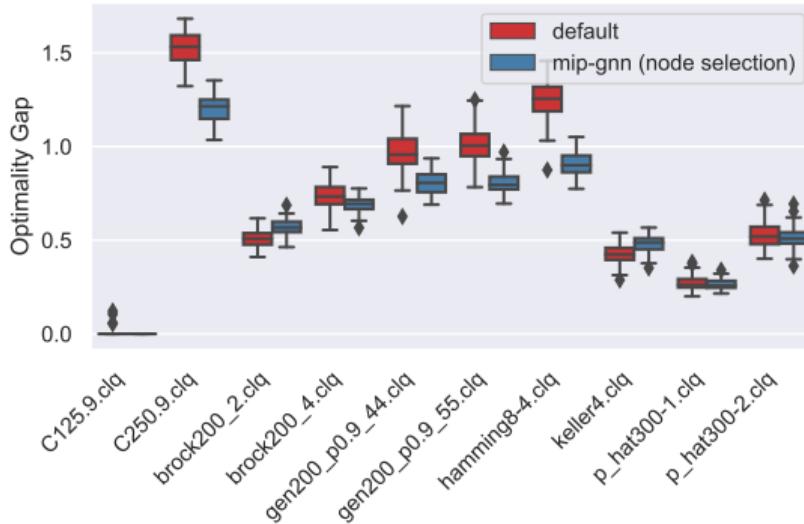
Distribution of **Primal Integrals** for ten problem sets of the Generalized Independent Set Problem, each with 100 unseen test instances.

MIP-GNN: Experimental results



Distribution of **Primal Integrals** for ten problem sets of the Generalized Independent Set Problem, each with 100 unseen test instances.

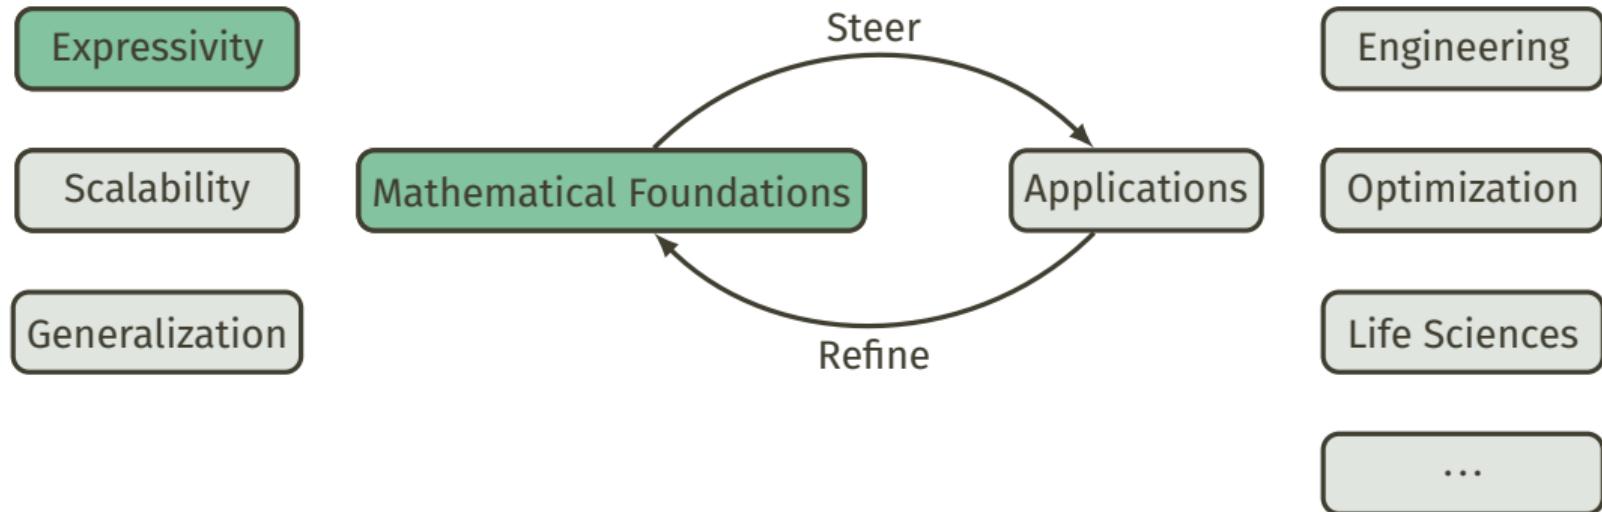
MIP-GNN: Experimental results: Limitations



Distribution of **Optimality Gaps** for ten problem sets of the Generalized Independent Set Problem, each with 100 unseen test instances.

GNNs have severe **limitations** in distinguishing graphs with different structures!

A practical theory of machine learning with graphs

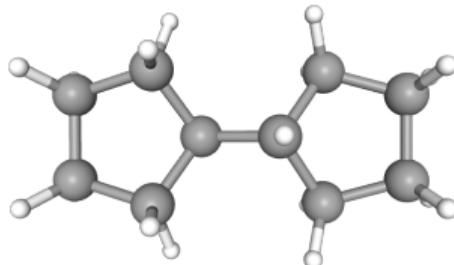


Expressivity of Graph Neural Networks

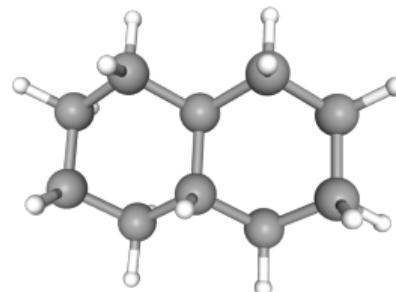
Insight

Any possible **GNN** architecture **misses crucial patterns** in the data!

Expressivity of Graph Neural Networks



(a) Bicyclopentyl



(b) Decalin

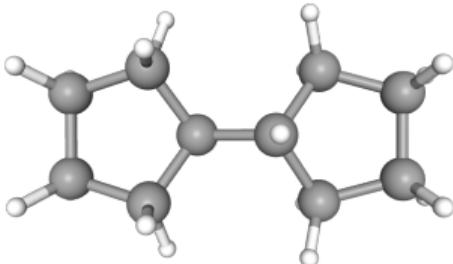
Insight

Any possible **GNN** architecture **misses crucial patterns** in the data!

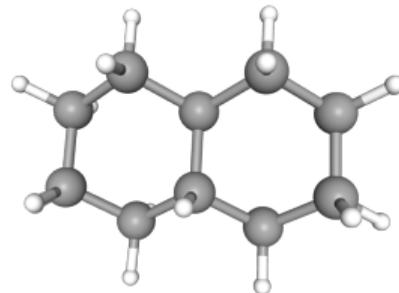
C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. Eric Lenssen, G. Rattan, and M. Grohe. "Weisfeiler and Leman Go Neural: Higher-order Graph Neural Networks". In: AAAI. 2019

C. Morris, G. Rattan, and P. Mutzel. "Weisfeiler and Leman Go Sparse: Higher-order Graph Embeddings". In: NeurIPS. 2020

A hierarchy of more powerful models



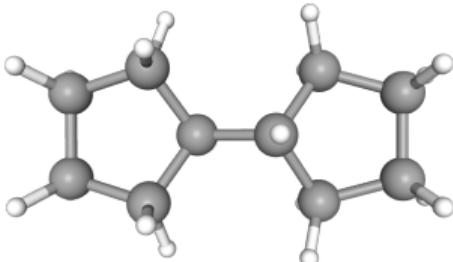
(a) Bicyclopentyl



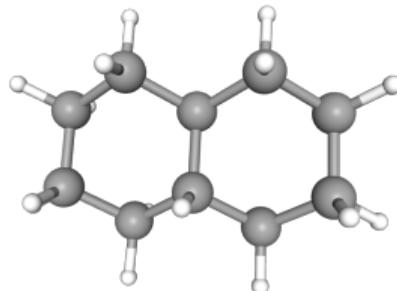
(b) Decalin



A hierarchy of more powerful models



(a) Bicyclopentyl



(b) Decalin



GNNs

2-GNNs

3-GNNs

...

k-GNNs



1-dim. Weisfeiler-Leman algorithm

1-dim. Weisfeiler-Leman algorithm

Heuristic for graph isomorphism testing

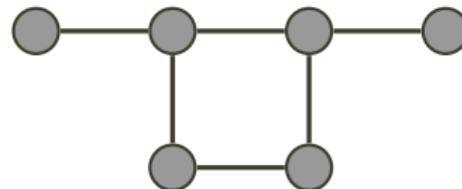
Iteration: Two vertices get **identical colors** iff their **colored neighborhoods are identical**

1-dim. Weisfeiler-Leman algorithm

1-dim. Weisfeiler-Leman algorithm

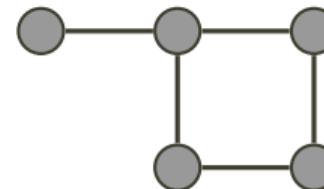
Heuristic for graph isomorphism testing

Iteration: Two vertices get **identical colors** iff their **colored neighborhoods** are **identical**



$$\phi(G_1) = (\quad)$$

(a) G_1



$$\phi(G_2) = (\quad)$$

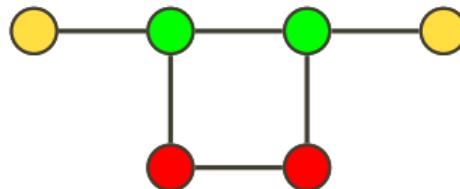
(b) G_2

1-dim. Weisfeiler-Leman algorithm

1-dim. Weisfeiler-Leman algorithm

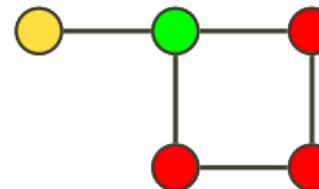
Heuristic for graph isomorphism testing

Iteration: Two vertices get **identical colors** iff their **colored neighborhoods** are **identical**



$$\phi(G_1) = (\textcolor{blue}{1}, \textcolor{red}{2}, \textcolor{blue}{2}, \quad)$$

(a) G_1



$$\phi(G_2) = (\textcolor{blue}{1}, \textcolor{red}{1}, \textcolor{blue}{3}, \quad)$$

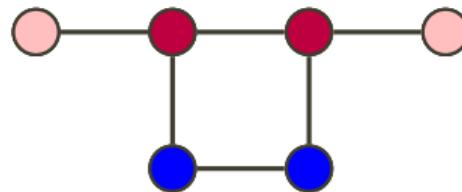
(b) G_2

1-dim. Weisfeiler-Leman algorithm

1-dim. Weisfeiler-Leman algorithm

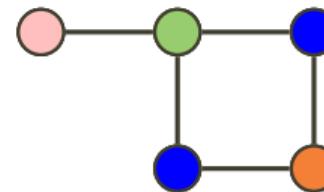
Heuristic for graph isomorphism testing

Iteration: Two vertices get **identical colors** iff their **colored neighborhoods** are **identical**



$$\phi(G_1) = (2, 2, 2, 2, 2, 2, 0, 0)$$

(a) G_1



$$\phi(G_2) = (1, 1, 3, 2, 0, 1, 1, 1)$$

(b) G_2

Relationship between 1-WL and GNN

1-WL coloring

$$c^{(t)}(v) = \text{recolor}\left(c^{(t-1)}(v), \{c^{(t-1)}(w) \mid w \in N(v)\}\right)$$

Relationship between 1-WL and GNN

1-WL coloring

$$c^{(t)}(v) = \text{recolor}\left(c^{(t-1)}(v), \{c^{(t-1)}(w) \mid w \in N(v)\}\right)$$

General form of GNNs

$$h^{(t)}(v) = f_{\text{merge}}^{W_1^{(t)}}\left(h^{(t-1)}(v), f_{\text{aggr}}^{W_2^{(t)}}(\{h^{(t-1)}(w) \mid w \in N(v)\})\right)$$

Relationship between 1-WL and GNN

1-WL coloring

$$c^{(t)}(v) = \text{recolor}\left(c^{(t-1)}(v), \{c^{(t-1)}(w) \mid w \in N(v)\}\right)$$

General form of GNNs

$$h^{(t)}(v) = f_{\text{merge}}^{W_1^{(t)}}\left(h^{(t-1)}(v), f_{\text{aggr}}^{W_2^{(t)}}(\{h^{(t-1)}(w) \mid w \in N(v)\})\right)$$

Theorem (Informal)

GNNs cannot be more expressive than 1-WL in terms of distinguishing non-isomorphic graphs.

C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. Eric Lenssen, G. Rattan, and M. Grohe. "Weisfeiler and Leman Go Neural: Higher-order Graph Neural Networks". In: AAAI. 2019

K. Xu, W. Hu, J. Leskovec, and S. Jegelka. "How Powerful are Graph Neural Networks?" In: ICLR. 2019

Relationship between 1-WL and GNNs

1-WL coloring

$$c^{(t)}(v) = \text{hash}\left(c^{(t-1)}(v), \{c^{(t-1)}(w) \mid w \in N(v)\}\right)$$

General form of GNNs

$$h^{(t)}(v) = f_{\text{merge}}^{W_1^{(t)}}\left(h^{(t-1)}(v), f_{\text{aggr}}^{W_2^{(t)}}\left(\{h^{(t-1)}(w) \mid w \in N(v)\}\right)\right)$$

Relationship between 1-WL and GNNs

1-WL coloring

$$c^{(t)}(v) = \text{hash}\left(c^{(t-1)}(v), \{c^{(t-1)}(w) \mid w \in N(v)\}\right)$$

General form of GNNs

$$h^{(t)}(v) = f_{\text{merge}}^{W_1^{(t)}}\left(h^{(t-1)}(v), f_{\text{aggr}}^{W_2^{(t)}}\left(\{h^{(t-1)}(w) \mid w \in N(v)\}\right)\right)$$

Theorem (Informal)

*There exists a **GNN architecture** and corresponding **weights** such that it has the **same power** as the **1-WL**.*

Relationship between 1-WL and GNNs

Theorem (Informal)

*There exists a **GNN architecture** and corresponding **weights** such that it has the **same power** as the **1-WL**.*

Relationship between 1-WL and GNNs

Theorem (Informal)

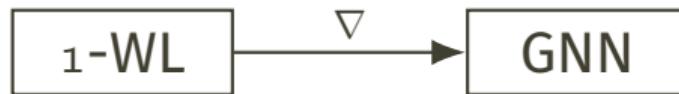
*There exists a **GNN architecture** and corresponding **weights** such that it has the **same power** as the **1-WL**.*



Relationship between 1-WL and GNNs

Theorem (Informal)

*There exists a **GNN architecture** and corresponding **weights** such that it has the **same power** as the **1-WL**.*



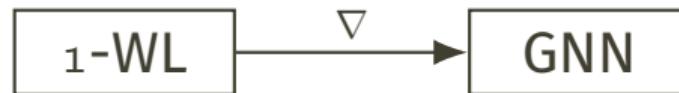
Take Away

GNNs have the **same power** as **1-WL** in distinguishing **non-isomorphic graphs**.

Relationship between 1-WL and GNNs

Theorem (Informal)

*There exists a **GNN architecture** and corresponding **weights** such that it has the **same power** as the **1-WL**.*



Take Away

GNNs have the **same power** as **1-WL** in distinguishing **non-isomorphic graphs**. Limits of 1-WL are **well understood**.

Limits of 1-WL and GNNs

Observation

GNNs cannot distinguish very basic graph properties, e.g.,

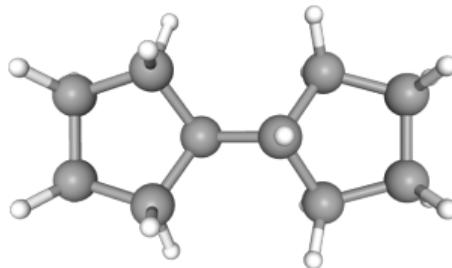
- Cycles of different lengths
- Triangle counts
- Bipartite graphs

Limits of 1-WL and GNNs

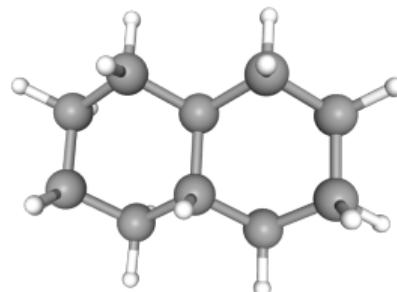
Observation

GNNs cannot distinguish very basic graph properties, e.g.,

- Cycles of different lengths
- Triangle counts
- Bipartite graphs



(a) Bicyclopentyl



(b) Decalin

Weisfeiler and Leman go Machine Learning: The Story so far

Christopher Morris¹, Yaron Lipman², Haggai Maron³, Bastian Rieck⁴, Nils M. Kriege⁵, Martin Grohe⁶, Matthias Fey⁷, and Karsten Borgwardt⁸

¹McGill University and Mila – Quebec AI Institute
²Department of Computer Science and Applied Mathematics, Weizmann Institute of Science

³NVIDIA Research

⁴AIDOS Lab, Institute of AI for Health, Helmholtz Zentrum München
⁵University of Vienna, Faculty of Computer Science, Vienna
⁶Department of Computer Science, RWTH Aachen University
⁷Department of Computer Science, TU Dortmund University
⁸Machine Learning & Computational Biology Lab, Department of Biosystems Engineering, ETH Zürich, and Swiss Institute of Bioinf.

Is the Weisfeiler-Leman algorithm the right yard stick to study GNNs for data-driven optimization?

Data-driven optimization via GNNs

Are GNNs powerful enough to solve LPs?

Multiplicative weight update algorithm

Is

$$Ax \geq b$$

feasible?

S. Arora, E. Hazan, and S. Kale. "The Multiplicative Weights Update Method: a Meta-Algorithm and Applications". In: *Theory of Computing* 8.1 (2012), pp. 121–164

Multiplicative weight update algorithm

Is

$$p^T Ax \geq p^T b$$

feasible?

S. Arora, E. Hazan, and S. Kale. "The Multiplicative Weights Update Method: a Meta-Algorithm and Applications". In: *Theory of Computing* 8.1 (2012), pp. 121–164

Multiplicative weight update algorithm

- ➊ Guess initial weighting, e.g., $p_i = 1/m$

Multiplicative weight update algorithm

- ① **Guess initial weighting**, e.g., $p_i = 1/m$
- ② **Solve** $p^T Ax \geq p^T b$

Multiplicative weight update algorithm

① Guess initial weighting, e.g., $p_i = 1/m$

② Solve $p^T A x \geq p^T b$

③ Compute error signal

$$e_j \leftarrow 1/\rho \left(\sum_{i \in N(j)} A_{ji} \mathbf{x}_i \right) - \mathbf{b}_j$$

Multiplicative weight update algorithm

① Guess initial weighting, e.g., $p_i = 1/m$

② Solve $p^T A x \geq p^T b$

③ Compute error signal

$$e_j \leftarrow 1/\rho \left(\sum_{i \in N(j)} A_{ji} \mathbf{x}_i \right) - \mathbf{b}_j$$

④ Update $w_i \leftarrow (1 - \eta e_j) w_i$

Multiplicative weight update algorithm

① Guess initial weighting, e.g., $p_i = 1/m$

② Solve $p^T A x \geq p^T b$

③ Compute error signal

$$e_j \leftarrow 1/\rho \left(\sum_{i \in N(j)} A_{ji} \mathbf{x}_i \right) - \mathbf{b}_j$$

④ Update $w_i \leftarrow (1 - \eta e_j) w_i$

⑤ Normalize weights $p_j \leftarrow w_j / \Gamma(t)$ for i in $[m]$, where $\Gamma(t) = \sum_{i \in [m]} w_i$

Multiplicative weight update algorithm

① Guess initial weighting, e.g., $p_i = 1/m$

② Solve $p^T A x \geq p^T b$

③ Compute error signal

$$e_j \leftarrow 1/\rho \left(\sum_{i \in N(j)} A_{ji} \mathbf{x}_i \right) - \mathbf{b}_j$$

④ Update $w_i \leftarrow (1 - \eta e_j) w_i$

⑤ Normalize weights $p_j \leftarrow w_j / \Gamma(t)$ for i in $[m]$, where $\Gamma(t) = \sum_{i \in [m]} w_i$

⑥ Update solution $\bar{\mathbf{x}} \leftarrow \bar{\mathbf{x}} + \mathbf{x}$

Multiplicative weight update algorithm

Theorem

Given a BIP $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$, its corresponding relaxation \hat{I} , and $\varepsilon > 0$, the MWU with

$$T = \left\lceil \frac{4\rho \ln(m)}{\varepsilon^2} \right\rceil,$$

outputs an ε -feasible solution for \hat{I} or determines that \hat{I} is not feasible.

GNNs and the multiplicative weight update algorithm

Theorem (Informal)

There exists a GNN architecture and corresponding weight assignments such it can simulate the MWU algorithm.

Data-driven optimization via GNNs

Can GNNs simulate more nature LP solver?

Conclusion

- ① MIP-GNN framework for data-driven optimization
- ② Limitations of GNNs
- ③ Preliminary results to better understand capabilities of GNNs for data-driven optimization

References i

-  S. Arora, E. Hazan, and S. Kale. “The Multiplicative Weights Update Method: a Meta-Algorithm and Applications”. In: *Theory of Computing* 8.1 (2012), pp. 121–164.
-  V. Arvind, J. Köbler, G. Rattan, and O. Verbitsky. “On the Power of Color Refinement”. In: *International Symposium on Fundamentals of Computation Theory*. 2015, pp. 339–350.
-  Q. Cappart, D. Chételat, E. Khalil, A. Lodi, C. Morris, and P. Velickovic. “Combinatorial optimization and reasoning with graph neural networks”. In: *International Joint Conference on Artificial Intelligence*. IJCAI, 2021.
-  M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. “Exact Combinatorial Optimization with Graph Convolutional Neural Networks”. In: *NeurIPS*. 2019.

References ii

-  E. Khalil, C. Morris, and A. Lodi. "MIP-GNN: A data-driven framework for guiding combinatorial solvers". In: AAAI. 2022.
-  C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. Eric Lenssen, G. Rattan, and M. Grohe. "Weisfeiler and Leman Go Neural: Higher-order Graph Neural Networks". In: AAAI. 2019.
-  C. Morris, G. Rattan, and P. Mutzel. "Weisfeiler and Leman Go Sparse: Higher-order Graph Embeddings". In: NeurIPS. 2020.
-  B. Weisfeiler. *On Construction and Identification of Graphs*. Lecture Notes in Mathematics, Vol. 558. Springer, 1976.
-  K. Xu, W. Hu, J. Leskovec, and S. Jegelka. "How Powerful are Graph Neural Networks?" In: ICLR. 2019.