THEMA
théorie économique,
modélisation et applications

THEMA Working Paper n°2026-07
CY Cergy Paris Université, France

Urban Emissions Modeling using the Metropolis Multi-Agent Framework: The Case of Riyadh City

F. Belaid, L. Yaseen, A. De Palma, M. Kilani

March 2026

CERGY PARIS UNIVERSITÉ
cnrs

Urban Emissions Modeling using the Metropolis Multi-Agent Framework: The Case of Riyadh City

F. Belaid, L. Yaseen, A. De Palma, M. Kilani

March 2026

Abstract

Urban mobility-related emissions are a growing concern in rapidly expanding cities, driving the need for robust methods to assess and mitigate their environmental impact. Here, we focus on developing a multi-agent model to estimate emissions in Riyadh city as the primary case study. Specifically, this study employs the dynamic traffic simulator, METROPOLIS2, to examine mobility-related emissions and their environmental impact. The framework integrates a cleaned road network, an origin–destination matrix for 162 zones with node–zone assignment, and a simplified metro layer. The model is calibrated to reproduce observed congestion patterns and verified through day-to-day dynamics—convergence in departure/arrival times, route reallocation under congestion, and distance–emissions co-movement under a homogeneous fleet. We report generalized cost and emissions indicators at the trip and network levels. The contribution is methodological: a transparent, reproducible baseline for Riyadh that enables credible scenario evaluation.

Keywords: Air pollution; Emission reduction strategies; Multi-agent simulation; Environmental impact; Policy interventions; Smart mobility solutions.

JEL: C63, R41, R48, D61

Public

1. Introduction

Saudi Arabia is undergoing a transformative phase of rapid urbanization, propelled mainly by its ambitious Vision 2030 initiative. This program, launched to diversify the economy beyond oil, includes numerous large-scale urban development projects that reflect the country's commitment to modern, sustainable cities. Riyadh stands out as the fastest-growing, with extensive urban interventions, road infrastructure expansion, and the introduction of Saudi Arabia's first mass transit metro system to support its evolving mobility landscape (Royal Commission for Riyadh City, 2025). Nevertheless, Riyadh's rapid growth and increasing motorization rates have resulted in pressing environmental and mobility challenges, notably severe congestion across critical corridors during peak periods (Yaseen et al., 2024). Given this context, Riyadh provides an ideal case study for assessing the impacts of urban transport policies on emissions and mobility outcomes.

Despite policy efforts, the absence of comprehensive, data-driven simulation tools has constrained the prediction of short-, medium-, and long-term environmental and operational impacts of evolving urban mobility strategies. This study addresses this gap by employing a multi-agent transport simulation (MATS) model to evaluate how targeted interventions—such as electric vehicle (EV) adoption, teleworking, and the introduction of low-emission zones—can shape Riyadh's emissions trajectory and contribute to more sustainable urban mobility. MATS is an innovative yet now widely adopted approach that leverages the principles of agent-based modeling to analyze and predict urban transport dynamics. In this framework, individual agents—representing vehicles, pedestrians, cyclists, and public transport users—interact within a simulated environment that mirrors real-world conditions. Agents follow specified behavioral rules, including route selection, speed regulation, and traffic signal compliance, facilitating detailed insights into modal interactions and their implications for urban traffic conditions. This approach also makes it possible to track each vehicle fleet, characterized by its type. Importantly, MATS also evaluates policies targeting environmental sustainability by assessing emission reductions and energy-efficiency improvements**.** By analyzing multiple policy scenarios, the model can support urban planners in devising effective strategies to mitigate environmental impacts.

A key application of MATS is scenario-based assessment of policy interventions targeting transport efficiency, emission reductions, and improved air quality. Urban planners can simulate scenarios involving public transport enhancements, congestion pricing, or adaptive traffic management. Adjusting simulation parameters allows stakeholders to quantify policy impacts not only on traffic flow and travel times but also emissions and energy use, core elements of sustainable urban development. Further, MATS can examine distributional implications of transport policies, identifying potential equity issues and promoting inclusive urban strategies. Visualization capabilities of the simulations enhance informed decision-making by providing dynamic representations of proposed interventions. Additionally, MATS can be instrumental in understanding the interactions between different transport modes and their collective impact on urban mobility and pollution levels. For example, the simulation can reveal how increased cycling infrastructure might affect car traffic or how pedestrian-friendly initiatives can enhance public transport usage. The model explicitly evaluates emissions-related policies such as electric vehicle adoption, low-emission zones,

and telecommuting incentives aimed at congestion and air-quality improvements. Capturing these interdependencies supports integrated solutions that simultaneously target congestion reduction, emissions mitigation, and improved accessibility.

Developing a transportation simulation model for Riyadh is useful for gaining deeper insights into the city's complex traffic dynamics. By replicating real-world conditions, the model allows researchers and policymakers to analyze traffic flows, identify bottlenecks, and evaluate strategies to reduce congestion and emissions. A Multi-agent simulator as METROPOLIS2 is particularly valuable as it captures interactions between diverse agents—such as drivers, public transport users, Critically, it quantifies how electrification, fuel-price reform, and fleet-renewal incentives alter vehicle kilometers travelled (VKT) and urban emissions.

The METROPOLIS2 dynamic traffic simulator offers a versatile tool capable of evaluating a wide range of urban and mobility policies. However, it focuses on short-term analyses, assuming fixed residential locations, job distributions, and land prices. Long-run feedbacks, therefore, require coupling with land-use-transport-interaction (LUTI) models. It is a mesoscopic model, which occupies an intermediate position between microscopic and macroscopic simulations. Building on agent-based methodologies, it captures heterogeneity on both the supply and demand sides while simplifying the modeling of congestion (e.g., omitting lane-changing or car-following behaviors). This simplification enhances scalability and makes them suitable for large-scale scenarios.

The first version of the model was presented in de Palma et al. (1997) and de Palma and Marchal (2002), with early applications to simplified networks (de Palma et al., 2005) and subsequent deployments across major metropolitan regions. It has been applied to Paris and Île-de-France (de Palma and Sanchez, 1998), Lausanne, Brussels, Québec, and Stockholm (Saifuzzaman et al., 2016), with ongoing work in Québec and Algiers. More recent studies extend its use to carpooling (de Palma et al., 2024) and to local and global pollution assessments in Île-de-France and Réunion Island (Le Frioux, 2024). Across these applications, METROPOLIS2 has been used to examine a wide spectrum of transport policies, including capacity expansion, traffic restrictions, incentives for public transport, pricing schemes, staggered schedules, low-emission zones, and carpooling, analyzing impacts on travel costs, user surplus, and pollution at multiple scales. While the model has been applied in other international contexts, this paper presents its first implementation for Riyadh, documenting the methodological adaptation and calibration required for a city facing rapid urbanization, severe congestion, and pressing air-quality challenges.

Given Riyadh's focus, this model can investigate city-relevant policies such as teleworking impacts, demographic and socioeconomic shifts influencing transport demand, and electric vehicle adoption scenarios. The model can facilitate detailed examinations of environmental benefits derived from urban policy interventions. Land-use planning strategies to reduce emissions and optimize urban energy use can similarly be explored through modified origin–destination matrices reflecting different spatial arrangements. Notably, the multi-agent approach employed in this study allows us to draw a complete emissions and energy consumption profile across all scenarios, providing key variables to feed a LUTI model. LUTI models can be used to better understand how these policies shape both urban operation efficiency and environmental outcomes, while describing their impacts on residential

Public

relocations and jobs. The quality of the results depends on the input data, particularly the origin-destination matrix, which provides critical insights into travel patterns. The fleet's composition—such as the types, sizes, and energy efficiencies of vehicles—affects emissions, energy consumption, and congestion levels, making it a key factor in the model's accuracy and relevance. Moreover, mobility surveys remain a cornerstone of any transport and urban analysis. It will provide necessary information to assess car ownership, mode choice, departure time choices, and more generally, activity patterns needed to build a synthetic population.

Moreover, the simulation incorporates Riyadh's newly operational metro line as a public transport component, recognizing it's its strategic importance for emission mitigation and urban mobility planning. Given data constraints, this evaluation currently employs a simplified approach. Understanding metro system impacts on mode choices, congestion, and emissions remains critical for guiding future public transport investments and supporting cleaner mobility strategies. As Riyadh's first urban emissions simulation framework, the model primarily aims to illustrate its broad applicability and potential.

This methodology paper introduces a comprehensive multi-agent simulation framework utilizing the METROPOLIS2 dynamic simulator, demonstrating how detailed input data—including traffic flows, infrastructure characteristics, population distribution, and behavioral parameters—can capture the complex interplay of urban mobility and emissions in Riyadh. Through the explicit integration of policy interventions such as EV adoption, teleworking, and modified travel matrices, the proposed model provides a robust analytical platform for assessing emission reduction strategies.

The rest of this report is structured as follows: Section 2 outlines the modeling framework. Section 3 presents the Riyadh context and input data, including the network configuration and the origin–destination matrix. Section 4 details the calibration procedure. Section 5 discusses the measurement of consumer surplus, and the integration of fuel consumption and pollution into the simulation framework. Section 6 concludes the paper.

Public

2. Methodology and simulator description

This section provides an overview of the METROPOLIS2, the theory underlying its different components, and the scope of its applications. The METROPOLIS2[1] simulator models movements of individual transport users and public transit in urban or peri-urban areas. This multi-agent explicitly captures individual user choices, with each traveler represented as an agent**.** It employs a dynamic approach, explicitly accounting for congestion effects continuously throughout the day**,** as a continuous variable represents time. By tracking vehicle trajectories at an individual level**,** the model facilitates detailed analyses of travel-related well-being, both for individual travelers and aggregated user groups. In this study, we specifically focus on weekday morning commuting patterns.

2.1. METROPOIS2 modules: overview

This subsection introduces the general simulation engine—its demand and supply components and the day-to-day iterative solution—and situates it within a five-module framework **(**Traffic/engine → Emissions → Dispersion → Exposure → Monetization**)** (Figure 2). The Traffic engine yields agent trajectories, link flows, and time-resolved speeds that feed the downstream modules. In this methodology paper we focus on the Traffic METROPOLIS2 emissions modules, while briefly outlining the other modules to situate the analysis within the broader framework.

The framework combines traffic outputs with environmental and health assessment tools to trace the causal chain from vehicle movements to pollutant emissions, population exposure, and associated costs. he integrated modeling chain, METRO-TRACE (Traffic-Related Air Pollution Cost Evaluation), was developed in accordance with the recommendations of Friedrich and Quinet (2011).

[1] The core of the METROPOLIS2 simulator is developed in Rust, while the input and output (I/O) modules are written in Python. Users have the flexibility to modify these external I/O modules as needed. The workflow of the model can be described through the diagram in Figure 1. In the first phase, input data is read, cleaned, and transformed into Parquet files, which are compatible with the METROPOLIS2 simulator. The simulator then processes this data through multiple iterations to generate output data. The stopping rule can either be exogenous (a predetermined number of iterations) or endogenous (based on some convergence criterion). Finally, the output data is transformed and sent to an output module that analyzes it.

Public

Figure 2: METRO-TRACE – Integrated modeling framework for traffic-related air pollution and cost evaluation

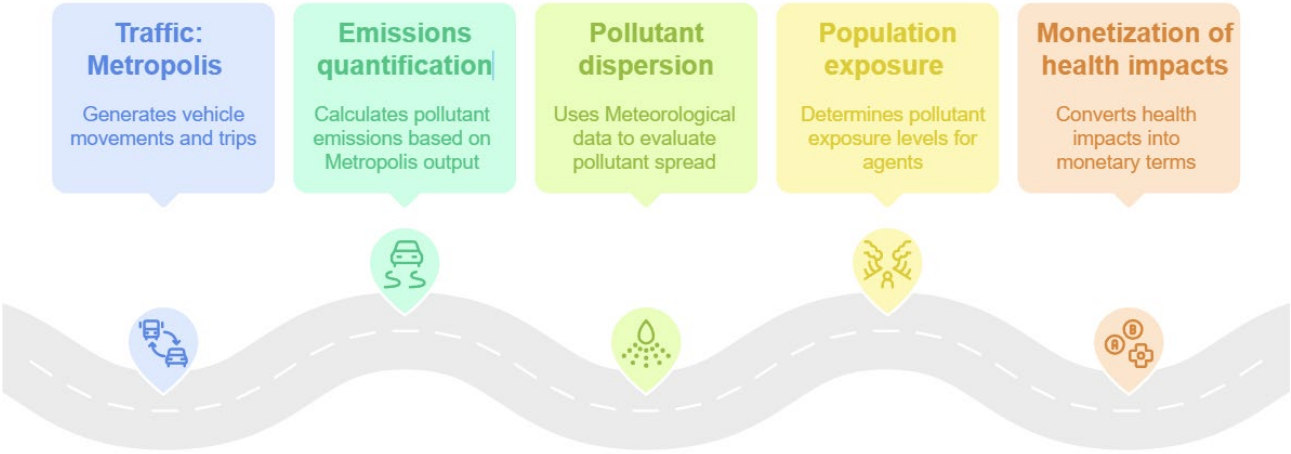

Source: Authors' design

The modules of the chain are described below:

- **Traffic Model:** This model generates agent trajectories, link flows, and time-resolved speeds for the morning peak using a nested choice structure (mode → departure time; route is deterministic with revisions under congestion) and a day-to-day learning process that converges to a stationary regime. These outputs are the primary inputs to emissions modeling.
- **Emission Module:** Maps simulated speeds, distances, and (assumed) fleet characteristics to per-link, time-resolved emissions. We report both **local pollutants** (e.g., NO_x, CO, $PM_{2.5}$) and greenhouse gases (**CO_2**)
- **Pollutant Dispersion Module:** Focused on local pollutants, this model uses Gaussian dispersion techniques for efficiency, primarily since it operates at the city scale. Parameters include variables such as temperature and wind speed.
- **Exposure Model:** This model determines pollutant exposure levels for different agents at any given moment, based on their locations. Since the transport model provides the precise location of each agent throughout the day (e.g., at home, at work, or traveling), exposure calculations are highly accurate, avoiding the assumption that individuals remain at their residence.
- **Monetization of Health Impacts:** Converts exposure changes into health damages using concentration–response functions and valuation parameters to produce external cost indicators.

2.2 Traffic module (METROPOLIS2 engine)

The METROPOLIS2 simulator models individual travelers' choices of mode, departure time, and route within a road network, all determined endogenously. Building on Vickrey's (1969, 1963) framework, road travelers are assumed to have individual-specific preferred arrival time windows at their destination and face schedule delay costs for arriving either earlier or later. The generalized systematic cost of public transport is assumed to be exogenous and independent of time of day. The generalized systematic cost of an auto trip is additively separable —that is, it can be written as the sum of travel time costs, schedule delay costs,

and flat taxes, and it is given by a piece-linear function. The demand model employs a two-stage nested logit model (Anderson et al., 1992), with mode choice represented in the outer nest (see Figure 1). At the second stage, users select the dynamic shortest paths (once the mode choice and departure time choice have been made).

The departure time choice is modeled within the inner nest, by a standard continuous logit specification. In the simulations presented in this report, the allowable departure time window spans 6:00 am to 12:00 am, as depicted. Route choice follows a heuristic derived from a generalization of Wardrop's minimum-cost principle, which closely aligns with the minimum generalized cost (de Palma et al., 2005).

On the supply side, congestion is essentially described as a series of journeys where speed depends on the level of road occupancy. Additionally, traffic slows down at specific points in the network, known as bottlenecks. The time required to clear a bottleneck depends on the ratio between the number of vehicles in the bottleneck and its capacity. A day-to-day adjustment process, incorporating exponential learning by drivers, governs changes in mode choice, departure time, and route choice, steering the system toward a stationary state. Further details on the learning process can be found in de Palma et al. (1997). Simulation experiments across various networks demonstrate that different initial conditions lead, after converging, to the same stationary regime.

Public

Figure 1: Workflow of the simulation process in METROPOLIS2 for the current project.

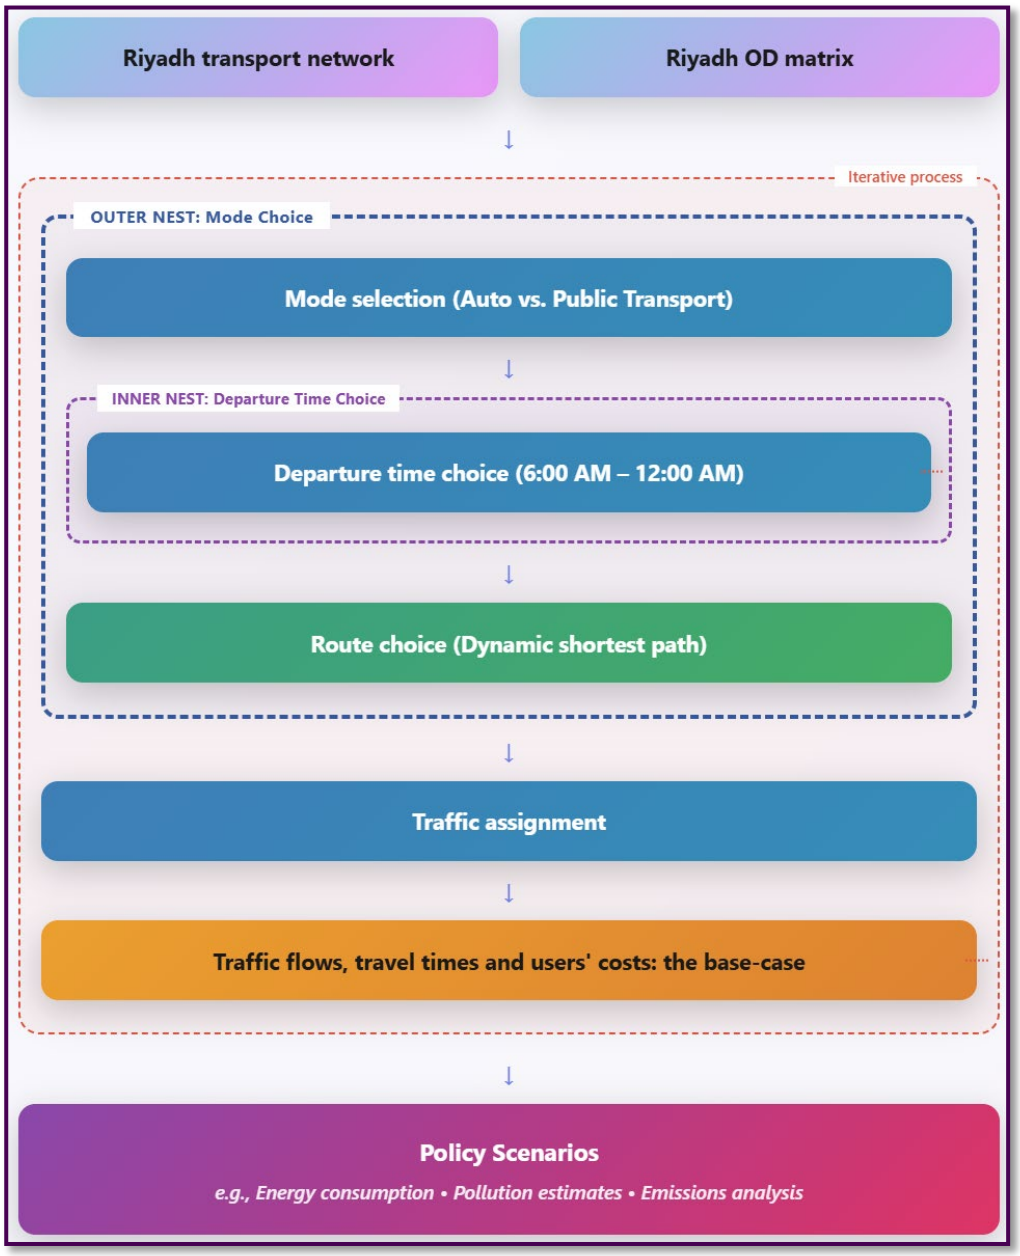

Note: Workflow of the simulation process in METROPOLIS2 for the current project. The two-stage nested logit structure shows mode choice in the outer nest, with departure time choice (6:00 AM–12:00 AM window) in the inner nest, followed by route choice based on dynamic shortest paths. The iterative process continues until convergence is reached. Policy scenarios include energy consumption analysis, pollution estimates, and emissions evaluation.Source: Authors' design.

Two types of data are required: on the supply side, the transport network; and on the demand side, the origin–destination matrix (along with behavioral parameters). This is followed by the users' decision tree: choice of transport mode (and, in the case of elastic demand, the decision whether or not to travel), choice of departure time, and choice of route. The simulation then produces travel times. By comparing simulated travel times with the times anticipated by users, individuals adjust their decisions on the following day until a stationary state is reached. This framework makes it possible to test various scenarios, such as the introduction of electric vehicles or changes in energy prices.

2.3. Demand modelling

To represent traveler behavior within the simulation, we define a generalized cost function that captures the main factors influencing individual travel choices.

We assume a cost function that represents the sum of several components: the travel time component, schedule delay component, and the pricing component. The generalized cost function for agent *k* is a sum of these components. For an agent *k*, departing at time *t*, facing travel time *tt(t),* the generalized cost is:

$$\underbrace{C_k\left(t;t_k^*\right)}_{\text{Generalized cost}} = \underbrace{\alpha_k tt^c\left(t\right)}_{\text{Travel time cost}} + \underbrace{\beta_k\left[t_k^* - t - tt^c\left(t\right)\right]^+}_{\text{Early schedule delay cost}} + \underbrace{\gamma_k\left[t_k^* - t - tt^c\left(t\right)\right]^+}_{\text{Late schedule delay cost}} + \underbrace{p_k\left(t\right)}_{\text{Toll}},$$

where $[a]^+ = a$ if and only if $a > 0$ and 0 otherwise, and where:

$$\begin{cases} \alpha_k =: & \text{Value of time} \\ \beta_k =: & \text{Unit cost for early arrival time} \\ \gamma_k =: & \text{Unit cost for late arrival time} \\ t_k^* =: & \text{Desired arrrival time at destination} \end{cases}$$

Here, [a]+ denotes the positive-part operator, defined as max(a, 0). It ensures that early and late arrival penalties are applied only when the corresponding deviation from the desired arrival time is positive. The term $p_k(t) > 0$ represents the toll incurred by individual *k*. for departure at time *t*. The no toll corresponds to $p_k(t) = 0$, and subsidies correspond to $p_k(t) < 0$.

Alternatively, the benefit of a trip is:

$$w_k = B_k - C_k\left(t;t_k^*\right),$$

where B_k represents the net benefit of the trip, performed by agent *k*. In this analysis, we set it to zero, a normalization factor.

In the public transportation model, the generalized cost comprises the access trip cost from the origin to the primary network (e.g., metro or subway), the travel time cost on the primary network, the egress trip cost (i.e., the final leg from the network exit point to the destination), and the schedule disutility associated with early or late departures).

This version assumes that the travel time for a given itinerary is constant and considers only the average waiting times for public transport modes (e.g., bus, metro, or subway). Congestion costs, which depend on the public transport occupancy rate and travel time uncertainty, are excluded in this approach.

Departure time choices are modeled using a standard continuous logit specification, as in the inner nest of Fig. 1. For this study, we select the day period 6:00 am to 12:00 pm.

More formally, the departure time model used here is a continuous logit. The cost function

Public

for a departure time at t is not known entirely by the modeler and described as a random variable as follows:

$$\tilde{C}_k\left(t;t_k^*\right) = C_k\left(t;t_k^*\right) + \mu\varepsilon ,$$

where $m > 0$ is a scale factor and ε are i.i.d. Gumbel-distributed random variables (with zero mean and unit variance). The CDF of ε_i are $F\left(x\right) = exp\left(-exp\left(-x\right)\right)$.

The probability of departure during time interval *(t, t+Dt),* given a continuous logit model, is:

$$P_k\left(t\right)\Delta t = \frac{\exp\left(-C_k\left(t;t_k^*\right)/\mu\right)}{\int_{T_i}^{T_f}\exp\left(-C_k\left(\tau;t_k^*\right)/\mu\right)d\tau}\Delta t ,$$

where T_i is the first departure time and T_f is the last departure time (here 6:00 am and 12:00 am, respectively).

The consumer surplus is given by standard "logsum formula" (Ben-Akiva and Lerman, 1985):

$$\Omega = \mu\log\left(\int_{T_i}^{T_f}\exp\left(-C_k\left(\tau;t_k^*\right)/\mu\right)d\tau\right)$$

2.4. The supply model

METROPOLIS2 route choice relies on routing algorithms known as Contraction Hierarchies (Batz et al., 2013; Geisberger and Sanders, 2010), which are more recent than those used in traditional transport simulators, offering improved computational speed. A two-level nested Logit model is employed to describe departure time and transport mode choices (Figure 3). At the upper level of the tree, the mode choice is analyzed (binary choice), while at the lower level, departure time is modeled by using a standard continuous Logit specification.

Each link consists of two segments: an initial segment followed by a final segment at the end of the link. Congestion is modeled at two levels within the links. In the first segment, speed is a decreasing function of the link's occupancy rate (flow congestion). Standard functional forms are provided to the user, who can also specify a custom function if desired. In the second, shorter segment, a bottleneck is present with an exit capacity. In a simplified representation, the travel time through this bottleneck is determined by the ratio of the number of users over the queue's capacity (homogeneity of degree zero in usage and capacity).

Public

Figure 3: An illustration of the nested logit model

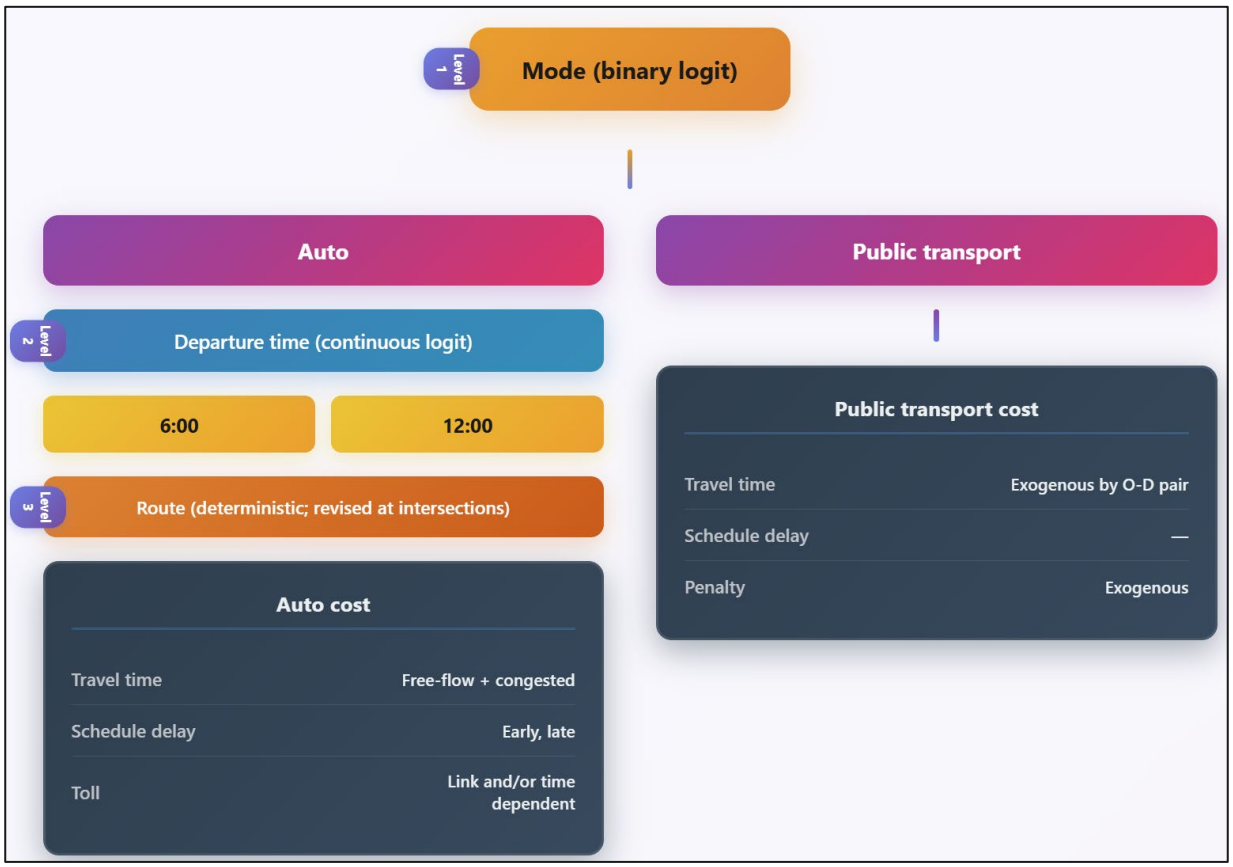

Note: At the top level, individuals choose between auto and public transport. Within the auto branch, departure time is determined by a continuous logit model, followed by deterministic route choice. Public transport costs are treated exogenously. Note: The Riyadh case study presented in this paper is unimodal (auto only). Public transport was included exogenously via service/timetable scenarios to assess its impact on fuel use and emissions. Source: Authors' design.

In addition to this static choice structure, METROPOLIS2 incorporates a day-to-day adjustment process to capture how travelers adapt over repeated interactions with the network.

This process, incorporating exponential learning by drivers, governs mode choice (using a standard binary logit model), departure time (based on a continuous time Logit model), and route selection changes (relying on a shortest path algorithm). These decisions are embedded in a hierarchical multinomial logit model. A day-to-day learning mechanism gradually guides the system toward a stationary state. A stationary solution is reached when the expected travel times (used the morning by the travelers) coincide with the simulated travel times (generated by the simulator). Further details about the decision processes and the learning process can be found in de Palma et al. (1997). Simulation experiments using different networks demonstrate that various adjustment processes consistently converge to the same stationary solution.

2.5. Public Transportation

To date, METROPOLIS2 relies on a simplified version of the OpenTripPlanner planning tool for public transport modeling (de Palma and Javaudin, 2024). However, this version does not yet account for train congestion or schedule variability, both of which significantly impact transit system performance and user choices. In this model, the total travel cost function is assumed to be the sum of four components:

1. Access cost – the cost incurred traveling from the origin to a transit station.
2. In-network travel cost – the cost of the trip within the primary transit network.
3. Scheduling cost, including the penalty for departing either too early or too late relative to the preferred schedule.
4. Egress cost – the cost of reaching the destination from the transit station.

In this analysis, we employ a simplified modelling framework in which the public transport network is superimposed onto the private transport network, thereby generating a super-network. This structure enables the representation of modal choice between private and public transport. The methodological and policy implications of this approach will be examined in greater detail in the policy discussion.

3. Riyadh Metropolitan Area: Emissions, Mobility, and Input Data for Emissions and Pollution Simulation Model

In transport simulation, as is standard practice, we require two primary inputs: the transportation network and the demand, which is typically represented by an origin-destination (OD) matrix. The transportation network describes the city's physical infrastructure, including roads, intersections, bus stops, metro stations, and other relevant elements that define how people and goods can move from one location to another. The OD matrix, on the other hand, represents the flow of trips between different zones or points in the city, capturing the number of trips generated from each origin and assigned to each destination during specific time periods. In this study, both the transportation network and OD matrix data have been collected specifically within the context of Riyadh city's metropolitan area.

3.1 Riyadh's Urban Mobility Landscape

Riyadh is the capital and largest city in Saudi Arabia. The city is experiencing rapid urbanization and significant population growth, transforming it into a dynamic hub for economic, social and cultural activity. Home to more than 7 million residents, Riyadh is evolving from a predominantly car-dependent toward a more sustainable and accessible urban environment. Historically, Saudi Arabia, including Riyadh city's mobility landscape, has been characterized by private vehicle dominance, resulting in considerable traffic congestion and challenges to urban livability (Al-Mosaind 1998). Moreover, rapid population growth in Riyadh has surged demand for personal transportation, with Riyadh province accounting for 30% of all registered vehicles in Saudi Arabia mainly concentrated in Riyadh city.(Zeaara 2022; Ikrami 2014)

Riyadh is located in the central part of the country. It is strategically positioned at the crossroads of major trade routes, making it a pivotal hub for commerce and culture in the

Arabian Peninsula. As of 2023, Riyadh boasts a population of approximately 7.5 million residents, making it the largest city in Saudi Arabia and one of the fastest-growing cities in the world.

The city is home to numerous corporate headquarters, financial institutions, and multinational companies, particularly in the King Abdullah Financial District (KAFD), which is designed to be a leading financial hub in the region. Riyadh's business environment is further enhanced by its modern infrastructure, including extensive road networks, a growing public transportation system, and state-of-the-art commercial facilities. Key locations for business activities include the Olaya district, known for its skyscrapers and corporate offices, and the Diplomatic Quarter, which houses embassies and international organizations. The city also hosts several major events and exhibitions throughout the year, attracting global business leaders and investors.

Riyadh's transportation infrastructure is continually evolving to meet the demands of its rapidly growing population and economic growth. Despite its extensive network of roads and highways designed to facilitate the movement of people and goods, the city faces substantial traffic congestion, especially during peak periods. Figure 4 illustrates the speed reduction index (SRI) over time (a measure of congestion level), indicating extreme lags in speed due to traffic congestion.

Figure 4: Average SRI level by Hour in Riyadh city highways based on October 2022 TomTom Data for weekdays (FRC=0;1)

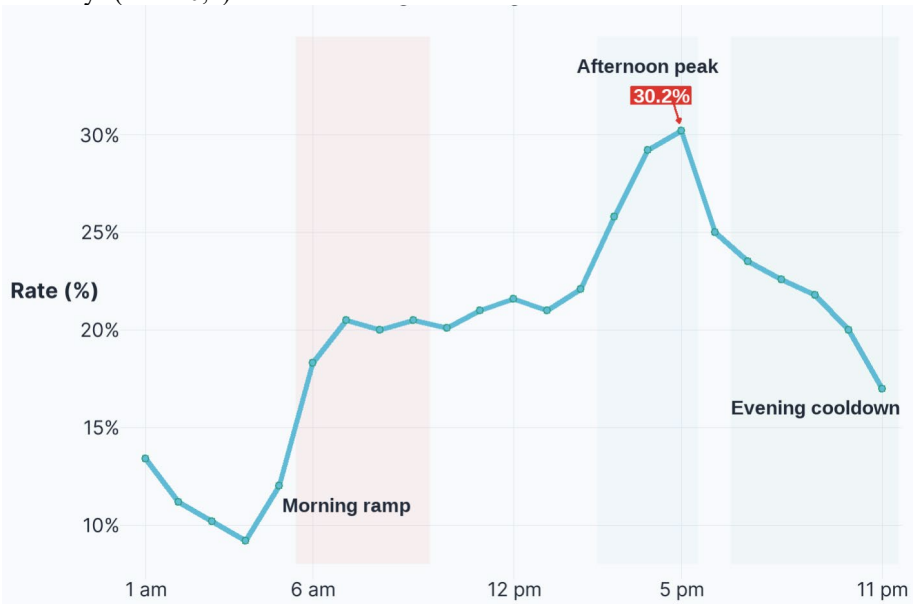

Source: Authors' design

The prevalence of car ownership in Riyadh can be attributed to several interconnected factors. Firstly, fuel prices in Saudi Arabia are among the lowest globally, significantly reducing the cost of vehicle operation. Many residents opt for loans to finance their vehicles, allowing them to manage costs over time despite high interest rates. The city's underdeveloped public transportation systems further contribute to this reliance on personal vehicles, as they fail to offer an efficient or accessible alternative. This high concentration underscores the strain placed on local infrastructure and resources. Aldalbahi & Walker

(2016) highlight challenges such as rising car ownership and insufficient public transport infrastructure. The Metropolitan Development Strategy (MEDSTAR) emphasized the need for high-quality public transit as part of Riyadh's urban planning since the early 2000s.

Sustainable transport strategies in Riyadh focus on reducing congestion and pollution through integrated metro and bus systems. The project is designed to handle up to 4.5 million passengers daily, connecting key areas such as universities, airports, and commercial hubs (Al Zohbi, 2021). Al Hosain & Alhussaini (2021) evaluated accessibility to Riyadh's planned public transport system, including 85 metro stations and a three-tier bus network.

It found that 74% of the population could drive to a metro station within 10 minutes, while only 14% could walk there within the same time frame. The study highlighted areas needing improvement to enhance accessibility and usability of the public transport system. Almatar (2022) establishes a links with transit oriented development (TOD), highlighting the objective of congestion and pollution through integrated metro and bus systems. To advance these goals, , several initiatives have been implemented. The King Abdulaziz Project for Riyadh Public Transport includes the first mass transit metro system in Saudi Arabia and one of the largest in the world, featuring six lines connecting key areas of the city, supported by an extensive bus network with three bus rapid transit (BRT) lanes. Launched in December 2024, the public transport system had already served approximately 18 million passengers by mid-February 2025, underscoring its immediate impact and strong public adoption.

Furthermore, the city is increasingly exploring micro-mobility solutions, such as bike-sharing and e-scooters, to provide effective last-mile connectivity and promote sustainable commuting practices. Introducing citywaid paid parking systems across the city is also helping manage traffic congestion, encourage public transportation, and improve overall urban mobility.

Additionally, Riyadh is implementing several mega-projects with large urban impacts, including The Sports Boulevard, King Salman Park, Diriyah Gate, and the Green Riyadh initiative, all designed to enhance urban connectivity, create pedestrian-friendly environments, and further support sustainable mobility (Royal Commission for Riyadh City 2025a; Dariyah Gate Development Authority 2025). In parallel, the city has invested heavily in public transit—most notably the Riyadh Metro—supported by an extensive bus network and BRT lanes. Complementary bus services and growing ride-sharing options expand coverage and convenience. These investments aim to build an integrated, sustainable system that reduces congestion and improves quality of life.

3.2 Road Network Data

In order to conduct a transport simulation using multi-agent models, a comprehensive dataset is essential to accurately represent the road network and the movement patterns of various agents, such as vehicles, pedestrians, and cyclists. The road network data should include detailed information about the geometry of the roads, including lane configurations, speed limits, traffic signals, and intersections. Additionally, attributes such as road capacity, types of road surfaces, and the presence of any physical barriers or features (like bridges and tunnels) are crucial for simulating realistic traffic flow. Geographic Information System (GIS) data can be particularly useful in providing spatial context and ensuring that the road network is accurately mapped.

The transportation network used in this study was imported from the OpenStreetMap database, a freely available and widely used spatial database that provides detailed road and geographic information for cities around the world. However, inherent inconsistencies in raw OpenStreetMap data—such as tiny edges, unconnected nodes, and rarely used minor links—made a cleaning phase essential. This process involved removing or correcting such anomalies to ensure the network's integrity and compatibility with the METROPOLIS2 simulator.

The city of Riyadh features a robust transportation network with 148,442 nodes and 375,708 edges, indicating extensive connectivity. The network density of 2.531 reflects an average degree of connections per node, suggesting a well-interconnected system. Geographically, the network spans a bounding box defined by WGS84 coordinates, covering specific areas within the city. With a total road length of 38,574.9 km and an average edge length of 0.102673 km, the network comprises mostly short segments, likely indicative of local streets or grid-like structures. The presence of both very short and longer (maximum 16.849 km) road segments suggests a mix of minor intersections and major thoroughfares within Riyadh's urban landscape. The traffic lights are not explicitly modeled in METROPOLIS2, but they are modeled using intersection penalties.

Table 1: Network Statistics and summary of Road Type Counts in Riyadh

Information and topology of the road network		**Statistics by road types**	
Features	**Values**	**Groups**	**Count**
Node Count	148,442	Residential	285,525
Edge Count	357,270	Tertiary	33,081
Network density (average degree)	2.531	Secondary	14,498
Bounding Box Area (WGS84)	46.463776, 24.4851, 46.9853, 24.9582	Unclassified	9,521
Total Road Length	38,574.9 km	Primary	7,262
Average Edge Length	0.102673 km	Motorway link	1,384
Maximum Edge Length	16.849 km	Motorway	1,384
		Primary link	1,124
		Tertiary link	1,663
		Secondary link	1,060
		Trunk	650
		Trunk link	111
		Service	4
		Pedestrian	3
		Total	357,270

Source: OpenStreetMap, after processing. Note: The network has been imported from OpenStreetMap, then processed to remove some inconsistencies (very short links, unconnected nodes).

Figure 5 illustrates Riyadh's road network alongside the six recently operational metro lines. Due to the network's dense connectivity, only nodes are displayed (edges omitted). The main panel on the left offers an overview of the complete network, while two panels on the right provide detailed views of selected urban zones. This visualization highlights the metro system's integration within the city's broader transportation infrastructure and demonstrates the distribution of mobility demand. The metro line geometries were obtained from the open data service provided by the Royal Commission for Riyadh City

(https://opendata.rcrc.gov.sa/pages/home-page/).

The distribution of road types in Riyadh reflects its predominantly residential character (**Table 1**). "Residential" streets dominate the city's network (285,525 links), underscoring their role as the primary urban fabric. The second-largest category, "unclassified" roads (9,521), likely comprises additional minor or local streets. The presence of significant highway infrastructure is indicated by 1,384 "motorway" segments and an equal number of "motorway links," supporting major transportation demands. Conversely, categories such as "pedestrian walkways" (3) and "service" roads (4) are notably limited, potentially due to OpenStreetMap (OSM) data constraints, and will therefore be excluded from this initial analysis. The relatively substantial count of "tertiary" roads (33,081) points to an extensive minor arterial network, whereas fewer "trunk" (650) and "trunk link" (111) segments suggest limited major highways interconnecting the city. Overall, Riyadh's road-type distribution reveals extensive local connectivity complemented by structured arterial routes, alongside minimal infrastructure supporting specialized transportation modes, such as pedestrian pathways.

To perform spatial operations effectively, nodes within the network must be matched precisely to their respective zones. This involves verifying whether a given node lies within or on the boundary of zone polygons, ensuring accurate spatial alignment of nodes and edges with defined zone boundaries. For technical details on the algorithm used to assign network nodes to the 162 traffic analysis zones, refer to **Appendix 2**.

Figure 5: The network imported from OSM (after cleaning).

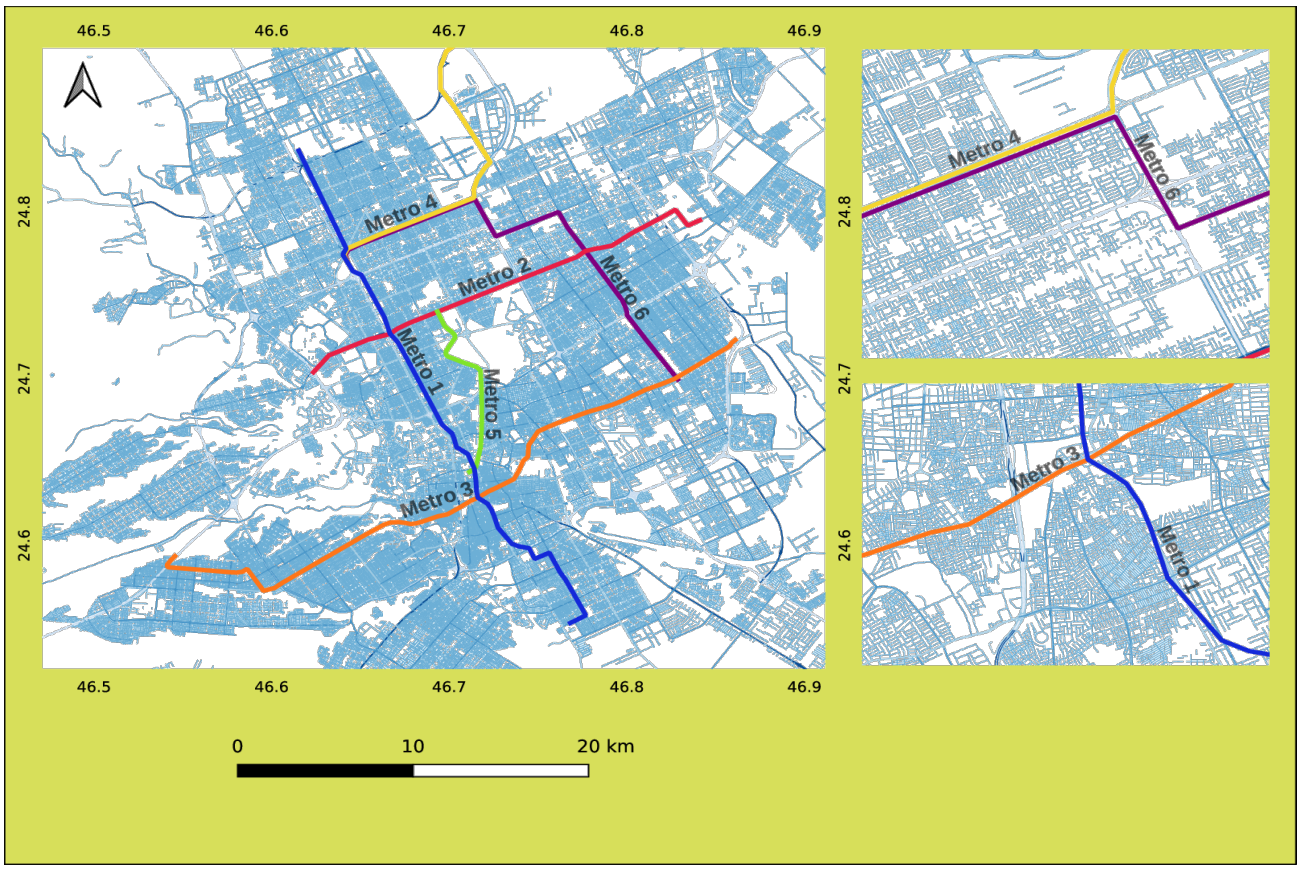

Source: Authors' design. Note: The figure shows the road links with the most relevant edges and the six metro lines that have recently been put into operation. The main panel provides an overview of the entire network,

while the two panels on the right focus on smaller, more detailed zones within the city.

3.3. The Origin-Destination (OD) Matrix and Travel Demand Patterns

Riyadh city is composed of 87 officially recognized neighborhoods (Riyadh Region Municipality 2025), which are further subdivided into around 177 subdistricts or localities for statistical and administrative purposes (Royal Commission for Riyadh City 2025b). The OD matrix for this study, captures travel patterns across 162 zones representing a subset of Riyadh's subdistricts. The data includes the subdistricts in the metro area covering areas with the highest population density and public transport accessibility (Figure 6a), providing a comprehensive view of trips within and between these zones. The dataset reports a sample of 24,159,087 trips recorded for October 2022 for the weekdays from 7:00 – 17:00, reflecting extensive mobility within the city.

The OD matrices quantify the travel demand between different zones within the study area. These matrices provide insights into the volume of trips generated from specific origins to various destinations, reflecting human mobility patterns.. By integrating detailed road network data with robust OD matrices, multi-agent models can effectively simulate traffic dynamics, allowing for the evaluation of different transportation scenarios and the development of targeted strategies to improve urban mobility. Figure 6b demonstrates a filtered sample of the OD Matrix showcasing flows between different zones, excluding intra-zonal trips.

Public

Figure 6: Population, inter-zonal flows, and origin trip density in Riyadh (Oct 2022, weekdays 07:00–17:00).

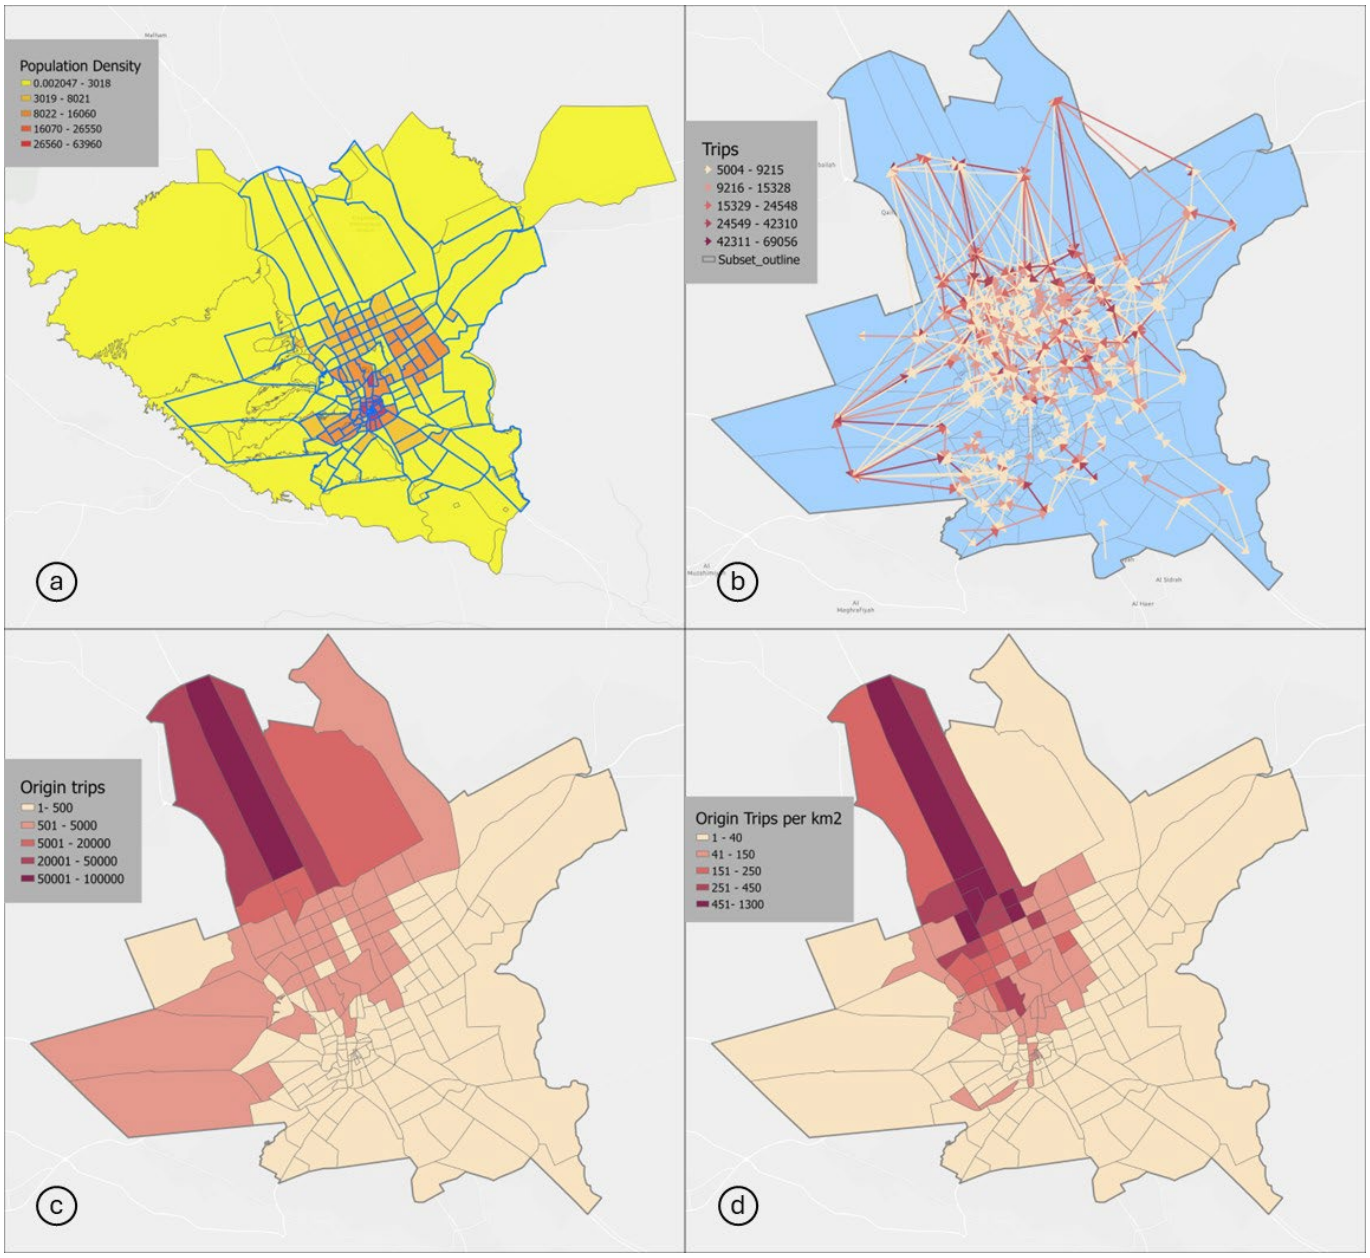

Note: **a,** Population density by subdistrict with road network overlay. **b,** Filtered inter-zonal OD flows (intra-zonal trips excluded); arrow width denotes trips. **c,** Origin-trip counts by zone. **d,** Origin-trip intensity (trips per km²). The OD matrix covers **162** zones in the metropolitan core and records **24,159,087** trips; demand is dominated by short-distance, intra-zonal travel (c–d). Source: Authors' design.

The observed density of trips in localized areas reflects the predominance of intra-zonal travel (Figure 6 c and d). (Figure 6c, 6d). This pattern is consistent with urban transportation trends, where short-distance trips typically account for the largest mobility share due to their frequent and essential nature. The dominance of intra-zonal trips underscores the importance of local accessibility and connectivity in Riyadh's urban structure.

Additionally, the OD dataset currently lacks socioeconomic characteristics, such as trip purpose, which were not captured in this OD subset but could be enriched through targeted survey efforts. In the interim, trip purposes will be estimated based on general assumptions, for example classifying flows to employment centers during morning peaks as work trips, return flows in the afternoon as homeward trips, and the remainder as education, shopping, or other activities.
To implement trips in the simulation, we assign trips from the OD matrix to nodes within the transportation network. Given that the OD matrix lacks detailed information regarding precise building-level locations within zones, we identify representative nodes within each zone to serve as origin or destination points. Although using precise building locations would increase accuracy, the current approach offers a reasonable approximation of urban travel patterns based on node distributions. The ray casting algorithm described earlier was employed at this stage to select representative nodes. Figure 6c shows traffic volumes distributed throughout the metropolitan area. However, when examining the flow density representation in Figure 6d, traffic concentrations are notably higher in central zones, while substantial traffic volumes persist across suburban areas.

Figure 7 presents a section of Riyadh city, highlighting two routes between two random locations from the Al Hamra (subject in Baladiyah al-Rawdah) and Al Murabba districts. The blue line represents the shortest path, which spans approximately 16 km, while the red line shows an alternative route that is longer but potentially less congested. (METROPOLIS2 computes these time-dependent paths with fast, parallelized dynamic shortest-path algorithms.) The visualization aims to demonstrate how overuse of the shortest path can lead to road congestion, slowing down travel times, and reducing its attractiveness for users. In such cases, alternative routes like the red one may become more appealing due to their ability to bypass heavily trafficked areas. Within the simulation model, at each iteration, the system evaluates available route options and selects the most convenient one for individual travelers. However, if no better alternatives emerge for a user, the chosen route remains consistent across iterations until traffic conditions change.

Figure 7: A route choice situation

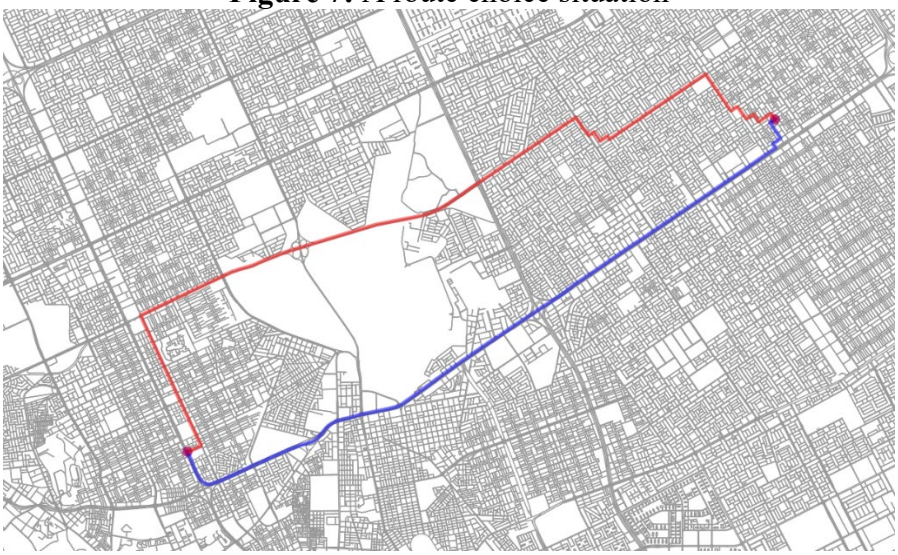

Source: Authors' design.

3.4. Filtering the OD Matrix: pre-processing the demand data for the calibration step

In order to prepare the ground for model calibration, we discuss here the operations performed on the data to produce a dynamic model. We adopted this approach, combining data on the city and the OD matrix of flows, because we found it to be the most relevant given the data at our disposal. Dynamic simulations rely on precise departure and arrival times for trips. However, since this temporal data is unavailable in the OD matrix, we analyze trips based on origin and destination zones instead. After preliminary testing, we filtered the OD matrix to better capture trip patterns through a structured approach involving several key steps:

1. **Identification of main activity locations**: A web search was conducted to identify major activity centers within Riyadh, such as business districts, educational institutions, and shopping areas (see Table 2 and Figure 8 below).
2. **Assumptions on trip purposes**: Based on the identified activity locations, we made assumptions regarding trip purposes. Trips with destinations in these zones were classified as home-to-work trips with a high probability, while trips in the opposite direction were categorized as work-to-home trips.
3. **Intra-zone trips**: We recognized that intra-zone trips, which account for approximately 30% of total trips, should be grouped together to simplify the model. These trips are those having departure and arrival nodes within the same zone. Given the large number of generated trips, it is possible to randomly draw the same node for departure and arrival locations. Our script checks and removes these odd trips. Intra-zone trips have short distances and are responsible for a large part of congestion in the secondary roads.
4. **Other trips**: A fourth category was established for "other trips," which encompasses various purposes such as shopping, leisure, or accompanying others.

Through a web search, we have identified a number of main activity areas that are reported in Table 3. Diverse activities are considered, including administration, education, commerce, and finance. Trips arriving at (*resp.* departing from) these districts are considered as home-to-work trips (*resp.* work-to-home) with a higher probability than the other trips. In the OD matrix considered, the proportion of home-to-work trips is 17.0% and the proportion of work-to-home trips is 16.3%.

Table 2: Main activity locations in the study area

Locations	Short description
King Abdullah Financial District (KAFD)	A business hub that houses numerous corporate offices, financial institutions, and commercial establishments. KAFD attracts a significant number of professionals during the morning rush as employees commute to work.
Olaya District	Commercial and retail activities, Olaya is home to many shopping centers, restaurants, and office buildings. The area is a popular destination for both workers and shoppers.
Diplomatic Quarter	Hosts various embassies, consulates, and international organizations. It attracts diplomats, government officials, and expatriates, leading to increased traffic as individuals commute to and from work.
Al-Malaz	A mix of residential and commercial areas, featuring government offices, educational institutions, and healthcare facilities. Al-Malaz sees a significant influx of commuters, particularly those traveling to government offices and schools.

Public

Riyadh City Center	Encompasses several key attractions, including cultural institutions, parks, and entertainment venues. It draws visitors and residents alike, contributing to morning traffic as people head to work or engage in leisure activities.
Universities and Educational Institutions	Areas surrounding major universities, such as King Saud University and Princess Nourah bint Abdulrahman University.

The map in Figure 8 illustrates zones with high job concentrations, corresponding to the data presented in Table 2. Within each of these zones, 100 nodes were randomly sampled (indicated by yellow dots) to validate the algorithm used for selecting representative origin and destination nodes.

Figure 8: Zones with High Employment Density and Randomly Sampled Nodes for Origin–Destination Validation

Note: The reported map was produced on the basis of web search focusing on the study area. The yellow dots are nodes drawn randomly in specific area. The orange dots correspond to nodes that do not belong to any zone and are not considered neither as origin nor as a destination. Source: Authors' design.

2. Illustrative simulations

In this section, we present illustrative results from the initial simulations conducted for the city of Riyadh using the METROPOLIS2 simulator. It is important to note that these early simulations are not intended to provide realistic representations of traffic dynamics; rather, they serve to illustrate the workflow and the iterative process involved in developing the

multi-agent transport model. At this stage, our focus is on demonstrating how the model evolves through successive iterations, highlighting the adjustments made to the parameters and the resulting changes in agent behavior and traffic flow. This foundational work lays the groundwork for the subsequent calibration phase, where we will refine the model to enhance its accuracy and alignment with real-world conditions.

In Figure 9, there are two plots. The left panel reports average "departure time" (purple) and average "arrival time" (green). Both curves are decreasing along the x-axis, indicating that as iterations progress, agents depart and arrive earlier. However, there is a notable difference in their rates of decrease. Initially, the arrival time curve decreases sharply due to agents' tendency to use the shortest routes. This overuse leads to congestion on heavily traveled links, particularly around the city center. Over iterations, they adapt by exploring alternative routes or adjusting their departure times to reduce travel costs. Over the iteration, we see that agents depart earlier, which helps alleviate congestion and further reduces arrival times (green curve). The "travel time", which corresponds to the vertical distance between the two curves, decreases significantly over iterations. This decline reflects the overall improvement in efficiency as agents optimize their routes and timing to minimize delays caused by congestion.

The right panel reports congestion level and utility of the agents. The utility can be understood as a measure of generalized costs (travel time and schedule delay costs) related to mobility. Congestion is measured as the difference between travel time observed in the simulation and free-flow travel time. It is clearly decreasing and, in these simulations, it converges to about 40% travel time. The utility costs increase because the user costs decrease due to the lower congestion and schedule delay costs.

Figure 9: Output of the illustrative simulations: departure and arrival times (the left panel) and utility (right panel); average value.

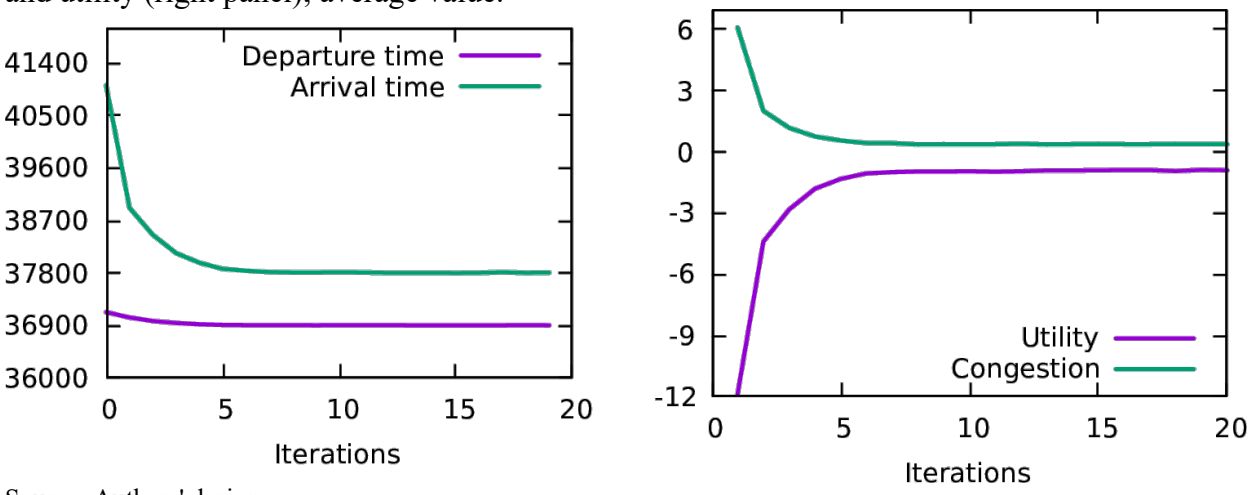

Source: Authors' design.

The simulation runs demonstrate stable model behavior and efficient convergence. Computation time scaled proportionally with the number of agents, remaining tractable even for scenarios involving 200,000–300,000 agents (approximately one hour per 50 iterations). This performance indicates the model's suitability for city-scale applications.

Figure 10a illustrates how the average travel distance evolves across iterations. Agents initially choose the shortest routes, generating congestion and delays. As congestion intensifies, rerouting

leads to longer travel distances before the system gradually stabilizes. At equilibrium, average distances remain above the initial values due to network constraints and decentralized routing decisions. Figure 10b presents the corresponding CO_2 emissions profile, which closely mirrors distance patterns under the homogeneous fleet assumption adopted here.

Figure 10: Simulation outputs showing convergence of average travel distance (a) and CO_2 emissions per trip (b) across iterations.

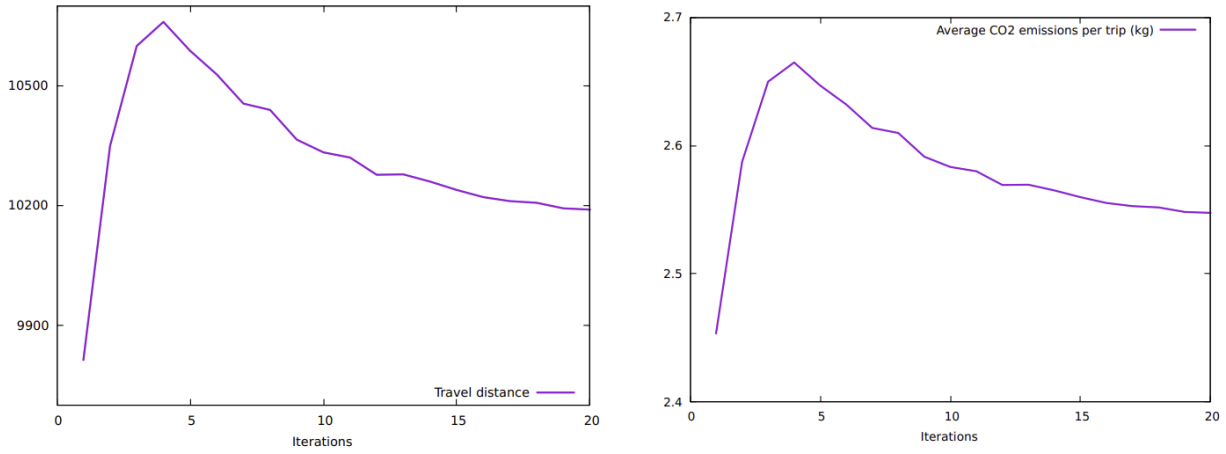

(a) Output of the first simulations: travel distance in meters.

(b) Average CO_2 emissions on an emissions ratio of 250g/km.

Source: Authors' design.

These results highlight the framework's ability to capture adaptive behavior, convergence dynamics, and the link between mobility and emissions. While further refinements—such as integrating detailed fleet composition—would improve the precision of environmental assessments, the present simulations already demonstrate the model's capacity to reproduce key traffic and emission dynamics in Riyadh.

In this methodology paper, best-guess parameters have been employed. A particularly critical parameter widely utilized in transportation modeling is the value of time (VoT), which can be derived either from econometric estimation or approximated using related socioeconomic indicators, such as average wage levels.

To convert the estimated VoT for transportation in Riyadh from Saudi Riyals (SAR) to US Dollars (USD), we use the exchange rate of 1 SAR $\simeq$ 0.27 USD. As mentioned above, in the absence of econometric studies, the VoT is typically based on wages, and we obtain using a back of the envelope computation: The average monthly salary in Saudi Arabia varies by sector, but a reasonable estimate is SAR 10,000 – 15,000 per month ($2,600 – $4,000). This translates to an hourly wage of around SAR 50 – 75.

Then, the VoT as a percentage of Wages will depend on travel motives:

- For commuting or personal travel, VoT is often 30% – 50% of the wage rate.
- For business trips, VoT can be 100% or more since time directly affects productivity.

Using these assumptions, we can consider the following values

Public

- Personal Travel: $4 – $11
- Work Commute: $5 – $14
- Business Travel: $14 – $20

This suggests that, on average, Riyadh residents value an hour of travel time at around $5/h – $14/H for commuting and potentially more for work-related travel. Based on the average number of cars and congestion produced by METROPOLIS2, we can directly assess the congestion costs at the city level.

4. Calibration

This section outlines the calibration process of the multi-agent simulation model we have implemented for Riyadh. We provide an overview of the procedure used, detailing the progressive steps taken to match simulated outputs with observed urban mobility and congestion patterns. This covers the experimental setup, OD matrix refinement, simulation parameterization, validation procedures, and integration of emissions and economic evaluation modules.

4.1. Network and OD Matrix Testing

Initial Experiments

In the initial experiments, we tested the network by simulating a small number of vehicles to verify their selection of optimized routes. Given the absence of congestion under these conditions, the simulation outcomes corresponded directly to the shortest-path solutions. Subsequently, we incorporated the origin–destination matrix but intentionally limited flows by focusing either on a single origin or a single destination. This targeted approach proved helpful in identifying the city's primary trip-generation and trip-attraction zones. By isolating these specific flows, we gained clearer insights into areas contributing most significantly to urban transport demand.

4.2. Simulation Parameters

For each simulation, we need to specify several parameters, providing:

- Start and end times for the simulated period: The simulations are run for the whole day, including the morning and the evening commute, from 6:00 am to 10:00 pm
- Value of time: In the simulator we use normalized monetary units and convert all costs to Saudi riyal for reporting. We set $\alpha = 10$ monetary units per time (mu/tu), where "mu" denotes monetary units and "tu" corresponds to the simulation's time step (minutes) (Small, K. A. (1982)).
- Schedule delay cost: Early arrival cost is $\beta = 7$ mu/tu and late arrival cost is $\gamma = 15$ mu/tu
- OD matrices with for specific time periods: The morning rush-hour ranges from 8:00 am to 11:00 am and the evening rush-hour from 3:00 pm to 5:00 pm.
- Scaling of the network: The capacities of each edge are rescaled to match the reduced size of the simulated population, ensuring that congestion levels remain consistent with real-world conditions.
- Iterations: About 150 iterations —equivalently, 150 simulated days—was necessary

Public

for convergence to a stationary regime.

Each simulation run incorporates four parameters corresponding to the importance of each trip category (intra-zone, home-to-work, work-to-home and other trips). Additionally, for each part of the OD matrix, we need to specify the time of the first and last departure. The whole OD matrix is obtained by adding trips from specific matrices over the simulation period (full day). For instance, during the morning rush hour, the majority of the trips are expected to be home-to-work trips. Therefore, the parameters for this group are set to higher values, while the other three groups are assigned relatively lower values. Conversely, during the evening rush hour, the proportion of home-to-work trips diminishes significantly, while a larger share is allocated to work-to-home trips. Notably, our analysis indicates that congestion levels are generally higher during the evening peak compared to the morning peak. To achieve comparable patterns to observed ones, we have adjusted the parameters for other trip purposes to larger values during the evening peak time.

4.3. Calibration and Validation Process

The calibration focuses mainly on the adjustment of traffic flows over several time segments covering the whole day. The distribution of the desired arrival times (the t^* parameter) for the trips is estimated by a piece-wise linear approximation of the observed congestion trend (observations are based on the *TomTom* website: https://www.tomtom.com/traffic-index/riyadh-traffic/). Further, the capacity of the network is adjusted to meet the demand for trips. To save computation time, we proceed with relatively large road capacities (low congestion), then reduce the capacities to obtain consistent travel time values.

More precisely, we proceed as follows. The day time is decomposed into a number of n time segments. Transport flows in each segment i are described by a column vector with six elements $v_i^T = (t_1^i, t_2^i, s_1^i, s_2^i, s_3^i, s_4^i)$ for $i = 1, .., n$. In time segment i =*1, .., n*, t_1^i and t_2^i, respectively, correspond to the start and end times of the time segment. s_1^i, s_2^i, s_3^i and s_4^i, respectively, correspond to the ratio of trips from four distinct parts identified in the population: the local, or within-zone, trips (s_1^i); the work-to-home trips (s_2^i); the work-to-home trips (s_3^i); and, the 'other' trips (s_4^i). The identification of the trip category follows the methodology described in Section 3.4. In our process, we have decomposed the simulated period (one day) into ten time periods (i.e., n = 10).

The objective of the calibration is to fit values in vectors $v_i, i = 1, .., n$, so that simulated flows correspond to observed flows. Explicitly, the quadratic loss function that we minimize (see De Lara et *al*., 2013, for a similar calibration approach) writes as

$$L(v_1, .., v_n) = w_t \frac{1}{N_t} \sum_{i=1,N_t} \left(\frac{t_{sim} - t_{obs}}{t_{obs}}\right)^2 + w_c \frac{1}{N_c} \sum_{i=1,N_c} \left(\frac{C_{sim} - C_{obs}}{C_{obs}}\right)^2,$$

Where 't' refers to travel time and 'C' for congestion level. A set of twenty trips were considered for the first term (travel times, i.e. N_t = 20) and a set of ten time of days were considered from the second term (congestion level, i.e. N_c = 10). The samples were drawn in distinct trips in the study area at distinct time of day. the 'sim' and 'obs' indexes refer, respectively, to simulated and observed values. w_t and w_c are weighting parameters, both normalized to one for this case.

Public

The iterative process used in the calibration works as follows. At each iteration, we proceed by evaluating L at many points, then select the one $\bar{v} = (v_1, .., v_n)$ that yields the minimum value (notice that at each time segment, we have four parameters s_1^i, s_2^i, s_3^i and s_4^i). We then consider a set of new points distributed around $\bar{v}$, and so on until the value of L no longer decreases significantly. The starting point uses the intuition that in the morning rush-hour most trips are home-to-work trips, and in the afternoon, most trips are predominantely work-to-home trips. The values obtained through the calibration process, which yielded a loss function at about 16%, are reported in Table 3.

Table 3. Parameter values obtained after calibration.

Time periods			**Adjusted during calibration**				
	t_1	t_2	s_1	s_2	s_3	s_4	**Comments**
1	08:20	12:30	0.005	0.090	0.005	0.008	Morning rush-hour
2	09:30	11:30	0.005	0.020	0.005	0.005	
3	11:00	12:00	0.000	0.000	0.000	0.005	
4	10:30	11:30	0.002	0.020	0.001	0.002	
5	12:00	14:00	0.020	0.005	0.020	0.020	Mid-day peak
6	13:30	14:30	0.010	0.002	0.010	0.010	
7	14:30	17:30	0.025	0.005	0.075	0.045	
8	15:00	19:00	0.005	0.015	0.085	0.015	Evening rush-hour
9	17:00	20:00	0.005	0.015	0.025	0.015	
10	08:00	21:30	0.010	0.003	0.003	0.023	
					Objective function		
					$L(v_1, .., v_n)$	16.3%	

Note: *The parameters that were adjusted during the calibration were the flows from the OD matrix for distinct time periods (ten are considered here), reflecting the main peak and off-peak periods of the day. The loss L in the initial runs (before calibration) was above 50%. Road capacities were (manually) adjusted to manage the overall congestion level in the outer loop. The calibrated model with these parameters has 108 thousand agents running on 6% scaled capacity network.*

4.4. Initial simulation output

After the calibration process, we conduct illustrative simulations of daily traffic flows in Riyadh. As anticipated, road usage is notably higher along the city's main axes and around the central zones. These usage patterns, aggregated over an entire day, are depicted in Figure 11.

Public

Figure 11: Calibrated Daily Traffic Volumes from Multi-Agent Transport Simulation for Riyadh

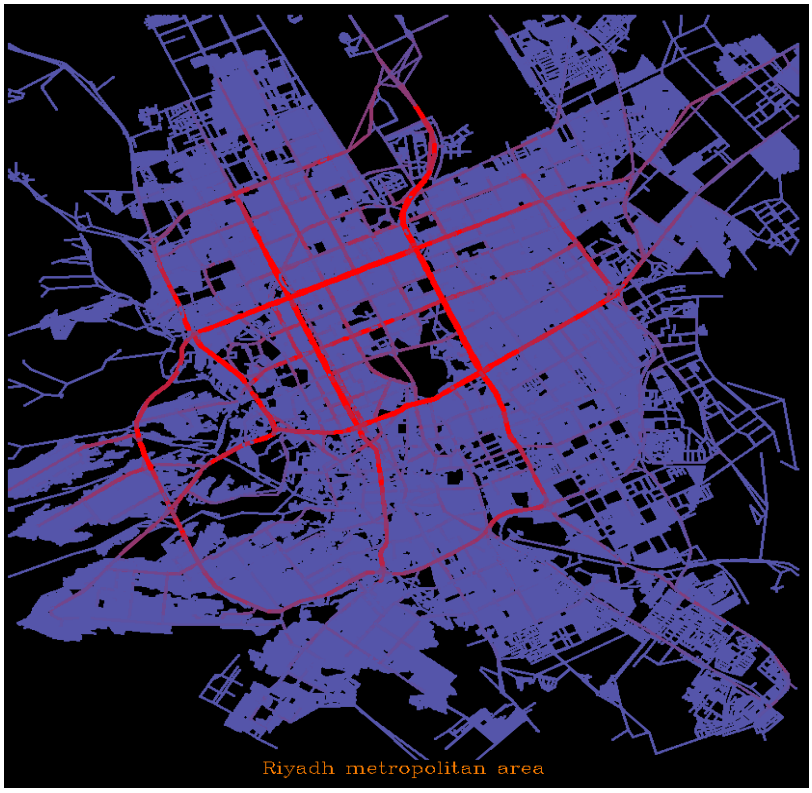

Source: Authors' design.

Figure 12 illustrates the dynamics of trip departures and arrivals as a function of time of day. In the upper panel, yellow dots represent departure times, while blue dots indicate arrival times. The y-axis denotes congestion levels—that is, the ratio of actual to free-flow travel time expressed as a percentage. Each trip could theoretically be depicted by a line connecting departure and arrival points; however, due to the large volume of trips, this simplified dot representation provides a clearer visualization. Although congestion can exceed 200%, the axis is shown up to 200% for readability.

The lower panel displays cumulative percentage of trips **0–1000–100%0–100** over time for departures (yellow) and arrivals (blue). In a first-in-first-out (FIFO) system, the horizontal distance between these cumulative curves corresponds directly to travel times; here, it serves as an approximate indicator. Instantaneous departure and arrival rates correspond to the slopes of the cumulative-percentage curves.

Public

Figure 12: Daily Dynamics of Trip Departures, Arrivals, and Congestion Levels

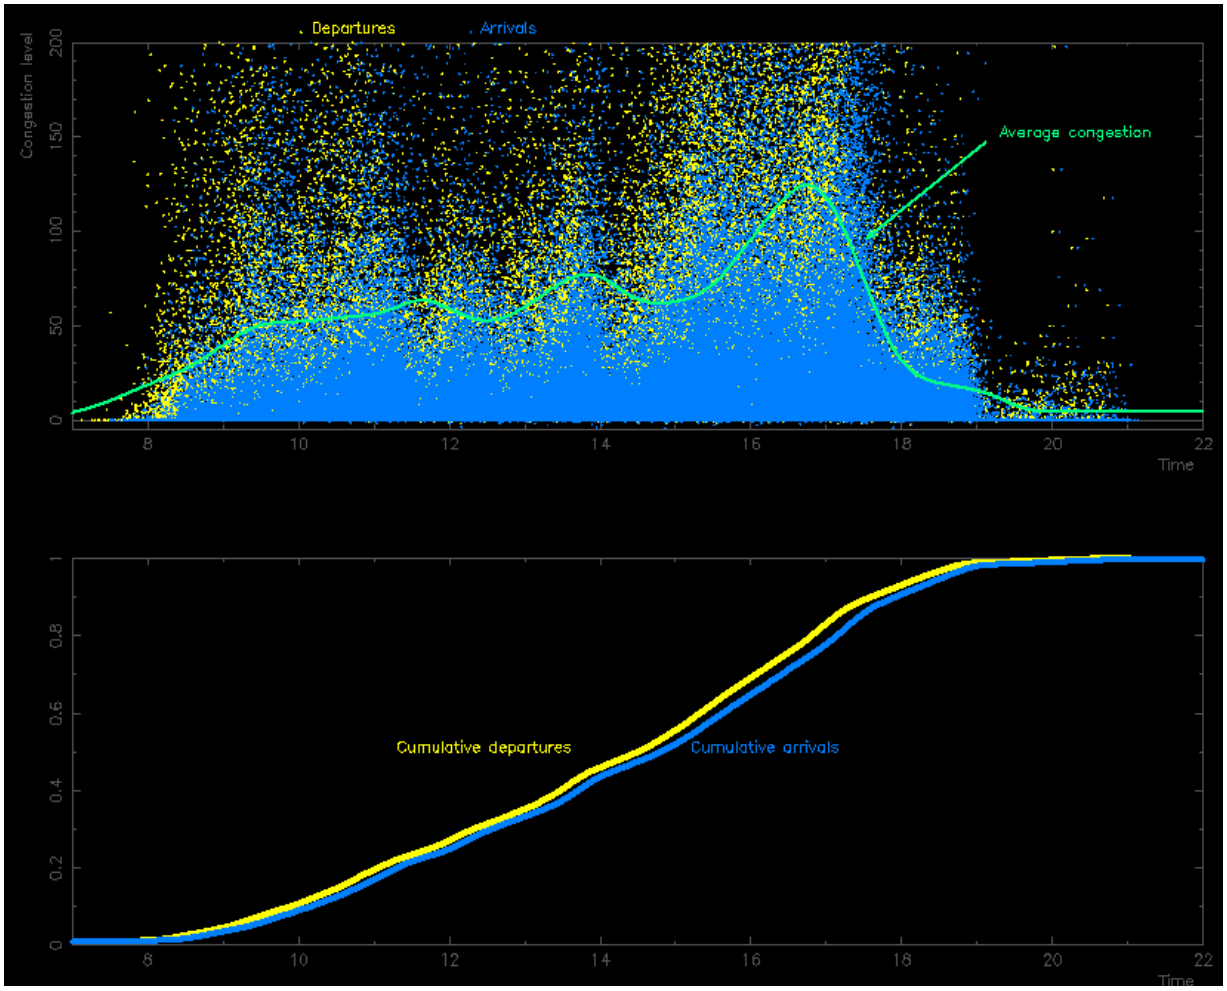

Source: Authors' design.

5. Consumers surplus and environmental measures

After briefly discussing how we measure consumer surplus, we discuss how fuel consumption and pollution can be included in the simulation framework.

5.1. Economic Analysis: Consumer Surplus

Transport simulation is valuable for examining mobility patterns and flows within the study area, but its primary utility lies in evaluating policy impacts. When economically comparing alternative scenarios, impacts are expressed in monetary terms via changes in money-metric consumer surplus**.** In METROPOLIS2 this is defined by the log-sum (inclusive value) of the nested logit (outer: mode; inner: departure time.

Let V_{kj} be systematic utility for traveler k, $\beta_{cost,k} < 0$ the cost coefficient, and $\lambda_g \in (0\ ,1)$ the dissimilarity parameter for nest g. Then:

$$IV_k = ln\left[\sum_g \left(\sum_{j \in g} exp(\lambda_g V_{kj}\right)^{1/\lambda_g}\right],\ CS_k = \frac{IV_k}{|\beta_{cost,k}|},\quad \Delta CS_k = \frac{{IV_k}^{(1)} - {IV_k}^{(0)}}{|\beta_{cost,k}|}.$$

Aggregating ΔCS_k over agents yields the total surplus change, which corresponds to the area under the demand curve (Ben-Akiva and Lerman 1985; and Train 2009).

Figure 13 provides an illustrative example using the calibrated model. Panel (a) presents the evolution of agents' utility values across simulation iterations. The upward trend demonstrates that, as the model converges toward equilibrium, users incrementally improve their utility by adapting routes and departure times in response to congestion patterns.

Panel (b) illustrates the distribution of consumer surplus sorted in ascending order across all agents. The curve reveals a gradual, steady increase for the majority, except for sharp fluctuations observed in the extreme lower and upper tails, representing less than 10% of agents. These outliers experience, respectively, exceptionally low and high congestion relative to most travelers.

This representation of consumer surplus will become crucial when comparing distinct policy scenarios. By generating surplus curves for each case, the welfare implications for individual agents can be quantified . Aggregating surplus across specific groups—such as those defined by location or departure time—allows identification of winners and losers from each proposed policy. Thus, consumer surplus, calculated as the integral beneath the demand curve, serves as a robust indicator of overall mobility-related well-being.

Figure 13. Overview of the consumer surplus of the calibrated model.

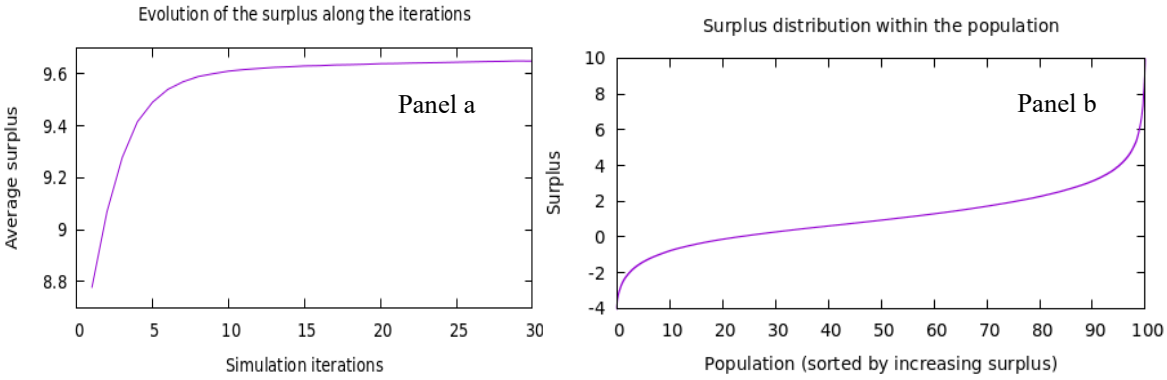

5.2. Fuel consumption and emissions as a function of travel speed

Although a comprehensive database on the precise distribution of vehicle attributes in Riyadh is currently unavailable, this methodology paper draws on existing market analyses, regional automotive trends, and publicly accessible data sources, such as Tomtom Riyadh[2]. Riyadh's vehicle fleet is notably diverse, comprising sedans, SUVs, pickup trucks, luxury vehicles, and an emerging segment of hybrid and fully electric vehicles (EVs). This version has no automobile ownership module, i.e., the fleet composition is exogenous.

Actual fuel consumption and emissions vary considerably due to driving behaviors (e.g., frequent idling, aggressive acceleration), traffic conditions, and maintenance practices. Given the lack of granular fleet-composition data, our current analysis adopts simplified

[2] https://www.tomtom.com/traffic-index/riyadh-traffic/

assumptions: an average fuel consumption rate of 10 L/100 km and average CO_2 emissions of 220 grams per kilometer (g/km). Future transport surveys and detailed vehicle fleet statistics would enable more precise estimations, accounting explicitly for variations across vehicle types, ages, and technologies, rather than relying on aggregate average values.

Fuel consumption significantly depends on travel speed, acceleration patterns, gear selection, and environmental conditions. Figure 14 reports the energy consumption factor (ECF)—defined as tank-to-wheel propulsion energy per vehicle-kilometre computed by COPERT—E(v) in MJ/km as a function of speed; for gasoline, a quick conversion is L/100 km = EEE×100/35**.** Energy use is U-shaped: it is elevated at low speeds (inefficient operation/stop-and-go) and rises again at high speeds due to increasing aerodynamic and rolling-resistance power demands.

Greenhouse gas (GHG) and particulate matter (PM) emissions generally correlate directly with fuel consumption, reflecting combustion efficiency. Each liter of gasoline combusted generates approximately 2.3 to 2.5 kg of CO_2, whereas diesel engines emit slightly more (2.6 to 2.8 kg per liter) due to higher carbon content.

Emissions (CO_2 and PM in g/km) are derived from the ECF. Like energy use, emissions typically reach a minimum in the "green-speed" range (~50–80 km/h), although the exact minimum depends on powertrain and driving cycle; outside this range, both energy use and emissions rise.

Figure 14. *Speed- energy consumption curve for medium cars. A* clear U-shape with a minimum at moderate speeds. Source: COPERT technical documentation.

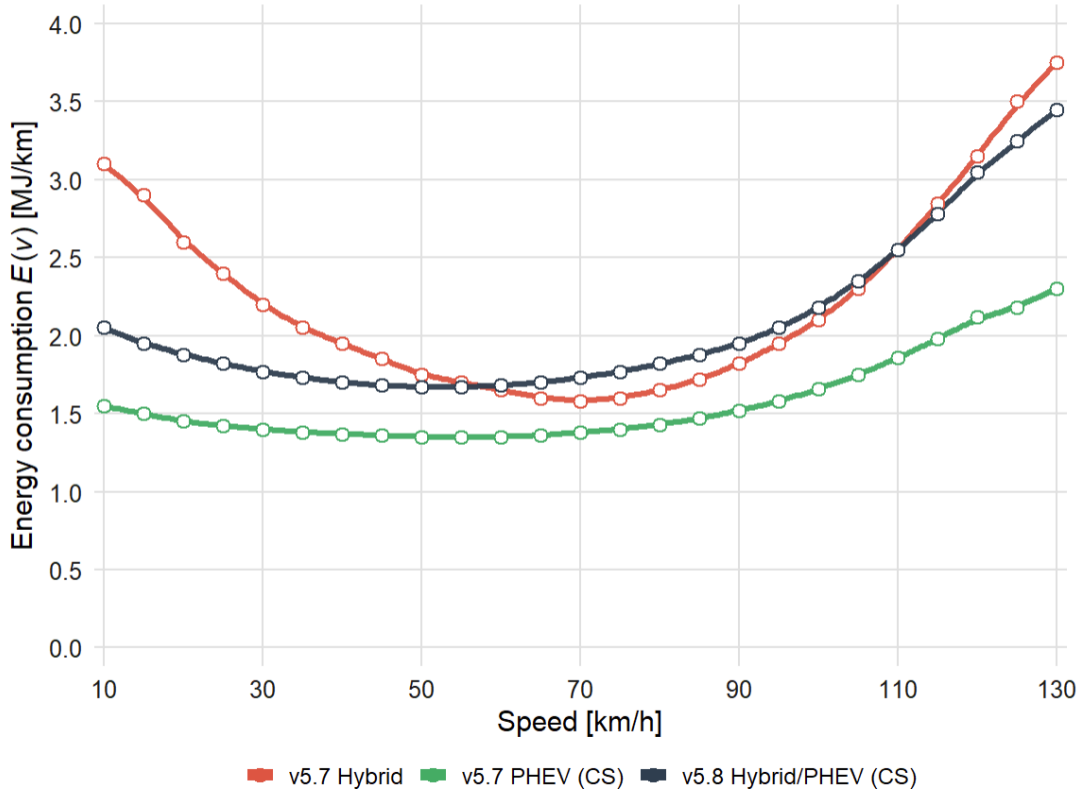

Note: The y-axis reports tank-to-wheel energy per kilometre (MJ/km). Curves correspond to COPERT v5.7 and v5.8 factors for medium petrol hybrids and

PHEVs in charge-sustaining mode. Emissions in g/km (e.g., CO_2, PM) can be computed downstream from these energy factors. Conversion to gasoline use: L/100 km = E(MJ/km) × 100 / 35.

6. Conclusion

This study developed and applied a dynamic traffic simulation framework for Riyadh to assess environmental externalities, providing a foundation for understanding how urban mobility patterns contribute to energy consumption, emissions, and air pollution. The framework demonstrates the feasibility of integrating dynamic transport modeling with environmental impact assessment. The core dynamic assignment is calibrated to private-vehicle (unimodal) travel data, while public transport (metro and buses) is represented exogenously—through timetable/capacity and service scenarios—to evaluate their implications for fuel consumption, emissions, and pollution.

The results show that the model can estimate travel times, congestion levels, and associated emissions with reasonable accuracy, enabling a systematic evaluation of pollution sources and mitigation strategies.

A key methodological contribution of this paperis the analytical capability to incorporate environmental considerations into urban mobility planning, particularly in fast-growing cities such as Riyadh, where private vehicle dependency, rapid urbanization, and extreme climatic conditions intensify emissions and air quality concerns. The insights gained here demonstrate how simulation-based approaches could potentially inform strategies to reduce emissions, mitigate localized pollution hotspots, and support transitions toward more sustainable transport systems.

While this work has established a calibrated methodological baseline, refinements are possible to increase accuracy and policy relevance. This includes incorporating real-world fleet composition and fuel consumption data, allowing for more precise quantification of local air pollutants such as nitrogen oxides (NO_x), particulate matter ($PM_{2.5}$), and carbon monoxide (CO), alongside greenhouse gas emissions. Further, refining the scheduling model to incorporate actual and desired arrival times at workplaces will improve assessments of peak-hour pollution concentrations and their potential health impacts. Another critical expansion involves examining the effects of electric vehicle adoption on urban air quality, using realistic penetration scenarios and electricity grid carbon intensity estimates.

Each of these developments presents technical and methodological challenges. Accurately modeling emissions requires integrating heterogeneous data sources, including on-road vehicle emissions inventories, meteorological conditions, and traffic flow measurements. Additionally, assessing the spatial distribution of pollutants necessitates incorporating atmospheric dispersion models, which will allow for a more granular evaluation of localized air quality effects. Electric vehicle adoption scenarios must also consider life-cycle emissions, grid capacity constraints, and potential behavioral changes in vehicle use, ensuring a realistic assessment of their environmental benefits.

Despite these challenges, the potential of this model extends beyond Riyadh. The framework can be applied to other cities facing similar environmental pressures, particularly those with

Public

high vehicular emissions, urban heat island effects, and growing energy demand for cooling systems. By integrating land-use data, socioeconomic indicators, and climate variables, the model can evolve into a comprehensive tool for assessing the broader ecological impacts of urban transport policies.

As Riyadh and other cities confront the escalating consequences of transport-related emissions and air pollution, dynamic modeling will play a crucial role in guiding environmental policy decisions. This initial version of the model lays the foundation for a more integrated approach to emission reduction, helping to shape a cleaner, healthier, and more sustainable urban future. Future developments will build on this work by incorporating multi-source environmental data, advanced air quality modeling, and climate impact assessments to enhance the model's predictive capabilities and policy relevance.

Finally, it is important to note that the optimization problem described in the calibration process (Section 5) has a relatively high dimension, particularly given the cost of evaluating the objective function. While the absence of detailed daily schedules (household and mobility surveys) constrains the number of parameters that can be fully identified in this first version—such as departure and arrival times—the approach remains robust for establishing a calibrated baseline model. The OD filtering method necessarily leaves a number of parameters to be inferred, which makes smooth optimization more challenging. Nevertheless, the resulting calibrated model is sufficiently rich to capture key dynamics and to generate meaningful insights on congestion, travel times, and emissions. With more comprehensive datasets, the calibration process can be extended to incorporate a wider set of behavioral parameters and achieve closer convergence with observed values, further strengthening the model's accuracy and policy relevance rather than replacing its current utility.

Acknowledgments: We would like to thank Lucas Javaudin for several insightful discussions and valuable advice on METROPOLIS.

Public

References

Al-Mosaind, M. (1998)

Freeway traffic congestion In Riyadh, Saudi Arabia: attitudes and policy implications, *Journal of Transport Geography Vol. 6 No. 4. pp, 263-272.*

Anderson, S. P., de Palma, A. & Thisse, J.-F. (1992),

Discrete choice theory of product differentiation, MIT press.

Batz, G. V., Geisberger, R., Sanders, P. & Vetter, C. (2013),

`Minimum time-dependent travel times with contraction hierarchies', *Journal of Experimental Algorithmics (JEA)* **18**, 1–1.

Ben-Akiva, M. and S. Lerman (1985),

`Discrete choice analysis: Theory and application to travel demand', *MIT Press google schola* **2**, 575–589.

Creutzig, Felix, Giovanni Baiocchi, R. Bierkandt, P. P. Pichler, and K. C. Seto. (2015),

"Global Typology of Urban Energy Use and Potentials for an Urbanization Mitigation Wedge." *Proceedings of the National Academy of Sciences* 112 (20): 6283–88. https://doi.org/10.1073/pnas.1315545112.

Diriyah Gate Development Authority. (2025),

"Diriyah Gate Development Authority. Accessed June 19, 2025. https://www.dgda.gov.sa/en/"

De Lara, M., De Palma, A., Kilani, M., & Piperno, S. (2013),

Congestion pricing and long-term urban form: Application to Paris region. *Regional Science and Urban Economics*, *43*(2), 282-295.

de Palma, A. & Javaudin, L. (2024),

Etude des extensions de la ligne 18 document confidentiel préparé pour la société des grands projets, Technical report, Sociťé des Grands Projets.

de Palma, A., Javaudin, L., Stokkink, P. & Tarpin-Pitre, L. (2024),

`Ride-sharing with inflexible drivers in the Paris metropolitan area', *Transportation* **51**(3), 963–986.

de Palma, A., Kilani, M. & Lindsey, R. (2005),

`Congestion pricing on road network: A study using the dynamic equilibrium simulator metropolis', *Transportation Research A* **39**(7-9), 588–611.

de Palma, A. & Lindsey, R. (2001),

`Optimal timetables for public transportation', *Transportation Research Part B: Methodological* **35**(8), 789–813.

de Palma, A., Lindsey, R. & Monchambert, G. (2017),

`The economics of crowding in rail transit', *Journal of Urban Economics* **101**, 106–122.

de Palma, A. & Marchal, F. (2002)

Real cases applications of the fully dynamic metropolis tool-box: an advocacy for large-scale

Public

mesoscopic transportation systems', *Networks and spatial economics* **2**, 347–369.

de Palma, A., Marchal, F. & Nesterov, Y. (1997),

`Metropolis: Modular system for dynamic traffic simulation', *Transportation Research Record* **1607**(1), 178–184.

de Palma, A., Stokkink, P. & Geroliminis, N. (2022),

`Influence of dynamic congestion with scheduling preferences on carpooling matching with heterogeneous users', *Transportation Research Part B: Methodological* **155**, 479–498.

Friedrich, R. & Quinet, E. (2011),

External costs of transport in europe, *in* `A handbook of transport economics', Edward Elgar Publishing.

Geisberger, R. & Sanders, P. (2010),

Engineering time-dependent many-to-many shortest paths computation, *in* `10th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS'10)', Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Intergovernmental Panel on Climate Change (IPCC). (2021),

"Climate Change 2021: The Physical Science Basis." Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press. https://www.ipcc.ch/ar6/wg1/.

Ikrami, Abdullah. (2014),

"18 Million Cars in Saudi Arabia… 356 Billion Riyals Imported in Six Years." Al-Eqtisadiah, October 22, 2014.

Javaudin, L. & A. de Palma (2024),

METROPOLIS2: Bridging Theory and Simulation in Agent-Based Transport Modeling. THEMA Working Paper No. 2024-03, CY Cergy Paris Université.

Kilani, M., de Palma, A. & Proost, S. (2017),

`Are users better-off with new transit lines?', *Transportation Research Part A: Policy and Practice* **103**, 95–105.

King Abdullah Financial District. (2025)

"King Abdullah Financial District. Accessed June 19, 2025. https://www.kafd.sa/en/"

Le Frioux, R., de Palma, A. & Blond, N. (2024),

Assessing the economic costs of road traffic-related air pollution in la reunion, Technical report, THEMA, Cergy Paris University.

Lopez-Aparicio, Susana, Henrik Grythe, Arkadiusz Drabicki, Konrad Chwastek, Kamila Tobola, Lidia Górska-Niemas, Urszula Kierpiec et al. (2025),

"Environmental sustainability of urban expansion: Implications for transport emissions, air pollution, and city growth." Environment International: 109310.

Luce, R., Bush, R. R. & Galanter, E. E. (1963),

`Handbook of mathematical psychology: I.'.

Monchambert, G. & de Palma, A. (2014),

`Public transport reliability and commuter strategy', *Journal of Urban Economics* **81**, 14–29.

Mulley, C., Nelson, J. & Ison, S. (2021),

Public

The Routledge handbook of public transport, Taylor & Francis.

Riyadh Region Municipality. (2025),

"The group of sub-municipalities within the boundaries of Riyadh City". Dataset. Accessed June 19, 2025

Ricardo Energy & Environment. (2025, March).

Speed Emission and Energy Consumption Curves for Ultra-Low Emission Vehicles (Issue 06)

Royal Commission for Riyadh City (RCRC). (2025a),

`Project and Programs, *Royal Commission for Riyadh City*, Accessed June 19, 2025. https://www.rcrc.gov.sa/en/projects/.

Royal Commission for Riyadh City (RCRC). (2025b),

'Number of Population of The City of Riyadh by Neighborhoods, Gender, and Nationality 1425H', Accessed March 19, 2025. https://www.rcrc.gov.sa/en/publications/riyadh-population-categorized-by-district-gender-nationality/

Saifuzzaman, M., Engelson, L., Kristoffersson, I. & de Palma, A. (2016),

`Stockholm congestion charging: an assessment with metropolis and silvester', *Transportation Planning and Technology* **39**(7), 653–674.

Small, K. A. (1982) The scheduling of consumer activities: work trips. *The American Economic Review*, *72*(3), 467-479.

Train, K. E. (2009),

Discrete choice methods with simulation, Cambridge university press.

Vickrey, W. (1969),

`Congestion theory and transport investment', *The American Economic Review* **59**(2), 251–260.

Vickrey, W. S. (1963),

`Pricing in urban and suburban transport', *The American Economic Review* **53**(2), 452–465.

Vision 2030. (2025),

`Vision 20230 Projects', *Saudi Arabia's Vision 2030*, Accessed March 19, 2025. https://www.vision2030.gov.sa/en/explore/projects

Vosough, S., de Palma, A. & Lindsey, R. (2022),

`Pricing vehicle emissions and congestion externalities using a dynamic traffic network simulator', *Transportation Research Part A: Policy and Practice* **161**, 1–24.

Yan, D., Zhao, Z., & Ng, W. (2012),

`Monochromatic and bichromatic reverse nearest neighbor queries on land surfaces. In *Proceedings of the 21st ACM international conference on Information and knowledge management* (pp. 942-951).

Yaseen, Lama, Nourah Al-Hosain, Ibrahem Shatnawi, and Abdelrahman Mohsen. (2024),

"Impact of Urban Traffic on Fuel Consumption Leveraging Iot Data: Case Study of Riyadh City." Available at SSRN 5018294.

Wolf, D. D., Diop, N. & Kilani, M. (2022),

`Environmental impacts of enlarging the market share of electric vehicles', *Environ Econ*

Public

Policy Stud .

Zeaara. (2022),

"Riyadh Car Stats." March 17, 2022. https://www.zeaara.com/blog/riyadhcarstats.html.

Public

Appendix

Appendix 1: Structure of the simulation work frame

The transportation simulation for Riyadh city is structured around a well-organized set of files that collectively define the operational framework, parameters, and dynamics of the system. The main configuration file resides in the root directory and serves as the central hub, containing essential parameters such as learning rate, number of iterations, logit model parameters, utility functions, and other critical settings necessary for the simulation to run effectively.

Within the net directory lies the network description, encapsulated in a CSV file where each row represents an edge with detailed attributes including start node, end node, capacity, type, and other relevant information. This file is fundamental as it defines the transportation infrastructure that agents will traverse during the simulation.

Additionally, the veh.csv file provides a comprehensive overview of the vehicle categories available within the simulation environment. Each row in this file corresponds to a distinct category of vehicles, which may influence how trips are executed based on the type and availability of these vehicles.

In the demand subdirectory, two key files are present: the agents.csv file, which lists all agents participating in the simulation along with their unique identifiers, and the trips.csv file, which enumerates every planned trip. Each row in the trips file is linked to a specific agent, detailing the origin and destination nodes for that particular trip. This setup ensures that each agent's movement is tracked precisely, reflecting real-world travel patterns within the city.

Overall, this structured file system allows for a detailed and dynamic simulation of transportation activities in Riyadh, enabling researchers to analyze various scenarios and optimize urban mobility solutions effectively.

The simulation results are stored in the 'output' directory, which contains various files summarizing both aggregate and detailed metrics. Aggregate result files provide high-level statistics such as average value of travel time, travel distance, departure time, and arrival time across all agents or trips. These files are typically generated to give a quick overview of system performance.

Detailed result files, on the other hand, contain granular information about individual agents' or trips' behaviors during the simulation. These files can be stored in either CSV or Parquet formats. While CSV is a standard format for data exchange, Parquet is often preferred for large-scale simulations due to its efficient compression and columnar storage, which help manage disk usage more effectively.

In Figure A1., we represent a schematic simplified description of directories related to simulations. It distinguishes between supply files (net), demand files (demand), both are inputs, and the output generated by METROPOLIS2. The general configuration file is usually located the root folder.

Public

Figure A1: the file structure of the project.

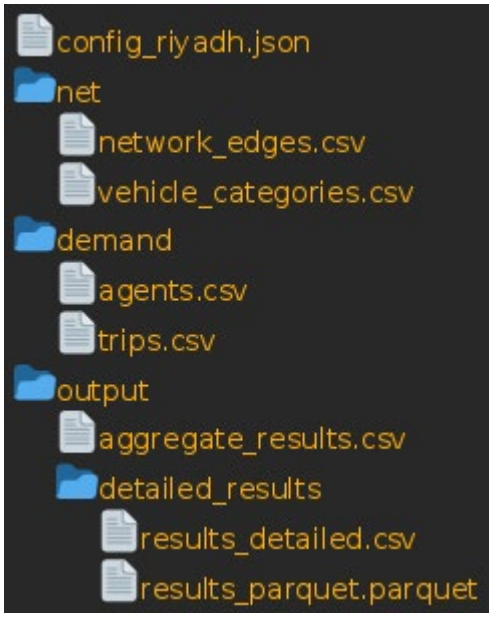

Source: Authors' design.

Appendix 2: Ray Casting Algorithm for Node-to-Zone Assignment

We used and programmed the *ray casting algorithm*, a standard method, for determining whether a given point lies inside or outside of a polygon. This algorithm works by drawing a semi-infinite line (often horizontal) from the point in question and counting how many times this line intersects with the edges of the polygon. If the number of intersections is odd, the point is inside the polygon; if even, the point is outside. An illustration is given in Figure A2. The main steps involve the following operations:

1. Choose a direction for the semi-infinite line: typically, a horizontal semi-line (to the right or left) is used because it simplifies calculations and avoids issues with vertical edges.
2. Iterate through each edge of the polygon: For each edge of the polygon, determine if it intersects with the chosen semi-infinite line.
3. Count intersections: for each intersection found, increment a counter.
4. Determine point location: if the total number of intersections is odd, the point is inside the polygon; If even, the point is outside the polygon.

The motivation for the algorithm can be seen from Figure 6. Particular cases (edge cases) can occur when the point lies exactly on an edge or a vertex. This can lead to ambiguous results because the point may lie on the boundary between inside and outside. Special handling is required for such cases.

Figure A2: The ray casting algorithm. Notice that some specific situations, with points on the edges, may complexify the problem.

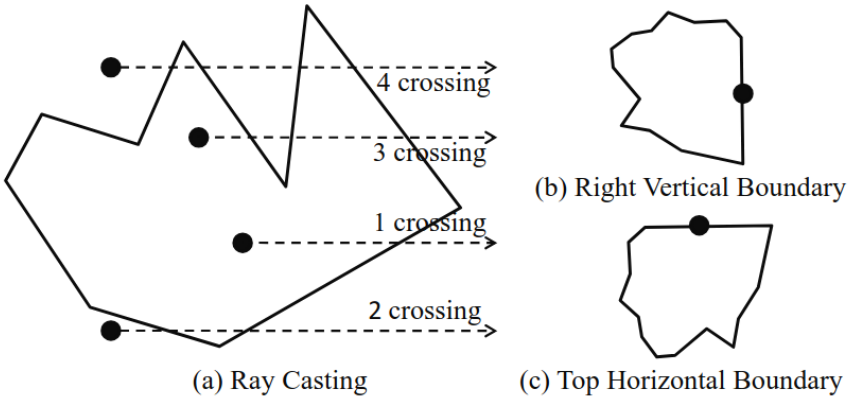

Source: Yan *et al. (2012)*.

The time complexity of the ray casting algorithm is *O(n)*, where *n* is the number of edges (or vertices) in the polygon. This makes it efficient for determining point-in-polygon relationships, especially for polygons with a large number of edges. It is widely applied in fields such as computer graphics, GIS applications, and computational geometry.

In this study, we implemented the ray casting algorithm to associate network nodes with their corresponding zones in the Riyadh Metropolitan Area. Specifically, over 130,000 nodes were spatially assigned to 162 zones**,** as defined by the provided OD matrix.