Cureus Journal of Computer Science

Open Access Original Article

An AI Framework for Generating and Simulating Public Transportation

Youssef Mohamed Abdelhamid Mohamed [1] , Zhiyuan Chen [1]

1. School of Computer Science, University of Nottingham Malaysia, Semenyih, MYS

Corresponding author: Zhiyuan Chen, zhiyuan.chen@nottingham.edu.my

Received 06/17/2025
Review began 06/23/2025
Review ended 09/12/2025
Published 09/15/2025

DOI:
https://doi.org/10.7759/s44389-025-07783-0

Abstract

The design of public transit networks is a complex and computationally challenging problem faced by urban planners worldwide, which traditionally focuses on incremental improvements or hypothetical replacements of existing public transit networks, neglecting data-scare and non-urbanised regions. This work introduces the Modular Adaptive Network for Transit Infrastructure Solutions (MANTIS), an artificial intelligence-based framework for transit network design, addressing the multi-objective challenge of optimising station locations, route structures, and service frequencies, with a focus on full control for urban developers, and adaptability for data-scarce regions. The proposed methodology integrates three key algorithmic components: K-means clustering for data-driven station placement, a gravity model for estimating travel demand between candidate stations, and a genetic algorithm for simultaneous optimisation of routes and frequencies across multiple lines, using a range of supported and optional geographic and passenger data. The framework is evaluated using agent-based simulation in Simulation of Urban MObility, enabling quantitative analysis of network performance. By combining clustering, demand modelling, and evolutionary optimisation, this approach provides a systematic and modular foundation for developing efficient, demand-responsive transit networks in rapidly evolving urban environments.

Categories: AI applications, Optimization Algorithms, Planning and Scheduling
Keywords: public transportation, transit network design, route optimisation, frequency setting, sumo

Introduction

Designing efficient public transit networks is a central challenge for urban planners and transportation authorities worldwide. The Transit Network Design and Frequency Setting Problem (TNDFSP) is a classic example of a large-scale, combinatorial optimisation problem, requiring the simultaneous determination of optimal station locations, route configurations, and service frequencies across complex urban environments. This problem is not only NP-hard, but also highly multi-objective, as it must balance user convenience, operational costs, and network coverage while adapting to evolving patterns of demand and urban growth [1].

As populations expand and mobility needs evolve, planners must create transit systems that are not only efficient and accessible but also adaptable to future changes. However, most existing approaches to transit network design rely heavily on historical travel data and incremental improvements to legacy systems, which can limit their applicability in developing or data-scarce regions. This gap is especially significant as many cities worldwide seek to build new networks from the ground up or overhaul outdated infrastructure to better serve contemporary demands.

This study aims to address these challenges by developing a flexible, data-driven framework for transit network design that does not depend on historical Origin-Destination (OD) matrices. Leveraging synthetic demand estimation and advanced optimisation techniques, our approach enables the joint optimisation of station locations, route structures, and service frequencies in diverse urban contexts. We apply this framework to both Greater London and Greater Kuala Lumpur, illustrating its adaptability and effectiveness across contrasting metropolitan environments. By systematically comparing different design priorities, such as demand alignment, population coverage, and connectivity to key destinations, our research provides actionable insights for planners seeking to create robust, user-centred transit systems for the cities of tomorrow.

TNDFSP is a foundational challenge in urban transportation planning, characterised by its combinatorial complexity and the need to balance operational efficiency, passenger satisfaction, and equity [1]. Early approaches to TNDFSP relied on heuristic and rule-based methods, but the limitations of these techniques in handling large-scale, multi-objective problems have led to the widespread adoption of metaheuristic and AI-based optimisation frameworks.

Genetic algorithms (GAs) have emerged as a particularly effective tool for transit network optimisation, owing to their ability to explore vast solution spaces and accommodate multiple, often conflicting objectives. Król and Król [2] demonstrated the application of GAs to metro network design, achieving significant reductions in construction costs and improvements in travel efficiency. Lammaihri et al. [3] further advanced this by integrating ant colony principles with GAs in a two-stage approach, achieving notable reductions in average trip durations. The flexibility of GAs to incorporate context-specific mutation operators and hybridise with other metaheuristics has been highlighted in several studies [4,5], with hybrid GA-particle swarm optimization and GA-simulated annealing models demonstrating superior performance in escaping local optima and improving solution quality.

Beyond GAs, nature-inspired algorithms such as Ant Colony Optimisation (ACO) and Bee Colony Optimisation (BCO) have shown promise in addressing the TNDFSP. Supangat and Soelistio [6] applied ACO to optimise bus routes in Jakarta, leveraging pheromone decay and Point of Interest (POI) density heuristics to reduce transfers and enhance connectivity. Nikolić and Teodorović [7] introduced a heterogeneous BCO variant that adaptively balanced exploration and exploitation, resulting in substantial reductions in passenger transfers in Belgrade's transit system. Comparative studies have found BCO to be particularly effective for minimising transfers [7], while GAs often excel in optimising fleet allocation and overall network structure [8].

Hybrid algorithms that synergise global and local search strategies have gained traction for their ability to address the multi-faceted nature of transit network design. Park et al. [9] combined a logit-based mode-choice model with evolutionary optimisation to enhance transit equity, while Owais et al. [10] used GIS-integrated GAs to improve network connectivity in Cairo by explicitly considering spatial constraints. These approaches underscore the importance of integrating demand responsiveness, equity, and spatial realism in network design methodologies.

Gravity models remain a cornerstone for estimating OD matrices, essential for both network design and evaluation. The classical gravity model, inspired by Newtonian physics, relates trip distribution to the attractiveness of origin and destination zones and the spatial impedance between them [11]. Recent advances have incorporated machine learning to enhance OD estimation, with Kalman filtering of Bluetooth-derived travel times enabling real-time updates [12], and deep learning models reducing prediction errors in large-scale urban rail systems [13]. Synthetic OD generation techniques have become increasingly important for data-scarce regions, enabling the application of advanced optimisation even where historical demand data is lacking [14].

Clustering algorithms, particularly K-means, play a vital role in determining optimal station locations by analysing spatial demand and activity patterns. Chen et al. [15] used K-means to segment transit travellers in Taipei, enabling targeted service improvements, while Supangat and Soelistio [6] combined mean shift clustering with ACO to optimise bus stop locations and routes in West Jakarta, achieving high accessibility and efficiency. Hierarchical and density-based clustering methods have also been explored for handling irregular spatial distributions and multi-scale network requirements [16].

Despite these methodological advances, several research gaps persist. Many studies continue to rely on static OD matrices derived from historical data, limiting adaptability in rapidly evolving or data-poor urban contexts [13]. The majority of research focuses on optimising or replacing existing networks rather than designing new systems from first principles, and benchmarking against real-world systems remains limited. There is a growing need for frameworks that integrate synthetic OD generation, adaptive optimisation, and robust simulation-based evaluation to enable the design of flexible, equitable, and future-ready transit networks, frameworks that can provide full practical control to city planners to meet their needs, without being restricted by historical data, allowing for operations on newly urbanising regions.

To address these gaps, this paper introduces Modular Adaptive Network for Transit Infrastructure Solution (MANTIS), a framework distinguished by several key contributions that differentiate it from prior work. Unlike studies that focus on isolated problems such as route optimisation or station placement, MANTIS provides a holistic, end-to-end pipeline that handles every stage from station placement and synthetic demand modelling to the joint optimisation of routes and frequencies. Its core novelty lies in this integrated approach, which is built upon a principle of data flexibility. By leveraging population or POI to synthetically model demand, MANTIS overcomes the critical barrier of dependency on historical OD, enabling network design in data-scarce or newly urbanising regions where such data is often unavailable. Furthermore, the framework is designed not as a rigid model, but as a modular, planner-centric tool that explicitly allows urban planners to prioritise different data inputs, tailoring the optimisation to specific local contexts, policy goals, and data availability. By integrating these features, MANTIS offers a practical and adaptable approach to designing effective, user-centred transit systems from first principles.

Materials And Methods

Dataset

This study leverages two comprehensive, multi-source datasets representing Greater London and Greater Kuala Lumpur to evaluate the proposed AI-driven transit network design framework. Each dataset integrates spatial, demographic, and mobility information, enabling robust benchmarking across diverse urban contexts.

The Greater London dataset is defined by the Greater London Authority boundary and incorporates:

1) Administrative boundaries from OpenStreetMap [17];

2) Population density on a hexagonal grid from the Kontur Population Dataset 2023 [18];

3) OD matrices using monthly rail journey counts from the Office of Rail and Road (2021-2022) [19];

4) POI data from the Overture Maps Foundation (2024) [20];

5) Travel times for all OD pairs, sourced from the Transport for London API [21].

The Greater Kuala Lumpur dataset is constructed using a custom-made bounding box encompassing all Rapid Rail stations and includes:

1) Administrative boundaries from OpenStreetMap [17];

2) Population density on a hexagonal grid (Kontur Population Dataset 2023) [18];

3) OD matrices from Prasarana Malaysia, with monthly and daily journey counts [22];

4) POI data exported from OpenStreetMap by the Humanitarian OpenStreetMap Team [23];

5) Travel times from the Moovit Journey Planner [24].

All spatial data layers were harmonised to a common coordinate reference system (EPSG:3857) and aggregated using the H3 hexagonal tessellation, which provides uniform spatial units for analysis. Feature engineering included spatial joins to assign administrative regions and POI counts, geocoding OD entries to hexagons, and z-score normalisation of population, POI, and demand values to ensure comparability and support machine learning algorithms.

This integrated, multi-source approach enables rigorous, spatially explicit evaluation of transit network designs and supports reproducibility and transferability of the methodology to other urban regions.

Proposed solution

In response to these limitations, this paper proposes the MANTIS framework for transit network design that integrates three core algorithmic components, in order to create full transit network solutions for the TNDFSP instead of individual aspects of a network.

MANTIS addresses these limitations by proposing a modular, highly adaptable AI-driven framework for transit network design that can utilise a wide range of data with full control being offered to the urban developers. This is done by creating a core modular geographic dataset, allowing urban developers to pick which data should be used, and utilising it throughout the entire process. Assuming all data is provided, this geographic dataset would include Population Density Data, POI Data, and Historical OD Matrix Data.

This modular database gives the freedom to urban developers to determine what is most important to them when designing a transit network. As long as at least one of these 3 datatypes are available, the framework is capable of generating a solution, instead of only being able to proposing solutions for existing urban regions. The public transportation OD Matrix can even be replaced with a Matrix for car traffic that shows the trips between important points in the region. Using this modular database, the framework then generates a transit network using:

1) K-means clustering for data-driven station placement, which can flexibly incorporate any combination of demand, population, or POI data, allowing planners to prioritise the data sources most relevant to their context, or make up for a lack of historical data.

2) Gravity model for robust, scalable demand estimation, capable of generating synthetic OD matrices when empirical data is unavailable, thus enabling application in both data-rich and data-poor regions.

3) GA for simultaneous route and frequency optimisation, supporting a wide variety of constraints

including minimum/maximum line lengths, station coverage requirements, transfer penalties, and operational limits.

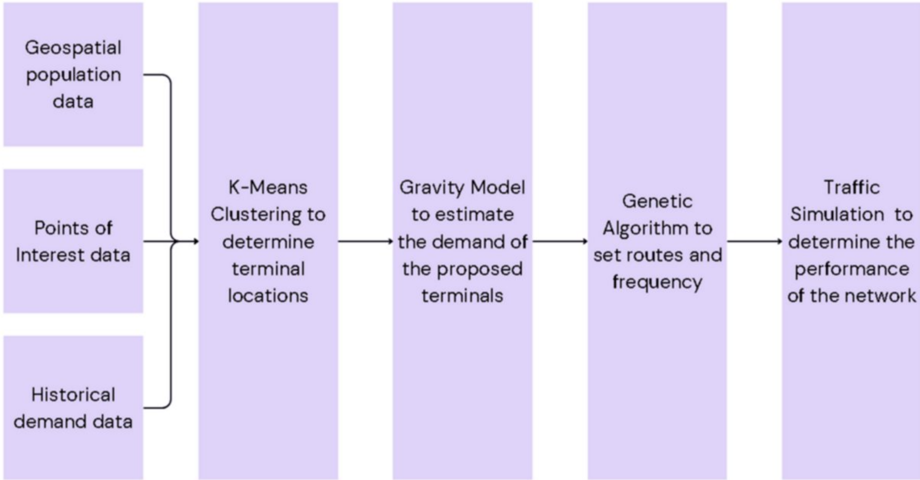

FIGURE 1: MANTIS Framework Design

MANTIS, Modular Adaptive Network for Transit Infrastructure Solution

As shown in Figure *1*, each layer of the framework uses the results of the layer before, merging to create a full network. The framework utilises each of the potential data sources, and begins with generating terminal locations, which are used to estimate the expected proportional demand of the network, and finalise the solution by setting routes and frequencies, ensuring that at least one possible path is guaranteed between each two terminals, incuding transfers. Afterwards, city planners can use their syntethic OD matrices and newly generated network to test the performance of the proposed solution on Simulation of Urban MObility (SUMO), using their own passenger figures.

K-means clustering

The core of the station siting process is a data-adaptive K-means clustering algorithm. K-means was selected not only for its extensive validation in spatial analysis literature for transit planning but also for its specific suitability to this problem. This allows the framework to build upon a proven method, focusing the contribution on the overall integrated pipeline. K-means also offers a combination of computational efficiency and scalability, which are critical for handling large-scale urban datasets. Furthermore, its resulting centroids have a direct and interpretable analogue as potential station locations. While other methods like the density-based DBSCAN can struggle with the varying densities inherent in urban centres versus suburban areas, and hierarchical clustering becomes computationally prohibitive for the number of data points in a metropolitan region, K-means provides a robust and practical foundation for the station placement module.

Algorithm 1 Calculate Regional K Values

```
Extract numerical values from feature series
MIN_K ← 5                                          ▷ Minimum number of clusters
MAX_K ← 50                                         ▷ Maximum number of clusters
if 97th percentile of feature values < 0.001 then
    Base K ← max(MIN_K, 15% of total data points in region)
else
    Base K ← max(MIN_K, 97th percentile × base multiplier)
end if
Start ← max(MIN_K, Base K −5)
End ← min(MAX_K, Base K +15)
Generate K list with values from Start to End, incrementing by 2
Remove K values that exceed number of data points in region
if K list is empty then
    return min(max(MIN_K, Base K), number of data points)
else
    return K list
end if
```

FIGURE 2: Regional K Calculation Pseudocode

Figure 2 demonstrates how the algorithm operates regionally, applying K-means independently within administrative or functional boundaries. This regionalisation ensures that clusters are sensitive to local demand structures and prevents high-density areas from monopolising station allocation, a common pitfall in global clustering. The number of clusters K is not arbitrarily fixed; instead, it is dynamically determined for each region using a data-driven heuristic that balances minimum service guarantees with empirical demand patterns. Combined with a guaranteed minimum number of centroids, this approach ensures balanced coverage across the region.

Algorithm 2 Spatial Distancing for Low-Demand Centroids

```
function SatisfiesDistance(point, existing_points, min_distance)
    for all p in existing_points do
        if (point.x − p.x)^2 + (point.y − p.y)^2 < min_distance then
            return false
        end if
    end for
    return true
end function
```

FIGURE 3: Spatial Distancing Pseudocode

To further guarantee practical service coverage and avoid inefficient clustering, the algorithm incorporates spatial distancing constraints (shown in Figure 3), particularly for low-demand centroids. The algorithm achieves this by calculating the Euclidean distance between each potential centroid location and all existing stations, rejecting any placement that falls within the predefined minimum distance threshold.. This ensures that stations are not over-concentrated and that even low-density areas receive equitable service.

Algorithm 3 Boundary Constraint for Centroids

```
function ConstrainToBounds(point, bbox)
    x ← point.x
    y ← point.y
    minx ← bbox[0], miny ← bbox[1], maxx ← bbox[2], maxy ← bbox[3]
    x ← max(minx, min(x, maxx))                ▷ Constrain x coordinate
    y ← max(miny, min(y, maxy))                ▷ Constrain y coordinate
    return (x, y)
end function
```

FIGURE 4: Boundary Constraint Pseudocode

Lastly, to ensure all centroids remain within the designated service area and maintain practical deployment feasibility, the algorithm implements boundary constraints (shown in Figure 4) that prevent station placement outside predefined geographical limits by clamping each centroid's coordinates to the nearest boundary edge when they exceed the defined limits, using a min-max operation that constrains both x and y coordinates to fall within the specified bounding box. This constraint is particularly important when centroids are iteratively updated during the optimization process, as it ensures stations remain within accessible and serviceable regions.

The strength of this approach lies in its adaptability and operational transparency. By allowing the inclusion or omission of any data layer, and by exposing all key parameters for user control, the methodology empowers urban developers to tailor the clustering to local priorities, data realities, and planning constraints. This adaptability, combined with rigorous normalisation, regionalisation, and constraint enforcement, ensures that station placement is both data-driven and policy-responsive, a critical advance over static or heuristic-based methods.

Gravity model

For passenger flow estimation, a gravity model was chosen due to its status as a cornerstone of transportation planning for estimating OD matrices. Its widespread adoption and well-understood principles allow this framework to build upon a reliable foundation, shifting the focus to MANTIS's primary goal of enabling network design in data-scarce regions. The model's strength lies in its interpretability and flexibility. By relating trip generation to the 'attractiveness' of zones and the 'impedance' of distance, parameters can be calibrated to diverse urban contexts. This approach supports the generation of synthetic OD matrices when historical data is unavailable, making it an indispensable component for fulfilling the framework's objective of operating in developing or previously un-serviced areas. The model predicts the volume of movement F_{ij} between every OD pair (i, j) as:

$$F_{ij} = k \times \frac{P_i^{\alpha} \times A_j^{\beta}}{d_{ij}^{\gamma}}$$

where P_i and A_j are the transformed demand scores of the origin and destination terminals, d_{ij} is the Euclidean distance in kilometres, *k* is a scaling factor, and *α*, *β*, and *γ* are exponents controlling the sensitivity to origin demand, destination demand, and distance decay, respectively. This formulation encapsulates the principle that interaction increases with terminal importance and decreases with distance, which is calculated using Euclidean distances, and its parameters provide the flexibility to adapt to a wide range of urban contexts and planning objectives.

Parameter calibration is performed through a two-stage process. Initial values for *α*, *β*, *γ*, and *k* are heuristically set based on the median nonzero distance in the study area and conservative bounds, reflecting diminishing returns for increased demand and superlinear distance decay. These are then refined using the Nelder-Mead optimisation algorithm, minimising the relative error between estimated flows and a reference OD matrix. When empirical OD data is unavailable, a synthetic matrix is generated using a simplified gravity formula:

$$\text{existing\_OD}_{ij} = 100 \times \frac{P_i \times A_j}{d_{ij}^{0.5}}$$

This approach ensures that calibration is always feasible, regardless of data constraints. The optimisation is bounded to maintain interpretability, and a relative error metric is used to prevent large-volume flows from dominating the calibration.

Once calibrated, the gravity model is systematically applied to all terminal pairs, generating a complete OD matrix. The methodology enforces zero diagonal entries and supports the extraction of key statistical summaries, such as system-wide trip totals and maximum flows. To bridge the gap between spatial abstraction and operational planning, the model incorporates a travel time estimation module. Base speeds are assigned according to trip length (20 km/h for <5 km, 35 km/h for <15 km, 50 km/h otherwise), and a demand-dependant congestion factor is applied:

$$\text{demand\_factor} = 1+0.3\times\frac{\log(1+\text{origin\_demand})+\log(1+\text{destination\_demand})}{10}$$

The final travel time in minutes is then:

$$\text{travel\_time} = \frac{\text{distance}}{\text{base\_speed}} \times 60 \times \text{demand\_factor}$$

with a minimum constraint of 1.2 minutes per kilometre to reflect operational realities.

By systematically addressing data normalisation, parameter calibration, and real-world constraints such as congestion, this gravity model delivers robust, transferable flow estimates that support both network optimisation and scenario analysis, regardless of the underlying data environment.

Genetic algorithm

GA is used to address the combinatorial complexity of TNDFSP. GAs are particularly well-suited for this task as their population-based search mechanism is highly effective at exploring the non-linear solution space of route and frequency combinations without requiring restrictive simplifications. This allows for the simultaneous optimisation of both route structures and service frequencies. A GA's ability to maintain a diverse population of candidate networks makes it less prone to premature convergence in local optima, ensuring a more robust search for globally competitive designs.

Within the GA, each chromosome encodes a complete transit network as a dual-layer structure: the first layer represents the spatial configuration of routes, where each route is an ordered sequence of station IDs, and the second layer specifies the service frequency for each line. For a network with *m* routes, the chromosome *C* is defined as $C = R, F$, where $R = [r_1, r_2, \ldots, r_m]$, and each $r_i = [n_{i,1}, n_{i,2}, \ldots, n_{i,l_i}]$, and *F* is the vector of frequencies. This representation enables the GA to independently manipulate both the structural and operational characteristics of the network, while maintaining their interdependence in the fitness evaluation.

Algorithm 4 Initial Population Generation (Backbone and Local Lines)

```
Initialise empty population P
for i = 1 to POPULATION_SIZE do
    Randomly determine number of lines L within allowed range
    Allocate ~70% of stations to N backbone lines
    for each backbone line do
        Saeed with backbone stations; connect using greedy or MST-based approach
        Assign random frequency within allowed range
    end for
    for each local line do
        Saeed with local stations; connect using nearest-neighbour or random approach
        Assign random frequency within allowed range
    end for
    Add any remaining uncovered stations to backbone/local lines or create new lines as needed
    Repair connectivity within each line
    Add chromosome to population P
end for
return population P
```

FIGURE 5: Population Generation Pseudocode

Initial population generation, as shown in Figure 5, has been designed to explicitly construct both backbone and local lines, reflecting real-world transit network design principles and improving the diversity and feasibility of candidate solutions. Backbone lines are seeded as high-capacity, long routes covering major corridors and a large proportion of stations, typically using greedy or MST-inspired approaches, while local lines are shorter, feeder routes initialised using randomised or nearest-neighbour strategies to ensure all areas are served. This backbone/local distinction is enforced throughout the evolutionary process, with all stations explicitly covered in the initial population. This approach ensures that every initial chromosome is both diverse and feasible, reduces the need for excessive repair operations, and accelerates convergence. The justification for this strategy is that backbone lines provide strong network connectivity and operational efficiency, while local lines enhance coverage and flexibility, together reflecting the hierarchical structure of effective real-world transit networks.

The fitness function is multi-objective, balancing operational efficiency, network simplicity, and passenger experience. For a solution with *n* lines, fitness is defined as:

$$\text{Fitness} = \frac{1}{n}\sum_{i=1}^{n}\left(0.5\cdot\frac{R_i}{T_i} + 0.2\cdot\frac{R_i}{E_i} + 0.3\cdot\frac{R_i}{W_i} \quad -P_i^{\max} - P_i^{\text{count}} - P_i^{\min}\right) + C_s + C_b - P_t - P_o - P_m$$

where R_i is the ridership, T_i is the travel time, E_i is the number of edges, W_i is the average waiting time, $P_i^{\max}$, $P_i^{\min}$, and P_i^{count} are station count penalties, C_s and C_b are coverage metrics, and P_t, P_o, and P_m represent penalties for transfers, overlap, and missing stations. This formulation ensures that solutions are not only efficient but also practical and passenger-friendly, with explicit penalties and bonuses guiding the search toward implementable networks.

Constraint handling is integral to the algorithm. All solutions are required to satisfy minimum and maximum bounds on the number of lines, stations per line, and service frequencies. Each line must form a connected subgraph, enforced by a repair function using a Union-Find structure and edge betweenness centrality. Isolated or missing stations are progressively connected, with increasing penalties to drive the population toward full coverage.

The multi-objective nature of this fitness function is designed to navigate the inherent trade-offs in transit network design, which are managed through the implicit weighting of each penalty and bonus term. A primary conflict exists between passenger convenience and operational cost. The optimisation seeks to improve the passenger experience by maximising Ridership and Coverage Bonus while minimising Travel Time, Average Waiting Time, and the Transfer Penalty. These goals, however, are in direct opposition to minimising the number of Edges and the Overlap Penalty, which represent the operational cost and physical complexity of the network. A highly convenient network often requires more routes and higher frequencies, thus increasing its cost. Similarly, the algorithm must balance network coverage against network efficiency. A high Coverage Bonus and a low Missing Station Penalty push the algorithm towards solutions that serve every potential station, ensuring equitable access across the region. This drive for comprehensive coverage can lead to longer, less direct routes that increase overall Travel Time and reduce the network's operational speed. Finally, the function manages the tension between route directness and network simplicity. Minimising the Transfer Penalty encourages the creation of direct, one-seat journeys, but this often results in a more complex network with significant route overlap and more Edges. Conversely, prioritising network simplicity with a high Overlap Penalty can create a leaner, more legible system at the cost of forcing passengers to make more transfers.

By defining these competing objectives within a single function, the genetic algorithm is tasked with finding solutions that represent a robust compromise, with the specific balance being tunable by a planner to reflect local policy priorities, operational goals, and stakeholder priorities, making it a versatile and robust tool for next-generation transit network design.

Results And Discussion

Experimental setup

The experiments, except for the SUMO simulation, were conducted using Google Colab with T4 GPU acceleration. All algorithms were implemented in Python and executed on the same computing instance to ensure consistent hardware specifications and fair comparison across the three algorithmic approaches and two urban contexts.

For the K-Means clustering algorithm, the algorithm was configured with a maximum of 300 iterations for initial optimization and 500 iterations for final clustering, with convergence tolerances set to 1e-5 for initial phases and 1e-6 for final clustering. The number of initializations was set to 1 to maintain computational efficiency while ensuring reproducible results, with a minimum and maximum centroids constraint for each region (5-30 for Greater Kuala Lumpur, 5-60 for Greater London), and a minimum distance of 3,000 meters between low-demand centroids to prevent over-concentration of stations in sparse areas.

The gravity model parameter optimisation employed was evaluated using multiple metrics including Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Percent RMSE. Model validation was performed through comparison of predicted versus observed average trip lengths and total flows. The default parameters for α, β, and γ were determined through iterative calibration specific to each urban context and weighting scheme.

For the genetic algorithm implementation, the maximum line length was restricted to 60 stations for Greater London and 30 for Greater Kuala Lumpur. Service frequency parameters were hierarchically distributed with maximum frequency set to 20 or 15 trips per hour for backbone lines for London and Kuala Lumpur respectively, and a minimum frequency of 10 trips per hour for feeder lines for both regions.

All experiments maintained identical hardware specifications and utilized default parameters throughout the testing process to ensure reproducibility and enable fair comparison across the three algorithmic approaches when applied to both urban contexts.

K-means clustering results

Type	Coverage Score
Demand	69.10%
Population	67.20%
Point of Interest	63.60%
Real Network	76.40%

TABLE 1: Greater London Coverage Scores

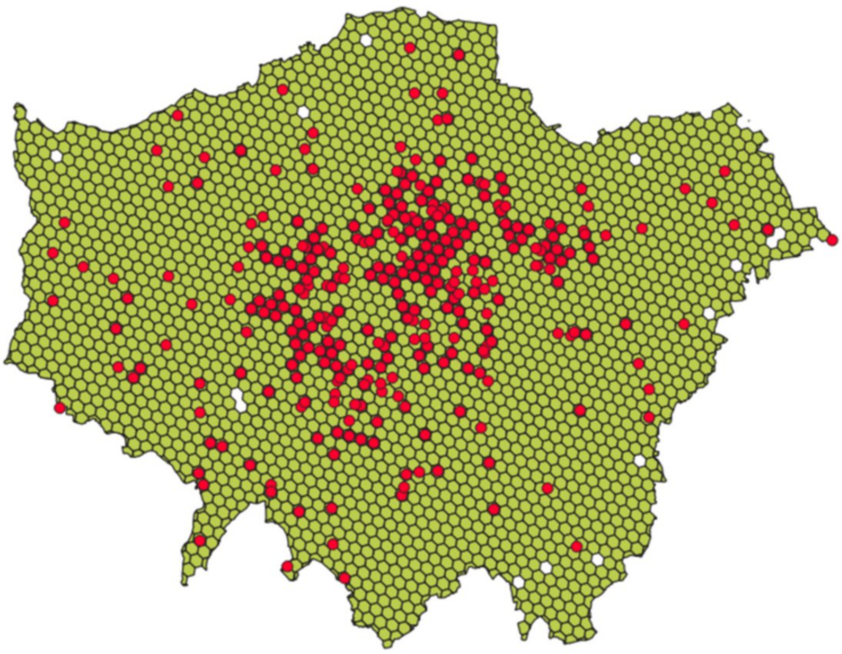

FIGURE 6: Greater London Demand-Focused Centroids

For the Greater London dataset, the demand-oriented clustering approach yielded the highest coverage score among the three algorithmic methods, achieving 69.1%. The population-based approach followed closely with a score of 67.2%, while the POI-based approach scored 63.6%. As detailed in Table *1*, all three algorithmically generated networks produced coverage scores lower than that of the existing real network, which has a score of 76.4%. The difference between the highest-scoring algorithmic approach (demand) and the real network was 7.3 percentage points. Figure 6 provides a visualisation of the station centroids generated by the demand-focused clustering algorithm. The map shows a high concentration of stations within the central urban core of Greater London, with a progressively sparser distribution of stations in the outer boroughs. The placement of centroids adheres to the predefined geographical boundaries of the study area, and a visual inspection confirms that the regionalisation and spatial distancing constraints successfully distributed stations across all administrative regions.

Type	Coverage Score
Demand	76.60%
Population	74.30%
Point of Interest	78.90%
Real Network	59.20%

TABLE 2: Greater Kuala Lumpur Coverage Scores

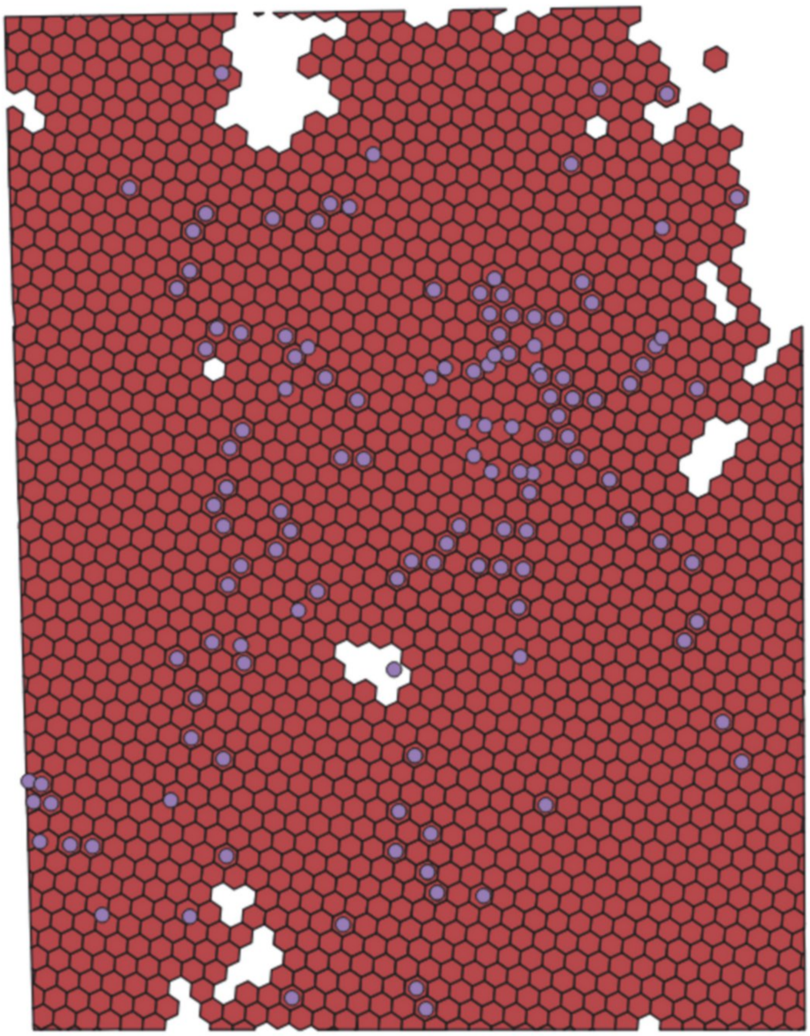

FIGURE 7: Greater Kuala Lumpur Centroids

In the Greater Kuala Lumpur content, shown in Table 2, the results showed a different pattern. All three algorithmic approaches produced networks with coverage scores that surpassed the existing real network's score of 59.2%. The POI-focused solution generated the highest-performing network with a coverage score of 78.9%, representing a 19.7-point improvement over the real network. The demand-based approach also performed strongly, achieving a score of 76.6%, while the population-focused network scored 74.3%. Figure 7 visualises the station centroids generated for Greater Kuala Lumpur. The map illustrates that the stations are distributed across the entire service area. Similar to the Greater London results, the centroids remain within the defined geographical bounding box, demonstrating the consistent application of the framework's constraints.

Gravity model results

The Gravity Model was calibrated and applied to both the Greater London and Greater Kuala Lumpur datasets, using demand, population, and POI as drivers of terminal attractiveness. Calibration was performed by iteratively optimising the α, β, γ, and k to minimise the discrepancy between observed and predicted flows. Model accuracy was evaluated using RMSE, MAE, and Percent RMSE, as well as by comparing predicted and observed average trip lengths and total flows.

Focus	Iteration	RMSE	MAE	Parameters [α, β, γ, k]
Demand	15	0.00015	0.000079	[0.3, 0.3, 0.516, 0.50]
Population	20	0.00014	0.00006	[0.3, 0.3, 0.616, 0.50]
Point of Interest	35	0.000309	0.000103	[0.3, 0.3, 0.609, 0.50]

TABLE 3: Greater London Calibration Results

MAE, Mean Absolute Error; RMSE, Root Mean Squared Error

The calibration results for Greater London are detailed in Table 3. The population-focused model achieved the lowest error metrics, producing an RMSE of 0.000140 and an MAE of 0.000060 after 20 iterations. The demand-focused model also performed with low error, resulting in an RMSE of 0.000150 and an MAE of 0.000079 in 15 iterations. The POI-focused calibration required the most iterations (35) and resulted in the highest error values, with an RMSE of 0.000309 and an MAE of 0.000103. The final optimised distance decay parameter, γ, varied across the models. The value for the population-focused model was 0.616, for the POI-focused model it was 0.609, and for the demand-focused model it was 0.516.

Focus	Iteration	RMSE	MAE	Parameters [α, β, γ, k]
Demand	25	0.000748	0.000399	[0.3, 0.3, 0.529, 0.50]
Population	25	0.000608	0.000281	[0.3, 0.3, 0.647, 0.50]
Point of Interest	20	0.000211	0.000136	[0.3, 0.3, 0.508, 0.50]

TABLE 4: Greater Kuala Lumpur Calibration Results

MAE, Mean Absolute Error; RMSE, Root Mean Squared Error

For the Greater Kuala Lumpur dataset, the calibration produced a different hierarchy of model performance, as shown in Table 4. The POI-focused model yielded the lowest error metrics, achieving an RMSE of 0.000211 and an MAE of 0.000136 in 20 iterations. The population-based model converged in 25 iterations with an RMSE of 0.000608 and an MAE of 0.000281. The demand-focused model produced the highest error values, with an RMSE of 0.000748 and an MAE of 0.000399, also after 25 iterations. The optimised distance decay parameter, γ, also showed variation. The value for the population-focused model was the highest at 0.647. The demand-focused model produced a γ value of 0.529, and the POI-focused model had the lowest γ value at 0.508.

Genetic algorithm results

Focus	Generation	Best Fitness	Stations Used
Demand	33	21,573.96	326
Point of Interest	37	22,728.82	303
Population	37	12,323.06	375

TABLE 5: Greater London Routes

Focus	Station Counts per Line	Total Non-Unique Stations
Demand	60, 60, 30, 30, 30, 30, 30, 30, 30, 2, 2, 2, 4, 2, 2	374
Point of Interest	60, 60, 30, 30, 30, 30, 11, 8, 8, 8, 8, 8, 8, 28, 8	395
Population	60, 60, 30, 30, 30, 30, 30, 30, 30, 12, 10, 14, 10, 27, 7	436

TABLE 6: Greater London Station Counts per Line and Overlap

The optimisation for the Greater London networks converged between 33 and 37 generations, as shown in Table 5. The population-focused solution achieved the lowest fitness score (12,323.06), while utilising all 375 available stations. In contrast, the demand-focused solution converged in 33 generations with a fitness score of 21,573.96 using 326 stations, and the POI-focused solution produced a fitness score of 22,728.82 using 303 stations after 37 generations. The structural details of the generated lines are presented in Table 6. A consistent structural pattern was observed across all three optimisation approaches: each solution generated two long "backbone" lines containing the maximum permitted 60 stations. These were complemented by a series of shorter "feeder" lines of varying lengths. For instance, the demand-focused solution included seven lines of 30 stations and six smaller lines with 2-4 stations each. A count of the total non-unique stations, a measure of line overlap, showed that the population-focused network was the most redundant with 436 station instances across all lines. The POI-focused network had 395 non-unique stations, and the demand-focused network had 374.

Focus	Generation	Best Fitness	Stations Used
Demand	30	8,178.19	160
Point of Interest	31	8,942.92	165
Population	21	7,001.46	188

TABLE 7: Greater Kuala Lumpur Routes

Focus	Station Counts per Line	Total Non-Unique Stations
Demand	20, 20, 20, 30, 30, 19, 9, 9, 18	175
Point of Interest	19, 20, 20, 30, 30, 28, 30, 30, 30	237
Population	20, 20, 20, 30, 30, 30, 29, 22, 6	207

TABLE 8: Greater Kuala Lumpur Station Counts per Line and Overlap

For the Greater Kuala Lumpur dataset, the GA converged more rapidly, with the best solutions appearing between 21 and 31 generations (Table 7). Similar to the London results, the population-focused approach yielded the lowest fitness score (7,001.46), achieved in 21 generations while using 188 stations. The demand-focused solution produced a fitness score of 8,178.19 (160 stations), and the POI-focused solution had a fitness of 8,942.92 (165 stations). Table 8 details the line structures for the Kuala Lumpur networks. The station counts per line were more evenly distributed compared to the London results, with most lines containing between 20 and 30 stations. For example, the POI-focused solution generated five lines with 30 stations each. The total non-unique station count was highest for the POI-focused network (237), followed by the population-focused network (207) and the demand-focused network (175). For all generated networks in both cities, the algorithm successfully met the constraint of including every candidate station, resulting in zero penalties for missing stations.

The GA simultaneously optimised service frequencies for each line, a critical advancement over approaches that treat route and frequency design separately. Frequencies (trips/hour) were bounded between 10 and 15, allowing the GA to balance service quality and operational cost.

Focus	Service Frequencies (trips/hour)
Demand	12, 12, 11, 10, 11, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10
Point of Interest	15, 15, 11, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10
Population	15, 15, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 10, 12, 10

TABLE 9: Greater London Service Frequencies per Line

The optimised service frequencies for the Greater London networks are presented in Table 9. For the demand-focused solution, frequencies were set to 12 trips/hour for two lines and 11 trips/hour for one line, with the remaining twelve lines operating at the minimum frequency of 10 trips/hour. In the POI-focused and population-focused solutions, the two backbone lines were both assigned the maximum frequency of 15 trips/hour. The remaining lines in the POI-focused solution all operated at 10 trips/hour. In the population-focused solution, one additional line operated at 12 trips/hour, while the other twelve lines were set to 10 trips/hour.

Focus	Service Frequencies (trips/hour)
Demand	12, 10, 10, 12, 12, 10, 10, 10, 10
Point of Interest	10, 10, 10, 12, 12, 10, 12, 10, 10
Population	10, 10, 10, 12, 12, 12, 10, 10, 10

TABLE 10: Greater Kuala Lumpur Service Frequencies per Line

Table 10 details the service frequencies for the nine-line networks designed for Greater Kuala Lumpur. The demand-focused solution assigned a frequency of 12 trips/hour to three of its lines, while the other six operated at 10 trips/hour. The POI-focused solution also assigned a frequency of 12 trips/hour to three lines, with the remaining six set to 10 trips/hour. The population-focused solution produced a similar pattern, with three lines operating at 12 trips/hour and six lines at 10 trips/hour. Across all solutions for Kuala Lumpur, no lines were assigned the maximum possible frequency.

SUMO simulation

The final evaluation of the optimised transit networks was conducted using SUMO, an open-source, microscopic traffic simulation platform. This step translated the abstract network designs into operational transit systems, enabling quantitative assessment under realistic demand scenarios. Optimised network designs from the genetic algorithm, including line structures, stop locations, and service frequencies, were exported and processed into SUMO-compatible XML files via Python scripts. OD matrices generated from the gravity model and clustering outputs defined realistic passenger trip distributions over a 12-hour simulation, with time-of-day weighting to capture peak and off-peak demand. The simulation ran locally on a Windows 10 machine, with 16 GB RAM and an 8-core CPU at 2GHz.

Key simulation setup features consist of:

1) Vehicles were scheduled at optimised frequencies, covering all stations in both directions.

2) Passenger trips were distributed according to a diurnal profile, with increased weighting during peak hours. Passenger numbers were taken directly from real life data, but can be replaced with an estimate by urban developers for data-scarce regions.

3) Passenger departure times were randomised within each hour to ensure realistic flow.

4) Transfer stations were explicitly linked for pedestrian transfers and associated delays.

5) Simulations ran for 12 hours, logging vehicle and passenger states, trip durations, waiting times, and network utilisation.

6) The pipeline is robust to missing or ambiguous data, automatically generating placeholder coordinates and default travel times as needed. SUMO's detailed event logging enabled extraction of comprehensive performance metrics for analysis.

Metric	Population	Demand	Point of Interest
Average Trip Length (m)	17,166.08	12,259.98	25,189.14
Average Trip Duration (s)	3,695.69	2,293.39	5,641.11
Average Trip Speed (m/s)	4.64	5.35	4.47
Average Waiting Time (s)	828.76	5.07	1,218.28
Total Network Length (km)	790	658	640
Number of Stations	375	326	303

TABLE 11: Greater London Networks Simuation Statistics

The simulation statistics for the Greater London networks are presented in Table *11*. The demand-focused network, which utilised 326 stations, produced the most efficient performance metrics. It had the lowest average trip duration at 2,293.39 seconds and the lowest average waiting time at just 5.07 seconds. It also achieved the highest average speed of 5.35 m/s. The population-focused network, which was the largest with 375 stations and a total length of 790.00 km, produced intermediate results. Its average trip duration was 3,695.69 seconds, with an average waiting time of 828.76 seconds and an average speed of 4.64 m/s. The POI-focused network, comprising 303 stations, recorded the longest average trip length (25,189.14 m) and the longest average trip duration (5,641.11 seconds). It also had the highest average waiting time at 1,218.28 seconds and the lowest average speed at 4.47 m/s.

Metric	Population	Demand	Point of Interest
Average Trip Length (m)	13,647.77	14,783.69	24,344.00
Average Trip Duration (s)	4,773.38	5,588.23	10,197.28
Average Trip Speed (m/s)	2.86	2.65	2.39
Average Waiting Time (s)	836.06	255.72	1,457.28
Total Network Length (km)	396	332	456
Number of Stations	188	160	165

TABLE 12: Greater Kuala Lumpur Networks Simulation Statistics

The simulation statistics for the Greater Kuala Lumpur networks, detailed in Table *12*, show different performance characteristics. The population-focused network, which used the most stations (188), achieved the shortest average trip duration at 4,773.38 seconds and the highest average speed at 2.86 m/s. Its average waiting time was 836.06 seconds. The demand-focused network produced the lowest average waiting time at 255.72 seconds. However, its average trip duration was higher at 5,588.23 seconds, and its average speed was 2.65 m/s. The POI-focused network recorded the longest average trip length (24,344.00

m), the longest average trip duration (10,197.28 seconds), and the highest average waiting time (1,457.28 seconds). It also had the lowest average speed at 2.39 m/s.

Discussion

The results of this study, generated through the sequential application of clustering, demand modelling, and genetic algorithm optimisation, highlight the adaptability of the MANTIS framework and reveal critical differences in transit planning requirements between mature and rapidly developing urban environments.

The clustering stage provided the first clear indication of the contrasting urban environments of Greater London and Greater Kuala Lumpur. In London, all algorithmically generated networks underperformed the existing real network in terms of coverage score (Table *1*) suggesting that its transit system is highly mature. The 7.3-percentage-point gap between the best algorithmic score and the real network likely represents decades of planning decisions influenced by political, financial, and historical factors that a purely data-driven model cannot capture. Furthermore, the small difference between the demand-focused (69.1%) and population-focused (67.2%) scores points to a well-established, likely monocentric urban form where population density and travel demand patterns are strongly correlated.

However, every MANTIS-generated network for Greater Kuala Lumpur significantly outperformed the existing real network's coverage score (Table *2*), with the POI-focused network showing a remarkable 19.7-point improvement. This large performance gap suggests that significant optimisation opportunities exist in rapidly developing cities where the transit network may not have kept pace with urban growth. The superiority of the POI-based approach in Kuala Lumpur, when it was the worst performing network London, strongly reflects its polycentric nature. In such a city, commercial hubs and distributed activity centres act as more powerful drivers of transit demand than simple population density.

The gravity model calibration results reinforce these findings. For London, the model calibrated most effectively (lowest RMSE/MAE) using population and demand data (Table *3*), which aligns with the idea of a mature city with predictable, population-driven commuting patterns. The higher distance decay (γ) values for population (0.616) and POI (0.609) models compared to the demand model (0.516) indicate that people are less willing to travel long distances for activities not related to their primary, established travel routes. For Kuala Lumpur, the gravity model calibrated best using POI data (Table *4*), confirming that activity centres are the strongest attractors. The extremely high γ value for the population model (0.647) suggests that residential travel patterns are highly localised, perhaps indicating that the current transit system is less effective at connecting disparate residential zones.

Furthermore, the GA consistently identified an efficient "backbone-and-feeder" structure in both cities, a key validation of MANTIS' ability to replicate sound, real-world network design principles. However, the specific optimisation trade-offs differed significantly between the two urban contexts.

In London, the algorithm generated two very long backbone lines at the maximum permitted length of 60 stations, supported by shorter feeder routes (Table *6*). This structure is ideal for a large, monocentric city, efficiently funnelling passengers along major corridors. The population-focused solution achieved the lowest fitness score (12,323.06) by creating the densest, most interconnected network that utilised all 375 candidate stations (Table *5*). This came at the cost of high redundancy (436 non-unique stations), illustrating a trade-off between maximising coverage and maximising operational simplicity. The hierarchical frequency allocation, with backbone lines receiving the maximum 15 trips/hour (Table *9*), further reinforces a design strategy focused on serving primary corridors.

In Kuala Lumpur, the faster convergence of the GA (21-31 generations vs. 33-37 in London) suggests a less complex optimisation landscape, possibly due to a more dispersed urban structure offering clearer optimal solutions. The resulting networks featured more balanced line lengths, with most lines containing 20-30 stations (Table *8*), and a more uniform frequency distribution (Table *10*). This indicates a network designed to serve a polycentric region, where multiple routes require a similarly high level of service to connect distributed activity centres, rather than concentrating service on only one or two primary corridors.

The final validation of these different design strategies was provided by the SUMO simulation, which translated the abstract network properties into tangible operational outcomes. These quantitative results allow for a deeper comparative analysis of the trade-offs inherent in each planning approach.

For Greater London, the simulation results demonstrate the superiority of a demand-focused strategy in a mature, monocentric city. The demand-focused network's average passenger waiting time of only 5.07 seconds (Table *11*) is exceptionally low, indicating that service frequencies and routes were almost perfectly aligned with the simulated travel patterns. This high level of service translated directly into the shortest average trip duration (2,293.39 seconds) and highest average speed (5.35 m/s). This outcome is critically important, as it suggests that the network's efficiency stems from providing direct, high-frequency service exactly where it is needed. In contrast, the population-focused network, despite using the most stations (375) and having the highest network density, resulted in an average waiting time of over 13 minutes (828.76 seconds). This implies a potential mismatch between coverage and service, where adding more stops without precisely matching demand can lead to underutilised segments and inefficient vehicle allocation, ultimately degrading the passenger experience. The POI-focused network performed worst on all user-centric metrics, confirming that simply connecting major destinations is an ineffective strategy if it fails to account for the origins and travel paths of the passengers who use them.

The simulation results for Greater Kuala Lumpur reveal a more complex set of trade-offs, due to city's polycentric environment. While the demand-focused network once again achieved the lowest average waiting time (255.72 seconds), indicating a highly responsive service, it did not deliver the fastest trips. Instead, the population-focused network produced the shortest average trip duration (4,773.38 seconds) and the highest average speed (2.86 m/s), despite a substantial average waiting time of nearly 14 minutes (836.06 seconds). This presents a city planners with a decision: should one prioritise minimising wait times or minimising in-vehicle travel time? The population-focused network likely provided more direct routes by ensuring broader geographical coverage, reducing the physical distance of many trips, but its service frequency was not sufficiently aligned with demand, leading to long waits. The demand-focused network provided frequent service where needed but may have resulted in more indirect routes or a higher number of transfers for passengers whose journeys did not align perfectly with the primary demand corridors. The poor performance of the POI-focused network on all metrics further underscores that in a sprawling context, connecting destinations without considering residential origins and overall network geometry can create a system that is functionally unusable, with average waits of over 24 minutes and trip durations nearly double that of the other networks.

The quantitative results from the SUMO simulation can be further contextualised by benchmarking them against the performance of the real-world rail networks in both cities. This comparison helps to validate the outputs of the MANTIS framework and grounds its theoretical performance within the scope of practical, operational reality.

For Greater London, the simulated demand-focused network shows a notable alignment with the mature, high-performance characteristics of the existing Transport for London system. The operational London Underground has 272 stations and a total network length of 402 km [25]. The largest network generated by MANTIS (population-focused) proposed a denser system with 375 stations and 790 km of track, whereas the most efficient demand-focused network was smaller, with 326 stations and 658 km of track. The average journey length on the real network is approximately 9-10 km [26], comparable to our simulated average of 12.26 km. Real-world average trip duration is reported to be around 30 minutes [27], while our simulation produced a slightly longer average of 38.2 minutes. The average operational speed of the Underground is approximately 33 km/h [28], which is higher than our simulated speed of 19.26 km/h (5.35 m/s); this difference can be attributed to the simulation's simplified assumptions about track geometry and idealised acceleration. Finally, while our simulated waiting time of 5 seconds is an idealised best case, it directionally reflects the real-world efficiency of London's high-frequency services, which aim for peak-hour headways of 3-4 minutes [29].

In the context of Greater Kuala Lumpur, the simulation results reflect the dynamics of a rapidly growing, polycentric network. The operational Klang Valley rail network consists of multiple lines, for instance, the Kajang and Putrajaya MRT lines together have a combined length of over 100 km and serve more than 60 stations [30,31]. This is smaller in scale than the simulated networks, which proposed between 160 and 188 stations, suggesting that MANTIS identifies potential for significant network expansion. Most significantly, the simulated average waiting time of 4.25 minutes for the demand-focused network is highly competitive and aligns closely with the real-world peak-hour train frequencies of 3-6 minutes on major lines [32]. The simulated average trip duration of 93.2 minutes is at the higher end of typical commute times, reflecting the model's generation of longer, more integrated routes to serve a sprawling region [33].

This multi-metric comparison between simulated and real-world data demonstrates that while the MANTIS framework necessarily simplifies some operational complexities, it produces network designs whose performance characteristics fall within a credible and realistic range across different urban contexts. Taken together, these detailed simulation findings underscore that there is no single optimal solution in transit network design. The quantitative performance metrics reveal that the effectiveness of any strategy is highly dependent on the city's unique urban form and developmental stage. The strength of the MANTIS framework lies in its ability to model these different strategies and provide planners with the evidence needed to weigh competing priorities, such as demand alignment, equitable coverage, and destination connectivity, and understand the direct impact of their choices on the passenger experience. This allows for the creation of networks that are not just theoretically optimal, but are truly adapted to the specific challenges and opportunities of their urban environment.

Conclusions

This research presents MANTIS, a comprehensive framework for automated public transit network design that generates synthetic demand patterns from population and point-of-interest data. The framework integrates regionally-aware K-means clustering for station placement, gravity model calibration for flow estimation, and genetic algorithm optimization for joint route structure and service frequency design. The primary contribution of this work is its delivery of a holistic, end-to-end design pipeline, which is made possible by its novel synthetic demand generation methodology. This methodology eliminates dependence on historical OD matrices by effectively utilising diverse data sources, enabling network design in data-scarce environments. This capability is what allows MANTIS to function as a complete solution, addressing station placement, demand modelling, and joint route and frequency optimisation in a single, integrated process, and in doing so, it addresses a critical barrier in transit planning where comprehensive travel data is often unavailable or outdated, particularly in rapidly developing cities. This work is further distinguished by its planner-centric modularity. The framework's multi-objective optimisation approach allows planners to prioritise demand patterns, population distribution, or activity centres based on local urban contexts. This provides flexibility in network design strategies and separates MANTIS from more rigid frameworks that offer less control to the end-user. Furthermore, several innovations emerged. The use of regionally-aware K-means clustering enabled more realistic and context-sensitive station placement, while the gravity model, calibrated with robust parameter optimisation, produced plausible flow estimates even in data-scarce environments. The genetic algorithm advanced the state of the art by jointly optimising both route structures and service frequencies, with a chromosome representation that maintained operational feasibility throughout the evolutionary process. Benchmarking AI-generated networks against real-world systems via SUMO simulation, rather than theoretical ideals, provided concrete evidence of the framework’s practical value and highlighted areas for further improvement. The key findings also demonstrate that AI-generated networks can outperform existing transit systems, particularly in rapidly developing urban regions. The data reveals that demand-focused approaches achieve superior performance in mature cities like London, while POI-focused strategies excel in polycentric cities like Kuala Lumpur. These context-sensitive network design insights show that optimal transit planning strategies vary significantly between urban forms, with different optimisation approaches required for monocentric versus polycentric cities.

While this research presents a robust framework for automated transit design, it is important to explicitly define the boundaries of its current applicability and to outline a clear roadmap for future enhancements. The primary limitation of the MANTIS framework in its current form is that it generates transit networks from scratch. This makes it a powerful tool for planning in new or unserved regions, which works well for urban regions without existing transit data as is intended, but it is not yet adapted for the more common and equally critical task of incrementally improving pre-existing transit infrastructure. Furthermore, the model's optimisations operate on a geospatial level and do consider physical or financial constraints for real-world deployment, such as traffic restrictions, existing land or budgetary limits on fleet size, and operational costs, meaning the resulting networks are operationally plausible but not yet budget-constrained. Future research will should aim to address these limitations and expand the framework’s practical utility. The most critical priority is to develop a new module that enables the intelligent integration of MANTIS-generated lines with existing transit systems. This would involve treating a city's current network as a fixed constraint while optimising new routes specifically to enhance connectivity, reduce redundancy, and fill service gaps. Furthermore, future work should focus on enhancing the model’s realism by incorporating real-world budgetary and physical constraints. This can be achieved by adding new penalty functions to the genetic algorithm that account for factors like fleet size limits and total kilometres, to better consider operational costs. Lastly, a longer-term goal is to explore the integration of dynamic, real-time data sources. Such an enhancement would allow the framework to evolve from a static planning tool into a dynamic operational one, providing the capability to suggest service frequency adjustments in response to real-time changes in urban demand, such as those caused by major public events or unexpected network disruptions. The MANTIS framework advances transit network design by providing a scalable, data-flexible methodology that produces operationally viable networks while accommodating diverse urban planning objectives and constraints. This research establishes a foundation for automated transit network design that can adapt to varying urban contexts and data availability conditions.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Zhiyuan Chen, Youssef Mohamed Abdelhamid Mohamed

Critical review of the manuscript for important intellectual content: Zhiyuan Chen

Supervision: Zhiyuan Chen

Acquisition, analysis, or interpretation of data: Youssef Mohamed Abdelhamid Mohamed

Drafting of the manuscript: Youssef Mohamed Abdelhamid Mohamed

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

References

1. Kepaptsoglou K, Karlaftis M: Transit route network design problem: Review. Journal of Transportation Engineering. 2009, 135:491-505. 10.1061/(asce)0733-947x(2009)135:8(491)
2. Król A, Król M: The design of a metro network using a genetic algorithm. Applied Sciences. 2019, 9:433. 10.3390/app9030433
3. Lammaihri H, Erramla M, Abdul Ghani N: Optimizing metro station locations and line layouts in selangor using genetic algorithm approach: Technical report. arXiv. 2024, 10.48550/arXiv.2411.03797
4. Feng X, Zhu X, Qian X, Jie Y, Ma F, Niu X: A new transit network design study in consideration of transfer time composition. Transportation Research Part D: Transport and Environment. 2019, 66:85-94. 10.1016/j.trd.2018.03.019
5. Zhang Y, Wang S, Ji G: A comprehensive survey on particle swarm optimization algorithm and its applications. Mathematical Problems in Engineering. 2015, 2015:931256. 10.1155/2015/931256
6. Supangat K, Soelistio YE: Bus stops location and bus route planning using mean shift clustering and ant colony in west Jakarta. arXiv. 2016, 10.48550/arXiv.1611.05579
7. Nikolić M, Teodorović D: Transit network design by bee colony optimization. Expert Systems with Applications. 2013, 40:5945-5955. 10.1016/j.eswa.2013.05.002
8. Kuan SN, Ong HL, Ng KM: Applying metaheuristics to feeder bus network design problem. Asia-Pacific Journal of Operational Research. 2004, 21:543-560. 10.1142/S0217595904000382
9. Park SJ, Kang S, Byon YJ, Kho SY: Multiobjective approach to the transit network design problem with variable demand considering transit equity. Journal of Advanced Transportation. 2022, 2022:1-17. 10.1155/2022/5887985
10. Owais M, Ahmed AS, Moussa GS, Khalil AA: Integrating underground line design with existing public transportation systems to increase transit network connectivity: Case study in Greater Cairo. Expert Systems with Applications. 2021, 167:114183. 10.1016/j.eswa.2020.114183
11. Abrahamsson T: Estimation of origin-destination matrices using traffic counts - a literature survey. IIASA Interim Report. IIASA, Laxenburg, Austria: IR-98-021. 1998, 1-32.
12. Barceló J, Montero L, Bullejos M, Serch O, Carmona C: Dynamic OD matrix estimation exploiting Bluetooth data in urban networks. 14th International Conference on Automatic Control, Modeling and Simulation, Saint Malo (France). 2012, 116-121.
13. Zhang J, Che H, Chen F, Ma W, He Z: Shortterm origin-destination demand prediction in urban rail transit systems: A channel-wise attentive split-convolutional neural network method. Transportation Research Part C: Emerging Technologies. 2021, 124:102928. 10.1016/j.trc.2020.102928
14. Dragu V, Roman EA: The origin-destination matrix development. MATEC Web of Conferences. 2019, 290:06010. 10.1051/matecconf/201929006010
15. Chen AHL, Liang YC, Chang WJ, Siauw HY, Minanda V: RFM model and K-means clustering analysis of transit traveller profiles: A case study. Journal of Advanced Transportation. 2022, 2022:1-14. 10.1155/2022/1108105
16. Chen S, Zhang F, Zhang Z, Yu S, Qiu A, Liu S, Zhao X: Multi-scale massive points fast clustering based on hierarchical density spanning tree. ISPRS International Journal of Geo-Information. 2023, 12:24. 10.3390/ijgi12010024
17. OpenStreetMap Foundation. (2025). Accessed: August 27, 2025: https://www.openstreetmap.org.
18. Kontur Population: Global population density for 400m H3 hexagons. (2023). Accessed: August 27, 2025: https://data.humdata.org/dataset/kontur-population-dataset.
19. Origin and destination matrix (ODM) 2021-22 (Version 1.0). (2021-2022). Accessed: August 27, 2025: https://raildata.org.uk/dataProduct/P-a9faf6fd-b31f-491c-935a-f4ab3ffa5890/overview.
20. Point of interest data for the United Kingdom. (2024). Accessed: August 27, 2025: https://data.geods.ac.uk/dataset/point-of-interest-data-for-the-united-kingdom.
21. Transport for London Unified API: Journey Planner. (2025). Accessed: August 27, 2024: https://api-portal.tfl.gov.uk/api-details#api=Journey&operation=Journey_JourneyResultsByPathFromPathToQueryViaQueryNa....
22. Rapid rail explorer. (2024). Accessed: August 27, 2024: https://data.gov.my/dashboard/rapid-explorer.
23. Malaysia points of interest (OpenStreetMap Export). (2025). Accessed: August 27, 2025: https://data.humdata.org/dataset/hotosm_mys_points_of_interest.
24. Moovit Public Transit Journey Planner. (2025). Accessed: August 27, 2025: https://moovitapp.com/.
25. Transport-for-London: What we do. Accessed: August 27, 2025: https://tfl.gov.uk/corporate/about-tfl/what-we-do.
26. London Assembly: Average distance travelled by person per mode in London . Accessed: August 27, 2025: https://www.london.gov.uk/who-we-are/what-london-assembly-does/questions-mayor/find-an-answer/average-distance-travel....

27. London Assembly: Tube journey times . Accessed: August 27, 2025: https://www.london.gov.uk/who-we-are/what-london-assembly-does/questions-mayor/find-an-answer/tube-journey-times.
28. Transport for London: FOI request detail. Average train speed. Accessed: August 27, 2025: https://tfl.gov.uk/corporate/transparency/freedom-of-information/foi-request-detail?referenceId=FOI-0228-1819.
29. Transport for London. (2024). Customer service and operational performance report. Quarter 4 2023/24 (10 December 2023-31 March 2024). Accessed: August 27, 2025: https://content.tfl.gov.uk/customer-service-and-operational-performance-report-quarter-4-2023-24-acc.pdf.
30. MRT Corp. Projects: Kajang Line. Accessed: August 27, 2025: https://www.mymrt.com.my/en/kajang-line/.
31. MRT Corp. Projects: Putrajaya Line. Accessed: August 27, 2025: https://www.mymrt.com.my/projects/putrajaya-line/.
32. MRT Corp. Train frequency. Accessed: August 27, 2025: https://mrt.com.my/g-around/Train_Frequency.htm.
33. The Edge Malaysia: MRT Corp's accumulated losses hit RM57.6b while MRT1, MRT2 fail to meet ridership, frequency targets — AG's report. Accessed: August 27, 2025: https://theedgemalaysia.com/node/717932.