

Combinatorial optimization and reasoning with GNNs: **Algorithmic reasoning**

Cappart, Chételat, Khalil, Lodi, Morris, Veličković

Guide to reading

Overview of current research in **neural algorithm reasoning**

Don't worry about understanding each and every direction 😊

The goal should be to get a broad sense of the field

Neural-algorithmic alignment

Many ways to attack combinatorial optimization with DL:

Why GNNS?

GNNs can execute poly-time **dynamic programming** algs

E.g. "Graph Neural Networks are Dynamic Programmers," Dudzik and Veličković, NeurIPS'22

Paradigm from which many poly-time algs can be constructed

Extrapolation

(G)NNs are traditionally powerful at ***interpolation***

- i.e., strong performance when test distribution \approx training distribution

(G)NNs typically struggle at ***extrapolation***

- i.e., evaluated out of distribution
- E.g., increasing number of nodes in input graph

Neural algorithmic alignment:

How to construct algorithmic reasoners that extrapolate?

Extrapolation

Key aspects of Veličković et al. [ICLR'20] enabling extrapolation:

1. Using the **encode-process-decode** framework
2. Favoring the **max-aggregation** function
3. Supervising with ground-truth alg's **execution traces**
4. Executing **multiple** related algorithms

Executing multiple related algorithms

The CLRS Algorithmic Reasoning Benchmark

Petar Veličković¹ Adrià Puigdomènech Badia¹ David Budden¹
Razvan Pascanu¹ Andrea Banino¹ Misha Dashevskiy¹ Raia Hadsell¹ Charles Blundell¹

Includes:

- Sorting
- Searching
- Dynamic programming
- Graph algorithms
- String algorithms
- Geometric algorithms

Outline

1. Introduction
- 2. Understanding max-aggregation**
3. Reasoning on natural inputs

Extrapolation error

[Xu et al., ICLR'21]

- $f: \mathcal{X} \rightarrow \mathbb{R}$ is a model trained on $\{(x_i, y_i)\}_{i=1}^n \subset \mathcal{D}$
 $y_i = g(x_i)$ for some ground-truth function g
- \mathcal{P} is a **distribution over $\mathcal{X} \setminus \mathcal{D}$**
- $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is a loss function
- **Extrapolation error:** $\mathbb{E}_{x \sim \mathcal{P}} [\ell(f(x), g(x))]$

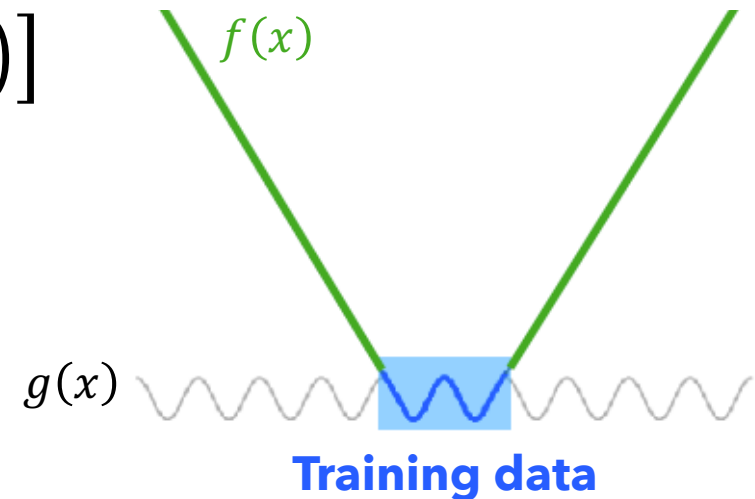
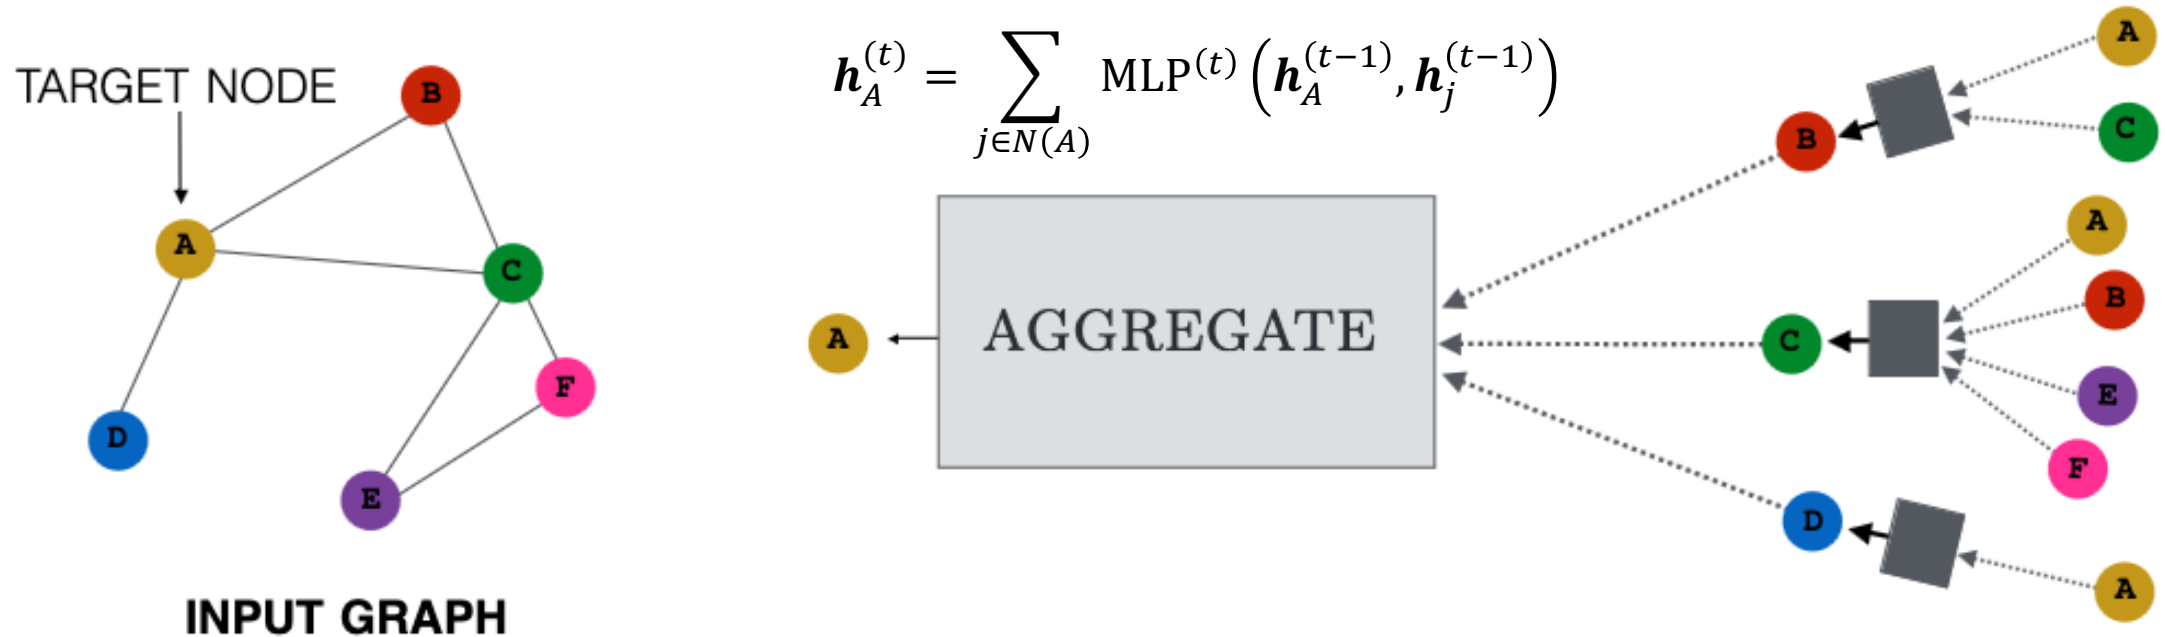


Figure by Xu et al., ICLR'21

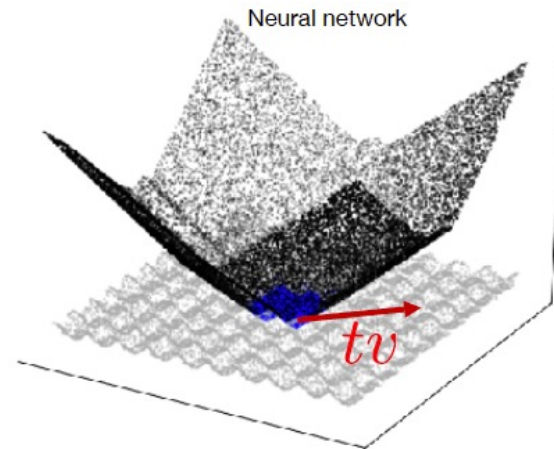
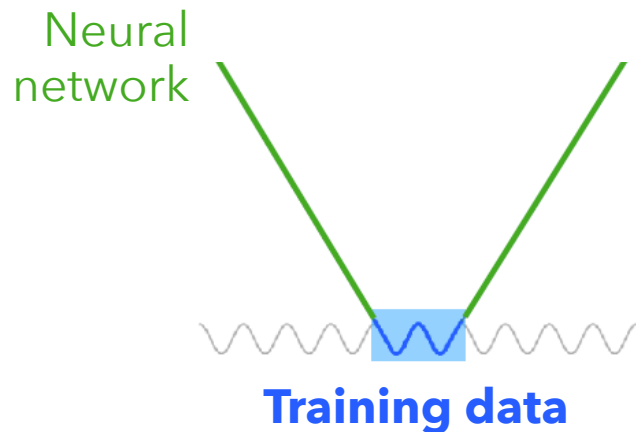
Aggregation functions



ReLU MLP extrapolate linearly

Theorem [Xu et al., ICLR'21, informal]:

- Let f be a 2-layer ReLU MLP trained with gradient descent
- Along any direction $\mathbf{v} \in \mathbb{R}^d$, f approaches a **linear** function
- I.e., let $\mathbf{x} = t\mathbf{v}$. Then $f(\mathbf{x} + h\mathbf{v}) - f(\mathbf{x}) \rightarrow \beta_{\mathbf{v}}h$ at a rate $O\left(\frac{1}{t}\right)$



Implications for GNNs

Shortest path: $x_i^{(t)} = \min \left\{ x_i^{(t-1)}, \min_{(j,i) \in E} x_j^{(t-1)} + e_{ji}^{(t-1)} \right\}$

GNN: $\mathbf{h}_i^{(t)} = \sum_{j \in N(i)} \text{MLP} \left(\mathbf{h}_i^{(t-1)}, \mathbf{h}_j^{(t-1)} \right)$

MLP must learn a **non-linearity**

Implications for GNNs

Shortest path: $x_i^{(t)} = \min \left\{ x_i^{(t-1)}, \min_{(j,i) \in E} x_j^{(t-1)} + e_{ji}^{(t-1)} \right\}$

GNN:


$$\mathbf{h}_i^{(t)} = \sum_{j \in N(i)} \text{MLP} \left(\mathbf{h}_i^{(t-1)}, \mathbf{h}_j^{(t-1)} \right)$$

GNN 2:

$$\mathbf{h}_i^{(t)} = \max_{j \in N(i)} \text{MLP} \left(\mathbf{h}_i^{(t-1)}, \mathbf{h}_j^{(t-1)} \right)$$

Predicting shortest path predecessor:

[Veličković et al. ICLR'20]



Outline

1. Introduction
2. Understanding max-aggregation
- 3. Reasoning on natural inputs**

Key question

Key question in neural algorithmic alignment:

If we're just teaching a NN to **imitate** a classical algorithm...

Why not just run that algorithm?

Why use GNNs for algorithm design?

Classical algorithms are designed with **abstraction** in mind
Enforce their inputs to conform to stringent preconditions

However, we design algorithms to solve **real-world** problems!

Our **goals** can be at odds with our **methods**
*Example: Harris and Ross's ['55] study of the **max flow** problem*

Original max flow study

SECRET

U. S. AIR FORCE
PROJECT RAND
RESEARCH MEMORANDUM

FUNDAMENTALS OF A METHOD FOR EVALUATING
RAIL NET CAPACITIES (U)

T. E. Harris
F. S. Ross

RM-1573

October 24, 1955

Copy No. 137

This material contains information affecting the national defense of the United States within the meaning of the espionage laws, Title 18 U.S.C., Secs. 793 and 794, the transmission or the revelation of which in any manner to an unauthorized person is prohibited by law.

SUMMARY

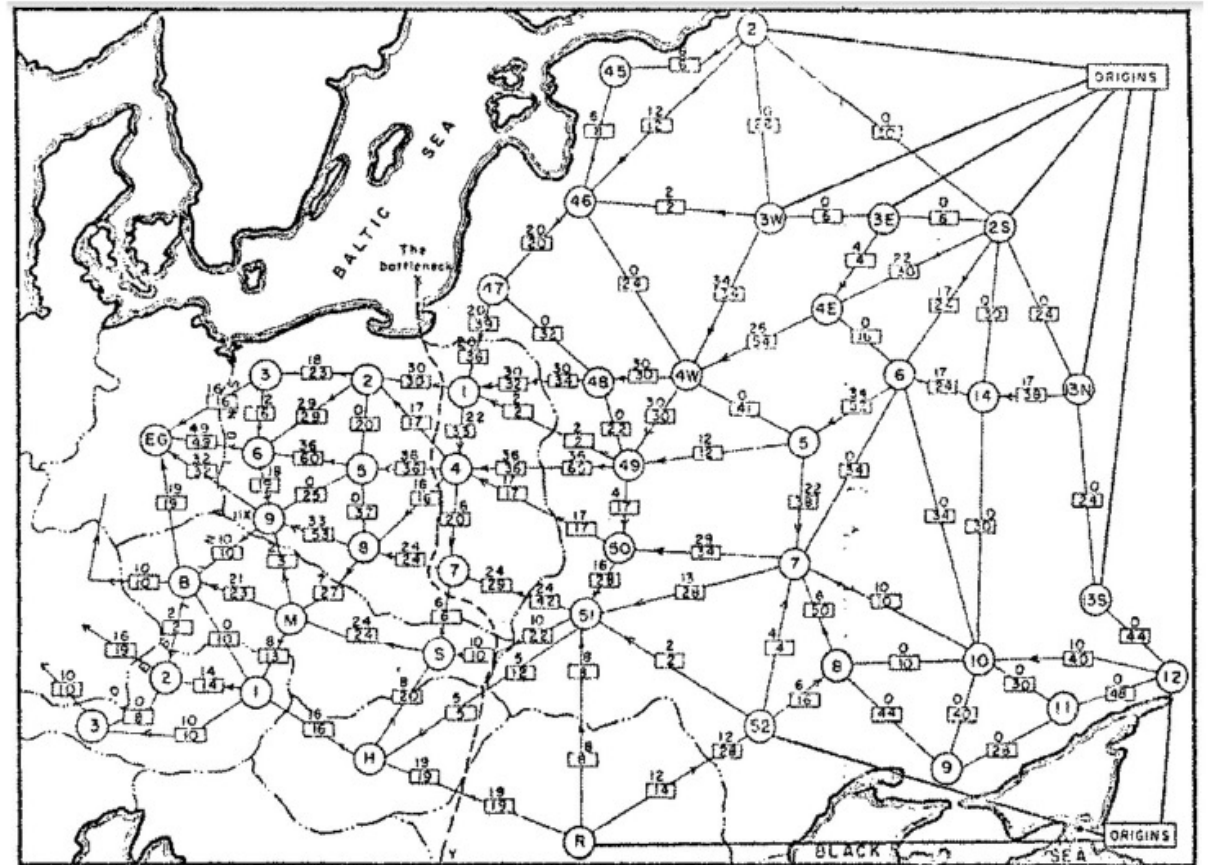
Air power is an effective means of interdicting an enemy's rail system, and such usage is a logical and important mission for this Arm.

As in many military operations, however, the success of interdiction depends largely on how complete, accurate, and timely is the commander's information, particularly concerning the effect of his interdiction-program efforts on the enemy's capability to move men and supplies. This information should be available at the time the results are being achieved.

The Warsaw Pact railway network

Goal: Find the "bottleneck"
i.e., the **minimum cut**

Equivalent to **max flow**



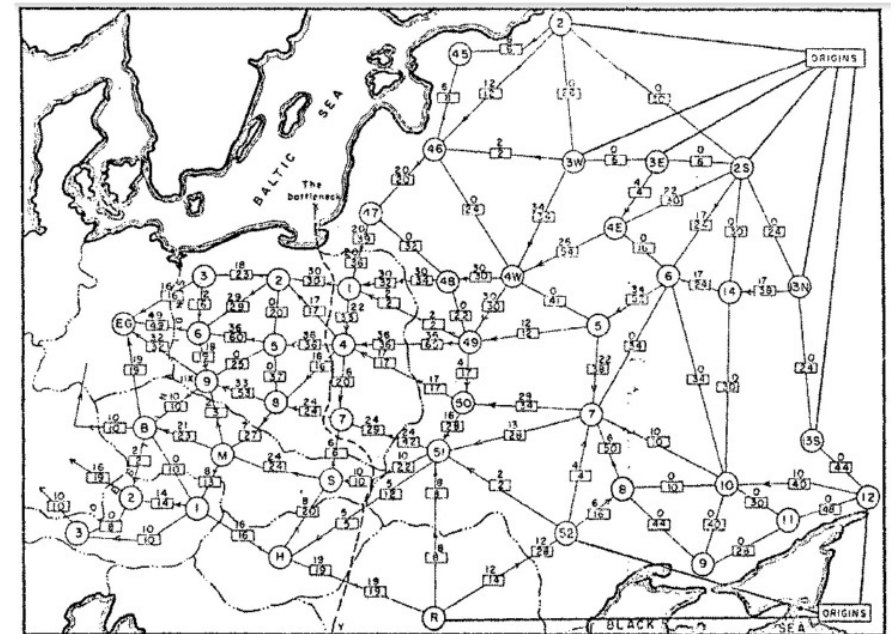
Abstractification \Rightarrow information loss

II. THE ESTIMATING OF RAILWAY CAPACITIES

The evaluation of both railway system and individual track capacities is, to a considerable extent, an art. The authors know of no tested mathematical model or formula that includes all of the variations and imponderables that must be weighed.* Even when the individual has been closely associated with the particular territory he is evaluating, the final answer, however accurate, is largely one of judgment and experience.

Abstractification \Rightarrow information loss

- Attaches a **single, scalar capacity** to an entire railway system
- Ignores a **wealth of information** from the underlying system



Amazon Last Mile Routing Challenge

“Important gap between **theoretical** route planning and **real-life** route execution”

...“In real-life operations, the quality of a route is not exclusively defined by its **theoretical length, duration, or cost**”


Many factors affect whether a driver “can **effectively, safely and conveniently** execute the planned route under real-life conditions.”

Abstractifying the core problem

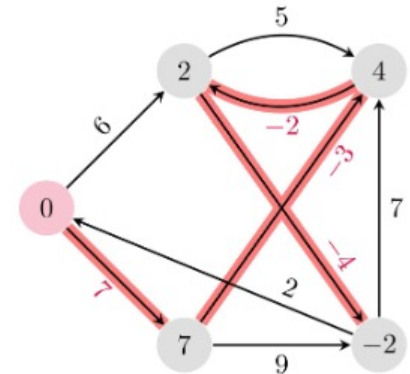
- Assume we have real-world inputs
 - ...but algorithm only admits abstract inputs
- Could try **manually** converting from one input to another



Natural inputs



Abstract inputs





Abstract outputs

Attacking the core problem


- Alternatively, **replace** human feature extractor with NN
 - Still apply same combinatorial algorithm
- Issue: algorithms typically perform **discrete optimization**
 - Doesn't play nicely with **gradient-based** optimization of NNs



Natural inputs



Abstract inputs



Abstract outputs

Algorithmic bottleneck

Second (more fundamental) issue: **data efficiency**

- Real-world data is often incredibly rich
- We still have to compress it down to scalar values


The algorithmic solver commits to using this scalar

Assumes it is perfect!

If there's insufficient training data to estimate the scalars:


- Alg will give a **perfect solution**
- ...but in a **suboptimal environment**

Neural algorithmic pipeline



1. On abstract inputs, learn encode-process-decode functions


Neural algorithmic pipeline



Processor P :


1. Is aligned with computations of target algorithm
2. Admits useful gradients
3. Operates over high-dim latent space (better use of data)

Neural algorithmic pipeline




2. Set up encode-decode functions for natural inputs/outputs

Neural algorithmic pipeline



3. Learn parameters using loss that compares $\tilde{g} \left(P \left(\tilde{f}(x) \right) \right)$ to y

Neural algorithmic pipeline



Note: Keep P frozen while learning \tilde{f} and \tilde{g}

Overview

Survey: broad overview of **neural algorithmic reasoning**

Among other topics, covers:

- Why **max-aggregation** allows for extrapolation
 - MLPs extrapolate linearly
 - Must hard-code algorithmic non-linearities
- The **neural algorithmic pipeline**
 - Allows us to solve natural, real-world instances