

Article

Not peer-reviewed version

---

# Small Language Models as Graph Classifiers: Towards Lightweight Model Adaptation for Graph-Structured Data

---

[Michal Podstawski](#) \*

Posted Date: 12 February 2026

doi: 10.20944/preprints202602.0946.v1

Keywords: graph classification; small language models; graph representation learning



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Small Language Models as Graph Classifiers: Towards Lightweight Model Adaptation for Graph-Structured Data

Michal Podstawski 

NASK National Research Institute, Warsaw, Poland; michal.podstawski@nask.pl

## Abstract

Graph classification is a fundamental task in graph representation learning and is traditionally addressed using graph neural networks that incorporate carefully designed inductive biases. This work focuses on small language models and investigates their suitability for graph classification tasks. By representing graphs in a textual form that captures both structural relationships and node features, we enable language models to operate directly on graph data without relying on graph-specific architectures. Across standard graph classification benchmarks, we show that small language models can learn meaningful graph representations and achieve competitive performance relative to established graph-based methods. Our findings highlight the potential of small language models as a flexible alternative for graph classification, particularly in rapidly changing settings where architectural simplicity and adaptability are critical.

**Keywords:** graph classification; small language models; graph representation learning

## 1. Introduction

Graph-structured data arise naturally in many domains such as chemistry, biology, and social networks. Graph classification, which assigns a label to an entire graph, is a central task in graph representation learning and has been extensively studied in recent years. Graph neural networks (GNNs) have become the dominant approach for this problem due to their ability to explicitly model relational structure through neighborhood aggregation and message passing [1,2].

Despite their success, GNNs rely on task-specific architectural design and strong inductive biases tailored to graph-structured data. In contrast, language models have emerged as general-purpose learners capable of solving a wide variety of tasks with minimal architectural modification. Recent studies have explored applying large language models to structured and non-textual data, including graphs [9], but these approaches often depend on very large models or complex prompting strategies, limiting their applicability in resource-constrained settings.

In this work, we investigate a simple yet underexplored question: can small language models act as effective graph classifiers? We focus on pretrained small language models and study their ability to perform graph classification when adapted to graph-structured inputs. Graphs are serialized into text using edge list-based representations that capture both structural relationships and node-level information, allowing language models to process graphs without relying on specialized graph-specific architectures. Through this formulation, we examine whether small language models can learn informative graph representations and support accurate graph-level prediction.

We evaluate our approach on the benchmarks from the TU dataset collection [3]. To efficiently adapt pretrained models to this task, we employ Low-Rank Adaptation (LoRA) [4], which enables stable fine-tuning with a small number of trainable parameters. Our results show that small language models can learn meaningful graph representations from serialized edge lists and achieve competitive performance under this structure-only setting.

### Contributions

The main contributions of this work are as follows:

- We investigate the use of small language models for graph classification, examining their ability to operate on graph-structured data without relying on graph-specific architectures.
- We introduce a textual representation of graphs that enables pretrained language models to perform graph-level prediction.
- We conduct a systematic empirical evaluation on standard graph classification benchmarks, highlighting both the strengths and limitations of small language models in this setting.

## 2. Method

### 2.1. Problem Definition

We consider the graph classification problem, where each input graph is defined as  $G = (V, E, X)$ , with node set  $V$ , edge set  $E$ , and node feature matrix  $X$ . Given a dataset of labeled graphs  $\{(G_i, y_i)\}_{i=1}^N$ , the objective is to learn a function  $f: \mathcal{G} \rightarrow \mathcal{Y}$  that maps each graph to a discrete graph-level class label  $y_i \in \mathcal{Y}$ . In this work, each graph is associated with a single graph-level label, and both graph topology and node features are used as input for classification.

### 2.2. Graph-to-Text Serialization

To enable the use of language models for graph classification, each graph is serialized into a structured textual prompt. Given a graph with  $|V|$  nodes and  $|E|$  edges, the prompt includes a task description, basic graph statistics, serialized node feature information, and an explicit edge list.

Nodes are indexed from 0 to  $|V| - 1$ . Node features are converted into a textual format and listed explicitly, allowing the language model to condition on both graph structure and node-level attributes. The edge list enumerates all edges as ordered source–target pairs, with a fixed ordering to ensure deterministic serialization.

This serialization exposes complete graph topology together with node features in a form that is directly consumable by transformer-based language models, without relying on graph-specific architectures or inductive biases.

### 2.3. Model and Fine-Tuning

We experiment with three state-of-the-art small language models (SLMs) with at most one billion parameters: Qwen3-0.6B [7], Llama-3.2-1B [6], and Gemma-3-1B [8]. These models are selected as they represent the strongest publicly available SLMs in this parameter regime, having demonstrated competitive performance across a wide range of language understanding and reasoning benchmarks. Using these models allows us to evaluate the graph classification capability of modern SLMs under a realistic and competitive setting. All models are adapted to the graph classification task using parameter-efficient fine-tuning.

Graph classification is formulated as a sequence-to-label prediction problem. Given a serialized graph as input, the model is trained to predict the corresponding graph label. We implement this using a causal language modeling objective, where the model is prompted to generate the class label as the next token following the serialized graph description. During training and inference, the model is instructed to output only the class label as a single digit, enabling direct graph-level prediction without modifying the underlying language modeling formulation.

### 2.4. Low-Rank Adaptation

We fine-tune all small language models using Low-Rank Adaptation (LoRA) [4]. Across all experiments, we use a LoRA rank of  $r = 16$ , scaling factor  $\alpha = 32$ , and dropout 0.1. LoRA modules are applied to both attention and feed-forward projections, specifically targeting `q_proj`, `k_proj`, `v_proj`, `o_proj`, `gate_proj`, `up_proj`, and `down_proj`. This configuration keeps the number of trainable parameters below 1% of the full model size in all cases.

### 3. Experiments and Results

#### 3.1. Experimental Setup

We evaluate the proposed approach on graph classification benchmarks from the TU dataset collection [3]. We fine-tune models using Low-Rank Adaptation (LoRA) as described in Section 2.

#### 3.2. GNN Baseline

As a baseline, we implement a simple graph neural network (GNN) for graph classification using standard message-passing components. The model consists of three graph convolutional layers (GCN) with a hidden dimensionality of 64. Each convolutional layer is followed by batch normalization and a ReLU activation. Dropout with a rate of 0.5 is applied after the first two convolutional layers and again after graph-level pooling to reduce overfitting.

Node embeddings produced by the final convolutional layer are aggregated into a graph-level representation using global mean pooling. This pooled representation is then passed through a linear layer to predict graph-level labels.

The model is trained end-to-end using cross-entropy loss and optimized with Adam using a learning rate of 0.001. Training is performed for 100 epochs with a batch size of 32. The dataset is randomly split into 80% training, 10% validation, and 10% test sets. The best-performing model is selected based on validation accuracy, and final results are reported on the test split.


This architecture serves as a standard GNN baseline, providing a straightforward and widely used point of comparison for evaluating the effectiveness of language-model-based approaches on graph classification tasks.

#### 3.3. Results

Table 1 and Figure 1 reports the graph classification accuracy of the evaluated small language models, along with the GNN baseline, across the considered datasets.

**Table 1.** Graph classification test accuracy.

Model	PROTEINS	MUTAG	ENZYMES
GNN	0.7857	0.7500	0.3333
Qwen3-0.6B	<b>0.8036</b>	0.8500	0.2000
Llama-3.2-1B	0.7946	0.8500	<b>0.3500</b>
Gemma-3-1B	0.7768	<b>0.9000</b>	0.3167



**Figure 1.** Graph classification test accuracy for the GNN baseline and three small language models across datasets.

Overall, the proposed approach achieves performance that is competitive with, and in some cases superior to, the GNN baseline across the evaluated graph classification tasks. These results indicate that small language models can effectively capture discriminative graph patterns from serialized representations, despite the absence of explicit graph-specific inductive biases. While performance varies across tasks—particularly for more challenging multi-class settings—the results consistently

demonstrate that small language models are capable of learning meaningful representations of graph topology and supporting accurate graph-level prediction when appropriately adapted.

## 4. Conclusions

In this work, we investigated whether small language models can be adapted for graph classification using only structural information. By fine-tuning models with parameter-efficient LoRA adapters, we demonstrated that pretrained language models can learn meaningful graph representations directly from serialized edge lists. Experiments benchmarks show that small language models achieve competitive performance despite lacking explicit graph inductive biases. These results highlight the potential of small language models as flexible and resource-efficient alternatives for graph classification in low-resource settings.

**Acknowledgments:** This manuscript acknowledges the use of ChatGPT [10], powered by the GPT-5 language model developed by OpenAI, to improve language clarity, refine sentence structure, and enhance overall writing precision.

## References

1. Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks. In: 5th International Conference on Learning Representations (ICLR) (2017).
2. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How Powerful Are Graph Neural Networks? In: 7th International Conference on Learning Representations (ICLR) (2019).
3. Morris, C., Ritzert, M., Fey, M., Hamilton, W.L., Lenssen, J.E., Rattan, G., Grohe, M.: TUDataset: A Collection of Benchmark Datasets for Learning with Graphs. In: ICML Workshop on Graph Representation Learning and Beyond (2020).
4. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Chen, W.: LoRA: Low-Rank Adaptation of Large Language Models. In: 10th International Conference on Learning Representations (ICLR) (2022).
5. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: 3rd International Conference on Learning Representations (ICLR) (2015).
6. Meta AI: The Llama 3 Herd of Models. arXiv preprint arXiv:2407.21783 (2024).
7. Qwen Team: Qwen3 Technical Report. arXiv preprint arXiv:2409.12191 (2024).
8. Google DeepMind: Gemma: Open Models Based on Gemini Research and Technology. arXiv preprint arXiv:2403.08295 (2024).
9. Chen, Z., Zhang, Q., Wang, Y., Li, Z.: Exploring the Potential of Large Language Models for Graph Learning. arXiv preprint arXiv:2307.03393 (2023).
10. OpenAI: ChatGPT <https://openai.com/chatgpt>.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.