Compositional Reasoning with Diffusion Models for Out-of-Distribution Generalization

Research in Data Science (6 KP)

Michal Tešnar
supervisors: Joël Mathys, Prof. Dr. Roger Wattenhofer

ETH Zürich

February 2, 2026

Context: Neural Algorithmic Reasoning

- Compression, memorization: generalization in distribution.
- Learn to discover and adapt algorithms.
- Test: algorithmic tasks, combinatorial puzzles OOD.

Figure: Deep Thought Computer[a]

[a] https://www.turbosquid.com/de/3d-models/3d-deep-thought-super-computer-scifi-gold-model-1813871

Previous Work: Inspiration

- Neural Algorithmic Reasoning (Veličković and Blundell, 2021):
  - formalization neural algorithms
  - Usage of GNNs (Veličković et al., 2019)
  - CLRS Benchmark (Veličković et al., 2022)
- Diffusion solvers for combinatorial graph optimization problems (Sun and Yang, 2023; Zhao et al., 2024).

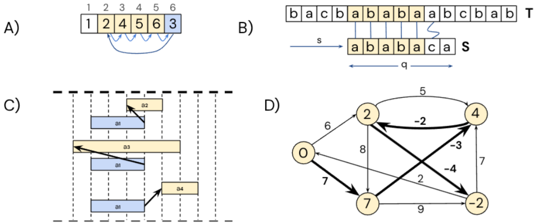

Figure 1. Example of four algorithms within CLRS-30. A) insertion sort; B) string matching; C) greedy task scheduling; D) shortest paths.

Problem: Tents Puzzle (Estermann et al., 2024)

- NxN grid: 'empty', 'grass', 'tent', 'tree' (also 'violated')
- 'tree' given, each 'tent' orthogonally adjacent to 'tree' (unique solution)
- # trees per column/row specified by numbers
- Metric: % fully solved puzzles from test set

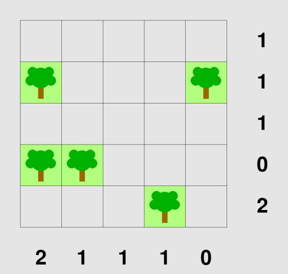

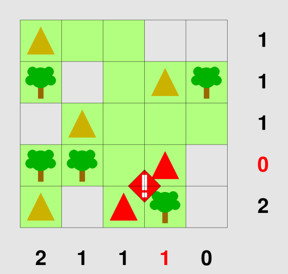

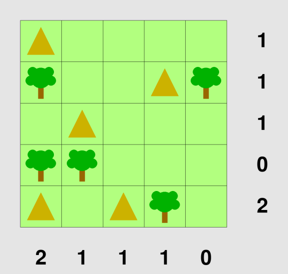

Data Representation: Localizing Constraints

1	2	3	4	5	31
6	7	8	9	10	32
11	12	13	14	15	33
16	17	18	19	20	34
21	22	23	24	25	35
26	27	28	29	30	

1	2	3	4	5
6	7	8	9	10
11	12	13	14	15
16	17	18	19	20
21	22	23	24	25

$\mathbb{1}$['empty']	$\mathbb{1}$['tree']	$\mathbb{1}$['tent']	$\mathbb{1}$['grass']	$\mathbb{1}$['violated']	$\mathbb{1}$['meta']	$\frac{\#\text{tents}}{\text{size}}\mathbb{1}$['meta']

$\mathbb{1}$['empty']	$\mathbb{1}$['tree']	$\mathbb{1}$['tent']	$\mathbb{1}$['grass']	$\frac{\#\text{tents in row}}{\text{row size}}$	$\frac{\#\text{tents in col}}{\text{col size}}$

Model: Transformer Encoder & RoPE

- 6 Encoder Layers (4-head, $d_{model} = 32$, $d_{ff} = 512$)
-
  - Attention Layer
  - MLP
- 2D Rotary Position Embeddings
- **Sampling**: Randomly crop 5×5 subgrids from larger grids during training.
- **Criterion:** Weighted Categorical CE

$$l_n = -w_{y_n} \log \frac{\exp(x_{n,y_n})}{\sum_{c=1}^{C} \exp(x_{n,c})} \mathbb{1}\{y_n \neq \text{MASK}\}$$

Training Strategy

Random 5×5 Window Sampling

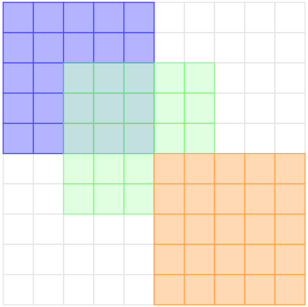

Methods: Encoder Model (No Decoding)

- Predict all cells simultaneously.
- Noise level = Corruption $\sim$ Difficulty (independently per cell)
- Loss only on corrupted, else copy
- **Train:** 5x5 grids with fixed corruption $c\%$
- **Test:** OOD sizes at same $c\%$

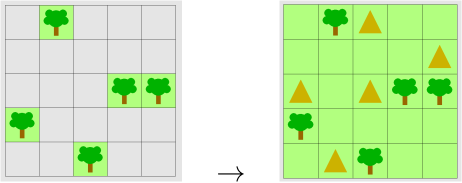

100% Corruption

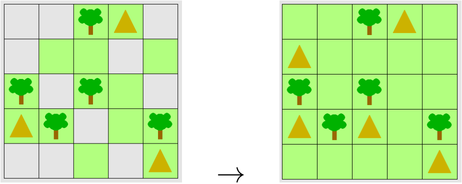

50% Corruption

Results: Encoder Model

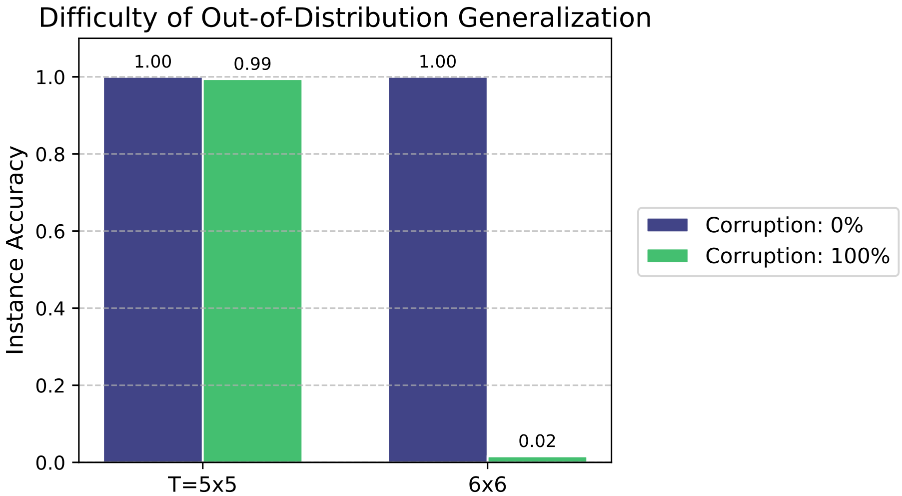

Results: Encoder Model

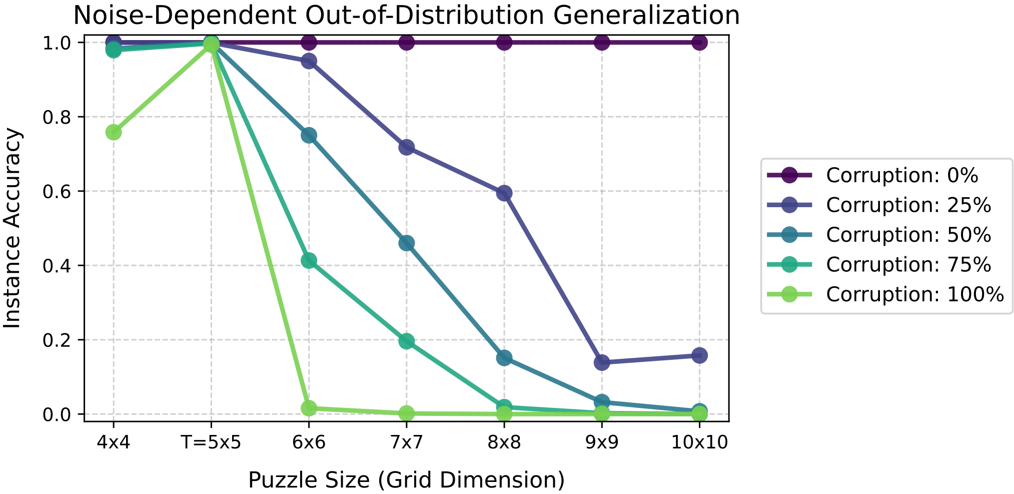

Discrete Diffusion for Puzzles Solving

- Diffusion = Iterative Refinement
- **Forward Process:** $q(z_t|x) = \text{Cat}(z_t; (1-t)x + t\text{MASK})$
- **Reverse Process:**
  - learn $p_\theta(x_0|x_t)$
  - sample $p_\theta(x_{t-1}|x_t) \propto \sum_{\tilde{x}_0} q(x_{t-1}, x_t|\tilde{x}_0)\tilde{p}_\theta(\tilde{x}_0|x_t)$
- Optimization objective (Sahoo et al., 2024)

$$\mathcal{L}^{\infty}_{\text{NELBO}} = \mathbb{E}_q \int_{t=0}^{t=1} \frac{-1}{t} \log\langle x_\theta(z_t, t), x\rangle \, dt$$

Methods: Decoding Model

- Train on 5x5 puzzles
- During training choose $t \in [0, 1]$ uniformly
- Test-time decoding
  - $\arg\max P$
  - one at the time
  - iterate until solved

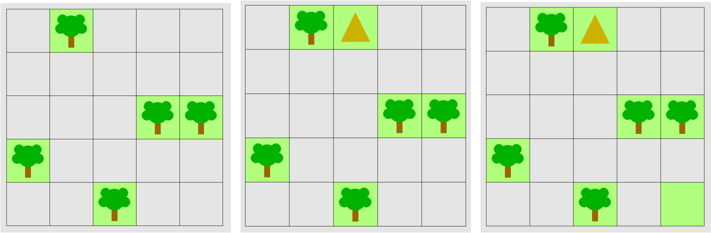

Step 1 **Step 2** **Step 3**

Results: Visual Hint

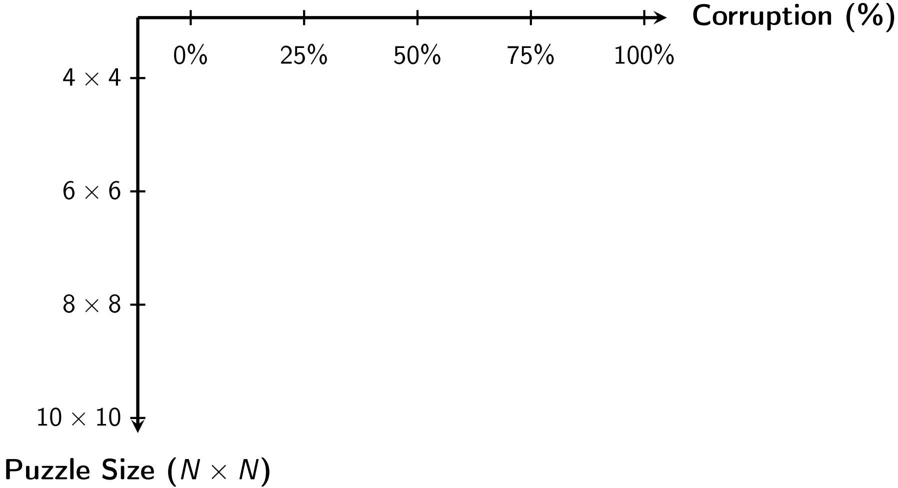

Results: Decoding Model

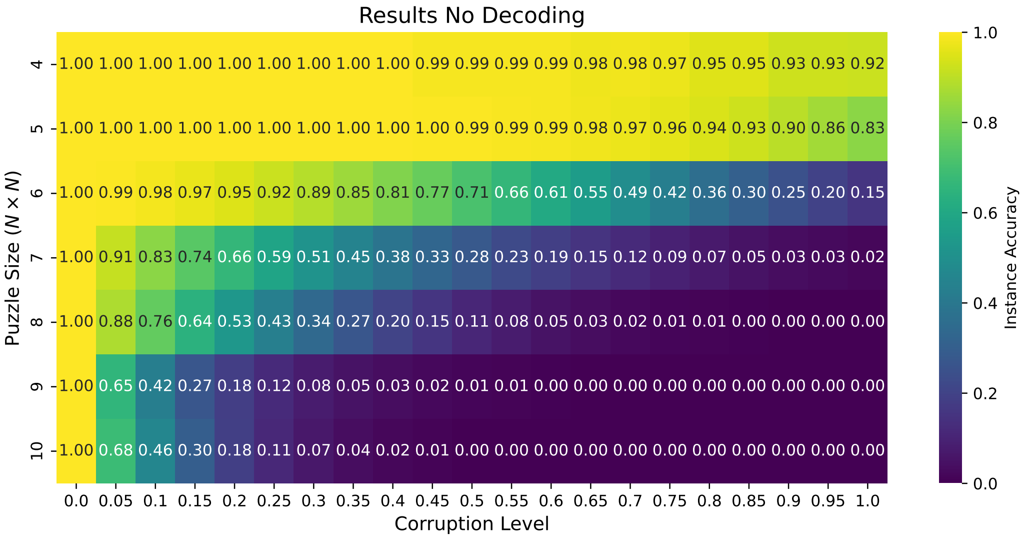

Results: Decoding Model

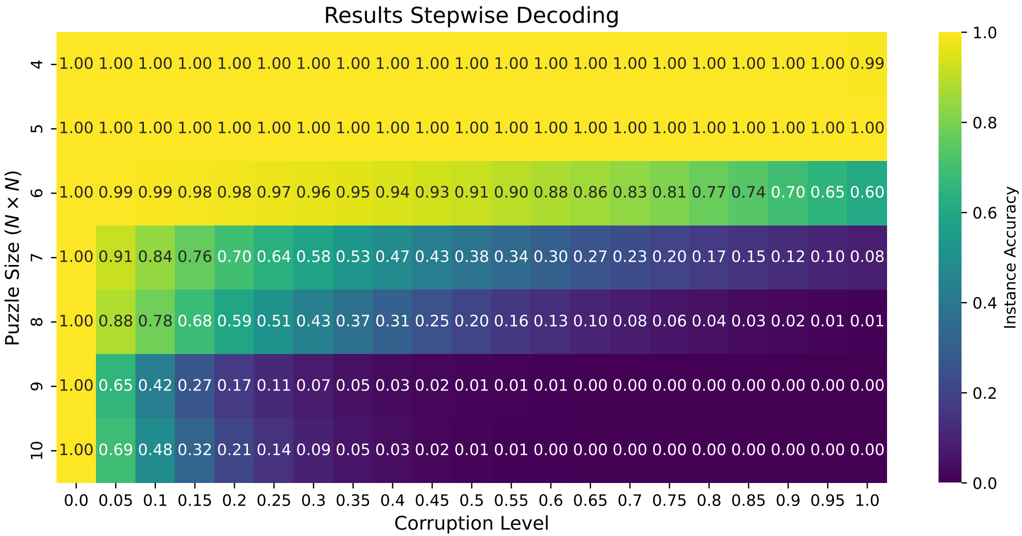

Results: Decoding Model

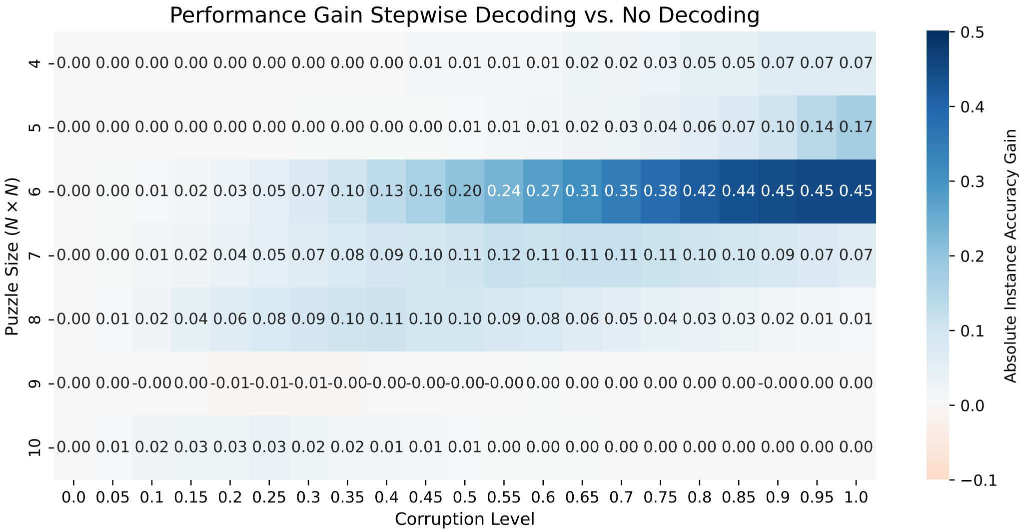

Reduction on false positive 'tent'.

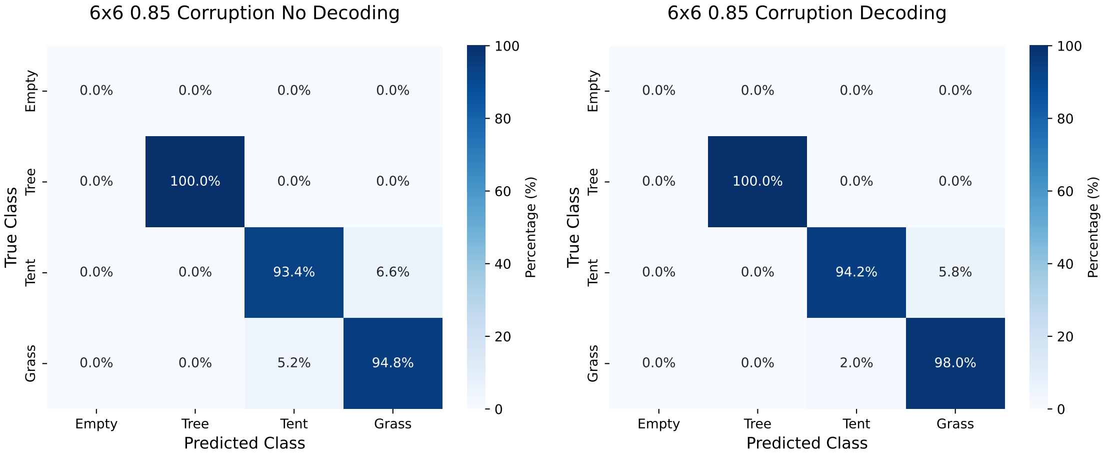

Compositional Reasoning: Diffusion Ensemble

Recombine prediction *in distribution* to predict on *OOD samples*.

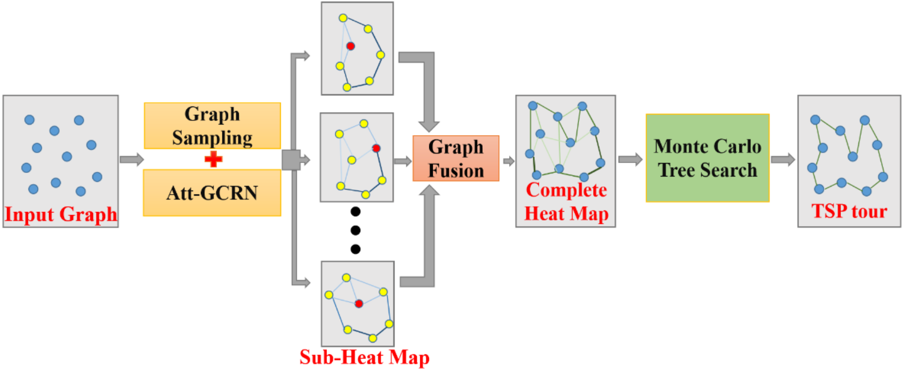

Figure: Figure from Fu et al. (2021)

Method: Diffusion Ensemble Inference

- Cut grids into overlapping 5x5 subgrids.
- Predict $p_\theta(x_0|x_t)$.
- Average the predicted probabilities in overlapping regions.
- Decode step by step.

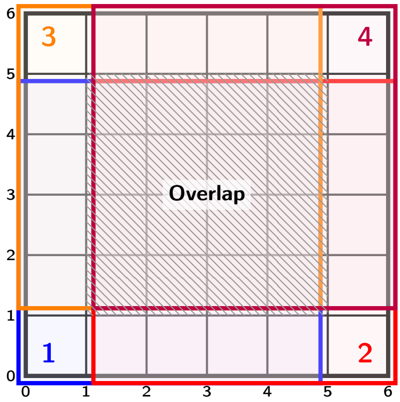

Results: Ensembles

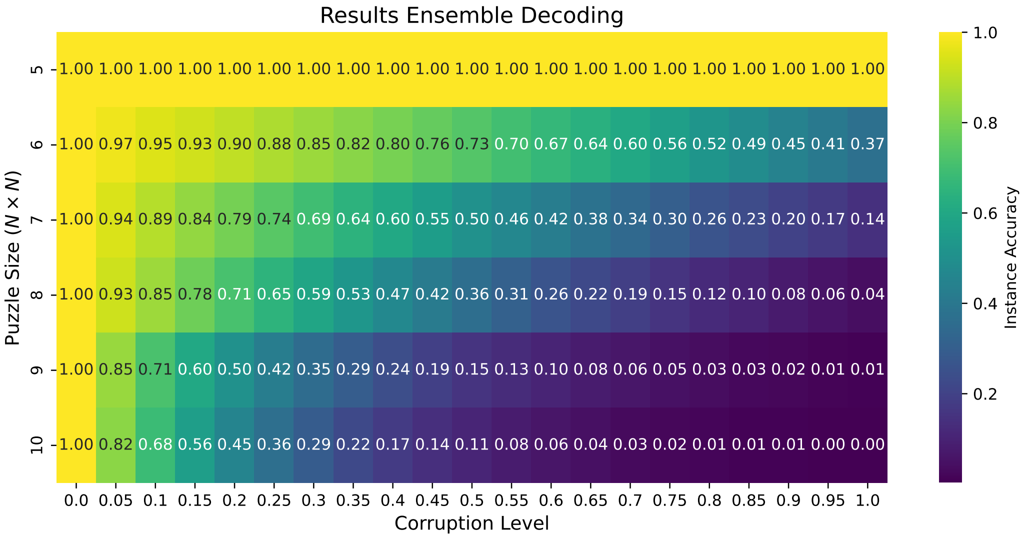

Results: Ensembles

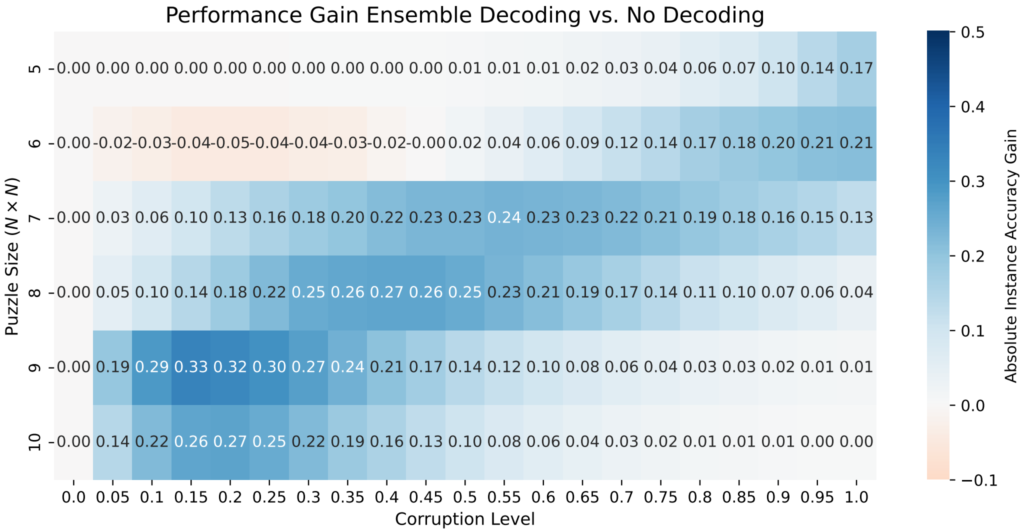

Discussion: Ensembles

Predicts more tents overall, increases grass misclassifications.

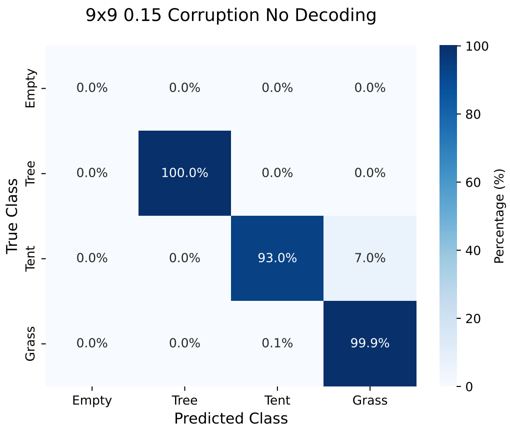

plots/confusion_matrix_9x9 0.15 Ensemble

Which one is better?

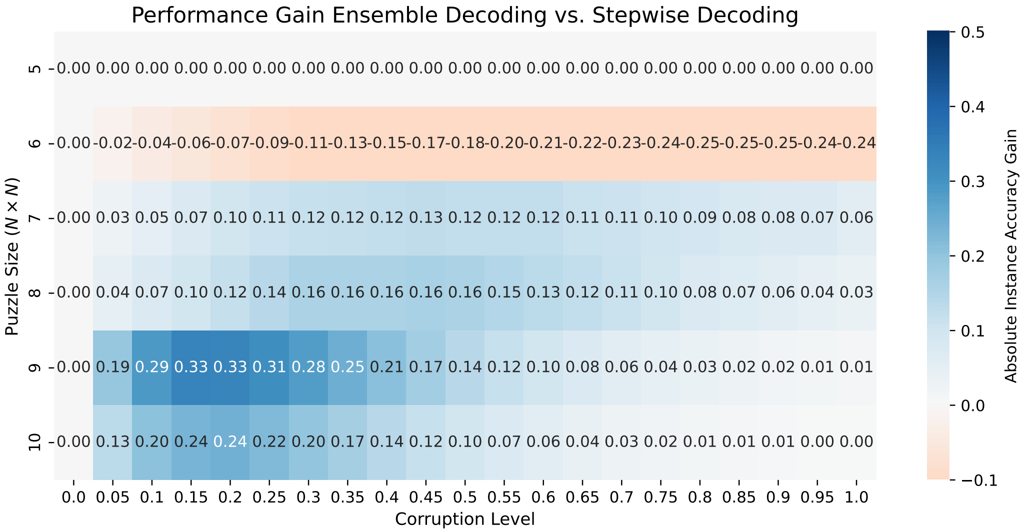

Conclusion & Future Work

Conclusions

- Localizing global constraints is vital for OOD generalization.
- Diffusion models: iterative solving, but have modest performance.
- Compositional ensembling boosts performance OOD.

Future Work

- Complex diffusion to allow "undoing" of mistakes.
- Different puzzles with compositional structure (Mosaic, Loopy).
- Analysis diffusion model ensembling.

Fertig lustig!

Thank you for your attention!

Bibliography

P. Veličković and C. Blundell, "Neural algorithmic reasoning," *Patterns*, vol. 2, no. 7, 2021.

P. Veličković, R. Ying, M. Padovano, R. Hadsell, and C. Blundell, "Neural execution of graph algorithms," *arXiv preprint arXiv:1910.10593*, 2019.

P. Veličković, A. P. Badia, D. Budden, R. Pascanu, A. Banino, M. Dashevskiy, R. Hadsell, and C. Blundell, "The clrs algorithmic reasoning benchmark," in *International Conference on Machine Learning*. PMLR, 2022, pp. 22 084–22 102.

Z. Sun and Y. Yang, "Difusco: Graph-based diffusion solvers for combinatorial optimization," *Advances in neural information processing systems*, vol. 36, pp. 3706–3731, 2023.

H. Zhao, K. Yu, Y. Huang, R. Yi, C. Zhu, and K. Xu, "Disco: Efficient diffusion solver for large-scale combinatorial optimization problems," *arXiv preprint arXiv:2406.19705*, 2024.

B. Estermann, L. A. Lanzendörfer, Y. Niedermayr, and R. Wattenhofer, "Puzzles: A benchmark for neural algorithmic reasoning," *Advances in Neural Information Processing Systems*, vol. 37, pp. 127 059–127 098, 2024.

S. Sahoo, M. Arriola, Y. Schiff, A. Gokaslan, E. Marroquin, J. Chiu, A. Rush, and V. Kuleshov, "Simple and effective masked diffusion language models," *Advances in Neural Information Processing Systems*, vol. 37, pp. 130 136–130 184, 2024.