Adaptive Urban Planning: A Hybrid Framework for Balanced City Development

Anonymous ACL submission

Abstract

Urban planning faces a critical challenge in balancing city-wide infrastructure needs with localized demographic preferences, particularly in rapidly developing regions. Although existing approaches typically focus on top-down optimization or bottom-up community planning, only some frameworks successfully integrate both perspectives. Our methodology employs a two-tier approach: First, a deterministic solver optimizes basic infrastructure requirements in the city region. Second, four specialized planning agents, each representing distinct sub-regions, propose demographic-specific modifications to a master planner. The master planner evaluates and integrates these proposals, ensuring both efficiency and adaptability. We validate our framework on a newly created dataset of detailed region and sub-region maps from three cities in India, focusing on areas undergoing rapid urbanization. The results show that it can improve planning efficiency, better address local demographic needs, and scale for real-world deployment. Our work also identifies key challenges, including the trade-off between system efficiency and adaptability, as well as the complexities of handling diverse urban datasets. These insights contribute to a broader understanding of how multi-agent systems can enhance large-scale urban planning.

1 Introduction

Urban planning is a critical challenge for rapidly growing cities, requiring solutions that balance infrastructure, housing, transportation, and community needs. Traditional top-down approaches often lack adaptability to localized demands, limiting the creation of inclusive urban environments. In developing regions, there is a growing need for planning methods that integrate both large-scale infrastructure goals and neighborhood-specific perspectives, while enabling efficient execution by industry stakeholders (Arnstein, 1969; Forester, 1982).

India exemplifies the challenges of modern urban planning due to its demographic diversity, population density, and history of unplanned growth (Ranjan, 2023). The coexistence of historic and modern layouts leads to congestion, inadequate green spaces, and uneven resource distribution. Urbanization has strained municipal authorities and infrastructure providers, highlighting the need for scalable, data-driven frameworks that meet both regulatory and community priorities (Kumar and Prakash, 2016).

Large Language Models (LLMs) present a promising solution by simulating stakeholder perspectives—community members, policymakers, and planners alike (Wang et al., 2024b). LLM-driven decision-making can enhance workflows, reduce manual effort, and support more adaptive urban strategies. This is particularly relevant in India's diverse cities, where local priorities often vary.

We introduce a hybrid urban planning framework that leverages deterministic solvers and LLM-driven agents (Wang et al., 2024a; Huang et al., 2024) to address the challenges of planning in rapidly urbanizing cities. Our proposed method consists of two primary components: First, the deterministic solver ensures an equitable distribution of essential infrastructure city-wide. Second, specialized LLM agents are designated to represent four distinct sub-regions, incorporating demographic-specific needs. A "master planner" then evaluates and integrates these suggestions, ensuring the final plan aligns with city-wide goals while accommodating local preferences.

We evaluate our framework using data from three rapidly urbanizing Indian cities. Our results indicate that LLM agents effectively capture community diversity, enabling infrastructure planning that is both inclusive and efficient.

We tested our framework using data from three fast-growing cities in India. The results show that

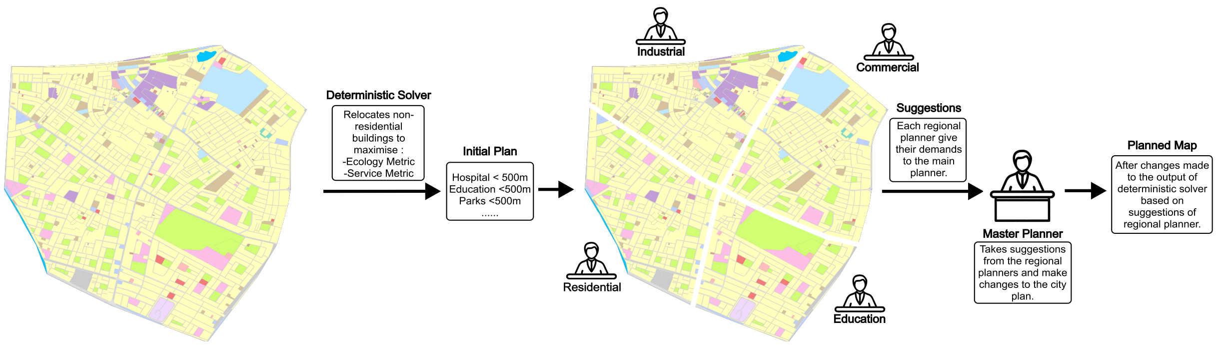

Figure 1: Workflow of the proposed urban planning framework. Integrating Deterministic Optimization, Regional Planner inputs, and Master Planner Coordination to achieve balanced and Area-Specific city layouts

our approach handles urban infrastructure needs while considering local demographic needs. Using LLM agents can represent diverse community needs, making urban planning more inclusive and efficient.

2 Background and Related Work

Large Language Models (LLMs) have demonstrated transformative potential in urban planning by automating complex tasks and facilitating participatory (Du et al., 2024) processes. Recent studies explore applying reasoning (Plaat et al., 2024) and planning (Valmeekam et al., 2023) capabilities in LLMs, such as GPT-4 (OpenAI et al., 2024) or Llama3 (Grattafiori et al., 2024), for specialized urban applications. For example, UrbanGPT (Li et al., 2024) integrates instruction-tuning and specialized decoders to enhance spatio-temporal forecasting, including traffic flow predictions.

One significant advancement in urban planning is a participatory framework that leverages LLMs, employing role-playing agents to emulate planners and residents (Zhou et al., 2024). In this framework, LLM-based agents collaboratively design land-use plans that balance community interests with expert constraints. The system includes an LLM agent acting as the planner and numerous agents representing residents with diverse profiles and backgrounds. The planner begins by proposing an initial land-use plan. Subsequently, a simulated fishbowl discussion mechanism is employed: a subset of residents actively discusses the plan while others act as listeners. The planner then revises the plan iteratively, incorporating resident feedback to achieve a more balanced and inclusive outcome.

Building on this framework, we propose an optimized approach that employs four regional agents (Chen et al., 2023) utilizing collaborative ability (Zhang et al., 2024) of LLMs, each representing a specific focus area that provides suggestions to a master planner. This master planner consolidates their suggestions and revises the city plan accordingly. By reducing the number of agents, our method significantly decreases computational overhead while maintaining robust decision-making. Furthermore, our approach prioritizes meeting fundamental needs before addressing region-specific demands, ensuring a more balanced and efficient planning process than previous methods.

3 Methodology

Our proposed methodology uses a two-tier framework for urban planning. In the first stage, a deterministic solver ensures that all residents have access to essential services and green spaces. The second stage introduces four region-specific planning agents, each advocating for the needs of their respective areas to a *master planner*. The master planner then evaluates these inputs and adjusts the city layout to harmonize fundamental requirements with demographic-specific needs. The overall structure of our pipeline is shown in Figure 1.

3.1 Deterministic Solver

The deterministic solver uses Genetic Algorithms (GA) (Forrest, 1996; Mirjalili and Mirjalili, 2019) to optimize urban layouts, ensuring equitable access to essential services and green spaces. The process begins with the original city plan, where essential services such as hospitals, schools, and businesses are assigned roles. To generate an initial configuration, a greedy solver creates an intermediate state by iteratively assigning elements to locations that maximize accessibility for residents. This intermediate layout provides a near-optimal starting point for the GA to refine further.

The GA uses two main steps: *mutation* and *selection*. Mutation creates new layouts by randomly swapping roles between locations, helping to explore different options and avoid premature convergence. For example, if a school and a hospital are swapped, it may improve *service accessibility* in one area while reducing it in another. Selection picks the best layouts using a tournament method. This favors high-performing layouts while maintaining diversity to explore alternative layouts.

Each layout's fitness is evaluated using two metrics:

- **Service Accessibility:** Measures the availability of schools, hospitals, and other essential services within 500 meters of residences.

- **Ecological Proximity:** Evaluates access to green spaces within 300 meters, reflecting urban livability and resident well-being.

The GA iteratively improves the layout across successive generations by ranking layouts based on fitness, retaining the top-performing configurations, and generating new ones through mutation. This process continues until fitness improvements plateau or a predefined number of generations is reached, indicating convergence.

The outputs of the deterministic solver include an optimized urban layout that maximizes accessibility to essential services and green spaces. Detailed formulation of the deterministic solver is given in Appendix C.

Algorithm 1 Deterministic Urban Layout Optimization

Input: Initial city layout L_0
Output: Intermediate layout L_{int}

$P \leftarrow$ InitializePopulation(L_0) ▷ Using greedy solver
for $g = 1$ to MaxGen **do**
  EvaluateFitness(P)
  $P' \leftarrow$ TournamentSelection(P)
  $P \leftarrow$ Mutate(P') ▷ Swap locations to generate new layouts
  if IsConverged(P) or $g =$ MaxGen **then**
    break
  end if
end for
return BestLayout(P)

3.2 Regional Adaptation via Dual-Planners

To refine the deterministic solver's output, we introduce a dual-planner approach that balances city-wide objectives with localized demographic needs. The city is divided into four sub-regions, each managed by a regional planner, an LLM tasked with advocating for its area-specific requirements. These regional planners generate proposals tailored to their assigned zone, which are then reviewed and integrated by a master planner. The master planner then reviews these proposals, integrating them into the city layout to balance local priorities with city-wide objectives.

3.2.1 Master Planner

It operates with a city-wide perspective, maximizing accessibility and achieving a balanced distribution of facilities. It prevents clustering in central areas and avoids over-dispersing facilities toward city edges, ensuring even coverage across the urban area.

Adhering to a minimal-change policy, the master planner makes essential layout adjustments only when necessary. These include reassigning vacant land for high-priority facilities, adding essential services in underserved areas, or swapping facility types to maintain efficient resource distribution. This strategy preserves the structural integrity of the city while meeting the overarching goals of urban planning.

3.2.2 Regional Planners

It complements the master planner by addressing the sub-region's specific demographic and functional needs. Each regional planner is designated to focus on one of four demographic roles: Industrial, Educational, Commercial, and Residential, chosen for their relevance to urban functionality. Further details about the demographics is shown in Appendix B.4.

4 Datasets

Our dataset consists of high-resolution thematic maps, essential for precise analysis of existing urban layouts. These maps are sourced from *Bhuvan AMRUT 4K* (Bhuvan, 2022) web services and provide detailed classifications of urban land use, including residential areas, government properties, commercial zones, transportation networks, green spaces, educational institutions, and other key infrastructure components.

Region	Metrics	Kanpur				Lucknow				Raipur			
		S1	**S2**	**S3***	**S3**	**S1**	**S2**	**S3***	**S3**	**S1**	**S2**	**S3***	**S3**
Region-1	**Service**	0.791	0.892	0.801	**0.916**	0.855	0.908	0.914	**0.943**	0.783	0.922	0.886	**0.948**
	Ecology	0.868	0.899	0.869	**0.899**	0.709	0.946	0.843	**0.946**	0.825	0.842	0.833	**0.842**
	Satisfaction	0.307	0.327	0.355	**0.489**	0.294	0.326	0.482	**0.683**	0.372	0.377	0.488	**0.615**
Region-2	**Service**	0.432	0.644	0.536	**0.710**	0.749	0.860	0.822	**0.895**	0.812	0.859	0.828	**0.926**
	Ecology	0.840	0.951	0.885	**0.951**	0.627	0.656	0.634	**0.656**	0.485	0.617	0.584	**0.617**
	Satisfaction	0.325	0.355	0.389	**0.507**	0.294	0.439	0.559	**0.765**	0.510	0.495	0.555	**0.653**

Table 1: Performance metrics are presented for multiple cities across various planning stages within two distinct geographic regions. The analysis spans four stages: Stage-1 (S1), the baseline configuration without optimization; Stage-2 (S2), an optimized layout derived from a deterministic solver; Stage-3* (S3*), outcomes obtained through AI-based planning; and Stage-3 (S3), the final integrated solution.

From the 238 available AMRUT city maps, we selected **Kanpur, Lucknow, and Raipur** for evaluation. These cities were chosen due to their diverse urban characteristics: *Kanpur* as a prominent industrial hub, *Lucknow* as an administrative and commercial center, and *Raipur* as a rapidly expanding urban region. Their unique planning challenges ensure that our framework is rigorously tested across different urban typologies, demonstrating its adaptability and effectiveness.

To extract the maps, we utilized the *Bhuvan API*, manually specifying coordinates for each target area. Subsequently, we applied connected component analysis with optimized parameters to segment regions based on land-use types. This structured extraction method allows us to generate high-resolution inputs for our hybrid planning model, ensuring accurate and scalable urban development analysis. Further details regarding the extraction are given in Appendix D.

5 Evaluation

To assess the effectiveness of our proposed framework, we use three key metrics: *Service Accessibility*, *Ecological Coverage*, and *Resident Satisfaction*. Together, these metrics assess the accessibility of public services, the availability of green spaces, and the fulfillment of residents' demographic-specific needs, emphasizing the framework's ability to create accessible, ecologically balanced, and resident-centric urban environments. These metrics are detailed in Appendix B

6 Results

Our proposed method demonstrates consistent improvements across all evaluated regions, as reflected in the three key metrics Service Accessibility, Ecological Coverage, and Resident Satisfaction—throughout the planning stages. The **Stage 1** represents the extracted baseline, identifying disparities in accessibility and sustainability, highlighting the necessity for an integrated planning approach. **Stage 2** introduces the deterministic solver, leading to notable gains in both Service Accessibility and Ecological Coverage. This step ensures a more balanced distribution of essential services and green spaces, forming a strong foundation for livable urban environments. In **Stage 3**, we incorporate inputs from specialized regional planning agents, with coordination by the master planner. This stage further refines all metrics by addressing localized demographic needs while maintaining city-wide balance, resulting in a substantial increase in Resident Satisfaction. To illustrate the necessity of the deterministic solver in our pipeline, we also evaluate **Stage 3***, where the LLM-based planner is applied directly to the baseline without the deterministic solver to show the importance of usage of AI based planning.

7 Conclusion

Our work presents a hybrid urban planning framework that optimizes city-wide infrastructure with localized demographic needs. By employing a two-tier methodology consisting of a deterministic solver and region-specific planning agents, our approach balances functional efficiency and community-specific requirements. Our results on diverse data from three rapidly urbanizing Indian cities demonstrates notable improvements in Service Accessibility, Ecological Coverage, and Resident Satisfaction across successive planning stages. The results highlight the advantages of combining systematic optimization with adaptive regional planning to create sustainable, inclusive, and livable urban environments.

8 Limitations

As the number of regional planners increases, conflict resolution between sub-regions becomes more complex, making it harder for the master planner to integrate diverse preferences effectively. Additionally, the computational demands of running multiple LLM agents may limit scalability, particularly for larger cities with many sub-regions. Furthermore, the framework's generalization ability may be constrained, as its effectiveness could vary across cities with different urban characteristics or data quality. Additionally, due to the unavailability of economic data for the regions, economic factors were not incorporated into the planning or evaluation processes, which may limit the comprehensiveness of the results.

References

Sherry R. Arnstein. 1969. A ladder of citizen participation. *Journal of the American Institute of Planners*, 35(4):216–224.

Bhuvan. 2022. Amrut.

Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2023. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. *Preprint*, arXiv:2308.10848.

Jiaxin Du, Xinyue Ye, Piotr Jankowski, Thomas W. Sanchez, and Gengchen Mai. 2024. Artificial intelligence enabled participatory planning: a review. *International Journal of Urban Sciences*, 28(2):183–210.

John Forester. 1982. Planning in the face of power. *Journal of the American Planning Association*, 48(1):67–80.

Stephanie Forrest. 1996. Genetic algorithms. *ACM computing surveys (CSUR)*, 28(1):77–80.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, and Akhil Mathur. 2024. The llama 3 herd of models. *Preprint*, arXiv:2407.21783.

Xu Huang, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang, Ruiming Tang, and Enhong Chen. 2024. Understanding the planning of llm agents: A survey. *arXiv preprint*, arXiv:2402.02716.

Ashok Kumar and Poonam Prakash. 2016. Public participation in planning in india.

Zhonghang Li, Lianghao Xia, Jiabin Tang, Yong Xu, Lei Shi, Long Xia, Dawei Yin, and Chao Huang. 2024. Urbangpt: Spatio-temporal large language models. *Preprint*, arXiv:2403.00813.

Seyedali Mirjalili and Seyedali Mirjalili. 2019. Genetic algorithm. *Evolutionary algorithms and neural networks: theory and applications*, pages 43–55.

Carlos Moreno, Zaheer Allam, Didier Chabaud, Catherine Gall, and Florent Pratlong. 2021. Introducing the “15-minute city”: Sustainability, resilience and place identity in future post-pandemic cities. *Smart Cities*, 4(1):93–111.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, and Florencia Leoni Aleman. 2024. Gpt-4 technical report. *Preprint*, arXiv:2303.08774.

Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, Niki van Stein, and Thomas Back. 2024. Reasoning with large language models, a survey. *Preprint*, arXiv:2407.11511.

Nalin Ranjan. 2023. Economic planning in practice: Indian experience and niti aayog.

Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. 2023. On the planning abilities of large language models : A critical investigation. *Preprint*, arXiv:2305.15771.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024a. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6).

Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Stephen W. Huang, Jie Fu, and Junran Peng. 2024b. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. *Preprint*, arXiv:2310.00746.

Jintian Zhang, Xin Xu, Ningyu Zhang, Ruibo Liu, Bryan Hooi, and Shumin Deng. 2024. Exploring collaboration mechanisms for llm agents: A social psychology view. *Preprint*, arXiv:2310.02124.

Zhilun Zhou, Yuming Lin, Depeng Jin, and Yong Li. 2024. Large language model for participatory urban planning. *arXiv preprint arXiv:2402.17161*.

A Implementation Details

We used GPT4o-Mini[1] for both master and regional planners. Together, the *master* and *regional* planners foster a coordinated city plan that upholds

[1]https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/

structural integrity and ecological balance while adapting to the particular needs of different areas. Due to financial constraints, all our results are evaluated for a single run.

B Evaluation Metrics

B.1 Service Accessibility

The *Service Accessibility* metric evaluates how efficiently essential services are distributed within residential areas. It measures the proportion of essential services (e.g., education, healthcare, workplaces, shopping, and recreation) accessible within a 500-meter radius of resident's homes, with values ranging from 0 to 1, where higher values represent better service accessibility.

The metric is computed as follows:

1. For each resident m, the minimum distance $d(m, j)$ to access a facility of type j is determined:

$$d(m, j) = \min_{1 \leq i \leq k_j} \textit{EucDis}(L_m, P_{j,i}) \quad (1)$$

where L_m is the resident's location, and $P_{j,k}$ denotes the k-th facility of type j.

2. The overall *Service Accessibility* metric aggregates these values for all residents n_m and service types n_j:

$$\text{Service} = \frac{1}{n_m} \sum_{m=1}^{n_m} \frac{1}{n_j} \sum_{j=1}^{n_j} \mathbb{1}[d(m, j) < 500] \quad (2)$$

where $\mathbb{I}[d(m, j) < 500]$ is an indicator function, returning 1 if the distance is less than 500 meters and 0 otherwise.

B.2 Ecological Coverage

The *Ecological Coverage* metric measures the availability of parks and green spaces, which play a critical role in promoting the health and well-being of urban residents. This metric evaluates the proportion of residents who live within a 300-meter radius of parks or open spaces, aligning with global standards for urban green accessibility.

The metric calculation is as follows:

1. The *Ecological Service Area (ESA)* is defined as the combined buffer zones extending 300 meters around each park or green space:

$$\text{ESR} = \bigcup_{i=1}^{k} \text{Buffer}(P_{\text{park},i}, 300) \quad (3)$$

where $P_{\text{park},k}$ represents the k-th park or green space.

2. The *Ecological Coverage* metric is computed as the proportion of residents L_m located within the ESA:

$$\text{Ecological Coverage} = \frac{1}{n_m} \sum_{m=1}^{n_m} \mathbb{1}[L_m \in \text{ESA}], \quad (4)$$

where $\mathbb{1}[L_m \in \text{ESA}]$ returns 1 if the resident is within the buffer zone and 0 otherwise.

B.3 Satisfaction

The *Satisfaction* metric evaluates how effectively the urban layout fulfills the specific needs of residents in different demographic sub-regions. Unlike the previous two metrics, this metric considers the unique requirements of each sub-region, such as educational facilities in academic zones or healthcare in posh neighborhoods, ensuring a more customized urban planning approach. This metric ranges from 0 to 1, with higher values indicating a better alignment between urban layouts and resident specific needs.

1. Each resident m in a sub-region is assigned a set of prioritized needs J_m, representing 3-5 most critical land-use categories for that demographic goal. The satisfaction level for an individual resident m is calculated as:

$$S_m = \frac{1}{n_j} \sum_{j \in J_m} \mathbb{1}[d(m, j) < 800], \quad (5)$$

where $d(m, j)$ is the minimum distance from the resident to a facility of type j, and $\mathbb{1}[d(m, j) < 800]$ indicates whether this distance is within 800 meters.

2. The overall *Satisfaction Metric* is then computed by aggregating the satisfaction values across all residents n_m in the region:

$$\text{Satisfaction} = \frac{1}{n_m} \sum_{m=1}^{n_m} S_m \quad (6)$$

Together, the *Service Accessibility*, *Ecological Coverage*, and *Satisfaction* metrics evaluate urban layouts by balancing accessibility, environmental sensitivity, and demographic inclusivity. These metrics demonstrate our framework's alignment with the concept of a "15-minute city" (Moreno et al., 2021), ensuring essential services and green spaces are within walking or cycling distance, fostering sustainable and resident-focused urban spaces for rapidly urbanizing regions.

B.4 Demographics

- **Industrial Zones:** Prioritize *factories, warehouses, and logistics hubs*, optimizing workforce accessibility and supply chain efficiency.
- **Educational Zones:** Ensure the strategic placement of *schools, universities, and student housing* to foster knowledge hubs.
- **Commercial Zones:** Designate spaces for *offices, retail hubs, and business infrastructure*, promoting economic activity.
- **Residential Zones:** Optimize *housing, community amenities, and daily services* to enhance urban livability.

C Formulation of Deterministic Solver

- S_0 : Initial game state (mapping of regions to roles, initially set to "None").
- P : Set of players, representing the non-residential types to be assigned to regions.
- C : Centroids dictionary (coordinates of the center of each region).
- L : Move limits dictionary, where $L[p]$ denotes the number of assignments allowed for player p.
- S_{final} : Final optimized game state after the genetic algorithm process.
- r^* : Region selected for assignment based on the highest return value in the greedy phase.
- $\mathcal{P}$: Population of layout configurations in the genetic algorithm.
- N : Population size for the genetic algorithm.
- G : Number of generations in the genetic algorithm.
- k : Number of top layouts selected for the next generation.
- S^* : Layout with the highest fitness value after the genetic algorithm optimization.
- `calculate_return` : Function used to calculate the return value for assigning a region to a player based on service and ecology metrics.
- `fitness_function` : Function that evaluates the fitness of a layout based on service accessibility and ecological proximity.
- `mutate` : Function that applies random swaps to create new variations of a layout.
- `initialize_population` : Function that generates the initial population for the genetic algorithm using random swaps.

1. **Input:**
   - Initial game state S_0 (mapping of regions to roles).
   - Players P (list of non-residential types to assign).
   - Centroids dictionary C (coordinates of region centers).
   - Move limits L (number of assignments allowed for each player).
2. **Phase 1: Greedy Assignment.**
   (a) Initialize the game state $S \leftarrow S_0$.
   (b) While unassigned regions remain or $\exists\, p \in P$ such that $L[p] > 0$:
      i. For each player $p \in P$:
         A. If $L[p] = 0$, continue to the next player.
         B. Find the region r^* that maximizes the return value using `calculate_return`.
         C. Assign r^* to p: $S[r^*] \leftarrow p$.
         D. Decrease the move limit: $L[p] \leftarrow L[p] - 1$.
   (c) Output intermediate layout S_{greedy}.
3. **Phase 2: Genetic Algorithm Optimization.**
   (a) Initialize population $\mathcal{P}$ of size N using S_{greedy} and random swaps.
   (b) For each generation $g \in \{1, 2, \ldots, G\}$:
      i. Evaluate the fitness of each layout $S \in \mathcal{P}$ using `fitness_function`.
      ii. Select the top k layouts to carry forward.
      iii. Mutate layouts to create $N - k$ new layouts and add to the next generation.

(c) Output the layout S^* with the highest fitness in $\mathcal{P}$.

4. **Output:** $S_{\text{final}} \leftarrow S^*$, the layout maximizing service accessibility and ecological proximity.

D Extraction of infomation from image

We used a predefined color legend (Table 2) to extract information from the image to categorize various land regions on a geographic map. Each land-use type, such as Residential, Business, Educational, and others, was associated with a specific color, enabling efficient map segmentation based on these color codes. The map, as illustrated in Figure 2, was first converted into the HSV (Hue, Saturation, Value) color space, facilitating easier color segmentation by defining precise color ranges for each land type. This transformation allowed for identifying pixels corresponding to specific regions, effectively distinguishing different land-use categories.

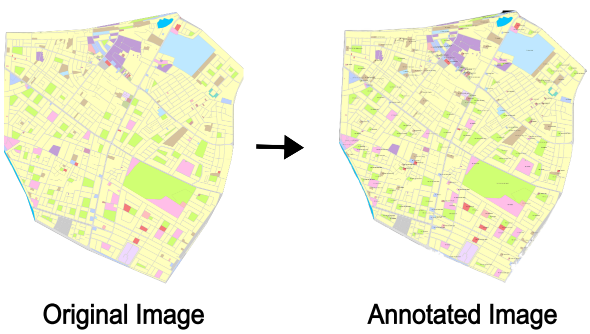

Figure 2: Conversion of the original map image into the annotated map image showcasing land-use categorization through color-based segmentation

We then performed connected component analysis on the map to identify distinct regions. Each identified region was labeled, and its area and centroid were calculated. Only regions with a minimum area threshold were considered for further analysis. For each valid region, a mask was generated, and relevant details, including the land type, label, area, and centroid, were recorded. The data was then organized into a structured format, making it suitable for further urban planning, environmental assessment, or other relevant applications. This process enabled the efficient extraction and categorization of land-use regions from the map, supporting various spatial analysis tasks.

Colour	Type	Structures
	Residential	Houses, Apartments, Villas
	State Govt. Property	Government offices, Emergency services (e.g., police, fire stations)
	Business	Commercial buildings, Office spaces, Retail stores
	Public Utilities	Water treatment plants, Sewage systems, Electricity stations
	Shops and Market	Markets, Grocery stores, Shopping malls
	Educational	Schools, Universities, Libraries, Educational centers
	Vacant Land	Open fields, Unused land
	Park and Open Space	Public parks, Playgrounds, Green spaces
	Hospital	Hospitals, Clinics, Healthcare facilities

Table 2: Pre-defined color legend for categorizing land-use types, associating each color with specific structures to support map segmentation and spatial analysis

E Sub-Region Extraction

A mask image representing predefined regions on a map is utilized to filter and validate centroids of land-use regions to extract sub-regions. The mask image, as shown in Figure 3, is loaded in grayscale, where white areas correspond to valid regions of interest. The dimensions of the mask image are verified to ensure proper alignment with the spatial data. A function is then defined to check if a given region's centroid falls within the mask's white area. This is done by converting the centroid's coordinates to integers and checking if they lie within the image boundaries and if the pixel at that location is white (indicating a valid region).

The filtering process is applied to the centroids of all regions in the dataset, and only those regions whose centroids fall within the white area of the mask are retained. This ensures that only relevant

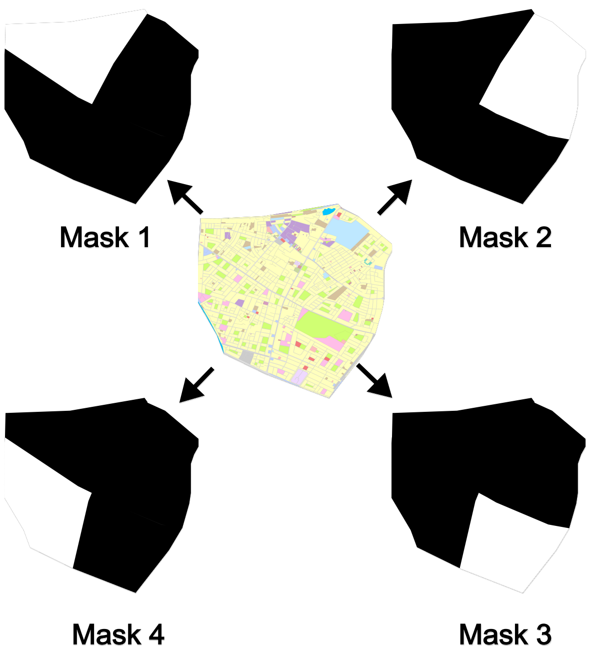

Figure 3: Utilization of mask images to validate centroids of land-use regions, showing the original map and corresponding masks defining valid sub-regions

regions located within predefined valid areas are considered for further analysis or processing. The result is a refined dataset containing only the regions that meet the criteria, enabling more focused and accurate urban or environmental assessments.