

An edge-aware graph autoencoder trained on scale-imbalanced data for traveling salesman problems

Shiqing Liu^a, Xueming Yan^b, Yaochu Jin^{a,c,*}

^a Nature Inspired Computing and Engineering, Faculty of Technology, Bielefeld University, Bielefeld, 33619, Germany

^b School of Information Science and Technology, Guangdong University of Foreign Studies, Guangzhou, 510006, China

^c School of Engineering, Westlake University, Hangzhou, 310030, China

ARTICLE INFO

Keywords:

Traveling salesman problem
Graph neural network
Link prediction
Neural combinatorial optimization

ABSTRACT

In recent years, there has been a notable surge in research on machine learning techniques for combinatorial optimization. It has been shown that learning-based methods outperform traditional heuristics and mathematical solvers on the Traveling Salesman Problem (TSP) in terms of both performance and computational efficiency. However, most learning-based TSP solvers are primarily designed for fixed-scale TSP instances, and also require a large number of training samples to achieve optimal performance. To fill this gap, this work proposes a data-driven graph representation learning method for solving TSPs with various numbers of cities. Specifically, we formulate the TSP as a link prediction task and propose an edge-aware graph autoencoder (EdgeGAE) model that can solve TSPs by learning from various-scale samples with an imbalanced distribution. A residual gated encoder is trained to learn latent edge embeddings, followed by an edge-centered decoder to output link predictions in an end-to-end manner. Furthermore, we introduce an active sampling strategy into the training process to improve the model's generalization capability in large-scale scenarios. To investigate the model's practical applicability, we generate a scale-imbalanced dataset comprising 50,000 TSP instances ranging from 50 to 500 cities. The experimental results demonstrate that the proposed edge-aware graph autoencoder model achieves a highly competitive performance among state-of-the-art graph learning-based approaches in solving TSPs with various scales, implying its remarkable potential in dealing with practical optimization challenges.

1. Introduction

Combinatorial optimization problems (COPs) exist in many practical applications of operational research (OR). Most COPs are NP-hard and difficult to solve within a polynomial time [1]. As an important field of research on COPs, routing problems arise in various application scenarios such as logistics, transportation, urban planning, and supply chain management. Generally, routing problems involve planning route scheduling solutions for a group of vehicles under constrained conditions to make deliveries or provide services to a set of locations efficiently. The target is to minimize the total traveling distance or overhead with all constraints being satisfied [2]. The traveling salesman problem (TSP) is one of the most intensively studied routing problems, which aims to find the shortest path to traverse a set of given locations once and return to the initial location.

As an NP-hard problem, it is non-trivial to find an efficient polynomial-time solution via an enumeration search for TSPs with even only hundreds of cities. Traditional methods for solving TSPs and other

optimization problems can be generally categorized into mathematical solvers [3–5] and hand-crafted heuristics [6–9]. A representative of exact solvers is Concorde [3], which incorporates both cutting-plane and branch-and-bound approaches for solving TSPs. For small and medium-sized problems, exact TSP solvers can find optimal solutions with guarantees. However, for large-scale instances, it will be intractable for exact solvers to get an exact solution within a reasonable amount of time due to the exponential growth of the computational complexity. To alleviate this limitation, evolutionary algorithms (EAs) are adopted to solve TSPs and proven to be well-suited for exploring a large search space. However, population-based EAs often require a large number of function evaluations, leading to huge computational costs, especially when the evaluations are time-consuming. Additionally, EAs are unable to learn from existing historical data, and even minor modifications to the problem require restarting the search from scratch.

An alternative, known as neural combinatorial optimization (NCO) [1,10], has been explored in the machine learning community from

* Corresponding author at: School of Engineering, Westlake University, Hangzhou, 310030, China.

E-mail address: jinyaochu@westlake.edu.cn (Y. Jin).

the perspective of data-driven optimization [1,11,12]. Instead of hand-crafted heuristics, NCO trains neural networks in an end-to-end manner to solve any combinatorial problems by learning directly from existing instances. Early work on learning-based approaches for routing problems can be traced back to Hopfield networks [13] and template models [14]. With the compelling performance of sequence-to-sequence models [15,16], Vinyals et al. proposed a Pointer Network [17] with an attention mechanism [18] to output TSP tours in a sequential manner. In a similar work proposed by Bello et al. [10], the Pointer Network is trained by an Actor-Critic reinforcement learning strategy instead of supervised learning [17].

With the development of graph representation learning [19–21], NP-hard routing problems are reconsidered as sequential decision tasks on graphs [22]. Consequently, there has been a recent surge of work on training graph neural networks (GNNs) [23–25] to solve COPs such as routing problems [22,26,27], facility location problems [28] and graph coloring problems [29]. Compared to OR solvers, GNN-based methods have achieved competitive performance on small problem instances [22,27,30,31]. Once trained, a GNN-based solver can generate approximate optimal solutions in a very short period of inference time, making it favorable for real-time decision-making. End-to-end learning for TSPs can be categorized into autoregressive and non-autoregressive approaches, depending on how solutions are generated. In autoregressive methods, a TSP tour is constructed regressively by identifying one city at a step [10,30,32,33]. By contrast, non-autoregressive approaches make predictions for arbitrarily large graphs in a one-shot fashion [22,26,27,34]. In an early work [26], Nowak et al. trained GNNs to output the prediction as an adjacency matrix of the input graph, which is further converted into feasible solutions by beam search. Joshi et al. [22] built on ideas from [26] and introduced the residual gated graph ConvNet [35] for solving TSP. By extending the work in [22], Fu et al. [36] replaced the beam search with reinforcement learning techniques. The output heatmaps are fed into the Monte Carlo tree search for TSP tours.

However, existing work on data-driven NCO approaches for TSPs suffers from distinct drawbacks. First, most learning-based solvers are only trained on fixed-sized TSP instances [17,22,32,33], which exhibit poor generalization ability on test cases with a number of cities different from that of the training data [27]. Secondly, training such a solver often requires millions of TSP instances with optimal solutions to achieve a competitive performance, assuming that the model has sufficient labeled data as its training dataset. However, it is non-trivial to satisfy this assumption in practice since NP-hard problems usually have a limited number of imbalanced optimized instances, which contain more small-scale cases and fewer large-scale cases.

In this paper, we propose an edge-aware graph autoencoder (EdgeGAE) framework for solving TSPs by learning from imbalanced data. In contrast to the traditional TSP solvers that take TSP as a route planning task and decide the next city to visit step by step, this work reformulates TSP as a link prediction task on the graph. Each city is treated as a node on the graph, and the optimal route can be represented by the connections between nodes. The encoder–decoder structure is widely adopted for traditional link prediction tasks in graph representation learning. Nevertheless, the conventional graph autoencoder (GAE) fails to solve TSP tasks, as its encoding–decoding process only considers node features, overlooking the importance of edge information for TSP. Considering that the quality of solutions to routing problems relies highly on the edge features (i.e., Euclidean distances), we modify the encoder with residual gated graph convolutions for the message-passing process to learn explicit edge embeddings in the latent space. To alleviate the limitation of inner products where only nodes are involved, we design an edge-centered decoder model that incorporates both the edge and corresponding node embeddings for the link prediction task. We adopt an active sampling (AS) strategy during the training process by incorporating *oversampling* and *undersampling* to deal with the imbalanced distribution of training data. The proposed framework

is evaluated on TSP instances with different sizes up to hundreds of cities and compared with state-of-the-art learning-based methods in terms of the optimal gap and the F1 score. Distinct from the traditional learning-based TSP solvers, the main contributions of this work are summarized as follows:

1. We first transform the TSP from a traditional path planning problem to a link prediction task in the field of graph representation learning. Drawing on this, we redesign the structure of the vanilla graph autoencoder and propose an edge-aware graph autoencoder model to accommodate the TSP characteristics.
2. We tackle the challenge of extremely imbalanced data distribution by introducing a random active sampling strategy into the end-to-end training process. The incorporation of *oversampling* and *undersampling* enables our model to retain strong generalization performance with only a few large-scale training samples.
3. To validate the model performance in scale-imbalanced optimization scenarios, we construct a benchmark dataset comprising 50,000 TSP instances with varying scales, ranging from 50 to 500 cities. The data follows an extremely imbalanced distribution where the number of instances is inversely proportional to the number of cities. Comprehensive experiments based on the benchmark demonstrate that the proposed EdgeGAE outperforms other learning-based methods for solving TSPs.

The remainder of this article is organized as follows. Section 2 gives a literature review of existing work on learning-based NCO methods for TSPs and their generalization ability. Section 3 provides the problem formulation and elaborates the proposed methodology in detail. The experimental results and analyses of these results are presented in Section 4. Section 5 concludes the paper and outlines future work.

2. Related work

Several attempts have been made to improve the generalization ability of machine learning methods for routing problems such as TSPs [36–39]. An intuitive approach is to decompose a large-scale problem into a group of sub-problems and solve them separately. Fu et al. [36] proposed to train a graph convolutional-based model on small-scale instances and generalize the pre-trained model to instances whose size is larger than that of the training data by using graph sampling and heatmap merging techniques. Specifically, a large number of instances with a fixed small size are first randomly generated and optimized by Concorde. Then a graph convolutional network [35] is trained on the generated data in a supervised manner, following the same model and framework in [22]. The trained model outputs a heatmap for the input graph as the probability prediction of each edge appearing in the optimal solution. When testing on TSP instances with a larger scale than those seen in the training dataset, the input graph is first decomposed into smaller subgraphs which consist of exactly the same number of vertices as the training data. The pre-trained model is adopted to output heatmaps for each subgraph separately, and then these heatmaps are merged into a complete heatmap as the edge prediction for the original input instance. Finally, the Monte Carlo tree search [40] strategy is adopted to convert the heatmaps into feasible solutions to TSPs. Similarly, Luo et al. [38] proposed a Light Encoder and Heavy Decoder (LEHD) model. The LEHD model adaptively captures the relationships among nodes of varying sizes dynamically, enhancing its ability to generalize across problems of diverse scales.

While such decomposition-based approaches are simple to implement, they also come with several limitations. The pre-trained model can only be trained on graphs of a fixed size with a predefined number of vertices (m), which may introduce an additional process to identify the optimal hyperparameter m . When tackling TSP instances of

an arbitrarily large size, it is required to decompose a large graph into subgraphs that have exactly the same number of m vertices as the training data. As a result, it can easily lead to overlaps between different subgraphs, posing difficulty in solving the original problem. Additionally, the complete heatmap is created by combining a number of sub-heatmaps that only consider the corresponding local information, overlooking the global knowledge of the original graph. Finally, it should be noted that the decomposition-based approach requires a huge amount of training data, for example, 990,000 TSP instances are generated for the pre-trained model in [36]. An alternative is to transfer the knowledge learned from small-scale instances into the optimization process of large-scale instances. For example, Zhang et al. [41] introduced progressive distillation into the training of NCO models by adopting curriculum learning to train TSP samples in an increasing order of the problem sizes.

Other work combines deep learning methods with traditional heuristics to achieve a good generalization ability on large-scale problems [1, 42, 43]. Traditional heuristic approaches are delicately designed with great interpretability based on decades of expert knowledge and experience [44]. With the advances in machine learning techniques, powerful network models are trained to learn inherent knowledge and extract complex patterns from known instances to enhance the performance of traditional heuristics [1]. Ye et al. [45] proposed a neural-enhanced meta-heuristic algorithm named DeepACO to leverage reinforcement learning in the heuristic designs of ant colony optimization. NeuroLKH is a representative work proposed by Xin et al. [42]. In NeuroLKH, a sparse graph network (SGN) is combined with the strong heuristic Lin-Kernighan-Helsgaun (LKH) [46, 47] for solving TSP instances. In the traditional LKH, an edge candidate set is pre-defined based on hand-crafted rules for the λ -opt searching process [48], and node penalty values are iteratively optimized via sub-gradient optimization to transform the edge distances for the searching process. NeuroLKH trains an SGN model to generate edge scores and node penalties simultaneously. Then the edge scores are used to create an edge candidate set instead of hand-crafted rules, and the node penalties are used to transform edge distances without performing iterative optimization for each instance as the traditional LKH does. The SGN model is trained on a wide range of instances with different sizes, which are optimized by the exact solver Concorde. When generalized to larger instances, the edge scores in SGN can be directly adopted without any modification, but a fine-tuning step is required to adapt the scale of the node penalties to large sizes. Similarly, VSR-LKH [43] combines three reinforcement learning strategies (Q-learning [49], Sarsa [50] and Monte Carlo [40]) with the traditional LKH algorithm for solving TSP. It replaces the fixed traversal operation in LKH and enables the algorithm to choose each search step by learning from reinforcement learning strategies. Inspired by the great success achieved by Large Language Model (LLM) [51] recently, Liu et al. [52] proposed an evolutionary optimization algorithm assisted by a Large Language Model by incorporating the LLM into the paradigm of evolutionary computation. The proposed method can automatically evolve the elite-guided local search algorithms which outperform human-designed algorithms in solving small-scale TSPs. Wang et al. [53] proposed a novel approach named Adaptive Staircase Policy Space Response Oracle to address the generalization issues of NCO, aiming to help neural solvers explore and adapt to different distributions and various problem scales.

A recent research in [27] brings together several recent works [22, 26, 32, 33] on learning-based TSP approaches into a uniform pipeline, with the aim to investigate various components in the algorithm that influence the generalization ability, including the model construction, learning paradigms and inductive biases [54]. To this end, the TSP is considered as a sequential decision-making task on graphs [22], and a unified end-to-end learning pipeline is constructed with five stages. The first stage is the problem definition, where the TSP is formulated as a fully connected graph [30]. The weight value of each edge can be directly determined by the Euclidean distance according

to the coordinates of the cities. For better computational efficiency and generalization, the graph can be sparsified via a priori knowledge such as k -nearest neighbors [22, 34, 36]. The second stage is graph embedding, where the GNN models are usually trained to output latent space embeddings for nodes and edges of the input graph [23, 24, 55]. Solution decoding and search process takes place in the next two stages. The latent embeddings are transformed into probability values for each edge belonging to the optimal solution set via either autoregressive [30, 32, 33, 56, 57] or non-autoregressive manners [22, 26, 34, 36, 58]. Then feasible solutions to the original problem are generated via graph search approaches such as beam search and greedy search [22, 59, 60]. Finally, the entire model is trained in an end-to-end manner by imitating an expert via supervised learning [22, 34], or by minimizing a cost function via reinforcement learning [10, 33]. The controlled experiments in [27] reveal several insights into zero-shot generalization. Concretely, the dominant evaluation approach tends to obscure the limited generalization ability, as the model performance is evaluated on fixed or excessively small TSP instances. The generalization performance of the GNN-based models benefits from dedicated redesign by considering shifting graph distributions. Compared to non-autoregressive decoding, the autoregressive manner improves generalization by imposing a sequential bias. However, it significantly increases the inference time, especially as the size of the instances grows.

3. The proposed methodology

In this section, we propose our learning-based EdgeGAE approach for solving TSPs of various sizes from a link prediction perspective. We first define the TSP on graphs and outline the overall framework. Then we develop a novel edge-aware residual graph autoencoder model with enhanced graph representation and edge-centered decoding scheme. Considering the various scales and imbalanced distribution of the training data, we design a training strategy with batch encoding and active sampling to improve the generalization ability. Finally, we construct a TSP benchmark dataset composed of 50,000 instances-solution pairs generated by Concorde to simulate the scale-imbalanced distribution in practical scenarios.

3.1. Overall framework

In this paper, we consider the 2-D Euclidean symmetric TSP, where each city has a two-dimensional coordinate, and the distance between two cities is independent of the traveling direction. It aims at finding the shortest route for a salesman to visit each city once and return to the departure city. Given a TSP instance with N cities, the input is a set of order-invariant nodes $S = \{x_i\}_{i=1}^N$ where $x_i \in \mathbb{R}^2$ represents the 2-D coordinates of the i th node. A feasible solution to the problem is a permutation of N nodes $\pi = \{\pi_1, \pi_2, \dots, \pi_N\}$ where π_i is the index of the i th node in the permutation. Assume that the salesman starts from city x_{π_1} and traverses all the cities in the order of the permutation π until x_{π_N} and finally returns to the departure city. The objective is to minimize the total length $L(\pi | S)$ of the tour:

$$L(\pi | S) = \left\| x_{\pi_N} - x_{\pi_1} \right\|_2 + \sum_{i=1}^{N-1} \left\| x_{\pi_i} - x_{\pi_{i+1}} \right\|_2, \quad (1)$$

where x_{π_i} denotes the 2-D coordinates of the i th visited city and $\|\cdot\|_2$ symbolizes the ℓ_2 norm.

In contrast to the autoregressive methods that consider TSPs as a sequential decision-making process, we formulate it as a binary-class link prediction task on graphs. The overall framework of the proposed methodology is illustrated in Fig. 1.

In earlier work [22, 33], TSP instances are represented as complete graphs where there are connections between each pair of nodes (cities) no matter what the Euclidean distance is. In a complete graph representation $\mathcal{G}(\mathcal{V}, \mathcal{E})$, \mathcal{V} denotes the set of N cities as nodes with $|\mathcal{V}| = N$, and \mathcal{E} consists of all undirected edges connecting any two nodes in the

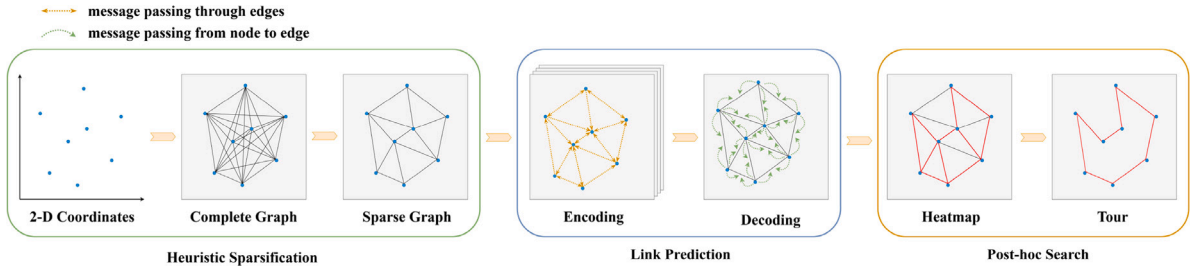


Fig. 1. Solving TSPs as a link prediction task on graphs. In the heuristic sparsification stage, the input coordinates of N cities are first converted into a complete graph where all nodes are directly connected, followed by heuristics such as k -nearest neighbors to get a sparse graph without redundant edges. In the link prediction stage, the sparse graph with node and edge features is input to the encoder to get latent graph embeddings via multiple layers of message passing through edges (in yellow arrows). Then the decoder performs link prediction on each edge by aggregating the embeddings of the source and target node into the edge embedding via an edge-centered message passing scheme (in green arrows). Finally, the output heatmaps are converted into feasible solutions through post-hoc graph search approaches.

set \mathcal{V} . The number of connections $|\mathcal{E}|$ increases exponentially with the number of cities, resulting in a high density of edges. Consequently, the pairwise computation for all nodes will be intractable as the graph becomes larger. Furthermore, according to the definition of the TSP, the connections in an optimal solution are more likely to exist between two nodes that are close to each other, as they usually result in a smaller objective value.

To improve the scalability of the model while preserving valuable information on the graph topology, we adopt k -nearest neighbors heuristics and convert the TSP instance from a complete graph to a sparse graph representation $\hat{\mathcal{G}}(\mathcal{V}, \hat{\mathcal{E}})$ with each node only connected to its k nearest neighbors:

$$\hat{\mathcal{E}} = \{(u, v) \mid u, v \in \mathcal{V}, v \in U_k\}, \quad (2)$$

where U_k is the set of all k -nearest neighbors of node u according to the Euclidean distances. The heuristic graph specification enables learning faster and scaling up to large instances by reducing the search space. In $\hat{\mathcal{G}}(\mathcal{V}, \hat{\mathcal{E}})$, the node features are defined as the normalized 2-D coordinates, and the edge features are the corresponding Euclidean distance between the source and target nodes. We transform the optimization of a TSP instance into a binary-class link prediction task and define the ground-truth label $Y(e)$ of each edge e as whether it exists or not in the optimal tour π^* .

$$Y(e \mid e \in \hat{\mathcal{E}}) = \begin{cases} 1, & \text{edge } e \text{ exists in } \pi^*; \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

In data-driven NCO, a neural network will be trained with labeled data (instances) in an end-to-end manner. After training, the model can output link predictions $\hat{Y}(e)$ on all edges of the input graph.

$$\hat{Y}(e \mid e \in \hat{\mathcal{E}}) \in [0, 1] \quad (4)$$

Finally, the output predictions are transformed into feasible solutions by graph search approaches. In this work, we employ the basic random search and 2-opt as our post-hoc search strategy [61]. Given the link predictions in the form of a heatmap, we construct a TSP tour sequentially by adding one node at each step. Starting from a randomly selected node, at each step we calculate the normalized prediction values of all unvisited nodes as the probability of being selected in the next step. Then the next node is selected via roulette and added to the partial tour. The selection continues recursively until all nodes are visited, followed by a 2-opt local search to further improve the quality of the solution. The 2-opt evaluates all possible edge swaps at each iteration, and selects the swap that can lead to the greatest reduction in the total length of the tour until no further improvements can be made. The post-hoc search algorithm with random sampling and 2-opt local search is shown in Algorithm 1.

Algorithm 1 Post-hoc Graph Search with 2-opt Heuristics

```

1: Input: Link predictions as a heatmap
2: Output: TSP tour
3: Initialize an empty tour tour
4: Randomly select a start node  $n_0$  from  $N$  nodes
5: Add  $n_0$  to tour
6: Mark  $n_0$  as visited
7: for  $i \leftarrow 1$  to  $N - 1$  do
8:   Normalize prediction values for all unvisited nodes
9:   Identify the next node  $n_i$  using roulette selection
10:  Add  $n_i$  to tour
11:  Mark  $n_i$  as visited
12: end for
13: improved  $\leftarrow$  True
14: while improved is True do
15:   improved  $\leftarrow$  False
16:   for  $i \leftarrow 1$  to  $N - 2$  do
17:    for  $j \leftarrow i + 1$  to  $N - 1$  do
18:     newTour  $\leftarrow$  reverse the order of nodes between  $n_i$  and  $n_j$ 
19:     in tour
20:     if  $\text{length}(\text{newTour}) < \text{length}(\text{tour})$  then
21:      tour  $\leftarrow$  newTour
22:      improved  $\leftarrow$  True
23:     end if
24:    end for
25:   end for
26: return tour

```

3.2. The edge-aware residual graph autoencoder

3.2.1. A residual gated encoder

The encoder module focuses on learning enhanced latent space graph embeddings via an efficient message-passing process. The message passing is a fundamental operator on graph-structured data in the encoding scheme, which enables effective information exchange and aggregation among nodes based on the topology of the graph. To capture complex graph patterns via information propagation, the message passing involves a recursive process of broadcasting and updating messages between related nodes in the graph. And the recursive message passing which can map an input graph to its hidden space representation is usually achieved by multiple layers of graph convolution-based operations. Specifically, at each layer, a node aggregates information from its neighbor nodes which have direct connections (edges) with itself, and then calculates and updates its representation according to the integrated information. Such an aggregation and update process occurs at all nodes simultaneously at each step, and this process is repeated for multiple iterations on the graph before generating the

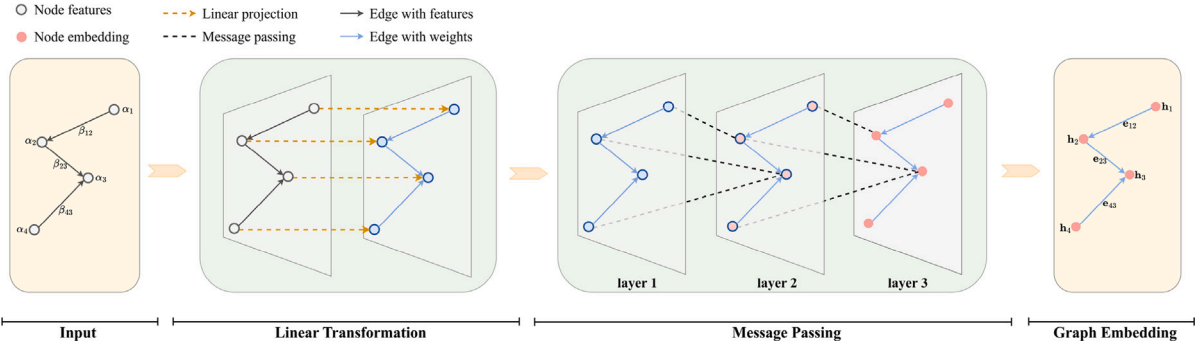


Fig. 2. The residual gated encoder. (1) The input to an encoder model is a directed sparse graph, where the input node features and edge features are represented as α and β , respectively. In the sparse graph representation of the TSP, each node is connected to its k nearest neighbors. (2) The node and edge features are mapped to high-dimensional embeddings through linear projections to enhance the expression power. (3) Multiple message-passing layers enable nodes and edges to learn information from local neighborhoods via residual gated graph convolutions. (4) The encoder outputs the graph with latent space embeddings which can capture inherent knowledge of the graph topology for downstream tasks.

final representation. The motivation of a recursive message-passing process is to enable every node to refine its embeddings for capturing progressively implicit patterns. Within one layer of the message passing, a node learns information from its 1-hop neighbors. After k layers of propagation, each node gets a reception field of its k -hop neighbor nodes. Therefore, the recursive message passing encourages the node embeddings to capture high-level dependencies by aggregating messages from distant nodes, and promotes the encoder to learn hierarchical representations for downstream tasks. Consequently, the depth of the encoder, measured by the number of layers, is crucial for effective feature extraction and message propagation.

As illustrated in Fig. 2, there are mainly three steps in the message passing process of the encoder: *transformation*, *aggregation*, and *update*. In the message transformation step, all neighboring nodes encode their features as messages, which may contain node-specific information or relationships. Then the embeddings are mapped to latent space representations via a linear or non-linear transformation. The mapped embeddings can have either the same or different dimensions as before. In the message aggregation step, each node aggregates the transformed information from its 1-hop neighbors according to the graph topology. Some common choices for aggregation functions include *sum*, *mean*, and *max*. More complex alternatives include weighted aggregation or attention mechanisms. The aggregation function should be permutation invariant to ensure that the aggregation result remains the same regardless of the processing order of nodes. Once the messages are aggregated, in the message update step, each node calculates a refined version of its representation via an update function that takes both the original embedding and the aggregated information as the inputs. Finally, the updated representations are taken as the inputs to the next layer of the encoder. The message-passing process can be generally formulated as:

$$\mathbf{h}_i \leftarrow \gamma \left(\mathbf{h}_i, \bigotimes_{j \in \mathcal{N}_i} \phi(\mathbf{h}_i, \mathbf{h}_j, \mathbf{e}_{j,i}) \right), \quad (5)$$

where \mathbf{h}_i denotes the embedding of the i th node in the graph and $\mathbf{e}_{j,i}$ denotes the embedding of a directed edge from node j to node i . ϕ is the message transformation function that takes the original node or edge embeddings as the inputs, and \bigotimes represents a permutation invariant function such as *sum*. \mathcal{N}_i consists of all 1-hop neighbors of node i , and γ is the update function which calculates the updated embedding \mathbf{h}_i according to the original embedding and the aggregated message from its first-order neighboring nodes.

As formulated in Eq. (5), the message exchange occurs at the local neighbors of each node in the graph, enabling them to receive inherent information about the structure and node features across the entire graph. The expressive capability of latent embeddings plays an important role in the following link prediction task, which requires a good understanding of the relationships and interconnections in the graph.

However, the original graph autoencoder framework [62] only takes a simple graph convolutional network (GCN) as its encoder model. The layer-wise propagation rule is as follows:

$$H^{(l+1)} = \sigma \left(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right) \quad (6)$$

$$\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}, \quad (7)$$

where $H^{(l)}$ denotes the node embeddings of the l th GCN layer. \tilde{A} and \tilde{D} are the adjacency matrix with self-connections and the degree matrix, respectively. σ represents the activation function and $W^{(l)}$ is the learnable weight matrix of the l th layer. In the GCN-based encoder model, only node embeddings are propagated and updated within message-passing layers, without accounting for any edge features. In addition, the same message transformation matrix $W^{(l)}$ is shared among all nodes, limiting the learning capability of the encoder model to identify fine-grained relationships between nodes. Furthermore, the iterative propagation in the standard GCN could result in over-smoothing, significantly reducing the diversity among different nodes after several message-passing layers. To overcome these limitations, we design an edge-aware encoder model with residual connections and a dense-attention mechanism built on ideas from graph ConvNets [35].

For a TSP instance with N cities, its graph representation $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is taken as the input of the encoder. We first map the original node and edge features to high-dimensional embeddings via linear transformations:

$$\mathbf{h}_i = \mathbf{W}_h \cdot \alpha_i + \mathbf{b}_h, \quad \alpha_i \in \mathcal{V} \quad (8)$$

$$\mathbf{e}_{i,j} = \mathbf{W}_e \cdot \beta_{ij} + \mathbf{b}_e, \quad \beta_{ij} \in \mathcal{E}, \quad (9)$$

where $\alpha_i \in \mathbb{R}^2$ is the input node feature of the i th city denoting the 2-dimensional coordinates, and $\beta_{ij} \in \mathbb{R}$ represents the Euclidean distance of the edge between cities i and j . $\mathbf{W}_h \in \mathbb{R}^{H \times 2}$, $\mathbf{W}_e \in \mathbb{R}^{H \times 1}$, $\mathbf{b}_h \in \mathbb{R}^H$ and $\mathbf{b}_e \in \mathbb{R}^H$ are the trainable weight parameters of the linear transformations. \mathbf{h}_i and \mathbf{e}_{ij} are H -dimensional node and edge embeddings, respectively. Since the performance of link prediction for the TSP relies heavily on the total length of selected edges, we construct explicit edge embeddings via an edge convolutional operation in each layer of the encoder:

$$\mathbf{e}_{ij}^{(k+1)} = \mathbf{C}^k \mathbf{e}_{ij}^{(k)} + \mathbf{D}^k \mathbf{h}_i^{(k)} + \mathbf{E}^k \mathbf{h}_j^{(k)}, \quad (10)$$

where $\mathbf{C}^k, \mathbf{D}^k, \mathbf{E}^k \in \mathbb{R}^{H \times H}$ are learnable parameters of the k th edge convolution layer, and the index k indicates that the weight parameters are not shared among different layers. The enhanced edge embeddings

are further adopted to generate a set of dense attention-based weight vectors for node convolution.

$$\omega_{ij}^k = \frac{\sigma(\mathbf{e}_{ij}^{(k)})}{\sum_{j \in \mathcal{N}_i} \sigma(\mathbf{e}_{ij}^{(k)}) + \delta}, \quad (11)$$

where $\sigma(\cdot)$ denotes the sigmoid function and $\delta \in \mathbb{R}^+$ is a small constant. The node convolutional operator in the k th encoder layer is written by:

$$\mathbf{h}_i^{(k+1)} = \mathbf{A}^k \mathbf{h}_i^{(k)} + \sum_{j \in \mathcal{N}_i} \omega_{ij}^k \odot \mathbf{B}^k \mathbf{h}_j^{(k)}, \quad (12)$$

where $\mathbf{A}^k, \mathbf{B}^k \in \mathbb{R}^{H \times H}$ are trainable parameters of the node convolution, and ω_{ij}^k is the gated dense-attention map vector derived from the edge embedding \mathbf{e}_{ij} . A batch normalization (BN) operator is incorporated into the message-passing process for both node and edge embeddings before the activation function ReLU. Finally, we add a residual connection within each encoder layer to alleviate the over-smoothing issue after multiple iterations of message passing. The overall message-passing process within each layer of our edge-aware encoder model is formulated as:

$$\mathbf{h}_i^{(k+1)} = \text{ReLU}(\text{BN}(\mathbf{A}^k \mathbf{h}_i^{(k)} + \sum_{j \in \mathcal{N}_i} \omega_{ij}^k \odot \mathbf{B}^k \mathbf{h}_j^{(k)})) + \mathbf{h}_i^{(k)}. \quad (13)$$

$$\mathbf{e}_{ij}^{(k+1)} = \text{ReLU}(\text{BN}(\mathbf{C}^k \mathbf{e}_{ij}^{(k)} + \mathbf{D}^k \mathbf{h}_i^{(k)} + \mathbf{E}^k \mathbf{h}_j^{(k)})) + \mathbf{e}_{ij}^{(k)}. \quad (14)$$

In contrast to the original GCN-based encoder, the edge-aware encoder model can generate explicit edge representations via multiple layers of an efficient message-passing process. This enables the model to learn distinctive edge embeddings directly from input features (distances). Since the objective function of the TSP is determined by the sum of distances in a tour, such an edge-aware encoding scheme is more effective for link predictions than the simple concatenation of source and target node representations. The graph encoding scheme for enhanced edge representations is illustrated in Algorithm 2.

Algorithm 2 Graph Encoding Scheme

- 1: **Input:** Graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with node features $\alpha_i \in \mathcal{V}$ and edge features $\beta_{ij} \in \mathcal{E}$, number of message passing layers K , latent space dimension H
 - 2: **Output:** latent node encodings $\mathbf{H} \in \mathbb{R}^{|\mathcal{V}| \times H}$ and latent edge encodings $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times H}$
 - 3: Map node features to latent space: $\mathbf{H} \leftarrow \alpha$ in Equation.(8)
 - 4: Map edge features to latent space: $\mathbf{E} \leftarrow \beta$ in Equation.(9)
 - 5: **for** $k = 1$ to K **do**
 - 6: **for** $i = 1$ to $|\mathcal{V}|$ **in parallel do**
 - 7: Edge convolution on $\{\mathbf{e}_{ij}^k \mid j \in \mathcal{N}_i\}$ in Equation.(10)
 - 8: Calculate gated weight vectors ω_{ij}^k in Equation.(11)
 - 9: Calculate node convolution on \mathbf{h}_i^k in Equation.(12)
 - 10: Batch normalization on both edges and nodes:
 - 11: $\mathbf{h}_i^{(k+1)} \leftarrow \text{BN}(\mathbf{h}_i^{(k)})$
 - 12: $\mathbf{e}_i^{(k+1)} \leftarrow \text{BN}(\mathbf{e}_i^{(k)})$
 - 13: Add residual connections after activation function:
 - 14: $\mathbf{h}_i^{(k+1)} \leftarrow \text{ReLU}(\mathbf{h}_i^{(k+1)}) + \mathbf{h}_i^{(k)}$
 - 15: $\mathbf{e}_i^{(k+1)} \leftarrow \text{ReLU}(\mathbf{e}_i^{(k+1)}) + \mathbf{e}_i^{(k)}$
 - 16: **end for**
 - 17: **end for**
 - 18: **return** $\mathbf{H} \in \mathbb{R}^{|\mathcal{V}| \times H}$ and $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times H}$
-

3.2.2. An edge-centered decoder

After encoding the node and edge representation in a latent space via the encoder model, the decoder aims to identify whether each edge exists in the optimal solution by making link predictions on the graph based on the encoding information. Generally speaking, a graph

autoencoder model performs link prediction by leveraging the learned node embeddings to estimate the likelihood of a connection between two nodes of the input graph. In the traditional link prediction tasks such as Cora [62], Citeseer [24] and PubMed [63], there is only one graph in the dataset with thousands of nodes. The encoder-decoder model is trained on an incomplete version of the original graph as training data, where part of the edges are removed from the graph with all node features remaining. The encoder learns to map the input node features into latent embeddings, and the decoder reconstructs the graph based on the learned embeddings. The training target is to minimize the reconstruction error of the input subgraph. Once trained, the node embeddings are extracted from the encoder model for link prediction, which is supposed to have captured the node attributes and relationships of the entire graph. With the extracted embeddings, the decoder model estimates the likelihood of edge existence by calculating the similarity between two nodes. The original graph autoencoder model employs a simple inner product decoder to measure the similarity between a pair of nodes [62]:

$$p(\mathcal{E}_{ij} = 1 \mid \mathbf{h}_i, \mathbf{h}_j) = \sigma(\mathbf{h}_i^\top \mathbf{h}_j), \quad (15)$$

where $p(\mathcal{E}_{ij}) \in [0, 1]$ denotes the probability of an edge \mathcal{E}_{ij} existing between node i and node j . $\mathbf{h}_i, \mathbf{h}_j \in \mathbb{R}^H$ represent the learned node embeddings from the encoder model, and $\sigma(\cdot)$ is the sigmoid function. The value of the inner product between two node vectors implies their similarity in the latent space. According to Eq. (15), the likelihood of the existence of an edge is determined by the degree of similarity between two node embeddings. A greater similarity between two nodes leads to a higher probability of a connection, and vice versa. This kind of similarity-based link prediction principle is based on the assumption that those nodes connected in a graph should have similar features and local neighborhood structures. This assumption applies to node-centered link prediction tasks like Cora, where similar documents do have a greater probability of cross-citation. However, it is not directly applicable to the optimal-tour prediction on graph-based routing problems like TSPs.

Based on the problem definition in Section 3.1, every node on a sparse TSP graph has the same degree determined by a heuristic value k , and the node features involve its two-dimensional coordinates in the Euclidean space. Therefore, all nodes on the graph have a similar local neighborhood structure independent of the node characteristics. Furthermore, there is no guarantee that an edge between two nodes with similar features (i.e., a similar horizontal or vertical coordinate) must exist in the optimal tour. The objective value of a TSP solution is determined by the distances of all edges in the tour, without direct relations to the coordinates of cities. In summary, the original GAE with an inner product decoder performs well on node-centered link prediction tasks, since it encourages similar nodes to have similar embeddings, which is suitable for the graph characteristics in transductive datasets. On the other hand, it overemphasizes the node similarity and proximity at the expense of edge information, leading to a poor generalization ability to link prediction on TSP, which requires the incorporation of both edge and node messages.

For the above reasons, we redesign the decoder structure and propose an edge-aware decoder integrating both node and edge information for efficient link predictions on TSPs. The decoder structure is shown in Fig. 3. The general idea is to decode an edge representation by incorporating its source and target node embeddings via an edge-centered message-passing process. In contrast to the encoding process where messages are passed from one node to another along the edges, in edge-centered decoding, messages on source and target nodes are aggregated into the edge embedding to update itself. To this end, the following three main steps are implemented in the decoding scheme:

- (a) Message transformation. To enhance the expressive power of learned graph representations, a shared linear transformation is performed on all source nodes, target nodes, and edges separately.

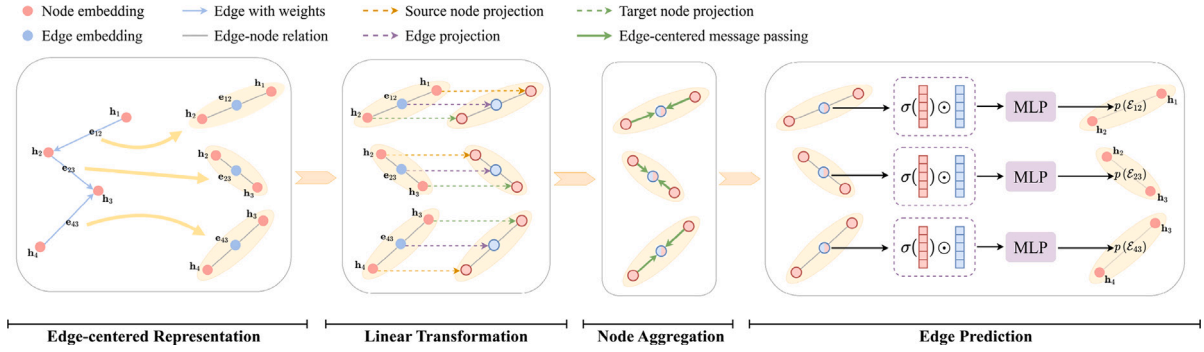


Fig. 3. The edge-centered decoder. (1) We formulate the input graph of the decoder into an edge-centered representation for a clear illustration, where each edge embedding is associated with its source and target node embeddings. (2) Three different linear projections are applied to all source nodes, target nodes, and edge embeddings, respectively. (3) Embeddings from the source and target nodes of each edge are summed up, and then aggregated into the edge-centered message-passing manner. (4) The inner product result of the sigmoid-transformed node embedding and the edge embedding is passed through a multi-layer perceptron (MLP) classifier to output the probability of the edge existing in the optimal tour.

- (b) Node aggregation. For each edge in the graph, the transformed information on its source and target nodes are aggregated.
- (c) Edge update. The aggregated node information is first normalized by the sigmoid function. Then, an inner product is performed on the node and edge embeddings.

Following the message passing scheme in the proposed edge-aware decoder, the decoding output \mathbf{d}_{ij} for an input edge embedding \mathbf{e}_{ij} is calculated by:

$$\mathbf{d}_{ij} = \sigma(\mathbf{F}\mathbf{h}_i + \mathbf{G}\mathbf{h}_j) \odot \mathbf{J}\mathbf{e}_{ij} \quad (16)$$

where $\mathbf{F}, \mathbf{G}, \mathbf{J} \in \mathbb{R}^{H \times H}$ denote learnable parameters of the linear transformation for target nodes, source nodes, and edges, respectively, and $\sigma(\cdot)$ is the sigmoid function. Finally, the H -dimensional decoding is converted into a probability prediction via a simple MLP model. The edge-centered decoding scheme for predicting the TSP edges is described in Algorithm 3.

$$p(\mathcal{E}_{ij} = 1) = \text{MLP}(\mathbf{d}_{ij}) \quad (17)$$

Algorithm 3 Edge-centered Decoding Scheme

1: **Input:** Graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, node embeddings $\mathbf{H} \in \mathbb{R}^{|\mathcal{V}| \times H}$, edge embeddings $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times H}$
 2: **Output:** link prediction results
 3: **for** $i = 1$ to $|\mathcal{V}|$ **in parallel do**
 4: **for** $j \in \mathcal{N}_i$ **in parallel do**
 5: Transform messages on nodes and edges:
 6: $\hat{\mathbf{h}}_i \leftarrow \mathbf{F}\mathbf{h}_i, \hat{\mathbf{h}}_j \leftarrow \mathbf{G}\mathbf{h}_j, \hat{\mathbf{e}}_{ij} \leftarrow \mathbf{J}\mathbf{e}_{ij}$
 7: Aggregation from source and target nodes:
 8: $\hat{\mathbf{h}}_{agg} = \hat{\mathbf{h}}_i + \hat{\mathbf{h}}_j$
 9: Update the edge embedding:
 10: $\mathbf{d}_{ij} = \sigma(\hat{\mathbf{h}}_{agg}) \odot \hat{\mathbf{e}}_{ij}$
 11: Calculate the existence probability:
 12: $p(\mathcal{E}_{ij} = 1) = \text{MLP}(\mathbf{d}_{ij})$
 13: **end for**
 14: **end for**
 15: **return** $\{p(\mathcal{E}_{ij} = 1) \mid \mathcal{E}_{ij} \in \mathcal{E}\}$

3.3. Training strategy

In data-driven NCO approaches, it is always challenging to handle data with various scales and imbalanced distributions. A typical scenario is combinatorial optimization in real-world applications, where the number of historical cases of small scales is much larger than the cases of large scales. Learning algorithms often face difficulties in generalizing inductive rules due to this kind of imbalance. Therefore, dealing with such problems requires the data-driven model to have the

ability to learn from scale-imbalanced data and generalize to instances of arbitrary scales. In general, a model trained on small-scale instances tends to have a significant degradation in performance on large-scale problems, since the implicit knowledge learned from small-scale data cannot be directly transformed to large-scale problems. It has been demonstrated that learning from data with variable sizes is beneficial for the model's generalization ability [27]. However, training a generalizable NCO model with limited large-scale instances while effectively leveraging knowledge from small-scale data remains an open question.

In this work, we introduce a simple yet effective training strategy called *random active sampling* [64] into our training process, in order to learn from scale-imbalanced TSP instances efficiently for link prediction tasks. The active sampling methods are originally designed to address imbalanced classification problems, where a majority class D_{maj} has significantly more data than other minority classes D_{min} . Since an equalized dataset generally improves the overall classification accuracy than an imbalanced dataset, an intuitive motivation of active sampling methods is to modify the dataset with customized sampling mechanisms to get a balanced distribution. Random *oversampling* and *undersampling* are two main strategies. The mechanism of random *oversampling* can be naturally implemented by adding a set of data from the minority class into the original dataset. Specifically, we randomly select S samples from the minority class D_{min} to duplicate and merge them with the original set. By doing so, the total number of minority class examples is augmented by S , leading to an adjustment in the class distribution balance. On the contrary, the random *undersampling* method contributes to the class balance by reducing the number of samples in the majority class. In particular, we randomly select S samples from the majority class D_{maj} and remove them from the dataset. As a result, the overall samples are adjusted manually for a balanced distribution via both *undersampling* and *oversampling* strategies.

In the link prediction task with scale-imbalanced TSP cases as the training data, we treat TSP instances with different numbers of cities as different classes. Hence for a TSP dataset with various sizes from 50 to 500 cities, there are 451 classes in total. Instead of removing or duplicating existing samples in the original dataset, we modify the traditional AS strategies by incorporating *oversampling* and *undersampling* into a *uniform-sampling* on classes, with the aim to improve the sampling fairness and the data utilization. Specifically, we follow a two-step sampling manner for generating each batch of data during the training process. In the first step, we select B classes uniformly from all classes of the number of cities, and B is the batch size. In the second step, for each of the B choices of the city number, we randomly sample an instance from the corresponding class, forming a batch of training data with a uniform distribution in the number of cities. In the original training process, all data from the training set are shuffled and divided into batches of data randomly. As a result, every batch of training data




Fig. 4. Training data distributions with different sampling strategies. The random sampling retains the original distribution of the training dataset, where small-scale data is the majority and large-scale data is the minority. The active sampling strategy can change the original distribution to a balanced distribution of the number of nodes.

will follow a similar distribution to the overall dataset. On the contrary, if each batch of data is generated via a class-uniform sampling strategy, then the overall training data will follow a balanced distribution of the number of nodes, regardless of the original distribution of data. The distribution of different training datasets obtained with different sampling strategies are shown in Fig. 4.

3.4. Dataset generation

To simulate the scale-imbalanced distribution in realistic optimization tasks, we create a TSP benchmark dataset with 50,000 instances as the training data. The dataset consists of graphs of different sizes from 50 to 500 nodes, and the number of graphs per size varies in inverse proportion to the number of nodes. As a result, there are 439 graphs with 50 nodes, and only 43 graphs have 500 nodes. Following the learning-based work in [22,27,34], the 2-dimensional locations $\{x, y\}$ of each node are uniformly sampled in the unit square $[0, 1]^2$ for each instance. The optimal solution to each instance is provided by Concorde in the form of the ground truth label of each edge existing in the optimal tour. To test the generalization ability of models at various scales, we further create five separate test datasets for graphs of size $\{50, 100, 300, 500, 700\}$ nodes, each including 1000 TSP instance-solution pairs. Fig. 5 illustrates the optimal tours obtained by Concorde.

4. Experimental results

In this section, we conduct systematic experiments on different amounts of training data, different scales of test cases, and different numbers of samples to compare the performance of our proposed approach with state-of-the-art learning-based NCO methods as well as the Concorde solver, in order to verify the effectiveness of the proposed method.

4.1. Experimental settings

As described in Section 3.1, this work considers the 2-dimensional symmetric TSPs with N cities randomly located in a unit square $[0, 1]^2$ for simplification. The objective is to find an optimal route for a salesman to visit all cities once and return to the starting point, in order to minimize the total length of the tour. We formulate a TSP instance as a sparse graph representation based on heuristic information, where

each node is a city and the edge weight represents the Euclidean distance between two cities. In the sparse graph representation, each node is only connected to its K nearest neighbors according to the distances, where K is a hyperparameter. The aim is to identify the shortest possible Hamiltonian cycle of the graph via link prediction and graph search methods. In our experimental setup, we configure four different levels of training data quantities with 1000, 5000, 10000, and 50000 instances. For each level of the training quantity, all instances are randomly sampled from the TSP benchmark dataset in Section 3.4, forming a scale-imbalanced distribution where the number of instances is inversely proportional to the number of nodes. The test dataset comprises a total of 5000 TSP cases, including five distinct scales of instances: TSP50, TSP100, TSP300, TSP500, and TSP700, with each scale containing 1000 instances. It is worth noting that all test cases are completely independent of the training dataset. The optimal solutions to all test cases are generated by the Concorde solver. During the test process, the trained model outputs the link-predicted results for each input case in the form of a heatmap, indicating the probability of each edge existing in the optimal tour. Subsequently, the probability predictions are converted into feasible solutions via post-hoc graph search approaches. We explore two different post-hoc search strategies, namely beam search and random search, to generate feasible TSP solutions. Beam search is a fixed-width breadth-first searching approach that is widely adopted in generating high-probability sequences [16]. Starting from the first node, the heatmap is explored recursively to construct b solutions in parallel. Specifically, at each iteration, the top- b edges with the highest probability in the heatmap will be selected and added to the b partial tours. For each partial tour, the selected nodes will be masked to ensure the feasibility of the solution. The recursive selection process will continue to expand the b tours until all nodes have been selected for each solution. The final prediction of beam search is determined by the tour with the minimum total length among the b complete tours. Random search constructs a feasible solution to a TSP instance sequentially by adding one node at each step in an iterative manner. In contrast to the beam search where fixed b nodes with the highest probabilities are selected at each step, random search determines the next node from all unvisited nodes via the roulette wheel selection, enabling the possibility of each node being selected with different probabilities. We adopt two search strategies in our experiments, and the results indicate that the random search achieves a better performance in terms of the solution quality. Therefore, we leverage the random search as our post-hoc graph search strategy in our experiments for all comparison methods, followed by a 2-opt local search to further enhance the model performance. We explore the effects of three different sampling quantities on solution performance by sampling 200, 500, and 1000 solutions for each test case. To investigate the effectiveness of the active sampling method introduced in Section 3.3, we train all learning-based models in our comparative experiments using two training strategies separately, with and without the *random active sampling* method.

4.2. Baseline methods

We compare the proposed methodology with both an exact solver and state-of-the-art NCO methods as baselines. All learning-based models are trained in a supervised learning manner for a fair comparison.

1. *Concorde* [65]: Concorde is an exact solver specially designed for solving TSP. It builds on ideas from branch-and-cut algorithms, and incorporates linear programming and cutting plane methods to find optimal solutions to NP-hard combinatorial optimization problems including TSPs. Concorde has demonstrated exceptional performance in finding optimal solutions to TSP instances up to thousands of cities.

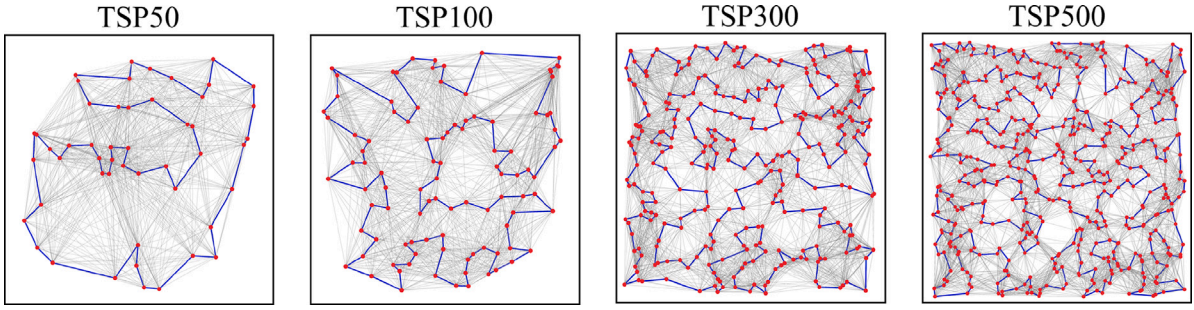


Fig. 5. TSP instances of different scales in the training dataset. Blue edges indicate the optimal tour obtained by the Concorde solver.

2. *Multi-layer Perceptron (MLP)* [66]: The MLP model takes the node features as inputs, with no consideration of the graph topology information. It first maps the node inputs to a high-dimensional latent space via multiple layers of linear projections, and generates edge embeddings by concatenating the two node embeddings of each edge. Finally, the edge embeddings are passed through a readout module to output probability predictions in the form of a heatmap.
3. *Graph Convolutional Network (GCN)* [23]: The vanilla GCN model takes the graph-structured data as input, which incorporates both the node and edge features and the topology information. The message passing occurs between neighbor nodes via multiple layers of graph convolutional operators, where each node aggregates information about its neighbors with the same weights.
4. *Graph Attention Network (GAT)* [24]: The GAT model introduces masked self-attention layers and enables nodes to attend over the features of its neighbors by implicitly allocating different weights to neighbors during the message aggregation. For each graph attention layer in the GAT model, the input node features are first projected to higher-level expressions via a shared linear transformation, and then attention coefficients between each pair of nodes are calculated by a shared attention mechanism. The coefficients are normalized using the softmax function to serve as the weights of a linear combination of all neighbor embeddings for each node. The edge embeddings are represented as the concatenation of node embeddings.
5. *Residual Gated Graph ConvNets (GatedGCN)* [35]: The GatedGCN model leverages an edge gating mechanism [67] to allocate different weight vectors to neighbor node embeddings for the aggregation operator. It also incorporates the residual connections and batch normalization in the graph convolutional operation. Compared to the vanilla GCN model, the GatedGCN introduces explicit edge feature representation and update except for the message passing of node embeddings. The edge gates serve as a soft attention map with respect to the original sparse attention mechanism [18,24].

4.3. Implementation details

All learning-based models are trained with the same hyperparameter settings. For the model configuration, each model consists of $L = 4$ message-passing layers with graph convolution-based operators, followed by an MLP block with 3 fully-connected layers to output link prediction values for all edges in the graph. The input node and edge features are 2 and 1 dimension, respectively, and the latent spaces of both node and edge embeddings have the same hidden dimension $H = 64$ for each layer. We set $k = 25$ in the k -nearest neighbor heuristic for all TSP instances of the training data. For the training process, we minimize the binary cross-entropy loss using the Adam optimizer, where the initial learning rate is $lr = 0.001$ without decay for all models. The batch size is set to $B = 32$ for all experiments, and

the maximum number of training epochs is 500. During the test process, we evaluate the trained model on five test datasets of different problem scales: TSP50, TSP100, TSP300, TSP500, and TSP700 with a batch size of 100. The learning-based methods are run on a single NVIDIA A40 GPU with 48 GB memory.

4.4. Performance indicators

Following the previous work in [22,27,34], we evaluate the model performance on solving TSP cases via three indicators: F1 score, ROC AUC score, and the optimal gap.

4.4.1. F1 score

It is commonly adopted as the performance indicator for binary classification tasks. It is particularly preferable over accuracy when dealing with uneven datasets such as the TSP, where one class significantly dominates the other. For example, in a complete graph with N^2 edges, only the $2N$ edges in the optimal tour have positive labels. In such cases, the accuracy alone may not represent the model performance precisely. For a binary classification task, the numbers of true positive (TP), false positive (FP), and false negative (FN) samples are used to calculate the following F1 score:

$$F1 = \frac{TP}{TP + \frac{1}{2}(FP + FN)}. \quad (18)$$

4.4.2. ROC AUC score

It is defined as the area under the receiver operating characteristic curve from prediction scores. The ROC AUC score represents the effectiveness of a model in distinguishing between the positive and negative classes. A higher ROC AUC score implies a better performance at the binary classification task. It is suitable for class-imbalanced datasets, as it is not affected by changes in the classification threshold. For a binary classification task, TPR and FPR are the true positive rate and false positive rate for different threshold values. The ROC AUC score is calculated as:

$$ROC\ AUC = \int_{t=0}^1 TPR(FPR^{-1}(t)) dt. \quad (19)$$

4.4.3. Optimal gap

It measures how close an obtained solution to the optimal one is. In our experiments, we define the optimal gap with respect to the Concorde solver, which calculates the normalized difference between the predicted tour length S and the optimal solution S_{best} provided by Concorde:

$$GAP = \frac{S - S_{best}}{S_{best}} \times 100\%. \quad (20)$$




Fig. 6. The optimal gap of different sample sizes on test cases.

Table 1

The average optimal gap (%) on 1,000 test cases w.r.t Concorde of learning-based models on TSP test datasets of different scales.

	Test Dataset	TSP50			TSP100			TSP300			TSP500			TSP700		
		Sample size	200	500	1000	200	500	1000	200	500	1000	200	500	1000	200	500
Train with 1k Data	MLP	1.854	1.180	0.731	4.718	4.047	3.517	8.226	7.711	7.362	9.473	8.979	8.635	10.003	9.575	9.310
	GCN	1.798	1.083	0.731	4.761	3.935	3.484	8.297	7.666	7.293	9.215	8.737	8.400	9.830	9.382	9.084
	GAT	1.955	1.234	0.796	4.584	3.867	3.494	8.188	7.479	7.168	9.158	8.641	8.278	9.625	9.289	8.99
	GatedGCN	1.627	1.079	0.774	4.473	3.737	3.227	7.973	7.523	7.124	9.161	8.717	8.290	9.703	9.244	9.042
	GatedGCN ^a	1.397	0.872	0.542	4.053	3.484	3.006	7.498	7.079	6.686	8.692	8.153	7.769	9.214	8.818	8.564
	MLP+AS	1.743	1.113	0.819	4.613	3.857	3.309	7.966	7.497	7.182	9.260	8.809	8.461	10.002	9.456	9.117
	GCN+AS	1.859	1.233	0.888	4.680	3.933	3.447	8.338	7.711	7.342	9.227	8.708	8.436	9.867	9.307	9.077
	GAT+AS	1.826	1.168	0.811	4.685	3.945	3.458	7.910	7.443	7.093	9.081	8.528	8.298	9.664	9.283	8.965
	GatedGCN+AS	1.764	0.960	0.686	4.420	3.526	3.065	8.049	7.457	7.103	9.089	8.696	8.420	9.764	9.398	9.116
	GatedGCN ^a +AS	1.261	0.723	0.434	3.835	3.084	2.620	7.136	6.662	6.324	8.181	7.793	7.533	8.720	8.376	8.079
EdgeGAE (ours)	1.148	0.642	0.435	3.893	3.111	2.660	7.149	6.487	6.100	8.010	7.646	7.375	8.629	8.289	8.032	
EdgeGAE+AS (ours)	1.120	0.735	0.424	3.652	3.091	2.680	7.182	6.612	6.250	8.176	7.669	7.451	8.796	8.347	8.029	
Train with 5k Data	MLP	1.768	1.134	0.834	4.603	3.719	3.251	8.181	7.669	7.184	9.402	8.887	8.528	9.916	9.499	9.254
	GCN	1.904	1.229	0.859	4.687	3.909	3.487	8.066	7.577	7.262	9.101	8.728	8.295	9.531	9.156	8.827
	GAT	1.823	1.069	0.773	4.677	4.061	3.420	7.990	7.486	7.142	8.933	8.483	8.072	9.494	9.145	8.830
	GatedGCN	1.388	0.781	0.519	3.725	3.221	2.727	7.313	6.829	6.475	8.383	7.873	7.648	9.088	8.712	8.506
	GatedGCN ^a	1.083	0.609	0.385	3.890	3.096	2.641	7.066	6.576	6.215	7.997	7.617	7.272	8.547	8.194	7.886
	MLP+AS	1.764	1.109	0.704	4.509	3.760	3.221	7.972	7.450	7.075	9.323	8.848	8.481	9.886	9.464	9.199
	GCN+AS	1.772	1.041	0.645	4.573	3.857	3.303	8.040	7.532	7.092	9.002	8.371	8.122	9.424	9.115	8.790
	GAT+AS	1.788	1.151	0.802	4.556	3.871	3.389	7.980	7.443	7.113	8.951	8.490	8.200	9.470	9.072	8.811
	GatedGCN+AS	1.435	0.901	0.638	4.361	3.642	3.065	7.844	7.291	6.926	8.888	8.398	8.055	9.453	9.114	8.810
	GatedGCN ^a +AS	1.190	0.568	0.315	3.719	3.033	2.576	7.113	6.513	6.208	8.195	7.668	7.437	8.816	8.426	8.184
EdgeGAE (ours)	0.958	0.539	0.280	3.568	2.768	2.382	6.864	6.425	5.992	8.100	7.562	7.224	8.438	8.108	7.837	
EdgeGAE+AS (ours)	0.949	0.525	0.280	3.466	2.856	2.354	6.719	6.220	5.767	7.759	7.325	6.939	8.318	7.908	7.697	
Train with 10k Data	MLP	1.933	1.208	0.811	4.572	3.783	3.223	8.188	7.635	7.284	9.454	9.008	8.632	9.934	9.519	9.203
	GCN	1.761	1.123	0.723	4.721	3.957	3.490	7.908	7.287	6.951	9.004	8.503	8.140	9.349	8.946	8.732
	GAT	1.930	1.268	0.805	4.677	3.979	3.526	7.917	7.325	7.019	8.963	8.536	8.241	9.433	9.098	8.766
	GatedGCN	1.544	0.923	0.589	4.281	3.482	2.975	7.572	6.933	6.667	8.451	7.944	7.713	9.030	8.640	8.370
	GatedGCN ^a	1.283	0.601	0.405	3.953	3.206	2.684	7.371	6.826	6.401	8.416	7.929	7.624	9.114	8.662	8.398
	MLP+AS	1.811	1.071	0.719	4.569	3.828	3.337	8.118	7.522	7.214	9.280	8.897	8.543	10.024	9.544	9.235
	GCN+AS	2.039	1.253	0.772	4.536	3.868	3.327	8.002	7.500	7.007	8.996	8.490	8.131	9.521	9.037	8.810
	GAT+AS	1.821	1.150	0.821	4.502	3.868	3.369	7.939	7.328	6.956	8.868	8.366	8.164	9.372	8.996	8.789
	GatedGCN+AS	1.518	0.882	0.586	4.417	3.743	3.186	7.744	7.105	6.772	8.716	8.306	8.005	9.273	8.977	8.697
	GatedGCN ^a +AS	1.092	0.582	0.271	3.742	3.046	2.632	7.223	6.566	6.239	8.012	7.626	7.260	8.642	8.287	8.014
EdgeGAE (ours)	0.753	0.366	0.247	3.253	2.638	2.257	6.760	6.246	6.009	7.821	7.437	7.074	8.369	7.971	7.718	
EdgeGAE+AS (ours)	0.769	0.365	0.208	3.352	2.626	2.233	6.825	6.301	5.946	7.897	7.375	7.128	8.392	8.082	7.748	
Train with 50k Data	MLP	1.623	1.052	0.729	4.461	3.805	3.325	8.102	7.634	7.182	9.198	8.841	8.500	9.919	9.513	9.215
	GCN	1.858	1.107	0.787	4.603	3.930	3.438	7.920	7.390	7.033	8.876	8.405	8.050	9.342	8.971	8.696
	GAT	1.701	1.174	0.802	4.507	3.751	3.288	7.889	7.379	6.916	8.799	8.439	8.078	9.405	8.951	8.744
	GatedGCN	1.285	0.759	0.484	4.318	3.555	3.010	7.532	7.098	6.646	8.503	8.077	7.753	9.007	8.69	8.378
	GatedGCN ^a	0.822	0.415	0.243	3.417	2.785	2.373	6.627	6.196	5.804	7.835	7.405	7.053	8.430	8.072	7.795
	MLP+AS	1.621	0.938	0.678	4.510	3.860	3.184	8.131	7.618	7.086	9.208	8.800	8.495	9.874	9.481	9.220
	GCN+AS	1.891	1.126	0.727	4.618	3.876	3.514	7.889	7.274	6.985	8.808	8.350	8.062	9.275	8.961	8.737
	GAT+AS	1.727	1.098	0.769	4.684	3.830	3.468	8.042	7.499	7.061	9.025	8.543	8.230	9.457	9.110	8.831
	GatedGCN+AS	1.281	0.694	0.521	4.124	3.375	3.002	7.530	6.951	6.654	8.462	8.113	7.811	8.972	8.667	8.347
	GatedGCN ^a +AS	0.886	0.486	0.247	3.463	2.812	2.303	7.052	6.483	6.068	7.904	7.400	7.106	8.459	8.105	7.857
EdgeGAE (ours)	0.723	0.311	0.137	3.025	2.330	2.056	6.652	6.071	5.640	7.564	7.112	6.814	8.026	7.785	7.502	
EdgeGAE+AS (ours)	0.694	0.323	0.168	3.069	2.463	2.028	6.572	6.089	5.772	7.529	7.170	6.904	8.161	7.754	7.499	

^a Indicates the usage of explicit edge representations in the message passing process.

The mark (+AS) indicates training the model with the *random active sampling* strategy in Section 3.3.

4.5. Results and discussion

4.5.1. Comparison against baseline models

The performance of the proposed EdgeGAE versus the baseline methods in terms of the optimal gap is shown in Table 1. There are four sets of experiments with different amounts of training data. For each set of experiments, the results of test cases of different sizes are evaluated on models trained on the same set of data. The proposed EdgeGAE shows superior performance compared to other learning-based methods in most cases of the training data amounts and problem sizes. In

the experiment with only 1000 training instances, EdgeGAE achieves the lowest optimal gap on TSP50, TSP500, and TSP700, regardless of the sample size. On the TSP100 test dataset, the proposed model exhibits a slightly inferior optimal gap compared to the GatedGCN model with explicit edge representations when the sample size of each case is set to 500 and 1000. One possible reason is that the small amount of training data is insufficient for the models to learn enough implicit knowledge, resulting in insignificant differences in performance between models at small and medium-scale problems. Among all the baseline methods, the MLP has the highest optimal

gap, as it is the only learning-based model that solely takes node features as its input without considering any topology information. GCN performs worst among all the message passing-based models, because there is no attention mechanism which can allocate learnable weights to different neighbors during the message-passing process. In various scales of problem sizes, the GatedGCN with explicit edge representations exhibits significant performance improvement compared to the original version, revealing the importance of effective edge embeddings in link prediction tasks. In other words, merely considering node embeddings in message passing is insufficient for the model to generalize well on problems of various scales. In the experiment with 5000, 10000, and 50000 training instances, both the EdgeGAE models trained with and without *random active sampling* strategy outperform other baseline methods at all settings of problem scales and sample sizes. The advantage becomes more notable as the number of training samples increases. The comparative results in terms of the optimal gap imply the effectiveness of the combination of the residual gated graph encoder and the edge-centered decoder model in link prediction for TSP tasks. Across all four experiment sets, the learning-based approaches exhibit a decline in performance with the increase of the problem sizes. The trained models achieve the highest average performance on TSP50 and the lowest on TSP700 test cases.

4.5.2. Sensitivity analysis of sample sizes

To investigate the influence of different sample sizes on the model performance, we sample 200, 500, and 1000 solutions for each test case in each group of experiments. Fig. 6 compares the average optimal gaps of EdgeGAE with different sample sizes under four experimental settings. It can be observed that a higher number of samples leads to a lower gap across all scales of test cases. It stands to reason that increasing the number of sampled solutions for a particular TSP instance increases the likelihood of obtaining better objective values. Nonetheless, it is worth noting that a greater number of samples generally comes with additional computational overhead. Consequently, it is crucial in real-world applications to balance the model performance and computational costs. Furthermore, the comparative results also show that the average optimal gap increases as the problem scale grows, which is irrelevant to the sample size.

4.5.3. Generalization ability

To demonstrate the ability of EdgeGAE to generalize over larger instances than it has ever seen, we evaluate the model performance on TSP700 cases after training on imbalanced data ranging from 50 to 500 cities. In Table 1, the optimal gap on TSP700 serves as the performance indicator of the proposed model and compared baselines. It can be observed that given a sufficient number of samples per case, the proposed model achieves an average optimal gap on TSP700 which is remarkably close to that of TSP500. Compared to other baseline methods, the proposed EdgeGAE exhibits a superior ability to generalize to large-scale problems. Furthermore, the model trained on a larger amount of small-scale instances tends to generalize better on large scales than those with limited training data. Hence, it can be concluded that the model can benefit from more training data even sampled from small-scale instances.

4.5.4. Inference time and complexity analysis

We further compare the inference time of the proposed EdgeGAE and the classical TSP solver on different scales of TSP instances, and the statistical results are presented in Fig. 7. For each problem size, the trained EdgeGAE model followed by graph search makes prediction on 1000 unseen TSP instances to obtain feasible solutions. The Concorde solver performs optimization on the same set of problems for a fair comparison. Fig. 7 illustrates the average inference times of two methods for each test case at five different scales. It is worth noting that on smaller-sized cases with no more than 100 cities, both the EdgeGAE and the Concorde solver require extremely short inference times with

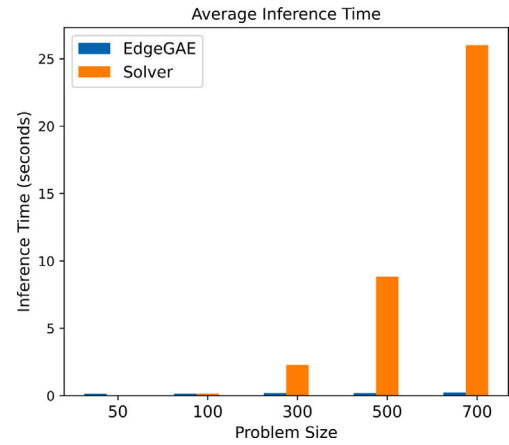


Fig. 7. Average inference time of the EdgeGAE and the classical TSP solver Concorde on TSP instances of different scales.

a negligible difference. However, as the problem size increases, the inference time taken by the solver to find an optimal solution grows exponentially. By contrast, the inference time of EdgeGAE remains basically constant, which is not affected by the problem size. This result implies that the EdgeGAE has a significant advantage over traditional solvers in terms of inference time, especially for solving large-scale optimization problems with real-time decision requirements. In the EdgeGAE model, the message passing process for nodes is calculated according to Eq. (13) with the time complexity of $O(d_n \cdot |\mathcal{N}|)$, where d_n is the dimension of the node representation, and $|\mathcal{N}|$ is the size of the neighborhood. Similarly, the message passing of edges in Eq. (14) has the time complexity of (d_e^2) , where d_e is the dimension of the edge representation. Therefore, the overall time complexity for each message passing layer in EdgeGAE is $O(d_n \cdot |\mathcal{N}| + d_e^2)$.

4.6. Ablation study

To investigate the effect of different sampling strategies during the training process on model performance, we conduct an ablation study to train each comparison model on the same set of training data using random shuffle sampling and random active sampling strategies separately, and evaluate them on different sizes of test cases. The model performances are compared in Table 1, where each subset of models without the mark (+AS) is trained with randomly shuffled data batches, and each subset of models with the mark (+AS) is trained with batches of data generated via the random active sampling strategy in Section 3.3. For a fair comparison, both sampling strategies share the same set of training data in each set of settings. The optimal gaps of EdgeGAE present noteworthy findings. When the training dataset is limited to only 1000 instances, the performance of random shuffle sampling surpasses that of active sampling. One reasonable explanation is that the active sampling from an extremely small dataset causes the model to focus solely on a portion of the data, thereby intensifying the model overfitting on training data. When the training dataset contains 5000 scale-imbalanced cases, the active sampling strategy outperforms random shuffle sampling significantly in any case of problem scales and sample sizes. The advantages of active sampling over random shuffle sampling become less noticeable as the quantity of training data continues to increase. These results indicate that having a considerable amount of training data can partially mitigate the effect of scale-imbalanced distributions on model performance, advocating the necessity of using active sampling strategies on small-scale constrained datasets with size-skewed distributions.

5. Conclusion

In this paper, we investigate solving scale-imbalanced TSPs as link predictions on graphs. To cope with this task, we propose a data-driven NCO framework named EdgeGAE, which enhances the original encoder–decoder framework with a residual gated encoder model and an edge-centered decoder model. The residual gated encoder model introduces explicit edge embeddings in the message-passing process for link prediction tasks. It incorporates residual connections and an edge gated attention mechanism into the vanilla graph convolution. The edge-centered decoder model replaces the inner product with an edge-aware decoding strategy, where each edge aggregates information from its source and target nodes via an edge-centered message-passing scheme. The proposed EdgeGAE framework enables the input graph to capture topology information and learn latent representations via message passing on both nodes and edges for link prediction on TSPs. To address the scale-imbalanced data distribution issues, we further introduce an active sampling strategy for the training process. We generate a scale-imbalanced TSP benchmark dataset with 50,000 instances and compare the proposed framework with state-of-the-art learning-based methods. Experimental results on various problem scales show that EdgeGAE outperforms other baselines in terms of solution quality, sample efficiency, and generalization ability.

Despite the competitive performance of the proposed method, it should be noted that there is still a gap between the learning-based NCO methods and exact solvers. To further improve the model performance with limited training data, we will consider incorporating semi-supervised learning and reinforcement learning into the learning process of data-driven NCO approaches. Moreover, existing NCO models are mostly trained in an offline manner and require all data to be centralized in a local storage. Considering the increasing demand for privacy protection, our future work will extend data-driven NCO to federated learning settings where privacy data is kept locally without sharing during the learning process. We will also consider NCO with online optimization to deal with time-varying problems in real-world applications.

CRedit authorship contribution statement

Shiqing Liu: Writing – original draft, Validation, Software, Investigation, Formal analysis, Data curation, Conceptualization. **Xueming Yan:** Investigation, Formal analysis, Conceptualization. **Yaochu Jin:** Writing – review & editing, Supervision, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant No. 62006053 and No. 62136003. All authors agree with the author list and the submission of the paper.

References

- [1] Y. Bengio, A. Lodi, A. Prouvost, Machine learning for combinatorial optimization: A methodological tour d'horizon, *European J. Oper. Res.* 290 (2) (2021) 405–421.
- [2] Y. Xu, M. Fang, L. Chen, G. Xu, Y. Du, C. Zhang, Reinforcement learning with multiple relational attention for solving vehicle routing problems, *IEEE Trans. Cybern.* 52 (10) (2022) 11107–11120.
- [3] D. Sanches, D. Whitley, R. Tinós, Improving an exact solver for the traveling salesman problem using partition crossover, in: *Proceedings of the Genetic and Evolutionary Computation Conference, 2017*, pp. 337–344.
- [4] D.J. Rosenkrantz, R.E. Stearns, P.M. Lewis, II, An analysis of several heuristics for the traveling salesman problem, *SIAM J. Comput.* 6 (3) (1977) 563–581.
- [5] K. Helsgaun, An extension of the Lin-Kernighan-Helsgaun TSP solver for constrained traveling salesman and vehicle routing problems, 2017.
- [6] J. Cui, L. Wu, X. Huang, D. Xu, C. Liu, W. Xiao, Multi-strategy adaptable ant colony optimization algorithm and its application in robot path planning, *Knowl.-Based Syst.* (2024) 111459.
- [7] S. Wang, B. Ding, Y. Jin, A multi-factorial evolutionary algorithm with asynchronous optimization processes for solving the robust influence maximization problem, *IEEE Comput. Intell. Mag.* 18 (3) (2023) 41–53.
- [8] S. Wang, Y. Jin, M. Cai, Enhancing the robustness of networks against multiple damage models using a multifactorial evolutionary algorithm, *IEEE Trans. Syst., Man, Cybern.: Syst.* (2023).
- [9] S. Wang, X. Tan, Solving the robust influence maximization problem on multi-layer networks via a memetic algorithm, *Appl. Soft Comput.* 121 (2022) 108750.
- [10] I. Bello, H. Pham, Q.V. Le, M. Norouzi, S. Bengio, Neural combinatorial optimization with reinforcement learning, in: *International Conference on Learning Representations, 2017*.
- [11] Y. Jin, H. Wang, T. Chugh, D. Guo, K. Miettinen, Data-driven evolutionary optimization: An overview and case studies, *IEEE Trans. Evol. Comput.* 23 (3) (2018) 442–458.
- [12] Y. Jin, H. Wang, C. Sun, *Data-Driven Evolutionary Optimization*, Springer, 2021.
- [13] J.J. Hopfield, D.W. Tank, “Neural” computation of decisions in optimization problems, *Biol. Cybern.* 52 (3) (1985) 141–152.
- [14] B. Angeniol, G.D.L.C. Vaubois, J.-Y. Le Texier, Self-organizing feature maps and the travelling salesman problem, *Neural Netw.* 1 (4) (1988) 289–293.
- [15] I. Sutskever, O. Vinyals, Q.V. Le, Sequence to sequence learning with neural networks, *Adv. Neural Inf. Process. Syst.* 27 (2014).
- [16] Y. Wu, M. Schuster, Z. Chen, Q.V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, et al., Google's neural machine translation system: Bridging the gap between human and machine translation, 2016, arXiv preprint arXiv:1609.08144.
- [17] O. Vinyals, M. Fortunato, N. Jaitly, Pointer networks, *Adv. Neural Inf. Process. Syst.* 28 (2015).
- [18] D. Bahdanau, K. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, 2014, arXiv preprint arXiv:1409.0473.
- [19] F. Chen, Y.-C. Wang, B. Wang, C.-C.J. Kuo, Graph representation learning: A survey, *APSIPA Trans. Signal Inf. Process.* 9 (2020) e15.
- [20] J. Zhao, K.H. Cheong, Obfuscating community structure in complex network with evolutionary divide-and-conquer strategy, *IEEE Trans. Evol. Comput.* (2023).
- [21] J. Zhao, Z. Wang, J. Cao, K.H. Cheong, A self-adaptive evolutionary deception framework for community structure, *IEEE Trans. Syst., Man, Cybern.: Syst.* (2023).
- [22] C.K. Joshi, T. Laurent, X. Bresson, An efficient graph convolutional network technique for the travelling salesman problem, 2019, arXiv preprint arXiv:1906.01227.
- [23] T.N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, in: *International Conference on Learning Representations, 2017*.
- [24] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, Y. Bengio, Graph attention networks, in: *International Conference on Learning Representations, 2018*.
- [25] K. Xu, W. Hu, J. Leskovec, S. Jegelka, How powerful are graph neural networks? 2018, arXiv preprint arXiv:1810.00826.
- [26] A. Nowak, S. Villar, A.S. Bandeira, J. Bruna, Revised note on learning quadratic assignment with graph neural networks, in: *2018 IEEE Data Science Workshop, IEEE, 2018*, pp. 1–5.
- [27] C.K. Joshi, Q. Cappart, L.-M. Rousseau, T. Laurent, Learning the travelling salesperson problem requires rethinking generalization, *Constraints* 27 (1–2) (2022) 70–98.
- [28] S. Liu, X. Yan, Y. Jin, End-to-end Pareto set prediction with graph neural networks for multi-objective facility location, in: *Proceedings of the 12th International Conference on Evolutionary Multi-Criterion Optimization, Springer, 2023*, pp. 147–161.
- [29] X. Wang, X. Yan, Y. Jin, A graph neural network with negative message passing for graph coloring, 2023, arXiv preprint arXiv:2301.11164.
- [30] H. Dai, E. Khalil, Y. Zhang, B. Dilikina, L. Song, Learning combinatorial optimization algorithms over graphs, *Adv. Neural Inf. Process. Syst.* 30 (2017).

- [31] Z. Li, Q. Chen, V. Koltun, Combinatorial optimization with graph convolutional networks and guided tree search, in: *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [32] M. Deudon, P. Cournut, A. Lacoste, Y. Adulyasak, L.-M. Rousseau, Learning heuristics for the TSP by policy gradient, in: *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 15th International Conference, CPAIOR 2018*, Springer, 2018, pp. 170–181.
- [33] W. Kool, H. van Hoof, M. Welling, Attention, learn to solve routing problems!, in: *International Conference on Learning Representations*, 2019.
- [34] V.P. Dwivedi, C.K. Joshi, T. Laurent, Y. Bengio, X. Bresson, Benchmarking graph neural networks, 2020, arXiv preprint arXiv:2003.00982.
- [35] X. Bresson, T. Laurent, Residual gated graph convnets, 2017, arXiv preprint arXiv:1711.07553.
- [36] Z.-H. Fu, K.-B. Qiu, H. Zha, Generalize a small pre-trained model to arbitrarily large tsp instances, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, (8) 2021, pp. 7474–7482.
- [37] J. Li, Y. Ma, Z. Cao, Y. Wu, W. Song, J. Zhang, Y.M. Chee, Learning feature embedding refiner for solving vehicle routing problems, *IEEE Trans. Neural Netw. Learn. Syst.* (2023).
- [38] F. Luo, X. Lin, F. Liu, Q. Zhang, Z. Wang, Neural combinatorial optimization with heavy decoder: Toward large scale generalization, in: *Thirty-Seventh Conference on Neural Information Processing Systems*, 2023.
- [39] J. Zhou, Y. Wu, W. Song, Z. Cao, J. Zhang, Towards omni-generalizable neural methods for vehicle routing problems, in: *The 40th International Conference on Machine Learning*, (ICML 2023), 2023.
- [40] C.B. Browne, E. Powley, D. Whitehouse, S.M. Lucas, P.I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, S. Colton, A survey of monte carlo tree search methods, *IEEE Trans. Comput. Intell. AI games* 4 (1) (2012) 1–43.
- [41] D. Zhang, Z. Xiao, Y. Wang, M. Song, G. Chen, Neural tsp solver with progressive distillation, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, (10) 2023, pp. 12147–12154.
- [42] L. Xin, W. Song, Z. Cao, J. Zhang, NeuroLKH: Combining deep learning model with Lin-Kernighan-Helsgaun heuristic for solving the traveling salesman problem, *Adv. Neural Inf. Process. Syst.* 34 (2021) 7472–7483.
- [43] J. Zheng, K. He, J. Zhou, Y. Jin, C.-M. Li, Combining reinforcement learning with Lin-Kernighan-Helsgaun algorithm for the traveling salesman problem, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, (14) 2021, pp. 12445–12452.
- [44] A. Colomi, M. Dorigo, F. Maffioli, V. Maniezzo, G. Righini, M. Trubian, Heuristics from nature for hard combinatorial optimization problems, *Int. Trans. Oper. Res.* 3 (1) (1996) 1–21.
- [45] H. Ye, J. Wang, Z. Cao, H. Liang, Y. Li, Deepaco: Neural-enhanced ant systems for combinatorial optimization, in: *Thirty-Seventh Conference on Neural Information Processing Systems*, 2023.
- [46] K. Helsgaun, An effective implementation of the Lin–Kernighan traveling salesman heuristic, *European J. Oper. Res.* 126 (1) (2000) 106–130.
- [47] K. Helsgaun, General k-opt submoves for the Lin–Kernighan TSP heuristic, *Math. Program. Comput.* 1 (2009) 119–163.
- [48] S. Lin, Computer solutions of the traveling salesman problem, *Bell Syst. Tech. J.* 44 (10) (1965) 2245–2269.
- [49] C.J. Watkins, P. Dayan, Q-learning, *Mach. Learn.* 8 (1992) 279–292.
- [50] R.S. Sutton, A.G. Barto, *Reinforcement learning: An introduction*, MIT Press, 2018.
- [51] X. Wu, S.-h. Wu, J. Wu, L. Feng, K.C. Tan, Evolutionary computation in the era of large language model: Survey and roadmap, 2024, arXiv preprint arXiv:2401.10034.
- [52] F. Liu, X. Tong, M. Yuan, X. Lin, F. Luo, Z. Wang, Z. Lu, Q. Zhang, An example of evolutionary computation+ large language model beating human: Design of efficient guided local search, 2024, arXiv preprint arXiv:2401.02051.
- [53] C. Wang, Z. Yu, S. McAleer, T. Yu, Y. Yang, ASP: Learn a universal neural solver!, *IEEE Trans. Pattern Anal. Mach. Intell.* (2024).
- [54] P.W. Battaglia, J.B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al., Relational inductive biases, deep learning, and graph networks, 2018, arXiv preprint arXiv:1806.01261.
- [55] J. Gilmer, S.S. Schoenholz, P.F. Riley, O. Vinyals, G.E. Dahl, Neural message passing for quantum chemistry, in: *International Conference on Machine Learning*, PMLR, 2017, pp. 1263–1272.
- [56] Q. Ma, S. Ge, D. He, D. Thaker, I. Drori, Combinatorial optimization by graph pointer networks and hierarchical reinforcement learning, 2019, arXiv preprint arXiv:1911.04936.
- [57] Y.-D. Kwon, J. Choo, B. Kim, I. Yoon, Y. Gwon, S. Min, Pomo: Policy optimization with multiple optima for reinforcement learning, *Adv. Neural Inf. Process. Syst.* 33 (2020) 21188–21198.
- [58] W. Kool, H. van Hoof, J. Gromicho, M. Welling, Deep policy dynamic programming for vehicle routing problems, in: *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 19th International Conference, CPAIOR 2022*, Springer, 2022, pp. 190–213.
- [59] M. Freitag, Y. Al-Onaizan, Beam search strategies for neural machine translation, 2017, arXiv preprint arXiv:1702.01806.
- [60] M. Huber, G.R. Raidl, Learning beam search: Utilizing machine learning to guide beam search for solving combinatorial optimization problems, in: *Machine Learning, Optimization, and Data Science: 7th International Conference*, Springer, 2022, pp. 283–298.
- [61] G.A. Croes, A method for solving traveling-salesman problems, *Oper. Res.* 6 (6) (1958) 791–812.
- [62] T.N. Kipf, M. Welling, Variational graph auto-encoders, in: *Advances in Neural Information Processing Systems*, Bayesian Deep Learning Workshop, 2016.
- [63] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, T. Eliassi-Rad, Collective classification in network data, *AI Mag.* 29 (3) (2008) 93.
- [64] H. He, E.A. Garcia, Learning from imbalanced data, *IEEE Trans. Knowl. Data Eng.* 21 (9) (2009) 1263–1284.
- [65] D.L. Applegate, R.E. Bixby, V. Chvátal, W.J. Cook, *The Traveling Salesman Problem: A Computational Study*, 2007.
- [66] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, *Neural Netw.* 2 (5) (1989) 359–366.
- [67] D. Marcheggiani, I. Titov, Encoding sentences with graph convolutional networks for semantic role labeling, in: *EMNLP 2017: Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 2017, pp. 1506–1515.