



Road Planning for Slums via Deep Reinforcement Learning

Yu Zheng*

Hongyuan Su*

Department of Electronic Engineering,
BNRist, Tsinghua University
Beijing, China

Depeng Jin

Department of Electronic Engineering,
BNRist, Tsinghua University
Beijing, China

Jingtao Ding

Department of Electronic Engineering,
BNRist, Tsinghua University
Beijing, China

Yong Li†

Department of Electronic Engineering,
BNRist, Tsinghua University
Beijing, China

ABSTRACT

Millions of slum dwellers suffer from poor accessibility to urban services due to inadequate road infrastructure within slums, and road planning for slums is critical to the sustainable development of cities. Existing re-blocking or heuristic methods are either time-consuming which cannot generalize to different slums, or yield sub-optimal road plans in terms of accessibility and construction costs. In this paper, we present a deep reinforcement learning based approach to automatically layout roads for slums. We propose a generic graph model to capture the topological structure of a slum, and devise a novel graph neural network to select locations for the planned roads. Through masked policy optimization, our model can generate road plans that connect places in a slum at minimal construction costs. Extensive experiments on real-world slums in different countries verify the effectiveness of our model, which can significantly improve accessibility by 14.3% against existing baseline methods. Further investigations on transferring across different tasks demonstrate that our model can master road planning skills in simple scenarios and adapt them to much more complicated ones, indicating the potential of applying our model in real-world slum upgrading. The code and data are available at <https://github.com/tsinghua-fib-lab/road-planning-for-slums>.

CCS CONCEPTS

• **Computing methodologies** → **Planning and scheduling**; **Reinforcement learning**.

KEYWORDS

road planning, slum upgrading, reinforcement learning

ACM Reference Format:

Yu Zheng, Hongyuan Su, Jingtao Ding, Depeng Jin, and Yong Li. 2023. Road Planning for Slums via Deep Reinforcement Learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*

*Both authors contributed equally to this research.

†Corresponding author (liyong07@tsinghua.edu.cn).



This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '23, August 6–10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0103-0/23/08.

<https://doi.org/10.1145/3580305.3599901>

(KDD '23), August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3580305.3599901>

1 INTRODUCTION

With rapid urbanization, currently about 4 billion people around the world live in cities, while 1 billion of them live in over 200,000 slums [49, 60]. The vast majority of slums suffer from poor accessibility, with internal places not connected to external road systems, and many places not even having addresses [10, 24]. Besides being unreachable by motor vehicles, urban services depending on road systems, such as piped services of water and sanitation buried under roads, cannot be delivered to places in slums, which leads to severe problems in public health, urban environment, *etc* [49]. To tackle these problems, local upgrading of slums has become the primary approach for the sustainable development of cities, rather than moving all the people to cities, due to the massive number of slum dwellers and the socio-economic costs [4, 24, 43, 59]. Particularly, improving the accessibility by planning roads plays an essential role in slum upgrading [9, 50].

Different from city-level road planning which grows a road network from the top down and arranges land functionalities accordingly [15], road planning for slums is a bottom-up process in which existing houses determine the possible forms of the road network [50]. Therefore, current city-level approaches cannot handle the *micro-level* road planning within a slum. Meanwhile, road planning for slums is challenging due to its large solution space. Take a moderate-size slum as an example, the solution space of planning 40 road segments from 80 candidate locations surpasses 10^{23} , which is too large for exhaustive enumeration. In practical slum upgrading, re-blocking [24, 37] strategy is adopted. It involves negotiations with multiple stakeholders and usually takes a long time for a specific case, thus it can not generalize globally to different slums. Given the enormous number of slums, it is necessary to develop a computational method that can automatically accomplish road plans with superior connectivity at minimal construction costs [9, 50]. Such a model can significantly benefit slum upgrading and eventually help achieve *cities without slums* [8, 9, 24, 59].

One pioneering work by Brelsford *et al.* [9] formulates road planning for slums as a constrained optimization problem, and proposes a heuristic search method to generate road plans. It makes this problem computationally solvable and has been adopted for slums in South Africa and India. Although the heuristic can be

applied to different slums, we empirically show that the quality of the obtained plans is not guaranteed, with the accessibility and construction costs far from optimal. Fortunately, with the rapid development of artificial intelligence (AI), it is promising to leverage AI to solve the problem of road planning for slums. First, data-driven parametric models have strong generalization ability, which can adapt to different scenarios [27, 57, 63]. In addition, AI models, especially deep reinforcement learning (DRL) algorithms, are good at searching in a large action space to optimize various objectives. The action space can be effectively eliminated by predicting rewards with a value network and sampling actions via a policy network [23, 30, 38, 47]. Particularly, DRL has been deployed in similar planning tasks, such as solving the vehicle routing problem [12, 41, 65] and designing circuit chips [3, 36, 44].

Inspired by the success of DRL, we propose a DRL-based method to solve this significant real-world problem, road planning for slums. Since slums are diverse in the original geometric space, *e.g.*, existing houses and paths can be in various irregular shapes, we propose a generic graph model to describe a slum, solving the problem from topology instead of geometry. The topology invariance of the graph model makes our method capable of generalizing to different slums of arbitrary forms. We further develop a policy network to select road locations and a value network to predict the performance of road planning based on a novel graph neural network (GNN), overcoming the difficulty of efficient search in the huge action space. We design a topology-aware message passing mechanism for GNN, which first gathers various topological information to edges from nodes, faces, and edges themselves, then broadcasts edge embeddings back to learn effective representations of roads and places in the slum. Furthermore, we develop a masked policy optimization method and connectivity-priority reward functions to optimize various objectives, including accessibility, travel distance, and construction costs. We conduct experiments on real-world slums to verify the effectiveness of our proposed model.

To summarize, the contributions of this paper are as follows,

- We formulate road planning for slums as a sequential decision-making problem, and propose a DRL-based solution.
- We develop a novel GNN and a multi-objective optimization method based on a generic graph model for slums. The proposed model can learn effective representations of places and roads in a slum, which enables superior road planning policy.
- We conduct extensive experiments on slums in different countries, and the results demonstrate the advantage of our proposed method against baseline methods. Our model can generate road plans with both higher accessibility and lower construction costs. Moreover, we also show the transferability of our model from small slums to large slums, indicating the potential of applying our method in real-world slum upgrading.

2 PROBLEM STATEMENT

From the perspective of connectivity, a slum can be decomposed into two categories of elements, *places* and *roads* [9]. Specifically, places are the houses and internal facilities of the slum, and roads are the street system that connects various external urban services. In most slums, a large fraction of places are disconnected from

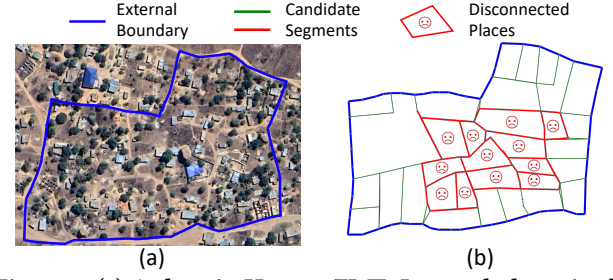


Figure 1: (a) A slum in Harare, ZWE. Internal places in the slum are not connected to the external road system, making urban services inaccessible to slum dwellers. (b) Geometric description of the slum. Red polygons are places disconnected to roads, and internal segments (green and red) are candidate locations for new roads. Best viewed in color.

roads, as shown in Figure 1. Such poor connectivity makes basic urban services inaccessible, *e.g.*, ambulances and fire fighting trucks cannot reach the disconnected places during emergencies; water and sanitation pipes buried under roads cannot be provided. Therefore, it is crucial to upgrade slums by planning more roads. To deliver basic urban services, a minimal road network needs to make all places directly adjacent to roads, which is called universal connectivity [9]. Besides the minimally necessary accesses, more roads are expected to promote internal transportation and reduce travel distance for slum dwellers. To minimize disruption to the slums, new roads are not allowed to pass through the middle of places, thus the candidate locations are restricted to the spacing between places. It is worth noting that each planned road segment also has a corresponding construction cost.

As illustrated in Figure 1(b), to describe the problem in geometric terms, a slum is a two-dimensional planar surface U whose exterior boundaries E are existing roads. The surface is filled by a tessellation of faces (polygons) P , where each polygon p_i is a place in the slum¹. Polygon boundaries in the interior of the surface represent the spacing between places, which form a collection of segments S and serve as the candidate locations for new roads. Road planning is to select a subset of these segments for construction as roads. Therefore, it can be formulated as follows:

Input: A planar surface U with exterior boundaries E for the slum, a collection of polygons P for places in the slum, a collection of segments S with their corresponding cost C for road construction, and the road planning budget K .

Output: A subset R of size K from S for construction as roads.

Objective: (1) Connecting all polygons in P to the road system $E \cup R$. (2) Minimizing the travel distance between any pair of polygons p_i and p_j over the road network $E \cup R$. (3) Minimizing the overall construction cost for the road plan $\sum_{i \in R} C_i$.

3 METHOD

3.1 Overall Framework

We formulate the road planning for slums as a sequential decision-making problem (see Section A.1 of the appendix for specific definitions of the Markov Decision Process (MDP)). As illustrated in

¹A planar surface is a graph which can be drawn in the plane without any edges crossing. When a planar graph is drawn with no crossing edges, it divides the plane into a set of regions, called faces.




Figure 2: Schematic of our approach. At each step, the agent receives states and rewards from the environment and outputs the road locations for the slum. Best viewed in color.

Figure 3a, given the planning budget, which is the total number of road segments, a road plan is accomplished through a sequence of location selection decisions, where at each step of the sequence, one new road segment is planned at a specific location. The goal of the sequential decision-making problem is to improve the connectivity and accessibility of the slum at minimal costs.

As shown in Figure 2, we develop an agent with a policy network and a value network to take actions and predict returns, respectively, and a shared GNN model as the state encoder. To address the challenge of geometrical diversity, we tackle road planning for slums at the level of topology instead of geometry with a generic graph model (Section 3.2). We then propose a novel GNN model to achieve a decent location selection policy on the graph (Section 3.3). In order to overcome the difficulty of multi-objective optimization in road planning, we further develop a masked policy optimization method with connectivity-priority reward functions (Section 3.4).

3.2 Graph Model

It is challenging to plan roads directly at the geometric level, since slums are very diverse in the original geometric space, *e.g.*, the polygons of places can be in various irregular shapes, and the segments can intersect at almost any angle. In addition, the spatial relationship between different geometries is more important for road planning than the specific shapes of geometries. In contrast to the diverse geometries, there exists certain invariance in the topology of places and roads in cities [9, 62], which can support the uniform modeling of different slums. Therefore, we solve the road planning problem from the topological viewpoint instead of the geometric one. Specifically, we construct a planar graph to represent a slum with the contained places and roads, transforming the geometries into elements on the graph, such as nodes, edges, and faces. In this way, we develop a generic graph model which can handle slums of arbitrary geometric forms at different scales with the same logic, solving the challenge of geometrical diversity.

The planar graph is constructed based on the original geometrical descriptions of the slum, including the surface, polygons, and segments. As shown in Figure 4(a), vertices and boundary segments of polygons become nodes and edges on the graph, respectively. Meanwhile, the original polygons naturally become faces surrounded by edges on the planar graph, where each face in the graph represents a place which is usually a house in the slum. Each edge has a *road* attribute indicating whether it is a road segment or not, and a *road*

segment can be either an existing external road or a planned new road. Moreover, we preprocess the transformed planar graph of the slum to remove redundant information, as illustrated in Figure 4(b-c). First, we merge multiple nodes/edges within a threshold distance as one node/edge, since they are supposed to share the same accessibility in the real space. We then delete nodes with degree 2 and merge the corresponding two edges (construction costs are added) to simplify the graph, which have no influence on road planning. Finally, we normalize the length of edges and align the coordinates, in order to support slums in different scales.

With the above generic graph model, road planning for slums is transformed into a sequential decision-making problem on a dynamic graph. Specifically, states are the information of the current graph, and actions for a road planning policy are edge selections on the graph. The graph also transits accordingly, *i.e.*, the *road* attribute of the selected edge changes from *False* to *True*, which in turn leads to subsequent changes in accessibility and travel distance of the slum. For example, with the newly planned road, some faces (places) are connected to the road system, and the travel distance between several faces is reduced. These changes are also reflected in the reward, which can be directly computed from the graph itself.

3.3 Planning with Graph Neural Networks

With the generic graph model of slums, we now introduce our proposed GNN model which performs road planning on the dynamic graph. As the task is to select edges, a policy needs to decide the probability of choosing different edges at each step. Since the topological information is critical to the effect of road planning, when computing the selection probability of each edge, it is necessary to consider its neighbors and even the whole graph, such as the travel distance of its neighboring faces. Thus, we adopt GNN in our policy because of its strong ability to extract topological information and fuse neighborhood features. As shown in Figure 2, we develop a GNN state encoder, which plays a fundamental role in the road planning agent. The learned representations from GNN are shared between the policy network and the value network, serving as the basis for policy making and return prediction.

To address the challenge of complex topology in road planning, we propose a novel GNN model which takes nodes, edges and faces into consideration. Figure 3b demonstrates our proposed road planning policy based on GNN. We first design rich features regarding accessibility, travel distance, and construction costs as the input of GNN. We then design a topology-aware message passing mechanism to learn effective representations of topological elements on the graph. Finally, we utilize an edge-ranking policy network to score edges based on the learned edge embeddings, supporting edge selection on the graph.

Input Features for Topological Elements. Topological features reflect the current state of road planning, serving as the original input for GNN to learn representations of topological elements. As illustrated in Table 1, we incorporate rich information about road planning into the designed features for nodes, edges, and faces. Specifically, there are static features that do not change with the actions of the agent, such as the coordinates and construction cost, while most of the features are dynamic and alter according to actions at each step. These meaningful features describe the current

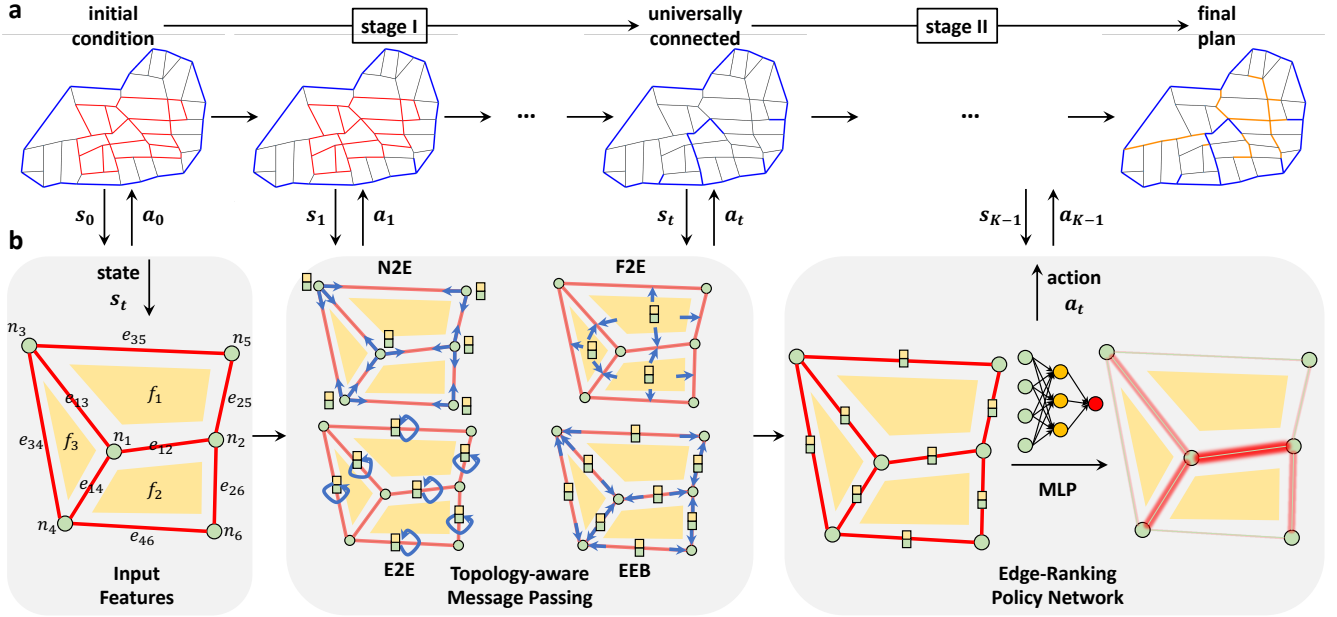


Figure 3: (a) Road planning for slums as a sequential decision-making problem, where one single road segment is planned at each step. In stage I, the agent plan roads (blue) to achieve universal connectivity, i.e., all disconnected places are connected to the road system. In stage II, the agent add road segments (orange) to reduce travel distance. (b) The proposed GNN model. We design rich features for nodes, edges and faces. Topology-aware message passing is proposed, which contains Node2Edge Propagation (N2E), Face2Edge Propagation (F2E), Edge Self-Propagation (E2E) and Edge Embedding Broadcast (EEB). Finally, an edge-ranking policy network is developed to sample actions of edge selection. Best viewed in color.




Figure 4: (a) The constructed graph transformed from the original geometrical descriptions of the slum. Faces are polygons (places). Nodes are vertices of polygons. Edges are polygon boundary segments. (b) We simplify the graph by merging nearby edges (top) and removing nodes with degree 2 (bottom). (c) The graph after preprocessing. Road planning solutions on the simplified graph can be easily mapped back to the original graph. Best viewed in color.

accessibility and travel distance of various places in the slum, which helps to decide which edges to plan as road segments. For example, *Connected* means whether a face is connected to the road system, thus building a road to an unconnected face can significantly improve the accessibility of the corresponding place to external urban services. Similarly, *Straightness* is the ratio of road network distance to the Euclidean distance of an edge, which directly indicates the travel distance between two places, and therefore selecting edges with high *Straightness* can substantially reduce long detours in the slum. These features support effective representation learning and subsequent decision-making, and details of all the designed topological features are introduced in Section A.2 of the appendix.

Table 1: Designed features for topological elements.

Topology	Feature	Dimension	Type
Node	Coordinates	2	Static
	Centrality	4	Static
	On Road	1	Dynamic
	Road Ratio	1	Dynamic
	Avg N2N Dis	1	Dynamic
Edge	Cost	1	Static
	Road	1	Dynamic
	Straightness	1	Dynamic
Face	Connected	1	Dynamic
	Avg F2F Dis	1	Dynamic
	F2E Dis	1	Dynamic

Topology-aware Message Passing. Since the policy selects edges on the graph to plan roads, we propose an edge-centric GNN to learn representations. We first encode the input topological features to dense embeddings with separate weight matrices as follows,

$$n_i^{(0)} = W_n^{(0)} A_{n_i}, \quad e_{ij}^{(0)} = W_e^{(0)} A_{e_{ij}}, \quad f_i^{(0)} = W_f^{(0)} A_{f_i}, \quad (1)$$

where A_{n_i} , $A_{e_{ij}}$ and A_{f_i} are input attributes for nodes, edges and faces, W_n , W_e and W_f are learnable embedding matrices.

To address the challenge of complex topological elements, we design a topology-aware message passing mechanism, which first *pulls* information from diverse topological elements into edges through *node-to-edge propagation*, *face-to-edge propagation*, and

edge self-propagation, and then *pushes* aggregated topological information back through *edge embedding broadcast*, as shown in Figure 3b. The edge embeddings are obtained as follows.

Node2Edge Propagation. For each edge, we take the embeddings of its two connected nodes and propagate them through a linear transformation layer and a non-linear activation layer. The node-to-edge message is computed as follows,

$$e_{ij,n \rightarrow e}^{(l+1)} = \tanh(W_{n \rightarrow e}^{(l+1)} (n_i^{(l)} \| n_j^{(l)})), \quad (2)$$

where $\|$ means concatenation, and $W_{n \rightarrow e}$ is a transformation layer.

Face2Edge Propagation. For each edge, we propagate the embeddings of its adjacent faces, and the face-to-edge message is computed as follows,

$$e_{ij,f \rightarrow e}^{(l+1)} = \tanh\left(\frac{1}{N_{ij}^f} \sum_{k \in F_{ij}} W_{f \rightarrow e}^{(l+1)} f_k^{(0)}\right), \quad (3)$$

where N_{ij}^f is the number of elements in F_{ij} , the set of adjacent faces for edge e_{ij} , and $W_{f \rightarrow e}$ is a linear transformation layer.

Edge Self-Propagation. Since each edge has its own attributes, we further include the propagation message from the edge itself, which is computed as follows,

$$e_{ij,e \rightarrow e}^{(l+1)} = \tanh(W_{e \rightarrow e}^{(l+1)} e_{ij}^{(0)}), \quad (4)$$

where a linear transformation matrix $W_{e \rightarrow e}$ is adopted.

The edge embedding is obtained by integrating the above three propagated messages as follows,

$$e_{ij}^{(l+1)} = \tanh(W_e^{(l+1)} (e_{ij,n \rightarrow e}^{(l+1)} \| e_{ij,f \rightarrow e}^{(l+1)} \| e_{ij,e \rightarrow e}^{(l+1)})), \quad (5)$$

where the three messages are concatenated and transformed with a linear layer $W_e^{(l+1)}$.

Edge Embedding Broadcast. We then *push* the edge embeddings back to nodes to update their embeddings as follows,

$$n_{i,e \rightarrow n}^{(l+1)} = \frac{1}{N_i} \sum_{j \in \mathcal{N}_i} e_{ij}^{(l+1)}, \quad (6)$$

$$n_i^{(l+1)} = n_i^{(l)} + n_{i,e \rightarrow n}^{(l+1)}, \quad (7)$$

where for each node, we average the embeddings of its connected edges and add it to the node embedding.

By stacking multiple layers of the above topology-aware message passing, each node or edge can exchange information with neighbors on the graph. We use the obtained embeddings at the last layer, $e_{ij}^{(L)}$ and $n_i^{(L)}$, as the final representations, where L is a hyper-parameter in our model. Through topology-aware message passing, the obtained edge representations can well capture the information about accessibility, travel distance, and construction costs of places and roads from its neighbors, which can effectively support the road planning policy.

Edge-ranking Policy Network. The policy must generate the probability of selecting different edges at each step. Therefore, we develop an edge-ranking policy network to score each edge, based on the obtained edge embeddings from GNN. The score is calculated with a multi-layer perceptron (MLP) as follows,

$$s_{ij} = \text{MLP}_p(e_{ij}^{(L)}). \quad (8)$$

The action of edge selection is sampled from a probability distribution over different edges according to their corresponding scores s_{ij} estimated by the policy network. Since the obtained edge embeddings contain rich topological information, the road planning action made by the policy network takes into account the accessibility, travel distance, and construction cost of the slum.

3.4 Multi-objective Policy Optimization

Among the three objectives, accessibility, *i.e.*, achieving universal connectivity for all places in the slum, is crucial for residents in the slum to access basic urban services, which is the primary target of road planning. Therefore, it is necessary to prioritize connectivity when optimizing the policy, and further reduce travel distance after universal connectivity is achieved. Meanwhile, for both connectivity and travel distance, it is desirable to optimize them at minimal construction cost. Towards this end, we propose a masked policy optimization method and connectivity-priority reward functions with two stages, as shown in Figure 3a. The optimization method encourages the policy to achieve universal connectivity in stage I, then reduce travel distance in stage II, preferring low construction cost in the whole process.

Stage I. The goal of this stage is to achieve universal connectivity as quickly as possible, making all places in the slum connected to the road system and accessible to urban services. Therefore, each new planned road segment is expected to connect more faces (places) that are not yet connected to any road segments. Meanwhile, since road planning is a gradual extension of the existing road system, a new road segment can not be created as a separate component without touching the already planned roads. We thus design an action mask to indicate feasible actions in this stage for the policy network, and the mask value of each edge is calculated as follows,

$$m_{ij} = \mathbb{1}[\text{On_Road}(n_i)] \wedge \mathbb{1}\left[\sum_{f \in F_{n_j}} (1 - \text{Connected}(f)) \geq 1\right], \quad (9)$$

where F_{n_j} represents all the faces that contain the node n_j . In other words, the action mask requires the selected edge to start from a *road node* and connect to at least one *unconnected face*. The mask value is multiplied over the obtained scores from the policy network in (8), which serves as the selection probability of different edges,

$$\text{Prob}(e_{ij}) = \frac{e^{s_{ij}}}{\sum_{(u,v) \in \mathcal{E}} e^{s_{uv}}} * m_{ij}, \quad (10)$$

where \mathcal{E} denotes all the edges on the graph. With the action mask, only those edges that start from existing roads and connect disconnected faces will be considered by the policy.

Besides the action mask, we also design a corresponding reward function in this stage, which is a weighted sum of the number of newly connected faces and the construction cost of the planned road. Given the action a_k at the k -th step selecting the edge e_{ij} , the reward is calculated as follows,

$$r_k = \alpha_1 \sum_{f \in F_{n_j}} (1 - \text{Connected}(f)) + \alpha_2 C_{e_{ij}}, \quad (11)$$

where $C_{e_{ij}}$ is the construction cost of the road segment specified by e_{ij} , and α_1 and α_2 are hyper-parameters in our model.




Figure 5: (a) In stage I, the planned roads (blue segments) connect all disconnected places. However, many places, even nearby places, still suffer from high travel distance. In this example, place A and B are next to each other, while it requires a long detour (green path) between them by vehicle. (b) In stage II, roads are planned (orange segments) to reduce travel distance. Now the trip from place A to B (green path) is much shorter. Best viewed in color.

Stage II. As shown in Figure 5(a), after the slum becomes universally connected in stage I, the generated road network looks like a tree with many dead-ends, which is undesirable in reality [2, 5]. Meanwhile, the traffic between some places is still poor and requires long detours, even for some nearby places. Therefore, stage II aims to add more roads to reduce travel distance within the slum, as shown in Figure 5(b). We still require the planned road segments to start from existing road nodes, and the mask value of different edges are calculated as follows,

$$m_{ij} = \mathbb{1}[\text{On_Road}(n_i)]. \quad (12)$$

The action probability is obtained in the same way as (10).

For the reward function given action a_k selecting edge e_{ij} , we compute the pairwise travel distance reduction of the slum, and combine it with construction cost,

$$D(k) = \frac{|F|(|F| - 1)}{2} \sum_{u=1}^{|F|-1} \sum_{v=u+1}^{|F|} d(f_u, f_v; k), \quad (13)$$

$$r_k = \alpha_1(D(k) - D(k+1)) + \alpha_2 C_{e_{ij}}, \quad (14)$$

where $d(f_u, f_v; k)$ denotes the travel distance between two faces, f_u and f_v , over the road network at the k -th step.

With the designed action mask and reward functions, the policy is guided to connect unconnected faces and reduce travel distance with low construction costs in the two stages, respectively.

Value Network and Optimization. Besides the policy network, we follow the actor-critic manner [30] and develop a value network to predict the effect of road planning. Since places and roads are captured with a graph, we compute graph-level representations to summarize the current state of the whole slum. Specifically, we take the average of all the node embeddings and edges embeddings, and also include a one-hot encoding of the stage as follows,

$$n_{avg} = \frac{1}{|\mathcal{N}|} \sum_{i=1}^{|\mathcal{N}|} n_i^{(L)}, e_{avg} = \frac{1}{|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} e_{ij}^{(L)}, \quad (15)$$

$$h_g = n_{avg} \| e_{avg} \| \text{one-hot}(stage), \quad (16)$$

where \mathcal{N} and \mathcal{E} are the sets of nodes and edges, and h_g is the graph representation. We utilize an MLP model to predict the return,

$$\hat{r} = \text{MLP}_v(h_g). \quad (17)$$

Table 2: Basic information of experimented slums. D.R. means the ratio of disconnected places.

Location	Place	Segment	D.R.	Solution
Harare, ZWE	32	85	37.5%	4×10^{20}
Cape Town, ZAF	34	91	44.1%	6×10^{25}
Cape Town, ZAF	59	164	59.3%	5×10^{47}
Mumbai, IND	92	208	58.7%	1×10^{60}

Finally, we adopt Proximal Policy Optimization (PPO) [45] to update the parameters of the policy network and value network, which encourages the agent to conduct safe and efficient exploration in the action space.

4 EXPERIMENTS

4.1 Experiment Settings

Slum Data. We conduct experiments on slums of different scales from different countries with publicly released data [9]. Table 2 shows the basic information of these slums, where we list the number of places and segments, as well as the size of the solution space. Notably, all the slums suffer from poor accessibility, with over 40% of places disconnected from road systems. More details of the data are introduced in Section B of the appendix.

Baselines. We compare our model with the following methods.

- **Random.** This method selects road segments randomly.
- **Greedy.** This method selects new road segments greedily according to accessibility (Greedy-A) and construction cost (Greedy-C).
- **Masked.** We add our proposed action mask to Random and Greedy baselines. Masked baselines select road segments that are True in the mask randomly (greedily).
- **Minimum Spanning Tree (MST).** A graph is built where nodes represent slums, edges represent road segments and edge weights represent road construction costs. We use Kruskal’s algorithm [31] to grow a minimum spanning tree.
- **Genetic Algorithm (GA)** [20]. This type of method is widely adopted in road planning. We include a generative version (GA-G) that adopts a linear layer as genes and builds one road at one step by multiplying edge features with a linear layer as sampling probability. We also include a swap version (GA-S) that directly uses the selection of road segments as genes and performs swapping between different solutions at each step.
- **Heuristic Search (HS-MC)** [9]. This recently proposed method formulates road planning for slums as a constrained optimization problem. It samples paths from external boundary roads to unconnected places using the Monte Carlo techniques [6].
- **DRL-MLP.** We implement a simplified DRL model by replacing the proposed GNN with an MLP, thus it ignores topological information when planning roads.

It is worthwhile to notice that Greedy-A, MST, GA, HS-MC and our DRL models are all with action masks themselves. We also include two generative models [14, 28], based on Generative Adversarial Networks (GAN) [22] and Variational Auto-Encoder (VAE) [29], though manual adjustments are required for these methods.

Table 3: Road planning performance comparison. Lower is better. F and INF means failing to achieve universal connectivity.

Method	Harare, ZWE			Cape Town, ZAF (A)			Cape Town, ZAF (B)			Mumbai, IND		
	NR	AD	SC	NR	AD	SC	NR	AD	SC	NR	AD	SC
Random	29	1.06	6.30	F	INF	10.83	F	INF	20.76	F	INF	26.05
Random (masked)	10	1.00	6.13	14	1.62	10.37	54	2.77	19.95	42	2.50	24.92
Greedy-A (masked)*	8*	0.63	5.04	13*	1.12	10.42	28*	1.66	18.91	29*	1.77	25.42
Greedy-C*	20	0.84	3.85*	35	1.83	7.03*	F	INF	14.10*	F	INF	19.45*
Greedy-C (masked)*	11	0.84	3.85*	14	1.81	7.23*	35	2.22	14.29*	45	2.81	19.28*
GAN (manually adjusted)	-	0.70	5.71	-	1.33	9.52	-	2.05	17.72	-	1.72	24.34
VAE (manually adjusted)	-	0.71	5.14	-	1.31	10.70	-	2.06	17.58	-	1.68	23.84
MST (masked)	<u>11</u>	0.59	5.57	<u>14</u>	1.17	8.75	35	<u>1.54</u>	17.16	45	1.63	<u>22.92</u>
GA-G (masked)	<u>11</u>	0.58	<u>4.60</u>	<u>14</u>	1.14	8.72	34	1.99	18.95	42	1.87	24.26
GA-S (masked)	-	0.58	5.25	-	1.21	8.44	-	1.89	17.72	-	1.88	23.22
HS-MC (masked)	13	0.62	5.31	16	1.09	9.09	37	1.55	16.98	43	1.61	23.00
DRL-MLP (ours, masked)	<u>11</u>	0.52	4.38	<u>14</u>	0.96	<u>8.28</u>	32	1.57	<u>15.66</u>	31	1.52	22.93
DRL-GNN (ours, masked)	9	0.50	<u>4.60</u>	13	0.93	8.24	31	1.51	15.62	29	1.51	22.82
impr% v.s. HS-MC	-25.0%	-19.4%	-17.5%	-18.8%	-14.7%	9.8%	-16.2%	-2.6%	-8.0%	-32.6%	-6.21%	-0.8%
Build All Roads	-	0.47	11.50	-	0.80	19.82	-	1.21	37.55	-	1.36	49.25

* Although they are equal to or even smaller than the **bolded** numbers, these methods exhibit imbalanced results with much worse performance on the other two metrics. Thus, the **bolded** and underlined numbers are assigned to **the lowest** and the second lowest values, *excluding greedy methods*.

Evaluation Metrics. As introduced in Section 2, we evaluate a road plan concerning accessibility, travel distance, and construction cost. The specific definitions are as follows,

- For accessibility, it is desired to achieve universal connectivity as early as possible, thus we calculate the number of road segments (NR) consumed to achieve universal connectivity.
- For travel distance, we compute the average distance (AD) between any pair of places in the slum over the road network.
- We define the construction cost of each road segment as its length, and calculate the sum of costs (SC) of all planned roads.

It is worth noting that all the metrics are *the lower the better*.

Model Implementation. We implement the proposed model with PyTorch [42], and all the codes and data to reproduce the results in this paper are released at <https://github.com/tsinghua-fib-lab/road-planning-for-slums>. We implement the greedy and GA baselines and integrate them into our framework. For the heuristic search baseline, we use the codes released in [9]. We carefully tune the hyper-parameters of our model, including learning rate, regularization, *etc.* For each road planning task, we collect millions of samples and train our model on a single server with an Nvidia GeForce 2080Ti GPU, which usually takes about 2 hours. A full list of hyper-parameters is provided in Section C of the appendix.

4.2 Performance Comparison

We set the planning budget (episode length) as 50% of the number of candidate segments. Results of our model and baselines are illustrated in Table 3, where we also include a reference model (*Build All Roads*) that sets 100% of candidate segments as roads. NR is not applicable to GA-S since it is not a generative method. From the results, we have the following observations,

- **Random and greedy algorithms are ineffective for road planning.** Randomly choosing locations fails to achieve universal connectivity in all slums except for the smallest one. Greedy-C

achieves the lowest construction cost for all four slums, while it fails to achieve universal connectivity in the two largest slums. Adding action masks can help these methods to achieve universal connectivity, however, the travel distance is still the worst. Similarly, Greedy-A is the earliest to achieve universal connectivity, however, the construction cost is the worst, and the travel distance is also much worse than other methods. Thus we do not consider these trivial methods in the following comparisons.

- **Generative models are not suitable to road planning for slums.** To obtain road plans for slums with the two generative models [14, 28], much of the work has to be conducted manually by human labor, which betrayed our original intention to automate the process of road planning. Not surprisingly, since they are not suitable to the sequential decision-making task, the performance of GAN and VAE falls far behind our proposed method and the HS-MC baseline.
- **DRL-based methods have significant advantages over other approaches.** DRL-MLP and DRL-GNN outperform GA-G, GA-S, and HS-MC on all metrics. The two DRL-based methods achieve much better road planning performance, with average reductions of about 23.2%, 10.7%, and 9.0% in NR, AD, and SC over the four slums. Baselines like GA and HS-MC fail to explore the solution space efficiently, making it difficult to obtain high-quality road plans. The performance gap verifies the strong ability of DRL to optimize multiple objectives in a large action space.
- **Our proposed model achieves the best performance.** Regarding accessibility, our model is the fastest to achieve universal connectivity for all slums, which is critical under tight planning budgets. Compared with HS-MC, our model connects all places with 3 fewer road segments (NR) for slums in Harare and Cape Town, and 14 fewer road segments in the largest slum in Mumbai, IND. Meanwhile, with respect to travel distance and construction cost, our model consistently outperforms baseline methods. Specifically, our method reduces AD by 19.4% and 14.7% for slums




Figure 6: The generated road plans for the slum in Cape Town, ZAF, and their corresponding travel distance matrices of (a) GA (b) HS-MC (c) DRL-GNN. Best viewed in color.




Figure 7: Performance of DRL-GNN and its variants that remove different features, including Centrality (F1), Road (F2) and Straightness (F3) for slums in (a) Harare, ZWE (b) Cape Town, ZAF. Best viewed in color.

in Harare and Cape Town, respectively, and reduces SC by 11.1% for slums in Cape Town. In particular, the road plan obtained by our method achieves a travel distance very close to that of *Build All Roads* at a much lower cost, making it more economical in real slum upgrading. Our model can capture topological information through the generic graph model and the novel GNN, and perform efficient searches in the large action space via masked policy optimization. These special designs enable our model to achieve superior performance in road planning for slums.

Figure 6 demonstrates the generated road plans of different models for the slum in Cape Town, ZAF, and their corresponding travel distance matrices. Although universal connectivity is achieved in all plans, the travel distance varies significantly across different methods. In the road plans of baselines, slum dwellers in some places have to travel a long detour to reach each other, which corresponds to several *hot* regions in the travel distance matrices as shown in Figure 6(a-b). In contrast, our method utilizes the progress in travel distance as the reward and optimizes it in stage II. Specifically, the proposed GNN model can detect places that suffer from long detours through topology-aware message passing on the graph, and add targeted roads to reduce travel distance effectively. Thus there are much fewer *hot* regions as shown in Figure 6(c). In addition, as demonstrated in Figure 6, our model is able to grow a road network in a less costly way, with the total length of planned roads much shorter than baselines by about 10%.




Figure 8: Performance of DRL-GNN and its variants that remove node2edge propagation (N2E), edge self-propagation (E2E) and face2edge propagation (F2E) for slums in (a) Harare, ZWE (b) Cape Town, ZAF. Best viewed in color.

4.3 Ablation Study

Graph Modeling. The spatial topological relationships between places and roads in a slum are crucial for road planning. The proposed graph modeling and GNN can capture such topological relationships, enabling decent location selection policies. Table 3 illustrates the performance of our method with and without graph modeling, *i.e.*, DRL-GNN and DRL-MLP. Specifically, it is easier for our graph model to perceive the currently disconnected regions, and layout corresponding road segments to connect them, leading to earlier universal connectivity in all 4 slums. The graph model can also capture the neighborhood information on travel distance and construction costs, leading to a more economical policy to reduce travel distance. As shown in Table 3, DRL-GNN outperforms DRL-MLP in AD and SC for 4 and 3 slums, respectively.

Topological Features. We investigate the role that the designed features for nodes, edges, and faces play in our model. We first obtain a well-trained model, then remove different features, *i.e.*, setting the feature values as 0, and evaluate its performance. Figure 7 demonstrate the performance of removing three features (F1: Centrality, F2: Road, F3: Straightness) compared with using all features. We can observe that feature Straightness brings the largest performance deterioration in travel distance, with 22.0% and 49.0% increases in Harare and Cape Town, respectively. This result is reasonable since *Straightness* is the ratio of road network distance to the Euclidean distance, which directly indicates long detours in the slum, thus this feature is critical to travel distance. In addition, feature *Road* also plays an important role, and removing it leads to a 20.2% and 13.0% increase in construction cost for the two slums, respectively. Our designed rich features describe the topological information of the slum, which is critical when selecting locations for new road segments.

Topology-aware Message Passing. In the proposed GNN model, we design various propagation messages to edges from different sources, including nodes, faces, and edges themselves. In this section, we study the effect of different propagation messages. Specifically, we design multiple variants of our GNN model, each of which blocks one single propagation message. We train these models and evaluate their road planning performance, as shown in Figure 8. We can find that deleting any propagation flow leads to the loss of topological information, and brings about a deterioration in performance. Specifically, deleting Node2Edge propagation makes travel distance worse by 10.2% in Harare and construction cost worse by 8.1% in Cape Town; deleting Edge Self-propagation increases

construction cost by 14.8% in Cape Town; and deleting Face2Edge propagation leads to a 3.0% increase in travel distance in Harare. The above results confirm the necessity of topology-aware message passing, which gathers diverse topological information to edges and makes our edge-centric GNN learn meaningful edge representations, enabling decent edge selection policies.

4.4 Analysis on Transferability

It is beneficial for a road planning model to generalize across different scenarios. On the one hand, the planning budgets may vary. We set the budget as 50% of candidate segments for training, and directly evaluate our model under different budgets. Figure 9(a) shows that DRL-GNN outperforms HS-MC under all different budgets, with more significant improvements under tight budgets, e.g., 10.3% travel distance reduction under 30% budgets, two times larger than 70% budgets. On the other hand, we study the transferability across different slums. We obtain a pretrained model on a small slum (Harare, ZWE), and finetune it on a large slum (Cape Town, ZAF). We compare the pretrained model with a model that is trained from scratch. Figure 9(b) demonstrates the travel distance at each step for the large slum, where the pretrained model is consistently better than the model trained from scratch in stage II. The above results verify that our model can learn universal road planning skills and successfully transfer them to scenarios of different budgets or different slums, which is crucial for practical slum upgrading.

5 RELATED WORK

Deep Reinforcement Learning for Planning. With the development of deep learning [32], utilizing deep neural networks (DNN) to achieve function approximation in reinforcement learning becomes the new state-of-the-art. Since the proposal of DQN [39, 40], DRL methods have achieved great success in complex planning tasks, such as the game of Go [47, 48], chemical synthesis [46], chip design [3, 36, 44], VRP [12, 41, 65], and solving mathematical problems [16]. Planning tasks usually have a huge action space, which can be effectively reduced by predicting rewards with a value network and sampling actions via a policy network [23, 30, 38, 47]. Recently, several works [13, 33, 35, 58] adopt GNN as policy and value networks to solve planning tasks on the graph [13, 35]. For example, Fan *et al.* [13] combine GNN with DQN to detect key nodes in complex networks. Meirom *et al.* [35] utilize GNN as a state encoder for DRL to solve the tasks of epidemic control and targeted marketing. In addition, GNN is leveraged to learn representations for urban regions and road networks [11, 25, 26, 34, 56, 61, 62, 64], and support downstream tasks like homogeneity analysis [62] and traffic prediction [11, 61]. However, they only study tasks on existing built roads, which is quite different from the task of planning new roads. Meanwhile, there have been some works utilizing DRL or generative models to accomplish city configuration and urban planning [14, 21, 28, 51–55], but they ignore the slums in cities which is an important issue regarding billions of population. In this work, we make the first attempt to plan new roads for slum upgrading with DRL and GNN.

Road Planning for Slums. Given the large number of slum dwellers and the economic costs, upgrading slums in situ has become the




Figure 9: (a) The travel distance of HSMC and DRL-GNN under different planning budgets. (b) The travel distance at each step of our DRL-GNN model. Best viewed in color.

primary strategy of urbanization, rather than relocating the population to cities. One primary goal of slum upgrading is to provide service access to every place in a slum by building more roads. The re-blocking method [8, 24, 37, 37, 43, 59] is widely adopted in practice, which reconfigures the space and adds road segments, to make each place connected to the road system. With more streets constructed, re-blocking has been shown to significantly reduce the cost of service provision for slums [1, 17]. However, it is not a computational method and requires negotiation with multiple stakeholders, so it is slow and case-by-case. A recent paper by Brelford *et al.* [9] formulated road planning for slums as a constrained optimization problem, making it computationally solvable. Specifically, they proposed a heuristic search approach, adding one path at a time to the least connected place, with the help of Monte Carlo sampling. Considering the huge solution space of this problem, it is difficult for heuristic methods to achieve optimal road planning performance. Different from heuristic search, in this work, we leverage the powerful DRL algorithm to search for optimal road plans in a data-driven way, improving accessibility at minimal costs.

6 CONCLUSION

In this paper, we investigate the problem of road planning for slums, a critical but little-studied issue in sustainable urban development. We formulate it as a sequential decision-making problem with a generic graph model, and propose a novel graph neural network to select locations for new road segments. The model is optimized to improve accessibility and reduce the travel distance of slum dwellers at minimal construction costs. We demonstrate that planning roads for slums through deep reinforcement learning is viable, effective, and can be migrated to real-world, large-scale scenarios. As for future work, we plan to develop a pre-trained model on a large amount of slum data to enable fast inference of road plans, which is beneficial for practical deployment.

ACKNOWLEDGMENTS

We sincerely thank all reviewers for their constructive feedback. We are grateful to Yuming Lin for valuable discussions. This work is supported in part by National Key Research and Development Program of China under grant 2022ZD0116402, National Natural Science Foundation of China under grant U22B2057, U21B2036, 62171260. This work is also supported in part by Beijing National Research Center for Information Science and Technology (BNRist).

REFERENCES

- [1] Alex Abiko, Luiz Reynaldo de Azevedo Cardoso, Ricardo Rinaldelli, and Heitor Cesar Riogi Haga. 2007. Basic costs of slum upgrading in Brazil. *Global Urban Development Magazine* 3, 1 (2007), 121–131.
- [2] Christopher Alexander et al. 2019. A city is not a tree.
- [3] Mohammad Amini, Zhanguang Zhang, Surya Penmetsa, Yingxue Zhang, Jianye Hao, and Wulong Liu. 2022. Generalizable Floorplanner through Corner Block List Representation and Hypergraph Embedding. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2692–2702.
- [4] Deepika Andavarapu, David J Edelman, et al. 2013. Evolution of slum redevelopment policy. *Current urban studies* 1, 04 (2013), 185.
- [5] Marc Barthélemy and Alessandro Flammini. 2008. Modeling urban street patterns. *Physical review letters* 100, 13 (2008), 138702.
- [6] Kurt Binder, Dieter Heermann, Lyle Roelofs, A John Mallinckrodt, and Susan McKay. 1993. Monte Carlo simulation in statistical physics. *Computers in Physics* 7, 2 (1993), 156–157.
- [7] Phillip Bonacich. 1987. Power and centrality: A family of measures. *American journal of sociology* 92, 5 (1987), 1170–1182.
- [8] Somsook Boonyabancha. 2005. Baan Mankong: Going to scale with “slum” and squatter upgrading in Thailand. *Environment and Urbanization* 17, 1 (2005), 21–46.
- [9] Christa Brelsford, Taylor Martin, Joe Hand, and Luis MA Bettencourt. 2018. Toward cities without slums: Topology and the spatial evolution of neighborhoods. *Science advances* 4, 8 (2018), eaar4644.
- [10] Jason Corburn and Irene Karanja. 2014. Informal settlements and a relational view of health in Nairobi, Kenya: sanitation, gender and dignity. *Health promotion international* 31, 2 (2014), 258–269.
- [11] Austin Derrow-Pinion, Jennifer She, David Wong, Oliver Lange, Todd Hester, Luis Perez, Marc Nunkesser, Seongjae Lee, Xueying Guo, Brett Wiltshire, et al. 2021. Eta prediction with graph neural networks in google maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 3767–3776.
- [12] Lu Duan, Yang Zhan, Haoyuan Hu, Yu Gong, Jiangwen Wei, Xiaodong Zhang, and Yinghui Xu. 2020. Efficiently solving the practical vehicle routing problem: A novel joint learning approach. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 3054–3063.
- [13] Changjun Fan, Li Zeng, Yizhou Sun, and Yang-Yu Liu. 2020. Finding key players in complex networks through deep reinforcement learning. *Nature machine intelligence* 2, 6 (2020), 317–324.
- [14] Zhou Fang, Ying Jin, and Tianren Yang. 2022. Incorporating planning intelligence into deep learning: A planning support tool for street network design. *Journal of Urban Technology* 29, 2 (2022), 99–114.
- [15] Reza Zanjirani Farahani, Elnaz Miandoabchi, Wai Yuen Szeto, and Hannaneh Rashidi. 2013. A review of urban transportation network design problems. *European journal of operational research* 229, 2 (2013), 281–302.
- [16] Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Francisco J R Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, et al. 2022. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* 610, 7930 (2022), 47–53.
- [17] Joe Flood. 2004. Cost Estimate for Millennium Development Goal 7, Target 11 on Slums, background report for UN Millennium Project Task Force on Improving the Lives of Slum Dwellers and UN-HABITAT. *Urban Resources, Elsternwick, Australia* (2004).
- [18] Linton C Freeman. 1977. A set of measures of centrality based on betweenness. *Sociometry* (1977), 35–41.
- [19] Linton C Freeman et al. 2002. Centrality in social networks: Conceptual clarification. *Social network: critical concepts in sociology*. Londres: Routledge 1 (2002), 238–263.
- [20] Ahmed Fawzy Gad. 2021. PyGAD: An Intuitive Genetic Algorithm Python Library. arXiv:2106.06158 [cs.NE]
- [21] Yuanzhe Geng, Erwu Liu, Rui Wang, Yiming Liu, Weixiong Rao, Shaojun Feng, Zhao Dong, Zhiren Fu, and Yanfen Chen. 2021. Deep reinforcement learning based dynamic route planning for minimizing travel time. In *2021 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 1–6.
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [23] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*. PMLR, 1861–1870.
- [24] UN Habitat. 2012. Streets as tools for urban transformation in slums: a street-led approach to citywide slum upgrading. 17 (2012), 2016.
- [25] Jilin Hu, Chenjuan Guo, Bin Yang, and Christian S Jensen. 2019. Stochastic weight completion for road networks using graph convolutional networks. In *2019 IEEE 35th international conference on data engineering (ICDE)*. IEEE, 1274–1285.
- [26] Tobias Skovgaard Jepsen, Christian S Jensen, and Thomas Dyhre Nielsen. 2019. Graph convolutional networks for road networks. In *Proceedings of the 27th ACM SIGSPATIAL international conference on advances in geographic information systems*. 460–463.
- [27] Kenji Kawaguchi, Leslie Pack Kaelbling, and Yoshua Bengio. 2017. Generalization in deep learning. *arXiv preprint arXiv:1710.05468* (2017).
- [28] Kira Kempinska and Roberto Murcio. 2019. Modelling urban networks using Variational Autoencoders. *Applied Network Science* 4, 1 (2019), 1–11.
- [29] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [30] Vijay Konda and John Tsitsiklis. 1999. Actor-critic algorithms. *Advances in neural information processing systems* 12 (1999).
- [31] Joseph B Kruskal. 1956. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society* 7, 1 (1956), 48–50.
- [32] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [33] Yu Lei, Hongbin Pei, Hanqi Yan, and Wenjie Li. 2020. Reinforcement learning based recommendation with graph convolutional q-network. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*. 1757–1760.
- [34] Zhenyu Mao, Ziyue Li, Dedong Li, Lei Bai, and Rui Zhao. 2022. Jointly Contrastive Representation Learning on Road Network and Trajectory. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 1501–1510.
- [35] Eli Meir, Haggai Maron, Shie Mannor, and Gal Chechik. 2021. Controlling graph dynamics with reinforcement learning and graph neural networks. In *International Conference on Machine Learning*. PMLR, 7565–7577.
- [36] Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Azade Nazi, et al. 2021. A graph placement methodology for fast chip design. *Nature* 594, 7862 (2021), 207–212.
- [37] Diana Mitlin and David Satterthwaite. 2012. *Urban poverty in the global south: scale and nature*. Routledge.
- [38] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*. PMLR, 1928–1937.
- [39] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [40] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [41] Mohammadreza Nazari, Afshin Oroojlooy, Lawrence Snyder, and Martin Takác. 2018. Reinforcement learning for solving the vehicle routing problem. *Advances in neural information processing systems* 31 (2018).
- [42] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [43] Sheela Patel, Carrie Baptist, and Celine d’Cruz. 2012. Knowledge is power—informal communities assert their right to the city through SDI and community-led enumerations. *Environment and Urbanization* 24, 1 (2012), 13–26.
- [44] Rajarshi Roy, Jonathan Raiman, Neel Kant, Ilyas Elkin, Robert Kirby, Michael Siu, Stuart Oberman, Saad Godil, and Bryan Catanzaro. 2021. Prefixrl: Optimization of parallel prefix circuits using deep reinforcement learning. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 853–858.
- [45] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [46] Marwin HS Segler, Mike Preuss, and Mark P Waller. 2018. Planning chemical syntheses with deep neural networks and symbolic AI. *Nature* 555, 7698 (2018), 604–610.
- [47] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
- [48] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *nature* 550, 7676 (2017), 354–359.
- [49] UN-Habitat. 2004. The challenge of slums: global report on human settlements 2003. *Management of Environmental Quality: An International Journal* 15, 3 (2004), 337–338.
- [50] UN-Habitat. 2014. A Practical Guide to Designing, Planning, and Executing Citywide Slum Upgrading Programmes.

- [51] Dongjie Wang, Yanjie Fu, Kunpeng Liu, Fanglan Chen, Pengyang Wang, and Chang-Tien Lu. 2023. Automated urban planning for reimagining city configuration via adversarial learning: quantification, generation, and evaluation. *ACM Transactions on Spatial Algorithms and Systems* 9, 1 (2023), 1–24.
- [52] Dongjie Wang, Yanjie Fu, Pengyang Wang, Bo Huang, and Chang-Tien Lu. 2020. Reimagining city configuration: Automated urban planning via adversarial learning. In *Proceedings of the 28th international conference on advances in geographic information systems*. 497–506.
- [53] Dongjie Wang, Kunpeng Liu, Pauline Johnson, Leilei Sun, Bowen Du, and Yanjie Fu. 2021. Deep human-guided conditional variational generative modeling for automated urban planning. In *2021 IEEE international conference on data mining (ICDM)*. IEEE, 679–688.
- [54] Dongjie Wang, Pengyang Wang, Kunpeng Liu, Yuanchun Zhou, Charles E Hughes, and Yanjie Fu. 2021. Reinforced imitative graph representation learning for mobile user profiling: An adversarial training perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4410–4417.
- [55] Dongjie Wang, Lingfei Wu, Denghui Zhang, Jingbo Zhou, Leilei Sun, and Yanjie Fu. 2022. Human-instructed Deep Hierarchical Generative Learning for Automated Urban Planning. *arXiv preprint arXiv:2212.00904* (2022).
- [56] Meng-Xiang Wang, Wang-Chien Lee, Tao-Yang Fu, and Ge Yu. 2020. On representation learning for road networks. *ACM Transactions on Intelligent Systems and Technology (TIST)* 12, 1 (2020), 1–27.
- [57] Song Wang, Kaize Ding, Chuxu Zhang, Chen Chen, and Jundong Li. 2022. Task-adaptive few-shot node classification. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1910–1919.
- [58] Tingwu Wang, Renjie Liao, Jimmy Ba, and Sanja Fidler. 2018. Nervenet: Learning structured policy with graph neural networks. In *Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada*, Vol. 30.
- [59] Jane Weru. 2004. Community federations and city upgrading: the work of Pamoja Trust and Muungano in Kenya. *Environment and Urbanization* 16, 1 (2004), 47–62.
- [60] Amy Wesolowski and Nathan Eagle. 2010. Parameterizing the dynamics of slums. In *2010 AAAI Spring Symposium Series*.
- [61] Ning Wu, Xin Wayne Zhao, Jingyuan Wang, and Dayan Pan. 2020. Learning effective road network representation with hierarchical graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 6–14.
- [62] Jiawei Xue, Nan Jiang, Senwei Liang, Qiyuan Pang, Takahiro Yabe, Satish V Ukkusuri, and Jianzhu Ma. 2022. Quantifying the spatial homogeneity of urban road networks via graph neural networks. *Nature Machine Intelligence* 4, 3 (2022), 246–257.
- [63] Rui Yang, Jie Wang, Zijie Geng, Mingxuan Ye, Shuiwang Ji, Bin Li, and Feng Wu. 2022. Learning Task-relevant Representations for Generalization via Characteristic Functions of Reward Sequence Distributions. *arXiv preprint arXiv:2205.10218* (2022).
- [64] Mingyang Zhang, Tong Li, Yong Li, and Pan Hui. 2021. Multi-view joint graph representation learning for urban region embedding. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*. 4431–4437.
- [65] Zefang Zong, Hansen Wang, Jingwei Wang, Meng Zheng, and Yong Li. 2022. RBG: Hierarchically Solving Large-Scale Routing Problems in Logistic Systems via Reinforcement Learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4648–4658.

APPENDIX

A RESEARCH METHODS

A.1 Markov Decision Process

We propose a DRL model to solve the sequential decision-making problem, where an intelligent agent learns to automatically select locations for road segments by interacting with a slum planning environment, as shown in Figure 2. From the perspective of DRL, the problem can be expressed as a Markov Decision Process (MDP), which contains the following critical components:

- **States** describe the current conditions of the slum, including both static and dynamic features for places and roads.
- **Actions** indicate the selected locations of new road segments.
- **Rewards** provide feedback for road planning actions, which consider the connectivity, travel distance, and construction cost to obtain a comprehensive evaluation.
- **Transitions** express the dynamic changes of the slum, such as the changes of segments from *candidates* to *roads*, and the resulting changes in accessibility and travel distance.

A.2 Definitions of Topological Features

We design rich features for topological elements on the graph, including nodes, edges and faces. These features are used as input of the proposed GNN model to learn representations. We include 11 categories of features as illustrated in Table 1. We now introduce the specific definitions of these features.

Node Features. Nodes represent junctions in a slum, *i.e.* *points* in the original geometric space. We include the following features,

- **Coordinates:** the Cartesian coordinates (x, y) indicating the location of the junction in the slum.
- **Centrality:** the network centrality metrics of the junction. We compute four centrality metrics including degree centrality, betweenness centrality [18], eigenvector centrality [7] and closeness centrality [19].
- **On Road:** a boolean feature indicating whether the junction is on a road, either external or planned.
- **Road Ratio:** the ratio of the number of adjacent road edges to the total number of adjacent edges. It is 0 when *On Road* is *False*.
- **Avg N2N Dis:** the average distance from the node to all other nodes over the constructed road network. It is set as a very large value if the node is not on a road.

Edge Features. Edges represent paths in a slum, *i.e.* *line segments* in the original geometric space. We include the following features,

- **Cost:** the construction cost of building the segment as a road, which is set as the length of the path.
- **Road:** a boolean feature indicating whether the path is a road or not. A road can be an existing one or a newly planned one.
- **Straightness:** the ratio of the road network distance to the euclidean distance between the two endpoints.

Face Features. Faces represent places in a slum, which are *polygons* in the original geometric space. We include the following features,

- **Connected:** a boolean feature indicating whether the place is connected to the road system.

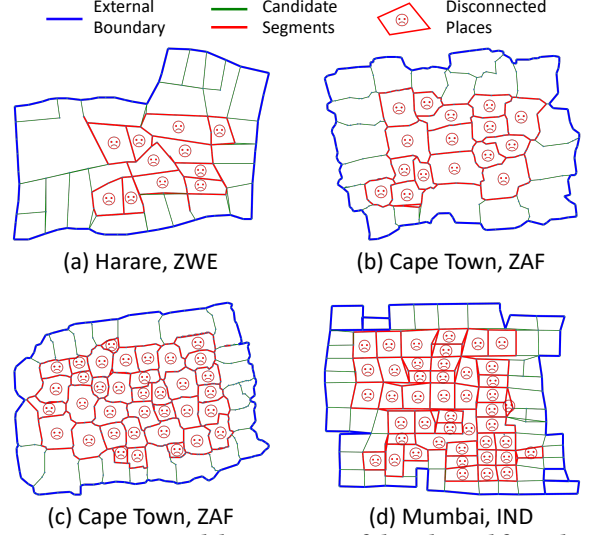


Figure 10: Geometrical descriptions of the adopted four slums in (a) Harare, ZWE (b) Cape Town, ZAF (c) Cape Town, ZAF (d) Mumbai, IND. All slums suffer from poor accessibility, with a large fraction of places disconnected to the road system. Best viewed in color.

Table 4: Designed features for topological elements.

Category	Hyper-parameter	Value
Network	GNN layer	2
	GNN node dimension	16
	Policy Head MLP_p	[32, 1]
	Value Head MLP_v	[32, 32, 1]
PPO	gamma	0.995
	tau	0
	Entropy Loss β	0.01
	Value Loss γ	0.5
Train	optimizer	Adam
	weight decay	0
	learning rate	0.0004

- **Avg F2F Dis:** the average distance from the place to all other places over the constructed road network. It is set as a very large value if the face is not connected to the road system.
- **F2E Dis:** the distance from the place to the external boundaries E of the slum.

B DETAILS OF SLUM DATA

As shown in Table 2, we conduct experiments on four real-world slums from three different countries, including Zimbabwe (ZWE), South Africa (ZAF) and India (IND). The specific locations of the four slums are Epworth (Harare, ZWE), Khayelitsha (Cape Town, ZAF) and Phule Nagar (Mumbai, IND). The digital maps and the geometrical descriptions of places and roads for the four slums are publicly released by [9], as shown in Figure 10.

C HYPER-PARAMETERS OF OUR MODEL

We implement the proposed model with PyTorch [42], and release the codes at <https://github.com/tsinghua-fib-lab/road-planning-for-slums>. We tune the hyper-parameters of our model carefully and list the adopted values in Table 4.