

A Context-Enhanced Framework for Sequential Graph Reasoning

Shuo Shi, Chao Peng*, Chenyang Xu* and Zhengfeng Yang

Shanghai Key Laboratory of Trustworthy Computing, Software Engineering Institute, East China Normal University, Shanghai, China

51255902122@stu.ecnu.edu.cn,
{cpeng, cyxu, zfyang}@sei.ecnu.edu.cn

Abstract

The paper studies sequential reasoning over graph-structured data, which stands as a fundamental task in various trending fields like automated math problem solving and neural graph algorithm learning, attracting a lot of research interest. Simultaneously managing both sequential and graph-structured information in such tasks presents a notable challenge. Over recent years, many neural architectures in the literature have emerged to tackle the issue. In this work, we generalize the existing architectures and propose a context-enhanced framework. The crucial innovation is that the reasoning of each step does not only rely on the outcome of the preceding step but also leverages the aggregation of information from more historical outcomes. The idea stems from our observation that in sequential graph reasoning, each step’s outcome has a much stronger inner connection with each other compared to traditional seq-to-seq tasks. We show that the framework can effectively integrate with the existing methods, enhancing their reasoning abilities. Empirical evaluations are conducted on the challenging CLRS Reasoning Benchmark, and the results demonstrate that the proposed framework significantly improves the performance of existing architectures, yielding state-of-the-art results across the majority of the datasets within the benchmark.

1 Introduction

Recent years have witnessed a steady growth in exploring the reasoning capacity of deep learning [He *et al.*, 2023; Jiang *et al.*, 2022; de Rezende *et al.*, 2021]. This research domain holds profound significance in the field of artificial intelligence — it enables AI not only to recognize patterns but also to deduce new knowledge from existing information, thereby broadening the capacity for complex problem-solving. As a variety of reasoning tasks can be formulated into graph problems, there is an increasing focus and attention on deep graph reasoning [Wang *et al.*, 2023; Ling *et al.*, 2022;

Zhu, 2022]. This trend delves into the study of graph-structured data and aims to build specialized network structures that can learn to reason over them.

In this paper, we consider *sequential graph reasoning*, one of the most challenging tasks in deep graph reasoning. Diverging from conventional reasoning tasks like graph (node) classification [Wu *et al.*, 2019], this type of task goes beyond the one-shot model and targets step-by-step reasoning over graphs. The applications of seq-graph reasoning span across various domains in artificial intelligence. For instance, in AI-powered math problem solving [Lample and Charton, 2020; Huang *et al.*, 2022; Drori *et al.*, 2020], neural networks are expected to provide the entire process of solving a mathematical problem, rather than a mere final answer. Another emerging application is neural algorithm learning [Diao and Loynd, 2023; Ibarz *et al.*, 2022; Jain *et al.*, 2023], where networks are anticipated to sequentially imitate individual steps of classical graph algorithms.

There have been many recent advances in seq-graph reasoning [Li *et al.*, 2023; Zhao *et al.*, 2024; Gerasimova *et al.*, 2023]. While the proposed neural architectures exhibit diversity, they generally adhere to a common sequential paradigm. In the framework, graph-structured hidden states are maintained in memory. At each step, the outcome of the preceding step serves as input, integrating with the hidden states for reasoning. Variations among existing works arise in the specifics of their integration processes. Numerous graph networks (e.g. GAT [Velickovic *et al.*, 2018] and MPNN [Gilmer *et al.*, 2017]) that excel at one-shot reasoning have been investigated.

The current paradigm is essentially a direct combination of the techniques for processing serialized data and those for handling graph-structured data. We notice that such a straightforward combination may fall short of yielding competitive reasoning capabilities due to the distinctive nature of seq-graph reasoning. Compared to traditional seq-to-seq tasks [Yu, 2017], each step’s outcome in seq-graph reasoning has a much stronger intrinsic connection with the others. For example, in automatic geometric problem solving [Schreck and Mathis, 2016], the outcome sequence forms an entire problem-solving process, which can even be represented by an extra reasoning graph. However, most existing works fail to capture these intrinsic connections.

*Corresponding authors.

1.1 Our Contributions

Motivated by the observation above, this paper proposes a framework to capture the contextual patterns of step outcomes in seq-graph reasoning. To evaluate the reasoning ability of this framework, we conduct empirical evaluations on the challenging CLRS Reasoning Benchmark [Velickovic *et al.*, 2022]. This benchmark consists of 30 different algorithmic tasks for probing the reasoning capabilities of graph-structured networks. The contributions of the paper are summarized as follows:

- We introduce a *Context-Enhanced* Framework (CEF) for seq-graph reasoning. The framework extends the existing paradigm and exhibits notable flexibility, allowing for adjustments based on available memory space during implementation.
- We show that the framework can effectively integrate with various existing architectures. Specifically, the paper categorizes existing methods into two types: GNN-based and Transformer-based, and provides detailed CEF implementations for both.
- We perform a diverse range of experiments on the CLRS benchmark. The results demonstrate that, for both types of existing architectures, our framework significantly enhances their performances across the majority of the datasets within the benchmark. When integrated with the latest architectures, we achieve new state-of-the-art results.

1.2 Related Work

Graph Neural Networks. Graph neural networks have emerged as a powerful tool for effectively capturing and leveraging the structural relationships present in graph-structured data. Numerous neural architectures have been proposed in this domain, such as Graph Convolution Networks (GCN) [Kipf and Welling, 2017], Graph Attention Networks (GAT) [Velickovic *et al.*, 2018], Message Passing Neural Networks (MPNNs) [Gilmer *et al.*, 2017] and so on. These architectures excel not only in one-shot graph reasoning [Ma *et al.*, 2023] but also demonstrate comparable performance in seq-graph reasoning tasks [Ramesh *et al.*, 2023]. Recently, [Ibarz *et al.*, 2022] introduced a novel GNN architecture called triplet-GMPNN. They augment the classical MPNN to perform message passing toward edges and have achieved remarkable results in neural algorithmic reasoning.

Graph Transformers. Transformers and its variants play a crucial role in domains such as natural language processing, computer vision, and time series analysis [Vaswani *et al.*, 2017]. Consequently, leveraging Transformers for tasks related to graph reasoning has emerged as a recent research focus, leading to the development of several variants of Transformers tailored for handling graph-structured data [Wu *et al.*, 2021; Hussain *et al.*, 2022]. A recent advancement in this direction is proposed by [Diao and Loynd, 2023]. They introduced Relational Transformers, which enhance attention computation by aggregating edge and graph feature information with node features.

Beyond One-Shot Graph Reasoning. The architectures mentioned above are fundamentally tailored for one-shot reasoning. When addressing seq-graph reasoning tasks, they necessitate the paradigm stated in the introduction. There have been some recent advances in exploring competitive paradigms for seq-graph reasoning [Numeroso *et al.*, 2023; Jain *et al.*, 2023; Jürß *et al.*, 2023]. Notably, the studies most closely related to our research were recently introduced by [Jain *et al.*, 2023] and [Jürß *et al.*, 2023], where the GNN architectures are augmented with a priority queue and a stack, respectively.

1.3 Paper Organization

The paper is organized as follows. Section 2 provides preliminaries, including a problem definition and a paradigm widely used in the literature. In Section 3, we introduce our framework, and subsequently, Section 4 details the concrete implementations of our framework for different types of existing approaches. Empirical evaluations are presented in Section 5. Finally, Section 6 concludes the paper.

2 Preliminaries

The section formalizes the problem setup and introduces a widely-used paradigm. In a seq-graph reasoning task, we are given graph-structured data and required to perform reasoning over it at most T times. The given data is represented by a graph $G = (V, E)$ with n nodes and m edges, where each node $v \in V$ has feature $x_v^{(0)}$ and each edge $e \in E$ has feature $x_e^{(0)}$. For the notational simplicity, define $X_V^{(0)} := \{x_v^{(0)}\}_{v \in V}$, $X_E^{(0)} := \{x_e^{(0)}\}_{e \in E}$, and $\mathbf{X}^{(0)} = \{X_V^{(0)}, X_E^{(0)}\}$. Further, for each training data, we are provided with the ground truth $\mathbf{Y}^{(t)}$ of each reasoning step $t \in [T]$.

A widely used paradigm in the literature is the *encode-process-decode* paradigm [Ibarz *et al.*, 2022]. In this paradigm, there are three modules: encoder \mathcal{E} , processor \mathcal{P} , and decoder \mathcal{D} . In each step $t \in [T]$, the input tensors $\mathbf{X}^{(t-1)}$ undergo three successive transformations $\mathcal{E}, \mathcal{P}, \mathcal{D}$ and then yield $\mathbf{X}^{(t)}$, which serves as the input for the next step.

Encoder Module. There are usually two encoder networks in this module¹, each tasked with encoding node features and edge features respectively. Thus, after feeding $\mathbf{X}^{(t-1)}$ to the encoder module, we obtain latent features $\mathbf{L}^{(t)}$ for nodes and edges in the graph:

$$L_V^{(t)}, L_E^{(t)} = \mathcal{E}(\mathbf{X}^{(t-1)}). \quad (1)$$

Processor Module. This module is the central component of the entire paradigm, where most prior works primarily vary. It consists of a (standard) graph-structured network (e.g., graph neural networks or graph transformers), and maintains a hidden state for each node and edge. Denote by $H_V^{(t-1)}, H_E^{(t-1)}$ the hidden states of the node set and the edge

¹Some prior works add one extra encoder network to handle the global feature of the graph.

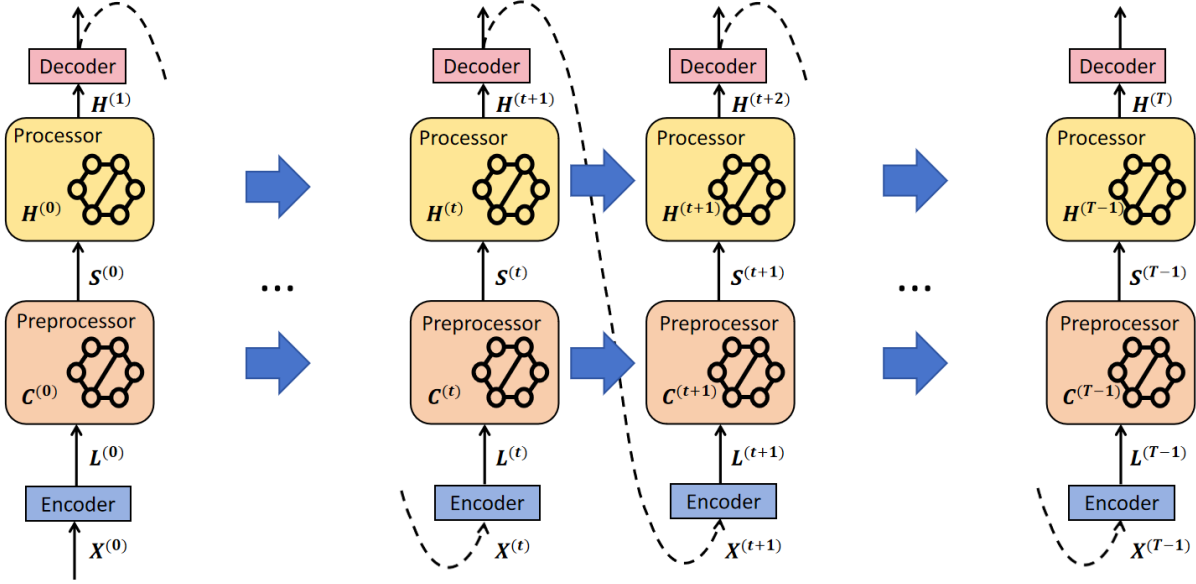


Figure 1: An illustration of the CEF framework. The solid and dashed arrows in the figure indicate the direction of data flow for the graph structure during the inference process. The blue arrows represent the historical information aggregation operation performed within the framework.

set at the beginning of step t respectively. First, we concatenate the latent feature and the hidden state of each node and edge, obtaining their new representations:

$$Z_V^{(t)} = L_V^{(t)} \parallel H_V^{(t-1)}, \quad Z_E^{(t)} = L_E^{(t)} \parallel H_E^{(t-1)}. \quad (2)$$

And then, feed these representations into the graph-structured network, obtaining new hidden states:

$$H_V^{(t)}, H_E^{(t)} = \mathcal{P}(Z_V^{(t)}, Z_E^{(t)}). \quad (3)$$

Decoder Module. Without loss of generality, we can assume that this module contains two networks to decode the hidden states of nodes and edges respectively:

$$X_V^{(t)}, X_E^{(t)} = \mathcal{D}(H_V^{(t)}, H_E^{(t)}). \quad (4)$$

Given the ground truth $\mathbf{Y}^{(t)}$ of each step, we can compute the total loss and thus, train the entire network.

3 CEF: A Context-Enhanced Framework

The section presents a framework that generalizes the standard encode-process-decode paradigm. Before stating the details, let us give some intuition first. Compared to the traditional seq-to-seq tasks, we observe that the ground truths in seq-graph reasoning have a much stronger inner connection with each other. For example, in neural graph algorithm learning, the goal is to imitate the execution of a classical graph algorithm like Prim’s algorithm [Prim, 1957]. In this scenario, each $\mathbf{Y}^{(t)}$ corresponds to an individual step of Prim’s algorithm, and the entire ground truth sequence represents the process of solving the minimum spanning tree prob-

lem by Prim’s algorithm, which is an optimization problem-solving procedure. Thus, similar to the consideration of momentum in historical gradients during gradient descent, in seq-graph reasoning, momentum concerning the outcome at each step should also be introduced and taken into account. This observation motivates our framework.

The framework builds upon the encode-process-decode paradigm and newly introduces a *Preprocessor* module between the encoder and processor. See Figure 1 for an illustration. This module is designed to capture dependencies and patterns on the latent feature sequence $\langle \mathbf{L}^{(1)}, \mathbf{L}^{(2)}, \dots, \mathbf{L}^{(t)}, \dots \rangle$, which are also referred to as “context” information. Hence, we term the new framework a *Context-Enhanced Framework* (CEF). More precisely, the module maintains a context state for each node and edge, storing the information aggregated from historical latent features. For each node v (resp. edge e), denote by $c_v^{(t)}$ (resp. $c_e^{(t)}$) the context state at the beginning of step t . Note that $c_v^{(t)}$ can be a direct concatenation of historical latent features or the outcome of applying some dimensionality reduction techniques (e.g. gating mechanisms) to them. For simplicity, $C_V^{(t-1)}$, $C_E^{(t-1)}$ and $\mathbf{C}^{(t-1)}$ are defined similarly.

The module involves two pivotal operations: context enhance function $\mathbf{f}_{\text{enhance}}$ and context update function $\mathbf{f}_{\text{update}}$. The application of $\mathbf{f}_{\text{enhance}}$ aims to augment latent features with contextual information, subsequently serving as input to the processor module:

$$\mathbf{S}^{(t)} = \mathbf{f}_{\text{enhance}}(\mathbf{L}^{(t)}, \mathbf{C}^{(t)}) \quad (5)$$

Concurrently, $\mathbf{f}_{\text{update}}$ is employed to update the context state

of each node and edge:

$$\mathbf{C}^{(t+1)} = \mathbf{f}_{\text{update}} \left(\mathbf{L}^{(t)}, \mathbf{C}^{(t)} \right) \quad (6)$$

The most critical aspect within the framework is how to set appropriate $\mathbf{f}_{\text{update}}$ and $\mathbf{f}_{\text{enhance}}$ so that they can integrate well with various types of processor modules. We notice that a popular approach to capturing patterns in historical latent features is through the use of the attention mechanism. Specifically, we can employ a node-wise (resp. edge-wise) concatenation function for $\mathbf{f}_{\text{update}}$ and a QKV attention network for $\mathbf{f}_{\text{enhance}}$. However, this method proves to be quite resource-intensive. The demand for memory significantly escalates with an increase in the total number of reasoning steps. An experimental study on this implementation is provided in the paper’s full version. In the main body, our focus shifts to leveraging *gating mechanisms* in the preprocessor module — a more space-efficient approach that demonstrates comparable empirical performance.

4 Applications

As outlined in Section 1.2, existing (processor) architectures fall within two categories: GNN-based and Transformer-based. This section considers both processor types and presents specific implementations of the framework for each.

4.1 Integration with GNN-based Processors

This subsection considers GNN-based processors. We follow the literature [Ibarz *et al.*, 2022] and formulate this type of processor into the model below. Note that the existing GNN-based processors only maintain hidden states for nodes.

In each step $t \in [T]$,

$$\begin{aligned} \forall v \in V : z_v^{(t)} = s_v^{(t)} \parallel h_v^{(t-1)} \\ r_v^{(t)} = \mathbf{f}_1 \left(z_v^{(t)} \right) \\ m_v = \bigoplus_{u:(u,v) \in E} \mathbf{f}_2 \left(r_v^{(t)}, r_u^{(t)}, s_{(u,v)}^{(t)} \right) \\ h_v^{(t)} = \mathbf{f}_3 \left(r_v^{(t)}, m_v \right) \end{aligned} \quad (7)$$

Various processors may utilize distinct \mathbf{f}_1 , \mathbf{f}_2 , \mathbf{f}_3 and aggregation function \bigoplus . Since only node hidden states are preserved and computed in such processors, we also focus on the transformation of node latent features in the preprocessor module and let $s_e^{(t)} = \ell_e^{(t)}$ directly for each edge $e \in E$. We introduce a novel gating mechanism to preprocess each node latent feature. Further, the weight sharing technique [Garland and Gregg, 2018] is employed to make the module more memory-efficient and reduce the risk of overfitting. We implement the preprocessor as follows:

$$\begin{aligned} \forall v \in V : \alpha_v = \text{ReLU} \left(\tanh \left(\mathbf{f}_{\text{linear}} \left(c_v^{(t)} \right) \right) \right) \\ s_v^{(t)} = c_v^{(t+1)} = \alpha_v \cdot c_v^{(t)} + (1 - \alpha_v) \cdot \ell_v^{(t)} \end{aligned} \quad (8)$$

$\mathbf{f}_{\text{linear}}$ is a linear layer that transforms $c_v^{(t)}$ into a one-dimensional scalar. Then, we apply function \tanh and ReLU sequentially, yielding $\alpha_v \in [0, 1]$. Interpreting α_v as a forget factor, $s_v^{(t)}$ and $c_v^{(t+1)}$ are obtained by a linear combination of the context state and latent feature. We provide an experimental study and discussion (in the full version of the paper) to explain why we chose the sigmoid function. In summary, empirical findings suggest that when integrating with GNN processors, the latent features in seq-graph reasoning exhibit distinctive patterns and prefer a higher probability of the forget factor being 0.

4.2 Integration with Transformer-based Processors

This subsection considers Transformer-based processors. They maintain hidden states for both nodes and edges. We propose a model below to capture the common structures of such processors. Use $\mathcal{N}(v) := \{u | (v, u) \in E\}$ to denote the neighbors of v . In each step $t \in [T]$,

$$\begin{aligned} \forall v \in V : \\ q_v = \mathbf{f}_{\text{query}} \left(z_v^{(t)}, \left\{ z_{(v,u)}^{(t)} \right\}_{u \in \mathcal{N}(v)} \right) \\ k_v = \mathbf{f}_{\text{key}} \left(\left\{ z_u^{(t)}, z_{(v,u)}^{(t)} \right\}_{u \in \mathcal{N}(v)} \right) \\ \vartheta_v = \mathbf{f}_{\text{value}} \left(\left\{ z_u^{(t)}, z_{(v,u)}^{(t)} \right\}_{u \in \mathcal{N}(v)} \right) \\ h_v^{(t)} = \mathbf{f}_{\text{node}} \left(z_v^{(t)}, q_v, k_v, \vartheta_v \right) \\ \forall (u, v) \in E : \\ h_{(u,v)}^{(t)} = \mathbf{f}_{\text{edge}} \left(h_u^{(t)}, h_v^{(t)}, z_{(u,v)}^{(t)}, z_{(v,u)}^{(t)} \right) \end{aligned} \quad (9)$$

The model extends the QKV attention mechanism, taking into account the incorporation and subsequent updating of edge features. Different processors employ different \mathbf{f} functions within the model. Certainly, we could adopt a context enhancement approach similar to the one in Section 4.1. However, we notice that, unlike existing GNN processors, the Transformer processors in the literature do not utilize gating mechanisms in the update of hidden states. Thus, to capture the contextual information of latent features and hidden states simultaneously, we propose a refined framework that seamlessly integrates the preprocessor and processor, allowing for mutual sharing of states. Initially, the preprocessor utilizes hidden states for context state updates.

$$\begin{aligned} \forall v \in V : \\ z_v^{(t)} = \ell_v^{(t)} \parallel h_v^{(t-1)} \\ \alpha_v = \text{sigmoid} \left(\mathbf{f}_{\text{linear-1}} \left(c_v^{(t)} \right) \right) \\ c_v^{(t+1)} = \alpha_v \cdot c_v^{(t)} + (1 - \alpha_v) \cdot z_v^{(t)} \\ \forall e \in E : \end{aligned} \quad (10)$$

$$\begin{aligned}
 z_e^{(t)} &= \ell_e^{(t)} \parallel h_e^{(t-1)} \\
 \alpha_e &= \text{sigmoid} \left(\mathbf{f}_{\text{linear-2}} \left(c_e^{(t)} \right) \right) \\
 c_e^{(t+1)} &= \alpha_e \cdot c_e^{(t)} + (1 - \alpha_e) \cdot z_e^{(t)}
 \end{aligned}$$

Subsequently, the processor employs the resulting context states to update hidden states.

$$\begin{aligned}
 \forall v \in V : \\
 q_v &= \mathbf{f}_{\text{query}} \left(z_v^{(t)}, \left\{ z_{(v,u)}^{(t)} \right\}_{u \in \mathcal{N}(v)} \right) \\
 k_v &= \mathbf{f}_{\text{key}} \left(\left\{ c_u^{(t+1)}, c_{(v,u)}^{(t+1)} \right\}_{u \in \mathcal{N}(v)} \right) \\
 \vartheta_v &= \mathbf{f}_{\text{value}} \left(\left\{ c_u^{(t+1)}, c_{(v,u)}^{(t+1)} \right\}_{u \in \mathcal{N}(v)} \right) \quad (11) \\
 h_v^{(t)} &= \mathbf{f}_{\text{node}} \left(z_v^{(t)}, q_v, k_v, \vartheta_v \right) \\
 \forall (u, v) \in E : \\
 h_{(u,v)}^{(t)} &= \mathbf{f}_{\text{edge}} \left(h_u^{(t)}, h_v^{(t)}, c_{(u,v)}^{(t+1)}, c_{(v,u)}^{(t+1)} \right)
 \end{aligned}$$

$\mathbf{f}_{\text{linear-1}}$ and $\mathbf{f}_{\text{linear-2}}$ are two linear layers that transform the context states of nodes and edges into one-dimensional scalars, respectively. Then we apply function sigmoid, yielding forget factors and updating context states. This process enables the context states to capture contextual information from both latent features and hidden states.

Next, we compute q_v, k_v, ϑ_v . While preserving the original computation for q_v , we strategically substitute $\mathbf{Z}^{(t)}$ with the context states $\mathbf{C}^{(t+1)}$ during the computation of k_v and ϑ_v . In essence, departing from conventional self-attention, we empower each node to aggregate attention information from the context states. Finally, we leverage q_v, k_v, ϑ_v to update the hidden state for each node and edge.

5 Experiments

This section validates the empirical performance of our framework and aims to investigate the following questions:

- Can the framework simultaneously enhance the reasoning capabilities of GNN-based and Transformer-based architectures? If so, is there a noticeable discrepancy in the performance improvements between these two architectural types?
- Consider an architecture faced with a reasoning task in which it inherently performs sub-optimally. If our framework is applied, can we expect an enhancement in performance? In simpler terms, does the framework push the original architecture towards a more “extreme” or improved state?

To this end, we apply the framework to two state-of-the-art architectures: Triplet-GMPNN [Ibarz *et al.*, 2022], a GNN-based model, and RT [Diao and Loynd, 2023], a Transformer-

based model. We refer to their implementations as CEF-GMPNN and CEF-RT, respectively. We remark that in the main body, we focus on the exploration of the aforementioned two questions; the ablation experiments regarding implementation details are deferred to the full version of the paper. The code is available at <https://github.com/Ghost-st/CEF>.

5.1 Setup

The experiments are conducted on a machine equipped with an i7-13700K CPU and an RTX 4090 GPU. The results are averaged over 5 runs with random seeds 5, 18, 25, 30, and 42.

Datasets. We use the CLRS Algorithmic Reasoning Benchmark [Velickovic *et al.*, 2022], a proven benchmark that offers a unified evaluation (micro-F1 score) for assessing the (seq-graph) reasoning capabilities of neural networks. It is derived from a foundational algorithms textbook “Introduction to Algorithms” [Cormen *et al.*, 2009] and consists of 30 algorithmic reasoning tasks, covering a wide range of algorithmic procedures including sorting, searching, dynamic programming, graph algorithms, string algorithms, and geometric algorithms. The inclusion of such a variety of tasks enables us to draw meaningful conclusions about the effectiveness of the proposed framework.

Baselines. The two most crucial baselines we compare against are Triplet-GMPNN [Ibarz *et al.*, 2022] and RT [Diao and Loynd, 2023], which play a pivotal role in addressing the aforementioned questions. Moreover, we include Memnet [Sukhbaatar *et al.*, 2015], MPNN [Gilmer *et al.*, 2017], PGN [Velickovic *et al.*, 2020], and NPQ [Jain *et al.*, 2023] in experiments as important references for comparison.

Training Details. In our implementations, $\mathbf{f}_{\text{linear}}$, $\mathbf{f}_{\text{linear-1}}$ and $\mathbf{f}_{\text{linear-2}}$ are individual linear layers. To ensure fair comparisons, we employ the experimental hyperparameter settings used in triplet-GMPNN [Ibarz *et al.*, 2022] and Relational Transformer [Diao and Loynd, 2023] for CEF-GMPNN and CEF-RT, respectively. Specifically, for CEF-GMPNN, we set the batch size to 32 and the network is trained for 10,000 steps by Adam optimizer with a learning rate of 0.001; while for CEF-RT, we set the batch size to 4 and the network is trained for 10,000 steps by Adam optimizer with a learning rate of 0.00025. We remark that with the introduction of a preprocessor module in our framework, the average training time increases by 3 to 5 minutes for both methods. This increment represents less than 10% of the total runtime, indicating the efficiency of our implementations. More details are deferred to the full version.

5.2 Comparison with GNN-based Processors

In this experiment, we compare our method CEF-GMPNN with Triplet-GMPNN across 30 reasoning tasks within the CLRS dataset. The results are shown in Figure 2. Bar charts illustrating average scores, with standard deviations denoted by black lines, are presented for each task. Additionally, the tasks are arranged in descending order of improvement magnitude to better showcase trends.

From Figure 2, we see that our framework demonstrates enhancement across all 30 reasoning tasks, implying that our framework pushes the Triplet-GMPNN architecture towards




Figure 2: Comparisons between our CEF-GMPNN and Triplet-GMPNN. The 30 tasks are arranged in descending order of improvement magnitude.




Figure 3: Comparisons between our CEF-RT and RT. The 30 tasks are arranged in descending order of improvement magnitude. Note that on the Knuth-Morris-Pratt reasoning task, both algorithms exhibit very poor performance. Refer to Table 1 for detailed scores.

an improved state. Notably, in tasks related to sorting and strings, such as bubble sort, quick select, Knuth-Morris-Pratt, and heap sort, the framework significantly improves the performance of Triplet-GMPNN.

5.3 Comparison with Transformer-based Processors

In this experiment, we compare our method CEF-RT with RT across 30 reasoning tasks within the CLRS dataset. The results are shown in Figure 3.

We observe a trend different from the previous experiment. The framework enhances the performance of RT across the majority of CLRS datasets. Specifically, in tasks related to graph optimization, such as strongly connected components (SCC) and depth-first search (DFS), there is a noteworthy increase in test scores of over 30%. However, for sorting tasks where the original RT approach exhibits poor average scores with very large derivations, our framework leads to a decrease in performance. This phenomenon indicates that for Transformer-based Processors, our framework steers them to-

ward a more “extreme” state.

5.4 Summarized Results and Analysis

We summarize the experiential results in Table 1. Due to space limitations, we use column “Prior Best” to represent the best results among Memnet [Velickovic *et al.*, 2022], PGN [Velickovic *et al.*, 2022], MPNN [Velickovic *et al.*, 2022], and NPQ [Jain *et al.*, 2023]. Note that [Jain *et al.*, 2023] proposed four variants of the NPQ approach, and for each task, we have chosen the best scores among these variants as reference items.

In the table, we present the specific test scores for each task, along with summarizing the average scores. Additionally, in line with prior work [Ibarz *et al.*, 2022], we provide the respective proportions of tasks achieving scores greater than 90%, 80%, and 60%. We see the following trends:

- Our framework yields state-of-the-art results across the majority of the CLRS benchmark, demonstrating higher average scores and lower deviations compared to prior work. The overall average score of CEF-GMPNN is

Tasks	Prior Best	Triplet-GMPNN	RT	CEF-RT (ours)	CEF-GMPNN (ours)
Activity Selector	83.36%±4.27	95.18%±0.45	87.72%±2.7	71.22%±3.51	95.58%±0.49
Articulation Points	50.91%±2.18	88.32%±2.01	34.15%±14.6	50.71%±1.42	96.51%±2.09
Bellman-Ford	92.99%±0.34	97.39%±0.19	94.24%±1.5	93.15%±0.68	97.80%±0.20
BFS	99.89%±0.05	99.73%±0.04	99.14%±0.7	99.99%±0.00	99.95%±0.09
Binary Search	76.95%±0.13	77.58%±2.35	81.48%±6.7	83.66%±4.18	83.91%±0.70
Bridges	72.69%±4.78	93.99%±2.07	37.88%±11.8	72.21%±2.24	99.36%±0.38
Bubble Sort	73.58%±0.78	67.68%±5.50	38.22%±13.0	4.35%±1.16	92.97%±1.73
DAG Shortest Paths	96.94%±0.16	98.19%±0.30	96.61%±1.6	97.35%±0.22	98.55%±0.19
DFS	13.36%±1.61	47.79%±4.19	39.23%±10.5	75.11%±10.23	64.17%±7.33
Dijkstra	91.50%±0.50	96.05%±0.60	91.2%±5.8	93.69%±2.14	97.89%±0.22
Find Max.Subarray	65.23%±2.56	76.36%±0.43	66.52%±3.7	96.63%±1.77	76.38%±3.04
Floyd-Warshall	28.76%±0.51	48.52%±1.04	31.59%±7.6	38.61%±2.42	55.42%±2.58
Graham Scan	91.04%±0.31	93.62%±0.91	74.15%±7.4	80.74%±2.97	96.27%±0.37
Heapsort	68.00%±1.57	31.04%±5.82	32.96%±14.8	5.57%±1.17	52.13%±4.91
Insertion Sort	71.42%±0.86	78.14%±4.64	89.43%±9.0	99.99%±0.00	92.41%±2.84
Jarvis' March	92.88%±2.87	91.01%±1.30	94.57%±2.2	88.93%±3.88	91.79%±2.83
Knuth-Morris-Pratt	3.91%±0.15	19.51%±4.57	0.03%±0.1	0.01%±0.02	42.99%±11.45
LCS Length	72.05%±5.72	80.51%±1.84	83.32%±4.1	86.21%±0.83	86.29%±1.35
Matrix Chain Order	83.91%±0.49	91.68%±0.59	91.89%±1.2	98.72%±0.69	92.96%±0.98
Minimum	87.71%±0.52	97.78%±0.55	95.28%±2.0	95.09%±8.50	98.50%±0.10
MST-Kruskal	70.97%±1.50	89.80%±0.77	64.91%±11.8	73.41%±0.98	91.04%±0.22
MST-Prim	69.08%±7.56	86.39%±1.33	85.77%±7.9	89.62%±2.80	93.08%±2.66
Naive String Matcher	4.24%±0.98	78.67%±4.99	65.01%±32.3	89.14%±3.71	98.72%±1.49
Optimal BST	72.03%±1.21	73.77%±1.48	74.4%±2.6	56.80%±0.89	79.43%±2.99
Quickselect	3.66%±0.42	0.47%±0.25	19.18%±17.3	1.42%±1.42	2.81%±0.54
Quicksort	73.10%±0.67	64.64%±5.12	39.42%±13.2	15.20%±0.60	88.56%±3.24
Segments Intersect	93.53%±0.88	97.64%±0.09	84.94%±2.6	86.96%±1.52	98.09%±0.13
SCC	32.19%±9.23	43.43%±3.15	28.59%±15.2	84.08%±5.70	51.93%±4.37
Task Scheduling	84.89%±0.91	87.25%±0.35	82.93%±1.8	85.69%±0.50	89.63%±0.39
Topological Sort	60.45%±2.69	87.27%±2.67	80.62%±17.5	98.46%±1.32	98.19%±1.03
Overall Average	66.04%	75.98%	66.18%	70.42%	82.68%
>90%	7/30	11/30	7/30	9/30	18/30
>80%	11/30	17/30	15/30	18/30	22/30
>60%	23/30	24/30	20/30	22/30	25/30

Table 1: The test score of CEF-GMPNN and CEF-RT on 30 reasoning tasks within CLRS. The best-performing results in each row are highlighted in bold. The column “Prior Best” in the table represents the best results among four approaches in the literature: Memnet [Velickovic *et al.*, 2022], PGN [Velickovic *et al.*, 2022], MPNN [Velickovic *et al.*, 2022], and NPQ [Jain *et al.*, 2023].

82.68%, which is a substantial improvement over previous results.

- The results show a broader range of tasks where we attain scores exceeding 60%, 80%, and 90%. Specifically, CEF-GMPNN yields a score of at least 90% in 18 tasks, 80% for 22 tasks, and 60% for 20 tasks.
- For both Triplet-GMPNN and RT, our framework enhances their performance in most tasks. However, the specific details of the enhancements differ. For Triplet-GMPNN, the improvements brought by the framework are comprehensive, showing enhancement across all tasks. In contrast, for RT, the framework exhibits more significant improvements in certain tasks compared to Triplet-GMPNN, but it also shows performance degradation in tasks where RT initially performs poorly. These results indicate that our framework pushes Triplet-GMPNN towards an improved state while making RT more “extreme” in its performance.

6 Conclusion

The paper studies seq-graph reasoning. We introduce a context-enhanced framework for such tasks and show that it can effectively integrate with various existing approaches. Our experiments reveal that the proposed framework consistently enhances the reasoning capabilities of both GNN-based and Transformer-based architectures, achieving new state-of-the-art results on the CLRS benchmark.

There are several intriguing avenues for future research. One interesting direction involves refining our current framework. As discussed in Section 1, the outcome sequence may form an entire problem-solving or optimization process, potentially represented by an additional reasoning graph. Therefore, it is worthwhile to explore a framework capable of efficiently capturing such graph structures hidden in the data. Another promising direction could be concentrating on the CLRS benchmark, aiming to design novel architectures that could yield further performance improvements.

Acknowledgements

This work is supported by the National Key Research Project of China under Grant No. 2023YFA1009402, the National Natural Science Foundation of China (No. 62302166), the Scientific and Technological Innovation 2030 Major Projects under Grant 2018AAA0100902, the Shanghai Science and Technology Commission under Grant No.20511100200, the Dean’s Fund of Shanghai Key Laboratory of Trustworthy Computing, ECNU, and the Key Laboratory of Interdisciplinary Research of Computation and Economics (SUFEC), Ministry of Education.

References

- [Cormen *et al.*, 2009] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, 3rd Edition*. MIT Press, 2009.
- [de Rezende *et al.*, 2021] Susanna F. de Rezende, Massimo Lauria, Jakob Nordström, and Dmitry Sokolov. The power of negative reasoning. In *CCC*, volume 200 of *LIPICs*, pages 40:1–40:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [Diao and Loynd, 2023] Cameron Diao and Ricky Loynd. Relational attention: Generalizing transformers for graph-structured tasks. In *ICLR*. OpenReview.net, 2023.
- [Drori *et al.*, 2020] Iddo Drori, Anant Kharkar, William R. Sickinger, Brandon Kates, Qiang Ma, Suwen Ge, Eden Dolev, Brenda Dietrich, David P. Williamson, and Madeleine Udell. Learning to solve combinatorial optimization problems on real-world graphs in linear time. In *ICMLA*, pages 19–24. IEEE, 2020.
- [Garland and Gregg, 2018] James Garland and David Gregg. Low complexity multiply-accumulate units for convolutional neural networks with weight-sharing. *ACM Trans. Archit. Code Optim.*, 15(3):31:1–31:24, 2018.
- [Gerasimova *et al.*, 2023] Olga Gerasimova, Nikita Severin, and Ilya Makarov. Comparative analysis of logic reasoning and graph neural networks for ontology-mediated query answering with a covering axiom. *IEEE Access*, 11:88074–88086, 2023.
- [Gilmer *et al.*, 2017] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 1263–1272. PMLR, 2017.
- [He *et al.*, 2023] Weinan He, Canming Huang, Zhanhao Xiao, and Yongmei Liu. Exploring the capacity of pre-trained language models for reasoning about actions and change. In *ACL (1)*, pages 4629–4643. Association for Computational Linguistics, 2023.
- [Huang *et al.*, 2022] Zeren Huang, Kerong Wang, Furui Liu, Hui-Ling Zhen, Weinan Zhang, Mingxuan Yuan, Jianye Hao, Yong Yu, and Jun Wang. Learning to select cuts for efficient mixed-integer programming. *Pattern Recognit.*, 123:108353, 2022.
- [Hussain *et al.*, 2022] Md. Shamim Hussain, Mohammed J. Zaki, and Dharmashankar Subramanian. Global self-attention as a replacement for graph convolution. In *KDD*, pages 655–665. ACM, 2022.
- [Ibarz *et al.*, 2022] Borja Ibarz, Vitaly Kurin, George Papanikarinos, Kyriacos Nikiforou, Mehdi Bennani, Róbert Csordás, Andrew Joseph Dudzik, Matko Bosnjak, Alex Vitvitskyi, Yulia Rubanova, Andreea Deac, Beatrice Bevilacqua, Yaroslav Ganin, Charles Blundell, and Petar Velickovic. A generalist neural algorithmic learner. In *LoG*, volume 198 of *Proceedings of Machine Learning Research*, page 2. PMLR, 2022.
- [Jain *et al.*, 2023] Rishabh Jain, Petar Velickovic, and Pietro Liò. Neural priority queues for graph neural networks. In *The 2023 ICML Workshop on Knowledge and Logical Reasoning in the Era of Data-driven Learning*, volume 202. PMLR, 2023.
- [Jiang *et al.*, 2022] Jinhao Jiang, Kun Zhou, Ji-Rong Wen, and Wayne Xin Zhao. \$great truths are always simple: \$ A rather simple knowledge encoder for enhancing the commonsense reasoning capacity of pre-trained models. In *NAACL-HLT (Findings)*, pages 1730–1741. Association for Computational Linguistics, 2022.
- [Jürß *et al.*, 2023] Jonas Jürß, Dulhan Hansaja Jayalath, and Petar Veličković. Recursive algorithmic reasoning. In *The Second Learning on Graphs Conference*, 2023.
- [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR (Poster)*. OpenReview.net, 2017.
- [Lample and Charton, 2020] Guillaume Lample and François Charton. Deep learning for symbolic mathematics. In *ICLR*. OpenReview.net, 2020.
- [Li *et al.*, 2023] Hongyang Li, Cui Zhu, Wenjun Zhu, Yunfeng Zhou, and Xinyu Zhang. Short-term sequential pattern for temporal knowledge graph reasoning. In *CAIBDA*, volume 373 of *Frontiers in Artificial Intelligence and Applications*, pages 728–739. IOS Press, 2023.
- [Ling *et al.*, 2022] Chen Ling, Tanmoy Chowdhury, Junji Jiang, Junxiang Wang, Xuchao Zhang, Haifeng Chen, and Liang Zhao. Deepgar: Deep graph learning for analogical reasoning. In *ICDM*, pages 1065–1070. IEEE, 2022.
- [Ma *et al.*, 2023] Ruixin Ma, Biao Mei, Yunlong Ma, Hongyan Zhang, Meihong Liu, and Liang Zhao. One-shot relational learning for extrapolation reasoning on temporal knowledge graphs. *Data Min. Knowl. Discov.*, 37(4):1591–1608, 2023.
- [Numeroso *et al.*, 2023] Danilo Numeroso, Davide Bacciu, and Petar Velickovic. Dual algorithmic reasoning. In *ICLR*. OpenReview.net, 2023.
- [Prim, 1957] R. C. Prim. Shortest connection networks and some generalizations. *The Bell System Technical Journal*, 36(6):1389–1401, 1957.
- [Ramesh *et al.*, 2023] Gowtham Ramesh, Makesh Narasimhan Sreedhar, and Junjie Hu. Single sequence prediction over reasoning graphs for multi-hop QA. In

- ACL (1)*, pages 11466–11481. Association for Computational Linguistics, 2023.
- [Schreck and Mathis, 2016] Pascal Schreck and Pascal Mathis. Automatic constructibility checking of a corpus of geometric construction problems. *Math. Comput. Sci.*, 10(1):41–56, 2016.
- [Sukhbaatar *et al.*, 2015] Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. End-to-end memory networks. In *NIPS*, pages 2440–2448, 2015.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, pages 5998–6008, 2017.
- [Velickovic *et al.*, 2018] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *ICLR (Poster)*. OpenReview.net, 2018.
- [Velickovic *et al.*, 2020] Petar Velickovic, Lars Buesing, Matthew C. Overlan, Razvan Pascanu, Oriol Vinyals, and Charles Blundell. Pointer graph networks. In *NeurIPS*, 2020.
- [Velickovic *et al.*, 2022] Petar Velickovic, Adrià Puigdomènech Badia, David Budden, Razvan Pascanu, Andrea Banino, Misha Dashevskiy, Raia Hadsell, and Charles Blundell. The CLRS algorithmic reasoning benchmark. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, pages 22084–22102. PMLR, 2022.
- [Wang *et al.*, 2023] Yuan Wang, Kun Yu, Chen Chen, Xiyuan Hu, and Silong Peng. Dynamic graph learning with content-guided spatial-frequency relation reasoning for deepfake detection. In *CVPR*, pages 7278–7287. IEEE, 2023.
- [Wu *et al.*, 2019] Jun Wu, Jingrui He, and Jiejun Xu. Demonet: Degree-specific graph neural networks for node and graph classification. In *KDD*, pages 406–415. ACM, 2019.
- [Wu *et al.*, 2021] Zhanghao Wu, Paras Jain, Matthew A. Wright, Azalia Mirhoseini, Joseph E. Gonzalez, and Ion Stoica. Representing long-range context for graph neural networks with global attention. In *NeurIPS*, pages 13266–13279, 2021.
- [Yu, 2017] Lei Yu. *Tackling sequence to sequence mapping problems with neural networks*. PhD thesis, University of Oxford, UK, 2017.
- [Zhao *et al.*, 2024] Zhenbing Zhao, Bin Liu, Yongjie Zhai, Wenqing Zhao, and Pan Su. Dual graph reasoning network for oil leakage segmentation in substation equipment. *IEEE Trans. Instrum. Meas.*, 73:1–15, 2024.
- [Zhu, 2022] Zihao Zhu. From shallow to deep: Compositional reasoning over graphs for visual question answering. In *ICASSP*, pages 8217–8221. IEEE, 2022.