

GALLa: Graph Aligned Large Language Models for Improved Source Code Understanding

Zi Yin Zhang^{12*} Hang Yu^{2*} Shijie Li² Peng Di² Jianguo Li^{2†} Rui Wang^{1†}

¹Shanghai Jiao Tong University

²Ant Group

{daenerystargaryen, wangrui12}@sjtu.edu.cn, {hyu.hugo, lijg.zero}@antgroup.com

Abstract

Programming languages possess rich semantic information - such as data flow - that is represented by graphs and not available from the surface form of source code. Recent code language models have scaled to billions of parameters, but model source code solely as text tokens while ignoring any other structural information. Conversely, models that do encode structural information of code make modifications to the Transformer architecture, limiting their scale and compatibility with pretrained LLMs. In this work, we take the best of both worlds with GALLa - Graph Aligned Large Language Models. GALLa utilizes graph neural networks and cross-modal alignment technologies to inject the structural information of code into LLMs as an auxiliary task during finetuning. This framework is both model-agnostic and task-agnostic, as it can be applied to any code LLM for any code downstream task, and requires the structural graph data only at training time from a corpus unrelated to the finetuning data, while incurring no cost at inference time over the baseline LLM. Experiments on five code tasks with seven different baseline LLMs ranging in size from 350M to 14B validate the effectiveness of GALLa, demonstrating consistent improvement over the baseline, even for powerful models such as LLaMA3 and Qwen2.5-Coder.

1 Introduction

In recent years, applying large language models (LLMs) to processing and generating source code has been a research topic of special interest in both natural language processing and software engineering community (Chen et al., 2021). However, unlike natural languages, programming languages have rich semantics besides the lexical representation of source code, such as the path of execution, the flow of data, and the dependency of

function calling. These semantics are represented by graph structures such as Abstract Syntax Tree (AST), Control Flow Graph (CFG), and Data Flow Graph (DFG). While enhancing LLMs with lexical representations of source code has the potential to boost their performance, the integration of these richer semantic constructs has yet to be fully realized (Zhang et al., 2024).

The major challenge of injecting code structures into code language models lies in the **incompatibility between structural graphs and large-scale pretrained language models**. On one end, following the scaling laws (Kaplan et al., 2020; Hoffmann et al., 2022), some works have tried to improve the capability of language models in code processing by increasing their model size and pretraining data size, leading to code LLMs with tens or even hundreds of billions of parameters such as Code LLaMA (Rozière et al., 2024) and DeepSeek-Coder (Guo et al., 2024; DeepSeek-AI et al., 2024). However, these models are standard decoder-only Transformer language models (Vaswani et al., 2017) trained with next token prediction, and are thus **unable to capture the semantics embedded in the structural graphs of code**.

On the other extreme, another line of research, represented by GraphCodeBERT (Guo et al., 2021) and TreeBERT (Jiang et al., 2021), has focused on injecting graph structures into Transformer-based code language models. However, these methods either linearize the graphs into text tokens and are thus only applicable to simple tree-based graph structures (Niu et al., 2022; Guo et al., 2022), or modify the Transformer architecture such as attention masks (Guo et al., 2021) and positional encodings (Peng et al., 2021) to encode graph information. **These modifications to the model structure make them incompatible with the large-scale pretrained decoder-only LLMs, and thus these experiments have been limited to a small scale.**

*Work done during Zi Yin's internship at Ant Group. The first two authors contributed equally.

†Corresponding authors.

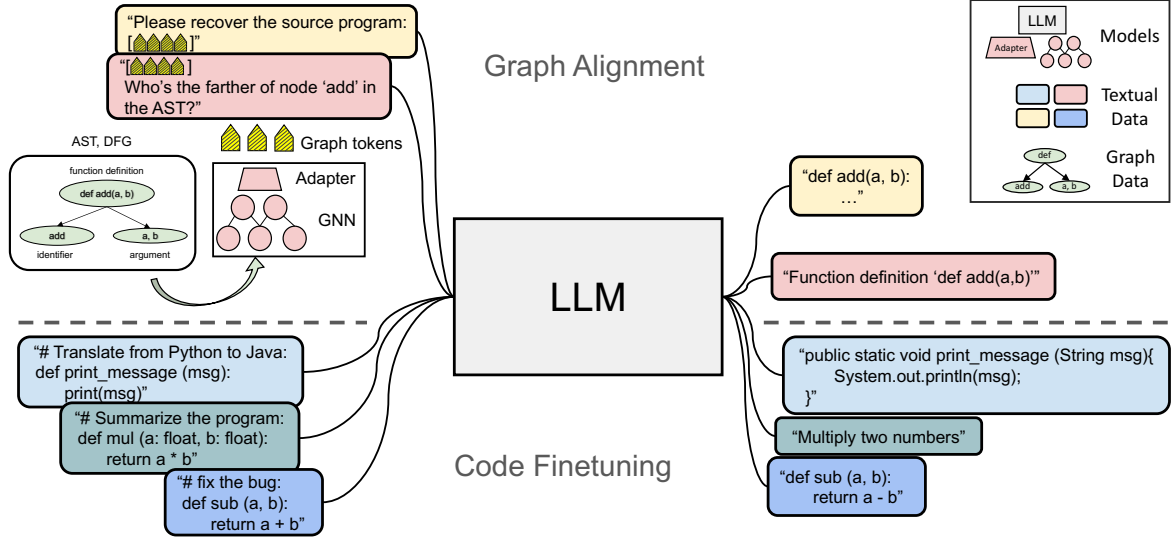


Figure 1: High-level overview of our method. The LLM is simultaneously trained on code downstream task data (such as code translation, code summarization, and code repair) and graph alignment data (graph-to-code generation, and question-answer pairs about the graph structures). The same-colored data on the left and right belong to the same input-output pair. The graph data is only required for the graph alignment tasks, but not the code finetuning tasks.

Interestingly, insights from the computer vision community suggest a promising avenue for bridging the gap between different modalities: by utilizing a light-weight adapter to project the output features of a non-textual input processing model (such as an image encoder or a graph encoder) into the embedding space of language models, they are able to ground LLMs’ understanding of text beyond the text modality, while in the meantime preserving LLMs’ capability acquired during text-only pretraining (Liu et al., 2023; Bai et al., 2023; Zhu et al., 2024).

Inspired by these works, we introduce GALLa: **Graph Aligned Large Language Models** for code. We utilize a graph neural network (GNN) to process ASTs and DFGs extracted from source code, which are then projected into the language model’s embedding space by a small adapter. The language model is trained to generate the source code conditioned on the graph information and to answer questions about the graph structure. These objectives align the language model’s representation of code to the graph structures and impart it with a deeper understanding of the source code. As demonstrated in Figure 1, our method is based on the transfer learning framework (Raffel et al., 2020) and separates graph alignment data from task-specific training data, thus preserving the general capability of the LLM acquired during pretraining and requiring no graph information about downstream task training

or test data.

Through extensive experiments on five code understanding and generation tasks, we validate the effectiveness of GALLa on seven distinct base LLMs ranging in size from 350M to 14B. GALLa brings consistent improvement over all baseline models, and even demonstrates abilities to generalize structural knowledge acquired during graph alignment to programming languages that are absent in the alignment data. All the data used in our experiments (both graph alignment and downstream tasks) are sourced from publicly available datasets¹, and the complete code for reproducing our results is released in <https://github.com/codefuse-ai/GALLa>.

2 Related Work

Existing works that utilize structural graphs to enhance code language models can be categorized into three types: 1) modifying attention masks to encode graph information, 2) integrating graphs into the textual input, and 3) enhancing positional encodings with graph information.

The first category is represented by GraphCodeBERT (Guo et al., 2021), which finetunes CodeBERT (Feng et al., 2020) on concatenated source code and DFGs. Attention masks in the model are

¹<https://huggingface.co/datasets/codefuse-ai/GALLa>

modified to reflect the graph structures: a node token v_a is allowed to attend another node token v_b , only if there is an edge from v_b to v_a in the graph, while a node v and a source code token c can attend to each other only if they correspond to each other. StructCoder (Tipirneni et al., 2024) also modifies attention masks similarly to encode the relations between source code tokens and AST, DFG tokens.

For the second category, TreeBERT (Jiang et al., 2021) is an encoder-decoder model where each input token to the encoder is a node’s constituent path in the AST, while the decoder is trained to generate the source code. Many other works in this category, including SynCoBERT (Wang et al., 2021), SPT-Code (Niu et al., 2022), and UniXcoder (Guo et al., 2022), map ASTs into text sequences (for example, by depth-first traversing) as model input. However, this method only applies to simple graph structures such as AST, but not more complex ones such as DFG, which may include loops.

In the third category, Peng et al. (2021) proposed TPTrans, which applies a recurrent sequence encoder to encode the relative paths between terminal nodes in ASTs and uses the results as relative positional encodings in self-attention modules, while the absolute paths from terminal nodes to the root node are similarly used as absolute positional encodings. Peng et al. (2022) used a list of 2-d coordinates and an embedding lookup table to describe the positions of AST nodes instead, and also used these embeddings to enhance both relative and absolute positional encodings in Transformer.

Despite their contributions, **all these approaches focus primarily on the encoder of the Transformer and are not fully compatible with decoder-only LLMs.** For instance, in the first group, the presence of cycles in the graphs means that the corresponding attention mask cannot retain a lower-triangular format, as required by LLMs. Adapting LLMs to these graph-structure-based attention masks would necessitate extensive retraining due to the mask’s inconsistency with the original causal language modeling objectives. In contrast, the proposed GALLa method processes graph information externally, allowing the LLM architecture to remain unmodified.

Apart from these graph-enhanced models, there is another model TransCoder-IR (Szafraniec et al., 2023), which utilizes LLVM intermediate representation (IR) to ground the model’s understanding of code by generating source code from IR and vice versa. Like our method, TransCoder-IR only

\mathcal{G}	a graph
v	a node in the graph
e	an edge in the graph
n_v	number of nodes in the graph
n_e	number of edges in the graph
d_{node}	node feature dimension of the graph data
d_{gnn}	GNN hidden and output dimension
d_{lm}	LLM hidden and embedding dimension
n_g	number of graph tokens in LLM’s input
n_t	number of text tokens in LLM’s input
V	node features (input to GNN, extracted by a text encoder)
E	edge indices (input to GNN)
H	contextual node features (output of GNN)
Q	query vectors (input to adapter)
X	embedding vectors (input to LLM)
Y	logits (output of LLM)

Table 1: Notations used in Section 3

uses IR for alignment at training time but does not need it at test time. However, intermediate representations are also text tokens and provide limited structural information.

3 Method

The objective of GALLa is to align language models to the implicit structures of code (represented by AST and DFG - see more details in Appendix A) when finetuning on code-related tasks, and an overview of the method is provided in Figure 1. In this section, we elaborate on two key challenges: **1) how to input graph information into LLMs**, and **2) how to design training tasks so that LLMs can learn about such graph information.** For the first challenge, we design a model that consists of three modules: GNN encoder, adapter, and LLM decoder (Section 3.1). For the second challenge, we propose a two-stage training scheme: graph encoder pretraining and graph-LLM alignment (Section 3.2 and Figure 2). The notations used in this section are presented in Table 1.

3.1 Model Architecture

3.1.1 GNN Encoder

To fully capture the rich information contained in the structural graphs such as loops, node degrees, and edge directions, we first process the graphs with a graph neural network (GNN) to extract node information. For a graph \mathcal{G} with n_v nodes and n_e edges, a text encoder is used to extract node features from the corresponding source code of each

node. These features are used to construct the node feature matrix $V \in \mathbb{R}^{n_v \times d_{\text{node}}}$, which are fed into the GNN along with the edge matrix $E \in \mathbb{Z}^{n_e \times 2}$, where each element is a node index. The output of the GNN is the contextual node representations $H \in \mathbb{R}^{n_v \times d_{\text{gnn}}}$.

In this module, the text encoder can be any code embedding model such as CodeT5+ (Wang et al., 2023). The GNN can be any convolution-based or self-attention-based directed GNN such as MagNet (Zhang et al., 2021) or DUPLEX (Ke et al., 2024), as DFG is a type of directed graph.

3.1.2 Adapter

The outputs of GNN are projected by an adapter into the LLM’s embedding space. Following Qwen-VL (Bai et al., 2023), we use a single layer of cross-attention as the adapter, where the GNN’s outputs H serve as keys and values, and the queries $Q \in \mathbb{R}^{n_g \times d_{\text{lm}}}$ are n_g learnable vectors:

$$X_g = \text{CrossAttn}(q = Q, k = H, v = H). \quad (1)$$

Alternative to the cross-attention layer, the adapter may also be a multi-layer perception (MLP), as used by LLaVA (Liu et al., 2023), which applies projection independently to each node. The main difference between cross-attention and MLP is that cross-attention allows for information exchange between the nodes, while MLP is applied independently to each node. Thus, we choose cross-attention as the adapter in the main results, and experiment with MLP in the ablation studies in Section 4.4.

3.1.3 LLM Decoder

The adapter’s outputs $X_g \in \mathbb{R}^{n_g \times d_{\text{lm}}}$ are n_g embedding vectors, which we dub “graph tokens”. Any other text in the LLM’s input (as shown at the top of Figure 1) is first tokenized and passed into the LLM’s embedding layer to obtain n_t text embeddings $X_t \in \mathbb{R}^{n_t \times d_{\text{lm}}}$, and then concatenated with the graph tokens to form the input to the LLM’s Transformer layers:

$$X = [X_g, X_t]^2 \in \mathbb{R}^{(n_g+n_t) \times d_{\text{lm}}}. \quad (2)$$

The LLM’s output logits $Y \in \mathbb{R}^{(n_g+n_t) \times d_{\text{lm}}}$ are then used to compute cross-entropy loss with next token prediction (i.e. causal language modeling). However, the loss is masked on the graph tokens and only computed on the text tokens.

²Depending on the actual input text, the graph tokens can also be placed after the text tokens or even inserted in the middle of text tokens, as illustrated in Figure 1.

3.1.4 Model Choice

Lastly, we emphasize that GALLa is a framework for bridging the text modality and graph modality, and each of the three modules in GALLa can be instantiated with different models. The choice of GNN - e.g. directed or undirected - depends on the properties of graph data, while the choice of LLM depends on application scenarios - e.g. monolingual or multilingual, general-purposed or domain-specialized.

3.2 Training Procedure

Motivated by the pretraining + instruction fine-tuning paradigm in LLMs, we divide the training of GALLa into two stages. The first stage is self-supervised (continue) pretraining, where only AST/DFG data extracted from source code are used. The second stage is supervised finetuning, where GraphQA data are collected by designing graph-related questions and manually crafting question-answer templates.

3.2.1 Stage 1: Graph Encoder Pretraining

In GALLa, the LLM is initialized from a pretrained LLM checkpoint such as LLaMA (Dubey et al., 2024) or StarCoder (Li et al., 2023), while both the GNN and the adapter are randomly initialized. Thus, to prevent the newly initialized GNN and adapter from disrupting the LLM’s pretrained representations, we fix the LLM’s weights in stage 1, and update only the GNN and the adapter. In this stage, the model is trained with graph-to-code generation (Graph2Code), where the model reconstructs a graph’s corresponding source code tokens X_t based on the graph tokens X_g by maximizing the probability $P(X_t|X_g)$. Similar to visual instruction tuning (Liu et al., 2023), this stage can be understood as training a “graph tokenizer” for the frozen LLM.

3.2.2 Stage 2: Graph-LLM Alignment

In the second stage, we aim to align the LLM’s pre-trained representations of source code to the structural graphs and deepen their understanding of code structures. In this stage, the LLM is unfrozen, and all three modules are updated together. The graph alignment tasks in this stage include Graph2Code (same as stage 1), and graph question answering (GraphQA). In GraphQA, the language model answers questions about a graph’s structures, such as predicting whether there is an edge between two given nodes, or predicting the children of a given

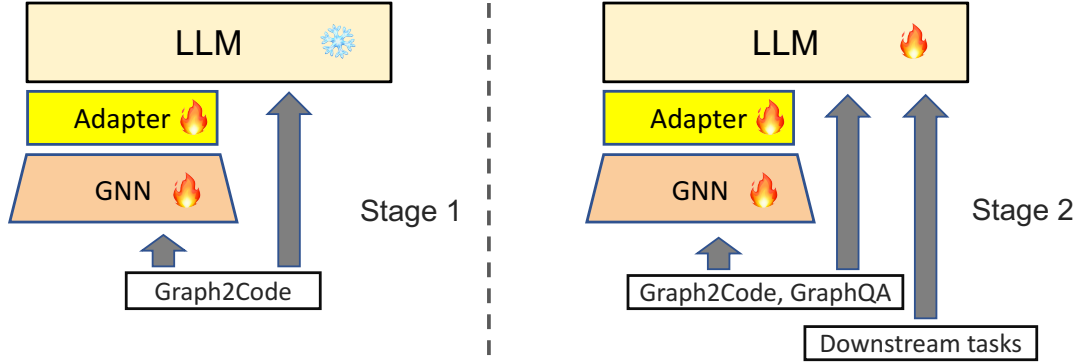


Figure 2: Illustration of the two-stage training schemes. In the first stage (left), the LLM’s weights are frozen, and the GNN and adapter are pretrained on the Graph2Code task. In the second stage (right), the LLM is unfrozen and trained together with the GNN and adapter on the graph alignment tasks, while also simultaneously trained along (without the GNN and adapter) on downstream finetuning data. For Graph2Code and GraphQA tasks, the graph in a data sample goes into GNN, while the texts (including instruction and answer) in the sample go directly into the LLM.

nodes. Formally, the model is trained to maximize the probabilities of the answer text tokens X_a conditioned on the graph tokens X_g and the question text tokens X_q : $P(X_a|X_g, X_q)$.

Since the ultimate goal of aligning LLMs to code graph structures is to enhance their performance on related downstream tasks, the model is simultaneously trained on the graph alignment data and downstream task data in this stage. However, we emphasize that the GNN and adapter are only used for the graph alignment tasks. No graph information is required of the downstream task data, as it goes directly into the LLM in the form of question-answer pairs, as shown at the bottom of Figure 1.

3.2.3 Inference

After aligning the source code representation of LLMs to code graph structures in training stage 2, at inference time the graph encoder and adapter are discarded, so that the LLM can respond to user queries using its internal knowledge without any loss of speed compared with the base LLM. We choose not to encode the code in user queries with GNN at inference time because 1) it would require extracting AST and DFG from the code first, resulting in an unreasonable inference latency; 2) extracting AST and DFG requires that the input code is syntactically correct and complete, which is not true for many downstream tasks such as defect detection and code repair.

4 Experiments

In this section, we first discuss the details of our experiments (4.1, 4.2), and then provide the results in 4.3 and 4.4.

4.1 Datasets

4.1.1 Graph Alignment

For graph alignment, we used 240K Python programs and 75K Java programs from CodeNet (Puri et al., 2021). For each program, we extracted one AST and one DFG using Program Structure Interface³, which resulted in 626K code-graph pairs after removing empty graphs and source code files that are longer than 4096 tokens. More details about these AST and DFG can be found in Appendix A.

For the Graph2Code task, we removed all Python programs where no main function can be found, resulting in 150K Java graph-code pairs and 81K Python graph-code pairs.

For the GraphQA task, we designed three types of questions:

- 1) Edge prediction: given the graph tokens and the source code of two nodes, the model is tasked to predict where there is an edge between them. This task is constructed only on DFG data.
- 2) Parent prediction: given the graph tokens and the source code of one node, the model is tasked to predict the node’s parent in the graph. This task is constructed on both AST and DFG data.

³<https://plugins.jetbrains.com/docs/intellij/psi.html>

	Translation	Clone Detection	Defect Detection	Summarization	Repair
Language	Py→Java/Java→Py	Java	C	Py/Java/JS	Java
Train Samples	44K/30K	300K	25K	265K/170K/62K	111K
Test Samples	499/164	10K	3K	3K/3K/3K	12K
Metric	pass@1	F1	Accuracy	BLEU	Exact Match

Table 2: Statistics of the downstream task datasets (Py: Python, JS: JavaScript).

3) Child prediction: given the graph tokens and the source code of one node, the model is tasked to predict the node’s children in the graph. This task is constructed on both AST and DFG data.

We sampled about 75K graph-question-answer tuples for each (language, graph type, question type) combination, which resulted in a total of 770K GraphQA data. Further details of these GraphQA tasks and their prompts are given in Appendix B.

4.1.2 Downstream Tasks

For downstream code tasks, we consider both discriminative and generative tasks, including code translation, clone detection, defect detection, code summarization, and code repair. For code translation, we use the CodefuseEval (Di et al., 2024) benchmark; for the other four tasks, we use the CodeXGLUE (Lu et al., 2021) benchmark. Among these tasks’ data, we downsampled the original data of code summarization and clone detection due to computation constraints, and filtered out 14K Java→Python samples in the training set of code translation where the target code does not follow Python coding conventions (e.g. no main function can be found). A summary of the final datasets is provided in Table 2. For all tasks we use the original metric for evaluation.

In the main experiments, we consider the setting of multi-task finetuning (MFT), where the model is simultaneously trained on all five downstream tasks. However, we demonstrate with a small-scale experiment that GALLa can be also used for single-task finetuning (SFT).

4.2 Models and Training

4.2.1 Model

We use a DUPLEX (Ke et al., 2024) with 1024 hidden states and 7M parameters as the GNN, and a cross-attention layer with randomly initialized learnable queries as the adapter. For the LLM, we consider six distinct models of varying sizes: CodeGen 350M (Nijkamp et al., 2023), Star-

Coder 1B (Li et al., 2023), Phi-1 1.3B (Gunasekar et al., 2023), LLaMA2 7B (Touvron et al., 2023) LLaMA3 8B (Dubey et al., 2024), and Qwen2.5-Coder 1.5B & 14B (Hui et al., 2024). All of these models are pretrained (partially) on source code data, and demonstrate strong performance on code-related downstream tasks.

4.2.2 Training

For the first stage training of GNN and adapter, we train the model on the graph data for 15 epochs with learning rate $1e-4$, 1K warmup steps, weight decay 0.1, 240 global batch size, AdamW optimizer (Loshchilov and Hutter, 2019), and ZeRO stage 2 (Rajbhandari et al., 2020). The training takes place on two machines, each equipped with 8 A100 80G GPUs.

The second stage of training largely follows the same setting but with a smaller learning rate ($5e-5$) and a smaller global batch size (96). The model is trained for 5 epochs on the mixture of downstream task data and graph alignment data in stage 2, and the checkpoint with the lowest validation loss on the downstream tasks is chosen for evaluation. Among the different LLMs, CodeGen, StarCoder, Phi-1, and Qwen2.5-Coder 1.5B are trained in full scale, while LLaMA2, LLaMA3, and Qwen2.5-Coder 14B are trained using LoRA (Hu et al., 2022) with rank 64. All training in this stage takes place on a single machine with 8 A100s.

As a baseline, the LLM is finetuned on only the downstream task data using the same hyperparameters as stage 2 training.

4.3 Results

The results of applying Graph2Code and GraphQA alignment to the code finetuning process of five models are presented in Table 3. Graph alignment brings consistent improvement over the baseline model, especially on the weaker backbones such as StarCoder, increasing the average performance on five tasks by up to 36%.

Notably, while our graph alignment data include only Python and Java programs, from Table 3 we

Model	Setting	Trans (Ja2Py/Py2Ja)	Clone	Repair	Sum (Java/Py/JS)	Defect	Avg
CodeGen 350M	Baseline	40.2/42.3	94.4	9.0	13.6/8.8/12.5	56.9	34.7
	G2C	42.7/43.3	94.6	8.5	13.9/9.5/12.8*	56.5	35.2 (+1%)
	G2C + GraphQA	50.0/45.3	93.9	8.6	14.0/9.5/12.7*	58.7	36.6 (+5%)
StarCoder 1B	Baseline	0.6/0.0	40.3	0.4	7.8/3.2/5.1	54.2	14.0
	G2C	6.1/1.2	50.4	0.6	9.4/4.3/6.8	54.8	16.7 (+20%)
	G2C + GraphQA	6.1/1.2	65.8	0.3	9.5/4.5/7.4	56.5	18.9 (+36%)
Phi-1 1.3B	Baseline	73.8/53.3	93.5	10.1	14.6/11.3/13.8	61.1	41.4
	G2C	67.7/64.3	94.3	10.9	15.1/11.1/13.4*	60.4	42.2 (+2%)
	G2C + GraphQA	72.0/66.3	94.9	10.2	15.2/11.7/13.6*	60.3	43.0 (+4%)
Qwen2.5-Coder 1.5B	Baseline	35.4/73.8	95.2	13.2	14.4/10.8/13.6	58.4	39.3
	G2C	64.0/77.4	95.5	15.7	14.7*/11.5/13.4	60.3	44.1 (+12%)
	G2C + GraphQA	65.2/76.0	95.7	15.0	15.0/11.3/14.3	61.3	44.2 (+12%)
LLaMA2 7B	Baseline	59.7/56.1	40.0	0.8	2.0/1.5/1.8	53.4	26.9
	G2C	69.5/59.1	41.5	1.0	2.4/1.7/1.9*	53.5	28.8 (+7%)
	G2C + GraphQA	64.0/53.3	39.0	1.6	2.2/1.6*/2.0	55.2	27.4 (+2%)
LLaMA3 8B	Baseline	80.5/74.8	94.7	12.1	14.2/11.2/12.3	56.7	44.6
	G2C	80.5/77.0	94.9	13.4	14.0*/11.7/11.8	56.7	45.0 (+1%)
	G2C + GraphQA	80.5*/78.2	94.9	13.6	14.5/11.7/12.9	57.1	45.4 (+2%)
Qwen2.5-Coder 14B	Baseline	38.4/82.1	94.3	12.7	18.0/20.2/14.5	58.5	42.3
	G2C	50.0/82.8	95.1	13.1	18.1*/19.9/15.7	59.1	44.2 (+4%)
	G2C + GraphQA	64.0/83.4	94.4	13.8	17.9*/20.5*/15.9	58.0	46.0 (+9%)

Table 3: Main results. For each model, the first row is baseline LLM finetuned on downstream task data only; the second row is GALLa finetuned on downstream task and Graph2Code data; the third row is GALLa finetuned on downstream tasks, Graph2Code, and GraphQA data. Relative performance increases w.r.t. baseline are given in parentheses in the last column. All differences from the baseline performance are statistically significant ($p < 0.1$) except for those marked with * (the complete results are given in Appendix C).

	Trans (Ja2Py/Py2Ja)	Clone	Repair	Sum (Java/Py/JS)	Defect	Avg _{all}	Avg _{java,py}
Baseline	76.8/66.1	95.5	13.5	15.2/11.3/14.2	65.7	44.8	46.43
G2C	79.3/68.9	95.8	15.1	15.2/11.7/14.1	64.8	45.6 _{+2%}	47.67 _{+3%}
G2C+GQA	75.6/69.5	95.7	13.6	15.3/11.7/13.9	59.9	44.4 _{-1%}	46.90 _{+1%}

Table 4: Results of single-task finetuning with Phi-1. Relative performance increases w.r.t. baseline are given in subscripts in the last two columns. The penultimate column is the average of all tasks, while the last column is the average of all Java and Python tasks. G2C: Graph2Code. GQA: GraphQA.

observe that they can even improve six of the seven models’ performance on tasks in other languages - code summarization in JavaScript, and defect detection in C. This showcases that the knowledge about code structures acquired in GALLa can be generalized across programming languages, as learning to align to Python and Java structural graphs improves the finetuning performance of downstream tasks in other languages.

In Table 4, we also present the results of single-task finetuning with Phi-1. In this setting, we find that Graph2Code still improves the average performance on all tasks by 2%, and by 3% on Python and Java tasks. On the other hand, GraphQA brings limited improvement, and even a large drop on the C defect detection task. We hypothesize that this

is because the diverse data in GraphQA serve a similar role to instruction tuning for LLMs (Sanh et al., 2022; Chung et al., 2022), where the benefits of cross-task transfer only start to manifest when the number of tasks is large. We find that models trained with single-task finetuning are prone to hallucinations, where they do not follow instructions in test samples but output answer templates learned from the GraphQA data instead (see Appendix D for examples).

4.4 Ablation Studies

Contribution of AST and DFG As an ablation study, we conducted experiments with Phi-1 and Graph2Code using either only the AST data or only the DFG data for graph alignment. The results

	Trans (Ja2Py/Py2Ja)	Clone	Repair	Sum (Java/Py/JS)	Defect	Avg
Baseline	76.8/66.1	95.5	13.5	15.2/11.3/14.2	65.7	44.8
G2C	79.3/68.9	95.8	15.1	15.2/11.7/14.1	64.8	45.6
G2C (AST only)	76.8/66.9	95.8	13.6	15.1/11.9/13.8	65.2	44.9
G2C (DFG only)	79.3/68.3	95.8	13.7	15.1/11.9/13.8	66.3	45.5
G2C (Code only)	78.7/69.3	95.1	12.4	15.2/12.3/14.0	62.9	45.0

Table 5: Ablation studies on the Graph2Code (G2C) task with Phi-1.

	Training	Trans (Ja2Py/Py2Ja)	Clone	Repair	Sum (Java/Py/JS)	Defect	Avg
MLP	G2C	73.2/66.9	95.2	11.7	15.1/11.6/14.1	61.6	43.7
	G2C + GQA	76.2/68.9	95.3	12.7	14.9/12.1/13.6	59.9	44.2
MagNet	G2C	75.6/65.3	94.9	11.2	14.5/11.3/13.9	63.1	43.7
	G2C + GQA	74.4/65.1	95.0	12.8	15.0/12.4/14.1	63.2	44.0

Table 6: Experiments with Phi-1 using an MLP instead of a cross-attention layer as the adapter (top), and using MagNet instead of DUPLEX as the graph encoder (bottom).

are shown in Table 5. We find that DFG brings more improvement compared with AST, which we contribute to the fact that AST is more closely related to the surface form of source code and thus provides the model with less additional structural information, while DFG includes more complex structures - such as loops - that are more informative.

Contribution of the Graph Information Observing the previous results, one may raise the question: are these improvements indeed attributable to the alignment of code representations to graphs, or are they simply a result of additional computation expense during the finetuning of LLMs?

To answer this question, we conducted a control experiment where we additionally trained a model on the mixture of downstream tasks and the source code from the graph alignment data, but did not provide it with the graph information - in other words, Graph2Code with only the code but not the graph, as illustrated in Figure 3, which is similar in essence to continual pretraining on the source code data. The results are given in the last row in Table 5, which suggests that the graph information is indeed helping the LLM to better understand programs.

Different Adapters and Graph Encoders To show that the proposed GALLa framework can be applied to various GNN encoders and adapter modules, we also conducted another two sets of experiments in the MFT setting with Phi-1: based on the settings in the main experiments, we 1) re-




Figure 3: The setting of control experiment, where the model is trained on the same source code data as the Graph2Code task in Figure 1 but the graph information is not given, i.e. Graph2Code without the graph.

placed the cross-attention adapter with a 3-layer MLP (which has a similar parameter count to the cross-attention layer), and 2) replaced the graph encoder with MagNet (Zhang et al., 2021). For these two experiments, we repeated the complete two stages of GALLa training, and the results are presented in Table 6.

Compared with the results using cross-attention from Table 3, we surprisingly find that the MLP adapter leads to a slightly better performance, which is counter-intuitive as MLP does not allow for information exchange between the nodes. We hypothesize that this is because with the same parameter count, MLP has more layers than cross-attention (3 vs. 1 in our case), thus having a stronger expressivity.

Similarly, when using MagNet (which is convolution-based) instead of DUPLEX (which is self-attention-based) for finetuning Phi-1, the model’s performance on downstream tasks is also

slightly improved. These results of using different adapters and graph encoders, together with the experiments of using different LLMs (Table 3), confirm that GALLa is indeed a flexible framework where each of the three modules - GNN encoder, adapter, and LLM decoder - can be replaced as more advanced models are proposed in the future, leading to better performance on downstream tasks.

5 Conclusion

In this work, we present the conceptual designs, implementation details, and experimental results of GALLa - Graph Aligned Large Language Models for improved source code understanding. Unlike previous works that modify code language models' internal structures to enhance them with graph information, GALLa follows the cross-modality alignment paradigm and leaves the structure of the language models intact, making it applicable to any off-the-shelf code LLMs. By integrating GALLa as an auxiliary task on a separate graph alignment dataset in the finetuning stage of code LLMs, we require no graph information of the task-specific finetuning or evaluation data, thus incurring no additional computation cost compared with the baseline LLM at inference time. Our experiments validate the effectiveness of GALLa on various code downstream tasks with base LLMs of varying sizes, paving road for a new paradigm in integrating code structural graphs into language models and providing insights for future research in developing structure-aware code LLMs.

Limitations

As we pointed out in the related work section, all existing methods that attempt to combine structural graphs with code language models focus on small, BERT-style models. Since our work is the first to utilize structural graphs to enhance decoder-only LLMs in understanding source code, we are unable to find any existing methods for fair comparison, and thus can only fine-tune the LLMs directly on source code as a baseline. However, our work can serve as a baseline for any future research in graph-enhanced (or more generally, structure-aware) code LLMs.

Due to our limited computational budget, we mostly focused on LLMs around 1B-10B parameters (especially in the ablation studies) and only verified the effectiveness of GALLa on one 14B

model. We leave the exploration of even larger structure-aware code LLMs to future work. Also due to the budget constraint, we mainly focused on multi-task training in our main experiments and did not conduct individual hyper-parameter searches for the five downstream tasks. However, we maintained a fair comparison between the baseline and the proposed methods.

Ethical Statements

All the datasets and models used in this paper are publicly available resources. We use them only for research purpose.

Acknowledgments

This paper is supported by Ant Group and the General Program of National Natural Science Foundation of China (62176153).

References

- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. [Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond](#).
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#).
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le,

and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).

DeepSeek-AI, Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y. Wu, Yukun Li, Huazuo Gao, Shirong Ma, Wangding Zeng, Xiao Bi, Zihui Gu, Hanwei Xu, Damai Dai, Kai Dong, Liyue Zhang, Yishi Piao, Zhibin Gou, Zhenda Xie, Zhewen Hao, Bingxuan Wang, Junxiao Song, Deli Chen, Xin Xie, Kang Guan, Yuxiang You, Aixin Liu, Qiushi Du, Wenjun Gao, Xuan Lu, Qinyu Chen, Yaohui Wang, Chengqi Deng, Jiashi Li, Chenggang Zhao, Chong Ruan, Fuli Luo, and Wenfeng Liang. 2024. [Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence](#).

Peng Di, Jianguo Li, Hang Yu, Wei Jiang, Wenting Cai, Yang Cao, Chaoyu Chen, Dajun Chen, Hongwei Chen, Liang Chen, Gang Fan, Jie Gong, Zi Gong, Wen Hu, Tingting Guo, Zhichao Lei, Ting Li, Zheng Li, Ming Liang, Cong Liao, Bingchang Liu, Jiachen Liu, Zhiwei Liu, Shaojun Lu, Min Shen, Guangpei Wang, Huan Wang, Zhi Wang, Zhaogui Xu, Jiawei Yang, Qing Ye, Gehao Zhang, Yu Zhang, Zelin Zhao, Xunjin Zheng, Hailian Zhou, Lifu Zhu, and Xianying Zhu. 2024. [Codefuse-13b: A pretrained multilingual code large language model](#). In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice, ICSE-SEIP 2024, Lisbon, Portugal, April 14-20, 2024*, pages 418–429. ACM.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lomakin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuen-

ley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Almeida, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Roman Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoqiang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collet, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delprat Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat

- Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsim-poukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiaojuan Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. [The llama 3 herd of models](#).
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. [Codebert: A pre-trained model for programming and natural languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 1536–1547. Association for Computational Linguistics.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks are all you need](#).
- Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. [Unixcoder: Unified cross-modal pre-training for code representation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2022, Dublin, Ireland, May 22–27, 2022, pages 7212–7225. Association for Computational Linguistics.
- Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin B. Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. 2021. [Graphcodebert: Pre-training code representations with data flow](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7, 2021*. OpenReview.net.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. 2024. [Deepseek-coder: When the large language model meets programming – the rise of code intelligence](#).
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katherine Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack W. Rae, and Laurent Sifre. 2022. [An empirical analysis of compute-optimal large language model training](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Kai Dang, An Yang, Rui Men, Fei Huang, Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and Junyang Lin. 2024. [Qwen2.5-coder technical report](#). *CoRR*, abs/2409.12186.
- Xue Jiang, Zhuoran Zheng, Chen Lyu, Liang Li, and Lei Lyu. 2021. [Treebert: A tree-based pre-trained model for programming language](#). In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021*, volume 161 of *Proceedings of Machine Learning Research*, pages 54–63. AUAI Press.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#).
- Zhaoru Ke, Hang Yu, Jianguo Li, and Haipeng Zhang. 2024. [DUPLEX: Dual GAT for complex embedding of directed graphs](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 23430–23448. PMLR.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy V, Jason T. Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvasi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2023. [Starcoder: may the source be with you!](#) *Trans. Mach. Learn. Res.*, 2023.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin B. Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie Liu. 2021. [Codexglue: A machine learning benchmark dataset for code understanding and generation](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2023. [Codegen: An open large language model for code with multi-turn program synthesis](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Changan Niu, Chuanyi Li, Vincent Ng, Jidong Ge, Liguang Huang, and Bin Luo. 2022. [Spt-code: Sequence-to-sequence pre-training for learning source code representations](#). In *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25-27, 2022*, pages 1–13. ACM.
- Han Peng, Ge Li, Wenhan Wang, Yunfei Zhao, and Zhi Jin. 2021. [Integrating tree path in transformer for code representation](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 9343–9354.
- Han Peng, Ge Li, Yunfei Zhao, and Zhi Jin. 2022. [Rethinking positional encoding in tree transformer for code representation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 3204–3214. Association for Computational Linguistics.
- Ruchir Puri, David S. Kung, Geert Janssen, Wei Zhang, Giacomo Domeniconi, Vladimir Zolotov, Julian Dolby, Jie Chen, Mihir R. Choudhury, Lindsey Decker, Veronika Thost, Luca Buratti, Saurabh Pujar, Shyam Ramji, Ulrich Finkler, Susan Malaika, and Frederick Reiss. 2021. [Codenet: A large-scale AI for code dataset for learning a diversity of coding tasks](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou,

- Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. [Zero: memory optimizations toward training trillion parameter models](#). In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2020, Virtual Event / Atlanta, Georgia, USA, November 9-19, 2020*, page 20. IEEE/ACM.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2024. [Code llama: Open foundation models for code](#).
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Marc Szafraniec, Baptiste Rozière, Hugh Leather, Patrick Labatut, François Charton, and Gabriel Synnaeve. 2023. [Code translation with compiler representations](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Sindhu Tipirneni, Ming Zhu, and Chandan K. Reddy. 2024. [Structcorder: Structure-aware transformer for code generation](#). *ACM Trans. Knowl. Discov. Data*, 18(3):70:1–70:20.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura,
- Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Xin Wang, Yasheng Wang, Fei Mi, Pingyi Zhou, Yao Wan, Xiao Liu, Li Li, Hao Wu, Jin Liu, and Xin Jiang. 2021. [Syncobert: Syntax-guided multi-modal contrastive pre-training for code representation](#).
- Yue Wang, Hung Le, Akhilesh Gotmare, Nghi D. Q. Bui, Junnan Li, and Steven C. H. Hoi. 2023. [Codet5+: Open code large language models for code understanding and generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 1069–1088. Association for Computational Linguistics.
- Xitong Zhang, Yixuan He, Nathan Brugnone, Michael Perlmutter, and Matthew J. Hirn. 2021. [Magnet: A neural network for directed graphs](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 27003–27015.
- Ziyin Zhang, Chaoyu Chen, Bingchang Liu, Cong Liao, Zi Gong, Hang Yu, Jianguo Li, and Rui Wang. 2024. [Unifying the perspectives of nlp and software engineering: A survey on language models for code](#).
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2024. [Minigpt-4: Enhancing vision-language understanding with advanced large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

A Introduction to Structural Graphs

In this work, we refer to Abstract Syntax Trees (ASTs) and Data Flow Graphs (DFGs) as program structural graphs. These are graph data structures (i.e. nodes and edges) that specify the internal logic of programs. Each node in these graphs is a snippet of code within the entire program.

A.1 AST

An AST is a tree representation of the syntactic structures of a program, with the root node being the entire program, and each leaf node being a single semantic unit in the program - such as an identifier (e.g. a variable, a class, a function) or an operator. Each non-leaf node in the tree is a combination of its children, indicating the structures of the code. As the name suggests, AST is an abstract representation of source code, in the sense that certain details such as white spaces, parentheses, and delimiters are omitted. Every node in the tree has an associated type attribute, such as *assignment expression* or *for loop*. A toy example in Python is provided below:

```
def add(a, b):
    return a+b
```

In the AST for this example, the root node is a *function expression* including the whole code snippet. The rest of the nodes are given in Table 7.

A.2 DFG

A DFG is a graph representation of the variable use-define chains within a program at execution time. For a given program, its DFG shares the same set of nodes with its AST, while the edges indicate the data flow between variable-involved nodes. An example DFG of the previously shown program is provided in Table 8.

A.3 Universal AST

Typically, ASTs are specific to the programming language in question. In other words, different programming languages would have different types of nodes in their ASTs - for example, a *for loop* node in Python may have different semantics from a *for loop* in C++. However, in the context of our work, such language-specific specifications may be detrimental to the cross-language alignment (e.g. the code translation task) or generalization capabilities (i.e. generalizing to languages not present in the training data) of the models.

idx	content	type	parent idx
1	a, b	<i>arguments</i>	0 (root)
2	return a+b	<i>return statement</i>	0 (root)
3	a	<i>arguments</i>	1
4	b	<i>arguments</i>	1
5	a+b	<i>binary expression</i>	2
6	a	<i>variable</i>	5
6	b	<i>variable</i>	5

Table 7: An illustrative example AST for the toy program in Appendix A.1. Each row is a node in the graph.

idx	from	to
1	<i>variable a</i>	<i>binary expression a+b</i>
2	<i>variable b</i>	<i>binary expression a+b</i>
3	<i>arguments a</i>	<i>variable a</i>
4	<i>arguments b</i>	<i>variable b</i>

Table 8: An illustrative example DFG for the toy program in Appendix A.1. Each row is an edge in the graph.

Thus, in this work we used a special type of AST (and DFG, as they share the same set of nodes): Universal AST (UAST). In UAST, the specifications of all node types are designed to be as language-independent as possible. Taking the nodes in Table 7 as an example, UAST abstracts language-specific node types into the most basic concepts shared by most programming languages, such as variable, binary expression, and return statement. In total, there are 43 node types in our graph data.

B GraphQA Prompts

In our main experiments, we used three types of questions for GraphQA: edge prediction, parent prediction, and child prediction. For each type of question, we wrote about ten question templates and ten answer templates, as shown in Figure 4 to 6. When constructing the data, one question template and one answer template are randomly chosen for each sample, as diverse question templates have been found to improve cross-task generalization capability in instruction finetuning (Sanh et al., 2022). The placeholders in the templates are replaced with the actual node contents and types, and the instantiated prompt is then randomly placed before or after the graph tokens.

In our preliminary experiments, we found edge prediction data constructed from AST to bring no

improvement, as connected nodes in the AST often have high textual overlap and can be trivially predicted. Thus we only use DFG data for the edge prediction task, while both AST and DFT data are used for the other two tasks. In early trials, we also experimented with and eventually discarded several other types of questions, including node classification (which can be easily done by only looking at the node’s source code) and some tasks that involve counting, such as counting the number of nodes in the graph, the number of edges in the graph, and the number of node types in the graph, which prove to be too difficult for the model.

C Statistical Tests

To verify the statistical significance of the main results, we conducted Chi squared tests on the four tasks with discrete performance metrics (i.e. code translation, code repair, clone detection, and defect detection) and Wilcoxon signed-rank test on the task with continuous performance metrics (i.e. code summarization). The results are presented in Table 9. Most of the differences are significant, except for CodeGen and Phi-1 on JavaScript summarization, and LLaMA3 on Java-to-Python translation and Python summarization.

D Example Outputs with Single-Task Finetuning

In Figure 7, we provide an example of model hallucination after training on only one downstream task in the GALLa framework. The model produces an answer template learned from the GraphQA task instead of responding to the actual question. This issue is mitigated in the multitask finetuning setting.

Question templates:

- In the graph, is there an edge from {node_type1} {node1} to {node_type2} {node2}?
- In the graph, is there an edge pointing from {node_type1} {node1} to {node_type2} {node2}?
- Please tell me if there is an edge pointing from {node_type1} {node1} to {node_type2} {node2} in this graph.
- Is there an edge from {node_type1} {node1} to {node_type2} {node2} in this graph?
- Does a connection exist from {node_type1} {node1} to {node_type2} {node2} in the graph?
- In this graph, do we have an edge leading from {node_type1} {node1} to {node_type2} {node2}?
- Is it true that {node_type1} {node1} is a predecessor of {node_type2} {node2} in this graph?

Answer templates 1 (positive):

- Yes, that is the case.
- Yes, there is an edge from {node_type1} {node1} to {node_type2} {node2}.
- Yes, there is an edge from {node_type1} {node1} to {node_type2} {node2} in this graph.
- Yes, there is an edge pointing from {node_type1} {node1} to {node_type2} {node2} in this graph.
- Affirmative, there exists an edge from {node_type1} {node1} to {node_type2} {node2}.
- Yes, that is the case. {node1} is directly connected to {node2}.

Answer templates 2 (negative):

- No, that is not the case.
- No, {node_type1} {node1} is not linked to {node_type2} {node2} by any edge in this graph.
- No, there is no edge from {node_type1} {node1} to {node_type2} {node2}.
- No, such an edge is absent from the graph.
- The graph does not show {node_type1} {node1} as a predecessor to {node_type2} {node2}.

Figure 4: Prompts for the edge prediction GraphQA task.

Model	Setting	Trans (Ja2Py/Py2Ja)	Clone	Repair	Sum (Java/Py/JS)	Defect
CodeGen 350M	G2C	0.0862/0.0000	0.0000	0.0000	0.0114/0.0000/0.3684	0.0000
	G2C + GQA	0.0000/0.0001	0.0000	0.0000	0.0111/0.0011/0.4332	0.0000
StarCoder 1B	G2C	0.0057/0.0140	0.0000	0.0000	0.0000/0.0000/0.0000	0.0000
	G2C + GQA	0.0057/0.0140	0.0000	0.0000	0.0000/0.0000/0.0000	0.0000
Phi-1 1.3B	G2C	0.0000/0.0000	0.0000	0.0000	0.0090/0.0595/0.5485	0.0000
	G2C + GQA	0.0009/0.0000	0.0000	0.0000	0.0058/0.0048/0.7818	0.0000
Qwen2.5-Coder 1.5B	G2C	0.0000/0.0000	0.0000	0.0000	0.2176/0.0001/0.0000	0.0000
	G2C + GQA	0.0286/0.0000	0.0000	0.0000	0.0008/0.0003/0.0001	0.0000
LLaMA2 7B	G2C	0.0037/0.0000	0.0000	0.0000	0.0195/0.0755/0.5454	0.0000
	G2C + GQA	0.0000/0.0000	0.0000	0.0000	0.0468/0.1406/0.0147	0.0000
LLaMA3 8B	G2C	0.0003/0.0000	0.0000	0.0000	0.5029/0.0000/0.0014	0.0000
	G2C + GQA	1.0000/0.0000	0.0000	0.0000	0.0019/0.0147/0.0001	0.0000
Qwen2.5-Coder 14B	G2C	0.0000/0.0000	0.0000	0.0000	0.9403/0.0213/0.0000	0.0000
	G2C + GQA	0.0509/0.0000	0.0000	0.0000	0.2508/0.6213/0.0000	0.0000

Table 9: Statistical significance (p values) of main results’ differences from the baselines.

Question templates:

- In the graph, what is the parent node of this {node_type} {node}.
- What is the parent of {node_type} {node} in this graph?
- What is the parent node of {node_type} {node} in the graph?
- Based on the graph, identify the parent of {node_type} {node}.
- Based on this graph, identify the parent of this {node_type} {node}.
- Identify the parent of {node_type} {node} in the graph.
- In the graph presented, what is the predecessor of {node_type} {node}?
- What node acts as the parent to {node_type} {node} in the graph displayed?
- Can you determine the parent node of {node_type} {node} in this graph?
- Which node is directly above {node_type} {node} in the hierarchy of the provided graph?
- What is the immediate ancestor of the {node_type} {node} in this graph?
- Regarding the graph, can you point out the parent of {node_type} {node}?
- In terms of graph theory, what is the parent of the {node_type} {node}?
- Who has the parental role for {node_type} {node} in the graph's topology?
- For {node_type} {node} in the given graph, which node supplies the incoming edge?

Answer templates 1 (has parent):

- In the given graph, the parent of the given {node_type} is {parent}, which is a {parent_type}.
- This {node_type}'s parent is the {parent_type} {parent}.
- The given {node_type}'s parent in the graph is the {parent_type} {parent}.
- The parent of {node_type} {node} in this graph is identified as {parent}, categorized as a {parent_type}.
- Node {parent}, a {parent_type}, serves as the parent to {node_type} {node} in the graph.
- As per the hierarchy, the {parent_type} node {parent} is the direct predecessor to {node_type} {node}.
- Upon inspection, it is clear that the parent of {node_type} {node} is the {parent_type} {parent}.
- The {node_type} {node} is immediately descended from {parent}, a {parent_type} in the graph.
- Within the nodal arrangement, {parent} is the progenitor to {node_type} {node}, having the classification of a {parent_type}.
- Tracing the edges leads to confirming {parent}, a {parent_type}, as the parent of {node_type} {node}."

Answer templates 2 (no parent):

- This {node_type} has no parent in the graph.
- There is no edge pointing to this {node_type} in the given graph. Therefore it does not have any parent.
- Within this graph, {node_type} {node} does not have a parent node.
- {node_type} {node} stands without a parent in the graph's existing structure.
- No parent node is associated with {node_type} {node} in the provided graph.
- A review of the graph establishes that there is no preceding node to {node_type} {node}; it has no parent.
- In this graph topology, {node_type} {node} is an orphan node with no parent.
- There is no edge incoming to {node_type} {node}, indicating the absence of a parent.
- After analyzing the graph, it becomes evident that {node_type} {node} lacks a directly linked parent node.
- As depicted in the graph, {node_type} {node} exists without a parent node.

Figure 5: Prompts for the parent prediction GraphQA task.

Question templates:

- In this graph, what are the children of this {node_type} {node}.
- Identify all children of {node_type} {node} in this graph.
- Find the child nodes of {node_type} {node} in the graph.
- In the graph, how many children does the {node_type} {node} have? What are they?
- How many children does {node_type} {node} have in this graph? What are they?
- Please find all children of {node_type} {node} in this graph.
- Can you find all children of {node_type} {node} in this graph?
- List all the descendant nodes of {node_type} {node} in this graph.
- What are the direct children of the {node_type} {node}?
- Can you enumerate the offspring of {node_type} {node} within this graph?
- Could you provide the list of child nodes attached to {node_type} {node}?
- Please identify the child nodes emanating from {node_type} {node}.
- Show me the child nodes of {node_type} {node}.
- What nodes are directly connected to {node_type} {node} as its children?
- I need to know all the child elements of {node_type} {node}. Can you provide that?
- Are there any nodes that directly derive from {node_type} {node} in this graph?
- Which nodes act as successors to the node tagged as {node_type} {node}?
- What are the adjacent nodes that are children of {node_type} {node}?
- Identify the nodes that are immediate successors of {node_type} {node} in this graph.
- Detail the nodes branching from {node_type} {node} in this graph structure.
- Reveal all nodes that are directly beneath {node_type} {node} in the hierarchy.

Answer templates 1 (has children):

- The given {node_type} has {child_num} children in the graph, they are: {child_nodes}
- This {node_type} has {child_num} children: {child_nodes}
- {node_type} {node} has a total of {child_num} children in this graph, which are: {child_nodes}
- There are {child_num} child nodes of {node_type} {node}, specifically: {child_nodes}
- As for the children of {node_type} {node}, you will find {child_num} direct descendants: {child_nodes}
- The count of {node_type} {node}'s children amounts to {child_num}. They include: {child_nodes}
- {node_type} {node} is parent to the following {child_num} nodes: {child_nodes}
- A list of the {child_num} children under {node_type} {node} is as follows: {child_nodes}
- Directly under {node_type} {node}, there are {child_num} children listed as: {child_nodes}
- {child_num} children spring from {node_type} {node}, which are given below: {child_nodes}

Answer templates 2 (no children):

- This {node_type} does not have any child nodes in the graph.
- This {node_type} does not have any children in the graph.
- There are no children of this {node_type} in the given graph.
- The given {node_type} does not have any children in the graph.
- After examining the graph, it's determined that this {node_type} has no children.
- I've checked the {node_type} {node} and found it has no direct descendants.
- There are no child nodes attached to {node_type} {node} in this graph.
- No descendants can be traced from this {node_type}.
- The {node_type} {node} is devoid of child nodes within the current graph structure.
- It appears {node_type} {node} has no children.

Figure 6: Prompts for the child prediction GraphQA task.

```

** Prompt **:
# Translate the following Java program to Python

# Java
public Object splitWords(String txt) {
    if (txt.contains(" ")) {
        return Arrays.asList(txt.split(" "));
    } else if (txt.contains(",")) {
        return Arrays.asList(txt.split("[\\s]"));
    } else {
        int count = 0;
        for (char c : txt.toCharArray()) {
            if (Character.isLowerCase(c) && (c - 'a') % 2 == 1) {
                count += 1;
            }
        }
        return count;
    }
}

# Python
def split_words(txt):

** Expected answer **:
    if " " in txt:
        return txt.split(" ")
    elif "," in txt:
        return txt.split(",")
    else:
        count = 0
        for c in txt:
            if c.islower() and (ord(c) - ord('a')) % 2 == 1:
                count += 1
        return count

** Model Response **:
There is 1 connected component in the program's DFG.

```

Figure 7: An example of hallucination after single-task finetuning.