


Large Language Models as Graph Computational Solvers via Topology-aware Residual Attention

Wei Zhuo Siqiang Luo

Nanyang Technological University

Motivation: graph serialization and reasoning

A six-node connectivity example from the paper



Original description

(0,1) (1,2) (3,4) (4,5)

Is there a path between **1** and **4**?

No correct

Relabeled + shuffled

(1,3) (2,4) (3,5) (4,0)

Equivalent query: path between **4** and **3**?

Yes incorrect

Reported outputs from Llama-3.1-8B-Instruct

Input symmetry

Equivalent graphs, consistent answers

Graph topology

Relations survive serialization


Legal execution

Each step follows the graph

TRACER: framework overview

Graph structure shapes attention and post-training shapes behavior

Graph + query Equivalent views



Input Align equivalent graph views


Representation Add structural attention

Reasoning Reward traces and answers

Topology-aware Residual Attention (TRA)

$$A = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_k}} + \alpha B \right)$$

$B \in \mathbb{R}^{T \times T}$ is a token-pair structural bias; $\alpha \geq 0$ is learned.



Local-hop formulation

For graph nodes mentioned by tokens i, j

Same node	φ_{same}
1–2 hops apart	$\varphi_{\text{dist}}(D_G)$
Otherwise	0

Graph-node token pairs only

If either token is non-node text, $B_{ij} = 0$.

Efficient TRA through query–key augmentation

$$B = E\Phi E^\top \approx (EL)(ER)^\top$$

E : token-to-node map Φ : node relations L, R : compact factors

$$\tilde{Q} = \left[\frac{Q}{\sqrt{d_k}}, \sqrt{\alpha} EL \right]$$

$$\tilde{K} = [K, \sqrt{\alpha} ER]$$

Content channels

Structure

Content channels

Structure

d_k content features + r structural features

Standard Flash-Attention


softmax scale = 1

No additional dense $T \times T$ bias

Exact if $LR^\top = \Phi$. Reported implementation: learned rank-8 landmark factors.

These factors approximate structural relations; non-node rows remain zero through E .

Multi-serialization supervised fine-tuning



Relabel the query, trace and answer consistently

Same graph, matched targets

Sample 4 serializations.

Relabel the query, trace and answer together.

Read a graph summary

Use the final-layer state h_i of the [GRAPH] token.

Train both behaviors

Predict the trace and answer.


Align summaries across serializations.

$$\mathcal{L}_{\text{SFT}} = \mathcal{L}_{\text{task}} + \lambda_{\text{repr}} \mathcal{L}_{\text{repr}}$$

$$\mathcal{L}_{\text{repr}} = \frac{1}{k-1} \sum_{i=2}^k (1 - \cos(h_1, h_i))$$

Reinforcement learning with trace-level rewards

1 Generate



Several traces per prompt

2 Score

Correctness

Final answer matches ground truth

Consistency

Canonical traces agree with correct peers

Process

Train on valid/invalid step labels
Average learned verifier scores

3 Update

GRPO


Center rewards
within each prompt
Update the policy

Consistency is rollout-specific

Compare canonical traces with correct peers from other serializations.
This signal can survive GRPO's within-prompt reward centering.

Main results on GraphInstruct

GraphInstruct, 400 test instances per task, exact-match answers



67.31%

TRACER average accuracy

+7.98 pp

over GraphWiz-DPO

59.33% average accuracy

63.31% before RL

TRACER-SFT

Full TRACER leads on **7 of 9** tasks in the reported comparison.

Serialization consistency and execution validity

■ GraphWiz-DPO ■ TRACER

Across equivalent serializations

Answer consistency



Equivalent answers across sampled serializations

Trace agreement



LCS similarity after canonicalizing node labels

Under symbolic trace replay

Step validity



Legal transitions under symbolic replay

All-step validity



Every step must be parsed and legal (ISR)

Trace agreement and legal execution measure different properties.

Ablation: architecture and training

Average exact-match accuracy, same Llama-3.1-8B backbone

SFT with and without each component

	Single serialization	Multi-serialization + alignment
Standard attention	54.56	60.72
TRA	56.67	63.31

+2.59 pp from TRA under matched multi-serialization SFT

Further gains from RL


63.31 %
TRACER-SFT

+4.00 pp

67.31 %
Full TRACER

Task transfer and graph-size scaling


Each target task held out of training



Mean: **32.25%** Naive SFT vs **50.00%** TRACER

Hamiltonian Path remains difficult: **33.50%** in the main benchmark.

Separate graph-size sweep



Accuracy falls as graphs grow.
Evaluation covers at most **100 nodes**.

Conclusion

- ▶ **Expose structure**


Topology-aware residual attention

- ▶ **Respect input symmetry**

Multi-serialization representation alignment

- ▶ **Train the reasoning process**

Consistency and process rewards with GRPO



github.com/JhuoW/TRACER

Code and research resources

Open challenge: reliable execution on larger graphs and harder search problems