

# GRAPH FOUNDATION MODELS: BRIDGING LANGUAGE MODEL PARADIGMS AND GRAPH OPTIMIZATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The pretrain-transfer paradigm, which underpins the success of large language models (LLMs), has demonstrated the immense power of creating foundation models that learn generalizable representations from vast datasets. However, extending this paradigm to Operations Research (OR) problems on graph structures remains challenging due to the fundamental conflict between the statistical flexibility of language and the strict combinatorial constraints of graphs. To bridge this gap, we introduce the Graph Foundation Model (GFM), the first framework capable of solving all distance-based optimization problems on graph structures. By introducing the LLM-like self-supervised pre-training paradigm on the paths generated from random walks in the graph, GFM is compelled to internalize the graph’s complex topological and combinatorial rules, where the connectivity of the structure itself can be treated as the supervisory signal. Unlike existing neural methods that learn complex and task-specific solving policies, our approach leverages the pre-trained GFM as a foundational model of the graph’s intrinsic structure, which in turn enables a simple generative heuristic to tackle a diverse range of optimization challenges effectively. Comprehensive experiments on networks ranging from 20 to 893 nodes demonstrate that GFM achieves competitive performance against specialized solvers across a variety of distinct optimization task classes, while maintaining significantly faster inference times. Our work establishes a new paradigm of adapting the pretrain-transfer framework to graph optimization, opening the door for applying foundation model innovations to operations research.

## 1 INTRODUCTION

The remarkable success of large language models (LLMs) has cemented the pretrain-transfer paradigm as a transformative shift in how we approach generative intelligence. From their reasoning capabilities to recent advances in solving complex mathematical problems, LLMs have demonstrated an unprecedented capability and generalizability in capturing intricate patterns across diverse domains (Brown et al., 2020; Achiam et al., 2023; Yang et al., 2025). This paradigm shift raises a compelling question: *Can the foundational principles that enable LLMs to excel at language understanding be extended to solve optimization problems on graphs?*

Graph problems have long been studied in the operations research (OR) community, such as transportation, logistics, and network design. Still, the NP-hard nature of these tasks remains a central bottleneck: exact methods provide optimal solutions at prohibitive cost, whereas heuristics achieve tractability with rare optimality guarantees at the expense of generality (Papadimitriou & Steiglitz, 1981). Recent learning-based efforts fall into two broad strands: (i) problem-specific deep learning (DL) models and (ii) LLM-centric approaches. From pointer networks (Vinyals et al., 2015) and reinforcement learning (RL)-trained attention models to GNN/Transformer variants (Bello et al., 2016; Kool et al., 2018; Kwon et al., 2020) and diffusion solvers, problem-specific DL methods typically rely on specialized decoders/pointer mechanisms (Sun & Yang, 2023) and achieve strong results on synthetic complete graphs. LLM-centric approaches either cast optimization as text generation or use LLMs as modeling/decomposition agents with tool calls (Yang et al., 2023a; Zhang & Luo, 2025; Huang et al., 2024a). However, both strands are trained/evaluated on synthetic complete graphs and

054 task-specific architectures. The synthetic setting will result in weak feasibility/robustness under  
055 real-world constraints; task-specific architectures will be limited in scalability to large instances  
056 and exhibit poor cross-task generalization. As a result, they fail to capture transferable structural  
057 priors and cannot effectively handle realistic road-network settings where sparsity, geometry, and  
058 constraints play a central role.

059 In this study, we propose the Graph Foundation Model (GFM), a pretraining framework that learns  
060 a transferable structural prior over graphs. Analogous to how LLMs learn linguistic priors, GFM  
061 is pretrained on large, heterogeneous graph corpora to internalize topological and geometric reg-  
062 ularities without using any problem-specific supervision. With this universal structural prior and  
063 light task specifications, GFM can solve distance-based graph optimization problems through down-  
064 stream generation strategies, rather than engineering a solver per task.

065 Concretely, we transform raw graphs into rich, self-supervised training signals via structure-aware  
066 walks inspired by Node2Vec (Grover & Leskovec, 2016). Synthetic trajectories generated from  
067 random walks can be used to pretrain the GFM through hidden segment reconstruction, enabling  
068 the GFM’s interpretability in terms of the graph’s connectivity, reachability, and path consistency.  
069 Such a pre-training process transfers the concept of words and paragraphs in language to nodes and  
070 trajectories under a graph structure, enabling the widely accepted LLM paradigm to be adopted in a  
071 graph setting and yielding a task-agnostic backbone that captures reusable graph semantics.

072 At the stage of inference, GFM functions as a distribution approximator over feasible structures. A  
073 simple, task-agnostic decoding, combined with lightweight projection, steers the pretrained prior to-  
074 ward concrete objectives, such as shortest paths, tour variants, and subset selection, while enforcing  
075 hard constraints. Crucially, no problem-specific feedback is used during pretraining, in contrast to  
076 many neural OR pipelines relying on supervised labels or RL rewards tied to a single problem (Bello  
077 et al., 2016). This separation, universal pretraining of the structure, and the subsequent thin task-  
078 specific generation, drives strong generalizability across disparate graph optimization tasks on real  
079 road networks.

080 Empirically, our pretrained GFM backbone adapts to diverse distance-based graph optimization  
081 tasks, including multiple NP-hard families, without any architectural modification. Across scales  
082 from small synthetic graphs ( $N=20$ ) to large real networks ( $N=893$ ), GFM achieves competitive  
083 solution quality compared to widely used OR methods such as LKH3 and OR-tools, demonstrating  
084 GFM’s capability for graph optimization.

085 In summary, we have the following contributions:

086 **(1) First graph foundation model for realistic, distance-based graph problems.** We first in-  
087 troduce GFM, a single pretrained model with lightweight constraint projection that excels across  
088 multiple graph tasks, spanning from shortest path to NP-hard routing variants, e.g., tour families, on  
089 real traffic networks, without any architectural changes.

091 **(2) Pretraining with random walks and insertion-based reconstruction.** We utilize structure-  
092 aware random walks to capture the graph topology and an insertion-based reconstruction curriculum  
093 to equip the model with holistic structural priors.

094 **(3) Transferring the LLM pretrain and adapt paradigm to OR Problems** This work explores a  
095 new research direction that transfers the powerful LLM paradigm to OR fields. By demonstrating  
096 that the foundational training principles of large language models can be effectively applied to graph-  
097 based optimization, we pave the way for leveraging the full spectrum of LLM innovations in graph-  
098 based optimization problems.

## 100 2 RELATED WORK

103 **Neural Combinatorial Optimization** Neural approaches to combinatorial optimization have  
104 evolved from early Hopfield networks (Hopfield & Tank, 1985) to sophisticated architectures lever-  
105 aging modern deep learning advances. The foundational work of Vinyals et al. (2015) introduced  
106 Pointer Networks, demonstrating that sequence-to-sequence models could generate near-optimal  
107 Traveling Salesman Problem (TSP) solutions through attention mechanisms. This paradigm was  
enhanced through RL by Bello et al. (2016) and attention-based encoders by Kool et al. (2018),

108 who achieved competitive performance on synthetic TSP instances. The POMO framework (Kwon  
109 et al., 2020) further improved solution quality through multi-start parallel decoding strategies.

110 The transformer revolution has significantly advanced neural combinatorial optimization. Bres-  
111 son & Laurent (2021) achieved remarkable 0.004% optimality gaps on 50-city TSP instances using  
112 standard transformer encoders without pointer mechanisms, while Yang et al. (2023b) addressed  
113 scalability by reducing attention complexity from  $O(n^2)$  to  $O(n \log n)$  through Sampled Scaled Dot-  
114 Product Attention. More recently, diffusion-based approaches such as DIFUSCO (Sun & Yang,  
115 2023) and general neural CO frameworks like Luo et al. (2023) have extended the paradigm to  
116 broader optimization settings. Recent structure-aware approaches like Zhao & Wong (2025) incor-  
117 porate graph centrality measures and spatial encodings to capture topological relationships. Zhou  
118 et al. (2024) proposed to introduce an instance conditional adaptation module in the neural routing  
119 solver, enabling the model to adjust according to the features of the input graph dynamically.

120 However, these methods face critical limitations: they are typically trained on synthetic complete  
121 graphs, rely on task-specific architectures that cannot transfer knowledge across different optimiza-  
122 tion problems, and are fundamentally task-driven rather than structure-driven—designing solvers  
123 narrowly around individual problems without cultivating transferable structural priors. Conse-  
124 quently, they struggle to generalize to real-world sparse road networks where sparsity, geometry,  
125 and topological constraints are central to optimization objectives.

126 **Foundation Models for Optimization** The success of LLM has inspired new paradigms for OR.  
127 Direct solving approaches include Yang et al. (2023a)’s OPRO framework, which places LLMs in  
128 iterative improvement loops for gradual solution refinement, and Ghimire et al. (2025)’s autoregres-  
129 sive TSP solver that treats tours as token sequences, achieving competitive performance on 100-node  
130 instances through Direct Preference Optimization. Multimodal approaches, such as Elhenawy et al.  
131 (2024), explore visual reasoning for TSP solving, while evolutionary frameworks, like Liu et al.  
132 (2023b), employ LLMs as variation operators within genetic algorithms.

133 LLM-assisted optimization represents an alternative paradigm where language models handle prob-  
134 lem formulation rather than direct solving. Zhang & Luo (2025) developed OR-LLM-Agent for  
135 structured prompting and tool use in OR, while Huang et al. (2024a) introduced ORLM for trans-  
136 lating natural language descriptions into executable optimization code. Ye et al. (2024) proposed  
137 ReEvo, where LLMs generate entire heuristic algorithms through evolutionary search with natural  
138 language feedback.

139 Despite promising results, existing approaches face fundamental challenges that limit their practical  
140 applicability. Neural methods trained on synthetic complete graphs struggle with realistic networks  
141 where edge connectivity and geometric constraints are paramount. LLM-based approaches suffer  
142 from representation mismatches between natural language and graph-theoretic properties, scalability  
143 limitations, and an inability to capture the topological and geometric relationships essential for real-  
144 world optimization problems. Our work addresses these limitations by developing a foundation  
145 model that directly operates on graph structures while preserving the transferability and scalability  
146 advantages of the pretrain-transfer paradigm.

### 148 3 GRAPH FOUNDATION MODEL

149 Our work aims to instantiate the successful pre-train and fine-tune paradigm (Devlin et al., 2019;  
150 Brown et al., 2020) on graph structures. We propose a learnable structural prior to capture the  
151 graph’s intrinsic topological properties. To achieve this, we devise a self-supervised method based  
152 on random walk context reconstruction. We then demonstrate how this powerful prior enables a  
153 simple ad-hoc approach to tackle complex OR problems effectively.

#### 154 3.1 METHODOLOGY OVERVIEW

155 We now introduce the methodology of our GFM, whose overall framework is shown in Figure 1.

156 Let a weighted graph be  $G = (V, E, w)$ , and let  $\mathcal{Y}(G)$  denote the space of all structured candidates  
157 (e.g., paths, tours, subgraphs). The GFM first uses a distance-biased random walk to sample a graph  
158 corpus, i.e., a trajectory dataset. We then employ a bidirectional Transformer Encoder to learn a  
159

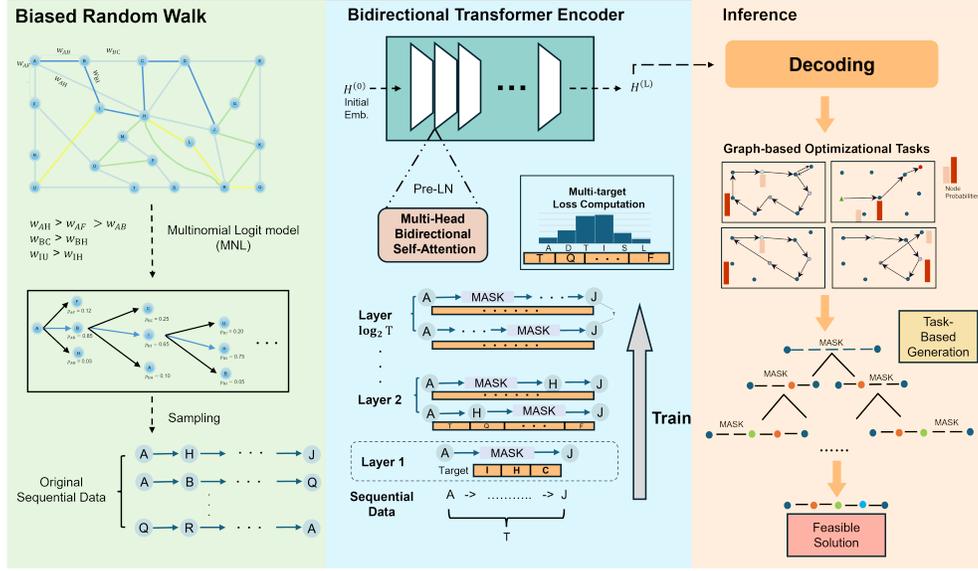


Figure 1: The overall framework of our proposed GFM.

Note. Progressive sample generation with distance-biased random walks provides training trajectories. Node and position embeddings are combined and passed into a Bidirectional Transformer Encoder, which captures graph context without causal masks. The model is optimized through a multi-target reconstruction process and adapted to various graph-based tasks through task-specific generation during inference.

structural prior distribution  $\pi(\mathbf{Y} | G)$  through an insertion-based trajectory reconstruction process. The learned structural prior  $\pi(\mathbf{Y} | G)$  is then used to generate required solutions. For any given task specification  $s$ , the feasible set  $\mathcal{F}(G, s) \subseteq \mathcal{Y}(G)$  contains only candidates that satisfy the task constraints, and  $J(\mathbf{Y}; G, s)$  denotes the associated objective:

$$\mathbf{Y}^* = \arg \min_{\mathbf{Y} \in \mathcal{F}(G, s)} J(\mathbf{Y}; G, s). \quad (1)$$

Rather than directly solving the exact objective above, our framework employs  $\pi(\mathbf{Y} | G)$  to generate task-consistent candidates  $\mathbf{Y} \in \mathcal{Y}(G)$ , by decoding from this prior under feasibility constraints.

### 3.2 RANDOM WALK TO GRAPH CORPORA

To establish a training signal for graphs, we transform graph topology into a sequential form through biased random walks. Specifically, given a sparse weighted graph  $G = (V, E, w)$ , we employ a Node2Vec-style biased random (Grover & Leskovec, 2016) walk with distance-aware utilities to generate path sequences that reflect both topology and metric relations. Formally, starting from node  $i = v_t$ , the probability of moving to a neighbor  $j$  given the previous node  $r = v_{t-1}$  is

$$\Pr(v_{t+1} = j | v_t = i, v_{t-1} = r) = \frac{\exp(-\beta \cdot w_{ij} \cdot b_{p,q}(i, j, r))}{\sum_{k \in \mathcal{N}(i)} \exp(-\beta \cdot w_{ik} \cdot b_{p,q}(i, k, r))}, \quad (2)$$

where  $w_{ij}$  is the edge weight,  $\beta > 0$  is a softmax scaling factor controlling the sharpness of the distribution, and  $b_{p,q}(i, j, r)$  denotes the Node2Vec bias term:  $p$  penalizes immediate backtracking to the previous node  $r$ , while  $q$  balances the likelihood of exploring outward versus staying local. We introduce a weight bias into the sample utility, allowing the frequency of trajectories to reflect their distance dissimilarities. The process ensures that both distance costs and structural exploration biases are taken into account during walk generation.

### 3.3 PROGRESSIVE TRAINING CURRICULUM

To learn the structure prior of the graph, we devise a self-supervised method based on the random walk context reconstruction.

**Algorithm 1** Progressive Curriculum Construction from Random Walks

---

**Require:** Random walk  $\mathbf{v} = [v_1, \dots, v_T]$  on graph  $G$ , max path length  $T_{\max}$ .  
**Ensure:** Training set  $\mathcal{D}$  of masked-prediction samples.

- 1: Initialize  $\mathcal{D} \leftarrow \emptyset$ , interior nodes  $\mathcal{V}_{\text{int}} = \{v_2, \dots, v_{T-1}\}$ .
- 2: Determine number of levels  $\ell_{\max} = \min(L_{\max}, \max(L_{\min}, \lceil \log_2 T \rceil))$ .
- 3: **Level 1 (global prior).** Construct  $\mathbf{x} = [v_1, \text{MASK}, v_T]$ , with admissible targets  $\mathcal{Y} = \mathcal{V}_{\text{int}}$ . Add  $(\mathbf{x}, \mathcal{Y}, \ell=1)$  to  $\mathcal{D}$ .
- 4: Initialize anchor set  $C^{(1)} = [v_1, v_T]$ . //  $C^{(\ell)}$ : anchor set at level  $\ell$
- 5: **for**  $\ell = 2$  **to**  $\ell_{\max}$  **do**
- 6:   Select a gap  $(a, b)$  from  $C^{(\ell-1)}$  and sample anchor  $u$  from  $\text{BETWEEN}(a, b \mid \mathbf{v})$  (if empty, sample from  $\mathcal{V}_{\text{int}}$ ).
- 7:   Insert  $u$  into  $C^{(\ell-1)}$  to form  $C^{(\ell)}$ . // refine anchors
- 8:   **for all** gaps  $(a', b')$  in  $C^{(\ell)}$  **do**
- 9:     Form query  $\mathbf{x}$  by inserting MASK between  $(a', b')$ .
- 10:     Admissible targets  $\mathcal{Y} = \text{BETWEEN}(a', b' \mid \mathbf{v})$ .
- 11:     Add  $(\mathbf{x}, \mathcal{Y}, \ell)$  to  $\mathcal{D}$ .
- 12:   **end for**
- 13:   Optionally mask additional random interior positions of  $C^{(\ell)}$  and generate samples analogously.
- 14: **end for**
- 15: **return**  $\mathcal{D}$

---

We adopt a hierarchical training curriculum that gradually refines supervision from coarse global signals to fine-grained local constraints. The purpose of this curriculum is to enable the model to learn the path context reconstruction. At the first level ( $k = 1$ ), only the endpoints are observed while the entire interior is masked,

$$\ell^1 = [v_1, \text{MASK}, v_T], \quad \mathcal{Y} = \{v_2, \dots, v_{T-1}\}, \quad (3)$$

$\ell^1$  is the first level of the curriculum  $\mathcal{Y}$  is the target set within the walk. This encourages the model to establish a global prior over feasible paths.

At higher levels ( $k > 1$ ), a subset of interior nodes  $a_1, \dots, a_m$  are revealed as anchors. The masked prediction is then restricted to the sub-path between consecutive anchors. For each adjacent anchor pair  $(a_j, a_{j+1})$  with  $j = 0, 1, \dots, m-1$ , we construct one masked sequence, with admissible target set  $\mathcal{Y}_j^k = \{v \in \mathbf{v} \mid a_j \prec v \prec a_{j+1}\}$ ,

$$\ell^k = [v_1, a_1, \dots, a_j, \text{MASK}, a_{j+1}, \dots, v_T], \quad (4)$$

where  $\prec$  denotes the ordering along the original walk. Hence, the level- $k$  training set is

$$\mathcal{L}^k = \{(\ell_j^k, \mathcal{Y}_j^k) \mid j = 0, \dots, k-1\}. \quad (5)$$

To learn from the path completions, we optimize a multi-target loss. For a masked query  $x$  with admissible targets set  $\mathcal{Y}_s$ , let  $z \in \mathbb{R}^{|\mathcal{V}|}$  be the logits and  $P = \text{softmax}(z)$ . The loss aggregates probability mass over all valid targets:

$$\ell_{\text{MT}}(s) = -\log \sum_{y \in \mathcal{Y}_s} P_y, \quad (6)$$

The final objective averages the multi-target loss across a mini-batch  $\mathcal{B}$ :

$$\mathcal{L}_{\text{prog}} = \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \alpha_{\ell(x)} \ell_{\text{MT}}(x), \quad (7)$$

where  $x$  indexes a query,  $\ell(x)$  is its curriculum level, and  $\alpha_{\ell}$  is a level-dependent weight increasing with  $\ell$ .

### 3.4 MODEL ARCHITECTURE

Our backbone is an encoder-only, pre-norm Transformer (Vaswani et al., 2017), aligned with BERT-style models (Devlin et al., 2019). Each block applies Layer Normalization, multi-head self-attention, and a feed-forward MLP with residual connections. Importantly, our self-attention is

bidirectional, allowing every token to access both its left and right context. The only masking applied is padding, which ensures that computation ignores padded positions while preserving global context.

Formally, let  $H \in \mathbb{R}^{T \times d}$  denote the hidden states at the input of a layer ( $T$ : sequence length,  $d$ : embedding dimension). Multi-head self-attention computes

$$\text{MHA}(H) = \left[ \text{softmax} \left( \frac{(HW_m^Q)(HW_m^K)^\top}{\sqrt{d_h}} + M_{\text{pad}} \right) (HW_m^V) \right]_{m=1}^M W^O, \quad (8)$$

where  $M$  is the number of heads,  $d_h = d/M$  is the head dimension,  $W_m^Q, W_m^K, W_m^V \in \mathbb{R}^{d \times d_h}$  are the learned projection matrices for queries, keys, and values, and  $W^O \in \mathbb{R}^{M d_h \times d}$  is the output projection. Here  $Q_m = HW_m^Q$ ,  $K_m = HW_m^K$ , and  $V_m = HW_m^V$ , i.e., the input hidden states  $H$  are linearly projected into query, key, and value spaces. The notation  $[\cdot]_{m=1}^M$  denotes concatenation over all heads. Finally,  $M_{\text{pad}} \in \mathbb{R}^{T \times T}$  is an additive mask with 0 for valid tokens and  $-\infty$  for padding. This bidirectional design enables the model to leverage full-sequence dependencies for predicting masked nodes.

### 3.5 DECODING FOR GRAPH OPTIMIZATION

At inference time, we obtain task-specific solutions by decoding from the learned structural prior under feasibility constraints. Formally, the final solution is generated as

$$\hat{Y} = \text{Decode}(\pi(\cdot | G), s), \quad (9)$$

where  $\hat{Y}$  denotes the final solution generated by our model and  $\pi(\cdot | G)$  provides the structural bias learned from the graph  $G$ ,  $s$  specifies task-dependent constraints (e.g., source/target in shortest path, required nodes in tours), and  $\text{Decode}(\cdot)$  denotes our decoding strategy that maps the prior distribution into a feasible solution.

## 4 EXPERIMENTS

We evaluate GFM on one synthetic graph and two real-world road networks across four problems: shortest path (SP) and three NP-hard tasks—Graphic Traveling Salesman Problem (Graphic-TSP), tour problem with the same origin and destination, and tour problem with different origins and destinations (Martin et al., 2022). The results highlight GFM’s cross-task capability and high-quality solving performance on graph optimization problems.

### 4.1 DATASETS AND PROBLEM CONSTRUCTION

**Datasets** To ground our evaluation on reproducible yet realistic scenarios, we adopt one controlled simulation graph and two real-world road networks, all are undirected connected graphs:

**(1) Simulation graph ( $N=20$ ).** A randomly generated graph with 20 nodes and 34 edges. The graph has a density of 0.1789 and an average degree of 3.40.

**(2) Chengdu–Longquanyi ( $N=132$ ).** A real road network from Longquanyi District, Chengdu, China, with 132 nodes and 222 edges. The graph has a density of 0.0257 and an average degree of 3.36. Original edge lengths are in meters (m) and converted to kilometers (km) for normalization.

**(3) Berkeley ( $N=893$ ).** A real road network from Berkeley, CA with 893 nodes and 1413 edges, density 0.0035, and average degree 3.16. The graph is obtained via OSMnx v2.0+ by clipping a 1.3 km radius around the landmark “Downtown Berkeley BART.”

**Problem Construction** Our goal is to achieve a city-scale GFM that is evaluated on real-world road graphs. Accordingly, we instantiate four distance-driven tasks on the graph  $G = (V, E, w)$ :

**(1) Shortest Path (SP).** Given  $s, t \in V$ , find a minimum-length  $s \rightarrow t$  path under edge lengths  $w$ .

**(2) Graphic Traveling Salesman Problem (Graphic-TSP).** Given a required node set  $R \subseteq V$ , find a minimum-length closed walk that visits all  $r \in R$  on the road graph. Unlike metric TSP on complete graphs, Graphic-TSP on general graphs naturally permits vertex/edge repetitions when necessary,

which matches urban networks with cul-de-sacs and limited connectivity, see approximation results and discussion in Sebö & Vygen (2012).

**(3) Tour Problem with Same Origin and Destination (TP-SOD)** Given a start/end depot  $o \in V$  and a required/attractive POI set  $R$ , produce a feasible tour starting and ending at  $o$  that visits all  $R$  while minimizing travel length. This is a closed special case of the Orienteering/Prize-Collecting family on road networks, widely studied in the orienteering literature and its variants (Gunawan et al., 2016). It falls under the broad class of tour problems (Martin et al., 2022), where the aim is to plan efficient sightseeing tours for visitors (Vansteenwegen et al., 2011).

**(4) Tour Problem with Different Origin and Destination (TP-DOD).** Given distinct  $o, d \in V$  and a required POI set  $R$ , find a minimum-length open  $o \rightarrow d$  walk that visits all  $R$ . Similar to TP-SOD, this also belongs to the family of tourist trip problems, modeling practical scenarios such as day trips between different locations (Vansteenwegen et al., 2011; Martin et al., 2022).

## 4.2 EXPERIMENTAL SETTINGS

**Hyperparameters** Our Graph Foundation Model (GFM) adopts Transformer backbones with task-specific settings. On the simulation graph ( $N=20$ ), we use a 6-layer, 6-head encoder ( $d=192$ , dropout 0.1, batch size 64, seq. length 10) trained for 15k steps with Adam ( $5 \times 10^{-4}$ ). For Chengdu ( $N=132$ ), we employ an 8-layer, 8-head encoder ( $d=256$ , batch size 32, seq. length 20) trained on 2.73M walks for 15k steps at  $3 \times 10^{-4}$ . For Berkeley ( $N=893$ ), we use a 6-layer, 6-head encoder ( $d=192$ , seq. length 200, batch size 64) for 150k steps at  $5 \times 10^{-4}$ . Parameter counts range 2.7M–6.7M. Baselines follow recommended settings: LKH3 with `MAX_TRIALS=10,000` and `RUNS=10`; OR-Tools with `PATH_CHEAPEST_ARC + GUIDED_LOCAL_SEARCH`; Gurobi with gaps 0.1–5% and limits 30–1800s; A\* with  $c_{\text{puct}} = 1.4$ ,  $\alpha = 1.8$ . Further details are in Appendix B and D.

**Runtime** All experiments were run on an Intel i5-12400F CPU, 32GB RAM, and RTX 4060 Ti (8GB) with CUDA 12.9. Classical solvers (Dijkstra, OR-Tools, LKH3, Gurobi) utilized official implementations, while LLM baselines (ChatGPT5, OR-LLM-Agent, Qwen3-235B) were accessed via APIs. GFM was trained and tested locally under the same environment, ensuring timing comparability across all methods. Detailed runtime protocols are in Appendix D.

**Baselines** We compare our approach against a broad set of baselines that cover exact solvers, classical heuristics, and modern learning-based approaches. For the shortest path (SP), we include the label-setting algorithm of Dijkstra (1959) as the ground-truth optimum, together with the A\* heuristic search and Google’s OR-Tools implementation. For tour-based problems, we utilize the state-of-the-art Lin–Kernighan–Helsgaun framework (LKH3) (Helsgaun, 2017), the Christofides approximation followed by 2OPT local refinement, and simple greedy strategies, such as Nearest Neighbor (NN). We also report Gurobi as a commercial MILP solver under fixed time limits, serving as a near-exact reference for NP-hard cases. On the learning side, we benchmark the Attention Model (AM) of Kool et al. (2018), an RL-trained neural combinatorial optimization architecture, under both greedy and sampling-based decoding, and the Instance-Conditioned Adaptation framework (ICAM) (Zhou et al., 2024), which adapts neural solvers to different instance scales via lightweight conditioning modules. We further evaluate OR-LLM-Agent (Zhang & Luo, 2025), which leverages structured prompting and tool use, and two large language models—ChatGPT5 and Qwen3-235B (Yang et al., 2025)—both in zero-shot settings with JSON-constrained outputs. Finally, we compare our proposed Graph Foundation Model (GFM) with this approach, which integrates curriculum-pretrained structural priors and task-specific decoding.

## 4.3 MAIN RESULTS

**Evaluation Metrics** We primarily report three metrics across all tasks: (**Obj**) the objective value, defined as the total path or tour length; (**Succ**) the success rate, i.e., the fraction of outputs that satisfy all task-specific feasibility constraints; and (**Time**) the average runtime per instance (s).

**Result Discussion** Table 1 reports results on simulation, Chengdu, and Berkeley networks. GFM achieves near-optimal objective values with almost perfect feasibility across all tasks.

Table 1: Performance across four routing problems on **Simulation** ( $N = 20$ ), **Chengdu** ( $N = 132$ ), and **Berkeley** ( $N = 893$ ). Entries report **objective (km)**, **success (%)**, and **time (s)**.

Method	Simulation ( $N = 20$ )			Chengdu ( $N = 132$ )			Berkeley ( $N = 893$ )				
	Obj (km)	Succ (%)	Time (s)	Obj (km)	Succ (%)	Time (s)	Obj (km)	Succ (%)	Time (s)		
SP	Dijkstra	6.09	100%	0.000	5.09	100%	0.000	2.23	100%	0.002	
	A*	6.09	100%	0.000	5.09	100%	0.000	2.23	100%	0.002	
	OR-Tools	6.09	100%	0.000	5.09	100%	0.000	2.23	100%	0.005	
	OR-LLM-Agent	15.22	30%	19.800	–	0%	45.100	–	0%	91.700	
	ChatGPT5	8.23	99%	3.394	8.85	62%	3.720	–	0%	–	
	Qwen3-235B	9.81	97%	1.89	11.98	45%	3.050	–	0%	3.820	
	<b>GFM</b>	<b>6.24</b>	<b>99%</b>	<b>0.046</b>	<b>5.19</b>	<b>100%</b>	<b>0.821</b>	<b>2.43</b>	<b>98%</b>	<b>23.109</b>	
Graphic-TSP	LKH3	39.48	100%	0.016	87.50	100%	2.020	84.75	100%	2328.000	
	Gurobi(30min)	39.48	100%	0.147	90.64	100%	46.423	120.00	100%	1800.000	
	OR-Tools	39.48	100%	15.000	88.12	100%	30.000	109.78	100%	690.000	
	Chr. + 2OPT	42.50	100%	0.007	97.78	100%	20.030	92.86	100%	606.502	
	Greedy	44.43	100%	0.001	111.51	100%	0.006	113.76	100%	0.213	
	OR-LLM-Agent	–	0%	246.830	–	0%	45.100	–	0%	16.530	
	AM (greedy)	100.91	100%	0.042	104.53	100%	0.020	–	–	–	
	AM (10)	82.16	100%	0.054	104.31	100%	0.017	–	–	–	
	AM (100)	74.83	100%	0.160	101.35	100%	0.017	–	–	–	
	AM (1000)	70.83	100%	1.366	100.36	100%	0.018	–	–	–	
	ICAM	41.14	100%	0.159	103.25	100%	1.982	122.51	100%	571.172	
	Qwen3-235B	103.95	41%	2.690	–	0%	22.410	–	0%	16.530	
	<b>GFM</b>	<b>39.48</b>	<b>100%</b>	<b>0.381</b>	<b>99.00</b>	<b>100%</b>	<b>6.405</b>	<b>100.87</b>	<b>100%</b>	<b>129.58</b>	
TP-SOD	LKH3	23.96	100%	0.003	22.55	100%	0.003	8.77	100%	0.042	
	Gurobi(30s)	23.96	100%	0.007	22.55	100%	0.021	8.77	100%	0.038	
	OR-Tools	23.96	100%	10.000	22.55	100%	10.000	8.77	100%	10.000	
	OR-LLM-Agent	20.10	76%	20.100	–	3%	72.700	–	0%	47.000	
	AM (greedy)	30.45	100%	0.033	27.78	100%	0.004	–	–	–	
	AM (10)	30.45	100%	0.039	25.46	100%	0.003	–	–	–	
	AM (100)	30.62	100%	0.040	25.20	100%	0.003	–	–	–	
	AM (1000)	30.62	100%	0.039	25.20	100%	0.003	–	–	–	
	Qwen3-235B	58.83	80.2%	2.656	–	0%	2.838	–	0%	9.830	
	<b>GFM</b>	<b>24.72</b>	<b>100%</b>	<b>0.03</b>	<b>20.31</b>	<b>100%</b>	<b>0.071</b>	<b>8.82</b>	<b>100%</b>	<b>0.077</b>	
	TP-DOD	LKH3	22.01	100%	0.002	19.99	100%	0.003	7.84	100%	0.061
		Gurobi(30s)	22.01	100%	0.004	19.93	100%	0.014	7.84	100%	0.064
		OR-Tools	22.01	100%	10.003	19.99	100%	10.006	7.99	100%	9.958
OR-LLM-Agent		21.40	40%	21.400	–	1%	68.300	–	0%	61.200	
AM (greedy)		31.24	100%	0.002	27.67	100%	0.004	–	–	–	
Qwen3-235B		38.04	41%	1.750	–	0%	8.305	–	0%	10.000	
<b>GFM</b>		<b>22.87</b>	<b>100%</b>	<b>0.037</b>	<b>20.31</b>	<b>100%</b>	<b>0.068</b>	<b>7.84</b>	<b>100%</b>	<b>0.104</b>	

Note. “–” has two meanings: when the **success rate** is 0%, the method produced outputs but none satisfied task constraints; when all three entries are “–”, the method could not generate solutions for that problem setting at all.

Compared to classical solvers, GFM provides competitive solutions with higher efficiency on large-scale problems. This efficiency advantage is most pronounced when solving Graphic-TSP problems. As the network size increases, the runtime of solver-based methods grows significantly, ranging from 606 to 2,328 seconds. In contrast, our GFM-based approach requires only 129 seconds on the same large-scale problems, while delivering solutions of comparable quality (100.87 km versus 84.75-120 km). This advantage stems from the fundamental mechanistic difference between the two approaches. Solver-based methods often rely on near-exhaustive search techniques to find high-quality solutions, which can be computationally intensive. GFM, however, generates solutions directly by leveraging a deep understanding of the graph structure. Our experiments validate the efficiency and effectiveness of the GFM approach.

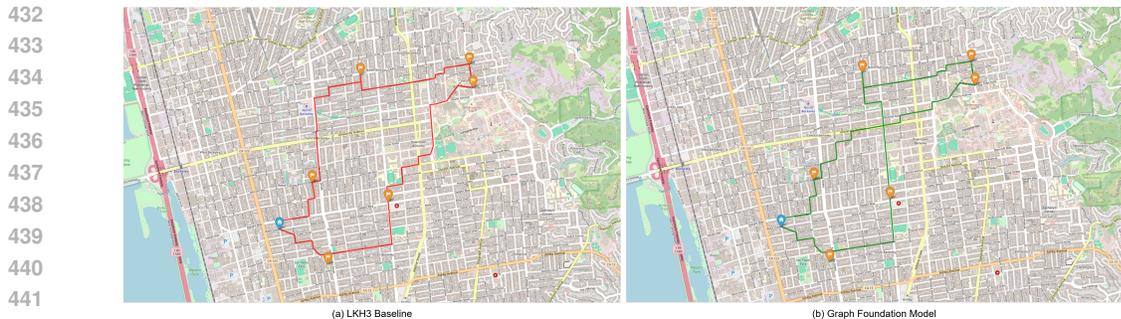


Figure 2: A comparative example of TP-SOD Solutions on the Berkeley road network

Note. GFM leverages learned structural priors to better capture the geometry and connectivity of real road networks, producing competitive solutions that demonstrate a deep understanding of graph structure.

Furthermore, we visualized the paths generated by LKH3 and GFM for the Tour-SOD problem on the Berkeley network, as shown in Figure 2. This visualization provides further evidence of the high quality of the solutions produced by GFM.

Compared with existing transformer-based methods, GFM can produce solutions with higher quality, e.g. shorter paths and tours. For example, when solving the Graphic-TSP problem on the simulation network, the path length produced by the AM method is approximately twice that of the GFM result. This is because Transformer-based methods are trained on task-specific synthetic graphs. The training setting significantly differs from real road networks that are sparse and topologically constrained, thus limiting their applicability to real mobility systems. In contrast, GFM excels at deeply exploring the network structure and fully perceiving the constraints of the real road network. This advantage enables it to provide higher-quality solutions for real-world graph-based problems.

Compared with current LLM-based baselines, GFM maintains solution validity and quality across different scales and optimization tasks. As network size increases from small to large scale, their success rate drops sharply (e.g., ChatGPT5 in the SP problem: 99% ( $N = 20$ ), 62% ( $N = 132$ ), and 0% ( $N = 893$ )). This infeasibility result highlights the inherent limitations of LLMs in directly or indirectly solving graph optimization problems. This discrepancy arises because natural language, which LLMs excel at processing, possesses a degree of inherent stochasticity. In contrast, graph optimization problems are governed by strict structural constraints and explicit objectives, a paradigm that falls outside the domain of what LLMs are adept at handling. Since GFM is trained to capture the network’s structure and connectivity, it consistently preserves solution quality as the graph size scales.

These results demonstrate the strength of our GFM on solution quality, efficiency, and stability, which validates the application potential of our GFM framework.

## 5 CONCLUSION

We presented the **Graph Foundation Model (GFM)**, which extends the pretrain–adapt paradigm of large language models to OR problems on graphs. By pretraining on structure-aware random walks with reconstructed trajectories training, GFM internalizes transferable structural priors that capture the geometry and topology of real road networks. This universal backbone enables lightweight solution generation across multiple NP-hard routing families, achieving competitive performance while maintaining efficiency. Our findings highlight the promise of foundation models for graphs: pretraining once on diverse networks can provide a reusable structural prior that generalizes broadly across optimization tasks, marking a step toward universal foundation-level solvers in operations research.

## 486 6 ETHICS STATEMENT

487

488 This work adheres to the ICLR Code of Ethics. In this study, no human subjects or animal ex-  
 489 perimentation was involved. All datasets used, including the dataset derived from OSMnx, were  
 490 collected in compliance with relevant usage guidelines and properly cited, ensuring no violation of  
 491 privacy. We have taken care to avoid any biases or discriminatory outcomes in our research process.  
 492 No personally identifiable information was used, and no experiments were conducted that could  
 493 raise privacy or security concerns. We are committed to maintaining transparency and integrity  
 494 throughout the research process.

495

## 496 7 REPRODUCIBILITY STATEMENT

497

498 We have made every effort to ensure that the results presented in this paper are reproducible. Upon  
 499 acceptance, we will release all code and datasets in a public repository, together with a full de-  
 500 scription of our contributions, experimental setup, model configurations, and hardware details. This  
 501 will facilitate replication and verification of our work. We believe these measures will enable other  
 502 researchers to reproduce our results and build upon our approach.

503

504

## 505 REFERENCES

- 506 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-  
 507 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and et al. Gpt-  
 508 4 technical report. 2023. URL [https://api.semanticscholar.org/CorpusID:  
 509 257532815](https://api.semanticscholar.org/CorpusID:257532815).
- 511 Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial  
 512 optimization with reinforcement learning. *ArXiv*, abs/1611.09940, 2016. URL [https://api.  
 513 semanticscholar.org/CorpusID:3649804](https://api.semanticscholar.org/CorpusID:3649804).
- 514 Xavier Bresson and Thomas Laurent. The transformer network for the traveling salesman prob-  
 515 lem. *ArXiv*, abs/2103.03012, 2021. URL [https://api.semanticscholar.org/  
 516 CorpusID:232110581](https://api.semanticscholar.org/CorpusID:232110581).
- 518 Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhari-  
 519 wal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and et al. Lan-  
 520 guage models are few-shot learners. *ArXiv*, abs/2005.14165, 2020. URL [https://api.  
 521 semanticscholar.org/CorpusID:218971783](https://api.semanticscholar.org/CorpusID:218971783).
- 522 Michel Deudon, Pierre Cournut, Alexandre Lacoste, Yossiri Adulyasak, and Louis-Martin  
 523 Rousseau. Learning heuristics for the tsp by policy gradient. In *Integration of AI and OR Tech-  
 524 niques in Constraint Programming*, 2018. URL [https://api.semanticscholar.org/  
 525 CorpusID:47017706](https://api.semanticscholar.org/CorpusID:47017706).
- 527 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep  
 528 bidirectional transformers for language understanding. In *North American Chapter of the Associ-  
 529 ation for Computational Linguistics*, 2019. URL [https://api.semanticscholar.org/  
 530 CorpusID:52967399](https://api.semanticscholar.org/CorpusID:52967399).
- 531 Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Math-  
 532 ematik*, 1:269–271, 1959. URL [https://api.semanticscholar.org/CorpusID:  
 533 123284777](https://api.semanticscholar.org/CorpusID:123284777).
- 534 Mohammed Elhenawy, Ahmad Abutahoun, Taqwa I. Alhadidi, Ahmed Jaber, Huthaifa I. Ashqar,  
 535 Shadi Jaradat, Ahmed Abdelhay, Sébastien Glaser, and Andry Rakotonirainy. Visual reasoning  
 536 and multi-agent approach in multimodal large language models (mllms): Solving tsp and mtsp  
 537 combinatorial challenges. *Mach. Learn. Knowl. Extr.*, 6:1894–1921, 2024. URL [https://  
 538 api.semanticscholar.org/CorpusID:270875399](https://api.semanticscholar.org/CorpusID:270875399).

- 540 Bishad Ghimire, Ausif Mahmood, and Khaled Elleithy. One-shot autoregressive generation of  
541 combinatorial optimization solutions based on the large language model architecture and learn-  
542 ing algorithms. *AI*, 2025. URL [https://api.semanticscholar.org/CorpusID:  
543 277385980](https://api.semanticscholar.org/CorpusID:277385980).
- 544 Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. *Proceedings  
545 of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*,  
546 2016. URL <https://api.semanticscholar.org/CorpusID:207238980>.
- 547 Aldy Gunawan, Hoong Chuin Lau, and Pieter Vansteenwegen. Orienteering problem: A survey  
548 of recent variants, solution approaches and applications. *Eur. J. Oper. Res.*, 255:315–332, 2016.  
549 URL <https://api.semanticscholar.org/CorpusID:33166987>.
- 550 William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large  
551 graphs. *ArXiv*, abs/1706.02216, 2017. URL [https://api.semanticscholar.org/  
552 CorpusID:4755450](https://api.semanticscholar.org/CorpusID:4755450).
- 553 Keld Helsgaun. An extension of the lin-kernighan-helsgaun tsp solver for constrained traveling  
554 salesman and vehicle routing problems: Technical report. In *Proceedings of the Conference*,  
555 2017. URL <https://api.semanticscholar.org/CorpusID:57634432>.
- 556 John J. Hopfield and David W. Tank. “neural” computation of decisions in optimization problems.  
557 *Biological Cybernetics*, 52:141–152, 1985. URL [https://api.semanticscholar.org/  
558 CorpusID:36483354](https://api.semanticscholar.org/CorpusID:36483354).
- 559 Chenyu Huang, Zhengyang Tang, Shixi Hu, Ruoqing Jiang, Xin Zheng, Dongdong Ge, Benyou  
560 Wang, and Zizhuo Wang. Orlm: A customizable framework in training large models for  
561 automated optimization modeling. *Operations Research*, 2024a. URL [https://api.  
562 semanticscholar.org/CorpusID:270067588](https://api.semanticscholar.org/CorpusID:270067588).
- 563 Zhehui Huang, Guangyao Shi, and Gaurav S. Sukhatme. From words to routes: Applying large  
564 language models to vehicle routing. *arXiv preprint arXiv:2403.10795*, 2024b.
- 565 Yan Jin, Yuandong Ding, Xuanhao Pan, Kun He, Li Zhao, Tao Qin, Lei Song, and Jiang Bian.  
566 Pointerformer: Deep reinforced multi-pointer transformer for the traveling salesman problem. In  
567 *AAAI Conference on Artificial Intelligence*, 2023. URL [https://api.semanticscholar.  
568 org/CorpusID:258212484](https://api.semanticscholar.org/CorpusID:258212484).
- 569 Chaitanya K. Joshi, Thomas Laurent, and Xavier Bresson. An efficient graph convolutional network  
570 technique for the travelling salesman problem. *ArXiv*, abs/1906.01227, 2019. URL [https://  
571 api.semanticscholar.org/CorpusID:174798181](https://api.semanticscholar.org/CorpusID:174798181).
- 572 Elias Boutros Khalil, Hanjun Dai, Yuyu Zhang, Bistra N. Dilkina, and Le Song. Learning com-  
573 binatorial optimization algorithms over graphs. *ArXiv*, abs/1704.01665, 2017. URL [https://  
574 api.semanticscholar.org/CorpusID:3486660](https://api.semanticscholar.org/CorpusID:3486660).
- 575 Thomas Kipf and Max Welling. Semi-supervised classification with graph convolutional net-  
576 works. *ArXiv*, abs/1609.02907, 2016. URL [https://api.semanticscholar.org/  
577 CorpusID:3144218](https://api.semanticscholar.org/CorpusID:3144218).
- 578 Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing prob-  
579 lems! In *International Conference on Learning Representations*, 2018. URL [https://api.  
580 semanticscholar.org/CorpusID:59608816](https://api.semanticscholar.org/CorpusID:59608816).
- 581 Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Seungjai Min, and Youngjune  
582 Gwon. Pomo: Policy optimization with multiple optima for reinforcement learning. *ArXiv*,  
583 abs/2010.16011, 2020. URL [https://api.semanticscholar.org/CorpusID:  
584 226222332](https://api.semanticscholar.org/CorpusID:226222332).
- 585 Hao Liu, Jiarui Feng, Lecheng Kong, Ningyue Liang, Dacheng Tao, Yixin Chen, and Muhan Zhang.  
586 One for all: Towards training one graph model for all classification tasks. *ArXiv*, abs/2310.00149,  
587 2023a. URL <https://api.semanticscholar.org/CorpusID:265871676>.

- 594 Shengcai Liu, Caishun Chen, Xinghua Qu, Ke Tang, and Yew Soon Ong. Large language models  
595 as evolutionary optimizers. *2024 IEEE Congress on Evolutionary Computation (CEC)*, pp. 1–8,  
596 2023b. URL <https://api.semanticscholar.org/CorpusID:264829031>.
- 597 Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial optimization  
598 with heavy decoder: Toward large scale generalization. *ArXiv*, abs/2310.07985, 2023. URL  
599 <https://api.semanticscholar.org/CorpusID:263909317>.
- 600 Sébastien Martin, Youcef Magnouche, Corentin Juvigny, and Jérémie Leguay. Constrained shortest  
601 path tour problem: Branch-and-price algorithm. *Comput. Oper. Res.*, 144:105819, 2022. URL  
602 <https://api.semanticscholar.org/CorpusID:248028823>.
- 603 Christos H. Papadimitriou and Kenneth Steiglitz. Combinatorial optimization: Algorithms and com-  
604 plexity. In *Proceedings of the 20th Annual Symposium on Foundations of Computer Science*,  
605 1981. URL <https://api.semanticscholar.org/CorpusID:265900001>.
- 606 András Sebő and Jens Vygen. Shorter tours by nicer ears:  $7/5$ -approximation for the graph-tsp,  $3/2$   
607 for the path version, and  $4/3$  for two-edge-connected subgraphs. *Combinatorica*, pp. 1–34, 2012.  
608 URL <https://api.semanticscholar.org/CorpusID:15165701>.
- 609 Zhiqing Sun and Yiming Yang. Difusco: Graph-based diffusion solvers for combinatorial opti-  
610 mization. *ArXiv*, abs/2302.08224, 2023. URL <https://api.semanticscholar.org/CorpusID:256900800>.
- 611 Jiabin Tang, Yuhao Yang, Wei Wei, Lei Shi, Lixin Su, Suqi Cheng, Dawei Yin, and Chao Huang.  
612 Graphgpt: Graph instruction tuning for large language models. *Proceedings of the 47th Inter-  
613 national ACM SIGIR Conference on Research and Development in Information Retrieval*, 2023.  
614 URL <https://api.semanticscholar.org/CorpusID:264405943>.
- 615 Pieter Vansteenwegen, Wouter Souffriau, and Dirk Van Oudheusden. The orienteering problem: A  
616 survey. *Eur. J. Oper. Res.*, 209:1–10, 2011. URL <https://api.semanticscholar.org/CorpusID:10893453>.
- 617 Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,  
618 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Neural Information Processing  
619 Systems*, 2017. URL <https://api.semanticscholar.org/CorpusID:13756489>.
- 620 Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio’, and Yoshua  
621 Bengio. Graph attention networks. *ArXiv*, abs/1710.10903, 2017. URL <https://api.semanticscholar.org/CorpusID:3292002>.
- 622 Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. *ArXiv*, abs/1506.03134,  
623 2015. URL <https://api.semanticscholar.org/CorpusID:5692837>.
- 624 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang  
625 Gao, Chengen Huang, Chenxu Lv, and et al. Qwen3 technical report. *ArXiv*, abs/2505.09388,  
626 2025. URL <https://api.semanticscholar.org/CorpusID:278602855>.
- 627 Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun  
628 Chen. Large language models as optimizers. *ArXiv*, abs/2309.03409, 2023a. URL <https://api.semanticscholar.org/CorpusID:261582296>.
- 629 Hua Yang, Minghao Zhao, Lei Yuan, Yang Yu, Zhenhua Li, and Ming Gu. Memory-efficient  
630 transformer-based network model for traveling salesman problem. *Neural networks : the offi-  
631 cial journal of the International Neural Network Society*, 161:589–597, 2023b. URL <https://api.semanticscholar.org/CorpusID:256975543>.
- 632 Haoran Ye, Jiarui Wang, Zhiguang Cao, and Guojie Song. Reevo: Large language models  
633 as hyper-heuristics with reflective evolution. *ArXiv*, abs/2402.01145, 2024. URL <https://api.semanticscholar.org/CorpusID:267406792>.
- 634 Bowen Zhang and Pengcheng Luo. Or-llm-agent: Automating modeling and solving of operations  
635 research optimization problem with reasoning large language model. *ArXiv*, abs/2503.10009,  
636 2025. URL <https://api.semanticscholar.org/CorpusID:276960951>.

648 Chun-Sheng Zhao and Li-Pei Wong. A transformer-based structure-aware model for tack-  
649 ling the traveling salesman problem. *PLOS One*, 20, 2025. URL [https://api.  
650 semanticscholar.org/CorpusID:277622775](https://api.semanticscholar.org/CorpusID:277622775).

651 Changliang Zhou, Xi Lin, Zhenkun Wang, Xialiang Tong, Mingxuan Yuan, and Qingfu Zhang.  
652 Instance-conditioned adaptation for large-scale generalization of neural routing solver. 2024.  
653 URL <https://api.semanticscholar.org/CorpusID:278768989>.

## 656 A DATA PROCESSING AND VISUALIZATION

### 657 A.1 ROAD-GRAPH CONSTRUCTION

658 **Datasets.** To ground our evaluation on reproducible yet realistic scenarios, we adopt one con-  
659 trolled simulation graph and two real-world road networks.

660 **Simulation graph ( $N=20$ ).** To emulate traffic properties while ensuring reproducibility, we con-  
661 struct a fixed sparse backbone of 20 nodes with several cross links. Edge weights are sampled once  
662 with a fixed random seed ( $s=42$ ): “arterial” roads from a uniform distribution  $U(1.0, 3.0)$  to model  
663 lower travel costs, and cross connections from  $U(3.5, 7.0)$  to represent less preferred detours. The  
664 topology is fixed, with only minor coordinate jitter for visualization, making the instance traffic-  
665 informed and exactly reproducible.

666 **Chengdu–Longquanyi ( $N=132$ ).** We extract the road network of Longquanyi District, Chengdu,  
667 as a self-collected dataset. The graph is preprocessed into an undirected primal form, and retains  
668 geographic coordinates for each node.

669 **Berkeley ( $N=893$ ).** The Berkeley network is obtained from OpenStreetMap using `OSMnx v2.0+`  
670 with `network_type=drive`, which filters for drivable roads. We clip a radius of 1300 m centered  
671 on the landmark “Downtown Berkeley BART,” yielding a sparse, connected graph with 893 nodes.  
672 Node IDs are remapped to  $0, \dots, N-1$ , and edge weights are computed as geometric distances  
673 (meters divided by 1000). All parameters (OSMnx version, query strings, radius, preprocessing)  
674 are fixed in our scripts to guarantee reproducibility. Raw OSM MultiDiGraphs are simplified into  
675 undirected simple graphs: we remove self-loops and duplicate edges, retain the largest connected  
676 component, and remap node IDs to  $0..(N-1)$  for efficient embedding. Edge weights are set to  
677 geometric lengths (meters) converted to kilometers, i.e., `length.km = length.m/1000`. Each  
678 node stores its GPS coordinates (lon, lat) for visualization and for computing path lengths consistent  
679 with our evaluation objective (km).

## 682 B BASELINE ADAPTATION ON SPARSE ROAD NETWORKS

683 Most classical solvers and some baselines (e.g., LKH3 (Helsgaun, 2017), OR-Tools, Gurobi,  
684 ICAM Zhou et al. (2024)) assume a complete graph input, where every pair of nodes has a di-  
685 rect edge. However, real road networks are sparse: many node pairs are not directly connected.  
686 To enable fair evaluation, we adopt two adaptation strategies: metric closure and virtual coordinate  
687 embedding.

688 **Metric closure.** Let  $G = (V, E, w)$  be the original weighted road network. For any pair  $u, v \in V$ ,  
689 define

$$690 d(u, v) = \min_{p \in \mathcal{P}(u, v)} \sum_{(i, j) \in p} w(i, j),$$

691 where  $\mathcal{P}(u, v)$  is the set of all paths between  $u$  and  $v$ . The metric closure is the complete graph  
692  $G^* = (V, E^*, d)$  with edge weights equal to shortest-path distances in  $G$ .

693 **Virtual edge construction.** For path problems with distinct start and end nodes, e.g., TP-DOD,  
694 solvers like LKH3 require a closed TSP formulation. We therefore introduce a virtual edge from the  
695 terminal node back to the start node with weight zero, transforming the open-path problem into a  
696 tour. After solving, this artificial edge is removed and the remaining sequence is expanded back into  
697 a feasible path in  $G$ .

702  
703  
704  
705  
706  
707  
708  
709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755

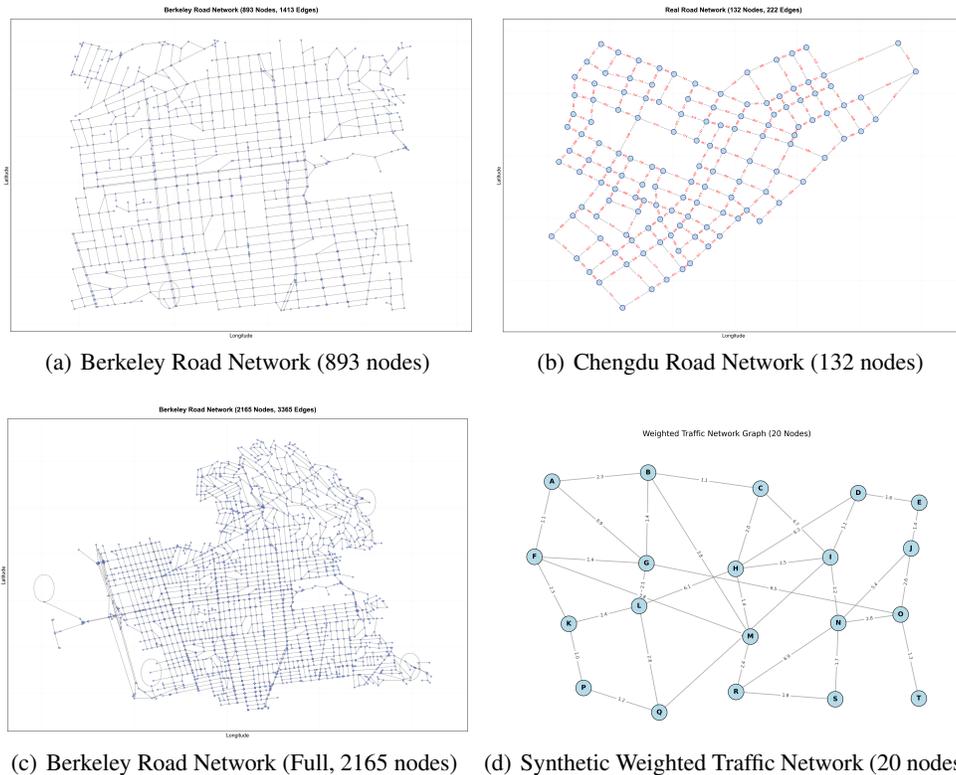


Figure 3: Examples of road network datasets used in our experiments. Real-world road networks are acquired from OpenStreetMap (**Berkeley**), while a synthetic traffic network provides controlled small-scale settings.

**Application.** For both strategies, the adapted representation (closure graph  $G^*$  or virtual edge construction) is fed to the baseline solver (e.g., LKH3, OR-Tools, AM, ICAM) to generate a high-level tour  $\pi^*$ . Each edge  $(u, v) \in \pi^*$  is then expanded back into its shortest path in the original graph  $G$ , guaranteeing feasibility while preserving the solver’s optimization logic.

## C RELATED WORKS (EXTENDED)

**Graph Neural Networks** The application of neural networks to combinatorial optimization can be traced back to the pioneering work of Hopfield & Tank (1985), who first employed Hopfield networks to solve the traveling salesman problem, establishing the foundation for using neural network dynamics in NP-hard optimization problems. Early graph neural network variants, including GCN Kipf & Welling (2016), GAT Velickovic et al. (2017), and GraphSAGE Hamilton et al. (2017), demonstrated significant success in learning relational data and solving graph-based problems.

The paradigm shift toward sequence-to-sequence models in combinatorial optimization began with Pointer Networks Vinyals et al. (2015), which utilized RNNs to encode city coordinates and employed attention mechanisms during decoding to directly generate city visiting sequences. This work established that end-to-end neural networks could learn to generate near-optimal TSP solutions without explicit algorithmic guidance. Building upon this foundation, Bello et al. (2016) introduced reinforcement learning with actor-critic methods to train pointer networks using negative tour distances as rewards, demonstrating that solution quality could be improved without supervised optimal solutions.

The integration of attention mechanisms further advanced the field when Kool et al. (2018) combined attention-based city encoding with pointer network selection mechanisms, training the resulting model through reinforcement learning. This was extended by Kwon et al. (2020), who proposed

756 POMO (Policy Optimization with Multiple Optima), employing multi-start parallel decoding to gener-  
757 ate multiple solutions and improve optimization quality. Concurrently, Joshi et al. (2019) applied  
758 graph neural networks directly to TSP by using graph convolutions to encode complete graph struc-  
759 tures combined with greedy selection strategies.

760  
761 **Transformer-Based Approaches** Following the transformer revolution Vaswani et al. (2017),  
762 several works have adapted transformer architectures for combinatorial optimization. Bresson &  
763 Laurent (2021) proposed using standard transformer encoders without pointer mechanisms, encod-  
764 ing city coordinate sequences as input tokens and outputting city-sorted logits representing visit  
765 orders through reinforcement learning optimization. Their implementation achieved remarkable  
766 performance with optimality gaps of 0.004% on 50-city instances and approximately 0.37% on 100-  
767 city problems after adequate training and beam search.

768 However, the quadratic complexity of standard attention mechanisms poses scalability challenges  
769 for large-scale problems. To address this limitation, recent works have focused on efficiency im-  
770 provements. Yang et al. (2023b) introduced TSPformer with Sampled Scaled Dot-Product Attention  
771 (SSDPA), reducing the standard attention complexity from  $O(n^2)$  to  $O(n \log n)$  while maintaining  
772 competitive performance on 1000-node TSP instances with minimal accuracy loss. Similarly, Jin  
773 et al. (2023) proposed Pointerformer, incorporating multi-pointer decoders and data augmentation  
774 techniques leveraging the rotation and reflection invariant properties of TSP solutions.

775 Alternative approaches have explored non-autoregressive formulations. Khalil et al. (2017) trained  
776 graph neural networks in a supervised manner to directly output tours as adjacency matrices, con-  
777 verting results to feasible solutions via beam search. However, this non-autoregressive approach  
778 achieved only 2.7% optimality gap for  $n=20$ , underperforming compared to autoregressive meth-  
779 ods. Deudon et al. (2018) developed a transformer-based model outputting fractional solutions to  
780 multiple TSP variants, treating results as linear relaxation solutions and employing beam search for  
781 integer feasibility.

782 Most recently, Zhao & Wong (2025) designed a structure-aware transformer incorporating closeness  
783 centrality rather than degree centrality for node representation, integrating spatial distance encoding  
784 to capture inter-node relationships. Their improved pointer mechanism incorporates information  
785 from all visited nodes into the context state, capturing the dynamic evolution of path construction and  
786 outperforming classical heuristics, benchmark methods, and previous learning-based approaches  
787 across diverse test scenarios. Also, Zhou et al. (2024) incorporate graph centrality measures and  
788 spatial encodings to capture topological relationships.

789  
790 **Large Language Models for Combinatorial Optimization** The emergence of large language  
791 models has opened new avenues for approaching combinatorial optimization problems. Recent  
792 research in this domain can be categorized into two primary paradigms: direct problem solving  
793 through language models and LLM-assisted optimization frameworks.

794 In the direct solving paradigm, several works have explored applying LLMs to graph problems  
795 through various architectural adaptations. Tang et al. (2023) developed GraphGPT, which embeds  
796 graph structures through graph encoders and aligns them with textual information, enabling simul-  
797 taneous utilization of semantic and structural knowledge in graph tasks. Liu et al. (2023a) proposed  
798 OFA, a unified framework that integrates graph tasks across different domains through Text Attribute  
799 Graphs and Graph Prompt Paradigms, achieving cross-task generalization capabilities.

800 For TSP-specific applications, Yang et al. (2023a) introduced OPRO (Optimization through  
801 Prompts), which places LLMs in an iterative improvement loop where existing candidate solutions  
802 and their costs enable the generation of better solutions. While not guaranteeing optimality, this ap-  
803 proach demonstrates the potential of language generation with feedback for gradual route improve-  
804 ment. Liu et al. (2023b) proposed LLM-driven evolutionary algorithms (LMEA), where LLMs  
805 serve as variation operators within evolutionary frameworks, selecting parent solutions and per-  
806 forming crossover/mutation operations through textual outputs. Although limited to small instances  
807 (up to 20 nodes), LMEA achieved competitive performance with traditional TSP heuristics.

808 Recent advances have explored more sophisticated approaches. Ghimire et al. (2025) developed a  
809 one-shot autoregressive TSP solver that treats tours as sequences of city tokens, training the model to  
predict tours autoregressively and applying Direct Preference Optimization for quality improvement.

This approach achieved tours within a few percent of optimal length on instances up to 100 nodes. Elhenawy et al. (2024) investigated visual reasoning with multimodal LLMs, where TSP instances are presented as images and the LLM reasons about efficient tours through visual intuition, achieving competitive performance with OR-Tools on smaller instances.

The second paradigm focuses on LLM-assisted optimization frameworks that leverage language models for problem formulation and modeling. Zhang & Luo (2025) developed OR-LLM-Agent, an agentic framework that models and solves operations research problems through structured prompting and tool use. Huang et al. (2024a) introduced ORLM, which trains specialized LLMs for optimization modeling, enabling the translation of natural language problem descriptions into formal models or executable code for traditional solvers.

Huang et al. (2024b) examined LLMs for vehicle routing problems described in natural language, demonstrating that GPT-4 can serve as a coding assistant for combinatorial solvers with self-debugging and refinement capabilities. Ye et al. (2024) proposed ReEvo (Reflective Evolution), where LLMs generate entire heuristic algorithms rather than individual solutions, employing evolutionary strategies with natural language feedback to evolve competitive algorithms across multiple NP-hard problems.

Despite these advances, LLM-based approaches face fundamental limitations including scalability constraints, the mismatch between natural language representations and graph-theoretic constraints, and the inability to capture geometric and topological properties essential for real-world networks. These challenges motivate the need for more principled approaches that bridge language modeling paradigms with graph optimization requirements.

## D IMPLEMENTATION DETAILS

**Hyperparameters** Our Graph Foundation Model (GFM) adopts Transformer backbones with task-specific configurations across three datasets. For the simulation graph ( $N=20$ ), we use a 6-layer, 6-head encoder with embedding dimension  $d=192$ , dropout 0.1, batch size 64, and sequence length 10, trained for 15k steps with Adam (learning rate  $5 \times 10^{-4}$ ). The Chengdu road graph ( $N=132$ ) employs a larger 8-layer, 8-head encoder with  $d=256$ , dropout 0.1, batch size 32, and sequence length 15, trained on 2.73M random walks for 15k steps at  $3 \times 10^{-4}$  learning rate. The Berkeley road graph ( $N=893$ ) uses a 6-layer, 6-head encoder with  $d=192$ , extended sequence length 200 to match longer walks, batch size 64, and a total of 150k training steps at  $5 \times 10^{-4}$ . Parameter counts range from 2.7M–6.7M, with Transformer blocks contributing over 95% of weights. For classical solvers and baselines, necessary adaptations such as metric closure were applied to ensure applicability on sparse road graphs; Adaptation details are provided in Appendix B. Baselines follow their recommended hyperparameter settings to ensure comparability. For LKH3 we set `MAX_TRIALS=10,000`, `RUNS=10`, and fixed random seeds, consistent with Helsgaun’s implementation (Helsgaun, 2017). OR-Tools solvers adopt `PATH_CHEAPEST_ARC` as first solution strategy and `GUIDED_LOCAL_SEARCH` for improvement, with a 300s cap for large graphs. Gurobi is tuned with gap tolerances from 0.1%–5%, heuristic strength 0.2–0.5, and runtime limits of 30–1800s depending on problem size. A\* search uses  $c_{\text{puct}} = 1.4$ ,  $\alpha = 1.8$ , and pruning optimizations (lower-bound cuts, transposition tables, progressive widening). Classical heuristics such as Nearest Neighbor and Greedy follow standard greedy insertion without additional tuning, while attention-based neural baselines adopt the publicly released settings of Kool et al. (Kool et al., 2018). These consistent hyperparameter choices ensure fairness across solvers and facilitate reproducibility.

**Runtime** All experiments were executed on a single workstation equipped with an Intel Core i5-12400F CPU (6 physical cores, 12 threads), 32GB DDR4 RAM, and an NVIDIA RTX 4060 Ti GPU with 8GB VRAM, running CUDA 12.9 and driver version 576.28. GPU utilization during training averaged below 20%, with peak memory footprint under 2GB, ensuring that models fit comfortably within consumer-grade hardware. To ensure fairness, each baseline solver was executed under identical hardware constraints and with fixed random seeds. Classical methods (Dijkstra, OR-Tools, LKH3, Gurobi) were executed using their official high-performance implementations with optimized C/C++ backends. LLM-based methods were evaluated through standardized APIs: ChatGPT5 and OR-LLM-Agent via the OpenAI API, and Qwen3-235B via the Together AI API.

864 GFM training and inference were conducted locally under the same environment. All timing results  
865 are therefore directly comparable across methods under uniform runtime protocols.  
866

## 867 E LLM USAGE 868

869 Large Language Models (LLMs) were used to aid in the writing and polishing of the manuscript.  
870 Specifically, we used an LLM to assist in refining the language, improving readability, and ensuring  
871 clarity in various sections of the paper. The model helped with tasks such as sentence rephrasing,  
872 grammar checking, and enhancing the overall flow of the text. Separately, we also employed some  
873 LLMs purely as baselines in our experiments, serving only as comparative references against our  
874 proposed method.  
875

876 It is important to note that LLMs were not involved in the ideation, research methodology, or ex-  
877 perimental design. All research concepts, ideas, and analyses were developed and conducted by the  
878 authors. The authors take full responsibility for the content of the manuscript, including any text  
879 generated or polished by the LLM. We have ensured that the LLM-assisted text adheres to ethical  
880 guidelines and does not contribute to plagiarism or scientific misconduct.  
881

882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917