

FAP-CD: Fairness-Driven Age-Friendly Community Planning via Conditional Diffusion Generation

Jinlin Li¹, Xintong Li², Xiao Zhou¹ *

¹ Gaoling School of Artificial Intelligence, Renmin University of China, China

² School of Statistics, Renmin University of China, China

jinlin2021@ruc.edu.cn, lixintong@ruc.edu.cn, xiaozhou@ruc.edu.cn

Abstract

As global populations age rapidly, incorporating age-specific considerations into urban planning has become essential to addressing the urgent demand for age-friendly built environments and ensuring sustainable urban development. However, current practices often overlook these considerations, resulting in inadequate and unevenly distributed elderly services in cities. There is a pressing need for equitable and optimized urban renewal strategies to support effective age-friendly planning. To address this challenge, we propose a novel framework, **F**airness-driven **A**ge-friendly community **P**lanning via **C**onditional **D**iffusion generation (FAP-CD). FAP-CD leverages a conditioned graph denoising diffusion probabilistic model to learn the joint probability distribution of aging facilities and their spatial relationships at a fine-grained regional level. Our framework generates optimized facility distributions by iteratively refining noisy graphs, conditioned on the needs of the elderly during the diffusion process. Key innovations include a demand-fairness pre-training module that integrates community demand features and facility characteristics using an attention mechanism and min-max optimization, ensuring equitable service distribution across regions. Additionally, a discrete graph structure captures walkable accessibility within regional road networks, guiding model sampling. To enhance information integration, we design a graph denoising network with an attribute augmentation module and a hybrid graph message aggregation module, combining local and global node and edge information. Empirical results across multiple metrics demonstrate the effectiveness of FAP-CD in balancing age-friendly needs with regional equity, achieving an average improvement of 41% over competitive baseline models.

Code — https://github.com/jinlin2021/FAP_CD

Introduction

Over the past two decades, global population aging has accelerated. By 2050, individuals aged 65 and older are projected to comprise 38% of the world’s population (Rudnicka et al. 2020), posing new demands on urban infrastructure and services. To address these challenges, the World Health Organization (WHO) introduced the concept of Age-friendly Cities in 2007 (Organization 2007), highlighting the

*Corresponding author

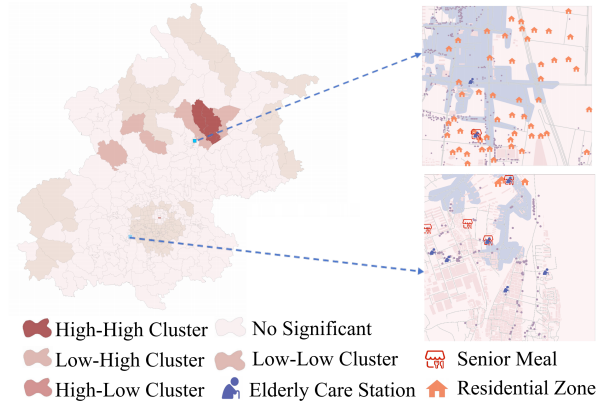


Figure 1: Local Moran’s Index analysis of the distribution of age-related facilities in Beijing, illustrating spatial disparities in resource allocation. Two regions with contrasting resource levels are magnified, with blue-shaded zones indicating 15-minute walkable areas around facilities.

need for developing Age-friendly Communities (AFCs) to support aging populations (Menec et al. 2011). AFCs aim to improve the quality of life and social participation of elderly residents through enhanced community services. Research further shows that most elderly individuals prefer to age in residential communities rather than rely on institutional care services (Sun and Yang 2023). In this context, studying and optimizing the distribution of aging-related facilities and services within communities is essential for the effective development of age-friendly cities.

Unfortunately, even in some of the most advanced and developed modern metropolises, the level of age-friendly infrastructure remains far from adequate (Wang, Gonzales, and Morrow-Howell 2017; Xie 2018). For instance, in Beijing, although 22.6% of the population was aged 60 and above by 2023, significant deficiencies in elderly care facilities persist. Our analysis, leveraging the 15-minute Walkability metric (Moreno et al. 2021)—which evaluates resource accessibility and captures the ease of access to community-based elderly services—shows that only 40.85% of communities have senior care centers within a 15-minute walking distance, and merely 24.51% offer meal assistance services for the elderly. Furthermore, an analysis us-

ing Local Moran’s Index reveals significant spatial inequalities in the distribution of elderly care facilities across Beijing’s neighborhoods, as illustrated in Figure 1. More critically, as the aging population grows, demands for medical, recreational, and educational facilities become increasingly diverse and urgent (Yung, Conejos, and Chan 2016), yet current facility allocations fail to meet these needs or support AFCs development and regeneration.

The planning of age-friendly communities is crucial for achieving active aging among community seniors and establishing sustainable communities. The core task is to create an optimized spatial layout that reflects the needs of the elderly in urban areas (Sun and Yang 2023). Traditional large-scale urban planning, from initial sketches to detailed designs, is not only complex and time-consuming but also incurs substantial costs. In today’s rapidly digitizing world, fueled by urban big data and advancements in artificial intelligence (AI), automated spatial layout design has not only become feasible but is also emerging as a prevailing trend.

In recent years, a few studies have attempted to generate planning schemes automatically through deep generative techniques. For instance, Wang et al.(2020) proposed a generative model for land use allocation, utilizing geographic and human mobility data to build spatial graph representations and training a generative adversarial network with positive and negative sample contrasts. Wang et al.(2023) developed an instruction-based hierarchical planner (IHPlanner) that captures the spatial hierarchy and planning dependencies between urban functional zones and land use allocation using human instructions and the surrounding environment, generating land use configurations for vacant areas. Zheng et al.(2023) proposed a reinforcement learning (RL)-based urban planning model, constructing an urban adjacency graph with urban geographic elements (blocks, roads) as nodes and framing urban planning as a sequential decision-making problem on the graph. Additionally, some studies (Ye, Du, and Ye 2022; Eskandar et al. 2023) have viewed urban areas as two-dimensional images, treating scheme generation as an image generation task.

Although these methods have advanced the field of automated urban planning scheme generation, they are limited in scope: either simulating urban development patterns at a macro level while neglecting detailed land use and building layouts, or coarsely updating land allocation without incorporating the urban multimodal data (Yong and Zhou 2024; Xu and Zhou 2024) of target communities. Moreover, no algorithms currently exist for generating planning schemes specifically tailored to age-friendly urban renewal. Addressing this critical technical gap is an urgent priority.

Recently, diffusion models (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020) have made significant advancements in the field of computer vision, becoming the most advanced deep generative models, surpassing other generative models such as Variational Autoencoders (VAEs) (Kingma and Welling 2013) and Generative Adversarial Networks (GANs) (Goodfellow et al. 2020). Diffusion models approximate complex data distributions by gradually removing small amounts of noise from a simple distribution. The refined denoising process enables these models

to effectively fit complex distributions and generate high-quality data. In this study, we propose a fairness-driven age-friendly community planning method (FAP-CD) based on the denoising diffusion model guided by stochastic differential equations (SDEs). Leveraging conditional diffusion generation, this approach frames the planning problem as a facility graph generation task. It utilizes complex walking graphs to capture spatial relationships among regional facilities, aiming to better meet the needs of elderly residents while promoting fairness across regions.

More specifically, we deploy a unified graph noise prediction model that intricately interacts with the node and edge representations in both real-valued matrices and discrete graph structures. Initially, we develop a fair-demand pre-training module, utilizing min-max optimization (Pardalos, Migdalas, and Pitsoulis 2008) to derive spatial region representations that achieve maximum minimum entropy (Han and Sung 2021; Razaviyayn, Hong, and Luo 2013). These representations are meticulously crafted to reflect residential demands within the region and to promote equitable characteristics across communities. Subsequently, we guide the diffusion process by carefully adjusting the mean of the introduced Gaussian noise. During the denoising phase, we implement a hybrid graph noise prediction model, which comprises a message-passing layer on the discrete graph and a graph transformer layer. The message-passing layer is responsible for modeling the dependencies among the edges of neighboring nodes, while the transformer layer is dedicated to extracting and disseminating global information. Furthermore, we utilize a straightforward walking capability graph to support conditional discrete graph denoising and employ m-step random walk matrices, derived from the discrete adjacency matrix, to enhance the representation learning of nodes and edges.

Our main contributions are summarized as follows:

- We propose an efficient conditional discrete graph diffusion framework. To the best of our knowledge, this is the first study on generating AFCs planning that consider both elderly needs and equity.
- We design a novel graph denoising diffusion network conditioned on fair-demand embedding and simple discrete structure to ensure practicality and fairness of the spatial layout in the generated grid layouts.
- Extensive experiments on real-world multi-modal urban data demonstrate that FAP-CD delivers superior performance and significantly outperforms contemporary state-of-the-art methods. Additionally, we provide in-depth analyses to highlight the superiority of our framework.

Preliminaries

Problem Statement

Given a city area divided into an $H \times W$ grid map based on latitude and longitude, with each grid of size $d \times d m^2$, the grid regions are denoted as $R = \{r_1, r_2, \dots, r_k\}$, where k is the number of grid regions in the city. In our study, d is set to 2 km, and each grid encompasses the following aspects:

Urban attributes. Urban attributes are the inherent social and geographic features of urban regions. A specific type

of regional attribute can be denoted as $A \in \mathbb{R}^{N \times f_a}$, where f_a is the region attribute dimensions. In our work, multiple region attributes, like demographics and socio-economic indicators, are considered.

Walking graph. For each grid region r_i , we construct an undirected weighted graph, denoted as $G = \{N, E, A\}$, where N , E are the node set and edge set, respectively. $A \in \{0, 1\}^{N \times N}$ is the adjacency matrix to describe the walkable accessibility between every two nodes in the region. $N \in \mathbb{R}^{D \times f_d}$, where D denotes the total number of facilities and f_d dictates the distinct categories of facilities within the grids R . Each category is encoded using a one-hot representation. Additionally, textual modal descriptions of each facility are captured as grid features $F \in \mathbb{R}^{N \times d}$.

Grids region demand. This includes the current status of facility distribution on each grid region, based on Beijing's standards for facilities per thousand residents¹ and the regulations for age-appropriate amenities², we evaluate the existing facility on each grid. We classify the state of area facilities into four categories: no supply, under supplied, appropriately supplied, and oversupplied. We standardize each demand representation $D \in \{0, 1, 2, 3\}^{N \times N}$ (i.e., high configuration, proper fit, low configuration, under allocation).

AFCs Planning Generation. Given a set of grid regions $R = \{r_1, r_2, \dots, r_k\}$, each characterized by grid demand D , urban region attribute A and grid feature F , our goal is to generate the AFCs planning that addresses equitable demands across various urban areas in different time intervals.

Preliminaries on Diffusion Process

Diffusion models are a class of generative models that generate samples by learning a distribution that approximates a data distribution. The first step in constructing diffusion models (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020) is to define a forward diffusion process. Assuming a continuous random variable $x_0 \in \mathbb{R}^d$, the diffusion models consist of a forward process that gradually add noise to perturbs x_0 until the output distribution becomes a known prior distribution within the specified time range t in the interval $[0, T]$, we have a Gaussian transition kernel as follows:

$$q_{0t}(x_t|x_0) = \mathcal{N}(x_t|\alpha_t x_0, \sigma_t^2 I), \quad (1)$$

where α_t and σ_t are time-dependent differentiable functions, usually chosen to ensure that the signal-to-noise ratio $\frac{\alpha_t^2}{\sigma_t^2}$ decreases. This ensures that $q_T(x_T) \approx \mathcal{N}(0, I)$, resulting in a tractable prior distribution with a low signal-to-noise ratio that is convenient for sampling. The forward process can be represented by the following SDE:

$$d\mathbf{x}_t = f(\mathbf{x}_t, t)dt + g(t)d\mathbf{w}_t, \quad (2)$$

where $f(\cdot, t)$ is the drift coefficient, $g(t)$ is the scalar diffusion terms (a.k.a. noise schedule function), and \mathbf{w}_t is a standard Brownian motions on $f(\cdot, t)$. The reverse-time SDEs from time T to 0 is denoted as:

$$d\bar{\mathbf{x}}_t = [f(t)\mathbf{x}_t - g^2(t)\nabla \log q_t(\mathbf{x}_t)]dt + g(t)d\bar{\mathbf{w}}_t, \quad (3)$$

¹<http://www.cacp.org.cn>

²<https://www.beijing.gov.cn/zhengce/zhengcefagui/>

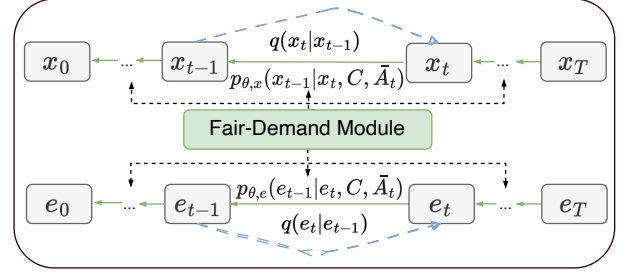


Figure 2: Framework overview of the proposed FAP-CD.

where $\nabla \log q_t(\mathbf{x}_t)$ is the score function and $d\bar{\mathbf{w}}_t$ is the reverse-time standard Wiener process. By learning to reverse this process, the diffusion model generates new samples from the data distribution.

In this work, to address the scarcity of elderly-friendly facilities and spatial inequity within contemporary urban structures, we adopt a diffusion model based on graph SDEs. Specifically, we abstract the spatial and facility relationships within grid region into a graphical representation, constructing a graph $G = (X, A)$ composed of N nodes with features $X \in \mathbb{R}^{N \times F}$ encoded using one-hot encoding to denote the types of facilities, and the adjacency matrix $A \in \mathbb{R}^{N \times N}$ (i.e., $A_{i,j} = 1$ if the walking time between node i and j is within 15 minutes).

Methodology

Figure 2 illustrates the overall framework of FAP-CD, which can be divided into two stages: forward process and reverse process with conditional graph denoising. The fair-demand module learns to ensure a valid fair representation optimized by the min-max principle across regions, which serves as the condition supporting the graph denoising diffusion process.

Fair-demand Module

To generate a target optimized spatial distribution map of facilities from noise in vector space, we design an attention mechanism-based fair-demand module as shown in Figure 3 (c). The denoising process utilizes the region-specific max-min fairness demand embeddings, which are learned through the fair-demand module, to guide the AFCs generation. Specifically, for each region r_i , let \bar{D}_i denotes the representation obtained by concatenating the urban attributes and demand representations, then use a linear projection to obtain the query vector $Q = \bar{D}_i W_q$. Subsequently, grid features are mapped through a fully connected layer to the same vector space to get the key vector $K = F_i W_k$ and the value vector $V = F_i W_v$. The importance of region demand representations and the output are calculated as follows:

$$\alpha_i = \text{softmax}\left(\frac{Q F_i^\top}{\sqrt{d_h}}\right), \quad (4)$$

$$H_i = \alpha_s V W_H, \quad (5)$$

where W_O is the output projection matrix.

In this study, we advance beyond simply concatenating various conditional embeddings for control conditions. We

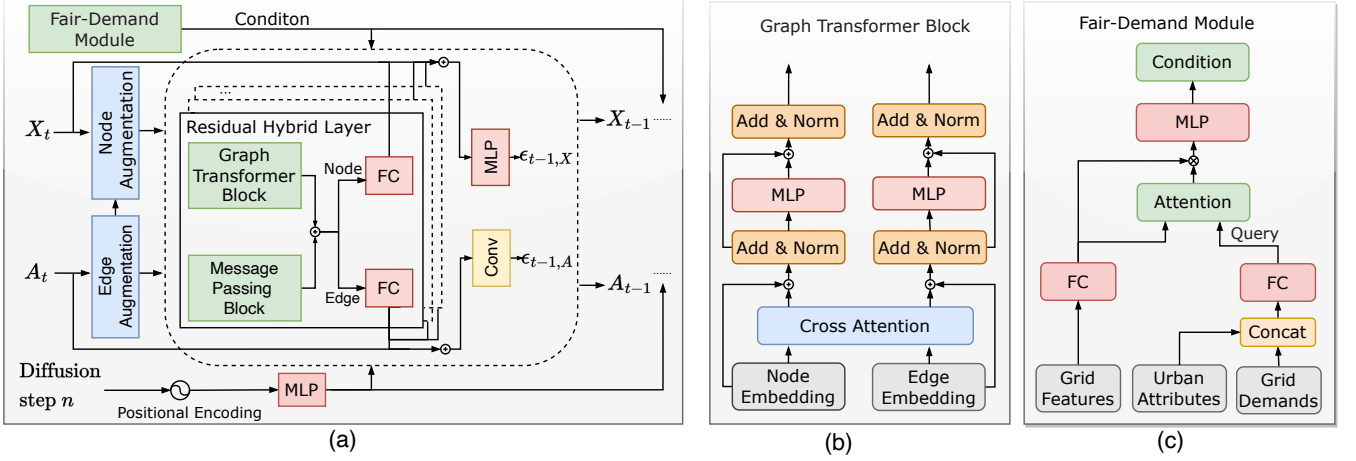


Figure 3: The architecture of (a) graph denoising network, (b) graph transformer block, and (c) fair-demand module.

focus on detailed modeling of demand interactions among facilities and the semantic characteristics of residential areas, while also addressing inter-regional fairness. Our refined conditional embedding, represented as $C_i = H_i \circ F_i$, empowers the model to adaptively leverage these representations for more nuanced and effective outcomes.

To generate a distribution of service facilities that meets the requirements of age-friendliness and spatial fairness, we develop a pre-trained model optimized based on the principle of max-min entropy. This framework has been widely applied across various optimization and control domains (Liu, Dai, and Luo 2013; Zhang, Lemoine, and Mitchell 2018; Wai et al. 2019). Our model is trained by minimizing the loss function:

$$L = \frac{1}{\min(-\bar{p}_i * \log(\bar{p}_i)) + 1}, i \in M, \quad (6)$$

where $\bar{p}_i = \sum_{j=1}^n p_{ij}$, represents the normalized probability distribution, while $p_{ij} = 0$ if $j \in R_i$. The term p_{ij} is the softmax output of O_i , n is the batch size, R_i denotes the index of the residence within the grid region. M represents the maximum number of facilities per region.

Discrete Graph Diffusion Model

In this work, we formulate the AFCs planning generation task in a conditional SDEs-based graph denoising framework. Inspired by (Jo, Lee, and Hwang 2022; Vignac et al. 2022) which treats (X, A) as a whole and omit the intrinsic graph structure. we diffuse separately on each node and edge feature of the constructed graph $G = (X, A)$. We apply the variance-preserving SDE introduced by (Song et al. 2020), which in its discrete-time manifestation adopts the framework of a denoising diffusion probabilistic model (Ho, Jain, and Abbeel 2020). Similarly to diffusion models for images, which apply noise independently on each pixel, we diffuse separately on each node and edge feature. Roughly speaking, the forward diffusion process for each graph sample can be described by the SDE with $t \in [0, T]$ as:

$$dG_t = f(G_t, t)dt + g(t)d(w_{X_t}, w_{A_t}), \quad (7)$$

where w_{X_t} and w_{A_t} are independent standard Wiener processes for the node and edge, respectively. This forward SDE shares the same transition distribution as in Eq. 1 to conveniently sample Gaussian noise $G_t = \alpha_t G_0 + \sigma_t \epsilon_G$ at any time t . For undirected graphs, noise is applied only to the upper triangular part of E and then symmetrized to maintain the undirected nature of the graph.

In the reverse process with conditional denoising, FAP-CD defines the conditional reverse diffusion process $p_\theta(G_{t-1}|G_t, C, \bar{A})$, which performs iterative denoising from pure Gaussian noise to generate optimized spatial facilities in vector space. This process is conditioned on the region fair-demand embeddings C and a discrete adjacency matrix \bar{A} , representing facility nodes within a 15-minute walk from residential areas. The corresponding reverse-time SDEs are given by:

$$\begin{aligned} dX_t &= [f(t)X_t - g^2(t)\nabla_X \log q_t(X_t, A_t)]dt + g(t)d\bar{w}_{X_t}, \\ dA_t &= [f(t)A_t - g^2(t)\nabla_A \log q_t(X_t, A_t)]dt + g(t)d\bar{w}_{A_t}. \end{aligned} \quad (8)$$

Then, we train two time-dependent score networks $\epsilon_{\theta, X}(\cdot)$ and $\epsilon_{\theta, A}(\cdot)$ i.e., the parametric score function, for estimating the score $\nabla_X \log q_t(X_t, A_t)$ and $\nabla_A \log q_t(X_t, A_t)$ for both (X_t, A_t) , separately. The denoising score matching objective (Vincent 2011) is modified for score estimation training as:

$$\min_{\theta} \mathbb{E}_t \{ \lambda(t) \mathbb{E}_{G_0|G_t} [\| \epsilon_X - \epsilon_{\theta, X}(G_t, C, \bar{A}_t, t) \|_2^2 + \| \epsilon_A - \epsilon_{\theta, A}(G_t, C, \bar{A}_t, t) \|_2^2] \}, \quad (9)$$

where $\lambda(t)$ is a given positive weighting function to keep the transition kernel $p_{0t}(x_t|x_0)$ to be a tractable Gaussian distribution. $\epsilon_X = \nabla_X \log p_{0t}(X_t|X_0)$ and $\epsilon_A = \nabla_A \log p_{0t}(A_t|A_0)$ are the sampled Gaussian noise. For the diffusion step t , we use a positional encoding method as described in Ho, Jain, and Abbeel(2020), followed by two fully connected layers before feeding into the denoising network.

Conditional Graph Denoising Network

The architecture of our denoising network is shown in Figure 3 (a). The conditional denoising network processes a

noisy graph represented by $G_t = (X_t, A_t)$, alongside the regional fair-demand embedding C and the discrete graph A as input conditions. The backbone network of our diffusion model integrates the graph transformer, as proposed by (Dwivedi and Bresson 2020), with standard message-passing networks to form a Residual Hybrid Layer. This design enhances the model’s ability to effectively process and integrate graph-based data. To incorporate time information, we normalize the time step to $[0, 1]$ and treat it as a global feature to add noise. It outputs tensors X and A in different time intervals t , which respectively represent the optimization distribution over clean graphs.

Graph Transformer Block. In each layer, the attention weights between every pair of nodes are calculated through self-attention. Notably, the calculation of attention weights for the propagation of information between nodes also incorporates the influence of edge features.

$$a_{i,j} = \text{softmax}\left(\frac{(\tanh((W_0^{k,l} E_{i,j}^l)) \cdot Q^{k,l}) K^{k,l}}{\sqrt{d_k}}\right), \quad (10)$$

$$M^{l+1} = \sum_{j=0}^{N-1} a_{i,j}^{k,l} (\tanh(W_1^{k,l} (E_{i,j}^l)) \cdot V^{k,l}), \quad (11)$$

where $E_{i,j}$ is the corresponding image-like tensor i.e. edge feature, Q, K, V are projected from the node features $H_{i,j}^{(l)}$, $H^{(l+1)} \in \mathbb{R}^{N \times d}$ represents the node outputs, $W_0^{(l)}$ and $W_1^{(l)}$ are learnable parameters, \tanh is the activation layer and l is transformer layer. Our transformer layers also feature residual connections and layer normalization.

Message Passing Block. Since the graph transformer captures global graph information, in our task, there is also a need to focus on neighboring nodes within a 15-minute walking range. Therefore, we rely on the decoded discrete graph structure and employ a standard message-passing layer based on the Graph Convolutional Network (GCN) to aggregate local neighbor edge features. Then, we feed the node outputs $\bar{H}^{(l+1)}$ into the FFN^(l), which comprises a multi-layer perceptron (MLP) and a normalization layer. Finally, we obtain the aggregated node output $\bar{H}_{i,j}^{(l+1)} \in \mathbb{R}^{N \times d}$ and the corresponding edge features $\bar{E}_{i,j}^{(l+1)} \in \mathbb{R}^{N \times N \times d}$.

$$H_{i,j}^{l+1} = \text{FFN}_0^l(\bar{H}^{l+1}), \quad (12)$$

$$E_{i,j}^{l+1} = \text{FFN}_1^l(\bar{H}_i^{l+1} + \bar{H}_j^{l+1}). \quad (13)$$

Next, node and edge outputs from the graph transformer and message passing blocks are independently aggregated, then each fed into a fully connected layer. The resulting outputs, $\epsilon_{\theta, X}$ and $\epsilon_{\theta, A}$, are computed by aggregating data via skip connections from each residual hybrid layer, followed by processing through an MLP and three Conv1x1 blocks.

Augmentation Module. We design a module to enhance node and edge features within graphs by incorporating statistical properties such as degree and centrality. Furthermore, using an m -step random walk derived from the discrete adjacency matrix, we develop m -step random walk matrices to employ the arrival probability vectors concatenated with original node features as enhanced node embeddings. Similarly, truncated shortest-path distances obtained

from the same matrix are used to augment edge embeddings. Through node and edge augmentation, the network can comprehensively model the noisy topology structure from both node and edge perspectives at each diffusion step.

Experiments

Experimental Setup

Data Description. We develop our model using real-world datasets, focusing on Beijing. We obtain geographic information from OpenStreetMap, including road networks and administrative divisions. Additionally, we source 149,049 Points of Interest (POIs) within the city from Baidu Maps³, categorized into 11 types such as food, medical services, etc. We also gather data on 568 senior meal programs and 1,472 senior care centers from⁴. In collaboration with the real estate agency⁵, we acquire comprehensive data on 8,201 residential areas, detailing housing prices, property fees, community populations, and elderly populations over age 60.

Evaluation Metrics. To comprehensively evaluate our model, we employ a set of evaluation metrics that take into account facility types, spatial distribution and fairness, which provide a multi-dimensional view of the effectiveness and equity of the allocations. The key evaluation metrics include: **Efficiency** = $\frac{1}{|R| \cdot |K|} \sum_{i \in R} \sum_{k \in K} \sum_{h \in H_i} \frac{\max_{n \in N_{ik}} A_{in,h}}{|N_{ik}|}$,

Diversity = $\frac{1}{|R|} \sum_{i \in R} \left(\frac{|\{C_{i,j} | j \in S_i \text{ and } C_{i,j} \neq R_i\}|}{N-1} \right)$, **Accessibil-**

ity = $\frac{1}{|R|} \sum_{i \in R} \sum_{h \in H_i} \frac{\max_{n \in N_{ik}} A_{in,h}}{P_i \cdot u}$, **Gini** = $1 - \frac{2 \cdot \sum_{i=1}^N \sum_{j=1}^i X_{(j)}}{N \cdot \sum_{i=1}^N X_i}$. In

this study, **Efficiency** refers to the distribution efficiency of life service and elderly care facilities, as measured by the Life Service and Elderly Care metrics. The **Average** metric assesses the overall performance on all metrics, where the Gini metric is negated.

Baseline Models. We compare our proposed FAP-CD with two classes of methods. The first type is **1) Traditional methods:** Walking-based (Song, Kong, and Cheng 2024). The Ant Colony Algorithm (ACA) (Yang, Yu, and Cheng 2007) generates an initial set of feasible metro lines and then optimizes community planning through crossover and mutation. The Genetic Algorithm (GA) (Jensen 2019) guides station selection with a pheromone-based rule and updates edge pheromones to optimize community planning. DRF (Ghodsi et al. 2011) dynamically allocates resources to the user (grid region) with the lowest dominant share until the user’s demand cannot be met. And the other **2) Deep generative models:** VGAE (Kipf and Welling 2016) uses GCN to encode nodes into a latent space and then reconstructs the graph by decoding the adjacency matrix from the node embeddings. GraphRNN (You et al. 2018) sequentially generates graphs by using recurrent neural networks to model the dependencies between nodes and edges. CondGEN (Yang

³<https://map.baidu.com>

⁴<https://www.beijingweilao.cn>

⁵<https://m.lianjia.com>

Model	Life Service \uparrow	Elderly Care \uparrow	Diversity \uparrow	Accessibility \uparrow	Gini \downarrow	Average \uparrow
Walking-based	0.419	0.223	0.600	0.167	0.710	0.140
ACA	0.484 \pm 0.003	0.501 \pm 0.004	0.749 \pm 0.001	0.111 \pm 0.000	0.869 \pm 0.008	0.195 \pm 0.001
GA	0.445 \pm 0.004	0.393 \pm 0.005	0.645 \pm 0.015	0.102 \pm 0.013	0.479 \pm 0.005	0.221 \pm 0.003
DRF	0.449 \pm 0.098	0.604 \pm 0.083	0.728 \pm 0.053	0.123 \pm 0.024	0.376 \pm 0.139	0.306 \pm 0.024
VGAE	0.180 \pm 0.020	0.263 \pm 0.125	0.218 \pm 0.001	0.356 \pm 0.086	0.701 \pm 0.033	0.063 \pm 0.040
GraphRNN	0.442 \pm 0.022	0.546 \pm 0.035	0.503 \pm 0.061	0.361 \pm 0.004	0.387 \pm 0.016	0.293 \pm 0.021
CondGEN	0.643 \pm 0.092	0.667 \pm 0.059	0.563 \pm 0.010	0.345 \pm 0.041	0.323 \pm 0.029	0.379 \pm 0.035
DDPM	0.564 \pm 0.028	0.689 \pm 0.030	0.820 \pm 0.015	0.402 \pm 0.020	0.394 \pm 0.027	0.416 \pm 0.014
EDGE	0.431 \pm 0.002	0.590 \pm 0.068	0.667 \pm 0.055	0.234 \pm 0.115	0.327 \pm 0.003	0.341 \pm 0.051
FAP-CD (Ours)	0.888\pm0.015	0.923\pm0.031	0.843\pm0.004	0.520\pm0.103	0.232\pm0.013	0.588\pm0.028

Table 1: Performance comparison of FAP-CD with baselines. The best results are in bold, and the second-best are underlined.

et al. 2019) is a variational graph generative model that integrates flexible structures while maintaining permutation invariance. DDPM (Ho, Jain, and Abbeel 2020; Wen et al. 2023) uses the UGnet framework to build a diffusion process applied to node representations and image-like adjacency matrices. EDGE (Chen et al. 2023) leverages graph sparsity during the diffusion process by focusing on a small subset of nodes, improving efficiency and addressing the challenges of large-scale graph generation. In this baseline, we use node category sequences as guidance instead of node degrees.

Implementation Details. We conduct experiments using two Nvidia A40 GPUs, configuring 200 diffusion steps and employing a 3rd-order DPM-Solver (Lu et al. 2022) along with a cosine noise scheduler (Nichol and Dhariwal 2021).

Experimental Results

Overall Comparison. Table 1 shows the overall comparison results in terms of all evaluation metrics in Beijing. Our observations based on the results are as follows. First, the AFCs planning generated by FAP-CD achieves the best realism in terms of both facility distribution efficiency and equity. Specifically, compared with the best baseline, the *Accessibility* is improved by over 29% and *Gini* index reduced over 28%, respectively. Moreover, our model outperforms all baselines on average by a large margin. Second, We discover that when categorizing facility allocations, traditional resource allocation models, particularly DRF, exhibit superior diversity. This superiority arises because the facility allocations in methods AC and GA are guided by a demand-driven pheromone update mechanism and fitness functions, respectively, and are also influenced by stochastic factors.

Our analysis further reveals that the overall performance of the EDGE is inferior to that of DDPM and our model. This indicates that while graph generation guided by node categorization presents significant advantages in terms of graph discretization and time efficiency, relying solely on a denoising network composed of MLPs and Gated Recurrent Units (GRUs) to extract neighborhood and global spatial information from the graph does not optimize graph generation. This underscores the relative advantages of our hybrid message passing and augmentation strategy.

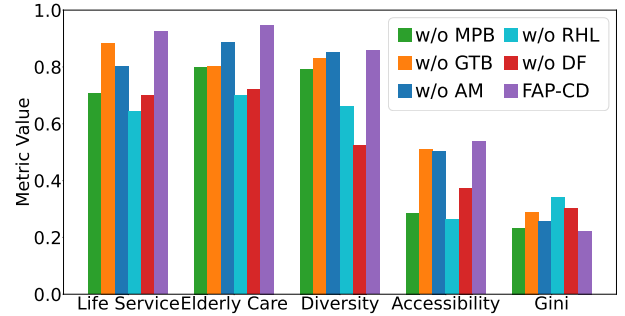


Figure 4: Performance Comparison in the Ablation Study.

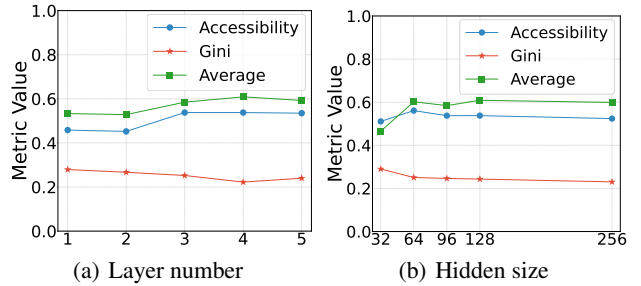


Figure 5: Analysis of Hyperparameters.

Ablation Study. We conduct an ablation study to validate the efficacy of the Residual Hybrid Layer. Results in Figure 4 indicate that incorporating the residual hybrid module (w/o RHL) within the denoising model enhances AFCs generation performance by up to 105.1% according to metric *accessibility*. Additionally, ablation experiments are conducted separately on the graph transformer block (w/o GTB) and message passing block (w/o MPB) within the denoising network, revealing that removing either mechanism causes performance degradation to varying degrees. We also evaluate the impact of the node and edge attribute enhancement module (w/o AM) within the SDE-based diffusion model. As shown in Figure 4, this design yields up to a 13.1% improvement in performance according to metric *Gini*, marking the most significant enhancement observed.

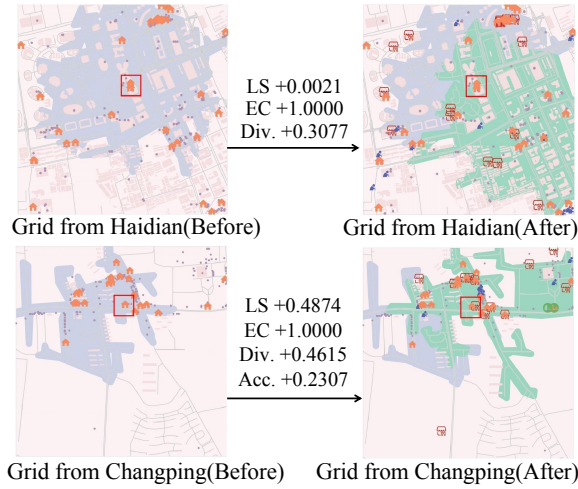


Figure 6: Comparison of Generated and Original AFCs.

Hyperparameters. We next conduct experiments on various hyperparameters, including the number of residual hybrid layers in the conditional graph denoising network and the dimension of hidden variables. As illustrated in Figure 5 (a), the results show that a shallow and appropriate number of layers (three in this case) is crucial for FAP-CD, as overly deep networks tend to overfit, leading to a decline in performance. Figure 5 (b) examines the impact of hidden variable dimension size. The results indicate that an optimal hidden size (e.g., 128) enhances model performance, whereas excessively large dimensions (e.g., 256) do not improve performance and increase the computational burden.

Visualization of Generated AFC Configurations. Figure 6 presents a visual comparison between the original and generated AFCs in Beijing, analyzing the facility distribution across various land parcels. The results show that the generated plan not only enhances community diversity but also significantly improves the efficiency of aging-friendly facility distribution. Furthermore, our model optimizes the spatial arrangement of these facilities. For example, in the lower grid, areas highlighted in green as 15-minute walkable zones and outlined by red boxes now enjoy improved access to previously hard-to-reach facilities, such as hospitals and parks. This enhancement is facilitated by reconfigured roadways, including major urban renewal projects like subways and overpasses. Thus, FAP-CD offers urban planning experts valuable insights for facility planning and clearly demonstrates its potential to optimize spatial allocations.

Related Work

Diffusion Models. Diffusion models, as an emerging generative technique, have demonstrated exceptional capabilities in fields such as computer vision (Dhariwal and Nichol 2021), natural language processing (Li et al. 2022), and graph generation (Niu et al. 2020; Vignac et al. 2022). These models operate through two main phases: an initial noise addition phase and a reverse denoising phase that restores noisy data to complex distributions. The most perti-

nent works involve the application of diffusion models to graph generation tasks (Niu et al. 2020; Vignac et al. 2022; Huang et al. 2023). For instance, Niu et al.(2020) used score-based models with Gaussian noise for permutation-invariant graph, while Jo, Lee, and Hwang(2022) and recent studies (Vignac et al. 2022; Haefeli et al. 2022) explored the potential of diffusion models for generating node and edge-enhanced graphs and using multinomial noise for discrete models that preserve graph sparsity and enhance quality. Moreover, incorporating additional information as conditioning cues is crucial for targeted generation. For example, Zhou et al.(2023) designed a novel diffusion process to generate city mobility flows, which are guided by a volume estimator and influenced by regional features.

Inspired by the significant success of diffusion models, our study integrates age-friendly needs and urban region characteristics into our model to optimize AFCs planning. This study explores the capabilities of graph diffusion models in strategic planning for AFCs, aiming not only to develop configurations for community revitalization but also to provide insights into optimal geographic locations.

Automated Planning Generation. Traditional urban planning methods have encountered many challenges in practice, including high cost, low efficiency, lack of skilled facilitators, and low interest to participate (Abas et al. 2023). Recently, the remarkable success of deep generative learning has led researcher about automated planning generation to improve the efficiency of urban planning. For example, Wang et al.(2023) developed an instruction-based hierarchical planner (IHPlanner), which automatically captured the spatial hierarchy and planning dependencies between urban functional zones and land use allocation based on human instructions. Zheng et al.(2023) proposed a RL-based urban planning model that trained a value network to enhance spatial planning quality focusing on the completeness of the 15-minute city concept (Moreno et al. 2021), with the aim of improving the efficiency of transportation, services, and ecological systems in urban regions. However, these plans are either completed based on single human instructions or remain inherently time-consuming. Furthermore, they lack scalability in age-friendly service facilities and do not address issues of planning efficiency and fairness. Compared to these efforts, the FAP-CD model is more advanced in automation and the facilitation of age-friendly urban planning.

Conclusion

In this work, we propose a novel conditional graph diffusion model FAP-CD, which transforms the planning issue into a problem of graph generation by constructing a detailed walking graph of spatial relationships among facilities within a region. This model leverages fair-demand guidance enhancement and discrete graph structure adjustments along with a finely tuned graph noise prediction network to generate AFCs planning during the reverse denoising process. Extensive experiments and case studies validate FAP-CD’s effectiveness, enhancing efficient and equitable age-friendly urban renewal planning and supporting sustainable urban development amid rapid global aging.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) (Grant No. 62106274) and the Fundamental Research Funds for the Central Universities, Renmin University of China (Grant No. 22XNKJ24).

References

- Abas, A.; Arifin, K.; Ali, M. A. M.; and Khairil, M. 2023. A systematic literature review on public participation in decision-making for local authority planning: A decade of progress and challenges. *Environmental Development*, 46: 100853.
- Chen, X.; He, J.; Han, X.; and Liu, L.-P. 2023. Efficient and degree-guided graph generation via discrete diffusion modeling. *arXiv preprint arXiv:2305.04111*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Dwivedi, V. P.; and Bresson, X. 2020. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*.
- Eskandar, G.; Farag, Y.; Yenamandra, T.; Cremers, D.; Guirguis, K.; and Yang, B. 2023. Urban-StyleGAN: Learning to Generate and Manipulate Images of Urban Scenes. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, 1–8. IEEE.
- Ghods, A.; Zaharia, M.; Hindman, B.; Konwinski, A.; Shenker, S.; and Stoica, I. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *8th USENIX symposium on networked systems design and implementation (NSDI 11)*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Haefeli, K. K.; Martinkus, K.; Perraudin, N.; and Wattenhofer, R. 2022. Diffusion models for graphs benefit from discrete state spaces. *arXiv preprint arXiv:2210.01549*.
- Han, S.; and Sung, Y. 2021. A max-min entropy framework for reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 25732–25745.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, H.; Sun, L.; Du, B.; and Lv, W. 2023. Conditional diffusion based on discrete graph structures for molecular graph generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4302–4311.
- Jensen, J. H. 2019. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chemical science*, 10(12): 3567–3572.
- Jo, J.; Lee, S.; and Hwang, S. J. 2022. Score-based generative modeling of graphs via the system of stochastic differential equations. In *International conference on machine learning*, 10362–10383. PMLR.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*.
- Li, X.; Thickstun, J.; Gulrajani, I.; Liang, P. S.; and Hashimoto, T. B. 2022. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35: 4328–4343.
- Liu, Y.-F.; Dai, Y.-H.; and Luo, Z.-Q. 2013. Max-min fairness linear transceiver design for a multi-user MIMO interference channel. *IEEE Transactions on Signal Processing*, 61(9): 2413–2423.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35: 5775–5787.
- Menec, V. H.; Means, R.; Keating, N.; Parkhurst, G.; and Eales, J. 2011. Conceptualizing age-friendly communities. *Canadian Journal on Aging/La revue canadienne du vieillissement*, 30(3): 479–493.
- Moreno, C.; Allam, Z.; Chabaud, D.; Gall, C.; and Pratloug, F. 2021. Introducing the “15-Minute City”: Sustainability, resilience and place identity in future post-pandemic cities. *Smart cities*, 4(1): 93–111.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.
- Niu, C.; Song, Y.; Song, J.; Zhao, S.; Grover, A.; and Ermon, S. 2020. Permutation invariant graph generation via score-based generative modeling. In *International Conference on Artificial Intelligence and Statistics*, 4474–4484. PMLR.
- Organization, W. H. 2007. *Global age-friendly cities: A guide*. World Health Organization.
- Pardalos, P. M.; Migdalas, A.; and Pitsoulis, L. 2008. *Pareto optimality, game theory and equilibria*, volume 17. Springer Science & Business Media.
- Razaviyayn, M.; Hong, M.; and Luo, Z.-Q. 2013. Linear transceiver design for a MIMO interfering broadcast channel achieving max–min fairness. *Signal Processing*, 93(12): 3327–3340.
- Rudnicka, E.; Napierała, P.; Podfigurna, A.; Meczekalski, B.; Smolarczyk, R.; and Grymowicz, M. 2020. The World Health Organization (WHO) approach to healthy ageing. *Maturitas*, 139: 6–11.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, L.; Kong, X.; and Cheng, P. 2024. Supply-demand matching assessment of the public service facilities in 15-minute community life circle based on residents’ behaviors. *Cities*, 144: 104637.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.

- Sun, W.; and Yang, W. 2023. The Application of Rehabilitation Garden Design in Elderly Care Communities in China Against the Background of Aging. *Journal of Innovation and Development*, 3(3): 17–21.
- Vignac, C.; Krawczuk, I.; Siraudin, A.; Wang, B.; Cevher, V.; and Frossard, P. 2022. Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*.
- Vincent, P. 2011. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7): 1661–1674.
- Wai, H.-T.; Hong, M.; Yang, Z.; Wang, Z.; and Tang, K. 2019. Variance reduced policy evaluation with smooth function approximation. *Advances in Neural Information Processing Systems*, 32.
- Wang, D.; Fu, Y.; Wang, P.; Huang, B.; and Lu, C.-T. 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.
- Wang, D.; Wu, L.; Zhang, D.; Zhou, J.; Sun, L.; and Fu, Y. 2023. Human-instructed deep hierarchical generative learning for automated urban planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4660–4667.
- Wang, Y.; Gonzales, E.; and Morrow-Howell, N. 2017. Applying WHO’s age-friendly communities framework to a national survey in China. *Journal of Gerontological Social Work*, 60(3): 215–231.
- Wen, H.; Lin, Y.; Xia, Y.; Wan, H.; Wen, Q.; Zimmermann, R.; and Liang, Y. 2023. Diffstg: Probabilistic spatio-temporal graph forecasting with denoising diffusion models. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, 1–12.
- Xie, L. 2018. Age-friendly communities and life satisfaction among the elderly in urban China. *Research on aging*, 40(9): 883–905.
- Xu, Z.; and Zhou, X. 2024. CGAP: Urban Region Representation Learning with Coarsened Graph Attention Pooling. *arXiv preprint arXiv:2407.02074*.
- Yang, C.; Zhuang, P.; Shi, W.; Luu, A.; and Li, P. 2019. Conditional structure generation through graph variational generative adversarial nets. *Advances in neural information processing systems*, 32.
- Yang, Z.; Yu, B.; and Cheng, C. 2007. A parallel ant colony algorithm for bus network optimization. *Computer-Aided Civil and Infrastructure Engineering*, 22(1): 44–55.
- Ye, X.; Du, J.; and Ye, Y. 2022. MasterplanGAN: Facilitating the smart rendering of urban master plans via generative adversarial networks. *Environment and Planning B: Urban Analytics and City Science*, 49(3): 794–814.
- Yong, X.; and Zhou, X. 2024. MuseCL: Predicting Urban Socioeconomic Indicators via Multi-Semantic Contrastive Learning. *arXiv preprint arXiv:2407.09523*.
- You, J.; Ying, R.; Ren, X.; Hamilton, W.; and Leskovec, J. 2018. Graphrnn: Generating realistic graphs with deep autoregressive models. In *International conference on machine learning*, 5708–5717. PMLR.
- Yung, E. H.; Conejos, S.; and Chan, E. H. 2016. Social needs of the elderly and active aging in public open spaces in urban renewal. *Cities*, 52: 114–122.
- Zhang, B. H.; Lemoine, B.; and Mitchell, M. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 335–340.
- Zheng, Y.; Lin, Y.; Zhao, L.; Wu, T.; Jin, D.; and Li, Y. 2023. Spatial planning of urban communities via deep reinforcement learning. *Nature Computational Science*, 3(9): 748–762.
- Zhou, Z.; Ding, J.; Liu, Y.; Jin, D.; and Li, Y. 2023. Towards generative modeling of urban flow through knowledge-enhanced denoising diffusion. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, 1–12.