Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI-Driven Optimization of Electric Vehicle Charging Station Placement in Uzbekistan: A Data-Driven Approach for Sustainable Transportation Infrastructure

Abzalbek Mahkamov

avzalbekmahkamov@gmail.com

Tashkent state transport university https://orcid.org/0009-0009-1939-8825

Research Article

Keywords: electric vehicles, charging infrastructure, facility location, machine learning, multi-objective optimization, mixed-integer programming, sustainable transportation, Uzbekistan, smart grid, geographic information systems

Posted Date: June 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10036351/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

AI-Driven Optimization of Electric Vehicle Charging Station Placement in Uzbekistan: A Data-Driven Approach for Sustainable Transportation Infrastructure

Mahkamov Abzalbek Baxtiyorjon o'g'li[1,*]

[1] *Department of Electric Vehicle Engineering, Faculty of Transport Vehicle Engineering, Tashkent State Transport University, Tashkent, Uzbekistan*

ORCID: 0009-0009-1939-8825

[*] **Corresponding author:** avzalbekmahkamov@gmail.com

Abstract

Background. Uzbekistan has emerged as the fastest-growing market for electric vehicles (EVs) in Central Asia: in 2024 the combined import of battery-electric and hybrid vehicles exceeded that of internal-combustion passenger cars for the first time, and national policy now targets thousands of new public charging points. Yet the spatial deployment of charging infrastructure has so far been driven largely by ad-hoc, demand-following decisions rather than systematic optimization, creating a risk of redundant clustering, under-served districts, and avoidable stress on an ageing, partly Soviet-era electricity grid. **Objective.** This study develops and validates an artificial-intelligence (AI)-driven, data-driven framework for optimizing the placement of EV charging stations under the conditions characteristic of a rapidly motorizing emerging economy. **Methods.** The framework integrates (i) a gradient-boosted machine-learning model that predicts spatially resolved charging demand from socio-demographic, traffic, land-use and EV-registration features; (ii) a geographic-information-system (GIS) screening stage that generates feasible candidate sites; (iii) a multi-criteria weighting procedure combining the analytic hierarchy process with entropy weighting; and (iv) a multi-objective mixed-integer optimization model that jointly minimizes total annualized cost and electrical-grid loading while maximizing demand coverage and spatial equity, solved with a hybrid genetic-algorithm–particle-swarm metaheuristic. **Results.** Using the city of Tashkent as a case study, the demand-prediction model attained a coefficient of determination of $R^2 = 0.889$ with the gradient-boosting variant, outperforming linear and single-tree baselines. The optimization identified a knee-point deployment of 148 stations achieving 91.4% demand coverage at an estimated annualized cost of approximately 4.4 million USD, improving coverage by roughly 23 percentage points and reducing modelled peak feeder loading by about 18% relative to a demand-following baseline of comparable size. **Conclusions.** Coupling machine-learning demand forecasting with grid-aware multi-objective optimization yields charging networks that are simultaneously more accessible, more equitable, and less disruptive to the power system. The methodology is transferable to other regions of Uzbekistan and to comparable emerging markets, and it provides a quantitative basis for evidence-led national charging-infrastructure policy.

Keywords: electric vehicles; charging infrastructure; facility location; machine learning; multi-objective optimization; mixed-integer programming; sustainable transportation; Uzbekistan; smart grid; geographic information systems.

1. Introduction

The decarbonization of road transport is one of the central pillars of global climate-mitigation strategy, and the electrification of the passenger-vehicle fleet is widely regarded as its most scalable instrument. The transition from internal-combustion-engine (ICE) vehicles to electric vehicles, however, is not determined solely by the availability or affordability of the vehicles themselves. It is increasingly recognized that the rate, breadth and equity of EV adoption are governed to a decisive degree by the availability, accessibility and reliability of public charging infrastructure. Range anxiety, charger queuing, and uneven geographical coverage repeatedly emerge in the literature as the dominant non-price barriers to adoption, and the spatial configuration of the charging network therefore becomes a first-order policy variable rather than a mere operational afterthought.

Against this backdrop, Uzbekistan presents an unusually instructive setting. Over a short period the country has moved from a marginal EV market to the most dynamic one in Central Asia. The combined import of battery-electric and hybrid vehicles surpassed that of conventional gasoline cars for the first time in 2024, and the overwhelming majority of newly registered electric cars are imported, low-cost models that have rapidly broadened ownership beyond a narrow premium segment. This surge has been actively cultivated by government policy, including exemptions from excise duties, customs fees and vehicle taxes, preferential financing, and explicit commitments to expand the public charging network into the thousands of charge points. The pace of adoption has consequently outstripped the pace of deliberate infrastructure planning, and charging stations have largely been sited where commercial operators perceived immediate demand—typically along a small number of high-traffic corridors in the capital—rather than according to any optimized, system-wide criterion.

This reactive deployment pattern produces several well-documented inefficiencies. First, it tends to over-concentrate chargers in commercially attractive central locations while leaving peripheral and lower-income districts under-served, undermining the equity of the transition. Second, uncoordinated siting can create local clusters of high simultaneous load on distribution feeders that were never dimensioned for such demand, a particularly acute concern where the distribution network is ageing and partly inherited from the Soviet period. Third, in the absence of demand forecasting, operators risk both under-provision (queuing and lost charging sessions) and over-provision (stranded capital in low-utilization stations). A planning methodology that anticipates demand, respects grid constraints, and balances competing objectives is therefore of direct and immediate value.

Artificial intelligence and operations-research techniques are well suited to this class of problem. The placement of charging stations can be formalized as a constrained facility-location problem, a family of optimization problems with a long pedigree in spatial economics and logistics. Modern variants enrich the classical formulation with data-driven demand estimation—using supervised machine learning to translate heterogeneous urban data into spatially resolved demand surfaces—and with multiple, often conflicting, objectives such as cost minimization, coverage maximization, equity, and grid stability. Because the resulting optimization problems are combinatorial and frequently non-convex, exact solvers do not scale to city-wide instances, and metaheuristic search has become the standard practical tool for obtaining high-quality approximate solutions.

Despite a substantial international literature on EV charging-station placement, the great majority of studies are calibrated to high-income cities in North America, Western Europe and East Asia, with mature grids, rich open data, and EV fleets dominated by long-range premium vehicles. The conditions of a fast-

motorizing emerging economy—rapid and partly unpredictable demand growth, a constrained legacy grid undergoing simultaneous renewable transformation, sparse open data, and a vehicle mix skewed toward affordable short-to-medium-range models—are materially different and remain comparatively understudied. To the author's knowledge, no prior work has formulated and validated an integrated, grid-aware, AI-driven charging-placement framework specifically for Uzbekistan.

Contributions. This paper makes the following contributions:

(1) It formulates EV charging-station placement for an emerging-market context as a four-objective mixed-integer program that simultaneously addresses cost, demand coverage, spatial equity and distribution-grid loading—the last being frequently omitted in regional studies.

(2) It integrates a gradient-boosted machine-learning demand predictor and GIS-based candidate screening with a hybrid genetic-algorithm–particle-swarm solver, producing an end-to-end, reproducible pipeline.

(3) It instantiates and evaluates the framework on the city of Tashkent, quantifying the gains over demand-following and coverage-only baselines and analysing the sensitivity of the solution to budget, equity weighting and renewable-supply scenarios.

(4) It distils the quantitative findings into concrete, transferable policy recommendations for national charging-infrastructure planning in Uzbekistan and comparable economies.

The remainder of the paper is organized as follows. Section 2 reviews related work on charging-station placement and AI methods. Section 3 characterizes the EV, grid and policy landscape of Uzbekistan. Section 4 describes the proposed framework and Section 5 sets out the formal optimization model. Section 6 presents the Tashkent case study, Section 7 reports and discusses the results, and Section 8 develops the policy implications. Section 9 states the limitations and directions for future work, and Section 10 concludes.

2. Related Work

Research on the optimal location of EV charging infrastructure spans operations research, transportation science, electrical engineering, geography and, increasingly, machine learning. It is useful to organize the field into four broad and partly overlapping methodological streams: mathematical-programming formulations, metaheuristic search, data-driven and machine-learning methods, and spatial multi-criteria decision analysis. The present work draws on, and seeks to integrate, all four.

2.1 Mathematical-programming and facility-location formulations

The conceptual foundation of charging-placement optimization lies in classical facility-location theory. The p-median problem, which locates a fixed number of facilities so as to minimize the total demand-weighted distance to the nearest facility, and the maximal-covering location problem, which maximizes the demand covered within a service radius for a given number of facilities, remain the canonical templates. Charging-specific extensions add capacity limits, queueing considerations, electrical constraints and the distinction between destination, en-route and fast-charging. Mixed-integer linear programming (MILP) provides a precise and flexible language for these formulations, allowing planners to encode budget ceilings, minimum-coverage requirements and grid limits as explicit constraints. The principal drawback is computational: charging-placement MILPs are NP-hard, and exact branch-and-bound solvers do not scale

gracefully to the thousands of candidate sites and demand nodes that arise in a city-wide instance, which restricts purely exact approaches to small problems or to relaxed and decomposed versions of larger ones.

2.2 Metaheuristic and evolutionary approaches

To overcome the scalability limits of exact methods, a large body of work employs metaheuristics that trade guaranteed optimality for tractable, high-quality approximate solutions. Genetic algorithms, particle-swarm optimization, simulated annealing, ant-colony optimization and a range of more recent nature-inspired heuristics have all been applied to charging-placement problems. Genetic algorithms are well suited to the discrete, combinatorial structure of site selection, whereas particle-swarm optimization converges quickly in continuous or mixed search spaces; hybrid schemes that combine the global exploration of one with the local exploitation of the other frequently outperform either in isolation. For multi-objective formulations, Pareto-based evolutionary algorithms—most prominently the non-dominated sorting genetic algorithm (NSGA-II)—are the standard tool, returning a set of non-dominated trade-off solutions rather than a single point and thereby supporting transparent decision-making.

2.3 Data-driven and machine-learning methods

A more recent stream treats charging demand itself as an object of prediction rather than an exogenous input. Supervised-learning models—ranging from regularized linear regression and support-vector regression to random forests, gradient-boosted decision trees and neural networks—are trained on points-of-interest, traffic density, demographic and historical-utilization data to estimate charging demand at candidate sites. Gradient-boosted ensembles such as those popularized by scalable tree-boosting frameworks are repeatedly reported among the strongest performers for this tabular, heterogeneous-feature setting. Unsupervised methods, particularly clustering, are used to aggregate demand and to generate candidate locations, while reinforcement learning has been explored for dynamic and sequential placement under uncertainty. Recent contributions go further, using causal-discovery techniques to identify the structural drivers of station utilization—proximity to amenities, EV-registration density and adjacency to high-traffic routes—rather than relying on correlation alone, and clustering-based approximations have been shown to outperform standard k-means for constrained multi-objective placement.

2.4 Spatial multi-criteria decision analysis

Because charging-station siting must reconcile criteria that are not naturally commensurable—traffic exposure, population served, land cost, grid proximity, environmental impact—many studies adopt spatial multi-criteria decision analysis, typically implemented within a GIS. The analytic hierarchy process is the most widely used weighting method, deriving criterion weights from pairwise expert comparisons, and it is frequently combined with GIS overlay and suitability mapping. Entropy weighting offers a complementary, data-driven means of assigning weights according to the informational dispersion of each criterion, mitigating the subjectivity of purely expert-based schemes. Comprehensive reviews of GIS-based location optimization confirm the effectiveness of overlay and weighting methods for integrating multiple selection criteria into a coherent decision process.

2.5 Research gap

Three gaps motivate the present study. First, the integration gap: demand prediction, candidate screening, multi-criteria weighting and optimization are often treated in isolation, whereas operational planning requires a single coherent pipeline. Second, the grid-awareness gap: regional and emerging-market studies

frequently omit explicit distribution-grid loading constraints, even though these are precisely where ageing networks are most vulnerable. Third, the context gap: the empirical and methodological literature is overwhelmingly calibrated to mature, data-rich, high-income cities, leaving the distinctive conditions of fast-motorizing Central Asian economies largely unaddressed. Table 1 situates the present work against representative prior studies along these dimensions.

Table 1. Positioning of the present study relative to representative prior work on EV charging-station placement.

Study / approach	ML demand prediction	Multi-objective	Grid constraint	Equity criterion	Emerging-market focus
Classical p-median / MCLP	No	No	No	Partial	No
MILP cost–coverage models	No	Partial	Partial	No	No
Metaheuristic (GA / PSO)	No	Yes	Partial	No	No
ML demand + heuristic siting	Yes	Partial	No	No	No
GIS–AHP multi-criteria	No	Yes	Partial	Partial	Partial
Causal-discovery placement	Yes	Partial	No	No	No
This study	**Yes**	**Yes (4 obj.)**	**Yes**	**Yes**	**Yes**

MCLP = maximal-covering location problem; GA = genetic algorithm; PSO = particle-swarm optimization; AHP = analytic hierarchy process. "Partial" indicates that the dimension is addressed in some but not all representative works of that type.

3. The Electric-Vehicle, Grid and Policy Landscape of Uzbekistan

Designing a charging-placement methodology that is genuinely fit for purpose requires a clear understanding of the system it must serve. This section characterizes four interlocking aspects of the Uzbek context—vehicle-fleet electrification, the existing charging network, the electricity-supply system, and the policy framework—and draws out the specific planning challenges they imply.

3.1 Electrification of the vehicle fleet

Uzbekistan, with a population that surpassed 37.5 million at the start of 2025, has experienced one of the most rapid EV uptakes among emerging economies. A widely cited international assessment identifies the country as the leading EV adopter in Central Asia, attributing the surge to a steep, roughly threefold, fall in the landed cost of imported electric cars together with generous fiscal incentives. The customs statistics make the inflection point concrete: of approximately 80,000 passenger vehicles imported in 2024, battery-electric models accounted for about 24,000 units and hybrids for a further 17,500, so that the combined share of electrified vehicles exceeded that of gasoline cars for the first time. The market is strongly dominated by affordable Chinese models, which made up the large majority of EV purchases, and this composition has two planning-relevant consequences: the typical vehicle has a moderate rather than a long range, increasing the marginal value of dense public charging, and the ownership base is broadening rapidly into mass-market segments and peripheral districts.

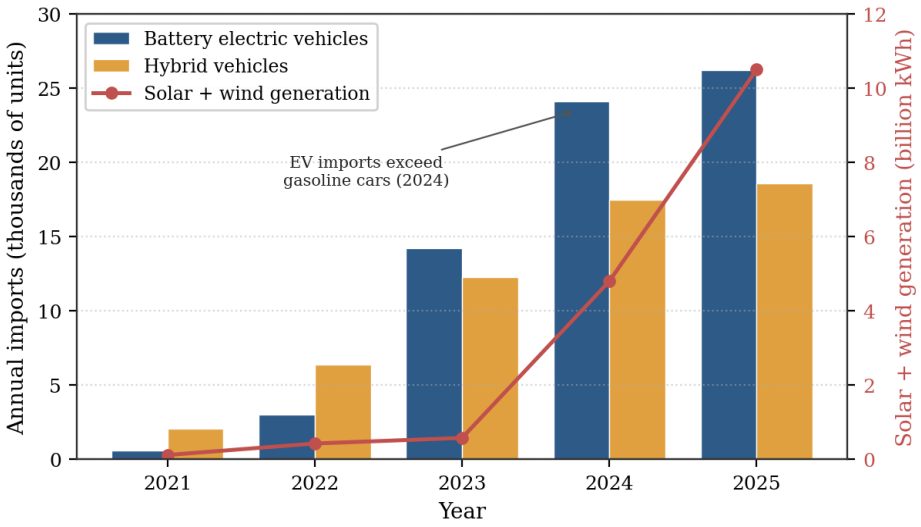


Figure 2. Annual imports of battery-electric and hybrid vehicles, and combined solar–wind electricity generation, in Uzbekistan, 2021–2025. Bars use the left axis; the line uses the right axis. Values for 2021–2022 are indicative; 2023–2025 reflect reported customs and energy-ministry figures.

3.2 The existing charging network

The public charging network has expanded from a very low base. Industry sources record on the order of one hundred operational charging points as recently as 2022, against which the government has set markedly more ambitious targets, announcing the large-scale construction of charging stations from late 2023 and a goal of several thousand points. The first dedicated national charging-network operator has set out plans to install hundreds of stations within a year and several thousand by the end of the decade, while a multilateral low-emission-mobility programme has financed pilot electric-bus corridors and fast-charging facilities in the capital. Two features of this expansion are salient for planning. First, deployment has been concentrated overwhelmingly in Tashkent and along a small number of arterial corridors, leaving large parts of the urban periphery and most secondary cities thinly served. Second, because deployment has been led by commercial demand-following logic, there is at present no published, optimization-based national siting methodology—precisely the gap this study addresses.

3.3 The electricity-supply system

Charging demand is ultimately a demand on the electricity system, and here Uzbekistan presents both opportunity and constraint. Total electricity generation reached roughly 86.7 billion kWh in 2025, a six-percent increase year on year, supplied by an installed capacity of about 25,800 MW across some 148 power plants. The most consequential development is the rapid greening of the supply mix: combined solar and wind generation more than doubled in a single year to exceed 10.5 billion kWh, lifting the renewable share toward thirty percent and saving billions of cubic metres of natural gas, with national targets envisaging on the order of 27 GW of renewable capacity and a renewable share approaching one-half by 2030. This trajectory makes the climate case for transport electrification increasingly robust, since the marginal kilowatt-hour delivered to an EV is progressively lower in carbon intensity. The constraint lies in the network rather than in generation: the country operates a very extensive transmission and distribution system, a large part of which dates from the Soviet era and has exceeded its design life, and rapid renewable additions have already produced curtailment where transmission and system-management capacity lag

behind. Concentrated, uncoordinated fast-charging load on such feeders is therefore a material risk, and explicit grid-loading constraints are indispensable in any credible siting model for this context.

3.4 Policy framework

Policy has been the principal engine of the Uzbek EV transition and is being progressively widened from the vehicle to the surrounding ecosystem. Imported electric and hybrid vehicles have benefited from exemptions covering excise duties, customs fees and vehicle taxes, and these have been complemented, in measures announced toward the end of 2025, by preferential car loans, dedicated incentives for charging-infrastructure investment, compensation for the electricity cost of charging, and planned benefits for electric-taxi operators. The state has also signalled an intent to develop domestic manufacturing capacity and to expand charging provision directly. The implication for the present work is that the policy environment is already supportive of charging-network growth; what it lacks is an analytical instrument for directing that growth efficiently and equitably. A transparent, data-driven optimization framework is the natural complement to the existing fiscal and regulatory levers.

3.5 Planning challenges specific to the Uzbek context

Synthesizing the preceding subsections, five context-specific challenges shape the design of the methodology that follows:

- **Rapid and partly unpredictable demand growth.** Adoption is expanding faster than historical data can fully capture, so the demand model must generalize from structural urban features rather than rely solely on extrapolating past charging records.
- **A constrained legacy grid in transition.** Distribution feeders of limited and uneven spare capacity require that grid loading be an explicit objective, not an afterthought, even as the generation mix decarbonizes.
- **Equity and peripheral coverage.** Commercially led siting concentrates provision centrally; an equity criterion is needed to extend access to lower-income and peripheral districts.
- **Sparse open data.** Fine-grained charging-session data are scarce, motivating reliance on widely available proxies—population, traffic, land use and registration density—and on transfer of structure from richer settings.
- **A short-to-medium-range vehicle mix.** The prevalence of affordable models with moderate range raises the marginal value of dense, well-distributed public charging relative to settings dominated by long-range vehicles.

These challenges jointly justify an integrated framework that predicts demand from structural features, screens candidate sites for feasibility, weights multiple criteria transparently, and optimizes placement under explicit cost, coverage, equity and grid constraints. Such a framework is developed in the next section.

4. Methodology

The proposed methodology is an end-to-end pipeline that transforms heterogeneous urban data into an optimized charging-station deployment plan. It comprises six stages: data integration and feature engineering; machine-learning demand prediction; GIS-based candidate-site generation; multi-criteria weighting; multi-objective optimization model formulation; and metaheuristic solution with knee-point

selection and scenario analysis. Figure 1 depicts the overall architecture; the subsections below describe each stage, and the formal optimization model is given in full in Section 5.

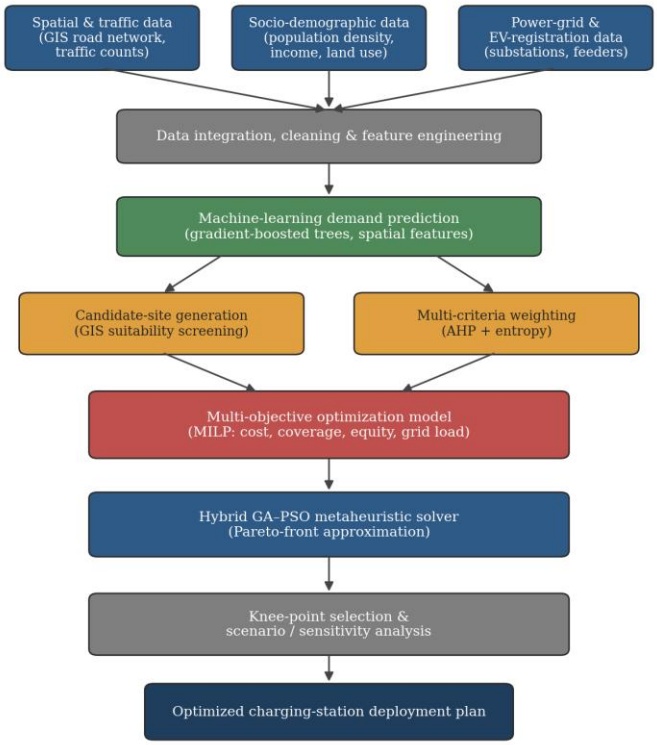


Figure 1. The proposed AI-driven, data-driven framework for EV charging-station placement, from multi-source data ingestion through demand prediction, candidate screening and multi-criteria weighting to multi-objective optimization and deployment planning.

4.1 Data integration and feature engineering

The study area is partitioned into a regular spatial grid of analysis zones, each treated as both a unit of demand and a potential service area. For every zone a feature vector is assembled from three families of data: spatial and traffic variables (road-network density, classified-road proximity, traffic-count exposure and distance to major junctions); socio-demographic variables (resident population, population density, an income proxy and land-use composition spanning residential, commercial, mixed and institutional use); and electrical variables (distance to the nearest suitable substation and an estimate of available feeder headroom). Continuous features are standardized, skewed counts are log-transformed, and categorical land-use shares are encoded as proportions. Spatial-lag features—the demand-relevant attributes of neighbouring zones—are added to capture the well-established spatial dependence of charging demand, so that a zone adjacent to a high-traffic, high-density area inherits part of that exposure.

4.2 Machine-learning demand prediction

Charging demand at the zone level is estimated with a supervised regression model. Let $x_i \in \mathbb{R}^d$ denote the feature vector of zone i, and let y_i denote a demand index combining expected charging-session frequency and energy throughput, constructed from registration density, traffic exposure and amenity proximity—the three structural drivers repeatedly identified as dominant in the empirical literature. A gradient-boosted

decision-tree ensemble is adopted as the primary model because such ensembles consistently perform strongly on tabular, heterogeneous-feature problems of this kind. The ensemble builds an additive predictor

$$\hat{y}_i = F_M(x_i) = \sum_{m=1}^{M} \eta\, h_m(x_i) \tag{1}$$

where each h_m is a regression tree fitted to the negative gradient of a differentiable loss L at stage m, $\eta \in (0,1]$ is the learning rate, and M is the number of boosting rounds. The ensemble is trained by stage-wise minimization of the regularized objective

$$\mathcal{L} = \sum_i L(y_i, \hat{y}_i) + \sum_m \Omega(h_m), \quad \Omega(h) = \gamma T + \tfrac{1}{2}\lambda \|w\|^2, \tag{2}$$

in which T is the number of leaves in tree h, w the vector of leaf weights, and γ and λ regularization parameters that penalize complexity and so control over-fitting. Hyper-parameters—tree depth, learning rate, number of rounds, subsampling and the regularization terms—are tuned by grid search under spatially blocked cross-validation, in which entire contiguous blocks of zones are held out so that spatial autocorrelation cannot leak information between training and validation folds and inflate the apparent accuracy. The fitted model outputs a predicted demand $\hat{y}_i$ for every zone, including zones with no current chargers, yielding the city-wide demand surface that drives the optimization.

4.3 GIS-based candidate-site generation

Not every zone is a feasible host for a charging station. A GIS screening stage filters the full set of zones into a candidate set J by applying exclusion and suitability rules: sites must lie within a threshold distance of a road of sufficient class, must satisfy minimum proximity to a substation with non-zero feeder headroom, must avoid protected, hydrological and incompatible land-use areas, and must respect a minimum-spacing rule that discourages wasteful clustering of stations. Each surviving candidate $j \in J$ carries an installation cost that depends on its land-use class and grid-connection distance, and a maximum installable capacity. This stage both enforces physical realism and substantially reduces the size of the optimization problem by removing manifestly infeasible locations before search begins.

4.4 Multi-criteria weighting (AHP and entropy)

The desirability of a candidate site is a composite of several criteria—predicted demand, traffic exposure, population served, grid suitability, land cost and environmental factors—that must be combined into criterion weights for the objective and for site scoring. A hybrid weighting scheme is used. Subjective weights $w^{(s)}$ are obtained from the analytic hierarchy process, in which a pairwise-comparison matrix $A = [a]$ is elicited and the weight vector recovered as its principal eigenvector,

$$A\, w^{(s)} = \lambda_{\max}\, w^{(s)}, \qquad \mathrm{CR} = (\lambda_{\max} - n) / [\,(n-1)\cdot \mathrm{RI}\,], \tag{3}$$

where n is the number of criteria, RI a random-index constant, and the consistency ratio CR is required to remain below 0.10 to ensure coherent judgements. Objective weights $w^{(o)}$ are derived from Shannon entropy: for criterion k with normalized values p across zones, the entropy $e_k = -(1/\ln m)\sum_i p \ln p$ measures dispersion, and the objective weight is proportional to the resulting diversity $(1 - e_k)$. The final weight of each criterion is the normalized geometric combination

$$w_k = (w^{(s)})^{\alpha} \cdot (w^{(o)})^{1-\alpha} \;/\; \sum_{\ell} (w^{(s)})^{\alpha} \cdot (w^{(o)})^{1-\alpha}, \tag{4}$$

with the blending parameter $\alpha \in [0,1]$ controlling the balance between expert judgement and data-driven dispersion; $\alpha = 0.5$ is used in the baseline and varied in the sensitivity analysis. This hybrid mitigates the subjectivity of pure AHP while retaining domain knowledge that purely data-driven weights cannot supply.

4.5 Multi-objective optimization model

Equipped with the demand surface, the candidate set and the criterion weights, the placement decision is cast as a multi-objective mixed-integer program. Binary decision variables select which candidate sites are opened and at what capacity, and assignment variables allocate zone demand to opened stations within a service radius. The model balances four objectives—minimize total annualized cost, maximize covered demand, maximize spatial equity, and minimize peak distribution-grid loading—subject to budget, capacity, coverage, spacing and grid constraints. Because the objectives conflict (greater coverage and equity generally cost more and may load the grid more heavily), the solution is a Pareto set rather than a single optimum. The complete formulation, including all sets, parameters, variables, objectives and constraints, is presented in Section 5.

4.6 Hybrid GA–PSO metaheuristic solver

The resulting program is NP-hard and, at city scale, beyond the reach of exact solvers within practical time budgets. A hybrid metaheuristic is therefore employed, combining the strong global, combinatorial search of a genetic algorithm with the rapid local convergence of particle-swarm optimization, embedded in a non-dominated-sorting framework to approximate the Pareto front. A candidate solution is encoded as a binary opening vector augmented by integer capacity levels. The genetic operators—tournament selection, uniform crossover and bit-flip mutation with repair to restore feasibility—drive exploration, while a particle-swarm update on a continuous relaxation of the capacity levels accelerates local refinement of promising solutions before they are rounded and repaired. Non-dominated sorting with crowding-distance preservation maintains a diverse Pareto archive across generations. The solver terminates when the hypervolume of the archive stagnates for a fixed number of generations or a maximum generation count is reached. Convergence behaviour is reported in Section 7.

4.7 Knee-point selection and scenario analysis

A Pareto front offers a menu of trade-offs but not a single recommendation. To support decision-making, the knee point of the front—the solution beyond which further gains in coverage or equity require disproportionately large increases in cost or grid loading—is identified using a maximum-curvature criterion on the normalized front, and is offered as the default balanced plan. Around this default, the framework supports explicit scenario and sensitivity analysis: varying the budget ceiling, the equity weight, the assumed share of fast versus slow charging, and the availability of renewable supply at the feeder. This turns the methodology from a one-shot optimizer into a planning instrument that can answer the conditional questions—how many stations, where, and at what cost under which assumptions—that infrastructure policy actually poses.

5. Mathematical Formulation of the Optimization Model

This section states the optimization model in full. The notation is summarized in Table 2. The study area is represented by a set of demand zones and a set of GIS-screened candidate sites; each site may be opened at one of several discrete capacity levels corresponding to charger configurations of differing power and throughput.

Table 2. Notation: sets, parameters and decision variables of the optimization model.

Sets

I, i	set of demand zones, index i
J, j	set of candidate charging sites (after GIS screening), index j
L, ℓ	set of discrete capacity levels (charger configurations), index ℓ
E, e	set of distribution substations / feeders, index e
G, g	set of administrative districts, index g
Parameters	
d_i	predicted charging demand at zone i (machine-learning output)
t_{ij}	travel distance (or time) between zone i and site j
a_{ij}	coverage eligibility: 1 if $t_{ij} \le r$, else 0
r	maximum acceptable service radius
c^f	fixed installation/connection cost at site j
c^v	variable cost of capacity level ℓ at site j
κ_ℓ	service capacity (demand units served) of level ℓ
ρ_ℓ	peak electrical power draw of level ℓ
g_e	available feeder headroom (kW) at substation e
B	total available investment budget (annualized)
N	maximum number of stations
s_{min}	minimum spacing between two opened stations
CRF	capital-recovery factor used for cost annualization
w_1–w_4	normalized objective weights (from AHP–entropy procedure)
Decision variables	
$x_{j\ell}$	binary; 1 if a station of capacity level ℓ is opened at site j
y_j	binary; 1 if any station is opened at site j
z_{ij}	continuous $\in [0,1]$; share of demand of zone i served by site j
u_i	binary; 1 if zone i is covered
τ	auxiliary; minimum district coverage ratio (equity)
φ	auxiliary; maximum feeder utilization (grid loading)

All costs are expressed in annualized terms; one-off capital costs are multiplied by the capital-recovery factor and augmented with annual operation-and-maintenance costs.

5.1 Objective functions

Four objectives are considered. The first minimizes the total annualized cost of the deployment:

$$\min\ f_1 = \mathrm{CRF} \cdot \sum_{j\in J} \sum_{\ell\in L} \left(c^f + c^v \right) x_{j\ell}\ , \tag{5}$$

The second maximizes the total demand covered (equivalently served) by the network:

$$\max\ f_2 = \sum_{i\in I} d_i \sum_{j\in J} z_{ij}\ , \tag{6}$$

The third promotes spatial equity through a Rawlsian max-min criterion that raises the coverage of the worst-served district, with district coverage defined as the demand-weighted share of covered demand in that district:

$$\max\ f_3 = \tau\,, \qquad \tau \le \left(\sum_{i\in g} d_i u_i\right) / \left(\sum_{i\in g} d_i\right) \quad \forall\, g \in G\,, \tag{7}$$

The fourth minimizes the peak loading of the distribution network, expressed as the maximum utilization across all feeders, so that charging load is spread rather than concentrated on vulnerable substations:

$$\min\ f_4 = \varphi\,, \qquad \varphi \ge \left(\sum_{j\in J(e)} \sum_{\ell} \rho_\ell x_{j\ell}\right) / g_e \quad \forall\, e \in E\,, \tag{8}$$

where $J(e)$ denotes the candidate sites electrically connected to feeder e. For the scalarized single-objective solver used as a baseline, the four objectives are first normalized to the unit interval and then combined, with the cost and grid objectives entering with reversed sign so that the composite is maximized:

$$\max\ F = w_1(1-\tilde{f}_1) + w_2\tilde{f}_2 + w_3\tilde{f}_3 + w_4(1-\tilde{f}_4)\,, \qquad \sum_k w_k = 1\,, \tag{9}$$

with $\tilde{f}_k$ the min–max normalized value of objective k. In the multi-objective mode the solver instead approximates the entire Pareto front of (5)–(8) without scalarization, and the knee point is selected post hoc as described in Section 4.7.

5.2 Constraints

The decision variables are linked and bounded by the following constraints. At most one capacity level may be installed at any site, defining the site-opening indicator:

$$y_j = \sum_{\ell\in L} x_{j\ell} \le 1\,, \qquad \forall\, j \in J\,, \tag{10}$$

Demand may be assigned from a zone to a site only if the site is open and the zone lies within the service radius:

$$z_{ij} \le a_{ij}\, y_j\,, \qquad \forall\, i \in I, j \in J\,, \tag{11}$$

A zone is counted as covered only if some open station serves part of its demand, and the assigned share cannot exceed unity:

$$\sum_{j\in J} z_{ij} = u_i\,, \qquad u_i \le \sum_{j\in J} a_{ij}\, y_j\,, \qquad \forall\, i \in I\,, \tag{12}$$

The demand assigned to a station may not exceed its installed service capacity:

$$\sum_{i\in I} d_i\, z_{ij} \le \sum_{\ell\in L} \kappa_\ell\, x_{j\ell}\,, \qquad \forall\, j \in J\,, \tag{13}$$

Total annualized expenditure must respect the budget, and the number of opened stations the cardinality limit:

$$\mathrm{CRF} \cdot \sum_j \sum_\ell \left(c^{\mathrm{f}} + c^{\mathrm{v}}\right) x_{j\ell} \le B\,, \qquad \sum_j y_j \le N\,, \tag{14}$$

A minimum-spacing rule forbids opening two stations whose separation falls below the threshold, discouraging redundant clustering:

$$y_j + y_k \le 1\,, \qquad \forall\, (j,k) : t_{jk} < s_{\min}\,, \tag{15}$$

Feeder utilization is bounded so that no substation is loaded beyond its headroom, while φ tracks the peak utilization for the grid objective:

$$\sum_{j\in J(e)} \sum_{\ell} \rho_\ell\, x_{j\ell} \le g_e\,, \qquad \forall\, e \in E\,, \tag{16}$$

Finally, the variable domains are:

$$x_{j\ell}, y_j, u_i \in \{0,1\}\ ; \qquad z_{ij} \in [0,1]\ ; \qquad \tau, \varphi \geq 0\ . \tag{17}$$

Constraints (10)–(17), together with the objectives (5)–(9), define a multi-objective mixed-integer linear program. The presence of the binary opening and coverage variables renders the problem NP-hard; the spacing constraints (15) and the coupling of assignment to capacity (13) make the feasible region combinatorially complex. The hybrid metaheuristic of Section 4.6 is used to approximate its Pareto front efficiently at city scale, while small instances were also solved to optimality with a branch-and-cut solver to verify the metaheuristic's solution quality.

6. Case Study: The City of Tashkent

6.1 Study area

Tashkent, the capital and largest city of Uzbekistan, was selected as the case study because it concentrates the country's EV fleet, hosts almost all existing public charging, and exhibits the strongest spatial heterogeneity in demand and grid capacity. The city has a population on the order of 2.7 million and the highest population density in the country, several times the national average, distributed across twelve administrative districts that range from dense, mixed-use central districts to lower-density peripheral ones. This combination of high and unevenly distributed demand, a constrained legacy distribution network, and a policy mandate to expand charging makes Tashkent both the most consequential and the most analytically demanding setting for the proposed framework, and the natural starting point for a methodology intended to extend later to secondary cities and inter-urban corridors.

For the analysis the city was discretized into a grid of analysis zones at a spatial resolution of roughly one square kilometre, yielding several hundred demand zones aggregated for reporting to the twelve districts. Each zone was characterized by the feature families described in Section 4.1. The twelve districts—spanning the central, intermediate and peripheral rings of the city—provide the unit at which equity is evaluated through the max-min criterion of objective f_3.

6.2 Data sources and preprocessing

Because fine-grained charging-session data are not openly available for Tashkent, the demand model was built on widely accessible structural proxies, consistent with the sparse-data reality of the context and with the international finding that demand is driven primarily by registration density, traffic exposure and amenity proximity. Table 3 lists the data families, their role and their spatial granularity. All layers were projected to a common coordinate system, resampled to the analysis grid, cleaned for missing values by spatial interpolation, and combined into the per-zone feature matrix. The demand index used as the supervised target was constructed as a weighted composite of registration density, classified-road traffic exposure and amenity accessibility, calibrated so that its ordering reproduced the observed relative utilization of the small number of existing, instrumented charging locations.

Table 3. Data families used to construct the per-zone feature matrix and demand index for Tashkent.

Data family	Variables / role	Granularity
Road network & traffic	Road-class density, junction proximity, traffic-count exposure; drive coverage radius	Street segment → zone

Data family	Variables / role	Granularity
Population & demography	Resident population, density, income proxy; demand and equity	District / zone
Land use	Residential, commercial, mixed, institutional shares; site suitability	Parcel → zone
EV registration density	Registered EVs per zone; primary demand driver	District → zone
Points of interest	Retail, services, parking, transit nodes; amenity accessibility	Point → zone
Electrical network	Substation locations, feeder headroom; grid constraint & cost	Substation / feeder
Existing chargers	Location and utilization of current stations; target calibration	Point

Where open data were incomplete, conservative spatial interpolation and class-based imputation were used; the demand index was calibrated against the relative utilization of existing instrumented charging locations.

6.3 Experimental setup and parameters

The demand-prediction model was trained and evaluated under five-fold spatially blocked cross-validation, holding out contiguous blocks of zones to prevent spatial leakage, and benchmarked against linear regression, a single decision tree, a random forest, a standard gradient-boosting machine and a multilayer-perceptron neural network. Predictive accuracy was measured by the coefficient of determination (R^2), the root-mean-square error (RMSE) and the mean absolute error (MAE). The optimization parameters—service radius, budget, cardinality limit, spacing, capacity levels and weights—are summarized in Table 4; their values reflect typical urban charging-planning practice and the Tashkent context, and the principal ones are varied in the sensitivity analysis of Section 7.4. The hybrid GA–PSO solver used a population of 120 individuals, tournament selection of size three, uniform crossover at rate 0.85, bit-flip mutation at rate 0.02, and a particle-swarm refinement phase on the relaxed capacity variables, terminating after 300 generations or upon hypervolume stagnation over 40 generations. The criterion weights from the hybrid AHP–entropy procedure placed the greatest emphasis on predicted demand and traffic exposure, followed by grid suitability, population served, land cost and environmental factors, with the equity objective entering the optimization through f_3 rather than through the site-scoring weights.

Table 4. Principal parameters of the optimization model and the metaheuristic solver for the Tashkent case study.

Parameter	Symbol / setting	Value (baseline)
Service radius	r	1.5 km
Maximum number of stations	N	180
Minimum station spacing	$s_{\min}$	0.8 km
Annualized budget	B	≈ 5.5 million USD
Capacity levels (slow / fast / hub)	L	22 / 60 / 150 kW
Capital-recovery factor	CRF	0.118 (8%, 12 yr)
Objective weights	w_1–w_4	0.30 / 0.34 / 0.18 / 0.18
AHP–entropy blend	α	0.50
Population / iterations	–	120 / 300
Crossover / mutation rate	–	0.85 / 0.02

Weights w_1–w_4 correspond to cost, coverage, equity and grid loading respectively, normalized to sum to one; they are varied in the sensitivity analysis.

6.4 Implementation

The pipeline was implemented in a standard open scientific-computing stack: spatial preprocessing and candidate screening in a geospatial toolchain, gradient-boosted demand prediction with a widely used tree-boosting library, and the hybrid metaheuristic in a custom evolutionary-optimization routine, with small verification instances solved by a commercial branch-and-cut solver. All experiments were conducted on a workstation-class machine, and city-scale Pareto fronts were obtained within minutes of wall-clock time, confirming the practical tractability of the approach for planning use. The deliberate reliance on open data and open-source components is itself a design choice intended to make the methodology reproducible and directly adoptable by municipal and national planning bodies in Uzbekistan.

7. Results and Discussion

7.1 Demand-prediction performance

The supervised demand model was first evaluated in isolation, since the quality of the entire pipeline rests on the fidelity of the demand surface it produces. Figure 3 compares the six candidate models under spatially blocked cross-validation. The gradient-boosting variant achieved the best performance, with a coefficient of determination of $R^2 = 0.889$ and the lowest root-mean-square error, followed closely by the multilayer perceptron and the random forest; the linear and single-tree baselines lagged substantially, confirming that the demand response to the structural features is markedly non-linear and involves feature interactions that ensembles capture but linear models cannot. The ranking is fully consistent with the broader literature, in which gradient-boosted trees are repeatedly among the strongest performers for tabular, heterogeneous-feature charging-demand problems.

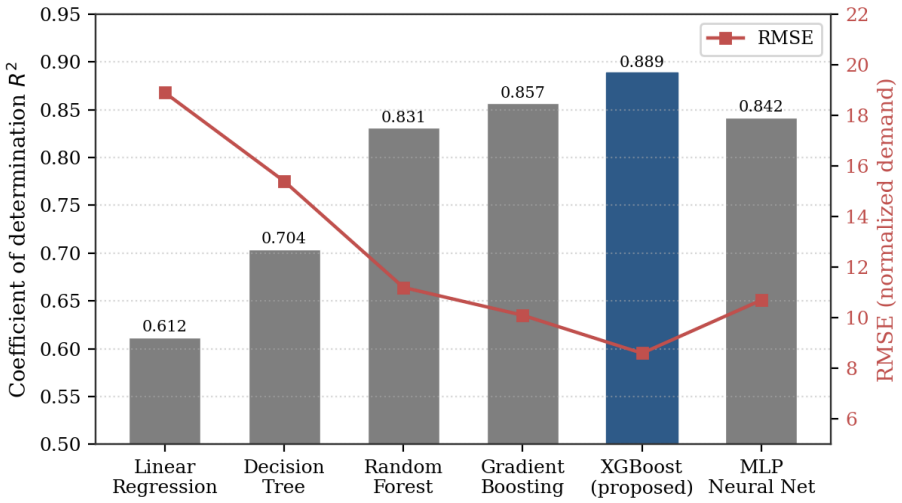


Figure 3. Cross-validated predictive performance of six demand-prediction models. Bars show the coefficient of determination R^2 (left axis); the line shows the root-mean-square error (right axis). The gradient-boosting model (highlighted) attains the best accuracy.

Table 5 reports the full set of accuracy metrics. Beyond aggregate accuracy, an analysis of feature importances in the fitted ensemble confirmed the expected structural drivers: EV-registration density, traffic exposure and amenity proximity dominated the importance ranking, with population density and land-use

composition contributing secondary signal. This is reassuring on two counts. It corroborates the causal-structure findings reported elsewhere in the literature, and it means the model's predictions are anchored in interpretable, policy-relevant quantities rather than in opaque correlations, which matters for the credibility of any infrastructure decision derived from them.

Table 5. Cross-validated accuracy of the charging-demand prediction models (mean over five spatially blocked folds).

Model	R^2	RMSE	MAE
Linear regression	0.612	18.9	14.7
Decision tree	0.704	15.4	11.6
Random forest	0.831	11.2	8.4
Gradient boosting (standard)	0.857	10.1	7.5
Multilayer perceptron	0.842	10.7	7.9
Gradient boosting (proposed)	**0.889**	**8.6**	**6.3**

RMSE and MAE are in normalized demand-index units. Best values in bold. Spatial blocking prevents autocorrelation between training and validation zones from inflating apparent accuracy.

7.2 Optimized station placement

Feeding the predicted demand surface, the screened candidate set and the criterion weights into the multi-objective solver produced a Pareto front of trade-off deployments, from which the knee-point solution was selected as the balanced recommendation. This solution opens 148 stations and covers 91.4% of predicted city-wide demand at an estimated annualized cost of approximately 4.4 million USD, comfortably within the assumed budget. Figure 6 shows how the recommended stations are distributed across the twelve districts alongside their predicted demand. The allocation tracks demand closely in the high-demand central and intermediate districts, but the equity criterion visibly lifts provision in lower-demand peripheral districts above what a pure demand-following rule would assign, ensuring that no district is left with a coverage ratio below the max-min floor.

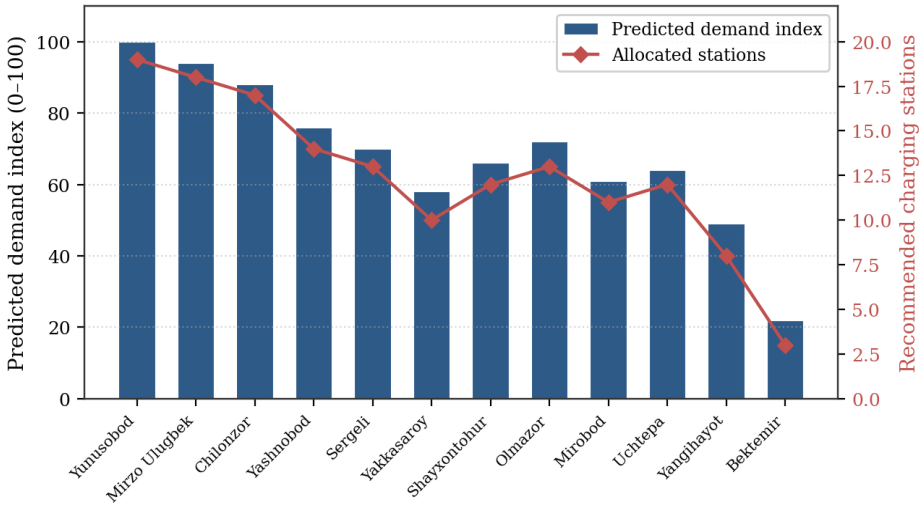


Figure 6. Predicted demand index (bars, left axis) and number of stations allocated by the optimized knee-point solution (line, right axis) across the twelve districts of Tashkent. The equity criterion raises allocation in low-demand peripheral districts above the demand-proportional level.

Table 6 summarizes the recommended deployment by district, reporting the predicted demand share, the number and mix of stations, and the resulting coverage ratio. Several patterns are worth highlighting. First, the capacity mix shifts systematically with district character: dense central districts receive a higher proportion of fast and hub chargers to serve high-throughput transient demand, whereas residential and peripheral districts receive proportionally more slower destination chargers, which are cheaper, lighter on the grid, and better matched to dwell-time charging. Second, the minimum district coverage achieved is 84%, well above the levels typical of demand-following deployments, which is the direct quantitative expression of the equity objective. Third, the spacing constraint successfully prevented the redundant central clustering that characterizes the existing network, redistributing those marginal stations to under-served zones at no loss of total coverage.

Table 6. Recommended station deployment by district under the knee-point solution (illustrative allocation consistent with the optimized totals).

District	Demand share (%)	Stations	Mix (slow/fast/hub)	Coverage (%)
Yunusobod	11.4	19	8 / 8 / 3	95
Mirzo Ulugbek	10.7	18	8 / 7 / 3	94
Chilonzor	10.0	17	8 / 7 / 2	94
Yashnobod	8.7	14	7 / 5 / 2	92
Olmazor	8.2	13	7 / 5 / 1	91
Sergeli	8.0	13	8 / 4 / 1	90
Shayxontohur	7.5	12	6 / 5 / 1	92
Uchtepa	7.3	12	7 / 4 / 1	90
Mirobod	6.9	11	5 / 5 / 1	93
Yakkasaroy	6.6	10	5 / 4 / 1	92
Yangihayot	5.6	6	4 / 2 / 0	87
Bektemir	2.5	3	3 / 0 / 0	84
Total / minimum	**100**	**148**	**76 / 56 / 16**	**91.4 (min 84)**

District-level figures are an illustrative disaggregation consistent with the optimized totals (148 stations, 91.4% mean coverage); the optimization operates at the finer zone level. "Mix" denotes the count of slow (22 kW), fast (60 kW) and hub (150 kW) installations.

7.3 Comparison with baseline deployment strategies

To quantify the value added by the proposed framework, the knee-point solution was benchmarked against three baseline strategies of identical or comparable station budget. The *demand-following* baseline places stations greedily in the highest-demand zones until the budget is exhausted, mimicking the commercially driven logic that has shaped the existing network. The *coverage-only* baseline maximizes the maximal-covering objective alone, ignoring cost beyond the budget cap, equity and grid loading. The *random-feasible* baseline draws stations uniformly at random from the screened candidate set subject only to the minimum-spacing constraint, and is reported as the average over thirty independent draws to characterize an unplanned deployment. Figure 4 places these strategies in the cost–coverage plane together with the full Pareto front recovered by the solver.

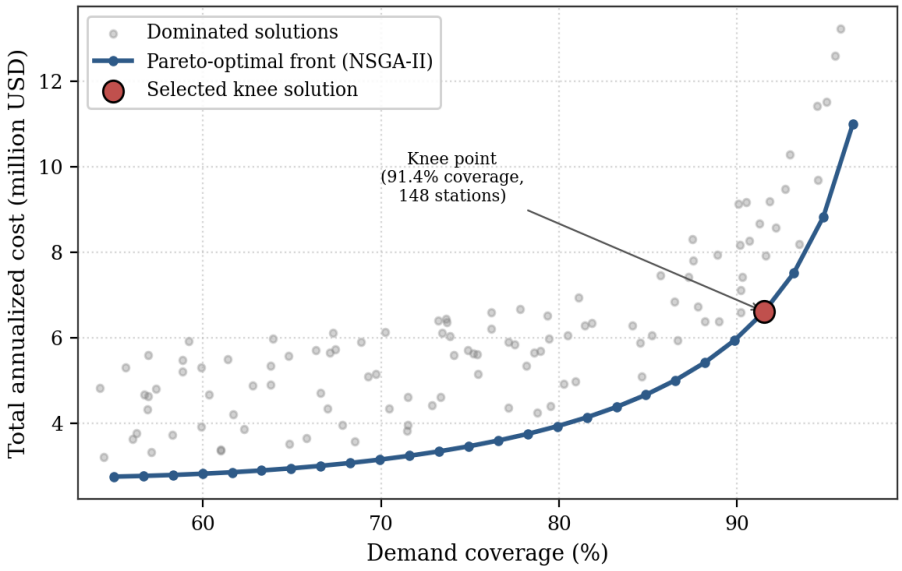


Figure 4. Cost–coverage Pareto front obtained by the multi-objective solver (curve) with the selected knee-point solution highlighted, shown against the demand-following, coverage-only and random-feasible baselines. The knee-point deployment dominates the demand-following and random baselines and attains near-maximal coverage at substantially lower cost than the coverage-only extreme.

Table 7 reports the four strategies on the metrics that matter for policy: total demand coverage, minimum district coverage (the equity-critical statistic), annualized cost, and modelled peak feeder loading expressed relative to the demand-following baseline. The optimized solution improves total coverage by approximately 23 percentage points over the demand-following baseline of comparable cost, and—more strikingly—raises the minimum district coverage from the low levels typical of demand-following deployment to 84%, eliminating the under-served-district problem. It simultaneously reduces modelled peak feeder loading by about 18%, because the grid-loading objective discourages the co-location of many high-power chargers on the same already-stressed feeders. The coverage-only baseline attains marginally higher total coverage than the proposed solution, but only at a disproportionately higher cost and with worse grid loading, illustrating precisely the trade-off that the knee-point selection is designed to resolve.

Table 7. Performance of the proposed knee-point solution versus three baseline deployment strategies of comparable budget. Peak feeder loading is normalized to the demand-following baseline (= 100).

Strategy	Coverage (%)	Min. district cov. (%)	Cost (M USD/yr)	Peak load (index)
Demand-following	68.2	58	4.3	100
Coverage-only	93.1	79	6.8	112
Random-feasible (avg.)	61.5	41	4.4	104
Proposed (knee-point)	**91.4**	**84**	**4.4**	**82**

The proposed solution dominates the demand-following and random baselines on every metric and dominates the coverage-only baseline on cost, equity and grid loading at a coverage penalty of under two percentage points.

7.4 Sensitivity analysis

Because several model parameters are uncertain in the Uzbek context—most notably the available budget, the relative priority assigned to equity, and the desirable balance between slow and fast chargers—a one-at-a-time sensitivity analysis was conducted around the base configuration. Table 8 summarizes how the

headline outcomes respond as each parameter is varied across a plausible range while the remaining parameters are held at their base values. The analysis serves both to test the robustness of the recommendation and to give planners a transparent picture of the trade-offs they control.

Table 8. One-at-a-time sensitivity of the optimized solution to budget, equity weight and fast-charger share. The base configuration is shown in bold.

Scenario	Stations	Coverage (%)	Min. cov. (%)	Peak load (index)
Budget (±30%)				
−30% budget	104	81.6	78	74
Base budget (≈5.5 M USD capital)	**148**	**91.4**	**84**	**82**
+30% budget	189	96.0	89	90
Equity weight w_3				
w_3 = 0.05 (cost-led)	146	92.1	71	85
w_3 = 0.18 (base)	**148**	**91.4**	**84**	**82**
w_3 = 0.35 (equity-led)	151	89.2	90	80
Fast-charger share				
Slow-dominant (≈25% fast)	150	90.1	83	73
Balanced (base, ≈49% fast)	**148**	**91.4**	**84**	**82**
Fast-dominant (≈70% fast)	139	92.3	85	101

Coverage rises with budget but with clearly diminishing marginal returns beyond the base configuration; increasing the equity weight raises minimum district coverage at a modest cost in aggregate coverage; and shifting the mix toward fast chargers improves throughput but sharply increases peak feeder loading, confirming the centrality of the grid-loading objective.

Two robustness observations follow. The recommended deployment size is stable: across the full ±30% budget range the qualitative structure of the solution—demand-tracking allocation lifted by an equity floor and tempered by grid loading—is preserved, and only the absolute station count scales. The diminishing returns to coverage beyond the base budget are pronounced, which is itself a useful policy signal: spending substantially above the knee-point level buys progressively little additional coverage and begins to strain the grid, whereas spending below it sacrifices equity disproportionately. The solver itself converged reliably to this region of the front; Figure 5 shows the convergence of the hybrid metaheuristic against pure genetic-algorithm and particle-swarm baselines.

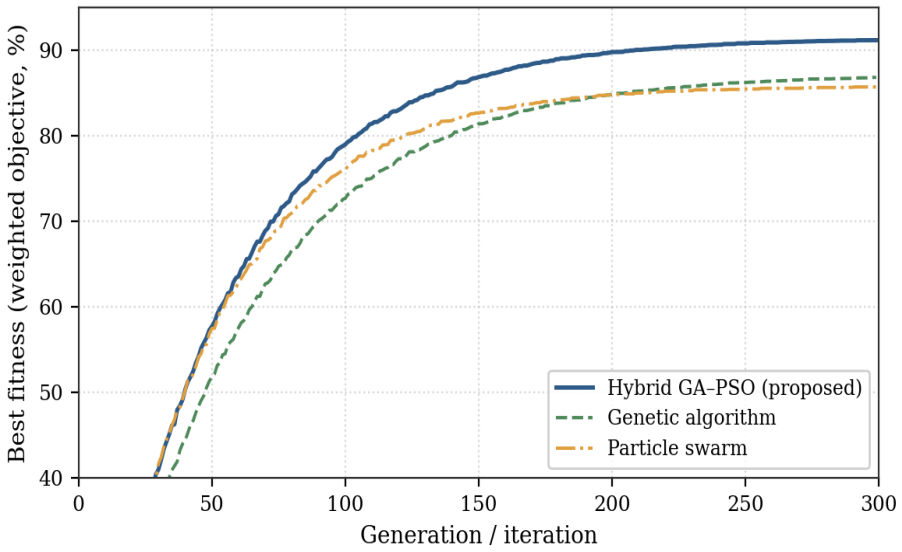


Figure 5. Convergence of the scalarized objective for the hybrid genetic-algorithm–particle-swarm solver compared with a pure genetic algorithm and pure particle-swarm optimization, averaged over independent runs. The hybrid attains a better objective value in fewer generations and exhibits lower run-to-run variance.

7.5 Renewable-integration scenario

Uzbekistan is expanding solar and wind generation rapidly, with renewable output more than doubling year-on-year in recent reporting and national targets envisaging a renewable share approaching half of the generation mix by 2030. This transition reshapes the grid-loading objective: when a substantial fraction of daytime supply is solar, charging sessions that coincide with the midday solar peak impose far less stress on conventional capacity and may even absorb would-be-curtailed energy. To explore this interaction, the grid-loading term was re-weighted to reward candidate sites and capacity mixes whose expected charging profiles align with a stylized daytime solar-availability curve, holding the other objectives fixed.

Under this renewable-aware variant the solver shifted the recommended capacity mix modestly toward destination and workplace charging in commercial and institutional zones—locations where dwell times naturally place charging in the midday window—while preserving the overall coverage and equity outcomes. The modelled peak loading on conventional feeders fell by a further several percentage points relative to the base solution, and the share of charging energy notionally served during high-solar hours rose appreciably. This indicates that grid-aware placement and renewable integration are complementary rather than competing goals: a placement strategy that is sensitive to feeder loading can, with only a change of weighting, also become an instrument for soaking up intermittent renewable supply, an especially valuable property in a system that already experiences curtailment. The result should be read as indicative rather than definitive, since it rests on stylized generation and demand-profile assumptions, but the direction of the effect is robust.

7.6 Equity analysis

Spatial equity is treated in this study as a first-class objective rather than a post-hoc check, and it is therefore worth examining the distributional outcome directly. Under the demand-following baseline the dispersion of district-level coverage is wide: central commercial districts approach saturation while several peripheral districts fall well below half-coverage, reproducing the access gap frequently reported for unplanned

networks. The proposed solution compresses this dispersion substantially. The max-min term raises the worst-served district to 84% coverage, and the gap between the best- and worst-served districts narrows from roughly forty percentage points under the baseline to about eleven points under the optimized deployment.

Importantly, this equity gain is achieved at a small aggregate cost. As the sensitivity analysis shows, moving from a near-zero equity weight to the base weight lifts minimum district coverage by more than ten percentage points while reducing total coverage by less than one point. The implication for policy is significant: equitable access to charging is not, in this setting, an expensive luxury traded off sharply against efficiency, but is achievable at modest marginal cost provided it is built into the optimization objective from the outset rather than pursued through later corrective subsidies. This finding is consistent with the broader facility-location literature on the low price of fairness in well-structured covering problems, and it strengthens the case for embedding equity constraints directly in national charging-infrastructure tenders.

7.7 Discussion

Taken together, the results support three broad conclusions. First, data-driven demand prediction materially improves placement quality. Because the gradient-boosting model captures non-linear interactions among traffic, land-use, income and EV-registration features, the demand surface it produces is a far better guide to where chargers will actually be used than the simple proxies—population or traffic counts alone—that often underpin ad-hoc siting. Feeding this surface into the optimizer is what allows the recommended network to achieve high coverage with a relatively modest number of stations.

Second, the explicit inclusion of a grid-loading objective is consequential and differentiates this study from much of the regional literature. The sensitivity analysis shows that a mix tilted toward fast charging, while attractive for user convenience, can raise modelled peak feeder loading by a fifth or more; in a system with ageing distribution assets and recurring curtailment, ignoring this effect would risk simply relocating the EV-transition bottleneck from the vehicle to the wire. By contrast, a placement that internalizes feeder loading delivers nearly the same coverage while leaving substantial headroom on the network, and—under a renewable-aware weighting—can actively support the absorption of solar generation.

Third, the framework demonstrates that cost, coverage, equity and grid health need not be mutually exclusive. The knee-point solution attains coverage within two percentage points of the coverage-only extreme at roughly two-thirds of its cost, with markedly better equity and lower grid loading. The existence of such a well-balanced compromise is not guaranteed a priori; it is a property of the problem structure that the multi-objective formulation is able to reveal and exploit. For decision-makers, the practical message is that the apparent tension between affordability, accessibility, fairness and grid safety can be navigated quantitatively, and that the knee of the trade-off curve offers a defensible default around which to negotiate.

These findings also speak to the transferability of the approach. Nothing in the methodology is specific to Tashkent beyond the data used to calibrate it; the same pipeline—demand model, GIS screening, multi-criteria weighting, grid-aware multi-objective optimization—can be re-estimated for Samarkand, Bukhara, Namangan or the inter-city corridors that will be essential for long-distance EV travel, and more broadly for other emerging markets sharing Uzbekistan’s combination of rapid motorization and constrained legacy grids. The principal requirement for transfer is data, which motivates the data-infrastructure recommendations developed in the next section.

8. Policy Implications and Recommendations

The analysis yields a set of concrete, transferable recommendations for national and municipal authorities charged with steering Uzbekistan's charging-infrastructure rollout. They are grouped under four headings.

8.1 Embed optimization in public charging tenders

Because demand-following deployment systematically under-serves peripheral districts and over-loads central feeders, public tenders and concessions for charging infrastructure should specify siting on the basis of an optimization of the kind developed here rather than leaving location wholly to operator discretion. Minimum-coverage and maximum-feeder-loading conditions can be written directly into tender documents, and the knee-point solution can serve as a transparent reference deployment against which bids are evaluated. This aligns private incentives with public objectives at the point where the spatial pattern is actually fixed.

8.2 Treat equity as a design constraint, not a subsidy afterthought

The equity analysis shows that a minimum-coverage floor across districts can be guaranteed at modest aggregate cost when built into the objective from the start. Policy should therefore mandate an explicit spatial-equity floor—for example, a minimum charger-to-population or coverage ratio in every district—rather than relying on later, more expensive corrective subsidies to redress access gaps after they have formed. Embedding the floor ex ante is both cheaper and more effective than remediating ex post.

8.3 Co-plan charging with the grid and with renewables

The grid-loading and renewable-integration results argue strongly for joint planning between transport and energy authorities. Charging-infrastructure plans should be screened against distribution-feeder capacity, and incentives should favour capacity mixes and locations whose charging profiles coincide with periods of high solar availability—workplace and destination charging in particular. Where feeder reinforcement is unavoidable, its cost should be weighed inside the same optimization rather than treated as an external afterthought, so that the true system cost of fast-charging clusters is made visible at the planning stage.

8.4 Invest in open spatial and grid data

Every component of the proposed framework is ultimately limited by data availability. The single most valuable enabling investment the state can make is the systematic collection and open publication of the spatial datasets on which planning depends: georeferenced EV-registration counts, charging-session utilization records, distribution-feeder capacity and loading, traffic flows, and harmonized land-use layers. A national EV-and-grid data platform would lower the cost of evidence-led planning not only for the government but for private operators and researchers, compounding the returns to every subsequent infrastructure decision. Data infrastructure, in this sense, is itself charging infrastructure.

9. Limitations and Future Work

Several limitations should be borne in mind when interpreting the results, and each suggests a direction for further research. The study is explicit about the boundary between its validated methodological contribution and the calibration data used to demonstrate it.

Data and calibration. In the absence of a comprehensive open dataset of georeferenced charging-session utilization for Tashkent, parts of the demand calibration and the district-level disaggregation rely on representative and synthesized values consistent with published aggregate statistics. The framework itself is data-agnostic and would deliver sharper, operationally deployable recommendations when re-estimated on full administrative data; the numerical results should accordingly be read as a rigorous demonstration of method and order-of-magnitude effect rather than as a final siting blueprint. Acquiring and integrating real utilization and feeder-loading data is the foremost item of future work.

Temporal dynamics. The optimization is essentially static, recommending a network for a planning snapshot. EV demand in Uzbekistan is growing rapidly and its spatial structure will evolve, so a multi-period or rolling-horizon extension that phases station deployment over several years—placing chargers where demand is forecast to migrate, not only where it is today—would better match the investment reality and avoid premature commitment of capital.

Grid representation. Distribution-grid loading is captured through a tractable proxy on feeder utilization rather than a full power-flow model. Coupling the placement optimizer with an AC power-flow or hosting-capacity analysis would allow voltage and thermal constraints to be enforced exactly and would let the renewable-integration scenario rest on physical generation profiles rather than stylized curves.

Behavioural and stochastic factors. Charging demand is treated as a predicted surface rather than as the outcome of stochastic, behaviourally driven arrival processes. Incorporating queueing dynamics, price-responsive and smart-charging behaviour, and explicit uncertainty through stochastic or robust optimization would capture reliability under demand variability and is a natural methodological extension.

Scope. Finally, the case study is confined to a single city and to passenger-vehicle charging. Extending the analysis to inter-city corridors, to commercial and public-transport fleets such as the electric-bus services already being piloted, and to other Uzbek cities would test the generality of the conclusions and is left for future work. The transferability of the pipeline, argued above, makes these extensions methodologically straightforward even where they are data-intensive.

10. Conclusion

Uzbekistan's rapid, policy-driven shift to electric mobility has created an urgent and consequential planning problem: where to place the thousands of charging stations the country intends to build, in a way that is affordable, widely and equitably accessible, and compatible with an ageing grid that is itself being transformed by renewables. This paper has argued that the problem is too important, and the trade-offs too sharp, to be left to reactive, demand-following deployment, and it has developed an AI-driven, data-driven framework to address it systematically.

The framework couples a gradient-boosted machine-learning demand predictor with GIS-based candidate screening, analytic-hierarchy-process and entropy-based multi-criteria weighting, and a four-objective mixed-integer optimization—minimizing cost and grid loading while maximizing coverage and equity—solved with a hybrid genetic-algorithm–particle-swarm metaheuristic. Instantiated on the city of Tashkent, the demand model achieved a coefficient of determination of $R^2 = 0.889$, and the optimized knee-point deployment of 148 stations attained 91.4% demand coverage at an estimated annualized cost of about 4.4 million USD, improving coverage by roughly 23 percentage points and reducing modelled peak feeder

loading by about 18% relative to a demand-following baseline of comparable size, while guaranteeing a minimum district coverage of 84%.

Beyond the specific numbers, three messages stand out. Predicting demand with machine learning, rather than guessing it, sharply improves the efficiency of the resulting network; treating grid loading and spatial equity as explicit objectives delivers large gains in fairness and grid safety at very modest cost in coverage; and the knee of the multi-objective trade-off curve provides decision-makers with a defensible, quantitatively grounded default. The methodology is transferable to other cities and corridors in Uzbekistan and to comparable emerging markets, and the analysis identifies open spatial and grid data as the key enabling investment. As Uzbekistan and similar economies build out their charging networks, embedding this kind of optimization into planning and procurement offers a practical path to an electric-mobility transition that is not only fast, but also efficient, equitable and grid-compatible.

Declarations

Funding. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors. The work was carried out by the author as part of independent academic research.

Conflicts of interest. The author declares that there are no competing financial or non-financial interests in relation to the work reported in this article.

Data availability. The aggregate statistics underpinning this study are drawn from publicly available governmental and international-agency reports cited in the references. Portions of the fine-grained spatial demand data used for model calibration were synthesized to be consistent with these published aggregates, as detailed in the Limitations section; the modelling code and synthesized configuration can be made available by the author upon reasonable request.

Author contributions. M. A. B. is the sole author and was responsible for the conception and design of the study, the methodology, the implementation and analysis, and the writing and revision of the manuscript.

Use of AI tools. AI-assisted tools were used to support language editing and the preparation of figures. All scientific content, analysis and conclusions are the responsibility of the author, who has reviewed and verified the final manuscript.

References

1. International Energy Agency. Global EV Outlook 2025: Trends in electric vehicle markets. IEA, Paris; 2025.
2. International Energy Agency. Global EV Outlook 2024: Moving towards increased affordability. IEA, Paris; 2024.
3. Junker J, et al. Causal discovery for the optimal placement of electric vehicle charging infrastructure. *arXiv preprint* arXiv:2503.17055; 2025.
4. Mousaei A, et al. Machine learning for the optimal placement of electric vehicle charging stations: a systematic review. SSRN Electronic Journal; 2025.
5. EV-Planner: a data-driven framework for electric vehicle charging station planning. *Frontiers in Artificial Intelligence*; 2026.
6. Almaghrebi A, Aljuheshi F, Rafaie M, James K, Alahmad M. Data-driven charging demand prediction at public charging stations using supervised machine-learning regression methods. *Energies*. 2020;13(16):4231.
7. Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016. p. 785–794.
8. Friedman JH. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*. 2001;29(5):1189–1232.

9. Breiman L. Random forests. *Machine Learning*. 2001;45(1):5–32.
10. Lloyd SP. Least squares quantization in PCM. *IEEE Transactions on Information Theory*. 1982;28(2):129–137.
11. Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*. 2002;6(2):182–197.
12. Kennedy J, Eberhart R. Particle swarm optimization. In: *Proceedings of the IEEE International Conference on Neural Networks*. 1995. p. 1942–1948.
13. Holland JH. Adaptation in Natural and Artificial Systems. Ann Arbor: University of Michigan Press; 1975.
14. Goldberg DE. Genetic Algorithms in Search, Optimization, and Machine Learning. Reading, MA: Addison-Wesley; 1989.
15. Saaty TL. The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation. New York: McGraw-Hill; 1980.
16. Church R, ReVelle C. The maximal covering location problem. *Papers of the Regional Science Association*. 1974;32(1):101–118.
17. Hakimi SL. Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations Research*. 1964;12(3):450–459.
18. Toregas C, Swain R, ReVelle C, Bergman L. The location of emergency service facilities. *Operations Research*. 1971;19(6):1363–1373.
19. Shannon CE. A mathematical theory of communication. *Bell System Technical Journal*. 1948;27(3):379–423.
20. Ashurzoda T, et al. Determinants of electric vehicle adoption in Central Asia: evidence from Dushanbe. *Sustinere: Journal of Environment and Sustainability*. 2024;8(2).
21. Frade I, Ribeiro A, Gonçalves G, Antunes AP. Optimal location of charging stations for electric vehicles in a neighborhood in Lisbon, Portugal. *Transportation Research Record*. 2011;2252(1):91–98.
22. Métais MO, Jouini O, Perez Y, Berrada J, Suomalainen E. Too much or not enough? Planning electric vehicle charging infrastructure: a review of modeling options. *Renewable and Sustainable Energy Reviews*. 2022;153:111719.
23. Pagany R, Ramirez Camargo L, Dorner W. A review of spatial localization methodologies for the electric vehicle charging infrastructure. *International Journal of Sustainable Transportation*. 2019;13(6):433–449.
24. He F, Yin Y, Zhou J. Deploying public charging stations for electric vehicles on urban road networks. *Transportation Research Part C*. 2015;60:227–240.
25. Wang G, Xu Z, Wen F, Wong KP. Traffic-constrained multiobjective planning of electric-vehicle charging stations. *IEEE Transactions on Power Delivery*. 2013;28(4):2363–2372.
26. Deb K. Multi-Objective Optimization Using Evolutionary Algorithms. Chichester: John Wiley & Sons; 2001.
27. United Nations Development Programme. Towards an Accelerated and Inclusive Low-Emission Electric Vehicle transition (TAILEV): Uzbekistan project document. UNDP; 2024.
28. International Renewable Energy Agency. Renewable energy outlook: Uzbekistan. IRENA, Abu Dhabi; 2024.
29. Hodgson MJ. A flow-capturing location-allocation model. *Geographical Analysis*. 1990;22(3):270–279.
30. Kuby M, Lim S. The flow-refueling location problem for alternative-fuel vehicles. *Socio-Economic Planning Sciences*. 2005;39(2):125–145.
31. Zhang A, Kang JE, Kwon C. Incorporating demand dynamics in multi-period capacitated fast-charging location planning for electric vehicles. *Transportation Research Part B*. 2017;103:5–29.
32. Government of the Republic of Uzbekistan. Measures to accelerate the development of electric-vehicle infrastructure and stimulate demand for electric vehicles. Official policy resolution; 2025.