energies

MDPI

Article

Graph Neural Networks and Bi-Level Optimization for Equitable Electric Vehicle Charging Infrastructure Planning

Javier Alexander Guerrero Silva [1], Jorge Ivan Romero Gelvez [1,*] and Sebastian Zapata [2]

[1] Faculty of Natural Sciences and Engineering, Universidad Jorge Tadeo Lozano, Bogotá 110311, Colombia; javiera.guerreros@utadeo.edu.co
[2] Escuela de Ingeniería de Antioquia, Universidad EIA, Envigado 055422, Colombia; sebastian.zapata23@eia.edu.co
* Correspondence: jorgei.romerog@utadeo.edu.co

Abstract

Equity-aware electric vehicle (EV) charging planning remains difficult in data-constrained cities. In this work, an integrated framework was developed by combining spatiotemporal graph neural networks (ST-GNNs), EVI-Pro Lite demand estimation, and lexicographic bi-level optimization, and was applied to Bogotá, Colombia (8.3 million inhabitants). Household travel survey data (12,500 households across 142 zones) were used to estimate zone-level priority scores and venue-specific temporal weights. EVI-Pro Lite simulations projected a 2025 requirement of 10,870 charging ports (7352 residential, 2739 workplace, and 779 public). In the allocation stage, Level 1 preserved priority-proportional targets, while Level 2 minimized inter-zonal inequality in Hansen accessibility subject to near-optimal Level-1 compliance. The final allocation retained strong priority alignment in installed ports (Spearman $\rho = 0.799$, $p < 10^{-31}$), while the priority–accessibility association was lower (Spearman $\rho = 0.320$, $p = 1.04 \times 10^{-4}$), consistent with second-stage equity redistribution. Equity outcomes also improved (Hansen Gini = 0.433; bottom-50% Lorenz share = 0.204). The mean Hansen accessibility reached 296.630 (standard deviation 248.099; minimum 1.126). These findings indicate that reproducible, equity-oriented EV infrastructure plans can be produced in cities where revealed charging microdata are limited.

Keywords: electric vehicle charging; infrastructure equity; graph neural networks; bi-level optimization; urban accessibility; Latin American cities; sustainable transport; facility location

1. Introduction and Literature Review

Electric vehicles are rapidly moving from niche technology to mass-market adoption. According to the IEA, global electric car sales surpassed 10 million in 2022, accounting for 14% of new car sales worldwide, and were projected to rise to 14 million in 2023, or about 18% of the global market [1]. That growth was confirmed in the following edition of the report, which showed that sales in 2023 had indeed reached nearly 14 million and that the global electric car stock had expanded to around 40 million vehicles [2]. The upward trend continued in 2024: sales exceeded 17 million, the market share of electric cars rose above 20% for the first time, and the global fleet approached 58 million vehicles. The IEA further expects sales in 2025 to surpass 20 million, representing more than one-quarter of all cars sold worldwide [3].

Charging infrastructure has expanded in parallel, but not always at the same pace or in the same places. Public charging points increased from 2.7 million worldwide at the

Check for updates

Academic Editor: Roberto Finesso

Received: 11 March 2026
Revised: 4 April 2026
Accepted: 10 April 2026
Published: 20 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

end of 2022 to 3.9 million at the end of 2023, and more than 1.3 million additional public chargers were installed in 2024 [2,3]. Even so, the spatial distribution of this infrastructure remains uneven. Accessibility has long been understood as an equity issue in transport planning [4]. The early concentration of chargers in higher-income areas can be partially explained by earlier EV adoption in those areas, but persistent concentration still creates structural access gaps for lower-income communities [5,6]. Recent evidence also highlights disparities in charger reliability—not only in charger counts—reinforcing the need to treat equity as a primary planning objective rather than an afterthought [7].

This pattern is particularly acute in Latin American cities, where institutional constraints, limited budgets, and incomplete spatial data complicate evidence-based planning [8]. Bogotá, Colombia, exemplifies these challenges despite an ambitious policy framework for sustainable mobility. The city's Sustainable and Safe Mobility Plan (Plan de Movilidad Sostenible y Segura), issued by the District Secretariat of Mobility (Spanish: Secretaría Distrital de Movilidad, SDM), sets out a long-term vision to 2035 centered on decarbonized transport, multimodal accessibility, safer travel, and reducing territorial, social, and gender inequalities [9].

Recent work on EV charging planning has moved from single-level facility-location models toward bi-level formulations that better capture the interaction between infrastructure decisions and user response [10]. In these formulations, planners determine locations/capacities at the upper level, while travelers respond through charging and routing behavior at the lower level. This structure is valuable, but two limitations remain in much of the literature: First, many studies still rely on demand inputs that are either highly aggregated or weakly connected to observed daily mobility in data-scarce contexts. Second, equity is often evaluated after optimization instead of being embedded directly in the objective structure.

Recent contributions have improved behavioral realism (e.g., waiting effects and charging-choice equilibrium) [11,12], and others have integrated transport and power-system interactions or modeled spatial–temporal charging-load uncertainty within bi-level planning frameworks [13–16]. However, evidence from equity-focused reviews and empirical analyses shows that access gaps can persist unless distributional criteria are explicitly protected in the optimization itself [17–19]. This gap motivates our design choice: combine mobility-derived prioritization with a second-stage equity objective that is solved as part of the model, not only reported ex post.

Optimal charging-infrastructure planning requires detailed information on driving patterns, parking locations, dwelling times, and charging behavior. In high-income countries with mature EV markets, these inputs are increasingly derived from large-scale GPS telemetry and connected-vehicle data [20]. Such data are rarely available in Latin American cities, where EV fleets are still limited, institutional data sharing is weak, and privacy constraints restrict access to vehicle-level traces. Under these conditions, household travel surveys remain the most informative source for reconstructing the spatial and temporal organization of mobility demand. Bogotá's 2023 Mobility Survey (Spanish: Encuesta de Movilidad, EM-2023) provides detailed trip records for approximately 12,500 households distributed across 142 Transport and Mobility Analysis Units (UTAMs), including trip origins and destinations, departure times, travel durations, and relevant socioeconomic characteristics, yet it does not contain direct observations of charging sessions, charger selection, or battery state-of-charge. Accordingly, observed charging demand is not predicted directly. Instead, EM-2023 is combined with an activity-based infrastructure estimation framework [21] and a spatiotemporal graph learning model to infer relative charging priority from mobility structure. In this way, behaviorally informed spatial allocation signals can be derived even in data-scarce environments where revealed charging data are not yet available.

Graph neural networks (GNNs) offer a solution to this data challenge. These architectures represent urban areas as graphs, with nodes corresponding to neighborhoods and edges representing spatial relationships, thereby enabling learning from network-structured data [22,23]. Graph convolutional networks aggregate information from spatial neighbors [23], Long Short-Term Memory (LSTM) layers capture temporal patterns [24], and attention mechanisms weight different hours differentially for different charging types [25,26]. The key insight is that household travel surveys contain sufficient signal to learn meaningful charging infrastructure priorities through appropriate analytical methods. An end-to-end methodology is developed to integrate spatiotemporal GNN outputs with equity-aware optimization. First, a GNN is constructed from Bogotá's household travel survey data to produce (i) a priority score for each of the 142 UTAM zones and (ii) venue-specific temporal weights derived from an attention mechanism. Second, EVI-Pro Lite simulations are calibrated for Bogotá's context to estimate aggregate charging-infrastructure requirements by venue type [21]. Third, a lexicographic bi-level optimization problem is solved: Level 1 allocates ports to follow the GNN-derived priority distribution, while Level 2 refines the solution to reduce inequality in Hansen accessibility across zones subject to near-optimal Level-1 compliance. The complete pipeline is fully reproducible and is specifically designed for data-constrained cities, where revealed charging-demand data are limited or unavailable.

Literature Review and Positioning

Research on electric vehicle (EV) charging infrastructure has evolved from technical characterization toward integrated planning frameworks. Early studies clarified the operational differences among Level 1, Level 2, and DC fast charging, as well as their implications for infrastructure design and grid interaction [27,28]. Empirical and review work further showed that charger type, fast-charging availability, and user preferences shape charging behavior and adoption [29,30]. Later syntheses placed these technologies within the broader transition to electric mobility and energy systems [31]. Building on this foundation, a large body of optimization literature addressed where charging stations should be located and how much capacity should be installed. Representative approaches include maximal coverage siting [32], flow-refueling formulations that account for route feasibility and range anxiety [33], and integrated location–sizing models that incorporate distribution network constraints [34]. Recent reviews show that the field now covers siting, sizing, upgrades, uncertainty, and operations, with increasing emphasis on hybrid frameworks that combine mathematical programming, GISs, and data-driven methods [35]. A related line of review work focused on charging-demand estimation also shows that infrastructure requirements are highly sensitive to assumptions about utilization, charging access, vehicle type, and local travel behavior, especially in underrepresented regions [36].

As a result, demand representation has become a central issue in EV infrastructure planning. Activity-based tools such as EVI-Pro Lite translate travel itineraries and vehicle attributes into charging events by location type and time of day, making them especially useful in contexts where large-scale telematics are unavailable [21]. More recent integrated frameworks link demand estimation directly to siting decisions and grid feasibility analysis, allowing forecasting, charger placement, and power system impacts to be evaluated within a single workflow [37,38]. In parallel, machine learning methods have expanded rapidly for charging-demand prediction. Recurrent models such as LSTM have been used for short-term station-level forecasting [39], while graph-based deep learning methods extend this approach by explicitly representing spatial dependence among stations or urban regions [40,41]. Complementary forecasting studies have also projected spatiotemporal adoption trends across heterogeneous user segments to support forward-looking

infrastructure planning [42]. A recent systematized review confirms both the rapid growth of this literature and its fragmentation in terms of data sources, targets, and evaluation settings [43]. For planning applications, this implies that predictive performance alone is not sufficient unless forecasts can be translated into transparent and spatially actionable allocation rules.

A second major strand of literature concerns transportation equity and accessibility. Accessibility evaluation emphasizes that infrastructure should be assessed not only by how much is provided but also by who can effectively benefit from it. Hansen's gravitational accessibility measure formalized the idea that access increases with nearby opportunities and decreases with distance [44], while later work refined accessibility analysis to account for land use, transport interactions, and alternative service metrics [45]. Empirical studies show that, without explicit equity objectives, transport investments can reproduce or intensify spatial inequalities [8]. In the EV domain, recent studies document disparities in charger density, reliability, and accessibility across income, race, and urban form gradients [7,46]. Systematic reviews likewise conclude that disadvantaged communities often face longer travel distances, lower station quality, and greater uncertainty in charger availability even where regional capacity appears sufficient [17]. This literature increasingly argues that equity should be embedded directly into siting and allocation models rather than treated only as an ex post diagnostic [18,19,47,48].

At the same time, graph neural networks (GNNs) have emerged as a powerful framework for modeling urban systems with explicit spatial structure [49]. GNNs learn from non-Euclidean data by updating node representations through neighborhood aggregation, making them well suited to road networks, transit systems, and zonal adjacency graphs [22,23]. Related graph-convolutional developments also addressed relational inference and matrix-completion settings, illustrating the flexibility of graph-based representation learning in sparse environments [50]. In transportation, the most influential line of work combines graph representations with temporal modeling. Diffusion Convolutional Recurrent Neural Networks (DCRNNs) coupled graph diffusion with recurrent sequence learning for traffic forecasting [51], Spatiotemporal Graph Convolutional Networks (STGCNs) introduced fully convolutional spatiotemporal blocks [52], and Graph WaveNet showed that adaptive dependency learning can capture latent spatial relations beyond fixed adjacency matrices [53]. Reviews of traffic forecasting and intelligent transportation systems consistently emphasize that GNNs are particularly effective when spatial dependence, temporal dynamics, and multi-source interactions must be modeled jointly [54–57]. More recently, GNNs have also been proposed as surrogates for strategic transport models, extending their relevance from short-term prediction to planning support [58].

These developments are increasingly relevant for EV applications. Spatiotemporal GNNs have been used to predict charging demand across urban regions [40]. Other recent work uses adaptive graph-recurrent architectures to learn latent station relationships directly from data [59], attention-based graph models to capture regional charging-load variation [60,61], and hybrid deep learning approaches to forecast charging demand from regional behavioral patterns [41]. Together, these studies suggest that graph-based learning is well suited to EV infrastructure planning because charging demand is inherently spatiotemporal and strongly shaped by network structure.

To position the present study more clearly, three complementary research streams are distinguished: The first stream focuses on demand forecasting and demand estimation, where the main objective is to predict charging loads, events, or aggregate infrastructure requirements from mobility and vehicle data. The second stream addresses facility siting and capacity allocation, where optimization models determine where and how much charging infrastructure should be deployed under technical and policy constraints. The third

stream emphasizes equity-aware planning, where accessibility and distributional justice are modeled explicitly as optimization objectives rather than only as ex post diagnostics.

Most existing contributions connect only two of these streams at a time. Forecasting studies often stop before spatial allocation decisions are made, while allocation studies frequently rely on exogenous or simplified demand inputs and then evaluate equity after optimization. The central contribution of this paper is to integrate all three components within one coherent workflow tailored to data-constrained cities: EVI-Pro Lite provides venue-level aggregate quotas [21], the ST-GNN learns spatial priority and temporal venue structure from EM-2023, and a lexicographic bi-level model preserves demand-related priorities while directly minimizing accessibility inequality. This positioning clarifies that the novelty is not a single model block in isolation but, rather, the explicit coupling between demand representation, allocation logic, and equity objective.

The remainder of this manuscript is organized as follows: Section 2 presents the data sources and methods. Section 3 reports the empirical findings and integrates the discussion. Section 4 closes with conclusions, practical implications, and concise limitations.

2. Materials and Methods

2.1. Study Area and Spatial Units

This study focuses on Bogotá, Colombia, a metropolitan area of approximately 1587 km^2 subdivided into 142 Transport and Mobility Analysis Units (Spanish: Unidades de Transporte y Análisis de Movilidad, UTAMs), which constitute the basic spatial units for analysis. The official UTAM 2023 polygonal layer provided by the District Secretariat of Mobility (Spanish: Secretaría Distrital de Movilidad, SDM) is used and is available through the EM-2023 zonification download page ("Shapefile UTAM") [62]. The dataset version used in this study corresponds to the EM-2023 release, accessed on 26 March 2026; availability is public through the SDM observatory portal, while long-term hosting and licensing conditions are governed by SDM terms. The file is loaded with `GeoPandas` (v0.14.4) [63] and then reprojected into two coordinate reference systems (CRSs): (i) WGS84 (EPSG:4326) for web-based visualization and map export, and (ii) MAGNA-SIRGAS/Colombia Bogotá (EPSG:3116) for metric calculations (areas, distances, and adjacency).

From the UTAM geometries in EPSG:3116, the area of each zone (km^2) and centroid coordinates are computed. A full pairwise distance matrix between UTAM centroids is then obtained using `scipy.spatial.distance_matrix` and converted to kilometers. This matrix $D = (d_{ij})$ is used in the accessibility function of the optimization model in Section 2.5, where d_{ij} denotes the network-agnostic straight-line distance between demand zone i and infrastructure zone j. Euclidean distance is used for two practical reasons: (i) a consistent citywide origin–destination travel-time matrix by venue and period is not publicly available at UTAM resolution, and (ii) Euclidean distances keep the kernel computation transparent and fully reproducible. This is therefore treated as a parsimonious approximation, and its limitations and extensions are explicitly discussed in Section 4.

Spatial adjacency is encoded using Queen contiguity: two zones are considered neighbors if they share at least one boundary point. This criterion is applied to the UTAM polygons to construct a binary adjacency matrix $\mathbf{A}$, which is subsequently normalized for graph convolutional processing (Section 2.4.1). Figure 1 shows the study area and the UTAM zoning used throughout the analysis.

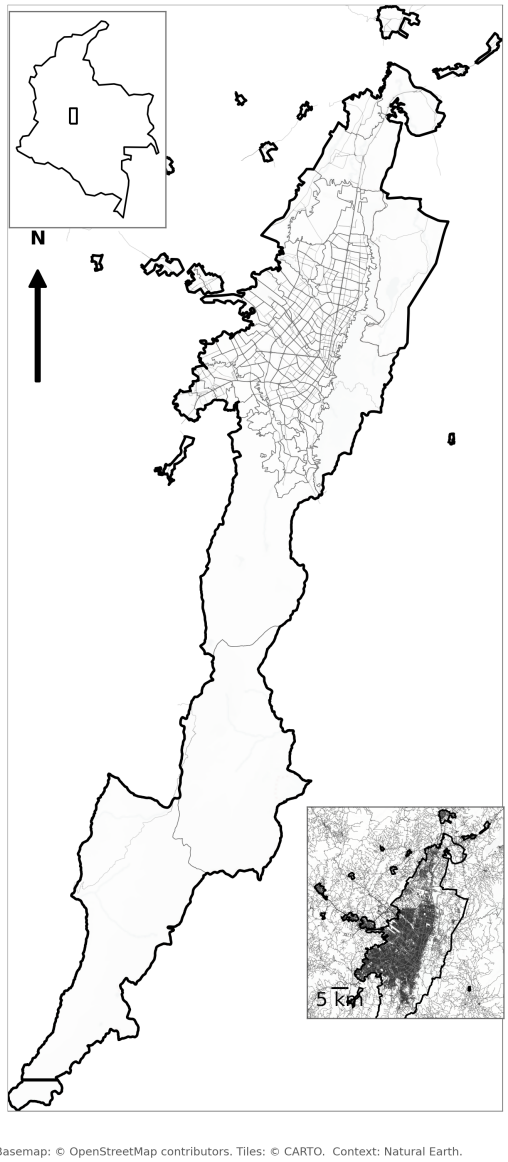


Figure 1. Bogotá study area and UTAM zoning. The 142 UTAM polygons form the nodes of the urban graph used in the graph neural network and optimization model.

2.2. Data Sources

Our methodology integrates three main data sources (Table 1):

- Household Travel Survey EM-2023: Bogotá's 2023 Mobility Survey (Spanish: Encuesta de Movilidad, EM-2023) provides detailed information on daily trips for approximately 12,500 households (around 45,000 individuals). Trips are geocoded to UTAMs and aggregated to obtain hourly origin and destination flows, which are subsequently normalized by population. EM-2023 also includes socioeconomic variables that are merged to the UTAM layer (population, income, employment, land-use indicators).
- ANDEMOS Vehicle Registry: Seven years of electric vehicle registration data from ANDEMOS (2017–2023) are used to characterize the temporal growth and spatial distribution of EV adoption in Bogotá. These data inform the EV fleet forecast and the calibration of EVI-Pro Lite (Section 2.3).
- GIS and Auxiliary Data: The UTAM 2023 shapefile (EM-2023 zonification, SDM observatory portal) provides the spatial boundaries and geometries required to build the urban graph, compute distances, and derive accessibility metrics [62]. Census,

cadastral, and tax records are used to enrich UTAM attributes with demographic and socioeconomic variables.

Table 1. Summary of key data sources used in the study.

Source	Description	Years	Spatial Unit	Main Variables Used
EM-2023 Household Travel Survey	Daily trips, purposes and modes; socioeconomic variables	2023	UTAM	Origin/destination flows by hour, population, income, employment, land use
ANDEMOS EV Registry	New EV registrations (BEV, PHEV) by municipality	2017–2023	Municipality/ locality	Annual EV counts, technology share, spatial distribution
UTAM GIS Layer	Official UTAM 2023 polygons (Bogotá)	2023	UTAM	Zone geometry, area, centroid, adjacency, distances

All datasets are merged by UTAM identifier to produce a consolidated GeoDataFrame that contains, for each zone, geometry, centroid coordinates, socioeconomics, and mobility indicators. These features constitute the input to the graph neural network and the optimization model.

2.3. EV Infrastructure Demand Estimation with EVI-Pro Lite

To derive citywide charging infrastructure requirements by venue type, the EVI-Pro Lite methodology developed by the National Renewable Energy Laboratory (Golden, CO, USA) is implemented [21]. EVI-Pro Lite is an activity-based simulation approach that translates daily travel itineraries into charging events for a synthetic EV fleet, accounting for vehicle characteristics, charging access, and time-of-day behaviors. The primary output used in this study is the required number of charging ports by venue (residential, workplace, public), which serves as an exogenous infrastructure "budget" that the subsequent allocation model must distribute spatially.

The simulation is calibrated to Bogotá's context using the following:

- An EV fleet forecast of 5200 vehicles by 2025, with a 60% BEV and 40% PHEV split;
- Daily travel patterns from EM-2023, including trip start and end times, trip durations and distances, and origin–destination UTAMs;
- Charging availability assumptions: 70% of vehicles with residential charging access, 40% with workplace charging access, and universal behavioral access to public charging;
- Representative battery capacities, efficiencies, and state-of-charge management strategies for the fleet;
- Seasonal temperature profiles for Bogotá (8–20 °C) to reflect climate impacts on energy consumption.

Because Bogotá does not yet publish a validated microdata series on private charger availability at the household and workplace levels, the 70%/40% access assumptions are treated as baseline planning priors rather than fixed empirical constants. These values are intentionally stress-tested in the sensitivity protocol (Section 2.3) to avoid overconfidence in a single access profile and to make the induced quota variability explicit [21,36].

To improve reproducibility, Table 2 reports the baseline EVI-Pro Lite parameterization used in this study.

These baseline values follow typical EVI-Pro Lite and charging infrastructure settings reported in the literature [21,27].

The simulation runs for a synthetic year, sampling daily itineraries with weekday/weekend distinctions and assigning charging events to residential, workplace, and public venues based on parking durations and access assumptions. Aggregating across simulated days yields the required number of charging ports by venue (and, when reported,

by powertrain). In this manuscript, these requirements are treated as hard quotas Q_v in the optimization stage: the bi-level allocation model must exactly satisfy the venue totals while deciding how to distribute ports across UTAM zones (Section 2.5). This separation of tasks—(i) estimating aggregate infrastructure needs through a behaviorally grounded simulation and (ii) allocating those needs spatially under explicit equity objectives—ensures that equity improvements arise from redistribution of access rather than from changing the overall infrastructure requirement.

Table 2. Baseline EVI-Pro Lite parameterization used for replication.

Parameter	BEV/General	PHEV
Usable battery capacity	50 kWh	14 kWh
Energy consumption rate	0.17 kWh/km	0.20 kWh/km (electric mode)
Residential charging power	7.2 kW (AC Level 2)	
Workplace charging power	7.2 kW (AC Level 2)	
Public charging power	50 kW (DC fast equivalent)	
Charging efficiency	90%	
Charging trigger state-of-charge	30%	
Target state-of-charge after session	90%	
Minimum reserve state-of-charge	15%	

Assumption Transparency and Sensitivity Protocol

Because EVI-Pro Lite outputs are assumption-dependent, a one-at-a-time sensitivity protocol is explicitly documented around the three inputs most directly linked to venue quotas: (i) total EV fleet size, (ii) BEV/PHEV composition, and (iii) access rates to residential and workplace charging. The protocol is defined as follows:

- Fleet-Size Scenarios: ±20% around the baseline forecast of 5200 vehicles;
- Powertrain-Mix Scenarios: BEV share varied from 50% to 70% (with PHEV share adjusted complementarily);
- Access Scenarios: Residential access varied from 60% to 80%, and workplace access from 30% to 50%.

For each scenario, the resulting venue-level quotas Q_v are passed unchanged to the bi-level optimization stage so that the effect of demand-side assumptions is not conflated with changes in allocation logic. As a first-order reference, if all other assumptions are held constant, the total quota scales approximately with fleet size: the baseline total of 10,870 ports implies about 8696 ports at −20% fleet and 13,044 ports at +20% fleet. This protocol improves transparency by making explicit which assumptions drive infrastructure totals before equity-oriented redistribution is applied.

Table 3 reports this first-order fleet-size sensitivity under a fixed venue-share approximation (residential 67.6%, workplace 25.2%, public 7.2%), using the baseline EVI-Pro composition. Figure 2 summarizes the EVI-Pro Lite calibration workflow and the venue-specific charging-port quotas used later as model constraints.

Table 3. First-order sensitivity of projected charging-port totals to EV fleet-size scenarios.

Scenario	Total Ports	Residential	Workplace	Public
−20% fleet (4160 EVs)	8696	5882	2191	623
Baseline (5200 EVs)	10,870	7352	2739	779
+20% fleet (6240 EVs)	13,044	8822	3287	935

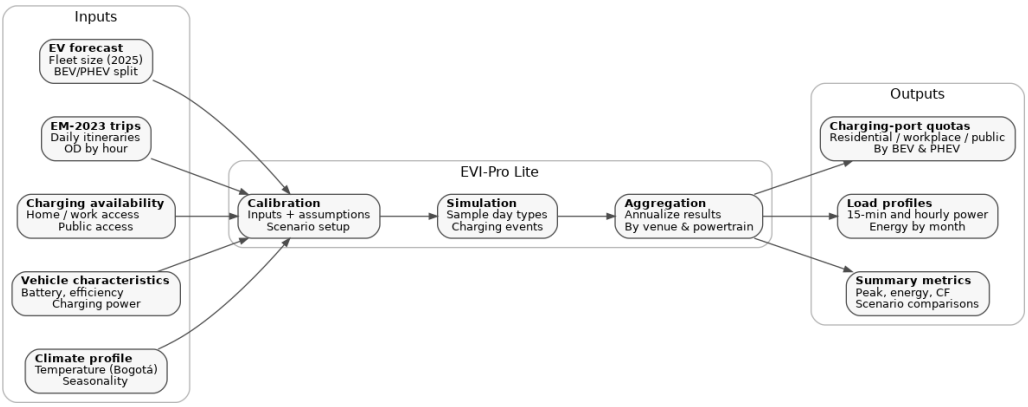


Figure 2. EVI-Pro Lite calibration workflow for Bogotá: input assumptions and venue-specific charging-port quotas used as bi-level constraints.

This sensitivity is not intended to replace a full multi-parameter uncertainty analysis. Instead, it provides a transparent quantitative envelope around the 10,870-port projection given current data constraints.

2.4. Graph Neural Network Model for Charging Priority

This study derives neighborhood-level charging priorities directly from the household travel survey by learning spatiotemporal patterns in UTAM-to-UTAM flows. A spatiotemporal graph neural network (ST-GNN) is implemented to combine graph convolutions for spatial dependence, a Long Short-Term Memory (LSTM) layer for temporal sequence modeling, and an additive attention mechanism for hour-level weighting. The model produces two outputs for each UTAM zone: (i) a scalar priority score, and (ii) a 24 h attention profile that is later aggregated into venue-specific weights used by the bi-level allocation model.

The proposed framework does not predict observed charging demand directly, since large-scale charging behavior data are not available for Bogotá. Instead, it learns a mobility-based proxy for relative charging priority across zones, using the spatial and temporal structure embedded in the origin–destination survey.

2.4.1. Graph Representation of Urban Structure

Bogotá's UTAM system is represented as a spatial graph:

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}),$$

where the node set $\mathcal{V}$ corresponds to UTAM zones and the edge set $\mathcal{E}$ captures spatial adjacency relationships between them. In the final implementation, the graph contains $|\mathcal{V}| = 142$ nodes, one for each UTAM zone in the study area.

Spatial edges are first constructed using Queen-type contiguity, so two zones are connected if their polygons touch or intersect. When contiguity is too sparse for stable message passing, a k-nearest-neighbors fallback on UTAM centroids ($k = 6$) is applied. This ensures a connected support for graph convolution while preserving neighborhood structure. Although the conceptual graph is undirected, implementation uses reciprocal directed edges in edge-index format, following standard GNN practice.

Following the standard graph convolution formulation, spatial propagation relies on the symmetrically normalized adjacency matrix with self-loops:

$$\tilde{\mathbf{A}} = \mathbf{D}^{-1/2}(\mathbf{A} + \mathbf{I})\mathbf{D}^{-1/2}, \quad (1)$$

where **A** is the adjacency matrix, **I** is the identity matrix, and **D** is the degree matrix. The normalization in Equation (1) prevents feature magnitudes from scaling with node degree and improves numerical stability during training [23].

2.4.2. Feature Construction from Survey Flows

Using the EM-2023 household travel survey, mobility activity is aggregated by UTAM origin, UTAM destination, and hour of day. For each zone i and hour $h \in \{0, \ldots, 23\}$, a compact feature vector is constructed

$$\mathbf{x}_{ih} \in \mathbb{R}^3,$$

defined as

$$\mathbf{x}_{ih} = \left[\text{orig}_{ih},\ \text{dest}_{ih},\ \text{total}_{ih}\right], \qquad \text{total}_{ih} = \text{orig}_{ih} + \text{dest}_{ih}, \tag{2}$$

where orig_{ih} denotes the number of trips originating in zone i at hour h, and dest_{ih} denotes the number of trips ending in zone i at hour h.

By stacking the 24 hourly feature vectors for each zone, the spatiotemporal input tensor is obtained

$$\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times 24 \times 3}, \tag{3}$$

which is used as the model input. Equations (2) and (3) define the full input representation used in training. Figure 3 visualizes the resulting UTAM graph, including the node set and the spatial adjacency relationships used in the ST-GNN.

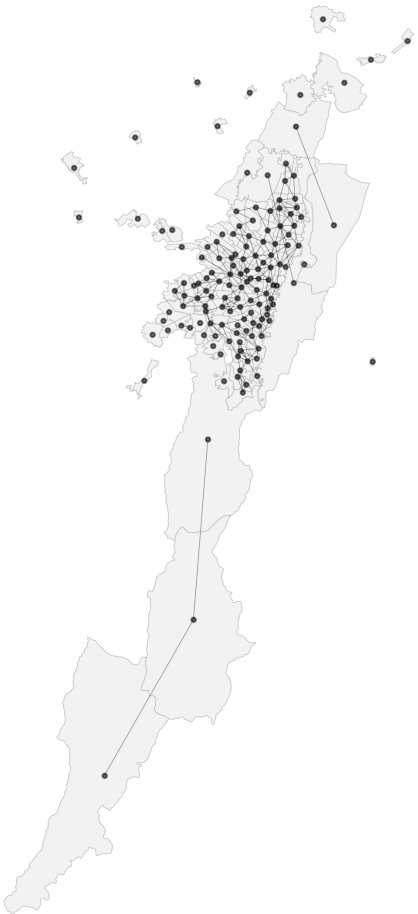


Figure 3. Graph representation of Bogotá's UTAM network. Points represent UTAM-zone nodes, and line segments represent adjacency edges under Queen contiguity, with a k-NN fallback on centroids when necessary. The map includes an approximate scale bar in kilometers.

Because Bogotá lacks observed charging demand data at sufficient spatial resolution, the model is trained using a mobility-derived proxy target. Specifically, for each zone i, a scalar intensity label is defined as the mean hourly total activity:

$$y_i = \frac{1}{24} \sum_{h=0}^{23} \text{total}_{ih}. \tag{4}$$

This target (Equation (4)) captures the average daily mobility intensity associated with each zone and provides a consistent supervision signal for learning spatial and temporal structure from the survey data. Accordingly, the model should be interpreted as a prioritization mechanism based on mobility patterns rather than a direct predictor of realized EV charging demand.

2.4.3. Spatiotemporal Graph Neural Network (ST-GNN) Architecture

The proposed ST-GNN framework, illustrated in Figure 4, adopts a hierarchical architecture designed to capture both spatial interdependencies among neighboring UTAM zones and temporal variation across the 24-h daily cycle. The workflow is organized into three sequential stages: spatial feature extraction, temporal sequence modeling, and attention-based score regression.

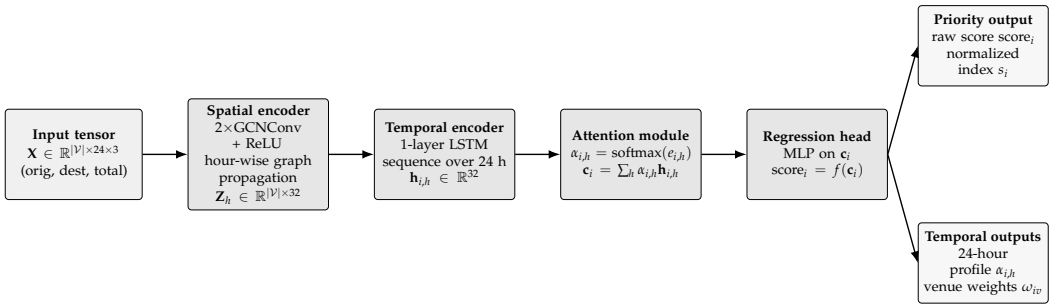


Figure 4. ST-GNN architecture overview. Spatial graph convolutions, temporal LSTM encoding, and attention-based regression generate UTAM priority and temporal-weight outputs.

I. Spatial Feature Extraction (GCN)

To account for the topological structure of the UTAM network, spatial message passing is applied independently at each hour $h \in \{0, \ldots, 23\}$. Let

$$\mathbf{X}_h \in \mathbb{R}^{|\mathcal{V}|\times 3}$$

denote the matrix of node features at hour h. For each hourly slice, the model applies a two-layer graph convolutional network (GCN) [23], implemented through `GCNConv` layers with ReLU activations:

$$\mathbf{Z}_h = \text{GCN}_2\big(\sigma\big(\text{GCN}_1(\mathbf{X}_h, \mathcal{G})\big), \mathcal{G}\big), \qquad \mathbf{Z}_h \in \mathbb{R}^{|\mathcal{V}|\times g}, \tag{5}$$

where $\sigma(\cdot)$ denotes the Rectified Linear Unit (ReLU), and $g = 32$ is the spatial hidden dimension. Equation (5) transforms the original mobility features into latent embeddings that encode neighborhood influence through the graph structure.

II. Temporal Sequence Modeling (LSTM)

The sequence of hourly spatial embeddings

$$\{\mathbf{Z}_0, \mathbf{Z}_1, \ldots, \mathbf{Z}_{23}\}$$

is then treated as a temporal signal for each zone. To capture inter-hour dependence and daily temporal structure, the model uses a one-layer Long Short-Term Memory (LSTM) network [24]. For each zone i, the LSTM processes the sequence of spatial representations

$$\{\mathbf{z}_{i,h}\}_{h=0}^{23}$$

and produces a corresponding sequence of hidden states

$$\mathbf{h}_{i,h} \in \mathbb{R}^t,$$

where $t = 32$ is the temporal hidden dimension. The full temporal output is therefore

$$\mathbf{H}_i = \left[\mathbf{h}_{i,0}, \mathbf{h}_{i,1}, \ldots, \mathbf{h}_{i,23}\right]. \tag{6}$$

III. Temporal Attention and Scalar Regression

Because not all hours contribute equally to charging priority, the model applies an additive attention mechanism [26] over the LSTM hidden states. This component learns hour-specific importance weights for each zone. The attention formulation is as follows:

$$\mathbf{a}_{i,h} = \tanh(\mathbf{W}\mathbf{h}_{i,h}), \tag{7}$$

$$e_{i,h} = \mathbf{u}^\top \mathbf{a}_{i,h}, \tag{8}$$

$$\alpha_{i,h} = \frac{\exp(e_{i,h})}{\sum_{h'=0}^{23} \exp(e_{i,h'})}, \tag{9}$$

where $\mathbf{W}$ and $\mathbf{u}$ are trainable parameters, $e_{i,h}$ is the unnormalized attention score, and $\alpha_{i,h}$ is the normalized importance assigned to hour h for zone i.

The attention-weighted context vector for each zone is computed as

$$\mathbf{c}_i = \sum_{h=0}^{23} \alpha_{i,h} \mathbf{h}_{i,h}. \tag{10}$$

This vector summarizes the zone's spatiotemporal mobility signature over the full day. Finally, the context vector is mapped to a scalar output through a small multilayer perceptron (MLP):

$$\text{score}_i = f(\mathbf{c}_i), \tag{11}$$

where $f(\cdot)$ denotes the final regression head. The resulting scalar score_i is interpreted as the raw charging-priority signal for zone i, while the 24 attention weights $\alpha_{i,h}$ provide an interpretable temporal importance profile. Equations (6)–(11) describe the temporal and attention components end-to-end.

To complement Figure 4, Table 4 summarizes the block-level architecture configuration used in implementation.

Table 4. ST-GNN block-level configuration.

Block	Configuration	Output Shape (per Zone)
Input tensor	24 hourly slices, 3 features/hour	24×3
Spatial encoder	$2 \times$ GCNConv + ReLU, hidden size $g = 32$	24×32
Temporal encoder	1-layer LSTM, hidden size $t = 32$	24×32
Attention module	Additive attention over 24 states	24 attention weights
Regression head	MLP on context vector $\mathbf{c}_i$	1 priority score

2.4.4. Priority and Venue-Weight Extraction

The final stage of the ST-GNN pipeline converts the model outputs into planning variables that can be used directly in the optimization model. Specifically, for each UTAM zone i, the network produces a scalar raw score score_i and a 24-h attention profile $\alpha_{i,h}$. These outputs are transformed into two quantities: (i) a normalized spatial priority index s_i, which captures the relative importance of each zone; and (ii) venue-specific weights ω_{iv}, which summarize the temporal structure of the learned mobility signal in a form suitable for infrastructure allocation.

I. Spatial Priority Normalization

To make scores comparable across the 142 UTAM zones, the raw output score_i is rescaled to the unit interval through min–max normalization:

$$s_i = \frac{\text{score}_i - \min_j \text{score}_j}{\max_j \text{score}_j - \min_j \text{score}_j + \epsilon}, \tag{12}$$

where $\epsilon > 0$ is a small constant introduced for numerical stability. The resulting index $s_i \in [0, 1]$ preserves the relative ordering of zones and is interpreted as the mobility-based charging priority used in the first level of the optimization model.

II. Temporal Aggregation into Venue Windows

The attention coefficients $\alpha_{i,h}$ indicate the relative contribution of each hour h to the final priority score of zone i. To connect these hour-level signals with the charging categories used in EVI-Pro Lite, the 24 h attention profile is aggregated into three operational windows associated with typical charging opportunities:

$$W_{i,\text{res}} = \sum_{h \in \{22,23,0,1,2,3,4,5\}} \alpha_{i,h}, \qquad \text{(overnight/residential)}, \tag{13}$$

$$W_{i,\text{work}} = \sum_{h=8}^{17} \alpha_{i,h}, \qquad \text{(daytime/workplace)}, \tag{14}$$

$$W_{i,\text{pub}} = \sum_{h=18}^{21} \alpha_{i,h}, \qquad \text{(evening/public)}. \tag{15}$$

These aggregated quantities provide an interpretable temporal signature for each zone and relate the learned attention structure to venue-specific charging contexts.

III. Venue-Weight Construction

Because the aggregated window scores are not constrained to sum to one across venue types, they are normalized row-wise to obtain a convex combination:

$$\omega_{iv} = \frac{W_{iv}}{\sum_{v' \in \mathcal{T}} W_{iv'}}, \qquad \mathcal{T} = \{\text{res}, \text{work}, \text{pub}\}. \tag{16}$$

The weight ω_{iv} therefore represents the relative importance of venue type v in zone i, conditional on the temporal mobility pattern learned by the ST-GNN.

Together, Equations (12)–(16) define how raw ST-GNN outputs are transformed into planning variables for optimization.

This transformation serves two purposes: First, it preserves inter-zonal heterogeneity by allowing different parts of the city to exhibit different venue profiles rather than imposing a uniform citywide split. Second, it provides a direct behavioral input to the bi-level allocation model: the normalized priority score s_i captures where infrastructure is most needed, whereas the venue weights ω_{iv} capture which type of charging opportunity is most consistent with the temporal structure of mobility in each zone.

2.4.5. Training Procedure

The ST-GNN is trained end-to-end on the full UTAM graph using the mobility-derived zone-level target y_i. The model is optimized with Adam, using a learning rate of 10^{-3} and weight decay of 10^{-4}. Training minimizes the mean absolute error (MAE) between the predicted priority score and the target value for each zone:

$$\mathcal{L} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} |\text{score}_i - y_i|. \tag{17}$$

The ST-GNN training objective is therefore fully specified by Equation (17).

The training procedure is run for 1500 epochs with a fixed random seed to ensure reproducibility. Because the model is estimated on the full UTAM graph and is not intended as a conventional out-of-sample forecasting model, its role in this study is to learn stable latent representations of spatiotemporal mobility intensity that can be translated into planning-oriented charging-priority indicators.

After convergence, the workflow exports two artifacts for subsequent analysis and optimization: (i) `priority_utam_scores.csv`, containing the raw score `priority_score` and the normalized index `priority_norm`; and (ii) `temporal_attention_by_utam.csv`, containing the 24-h attention profile for each UTAM zone. These outputs are then used to derive the venue weights and priority targets required by the bi-level charging infrastructure allocation model.

2.4.6. Notation Summary

Table 5 defines the main symbols and parameters used in the optimization formulation.

Table 5. Main notation and parameter definitions.

Symbol	Definition
I	Set of UTAM zones ($\|I\| = 142$)
V	Set of venue types $\{\text{res}, \text{work}, \text{pub}\}$
n_{iv}	Integer number of ports allocated to zone i, venue v
Q_v	Exogenous quota of ports for venue v (from EVI-Pro Lite)
s_i	Normalized ST-GNN priority score for zone i
ω_{iv}	ST-GNN venue weight for zone i, venue v
π_{iv}	Priority-proportional allocation share
T_{iv}	Priority-proportional target ports ($T_{iv} = \pi_{iv} Q_v$)
d_{ij}	Euclidean centroid distance between zones i and j (km)
λ_v	Venue-specific distance-decay parameter (km)
α	Priority amplification parameter in Level 1
ε	Relaxation parameter in lexicographic coupling
$U^\star$	Optimal Level-1 deviation objective value

2.5. Bi-Level Lexicographic Optimization Model

Charging ports are allocated across Bogotá's UTAM zones through a lexicographic bi-level optimization model that combines demand fidelity and equity improvement. Let I denote the set of UTAM zones, with $|I| = 142$, and let V denote the set of venue types:

$$V = \{\text{res}, \text{work}, \text{pub}\}.$$

The decision variable is the integer number of charging ports installed in zone $i \in I$ for venue type $v \in V$:

$$n_{iv} \in \mathbb{Z}_+.$$

The model is designed to solve two planning objectives sequentially. The first level preserves the mobility-derived priority structure learned by the ST-GNN, while the second level improves the equity of accessibility across zones without substantially departing from the first-level solution.

Hard Venue-Specific Quotas:

The aggregate charging requirements estimated with EVI-Pro Lite are imposed as hard constraints. For each venue type v, the total number of allocated ports must equal the exogenous quota Q_v:

$$\sum_{i \in I} n_{iv} = Q_v \qquad \forall v \in V. \tag{18}$$

These constraints (Equation (18)) ensure that the optimization redistributes charging infrastructure spatially without altering the citywide infrastructure totals determined in the demand-estimation stage.

Level 1: Priority-Preserving Allocation:

The ST-GNN provides two inputs to the allocation model: a normalized zone-level priority score $s_i \in [0, 1]$, and a set of venue-specific temporal weights ω_{iv}, where $\sum_{v \in V} \omega_{iv} = 1$ for each zone i. Using these quantities, a priority-proportional allocation share is defined for each zone and venue:

$$\pi_{iv} = \frac{\max(10^{-6}, s_i)^{\alpha} \omega_{iv}}{\sum_{k \in I} \max(10^{-6}, s_k)^{\alpha} \omega_{kv}}, \qquad \forall i \in I,\ v \in V, \tag{19}$$

where α is a tuning parameter that controls the strength of priority amplification. The corresponding target number of ports is

$$T_{iv} = \pi_{iv} Q_v. \tag{20}$$

The first optimization level minimizes the total absolute deviation between the integer allocation n_{iv} and these target values:

$$\min \sum_{i \in I} \sum_{v \in V} |n_{iv} - T_{iv}|. \tag{21}$$

Equations (19) and (20) define the priority-proportional targets, and Equation (21) defines the Level-1 deviation metric.

To obtain a linear mixed-integer formulation, auxiliary non-negative variables $u_{iv} \geq 0$ are introduced such that

$$u_{iv} \geq n_{iv} - T_{iv}, \qquad u_{iv} \geq T_{iv} - n_{iv}, \qquad \forall i \in I,\ v \in V. \tag{22}$$

The Level-1 problem can then be written as

$$\min \sum_{i \in I} \sum_{v \in V} u_{iv}. \tag{23}$$

Equation (22) linearizes the absolute deviations, and Equation (23) is the solved MILP objective. Let $U^{\star}$ denote the optimal value of this first-level objective.

Level 2: Equity Improvement through Hansen Accessibility:

After obtaining the best priority-preserving allocation, the second level seeks to reduce inter-zonal inequality in accessibility. For each demand zone i, Hansen-type accessibility is defined as

$$A_i = \sum_{v \in V} \omega_{iv} \sum_{j \in I} \exp(-d_{ij}/\lambda_v)\, n_{jv}, \tag{24}$$

where d_{ij} is the straight-line distance in kilometers between the centroids of zones i and j, and λ_v is a venue-specific distance-decay parameter. Equation (24) follows the Hansen accessibility logic [44] and captures the fact that charging infrastructure in one zone may also benefit nearby zones through spatial spillovers, with the magnitude of that benefit decreasing with distance.

Equity is operationalized by minimizing the sum of pairwise absolute differences in accessibility:

$$\min \sum_{i<j} |A_i - A_j|. \tag{25}$$

This objective is proportional to the numerator of the Gini coefficient and, therefore, provides a direct measure of accessibility inequality across the urban system.

To linearize the objective, auxiliary variables $z_{ij} \geq 0$ are introduced for all $i < j$, subject to

$$z_{ij} \geq A_i - A_j, \qquad z_{ij} \geq A_j - A_i, \qquad \forall i < j. \tag{26}$$

The Level-2 objective then becomes

$$\min \sum_{i<j} z_{ij}. \tag{27}$$

Equation (25) defines the inequality criterion, while Equations (26) and (27) provide its linearized implementation.

Lexicographic Coupling:

The second level is solved subject to a near-optimality condition that preserves the priority structure obtained in Level 1. Specifically, the total deviation from the priority-proportional targets is constrained by

$$\sum_{i \in I} \sum_{v \in V} u_{iv} \leq (1 + \varepsilon) U^{\star}, \tag{28}$$

where $\varepsilon > 0$ is a small relaxation parameter. This condition allows the second level to trade a limited deterioration in priority compliance for a potentially larger improvement in accessibility equity.

In the implementation reported here, $\varepsilon = 0.05$ is set, which limits the second-stage deterioration in Level-1 fidelity to 5%. This value was selected as a policy-interpretable compromise: large enough to allow meaningful equity gains, but small enough to preserve the original priority ordering as the dominant signal. The coupling constraint is enforced through Equation (28).

The lexicographic structure therefore separates two policy-relevant aims: first, allocate infrastructure in accordance with the mobility-derived demand signal; second, redistribute accessibility more fairly across zones while remaining close to that initial allocation logic.

Implementation:

The model is implemented in `Pyomo` [64] and solved with open-source MILP solvers, using CBC as the main solver and GLPK as a fallback. Distance kernels are computed from the UTAM centroid matrix and venue-specific decay parameters:

$$\lambda_{\text{res}} = 3 \text{ km}, \qquad \lambda_{\text{work}} = 5 \text{ km}, \qquad \lambda_{\text{pub}} = 2 \text{ km}.$$

These values reflect different effective catchment ranges for residential, workplace, and public charging opportunities. Interpreted as exponential kernels, they correspond to half-decay distances of $\lambda_v \ln 2$: about 2.08 km (residential), 3.47 km (workplace), and 1.39 km (public). The larger workplace parameter captures longer routine commuting catchments, while the shorter public parameter represents more localized opportunistic charging behavior. The final output of the model is the integer allocation matrix (n_{iv}), which specifies the number of charging ports assigned to each UTAM zone and venue type. Figure 5 summarizes the resulting bi-level optimization structure, including inputs, decision levels, and outputs.

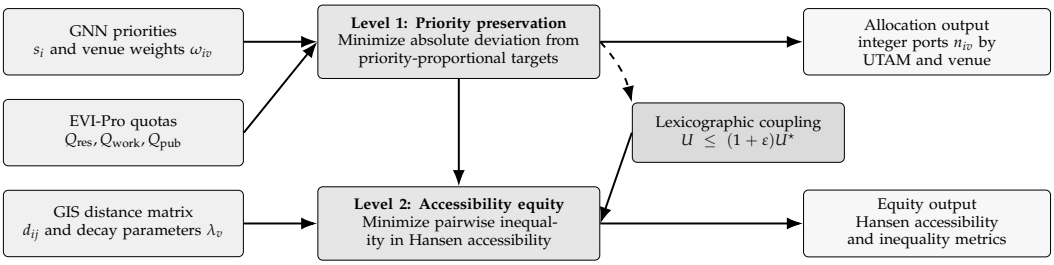


Figure 5. Structure of the bi-level optimization model. Inputs include GNN-derived priorities and venue weights, EVI-Pro venue quotas, and GIS-based distances; outputs are the number of charging ports by UTAM and venue.

2.6. Implementation Workflow and Reproducibility

Figure 6 summarizes the end-to-end workflow. EM-2023, ANDEMOS, and UTAM GIS data are cleaned, merged, and aggregated at the UTAM level. The resulting GeoDataFrame is used both to run the EVI-Pro Lite simulation (yielding citywide charging quotas) and to train the spatiotemporal GNN (yielding charging-priority scores and venue-specific weights). These outputs feed the bi-level optimization model, which produces an allocation that follows priorities (Level 1) and then improves accessibility equity (Level 2). All code is implemented in Python using open-source libraries (GeoPandas, PyTorch, PyTorch Geometric, Pyomo).

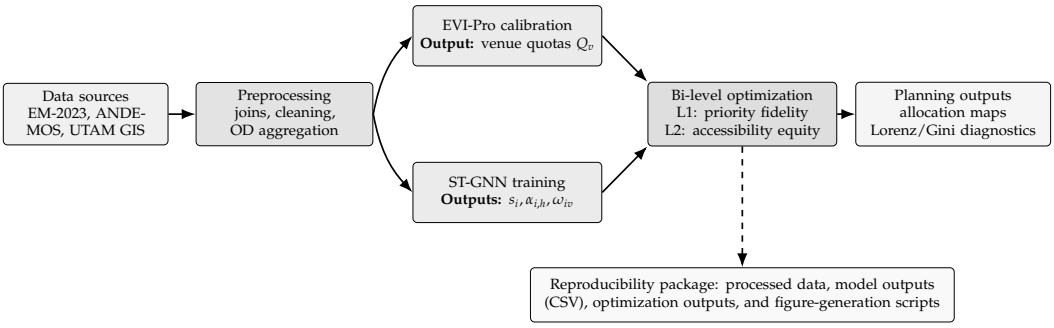


Figure 6. Methodological pipeline: data acquisition and GIS preprocessing, EVI-Pro Lite demand estimation, GNN-based priority learning, and bi-level optimization of charging infrastructure.

3. Results and Discussions

3.1. ST-GNN Training and Priority Extraction Results

Figure 7 shows the evolution of the training loss across epochs. The MAE decreases steadily during the early and intermediate stages of optimization and then transitions to a smoother convergence regime. Although minor oscillations remain, the overall trajectory indicates that the model learns stable spatial and temporal regularities from the UTAM-level mobility tensor rather than fitting a noisy or unstable representation.

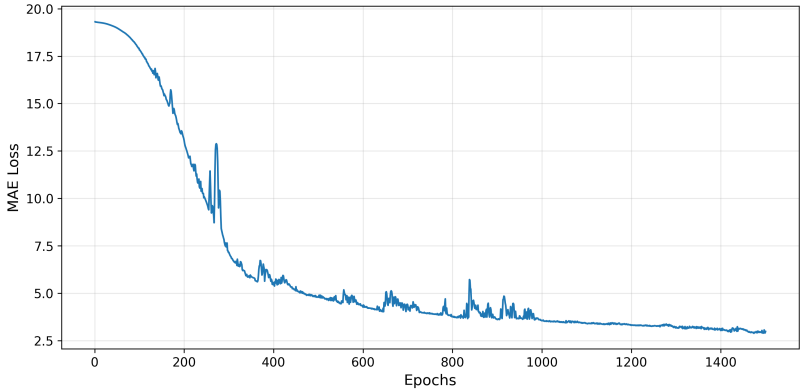


Figure 7. Training loss across epochs. Mean absolute error (MAE) during 1500 training epochs, showing stable convergence.

Because independent charging-session labels are unavailable, validation of ST-GNN outputs is performed against the observed mobility ground truth from EM-2023, rather than against observed charging demand. Pattern-level consistency checks are therefore used: (i) spatial consistency with high observed origin–destination activity zones (Equation (4)), (ii) temporal consistency between learned attention peaks and observed high-activity windows, and (iii) comparison against simpler mobility baselines. This should be interpreted as structural validation of mobility-derived priorities, not as direct validation of realized charging sessions.

Beyond convergence diagnostics, the role of the ST-GNN is clarified through comparisons with simpler reference formulations derived from the same EM-2023 tensor: (i) a direct mobility-intensity index based only on aggregated trip totals, and (ii) a fixed-window temporal heuristic using citywide venue shares without UTAM-specific attention. These baselines are useful as transparency checks, but they cannot jointly represent neighborhood propagation on the UTAM graph and zone-specific hour weighting. The ST-GNN is therefore used not merely as a smoother of trip intensity but as a representation learning step that generates two coupled planning signals: a spatial priority score and heterogeneous venue weights.

Accordingly, the model should be interpreted as a planning-oriented extractor of relative structure rather than as a black-box demand forecaster. This distinction is important in data-constrained settings: the objective is to recover actionable spatiotemporal patterns from survey data with limited direct charging observations, and then propagate those patterns into the optimization stage in a controlled and interpretable way.

After training, the model yields two outputs for each UTAM zone: a scalar priority score and a 24 h temporal attention profile. The attention vectors provide an interpretable representation of hour-level importance, allowing us to examine when mobility activity contributes most strongly to the learned charging-priority signal. Figure 8 reports the attention profiles for the highest-priority UTAM zones. In Bogotá, the most pronounced peaks are concentrated in late-evening and night-time periods, suggesting that the learned signal is closely associated with residential and end-of-day dwelling opportunities.

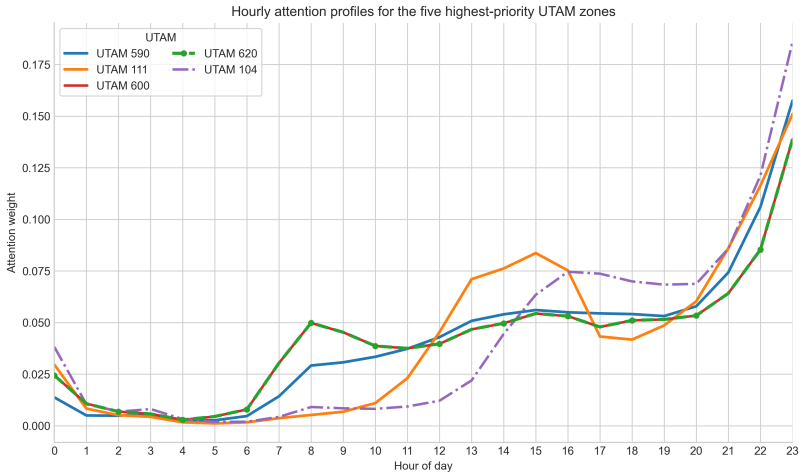


Figure 8. Hourly attention profiles for the five highest-priority UTAM zones. Learned 24-h weights indicate when mobility contributes most to charging priority. UTAM620 is shown with markers because its profile overlaps closely with UTAM600.

To enable spatial comparison across the study area, the raw priority scores are normalized through min–max scaling to obtain the index $s_i \in [0, 1]$. Figure 9 presents the resulting normalized UTAM priority map. The spatial pattern is clearly uneven: a relatively small subset of zones concentrates the highest values, while a large share of the city remains near the lower end of the scale. This indicates that the ST-GNN does not infer a uniform demand surface but instead identifies a strongly differentiated geography of charging priority from the observed spatiotemporal mobility structure.

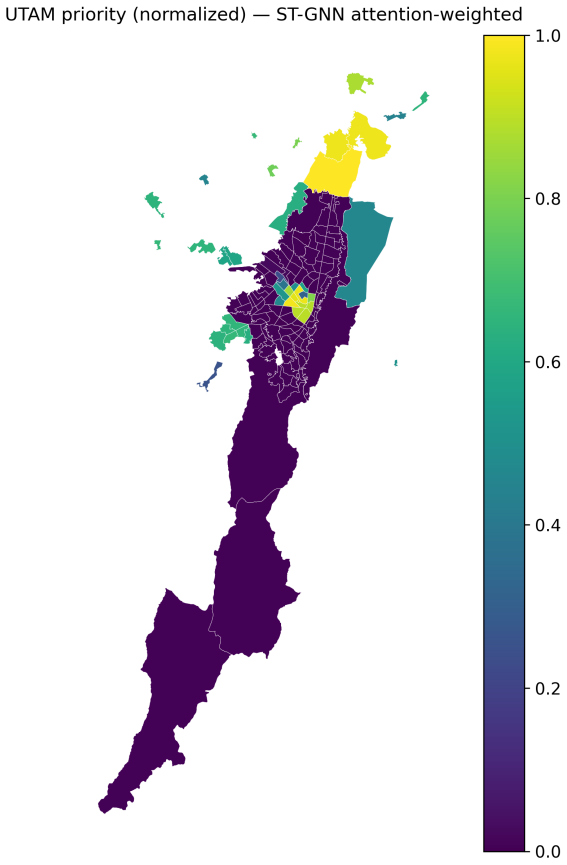


Figure 9. Normalized UTAM charging-priority map. Distribution of normalized ST-GNN priority scores across Bogotá's UTAM zones (approximate kilometer scale bar included).

The highest-priority zones are concentrated mainly in the central and northern parts of Bogotá, consistent with the mobility intensity captured in the survey data and with the temporal weighting learned by the network. These normalized scores constitute the principal demand-side signal used in the subsequent optimization stage, where they are preserved, subject to controlled relaxation, in the first level of the lexicographic allocation model.

3.2. Bi-Level Allocation and Equity Outcomes

The lexicographic optimization distributes the EVI-Pro quotas exactly across venue types, assigning 7352 residential ports, 2739 workplace ports, and 779 public ports, for a total of 10,870 charging ports citywide. The first optimization level preserves the ST-GNN priority structure through priority-proportional targets, while the second level redistributes accessibility through spatial spillovers in order to reduce inequality in the Hansen accessibility surface.

The final allocation remains strongly aligned with the learned priority scores in terms of installed capacity. As shown in Figure 10, the Spearman rank correlation between s_i and total installed ports is $\rho = 0.799$ ($p = 9.99 \times 10^{-33}$), indicating that the optimized allocation continues to reflect the mobility-derived priority ordering. However, the corresponding correlation between priority and Hansen accessibility drops to $\rho = 0.320$ ($p = 1.04 \times 10^{-4}$; Figure 11), which reveals the effect of the second-stage equity adjustment: accessibility is no longer a simple mirror of the priority ranking, because the model accounts for venue-specific distance decay and inter-zonal spillovers.

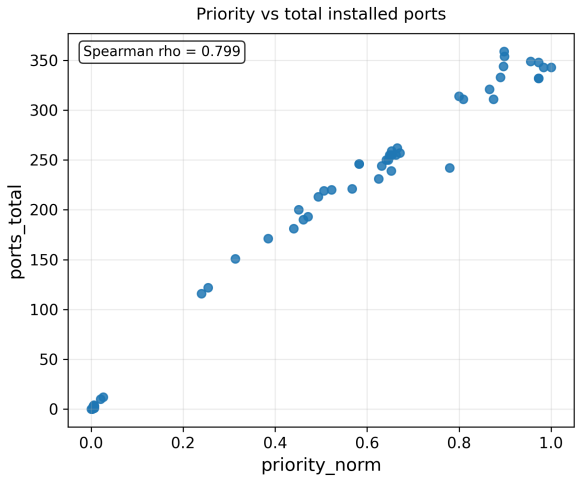


Figure 10. Priority versus total installed ports under the bi-level solution ($\rho = 0.799$, $p = 9.99 \times 10^{-33}$).

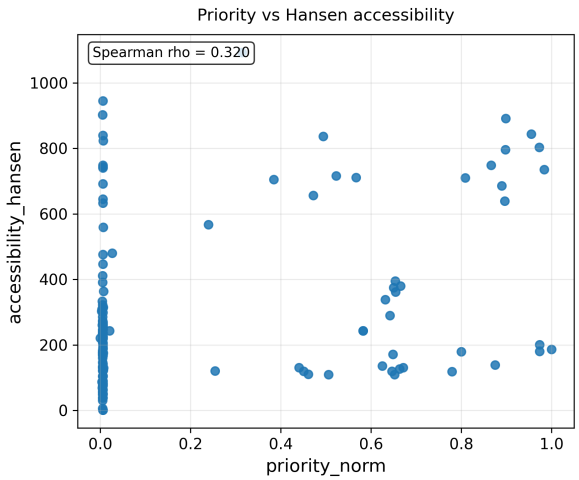


Figure 11. Priority versus Hansen accessibility under the bi-level solution ($\rho = 0.320$, $p = 1.04 \times 10^{-4}$).

This correlation is computed on total installed ports (all venue types combined). It should not be interpreted as implying identical rank retention by venue category. Because quota magnitudes differ substantially across residential, workplace, and public charging, venue-specific correlations can vary, especially for sparse categories such as public charging.

The resulting accessibility distribution remains unequal but substantially more diffused than the raw port allocation. The Hansen accessibility Gini coefficient is 0.433, and the Lorenz curve in Figure 12 shows that the bottom 50% of zones account for 0.204 of total accessibility. Table 6 summarizes the main quantitative outcomes. The mean Hansen accessibility reaches 296.630 (95% CI for the mean: 255.823–337.437), with a standard deviation of 248.099 and a minimum of 1.126. By comparison, the total port distribution is more concentrated, with a Gini coefficient of 0.733, indicating that the accessibility objective smooths part of the inequality induced by the physical concentration of installations.

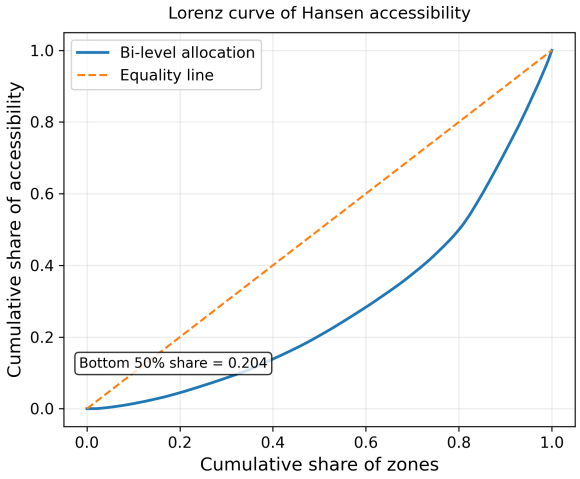


Figure 12. Lorenz curve of Hansen accessibility under the bi-level allocation. The bottom 50% of zones account for 0.204 of total accessibility.

Table 6. Summary metrics for the bi-level solution (Hansen accessibility, priority alignment, and uncertainty margins).

Metric	Value	Metric	Value
Min Hansen accessibility	1.126	Spearman (s_i, ports)	0.799
Mean Hansen accessibility	296.630	Spearman (s_i, Hansen access.)	0.320
SD Hansen accessibility	248.099	Gini (Hansen accessibility)	0.433
CV (Hansen accessibility)	0.836	Bottom 50% Lorenz share (Hansen)	0.204
Gini (ports_total)	0.733	Number of UTAM zones	142
p-value (Spearman s_i, ports)	9.99×10^{-33}	p-value (Spearman s_i, Hansen)	1.04×10^{-4}

At the same time, the large dispersion (SD/mean = 0.836) indicates that substantial heterogeneity remains after optimization. This appears to reflect both the inherent core–periphery spatial structure in Bogotá and modeling simplifications (fixed venue quotas, Euclidean impedance, and omitted grid/land-cost constraints).

A useful implication of these results is that the zone receiving the highest accessibility is not necessarily the zone receiving the highest number of chargers. In the final solution, the largest installed totals are concentrated in a handful of high-priority UTAMs, but the highest Hansen accessibility is achieved in centrally positioned zones that benefit from cumulative spillovers from neighboring installations. This confirms that the second-stage accessibility objective changes the spatial meaning of the solution from "where chargers are placed" to "where access is effectively experienced".

Uncertainty margins for the main statistics are as follows: Spearman (s_i, ports) = 0.799 (95% CI: 0.730–0.852; $p = 9.99 \times 10^{-33}$); Spearman (s_i, Hansen accessibility) = 0.320 (95% CI: 0.164–0.460; $p = 1.04 \times 10^{-4}$); mean Hansen accessibility = 296.630 (95% CI for the mean: 255.823–337.437).

3.3. Spatial Allocation Patterns

Figure 13 shows the spatial distribution of the final allocation by venue type. The total port map reveals a strong concentration of infrastructure in the north–central urban corridor and in a compact central cluster of high-priority zones. This pattern is consistent with the learned ST-GNN priorities, but the decomposition by venue type provides additional structure. Residential ports are the most spatially widespread category, workplace ports are more concentrated in high-access and employment-oriented zones, and public ports are assigned more selectively, reflecting both the smaller quota and the sharper venue-specific weighting induced by the optimization model.

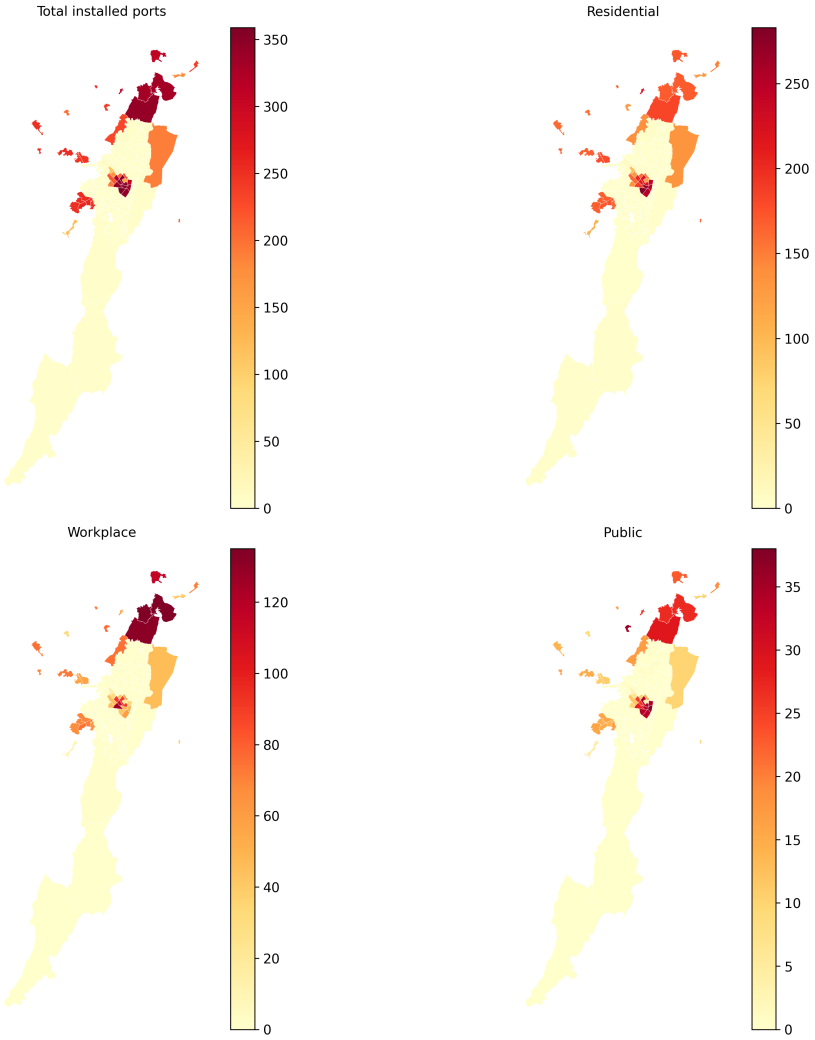


Figure 13. Bi-level UTAM allocation by venue: total, residential, workplace, and public ports (approximate kilometer scale bar included).

The concentration of installed capacity is nontrivial: the ten highest-allocation UTAM zones account for approximately 31.6% of all ports. At the same time, the final solution is not purely concentrated in a few locations. Residential ports appear in 140 of the 142 zones, workplace ports in 80 zones, and public ports in 47 zones, indicating that the optimization balances concentration and territorial coverage rather than collapsing entirely into a

small core. Figure 14 shows the resulting geography of realized accessibility, combining local and spillover effects under the final bi-level allocation.

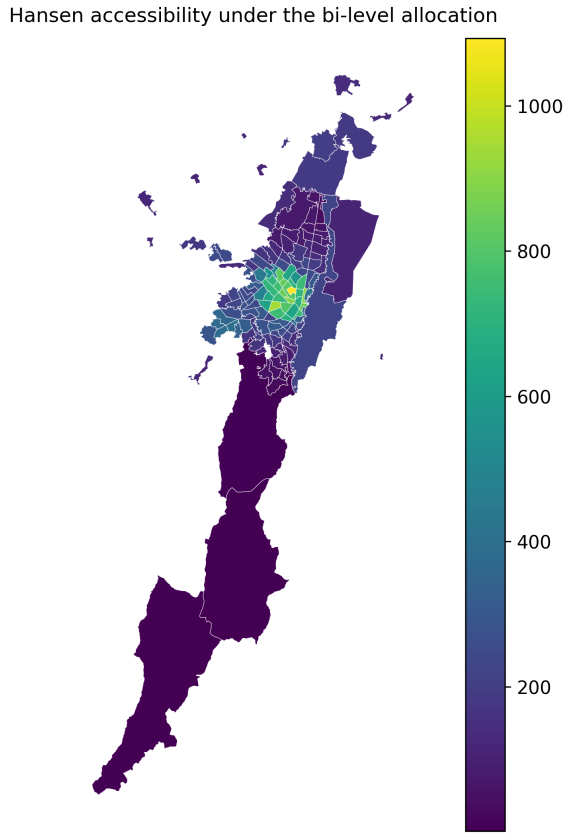


Figure 14. Hansen accessibility under the bi-level allocation, including local and spillover effects (approximate kilometer scale bar included).

3.4. Integrated Discussion

The proposed framework separates two planning objectives that are often conflated in EV infrastructure studies: aligning charging capacity with a mobility-derived signal of relative need, and improving the territorial equity of access once that signal has been respected. The lexicographic structure makes this distinction explicit. In the present case, the final solution preserves a strong monotonic relationship between the learned priority scores and installed charging capacity, as reflected by the Spearman correlation of $\rho = 0.799$ between s_i and total ports. This confirms that the first optimization level is effective in retaining the spatial structure learned by the ST-GNN, even after the second-stage equity refinement is imposed.

At the same time, preserving priority is not equivalent to preserving the same ranking in realized service access. The correlation between priority and Hansen accessibility falls to $\rho = 0.320$, indicating that the second optimization level reshapes the geography of effective access. If the objective only followed priority-proportional targets, the final map would largely reproduce the concentration pattern of the ST-GNN outputs. Once accessibility equalization is introduced, the solution becomes more territorially diffused because each installed port contributes both to its host zone and to surrounding zones through the distance-decay kernel. The optimization therefore determines not only where chargers are placed but also how accessibility benefits are redistributed.

This distinction is reflected in the inequality metrics. The final Hansen accessibility Gini coefficient is 0.433, while the Gini coefficient of total port distribution is 0.733, indicating that physical infrastructure remains more concentrated than the accessibility it generates.

This gap confirms that accessibility-based evaluation is essential for equity analysis, because centrally located chargers can improve access in multiple neighboring zones, and those spillovers are missed by simple zone-level port counts. The Lorenz result is consistent: the bottom 50% of zones account for 0.204 of total Hansen accessibility. Although inequality remains, accessibility is still less concentrated than raw port allocation, suggesting that the second-stage objective partially compensates for polarization in learned priorities.

The spatial results also reveal a meaningful tension between concentration and coverage. A disproportionate share of ports is assigned to a relatively small subset of high-priority UTAM zones, especially in the north–central and central urban structure, consistent with the ST-GNN signal and observed mobility intensity. However, the venue-specific decomposition shows that the model does not collapse into a purely core-based strategy. Residential ports are distributed across nearly the entire city, while workplace and public ports are more selective and concentrated. This pattern is behaviorally plausible and suggests that temporal-attention venue weights do more than scale quotas: they introduce behavioral differentiation into the final allocation.

From a methodological standpoint, the results support the value of combining survey-based representation learning with accessibility-aware optimization in data-constrained environments. Bogotá does not yet have large-scale revealed charging-demand data, and this limitation would normally prevent planners from constructing detailed demand surfaces. The ST-GNN partially addresses this constraint by extracting relative priority patterns from origin–destination flows and time-of-day structure. The optimization model then translates those learned priorities into a policy-relevant infrastructure plan under fixed venue quotas. The resulting pipeline is especially useful for cities where detailed EV telemetry is unavailable but household travel surveys and GIS layers are accessible.

The revision also strengthens the interpretation of model robustness at two levels. On the demand-estimation side, the scenario protocol in Section 2.3 clarifies how fleet size, BEV/PHEV composition, and charging-access assumptions can shift venue quotas before spatial allocation. On the optimization side, parameter choices are explicitly policy-interpretable: ε controls the acceptable deviation from priority fidelity, while λ_v encodes venue-specific spatial reach in the accessibility kernel. This separation helps distinguish uncertainty in total infrastructure need from normative choices in redistribution.

At the same time, the results also underline the normative nature of the planning problem. A stronger emphasis on Level-1 fidelity would produce a more sharply concentrated allocation aligned with the learned priorities, whereas a stronger emphasis on Level-2 accessibility equalization would further diffuse access across zones. The relaxation parameter in the lexicographic coupling therefore has a clear policy interpretation: it defines how much deviation from the estimated need signal is acceptable in exchange for a fairer accessibility distribution. In practical planning settings, this parameter could be selected through stakeholder negotiation rather than purely technical tuning.

Overall, the findings suggest that equitable EV charging deployment should not be framed as a choice between efficiency and fairness in the abstract. Instead, it should be understood as a structured trade-off between preserving a behaviorally grounded demand proxy and redistributing the accessibility generated by infrastructure across the urban network. The bi-level framework developed here offers a transparent way to formalize that trade-off and to produce allocations that remain interpretable both from a mobility analysis perspective and from an equity perspective.

3.5. Minimum Data Requirements and Transferability

To strengthen reproducibility claims for other data-constrained cities, a practical minimum-input checklist is summarized. At minimum, the framework requires (i) a zonal

polygon layer with centroids/adjacency, (ii) household travel survey data or equivalent OD activity by hour, (iii) a city-level EV adoption projection with at least BEV/PHEV shares, and (iv) context-specific assumptions on charging access by venue. Without these four elements, the full ST-GNN + EVI-Pro + bi-level pipeline cannot be operationalized consistently.

Transfer to another city can be structured in five steps: (1) harmonize zoning and survey records into a common geospatial key, (2) construct UTAM-equivalent graph/tensor inputs, (3) calibrate EVI-Pro assumptions to local fleet and charging-access conditions, (4) train the ST-GNN to derive local priority and venue weights, and (5) solve the lexicographic allocation under locally selected equity and policy parameters. This sequence clarifies which components are portable by design and which require local recalibration.

4. Conclusions

This study presents a reproducible framework for equitable EV charging planning in resource-constrained cities. The approach combines ST-GNN prioritization, EVI-Pro Lite demand estimation, and lexicographic bi-level optimization. Applied to Bogotá, the framework preserves mobility-derived priority structure in installed capacity (Spearman $\rho = 0.799$, $p < 10^{-31}$), while the priority–accessibility relationship is weaker after equity refinement (Spearman $\rho = 0.320$, $p = 1.04 \times 10^{-4}$). Accessibility equity is improved overall (Hansen Gini $= 0.433$; bottom-50% Lorenz share $= 0.204$).

The main practical implication is that planners do not need to choose between demand fidelity and equity in purely qualitative terms: the trade-off can be made explicit, parameterized, and auditable. This is especially relevant for Latin American cities where household survey data are available but revealed charging microdata remain limited.

Several limitations should be considered when interpreting these findings. Priority signals remain mobility-derived proxies rather than observed charging sessions; infrastructure deployment is modeled as a static one-step decision; Euclidean impedance is used instead of network travel time; and implementation constraints such as grid capacity, land availability, zoning, and installation costs are not explicitly modeled. Future work should therefore prioritize multi-period deployment, network-based impedance calibration, and integration of feasibility constraints and social vulnerability indicators.

Supplementary Materials: Processed datasets, trained model weights, optimization code, and figure generation scripts are available in a public GitHub repository (project repository, https://github.com/jorgeiv500/bogota-ev-charging-gnn-bilevel accessed on 9 April 2026). Raw household travel survey data subject to privacy protections cannot be publicly shared but may be requested from Bogotá's District Secretariat of Mobility (Spanish: Secretaría Distrital de Movilidad) for research purposes.

Author Contributions: All authors contributed equally to the conceptualization, methodology, analysis, writing, and revision of this manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Universidad de Bogotá Jorge Tadeo Lozano. The APC was funded by Universidad de Bogotá Jorge Tadeo Lozano.

Data Availability Statement: The original contributions presented in the study are included in the article/Supplementary Materials, further inquiries can be directed to the corresponding authors.

Acknowledgments: We acknowledge Bogotá's District Secretariat of Mobility (Spanish: Secretaría Distrital de Movilidad) for providing access to EM-2023 household travel survey data, and ANDEMOS for electric vehicle registration statistics.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BEV	Battery Electric Vehicle
PHEV	Plug-in Hybrid Electric Vehicle
GNN	Graph Neural Network
GCN	Graph Convolutional Network
LSTM	Long Short-Term Memory
MILP	Mixed-Integer Linear Programming
UTAM	Transport and Mobility Analysis Unit
EM-2023	Bogotá 2023 Mobility Survey (Encuesta de Movilidad 2023)
ANDEMOS	National Association for Sustainable Mobility (Asociación Nacional de Movilidad Sostenible)
EVI-Pro	Electric Vehicle Infrastructure Projection Tool

References

1. International Energy Agency (IEA). *Global EV Outlook 2023*; IEA: Paris, France, 2023. Available online: https://www.iea.org/reports/global-ev-outlook-2023 (accessed on 10 March 2026).
2. International Energy Agency (IEA). *Global EV Outlook 2024*; IEA: Paris, France, 2024. Available online: https://www.iea.org/reports/global-ev-outlook-2024 (accessed on 10 March 2026).
3. International Energy Agency (IEA). *Global EV Outlook 2025*; IEA: Paris, France, 2025. Available online: https://www.iea.org/reports/global-ev-outlook-2025 (accessed on 10 March 2026).
4. Lucas, K.; Mattioli, G.; Verlinghieri, E.; Guzman, A. Transport poverty and its adverse social consequences. *Proc. Inst. Civ. Eng. Transp.* **2016**, *169*, 353–365. [CrossRef]
5. Carlton, G.J.; Sultana, S. Electric vehicle charging equity and accessibility: A comprehensive United States policy analysis. *Transp. Res. Part D Transp. Environ.* **2024**, *129*, 104123. [CrossRef]
6. Canepa, K.; Hardman, S.; Tal, G. An early look at plug-in electric vehicle adoption in disadvantaged communities in California. *Transp. Policy* **2019**, *78*, 19–30. [CrossRef]
7. Yu, Q.; Que, T.; Cushing, L.J.; Pierce, G.; Shen, K.; Kejriwal, M.; Yao, Y.; Zhu, Y. Equity and reliability of public electric vehicle charging stations in the United States. *Nat. Commun.* **2025**, *16*, 5291. [CrossRef]
8. Delmelle, E.C.; Casas, I. Evaluating the spatial equity of bus rapid transit-based accessibility patterns in a developing country: The case of Cali, Colombia. *Transp. Policy* **2012**, *20*, 36–46. [CrossRef]
9. District Secretariat of Mobility (Secretaría Distrital de Movilidad, SDM). *Sustainable and Safe Mobility Plan (Plan de Movilidad Sostenible y Segura)*; District Secretariat of Mobility: Bogotá, D.C., Colombia, 2025. Available online: https://www.movilidadbogota.gov.co/web/plan_de_movilidad_sostenible_y_segura (accessed on 10 March 2026).
10. ReVelle, C.S.; Eiselt, H.A. Location analysis: A synthesis and survey. *Eur. J. Oper. Res.* **2005**, *165*, 1–19. [CrossRef]
11. Zhang, B.; Zhao, M.; Hu, X. Location planning of electric vehicle charging station with users' preferences and waiting time: Multi-objective bi-level programming model and HNSGA-II algorithm. *Int. J. Prod. Res.* **2023**, *61*, 1394–1423. [CrossRef]
12. Qiao, D.; Wang, G.; Xu, M. Fast-charging station location problem: A two-phase approach with mathematical program with equilibrium constraints considering charging choice behaviour. *Sustain. Cities Soc.* **2023**, *96*, 104678. [CrossRef]
13. Zhou, G.; Dong, Q.; Zhao, Y.; Wang, H.; Jian, L.; Jia, Y. Bilevel optimization approach to fast charging station planning in electrified transportation networks. *Appl. Energy* **2023**, *350*, 121718. [CrossRef]
14. Li, H.; He, Y.; Fu, W.; Li, X. Bi-level planning of electric vehicle charging station in coupled distribution-transportation networks. *Electr. Power Syst. Res.* **2024**, *232*, 110442. [CrossRef]
15. Gan, H.; Ruan, W.; Wang, M.; Pan, Y.; Miu, H.; Yuan, X. Bi-Level Planning of Electric Vehicle Charging Stations Considering Spatial–Temporal Distribution Characteristics of Charging Loads in Uncertain Environments. *Energies* **2024**, *17*, 3004. [CrossRef]
16. Sun, Z.; He, Y.; Wu, H.; Wu, A.Y. Bi-level planning of electric vehicle charging stations considering charging demand: A Nash bargaining game approach. *Energy* **2025**, *332*, 137137. [CrossRef]
17. Varghese, A.M.; Menon, N.; Ermagun, A. Equitable distribution of electric vehicle charging infrastructure: A systematic review. *Renew. Sustain. Energy Rev.* **2024**, *206*, 114825. [CrossRef]
18. Hopkins, E.; Potoglou, D.; Orford, S.; Cipcigan, L. Can the equitable roll out of electric vehicle charging infrastructure be achieved? *Renew. Sustain. Energy Rev.* **2023**, *182*, 113398. [CrossRef]
19. Jiao, J.; Choi, S.J.; Nguyen, C. Toward an equitable transportation electrification plan: Measuring public electric vehicle charging station access disparities in Austin, Texas. *PLoS ONE* **2024**, *19*, e0309302. [CrossRef]

20. Lee, J.H.; Chakraborty, D.; Hardman, S.J.; Tal, G. Exploring electric vehicle charging patterns: Mixed usage of charging infrastructure. *Transp. Res. Part D Transp. Environ.* **2020**, *79*, 102249. [CrossRef]
21. Wood, E. Evi-Pro Lite API. Dataset. United States. 2022. Available online: https://www.osti.gov/biblio/1855624 (accessed on 10 March 2026).
22. Scarselli, F.; Gori, M.; Tsoi, A.C.; Hagenbuchner, M.; Monfardini, G. The graph neural network model. *IEEE Trans. Neural Netw.* **2009**, *20*, 61–80. [CrossRef]
23. Kipf, T.N.; Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv* **2016**, arXiv:1609.02907. [CrossRef]
24. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]
25. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. *arXiv* **2017**, arXiv:1706.03762. [CrossRef]
26. Bahdanau, D.; Cho, K.; Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv* **2014**, arXiv:1409.0473. [CrossRef]
27. Yilmaz, M.; Krein, P.T. Review of battery charger topologies, charging power levels, and infrastructure for plug-in electric and hybrid vehicles. *IEEE Trans. Power Electron.* **2013**, *28*, 2151–2169. [CrossRef]
28. Mwasilu, F.; Justo, J.J.; Kim, E.-K.; Do, T.D.; Jung, J.-W. Electric vehicles and smart grid interaction: A review on vehicle-to-grid and renewable energy sources integration. *Renew. Sustain. Energy Rev.* **2014**, *34*, 501–516. [CrossRef]
29. Hardman, S.; Jenn, A.; Tal, G.; Axsen, J.; Beard, G.; Daina, N.; Figenbaum, E.; Jakobsson, N.; Jochem, P.E.P.; Kinnear, N. A review of consumer preferences of and interactions with electric vehicle charging infrastructure. *Transp. Res. Part D Transp. Environ.* **2018**, *62*, 508–523. [CrossRef]
30. Neaimeh, M.; Salisbury, S.D.; Hill, G.A.; Blythe, P.T.; Scoffield, D.R.; Francfort, J.E. Analysing the usage and evidencing the importance of fast chargers for the adoption of battery electric vehicles. *Energy Policy* **2017**, *108*, 474–486. [CrossRef]
31. Noel, L.; Zarazua de Rubens, G.; Sovacool, B.K.; Kester, J. Fear and loathing of electric vehicles: The reactionary rhetoric of range anxiety. *Energy Res. Soc. Sci.* **2019**, *48*, 96–107. [CrossRef]
32. Frade, I.; Ribeiro, A.; Gonçalves, G.; Antunes, A.P. Optimal location of charging stations for electric vehicles in a neighborhood in Lisbon, Portugal. *Transp. Res. Rec.* **2011**, *2252*, 91–98. [CrossRef]
33. He, F.; Wu, D.; Yin, Y.; Guan, Y. Optimal deployment of public charging stations for plug-in hybrid electric vehicles. *Transp. Res. Part B Methodol.* **2013**, *47*, 87–101. [CrossRef]
34. Dong, J.; Liu, C.; Lin, Z. Charging infrastructure planning for promoting battery electric vehicles: An activity-based approach using multiday travel data. *Transp. Res. Part C Emerg. Technol.* **2014**, *38*, 44–55. [CrossRef]
35. Guerrero-Silva, J.A.; Romero-Gelvez, J.I.; Aristizábal, A.J.; Zapata, S. Optimization and trends in EV charging infrastructure: A PCA-based systematic review. *World Electr. Veh. J.* **2025**, *16*, 345. [CrossRef]
36. Xylia, M.; Olsson, E.; Macura, B.; Nykvist, B. Estimating charging infrastructure demand for electric vehicles: A systematic review. *Energy Strategy Rev.* **2025**, *59*, 101753. [CrossRef]
37. Wang, J.; Kaushik, H.D.; Jacob, R.A.; Zhang, J. Spatiotemporal planning of electric vehicle charging infrastructure: Demand estimation and grid-aware optimization under uncertainty. *iScience* **2025**, *28*, 113368. [CrossRef]
38. Wang, X.; Jiang, J.; Zhou, Y.; Wang, X. Integrated infrastructure planning of charging and electricity generation. *Transp. Res. Part Transp. Environ.* **2023**, *122*, 103807. [CrossRef]
39. Wang, S.; Zhuge, C.; Shao, C.; Wang, P.; Yang, X.; Wang, S. Short-term electric vehicle charging demand prediction: A deep learning approach. *Appl. Energy* **2023**, *340*, 121032. [CrossRef]
40. Wang, S.; Chen, A.; Wang, P.; Zhuge, C. Predicting electric vehicle charging demand using a heterogeneous spatio-temporal graph convolutional network. *Transp. Res. Part C Emerg. Technol.* **2023**, *153*, 104205. [CrossRef]
41. Çavuş, M.; Ayan, H.; Dissanayake, D.; Sharma, A.; Deb, S.; Bell, M. Forecasting electric vehicle charging demand in smart cities using hybrid deep learning of regional spatial behaviours. *Energies* **2025**, *18*, 3425. [CrossRef]
42. Voinea, G.-D.; Gîrbacia, F.; Duguleană, M.; Postelnicu, C.-C. Planning future EV charging infrastructure by forecasting spatio-temporal adoption trends across heterogeneous user segments. *Information* **2025**, *16*, 933. [CrossRef]
43. Alaraj, M.; Radi, M.; Alsisi, E.; Majdalawieh, M.; Darwish, M. Machine learning-based electric vehicle charging demand forecasting: A systematized literature review. *Energies* **2025**, *18*, 4779. [CrossRef]
44. Hansen, W.G. How accessibility shapes land use. *J. Am. Inst. Plann.* **1959**, *25*, 73–76. [CrossRef]
45. Geurs, K.T.; van Wee, B. Accessibility evaluation of land-use and transport strategies: Review and research directions. *J. Transp. Geogr.* **2004**, *12*, 127–140. [CrossRef]
46. Chen, B.; Moore, A.; Jia, B.; Zhang, K.; Cao, M. Equity considerations in public electric vehicle charging: A review. *World Electr. Veh. J.* **2025**, *16*, 553. [CrossRef]
47. Lezcano, C.M.S.; Harper, C.D.; Nock, D. Siting for demand and equity: Optimizing level 2 electric vehicle charger placement. *J. Transp. Geogr.* **2025**, *128*, 104369. [CrossRef]

48. Soltani Mandolakani, F.; Singleton, P.A. Electric vehicle charging infrastructure deployment: A discussion of equity and justice theories and accessibility measurement. *Transp. Res. Interdiscip. Perspect.* **2024**, *24*, 101072. [CrossRef]
49. Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; Yu, P.S. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *32*, 4–24. [CrossRef]
50. van den Berg, R.; Kipf, T.N.; Welling, M. Graph convolutional matrix completion. *arXiv* **2017**, arXiv:1706.02263. [CrossRef]
51. Li, Y.; Yu, R.; Shahabi, C.; Liu, Y. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In Proceedings of the International Conference on Learning Representations (ICLR), Vancouver, BC, Canada, 30 April–3 May 2018. Available online: https://openreview.net/forum?id=SJiHXGWAZ (accessed on 10 March 2026).
52. Yu, B.; Yin, H.; Zhu, Z. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, Stockholm, Sweden, 13–19 July 2018; pp. 3634–3640. [CrossRef]
53. Wu, Z.; Pan, S.; Long, G.; Jiang, J.; Zhang, C. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, Macao, China, 10–16 August 2019; pp. 1907–1913. [CrossRef]
54. Jiang, W.; Luo, J. Graph neural network for traffic forecasting: A survey. *Expert Syst. Appl.* **2022**, *207*, 117921. [CrossRef]
55. Jiang, W.; Luo, J.; He, M.; Gu, W. Graph neural network for traffic forecasting: The research progress. *ISPRS Int. J. Geo-Inf.* **2023**, *12*, 100. [CrossRef]
56. Rahmani, S.; Baghbani, A.; Bouguila, N.; Patterson, Z. Graph neural networks for intelligent transportation systems: A survey. *IEEE Trans. Intell. Transp. Syst.* **2023**, *24*, 8846–8885. [CrossRef]
57. Xue, R.; Zhang, J.; Li, Y.; Wang, H. Dynamic spatio-temporal networks with graph neural architectures for traffic forecasting. *Data Sci. Transp.* **2025**, *7*, 10. [CrossRef]
58. Narayanan, S.; Makarov, N.; Antoniou, C. Graph neural networks as strategic transport modelling alternative–A proof of concept for a surrogate. *IET Intell. Transp. Syst.* **2024**, *18*, 2059–2077. [CrossRef]
59. Wang, S.; Li, Y.; Shao, C.; Wang, P.; Wang, A.; Zhuge, C. An adaptive spatio-temporal graph recurrent network for short-term electric vehicle charging demand prediction. *Appl. Energy* **2025**, *383*, 125320. [CrossRef]
60. Chen, Y.; Tang, Z.; Cui, Y.; Rao, W.; Li, Y. Electric Vehicle Charging Demand Prediction Model Based on Spatiotemporal Attention Mechanism. *Energies* **2025**, *18*, 687. [CrossRef]
61. Jiang, D.; Gong, X.; Wei, Y.; Peng, B.; Xu, Z. An electric vehicle charging demand prediction approach based on a Graph-based Spatio-temporal Attention Network. *Sustain. Energy Grids Netw.* **2025**, *44*, 101975. [CrossRef]
62. District Secretariat of Mobility (Secretaría Distrital de Movilidad, SDM). *Bogotá 2023 Mobility Survey: UTAM Zoning and Cartography Downloads (Encuesta de Movilidad 2023: Descargas de zonificación y cartografía UTAM)*; Mobility Observatory (Observatorio de Movilidad): Bogotá, D.C., Colombia, 2023. Available online: https://observatorio.movilidadbogota.gov.co/publicaciones/resultados/encuesta-de-movilidad-2023 (accessed on 26 March 2026).
63. GeoPandas Development Team. GeoPandas: Python Tools for Geographic Data. Version 0.14.4. 2024. Available online: https://geopandas.org (accessed on 26 March 2026).
64. Bynum, M.L.; Hackebeil, G.A.; Hart, W.E.; Laird, C.D.; Nicholson, B.L.; Siirola, J.D.; Watson, J.-P.; Woodruff, D.L. *Pyomo—Optimization Modeling in Python*, 3rd ed.; Springer: Cham, Switzerland, 2021. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.