ICLR 2026 STRUCTURED REASONING FOR LLMS: A UNIFIED FRAMEWORK FOR EFFICIENCY AND EXPLAINABILITY ICLR 2026

Yubo Dong, Hehe Fan, Linchao Zhu

1. Structured Reasoning Data & Tuning

(based on <IMAGE-0> - Data Collection)

Input: What is √53 in simplest radical form? Please reason step by step using structured reasoning tags.

Reasoning Steps:

<assumption>
To find the simplest radical form of √53. I need to check if S3 has any perfect square factors.

<decompose>
Lett a prime number with factors 1 and 53.

<verify>
Since 53 is prime, no perfect squares greater than 1 divide it.

<inference>
Therefore, √53 cannot be simplified further, √53 is siready in its simplest form.

<verify>
Double checking: If we had √53 = a/b where b is square free, then a'b = 53.

<conclusion>
The simplest radical form is √53.

Output: $\sqrt{53}$

Output: $\sqrt{53}$

Data Pipeline:

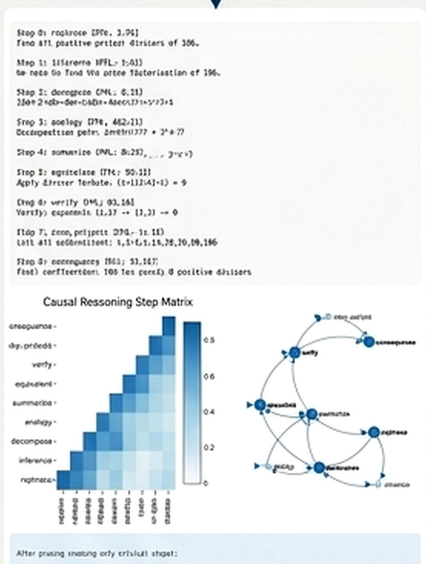

Structured Tuning Dataset:

- 500 high-quality samples from S1k, DeepScaleR.
- 23 Reasoning Tags.

2. Methods - Traceability & Optimization

(based on <IMAGE-3>, <IMAGE-4>, and <IMAGE-2> Step dependency/Optimization

Layer-wise step Attention (Left)

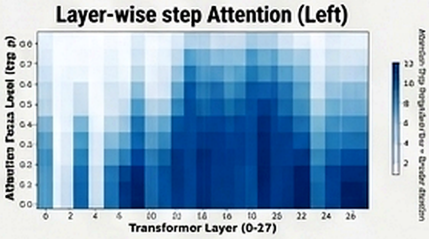

1.5B Models (Right)

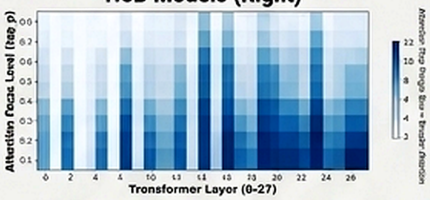

local transition to global transition Local at Layer 14.

Middle layers integrate broader reasoning context.

A. MAX-Flow Algorithm: sparse graphs via ΔF_k (node importance)

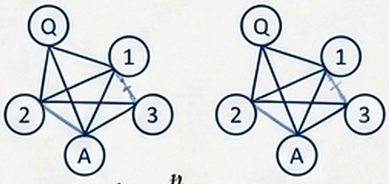

$$\Delta F_k = \frac{1}{HT_i} \sum_{h=1}^{n} \sum_{acT_i} \max_{bcT_i} A_{h,a,b}$$

B. LCS Algorithm: optimal common subsequences from multiple generated responses

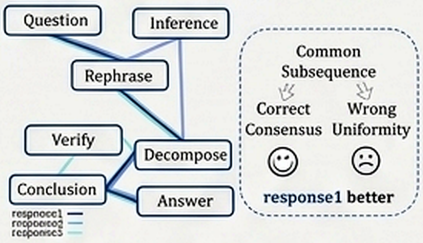

MAX-Flow ΔF_k outperforms other filters in EFE.

3. General Benchmarks

(based on <IMAGE-0>

A) Efficiency (1.5B)

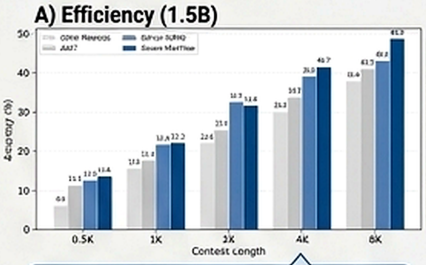

better performance with fewer steps at longer context lengths.

A) Efficiency (7B)

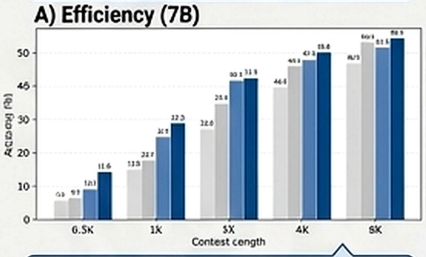

better performance engths.
MATH500 7B MaxFlow: 92.67% accuracy

B) Stability (1.5B)

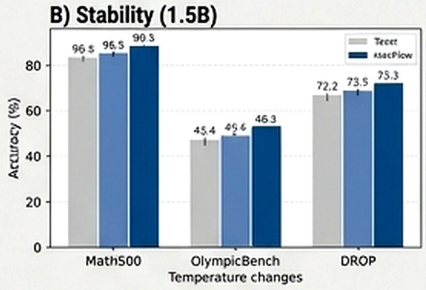

C) Interpretability (7B)

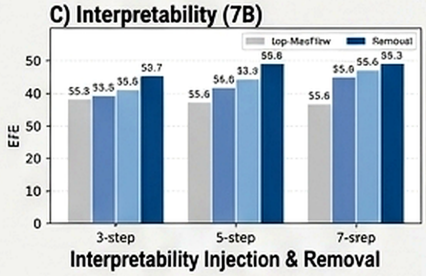

Interpretability Injection & Removal

top-matrix filter outperforms random and perplexity filter.

Conclusion

Structured Reasoning improves efficiency, stability, and interpretability.

Future: redundant layer pruning.

Project: https://cnsdqd-dyb.github.io/structured-reasoning/