

IMPROVING REACHABILITY ON REASONING PUZZLES

Anonymous authors

Paper under double-blind review

ABSTRACT

Recent work has observed the accuracy collapse of large language and reasoning models on moderately large puzzles. Apriori, it is unclear if this failure is due to the inability to find an algorithm or the inability to execute a given algorithm faithfully or simply due to the large number of execution steps. Many puzzles such as Tower of Hanoi are essentially graph reachability tasks, whose graph size increases exponentially. However, symmetry and structure of these graphs allow simple low-complexity algorithms to generate the solution, compared to general-purpose algorithms for reachability that require more compositions, recursion, memory. We theoretically show that faithful execution of certain low-complexity algorithms for reachability are easier to express for small transformer models, compared to general-purpose algorithms for reachability. We empirically observe that explicitly providing such algorithmic hints to language and reasoning models leads to substantially improved performance and pushes the frontier of their accuracy collapse.

1 INTRODUCTION

Recent years have witnessed the emergence of a new class of models, termed Large Reasoning Models (LRMs), which are explicitly designed to handle complex reasoning tasks (OpenAI, 2024; Guo et al., 2025a; Anthropic, 2025). These models typically employ Chain-of-Thought (CoT) reasoning, enabling them to explore multiple intermediate reasoning paths before producing a final answer. Despite achieving state-of-the-art performance on a wide range of reasoning benchmarks, recent evidence suggests that LRMs struggle with long-horizon reasoning, including comparatively simple reasoning puzzles whose solutions require coherent answers with a large number of steps (Shojaee* et al., 2025).

In this work, we study the algorithmic structure underlying reasoning tasks and argue that the algorithm, rather than task complexity that is believed to be the limiting factor, plays a decisive role in whether reasoning models succeed or fail. We focus on puzzles and reachability-style problems whose state spaces grow exponentially with input size, such as the Tower of Hanoi, river crossing puzzles, and checker-jumping games. These problems admit multiple correct solution strategies: some rely on recursive or sequential reasoning, while others admit explicit, closed-form descriptions that compute the desired outcome directly from the input. We naturally use results on transformer expressibility as the lens through which we differentiate various solution strategies.

1.1 RELATED WORKS

There has been substantial effort to analyze the expressive power of transformer architectures through a classical theoretical lens (Strobl et al., 2024). Recent works have shown that augmenting transformers with additional computational resources such as Chain-of-Thought Merrill & Sabharwal (2024), padding tokens Merrill & Sabharwal (2025a), and increased depth Merrill & Sabharwal (2025b) can significantly boost their expressive power.

Additionally, numerous works have focused on classifying transformer architectures in terms of classical complexity classes, with Strobl (2023) showing that with $O(\log n)$ bits of floating-point precision, average-hard attention transformers (AHATs) can be approximated in L-uniform TC^0 . Later, Merrill & Sabharwal (2025c) showed that softmax-attention transformers (SMATs) can be approximated in DLOGTIME-uniform TC^0 . These results were further improved by Chiang (2025),

054 who showed that SMATs with $O(\text{poly}(n))$ bits of floating-point precision, and AHATs with no
 055 approximations, are in DLOGTIME-uniform TC^0 .
 056

057 While expressibility is important from a theoretical standpoint, learnability is crucial for real-world
 058 applications. Many studies have attempted to bridge this gap with Yang et al. (2025) showing that
 059 a 2-layered decoder only transformer not only expresses but also learns tree reachability. Ye et al.
 060 (2025) show that the learned strategy for graph reachability hinges on whether most training instances
 061 are within the models capacity or not.

062 Shojae* et al. (2025) Schnabel et al. (2025) introduce the bounded attention prefix oracle (BAPO)
 063 model, a new computational framework that models bandwidth constraints on attention heads, the
 064 mechanism for internal communication in LLMs to show that several important reasoning problems
 065 like graph reachability require high communication bandwidth for BAPOs to solve.

066 1.2 OUR CONTRIBUTIONS

067 We view reasoning on puzzle graphs through the lens of reachability in the associated configuration
 068 graph, and observe that many such graphs exhibit a structure that can be exploited to have *paralleliz-*
 069 *able* algorithms (more formally, TC^0 circuits). Under this perspective, we focus on the Tower of
 070 Hanoi, a canonical puzzle that has been shown to be challenging for large reasoning models (LRMs)
 071 in recent work.
 072

- 073 1. We theoretically prove that a simple 1-layer, 2-head transformer can express the solution
 074 for Tower of Hanoi 3.1 when given the configuration graph as input. In addition, we also
 075 compile a transformer object using RASP that realizes the closed form TC^0 solution.
 076
- 077 2. Guided by insights on transformer expressibility through TC^0 computation, we perform
 078 experiments on LRMs. Specifically, for Tower of Hanoi, we showcase an increase in
 079 accuracy using two techniques (3): (i) passing the configuration graph to the LRM, (ii) using
 080 the non-sequential TC^0 . A detailed account of the experiments can be found in section 4.
 081

082 2 PRELIMINARIES

083 2.1 RELEVANT COMPLEXITY CLASSES

084 A Boolean circuit is a directed acyclic graph whose internal nodes are gates and whose leaves are
 085 input variables or constants. The size of a circuit is the number of gates, and the depth is the length
 086 of the longest path from an input to the output. A circuit family $\{C_n\}_{n \in \mathbb{N}}$ is *uniform* if there exists
 087 a logspace Turing machine that, on input 1^n , outputs a description of C_n . We define here all the
 088 complexity classes referred to in this paper.
 089

- 090 1. AC^0 consists of all Boolean functions computable by families of polynomial-size, depth-
 091 $O(1)$ circuits with unbounded fan-in $\{\text{AND}, \text{OR}, \text{NOT}\}$ gates. Negations are allowed only
 092 at the inputs.
 093
- 094 2. TC^0 is defined similarly, except that circuits may also use unbounded fan-in *threshold gates*.
 095 A threshold gate outputs 1 if the weighted sum of its inputs exceeds a given threshold.
 096
- 097 3. NC^1 consists of all Boolean functions computable by families of polynomial-size, depth-
 098 $O(\log n)$ circuits with $\{\text{AND}, \text{OR}, \text{NOT}\}$ gates with fan-in 2. Negations are allowed only
 099 at the inputs.

100 It is known that $\text{AC}^0 \subsetneq \text{TC}^0 \subseteq \text{NC}^1 \subseteq \text{L}$. Intuitively, L captures problems that can be solved while
 101 storing only a logarithmic number of bits of information about the input. This amount of memory is
 102 sufficient to maintain a constant number of pointers or counters into the input, but not enough to store
 103 large intermediate results.
 104

105 2.2 THEORETICAL LIMITATIONS OF TRANSFORMERS

106 Chiang (2025) showed that softmax-attention transformers (SMATs) with $O(\text{poly}(n))$ bits of floating-
 107 point precision, and average-hard attention transformers (AHATs) with no approximations are in

DLOGTIME-Uniform TC^0 , improving over the previous results from Strobl (2023) and Merrill & Sabharwal (2025c) who respectively showed that with $O(\log n)$ bits of floating-point precision AHATs can be approximated in L-uniform TC^0 , and SMATs can be approximated in DLOGTIME-uniform TC^0 .

SMATs are shown to exist between AC^0 and TC^0 , i.e.:

$$NC^0 \subsetneq AC^0 \subsetneq SMAT \subseteq TC^0 \subseteq NC^1.$$

2.3 GRAPH REACHABILITY

Given a graph $G = (V, E)$ and vertices $s, t \in V$, the *reachability problem* asks whether there exists a path from s to t in G .

Reachability in Special Graphs The computational difficulty of reachability depends crucially on the structure of the underlying graph. In general directed graphs, the reachability problem is complete for nondeterministic logspace (NL Jones (1975)). However, imposing structural constraints on the graph can dramatically reduce complexity. A classic example would be that the reachability of the undirected graph is L complete Reingold (2005), which is spatially more efficient (under the conjecture that $NL \neq L$). Allender et al. (2006) shows that reachability on a special case of grid graphs is TC^0 complete and we know that $TC^0 \subsetneq NC^1 \subseteq L$ making these much easier to solve.

Tower of Hanoi as reachability problem

2.4 RESTRICTED ACCESS SEQUENCE PROCESSING (RASP)

Restricted Access Sequence Processing (RASP) Weiss et al. (2021) is a small domain-specific language (DSL) designed to describe computations that are *native* to Transformer architectures. Rather than providing general-purpose control flow, RASP restricts the programmer to a handful of primitives that mirror common Transformer building blocks: elementwise transformations over token-wise features, comparisons that induce attention-style relations, and aggregation operators that combine information across sequence positions. As a result, a RASP program can often be viewed as a high-level, architecture-aligned specification of what a Transformer *could* compute.

Programs written in RASP can be compiled by tools such as Tracr Lindner et al. (2023) into concrete Transformer objects whose forward pass implements the program’s logic. Consequently, exhibiting a RASP program for a task serves as an existence proof: there exists a Transformer (possibly large and carefully constructed) that solves the task. Since expressing programs in RASP is a sufficient but not necessary requirement, $RASP \subseteq SMAT$.

3 REACHABILITY PROBLEMS WITH AND WITHOUT SEQUENTIAL STRUCTURE

3.1 SOLVING PUZZLES IS SOLVING REACHABILITY

Many combinatorial puzzles, such as the Tower of Hanoi, can be modeled as reachability problems in an implicitly defined graph, often called a configuration graph or a state graph. For example, the Tower of Hanoi puzzle can be formalized as a reachability problem on a directed graph $G_n = (V_n, E_n)$. The vertex set V_n consists of all legal configurations of n disks distributed among a fixed set of towers. There is a directed edge $(u, v) \in E_n$ if and only if configuration v can be obtained from u by a single legal move of one disk. Given a designated start configuration $s \in V_n$ and a target configuration $t \in V_n$, the Tower of Hanoi problem asks whether t is reachable from s in G_n . Other puzzle graphs such as River Crossing, Checker Jumping etc as well translate naturally in graph problems.

As we have discussed in the preliminaries solving reachability in the directed graph is NL complete, but special structural properties of the graph allow for easier algorithms. For our puzzle graphs, where usually the configuration graph is exponential in the size of the input, graph reachability would appear very hard to solve. In fact, Hearn & Demaine (2005) actually shows that some puzzle problems, such as the sliding block puzzle, are PSPACE-complete, making them very hard to compute. However,

162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215




Figure 1: Reformulation of Tower of Hanoi

unlike general puzzle reachability problems captured by the above problem, although $|V_n|$ for Tower of Hanoi grows exponentially with n , the graph G_n has a highly regular and recursive structure, making it a canonical example of reachability in an implicitly defined and structurally constrained graph.

3.2 TOWER OF HANOI : RECURSIVE VS TC^0 ALGORITHM

The classical recursive algorithm for solving the Tower of Hanoi produces an optimal sequence of moves but is inherently sequential: each move depends on the previously executed moves.

Surprisingly, there also exists a closed-form arithmetic description of individual moves in the optimal solution. This allows the computation of any constant number of moves using highly parallel computation, placing this task within the complexity class TC^0 .

Algorithm 1 TOWER-OF-HANOI (Arithmetic / TC^0)

Require: Move index i , number of disks n

- 1: $d \leftarrow v_2(i) + 1$
- 2: $t \leftarrow \left\lfloor \frac{i}{2^{d-1}} \right\rfloor$
- 3: **if** $n - d$ is even **then**
- 4: $\delta_d \leftarrow +1$
- 5: **else**
- 6: $\delta_d \leftarrow -1$
- 7: **end if**
- 8: $\text{src}(i) \leftarrow \delta_d \cdot (t - 1) \bmod 3$
- 9: $\text{dst}(i) \leftarrow \delta_d \cdot t \bmod 3$
- 10: Move disk d from peg $\text{src}(i)$ to peg $\text{dst}(i)$

Algorithm 2 TOWER-OF-HANOI (Recursive)

Require: Number of disks n , pegs (s, a, t)

- 1: **if** $n = 1$ **then**
- 2: Move disk 1 from s to t
- 3: **else**
- 4: TOWER-OF-HANOI($n - 1, s, a, t$)
- 5: Move disk n from s to t
- 6: TOWER-OF-HANOI($n - 1, a, s, t$)
- 7: **end if**

TC^0 Complexity Class This captures computations that can be performed by highly parallel, constant-depth circuits that allow threshold gates. From an algorithmic perspective, this characterization emphasizes that such algorithms perform operations such as indexing, thresholding, parity, and constant-size arithmetic. In contrast, TC^0 circuits are generally not believed to support unbounded sequential processes such as recursion or general graph exploration.

3.3 TRANSFORMER EXPRESSIVITY OF TOWER OF HANOI

Prior literature has mapped the transformer architecture to classical models of computation Strobl et al. (2024), within which Boolean circuits are the most natural due to the inherent parallelizability of attention, and also analyzed how adding computation resources like padding tokens Merrill & Sabharwal (2025a), depth scaling Merrill & Sabharwal (2025b), etc. increase the expressive power of these models. Recent expressibility results Merrill & Sabharwal (2025b) Merrill & Sabharwal (2024) place a class of transformers with logarithmic CoT within uniform TC^0 . We adopt constant-depth threshold circuits (TC^0) as a convenient abstraction for the algorithmic capabilities of small

transformer models. Consequently, problems whose solutions can be expressed by direct closed-form arithmetic computations are well aligned with the expressive power of these transformers, while tasks that are inherently sequential would not.

We show the explicit construction of a simple transformer that expresses the path for Tower of Hanoi graph given the configuration graph G_n

Theorem 3.1. *There exists a 2-head 1-layer decoder only transformer with embedding dimension $O(\exp(n))$ that given the Tower of Hanoi graph G can output the path needed to move the configuration with all disks on 1 tower to all disks on another*

ToH in RASP We show that the TC^0 closed form algorithm can be expressed in RASP and compile it to a Transformer object with 4 heads and 7 layers, where the width of the MLP blocks scales quadratically with n , thus proving that the algorithm is expressible with the attention model. The compiled model also only requires attention in the first layer, followed by a series of MLP operations.

Other architectural and implementation details can be found in Section D.1, D.2 and the code repository.

4 EXPERIMENTS

Building on the distinction between sequential and parallelizable computation discussed above, we design experiments that probe the limits of large reasoning models under structured algorithmic tasks. Compared to prior work such as the Shojaei* et al. (2025), our evaluations explicitly leverage closed-form and TC^0 -style computations as guiding principles. We examine whether this alignment improves robustness and increases the failure threshold beyond which model performance degrades.

4.1 INITIAL EXPERIMENTS TO IMPROVE FAILURE THRESHOLD

While Shojaei* et al. (2025) reports that providing explicit algorithms to reasoning models does not lead to substantial accuracy gains, we show that modification such as prompt refinements, use of in-context examples, improved algorithm selection, and modifications to the output format can result in significant performance improvements on certain puzzles. A detailed description of the experimental setup is provided in Appendix ??

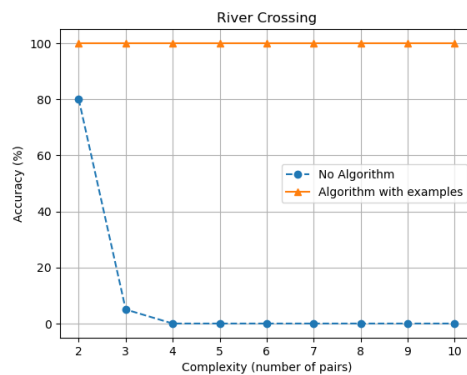


Figure 2: River Crossing with Algorithm

4.2 GRAPH AS INPUT

Since the tower of Hanoi problem can be represented as a graph reachability instance, we programmatically construct a graph which captures the given problem. We then pass the graph, along with the start and end node to the LRM. The following plot indicates the improvement in the performance.

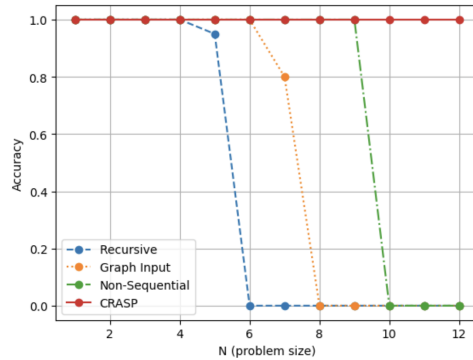


Figure 3: Accuracy across algorithms

Although the LRM is able to solve the problem for instances less than 8, due to the exponential increase in the size of the graph with the increase in n , it becomes impractical to feed the entire graph to the model.

4.3 NON-SEQUENTIAL (TC^0) ALGORITHM

To combat the issue arising due to the increase in graph size, we use the efficient TC^0 algorithm to split the long horizon reasoning task into non-sequential sub-tasks. Due to the independence of the current step from the previous step we are able to prevent both the accumulation of error as well as increase the accuracy significantly. Using this method, the accuracy now shoots up to a 100% till $n = 9$ (figure 3).

Other Puzzle Problems with TC^0 algorithms We also note that the checkers jumping problem also has a TC^0 algorithm that improves the failure threshold (figure 4)

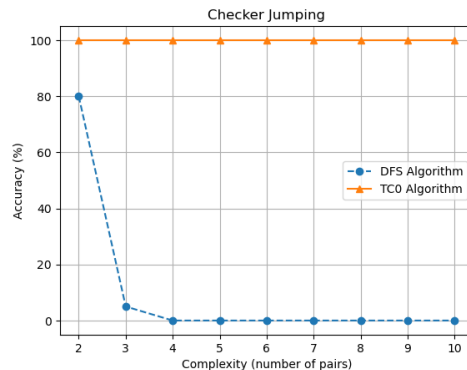


Figure 4: Checker Jumping with different algorithms

4.4 TRAINING A TRANSFORMER

We evaluate the transformer on the Tower of Hanoi reachability task with training instances up to $n = 10$ disks and test instances up to $n = 13$. The model reliably learns to predict the identity of the disk moved at each time step, generalizing almost perfectly to larger n well beyond the training range. In contrast, predicting the source and destination pegs is significantly harder: while accuracy is high for small n , it degrades rapidly as n increases. This gap suggests that the model captures the structure governing which disk moves but struggles to infer the finer-grained tower dynamics required to determine where the disk moves for longer horizons.

Architectural details We use a decoder-only Transformer with 4 layers, 8 self-attention heads per layer, and a model dimension of 512, with a feedforward dimension of 2048. Inputs are tokenized representations of problem size n , move index t , and disk index, embedded using learned embeddings and combined with learned positional embeddings. The encoder outputs are pooled by selecting the final token representation, which is then fed into task-specific linear heads. One head predicts the disk moved, while two separate heads predict the source and destination towers. All components are trained jointly using cross-entropy loss with AdamW optimization.




Figure 5: Performance of a simple transformer

5 CONCLUSION

We investigated the failure modes of large reasoning models (LRMs) on combinatorial puzzle problems through the lens of transformer expressibility. Although prior work attributes performance degradation to exponential growth of the underlying state space, our results suggest that the algorithmic structure of the solution plays a decisive role.

By reframing puzzles such as Tower of Hanoi and Checker Jumping as reachability problems on configuration graphs, we distinguished between inherently sequential solution strategies and closed-form, highly parallelizable algorithms. We construct small constant-depth transformers can express the closed-form solution for the Tower of Hanoi problem. Empirically, we demonstrated that explicitly providing this TC^0 algorithm substantially improves the failure threshold of LRMs.

Two natural directions follow. First, it would be valuable to explore this expressibility gap across a broader class of puzzle reasoning problems. Second, beyond expressibility, an important open question is learnability: how architectural biases and structural alignment affect what models can reliably acquire under gradient-based training Yang et al. (2025).

REFERENCES

- Eric Allender, David A. Mix Barrington, Tanmoy Chakraborty, Samir Datta, and Sambuddha Roy. Grid graph reachability problems. In *Proceedings of the 21st Annual IEEE Conference on Computational Complexity*, CCC '06, pp. 299–313, USA, 2006. IEEE Computer Society. ISBN 0769525962. doi: 10.1109/CCC.2006.22. URL <https://doi.org/10.1109/CCC.2006.22>.
- Anthropic. Claude 3.7 sonnet, Feb 2025. <https://www.anthropic.com/index/claude-3-7-sonnet>.
- David Chiang. Transformers in uniform tc^0 , 2025. URL <https://arxiv.org/abs/2409.13629>.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint*, arXiv:2501.12948, 2025a.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms

- 378 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025b. URL <https://arxiv.org/abs/2501.12948>.
379
380
- 381 Robert A. Hearn and Erik D. Demaine. Pspace-completeness of sliding-block puzzles and other
382 problems through the nondeterministic constraint logic model of computation. *Theoretical*
383 *Computer Science*, 343(1):72–96, 2005. ISSN 0304-3975. doi: [https://doi.org/10.1016/j.tcs.](https://doi.org/10.1016/j.tcs.2005.05.008)
384 2005.05.008. URL [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S0304397505003105)
385 S0304397505003105. Game Theory Meets Theoretical Computer Science.
- 386 Andreas M. Hinz, Sandi Klavžar, Uroš Milutinović, and Ciril Petr. *The Tower of Hanoi – Myths and*
387 *Maths*. Birkhäuser, 2013 (2nd ed. 2018).
388
- 389 Neil D. Jones. Space-bounded reducibility among combinatorial problems. *Journal of Com-*
390 *puter and System Sciences*, 11(1):68–85, 1975. ISSN 0022-0000. doi: [https://doi.org/10.](https://doi.org/10.1016/S0022-0000(75)80050-X)
391 1016/S0022-0000(75)80050-X. URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S002200007580050X)
392 article/pii/S002200007580050X.
- 393 David Lindner, János Kramár, Sebastian Farquhar, Matthew Rahtz, Thomas McGrath, and Vladimir
394 Mikulik. Tracr: Compiled transformers as a laboratory for interpretability, 2023. URL <https://arxiv.org/abs/2301.05062>.
395
396
- 397 William Merrill and Ashish Sabharwal. The expressive power of transformers with chain of thought,
398 2024. URL <https://arxiv.org/abs/2310.07923>.
399
- 400 William Merrill and Ashish Sabharwal. Exact expressive power of transformers with padding, 2025a.
401 URL <https://arxiv.org/abs/2505.18948>.
- 402 William Merrill and Ashish Sabharwal. A little depth goes a long way: The expressive power of
403 log-depth transformers, 2025b. URL <https://arxiv.org/abs/2503.03961>.
404
- 405 William Merrill and Ashish Sabharwal. A logic for expressing log-precision transformers, 2025c.
406 URL <https://arxiv.org/abs/2210.02671>.
407
- 408 OpenAI. Introducing openai o1, Jan 2024. [https://openai.com/research/](https://openai.com/research/introducing-openai-o1)
409 introducing-openai-o1.
- 410 Omer Reingold. Undirected st-connectivity in log-space. In *Proceedings of the Thirty-Seventh Annual*
411 *ACM Symposium on Theory of Computing*, STOC ’05, pp. 376–385, New York, NY, USA, 2005.
412 Association for Computing Machinery. ISBN 1581139608. doi: 10.1145/1060590.1060647. URL
413 <https://doi.org/10.1145/1060590.1060647>.
414
- 415 Tobias Schnabel, Kiran Tomlinson, Adith Swaminathan, and Jennifer Neville. Lost in
416 transmission: When and why llms fail to reason globally. In *NeurIPS 2025*, Decem-
417 ber 2025. URL [https://www.microsoft.com/en-us/research/publication/](https://www.microsoft.com/en-us/research/publication/lost-in-transmission-when-and-why-llms-fail-to-reason-globally/)
418 lost-in-transmission-when-and-why-llms-fail-to-reason-globally/.
- 419 Parshin Shojaee*, Iman Mirzadeh*, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad
420 Farajtabar. The illusion of thinking: Understanding the strengths and limitations of reasoning
421 models via the lens of problem complexity. In *NeurIPS*, 2025. URL [https://arxiv.org/](https://arxiv.org/abs/2506.06941)
422 abs/2506.06941.
423
- 424 Lena Strobl. Average-hard attention transformers are constant-depth uniform threshold circuits, 2023.
425 URL <https://arxiv.org/abs/2308.03212>.
- 426 Lena Strobl, William Merrill, Gail Weiss, David Chiang, and Dana Angluin. What formal languages
427 can transformers express? a survey. *Transactions of the Association for Computational Linguistics*,
428 12:543–561, 2024. ISSN 2307-387X. doi: 10.1162/tacl.a.00663. URL [http://dx.doi.org/](http://dx.doi.org/10.1162/tacl.a.00663)
429 10.1162/tacl.a.00663.
430
- 431 Gail Weiss, Yoav Goldberg, and Eran Yahav. Thinking like transformers, 2021. URL <https://arxiv.org/abs/2106.06981>.

432 Tong Yang, Yu Huang, Yingbin Liang, and Yuejie Chi. Multi-head transformers provably
 433 learn symbolic multi-step reasoning via gradient descent. 2025. URL <https://api.semanticscholar.org/CorpusID:280566397>.

434 Qilin Ye, Deqing Fu, Robin Jia, and Vatsal Sharan. When do transformers learn heuristics for graph
 435 connectivity?, 2025. URL <https://arxiv.org/abs/2510.19753>.

438 A TOWER OF HANOI IN L - UNIFORM TC^0

439 A.1 INDEX OF WHICH DISK TO MOVE IN THE OPTIMAL SOLUTION PATH AT STEP i ALGORITHM

440 **Theorem A.1.** *In the optimal Tower of Hanoi solution with n disks, the disk moved at step i (for
 441 $1 \leq i \leq 2^n - 1$) is the disk whose index equals the position of the least significant 1 bit in the binary
 442 representation of i . Equivalently,*

$$443 \text{disk}(i) = v_2(i) + 1,$$

444 where $v_2(i)$ denotes the largest power of 2 dividing i (i.e., $v_2(i) = \max\{k : 2^k \mid i\}$).

445 **Lemma A.2.** *Let $n \geq 1$ and let $1 \leq j \leq 2^{n-1} - 1$. Then*

$$446 v_2(2^{n-1} + j) = v_2(j).$$

447 *Proof.* Write $j = 2^r \cdot t$ with t odd and $r = v_2(j) \geq 0$. Since $1 \leq j \leq 2^{n-1} - 1$ and j is a positive
 448 integer less than 2^{n-1} , we must have $r \leq n - 2$. (If $r \geq n - 1$, then $j \geq 2^{n-1}$, contradicting
 449 $j \leq 2^{n-1} - 1$.)

450 Therefore,

$$451 2^{n-1} + j = 2^{n-1} + 2^r \cdot t = 2^r(2^{n-1-r} + t).$$

452 Since $r \leq n - 2$, we have $n - 1 - r \geq 1$, which means 2^{n-1-r} is even. Because t is odd, the sum
 453 $2^{n-1-r} + t$ is odd (even + odd = odd).

454 Thus, 2^r is the exact power of 2 dividing $2^{n-1} + j$, and so

$$455 v_2(2^{n-1} + j) = r = v_2(j). \quad \square$$

456 *Proof of the theorem.* We proceed by induction on n .

457 **Base case ($n = 1$).** For $n = 1$ there is exactly one move $i = 1$, and $v_2(1) = 0$, so $\text{disk}(1) =$
 458 $v_2(1) + 1 = 1$.

459 **Inductive hypothesis.** Assume the statement holds for $n - 1$: for all $1 \leq j \leq 2^{n-1} - 1$, $\text{disk}(j) =$
 460 $v_2(j) + 1$ in the recursive solution for $n - 1$ disks.

461 **Inductive step.** The recursive solution for n disks is

$$462 S_n = S_{n-1}, \quad \text{move disk } n, \quad S_{n-1},$$

463 with total moves $1, 2, \dots, 2^n - 1$.

- 464 1. If $1 \leq i \leq 2^{n-1} - 1$, the i -th move is in the first copy of S_{n-1} , so by the inductive
 465 hypothesis $\text{disk}(i) = v_2(i) + 1$.
- 466 2. If $i = 2^{n-1}$, this is the single move of disk n . Since $v_2(2^{n-1}) = n - 1$, we get $\text{disk}(i) =$
 467 $v_2(i) + 1 = n$.
- 468 3. If $2^{n-1} + 1 \leq i \leq 2^n - 1$, write $i = 2^{n-1} + j$. Then necessarily

$$469 1 \leq j \leq 2^{n-1} - 1.$$

470 (Indeed, the maximum possible j in this decomposition is $j = 2^{n-1} - 1$, and $2^{n-1} - 1 <$
 471 $2^n - 1$, so j is strictly less than $2^n - 1$; this ensures i remains within the range $1, \dots, 2^n - 1$.)

472 By the recursive structure the i -th move of S_n is the same disk as the j -th move of S_{n-1} .
 473 By the lemma above we have

$$474 v_2(i) = v_2(2^{n-1} + j) = v_2(j).$$

486 Applying the inductive hypothesis to S_{n-1} gives

$$487 \text{disk}(i) = v_2(i) + 1 = v_2(j) + 1 = \text{disk}(j),$$

488 which is exactly the disk moved in the j -th move of S_{n-1} .

489 Thus the formula holds for n . By induction it holds for all $n \geq 1$. \square

490 **Theorem A.3.** *In the optimal Tower of Hanoi solution with n disks, the disk moved at step i*
 491 *($1 \leq i \leq 2^n - 1$) is exactly the disk whose index equals the position of the least significant 1 bit of i .*
 492 *Equivalently,*

$$493 \text{disk}(i) = v_2(i) + 1,$$

494 where $v_2(i) = \max\{k : 2^k \mid i\}$. Moreover, the predicate $\text{disk}(i) = d$ is definable in uniform TC^0 .

495 *Proof.* The classical recursive solution to the Tower of Hanoi problem has the following well-known
 496 characterization: disk d is moved exactly once every 2^{d-1} steps, and it is never moved on any step
 497 divisible by 2^d . Hence disk d is moved at step i if and only if

$$500 2^{d-1} \mid i \quad \text{and} \quad 2^d \nmid i.$$

501 Equivalently, $d - 1 = v_2(i)$, and therefore $\text{disk}(i) = v_2(i) + 1$.

502 We now show that the predicate $\text{disk}(i) = d$ is definable in uniform TC^0 .

503 Represent i in binary as $i = (b_{m-1} \dots b_1 b_0)_2$ with $m = O(n)$. Then $v_2(i) = k$ holds exactly when

$$504 b_0 = b_1 = \dots = b_{k-1} = 0 \quad \text{and} \quad b_k = 1.$$

505 Thus $\text{disk}(i) = d$ holds iff

$$506 b_0 = \dots = b_{d-2} = 0 \quad \wedge \quad b_{d-1} = 1.$$

507 This predicate is computable by a constant-depth, polynomial-size threshold circuit: we check with
 508 an AND gate that the first $d - 1$ bits are all 0, and with a single wire that bit $d - 1$ equals 1. Such bit
 509 tests and bounded conjunctions are in $\text{AC}^0 \subseteq \text{TC}^0$, uniformly constructible from n and d .

510 Therefore the function $\text{disk}(i)$ is computable in uniform TC^0 , and the stated formula holds. \square

511 A.2 PEG NUMBER FOR A GIVEN DISK AT STEP i , I.E., STATE CONFIGURATION AT STEP i ALGO

512 **Theorem A.4.** *Let n be the number of disks and $1 \leq i \leq 2^n - 1$ a move index. Let*

$$513 d = \text{disk}(i) = v_2(i) + 1.$$

514 Define

$$515 t(i, d) = \left\lfloor \frac{i}{2^{d-1}} \right\rfloor, \quad \delta_d = \begin{cases} +1 & \text{if } n - d \text{ is even,} \\ -1 & \text{if } n - d \text{ is odd,} \end{cases}$$

516 with arithmetic modulo 3 on peg labels $\{0, 1, 2\}$.

517 Then the source and destination pegs of the move at step i are

$$518 \begin{aligned} \text{src}(i) &= \delta_d \cdot (t(i, d) - 1) \pmod{3}, \\ \text{dst}(i) &= \delta_d \cdot t(i, d) \pmod{3}. \end{aligned}$$

519 Moreover, the functions $\text{disk}(i)$, $\text{src}(i)$, and $\text{dst}(i)$ are computable in uniform TC^0 .

520 *Proof.* We use the standard structural properties of the optimal Tower of Hanoi solution.

540 **Step 1: Identification of the moving disk.** It is classical that disk d is moved exactly at those times
 541 i such that

$$542 \quad 2^{d-1} \mid i \quad \text{and} \quad 2^d \nmid i,$$

543 equivalently $d = v_2(i) + 1$. Thus $\text{disk}(i) = d$.
 544

545 **Step 2: Frequency of motion.** Disk d moves once every 2^{d-1} steps. Hence before step i , disk d
 546 has moved

$$547 \quad t(i, d) - 1 = \left\lfloor \frac{i}{2^{d-1}} \right\rfloor - 1$$

548 times, and after step i it has moved $t(i, d)$ times.
 549
 550

551 **Step 3: Direction of motion.** In the optimal solution, each disk moves monotonically in a fixed
 552 cyclic direction around the three pegs. The direction depends only on the parity of $n - d$: if $n - d$ is
 553 even, disk d moves clockwise; if $n - d$ is odd, it moves counterclockwise. This direction is encoded
 554 by $\delta_d \in \{+1, -1\}$.
 555

556 **Step 4: Peg positions.** Each move advances disk d by one peg in its fixed direction. Starting from
 557 peg 0, after k moves disk d is on peg

$$558 \quad \delta_d \cdot k \pmod{3}.$$

559 Therefore:
 560

- 561 • Before step i , disk d is on peg

$$562 \quad \delta_d \cdot (t(i, d) - 1) \pmod{3},$$

- 564 • After step i , disk d is on peg

$$565 \quad \delta_d \cdot t(i, d) \pmod{3}.$$

566 Thus the move at time i is exactly from
 567

$$568 \quad \text{src}(i) = \delta_d(t(i, d) - 1) \pmod{3} \quad \text{to} \quad \text{dst}(i) = \delta_d t(i, d) \pmod{3}.$$

569 This is the unique legal optimal move, proving correctness. □
 570
 571

572 *Complexity Analysis.* We show that $\text{disk}(i)$, $\text{src}(i)$, and $\text{dst}(i)$ are computable in uniform TC^0 .
 573

574 **Disk index.** The predicate $v_2(i) = k$ is definable by testing

$$575 \quad b_0 = \dots = b_{k-1} = 0 \wedge b_k = 1,$$

576 which is in $\text{AC}^0 \subseteq \text{TC}^0$. Hence $d = v_2(i) + 1$ is in uniform TC^0 .
 577

578 **Move count.** The quantity $t(i, d) = \lfloor i/2^{d-1} \rfloor$ is obtained by shifting i right by $d - 1$ bits. Bit
 579 selection and shifts by a TC^0 -computable index are in uniform TC^0 .
 580

581 **Direction bit.** The parity of $n - d$ is computable in AC^0 , hence $\delta_d \in \{+1, -1\}$ is definable in
 582 uniform TC^0 .
 583

584 **Modulo 3 reduction.** Multiplication by ± 1 and reduction modulo 3 are computable by constant-
 585 depth threshold circuits; $\text{mod}3$ is in uniform TC^0 .
 586

587 **Conclusion.** All operations in the definitions of $\text{src}(i)$ and $\text{dst}(i)$ — bit tests, shifts, parity, multi-
 588 plication by ± 1 , and reduction modulo 3 — lie in uniform TC^0 . Therefore the full transducer

$$589 \quad i \mapsto (\text{disk}(i), \text{src}(i), \text{dst}(i))$$

590 is computable in uniform TC^0 . □
 591
 592

593 Since TC^0 is an upper bound on Softmax Attention Transformers, we also prove that the aforemen-
 tioned formula is computable by a transformer by writing it in RASP.

B OTHER PUZZLE PROBLEMS IN L - UNIFORM TC^0

B.1 RIVER CROSSING

Algorithm 3 RIVER-CROSSING($\{(a_i, A_i)\}_{i=1}^n$)

```

1: for  $i = 1$  to  $n$  do
2:   Send  $(a_i, A_i)$  across the river
3:   Return  $a_i$  alone
4:   Send  $(a_i, a_{i+1})$  across the river
5:   Return  $a_{i+1}$  alone
6: end for

```

B.2 CHECKERS JUMPING

Algorithm 4 Recursive Backtracking Solver for Checker Jumping

Require: Current board $board$, target configuration $goal$, list $moves$

Ensure: Sequence of moves reaching $goal$, or **null** if none exists

```

Solveboard, goal, moves
1: if  $board = goal$  then
2:   return  $moves$ 
3: end if
4: for  $i = 1$  to  $length(board)$  do
5:   if  $board[i] = R$  then
6:      $result \leftarrow TryMoveR, i, i + 1, board, goal, moves$ 
7:     if  $result \neq null$  then
8:       return  $result$ 
9:     end if
10:     $result \leftarrow TryMoveR, i, i + 2, board, goal, moves$ 
11:    if  $result \neq null$  then
12:      return  $result$ 
13:    end if
14:  end if
15:  if  $board[i] = B$  then
16:     $result \leftarrow TryMoveB, i, i - 1, board, goal, moves$ 
17:    if  $result \neq null$  then
18:      return  $result$ 
19:    end if
20:     $result \leftarrow TryMoveB, i, i - 2, board, goal, moves$ 
21:    if  $result \neq null$  then
22:      return  $result$ 
23:    end if
24:  end if
25: end for
26: return  $null$ 
TryMovepiece, from, to, board, goal, moves
27: if  $to$  is a valid position AND  $board[to] = \_$  then
28:    $new\_board \leftarrow copy\ of\ board$ 
29:    $new\_board[from] \leftarrow \_$ 
30:    $new\_board[to] \leftarrow piece$ 
31:    $new\_moves \leftarrow moves \cup \{(from, to)\}$ 
32:   return  $Solve(new\_board, goal, new\_moves)$ 
33: end if
34: return  $null$ 

```

Algorithm 5 Move Type at Index i

Require: Move index i
Ensure: Symbol at move i

- 1: For all r in parallel:
- 2: $d_r \leftarrow |pos_B[r] - pos_A[r]|$
- 3: Compute in parallel prefix sums:
- 4: $S_r \leftarrow \sum_{t \leq r} d_t$
- 5: $S_{r-1} \leftarrow S_r - d_r$
- 6: For all r in parallel:
- 7: $owner_r \leftarrow (S_{r-1} < i \leq S_r)$
- 8: Let r^* be the unique r with $owner_r = 1$
- 9: $offset \leftarrow i - S_{r^*-1}$
- 10: **if** $offset \leq \lfloor d_{r^*}/2 \rfloor$ **then**
- 11: **if** piece r^* is R **then**
- 12: **return** J
- 13: **else**
- 14: **return** j
- 15: **end if**
- 16: **else**
- 17: **if** piece r^* is R **then**
- 18: **return** M
- 19: **else**
- 20: **return** m
- 21: **end if**
- 22: **end if**

C TOWER OF HANOI TRANSFORMER

C.1 EDGE PATTERN ALGO

An iterative solution Hinz et al. (2013 (2nd ed. 2018)) to the Tower of Hanoi problem can be described as the repeated execution of the following sequence of steps until the goal configuration is reached:

- Move one disk between pegs A and B , in whichever direction the move is legal.
- Move one disk between pegs A and C , in whichever direction the move is legal.
- Move one disk between pegs B and C , in whichever direction the move is legal.

Following this procedure, the stack of disks will end up on peg B if the number of disks is odd, and on peg C if the number of disks is even.

The reachability version of this uses the following algorithm when the moves $A \leftrightarrow B, B \leftrightarrow C, C \leftrightarrow A$ are labelled as $\{1, 2, 3\}$

Algorithm 6 Follow Edge Pattern in a Colored Graph

- 1: **Input:** Graph $G = (V, E)$ with edge colorings $\{1, 2, 3\}$
- 2: **Initialize:** Start at the initial vertex v_0
- 3: **while** $v \neq$ final state **do**
- 4: **for** color c in sequence $[1, 2, 3]$ **do**
- 5: Follow the edge from current vertex v with color c
- 6: Update current vertex v to the reached vertex
- 7: **if** v is the final state **then**
- 8: **break**
- 9: **end if**
- 10: **end for**
- 11: **end while**

Yang et al. (2025) shows explicit construction of a 2 layered transformer with multiple heads that can be trained to solve the reachability problem on trees. We show that a modified construction can be used to solve the Tower of Hanoi with the above algorithms.

C.1.1 EQUIVALENCE TO PATTERNED GRAPH REACHABILITY

The above algorithm shows that solving graph reachability where the path is just $(abc)^n$ with a, b, c being colored edges. Hence out path finding task for solving Tower of Hanoi is the following:

Path finding

Embedding For each node i , we represent by $a_i \in \mathbb{R}^{d_1}$ to be its vertex encoding. Suppose that $\{1, 0, 0\}, \{0, 1, 0\}, \{0, 0, 1\}$ is the color encoding of the 3 edge colors, we now represent the edge $e = (p(i), i)$ with color c_i by encoding it as $(a_{p(i)}^\top, 000, a_i^\top, c_i^\top)^\top$

C.1.2 TRANSFORMER ARCHITECTURE

As in Yang et al. (2025) the architecture for a single layer transformer with $H \geq 1$ heads with trainable parameters θ is as follows:

Given the input embedding matrix $E_{d \times n}$ for every head $h \in [H]$ we have the query, key, value $W_{d \times d}^K, W_{d \times d}^Q, W_{d \times d}^V$ matrices that give the the head output:

$$\text{head}_h(E) = W^V E \cdot \text{softmax}\left(E^\top W_h^{K^\top} W_h^Q E\right)$$

$\text{softmax}\left(E^\top W_h^{K^\top} W_h^Q E\right)$ gives the attention matrix on the previous tokens (each of the n d-dimensional tokens).

$$f(E; \theta) = W^O \begin{pmatrix} \text{head}_1(E) \\ \vdots \\ \text{head}_H(E) \end{pmatrix},$$

We consider step-by-step reasoning in an autoregressive manner. At each reasoning step $k > 0$, given the input embedding $E^{(k)}(G)$, the output vector is taken as the last column of the output matrix, which is then concatenated with the input embedding to form the new input to the next reasoning step,

$$\hat{\delta}^{(k+1)}(G, \theta) = f(E^{(k)}(G), \theta)_{:-1}, \quad E^{(k+1)}(G) = (E^{(k)}(G), \hat{\delta}^{(k+1)}(G, \theta)) \quad (1)$$

C.1.3 EXPRESSIBILITY FOR BACKWARD REASONING

Goal to root:

For the backward reasoning task, we are interested in outputting the path from the goal node to the root node (g2r). Let the input embedding of the transformer be

$$E = E_{\text{g2r}}(\mathcal{G}) = \begin{pmatrix} X(\mathcal{G}) \\ Y(\mathcal{G}) \end{pmatrix} = \begin{pmatrix} x_1 & \cdots & x_{l(\mathcal{G})} & a_{g(\mathcal{G})} \\ 0 & \cdots & 0 & c_g \\ y_1 & \cdots & y_{l(\mathcal{G})} & a_0 \\ c_1 & \cdots & c_{l(\mathcal{G})} & 0 \end{pmatrix} \in \mathbb{R}^{2d_l \times (l(\mathcal{G})+1)}, \quad (2)$$

where $a_0 \in \mathbb{R}^{d_l}$ is used to fill the empty slots. Note that c is the embedding of the color and this can be of size 3 since there are only 3 possible colors We set the ground label of \mathcal{T} as the embedding of the reversed path:

$$O_{\text{g2r}}(\mathcal{G}) = \begin{pmatrix} a_{p^1(g(\mathcal{T}))} & a_{p^2(g(\mathcal{T}))} & \cdots & a^{r(\mathcal{T})} \\ c_{p^1(g(\mathcal{T}))} & c_{p^2(g(\mathcal{T}))} & \cdots & c_{r(\mathcal{T})} \\ a_{g(\mathcal{T})} & a_{p(g(\mathcal{T}))} & \cdots & a_{p_{m(\mathcal{T})-1}(g(\mathcal{T}))} \\ 0 & \cdots & 0 & 0 \end{pmatrix} \in \mathbb{R}^{2(d_l+3) \times m(\mathcal{T})}. \quad (3)$$

We consider a one-layer single-head transformer, where $H = 1$. We impose the following parameters: let $W_1^V = W_1^K = I_{2d_l}$, and set

$$W^O = \begin{pmatrix} I_{d_l} & 0_{d_l \times 3} & 0_{d_l \times d_l} & 0_{d_l \times 3} \\ 0_{3 \times d_l} & 0_{3 \times 3} & 0_{3 \times d_l} & P_{3 \times 3} \\ 0_{d_l \times d_l} & 0_{d_l \times 3} & I_{d_l} & 0_{d_l \times 3} \\ 0_{3 \times d_l} & 0_{3 \times 3} & 0_{3 \times d_l} & 0_{3 \times 3} \end{pmatrix}, P = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \quad (4)$$

$$W_1^Q = \begin{pmatrix} 0_{(d_l+3) \times (d_l+3)} & 0_{(d_l+3) \times (d_l+3)} \\ B^* & 0_{(d_l+3) \times (d_l+3)} \end{pmatrix}, B^* = \begin{pmatrix} B_{d_l \times d_l} & 0_{d_l \times 3} \\ 0_{3 \times d_l} & I_3 \end{pmatrix} \quad (5)$$

where $B \in \mathbb{R}^{d_l \times d_l}$ is trainable. In this case, $\theta = \{B\}$.

Multi-step reasoning. We set the first step $E_{g_{2r}}^{(0)} = E_{g_{2r}}(\mathcal{T})$, and let

$$X^{(k-1)} := E_{g_{2r}}^{(k-1)}(1 : d_l + 3, :) \in \mathbb{R}^{(d_l+3) \times (l(\mathcal{T})+k)}$$

$$Y^{(k-1)} := E_{g_{2r}}^{(k-1)}(d_l + 4 : 2d_l + 6, :) \in \mathbb{R}^{(d_l+3) \times (l(\mathcal{T})+k)}.$$

Then according to the transformer architecture described above, we have the output $\hat{o}^{(k)}$ at reasoning step k as

$$\hat{o}^{(k)} = W^O \begin{pmatrix} X^{(k-1)} \\ Y^{(k-1)} \end{pmatrix} \cdot \text{softmax}\left(Y^{(k-1)T} B x_{l-1}^{(k-1)}\right), \quad (6)$$

where $x_{l-1}^{(k-1)}$ is the last column of $X^{(k-1)}$. Reasoning recursively for $m(\mathcal{T})$ steps according to (2), we obtain the output

$$\hat{O}_{g_{2r}}(\mathcal{G}; \theta) = (\hat{o}^{(1)}, \dots, \hat{o}^{(m(\mathcal{G}))}).$$

We make the following assumption on the node embeddings for the backward reasoning case.

Assumption 1 (Linear independent embeddings (backward reasoning)). *Suppose $a_0 = 0_{d_l}$, and $\{a_i\}_{i \in [S]}$ are linearly independent.*

Let $A := (a_1, \dots, a_S) \in \mathbb{R}^{d_l \times S}$ be the embedding matrix. The following theorem provides a construction of the transformer that solves backward reasoning.

Theorem C.1 (Construction for backward reasoning). *Under Assumption 1, for any $\alpha \in \mathbb{R}$, there exists $B = B_\alpha \in \mathbb{R}^{d_l \times d_l}$ such that*

$$A^\top B A = \alpha I_S. \quad (7)$$

Let $\theta = \{B_\alpha\}$, then for any tree \mathcal{T} , we have

$$\hat{O}_{g_{2r}}(\mathcal{G}; \theta) \rightarrow O_{g_{2r}}(\mathcal{G}) \quad \text{as } \alpha \rightarrow +\infty.$$

The above when $\alpha \rightarrow +\infty$ ensures that

$$\begin{pmatrix} X^{(k-1)} \\ Y^{(k-1)} \end{pmatrix} \cdot \text{softmax}\left(Y^{(k-1)T} B^* x_{:-1}^{(k-1)}\right) = \begin{pmatrix} a_{\rho^{(k)}(g)} \\ 0 \\ a_{\rho^{(k-1)}(g)} \\ c_{\rho^{(k-1)}(g)} \end{pmatrix} \quad (8)$$

Note that W^O performs the affine transform needed to give

$$W^O \begin{pmatrix} a_{\rho^{(k)}(g)} \\ 0 \\ a_{\rho^{(k-1)}(g)} \\ c_{\rho^{(k-1)}(g)} \end{pmatrix} = \begin{pmatrix} a_{\rho^{(k)}(g)} \\ c_{\rho^{(k)}(g)} \\ a_{\rho^{(k-1)}(g)} \\ 0 \end{pmatrix} \quad (9)$$

D EXPERIMENTAL DETAILS

We provide the exact prompts used in all experiments to ensure full reproducibility. All experiments were conducted using DeepSeek-R1 Guo et al. (2025b).

Prompt for the Non-recursive Algorithm - Tower of Hanoi

You are a deterministic Python code execution engine.
Your task is to EXECUTE the Python code below exactly as written and return the FINAL OUTPUT.

STRICT RULES (must follow all):

- Do NOT explain the code.
- Do NOT add any commentary, text, or formatting.
- Do NOT modify the code.
- Do NOT abbreviate or partially print the output.
- Assume all required imports (e.g., math) are available.
- Follow Python operator precedence exactly.
- Return ONLY the final Python object produced by the function call.
- The output must be exactly what Python would produce.

Code:

```
def get_right_most_bit(n):
    if n == 0:
        return 0
    res = n & (~n + 1)
    return int(math.log2(res)) + 1

def to_h(start_step, end_step):
    steps = []
    for n in range(start_step, end_step + 1):
        num_1 = get_right_most_bit(n)
        num_2 = (n & (n - 1)) % 3
        num_3 = ((n | (n - 1)) + 1) % 3
        steps.append([num_1, num_2, num_3])
    return steps
```

Function call:

```
to_h(3, 4)
n = 3
Binary: 11
get_right_most_bit(3):
~3 + 1 = -3
3 & (-3) = 1
log2(1) + 1 = 1
num_1 = 1
num_2:
n & (n - 1) = 3 & 2 = 2
2 % 3 = 2

num_3:
n | (n - 1) = 3 | 2 = 3
(3 + 1) % 3 = 1
```

Result for $n = 3$:

```
[1, 2, 1]
```

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Prompt for the Graph Algorithm - Tower of Hanoi

You are a BFS agent that continues a breadth-first search (BFS) algorithm.

You will be given a HISTORY in the following format:

```
{
  "queue": [...], # frontier as (node, path) pairs,
  "visited": [...], # nodes already visited,
  "parent_map": {...}, # child → parent mapping (if any),
  "adj_retrieved": {...} # adjacency lists fetched so far
}
```

You have access to TOOLS that can fetch adjacency lists of nodes.

YOUR JOB:

- Continue the BFS algorithm step-by-step.
- Use tools only when adjacency info is missing.
- Update the BFS state deterministically.
- When the goal node is found, reconstruct and output the path.

YOUR OUTPUT FORMAT RULES:

1. If BFS is still ongoing:

- Return ONLY the updated state JSON (no extra text).
- The format must be exactly:

```
{
  "queue": [...],
  "visited": [...],
  "parent_map": {...},
  "adj_retrieved": {...}
}
```

2. If BFS is finished and a path is found:

- Return ONLY:
FINAL_PATH:[v1, v2, v3, ...]

3. Do NOT include any explanations, reasoning, or commentary.

4. Do NOT output markdown, quotes, or prose — only raw JSON or FINAL_PATH line.

You must always return one of the two:

- The new BFS state JSON, or
- The final path (if done).

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

Prompt for River crossing

You are given a river crossing puzzle involving multiple agents (A_1, A_2, \dots) and their corresponding actors (a_1, a_2, \dots).

Your task is to compute a correct list of boat moves that solves the puzzle, following the rules below:

- Each agent A_i and actor a_i must end up on the right side of the river.
- The boat can carry at most two individuals at a time.
- The boat must always have at least one person on board to move.
- No actor must be left with agents other than its own without its own agent present, even if there are other actors present on either of the shores or the boat.

For example, a_1 cannot be left on the shore with (A_2, a_2, A_3) as A_1 is not present.

- Actors on the boat do not have to get down unless explicitly dropped off. For example, in the sequence

```
[["A_2", "a_2"], ["a_2"], ["a_2", "a_1"], ...]
```

no rule is violated as a_2 never gets off the boat after the second step.

2. Entity Naming

- Actors: a_1, a_2, \dots, a_n
- Agents: A_1, A_2, \dots, A_n

3. Output Format

Your final output must be a Python-style list of lists of strings. For example:

```
[["A_2", "a_2"], ["a_2"], ["a_2", "a_1"], ["a_1"], ["A_1", "a_1"]]
```

4. Algorithm to Follow

```
moves = [["A_1", "a_1"]]
for i in range(1, n):
    moves.append([f"a_{i}"])
    moves.append([f"a_{i}", f"a_{i+1}"])
    moves.append([f"a_{i+1}"])
    moves.append([f"A_{i+1}", f"a_{i+1}"])
return moves
```

This means the general sequence will look like:

```
[["A_1", "a_1"], ["a_1"], ["a_1", "a_2"], ["a_2"], ["A_2", "a_2"], ["a_2"], ["a_2", "a_3"], ["a_3"], ["A_3", "a_3"], ...]
```

FOLLOW THE ALGORITHM EXACTLY.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

Prompt for Checkers

You are a helpful assistant. Solve this puzzle for me.

On a one-dimensional board, there are red checkers ('R'), blue checkers ('B'), and one empty space ('_').

A checker can move by either:

1. Sliding forward into an adjacent empty space, or
2. Jumping over exactly one checker of the opposite color to land in an empty space.

The goal is to swap the positions of all red and blue checkers, effectively mirroring the initial state.

Example: If the initial state is ['R', '_', 'B'], the goal is to reach ['B', '_', 'R'].

Your output should be a string or list of move symbols, where each symbol corresponds to a valid move made:

'S' = red checker (to immediate left of blank) slides right.
's' = blue checker (to immediate right of blank) slides left.
'J' = red checker (to the 2nd left of the blank) jumps right over blue.
'j' = blue checker (to the 2nd right of the blank) jumps left over red.

Example solution for the initial state ['R', '_', 'B'] might be:

```
moves = ['S', 'j', 'S']
```

This means:

'S': Red checker steps into the blank on its right.
'j': Blue checker jumps left over red into the new blank.
'S': Red checker steps into the blank on its right.

Your goal is to output the full sequence of symbolic moves to achieve the mirrored configuration.

Ensure the solution is correct and minimal (i.e., optimal number of moves).

D.1 ARCHITECTURAL DETAILS OF THE RASP MODEL

Table 1: Architectural details for the RASP compiled model at $n = 20$

Parameter	Value	Notes
num_layers	7	One active attention layer and six additional blocks implementing arithmetic in the MLPs.
num_heads	4	Four-head attention used to broadcast metadata tokens in layer 0.
key_size	27	Scales approximately linearly with max_disks ($\approx \text{max_disks} + 7$).
mlp_hidden_size	1804	Dominant width term, scaling quadratically with max_disks.
residual_dim	190	Dimension of the joint residual feature space.

Key dimensions as max_disks grows (single-step compilation).

	max_disks	Residual dim	key_size	mlp_hidden_size
1026				
1027	1	61	8	28
1028	2	63	9	40
1029	3	70	10	70
1030	5	85	12	154
1031	10	120	17	504
1032	20	190	27	1804

1033 **Closed-form trends (matching observed data for `max_disks` ≥ 2).**

$$1034 \text{ residual_dim}(n) \approx 7n + 50,$$

$$1035 \text{ key_size}(n) = n + 7,$$

$$1036 \text{ mlp_hidden_size}(n) = 4n^2 + 10n + 4.$$

1037 The quadratic MLP width dominates parameter growth; the residual dimension and attention key size
1038 increase only linearly to accommodate additional scratch slots and positional bases.

1042 D.2 RASP CODE

1043 `"""Compile Tower of Hanoi peg configuration into a Tracr transformer.
1044 Compiles into transformer with 1 attention layer with 4 heads and
1045 6 MLP layers. No layer norm, dropout and non causal masking.
1046 """`

1047 The RASP program implements the formula for determining the peg index of each disk
1048 at a given step in the optimal solution to the Tower of Hanoi problem.
1049

1050 The formula is:

$$1051 s_d \leftarrow (\text{ell}/2^d + 1/2 \times (j - i) \times ((n - d) \bmod 2 + 1) + i) \bmod 3$$

1052 Where:

- 1053 - `s_d`: peg index for disk `d` at step
- 1054 - `ell`: step index
- 1055 - `d`: disk number
- 1056 - `n`: total number of disks
- 1057 - `i`: starting peg index
- 1058 - `j`: target peg index
- 1059 - `x`: floor function

1060 `"""`

1061 `BOS_TOKEN = "BOS"`

1062 `# Positions 0-4 hold metadata tokens; everything from index 5 onward is scratch
1063 # space used to write the peg index for each disk.`

1064 `SCRATCH_OFFSET = 5`

1065 `def _broadcast_value(position: int) -> rasp.SOp:`

1066 `"""Broadcast a token at `position` across the entire sequence."""`

1067 `selector = rasp.Select(
1068 rasp.indices,
1069 rasp.indices,
1070 lambda key, unused_query, pos=position: key == pos)
1071 return rasp.Aggregate(selector, rasp.tokens)`

1072 `def _disk_number_sop() -> rasp.SOp:`

1073 `"""Return disk numbers (1-indexed) for scratch-pad positions, else None."""`

1074 `indices = rasp.indices
1075 return rasp.Map(
1076 lambda idx: (idx - (SCRATCH_OFFSET - 1)) if idx >= SCRATCH_OFFSET else None,
1077 indices)`

```

1080     indices)
1081
1082
1083 def build_hanoi_program() -> rasp.SOp:
1084     """Construct a RASP program that writes peg indices for each disk."""
1085
1086     n_tokens = _broadcast_value(0).named("n")
1087     start_peg = _broadcast_value(1).named("start_peg")
1088     target_peg = _broadcast_value(2).named("target_peg")
1089     step_index = _broadcast_value(3).named("step_index")
1090
1091     disk_numbers = _disk_number_sop().named("disk_number")
1092
1093     # Compute round(step_index / 2^disk) modulo 3 using bit operations.
1094     upper_bits_mod3 = rasp.SequenceMap(
1095         lambda ell, disk: ((int(ell) >> int(disk)) % 3)
1096         if isinstance(ell, int) and isinstance(disk, int) else None,
1097         step_index,
1098         disk_numbers,
1099     ).named("upper_bits_mod3")
1100
1101     half_threshold_bit = rasp.SequenceMap(
1102         lambda ell, disk: ((int(ell) >> (int(disk) - 1)) & 1)
1103         if isinstance(ell, int) and isinstance(disk, int) and int(disk) > 0 else None,
1104         step_index,
1105         disk_numbers,
1106     ).named("half_threshold_bit")
1107
1108     rounded_mod3 = rasp.SequenceMap(
1109         lambda upper, half: (upper + half) % 3 if upper is not None and half is not None
1110         else None,
1111         upper_bits_mod3,
1112         half_threshold_bit,
1113     ).named("rounded_mod3")
1114
1115     peg_delta_mod3 = rasp.SequenceMap(
1116         lambda target, start: (int(target) - int(start)) % 3
1117         if isinstance(target, int) and isinstance(start, int) else None,
1118         target_peg,
1119         start_peg,
1120     ).named("peg_delta_mod3")
1121
1122     n_minus_d_parity = rasp.SequenceMap(
1123         lambda n_val, disk: (int(n_val) - int(disk)) & 1
1124         if isinstance(n_val, int) and isinstance(disk, int) else None,
1125         n_tokens,
1126         disk_numbers,
1127     ).named("n_minus_d_parity")
1128
1129     parity_plus_one_mod3 = rasp.Map(
1130         lambda parity: (int(parity) + 1) % 3 if isinstance(parity, int) else None,
1131         n_minus_d_parity,
1132     ).named("parity_plus_one_mod3")
1133
1134     parity_term_mod3 = rasp.SequenceMap(
1135         lambda base, start: (int(base) + int(start)) % 3
1136         if isinstance(base, int) and isinstance(start, int) else None,
1137         parity_plus_one_mod3,
1138         start_peg,
1139     ).named("parity_term_mod3")

```

```
1134
1135 first_product_mod3 = rasp.SequenceMap(
1136     lambda a, b: (int(a) * int(b)) % 3
1137     if isinstance(a, int) and isinstance(b, int) else None,
1138     rounded_delta_mod3,
1139     peg_delta_mod3,
1140 ).named("first_product_mod3")
1141
1142 mod_three = rasp.SequenceMap(
1143     lambda a, b: (int(a) * int(b)) % 3
1144     if isinstance(a, int) and isinstance(b, int) else None,
1145     first_product_mod3,
1146     parity_term_mod3,
1147 ).named("mod_three")
1148
1149 is_scratch = rasp.Map(
1150     lambda idx: idx >= SCRATCH_OFFSET, rasp.indices).named("is_scratch")
1151
1152 scratch_values = rasp.SequenceMap(
1153     lambda scratch, value: value if scratch else None,
1154     is_scratch,
1155     mod_three,
1156 ).named("scratch_values")
1157 final_output = rasp.SequenceMap(
1158     lambda new, original: new if new is not None else original,
1159     scratch_values,
1160     rasp.tokens,
1161 ).named("hanoi_configuration")
1162
1163 return final_output
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
```