

A Comprehensive Framework for Municipal Solid Waste Forecasting and Route Optimization: A Comparative Analysis of Different Machine Learning Models for the City of Surat

A Synopsis of the Ph.D. Thesis submitted to
Gujarat Technological University, Ahmedabad
for the Degree of
Doctor of Philosophy
in
Civil Engineering



by
Rathod Abhijit Ramjibhai
Enrollment No.: 209999912005

October 2026

Supervisor
Dr. Vinodkumar Manilal Patel

**Professor, Civil Engineering,
S. S. Engineering College, Bhavnagar, Gujarat, India**

Doctoral Progress Committee (DPC) Members:

Dr. Umesh Kumar Khare

Professor in Civil Engineering,
Government Engineering College, Modasa

Dr. Yogeshkumar Shankarlal Patel

Professor in Civil Engineering,
Sankalchand Patel University, Visnagar

Contents

1. Abstract	1
2. Brief Description and Introduction	2
3. Need, Objectives and Scope of Research	4
4. Literature Review and Research Gaps	7
5. Study Area and Data Collection	10
6. Research Methodology, Data Analysis and Results	13
6.1 Overall Research Design	13
6.2 Feature Engineering Framework	15
6.3 Forecasting Framework (RO-1).....	15
6.4 Route Optimisation Framework (RO-2).....	20
6.5 Decision Support System Framework (RO-3).....	24
6.6 Integrated Framework Synthesis, Comparative Positioning, and Robustness.....	26
7. Major Findings, Contributions and Validation	29
7.1 Objective-wise Findings	29
7.2 Contributions of the Research.....	30
7.3 Validation.....	31
8. Conclusions, Limitations and Future Recommendations	31
8.1 Conclusions.....	31
8.2 Recommendations for Municipal Authorities.....	33
8.3 Limitations	33
8.4 Future Scope	34
9. List of Papers Published / Presented.....	35
10. References.....	35

1. Abstract

Indian cities are urbanising faster than their municipal solid waste (MSW) systems can absorb. Collection infrastructure lags behind generation rates, forecasting methods remain poorly adapted to the data conditions found in Indian municipalities, and route optimisation studies routinely assume levels of data availability and GPS accuracy that most Urban Local Bodies (ULBs) simply do not have. This synopsis sets out a three-part framework covering forecasting, route optimisation, and decision support, built and tested on data from the Surat Municipal Corporation (SMC), Gujarat, and designed to close all three gaps within a single empirically grounded programme of research.

The forecasting component tests seventeen models, drawn from linear and regularised methods, kernel methods, ensembles, deep learning, and classical time-series approaches, against a feature pipeline of 419 engineered variables built from 186 months of SMC waste generation records (June 2009 to December 2024). Across an eight-metric evaluation protocol, Ordinary Least Squares regression produced the strongest overall result ($R^2 = 0.9972$, $\text{MAPE} = 1.43\%$), outperforming every deep learning architecture tested, including LSTM, under identical tabular conditions. A second analysis, an interrupted time-series study on an independent 196-month raw dataset, isolated the effect of the COVID-19 disruption: linear models that included a pandemic indicator showed up to a 16% reduction in RMSE.

The route optimisation work unfolds across three distinct stages. The first, published in the *International Journal of Applied Mathematics* (IJAM, 2025), couples an LSTM forecaster with a CVRP solver as a proof of concept on a synthetic network scaled to Surat, and reduces both collection distance and fleet size substantially. The second stage feeds the forecasting output into a network of eight GPS-verified transfer stations plus the Khajod disposal site, and finds an 82% theoretical efficiency potential against current operations. The third stage works directly with 662,803 real operational records from October 2025 and uncovers a serious geocoding constraint: only 8.6% of unique collection point names could be resolved to coordinates. This constraint motivates IMPACT, a GPS-free framework built from five components and validated on the complete operational dataset. Using an XGBoost classifier

with SHAP attribution, IMPACT shows that the dominant cause of missed collection stops across all nine zones is contractor non-compliance, not insufficient fleet capacity.

The decision support component operationalises the forecasting and optimisation outputs within a four-layer architecture (Data Ingestion, Model Execution, Analytics, User Interface), implemented as a Streamlit proof of concept with multiple stakeholder dashboards and verified data sources. The system is explicitly designed to be GPS free and IoT free, a design choice motivated by the Stage 3 geocoding finding and intended to be transferable to the substantial majority of Indian ULBs that lack GPS tracked fleets.

Collectively, the research delivers eighteen original contributions (methodological, empirical, and practical) mapped to three research objectives and addressing five major literature gaps across forecasting, feature engineering, routing, GIS, and integrated decision support. It yields a compliance first governance recommendation validated on large scale operational data for an Indian tier II municipality, and the framework is explicitly aligned with the Smart Cities Mission, Swachh Bharat Mission 2.0, and Sustainable Development Goals 11, 12, and 13, with design choices intended to support replication in comparable Indian and developing economy urban contexts.

2. Brief Description and Introduction

Global municipal solid waste generation is estimated at about 2.01 billion tonnes per year and is projected to rise to 3.40 billion tonnes per year by 2050. India alone generates roughly 62 million tonnes of MSW annually across more than 4,041 Urban Local Bodies, a burden that is intensifying with urbanisation, rising consumption, and the growth of tier II and tier III cities. Surat, Gujarat's second largest city and a Smart Cities Mission beneficiary, exemplifies this trajectory: waste handled by the Surat Municipal Corporation (SMC) has increased from around 3,800 tonnes per day in June 2009 to about 6,200 tonnes per day by December 2024, placing sustained pressure on a nine zone, thirty ward collection and disposal system anchored by eight transfer stations and the Khajod regional disposal facility.

Municipal waste management is conventionally organised around three sequential decisions: how much waste will be generated (forecasting), how collection vehicles should be routed to collect it (optimisation), and how these outputs should be presented to decision makers (decision support). Each decision, taken alone, has been studied internationally - forecasting through classical statistical, machine learning, and deep learning approaches [1]; routing

through exact, heuristic, and metaheuristic formulations; and decision support through dashboard based tools and IoT integrated digital twin platforms. In both the published literature and Indian municipal practice, however, these three decisions are typically addressed as independent problems, evaluated with different rigour on different datasets and rarely validated together against the operational conditions of a single real municipal system at scale. Forecasting studies for Indian cities typically compare a small number of models (often fewer than five) without systematic feature engineering. Route-optimisation studies for Indian MSW systems rely on simplified or synthetic infrastructure data rather than large-scale real operational records, and universally assume the availability of GPS-quality spatial data. It's an assumption that does not hold for the majority of Indian ULBs, where collection points are identified by informal, landmark-based addressing rather than structured geocodes. Decision-support system (DSS) research for waste management, where it exists, is rarely grounded in real operational validation and seldom accounts for the GPS-infrastructure gap that characterises Indian municipal data environments.

This research addresses these three deficiencies as a single, integrated problem rather than as three separate ones. It develops (i) a systematically engineered, comparatively evaluated forecasting framework for SMC's monthly waste generation; (ii) a three-stage route-optimisation research programme that evolves in direct response to the data constraints actually encountered, culminating in a GPS-free operational optimisation framework validated on 662,803 real operational records; and (iii) an integrated decision-support system that operationalises both outputs for SMC's operational and strategic stakeholders. The research draws on three peer-reviewed publications - an IJAM paper (2025, published) [2] establishing the Stage 1 proof-of-concept and DSS architectural foundation, a JSWTM paper (2026, published) [3] establishing the forecasting methodology and COVID-19 interrupted-time-series analysis, and a JETS paper (2026, accepted, Scopus Q3) [4] establishing the seventeen-model comparative forecasting framework. Putting together with substantial thesis-original analysis extending each component to real, large-scale operational data.

The three components are deliberately sequenced rather than merely juxtaposed. The forecasting component establishes *how much* waste SMC's system must handle and *when*; the route-optimisation component establishes *how* that demand should be collected given the city's actual infrastructure and operational constraints; and the decision-support component establishes *who* within SMC's institutional structure needs which output, at what frequency, to act on it. This sequencing is not incidental. It reflects a deliberate research design choice

to treat forecasting accuracy, routing efficiency, and decision-support usability as interdependent rather than independently optimisable properties of a single municipal waste-management system. The research programme also evolves adaptively in response to what the data itself reveals: the transition from a synthetic proof-of-concept (Stage 1) to real infrastructure analysis (Stage 2) to a documented data-infrastructure failure that motivates an entirely GPS-free architecture (Stage 3/IMPACT) is presented, and defended methodologically, as evidence-driven adaptive research design. The research process responds honestly to the constraints it encounters rather than assuming away the data conditions that actually characterise Indian municipal practice.

3. Need, Objectives and Scope of Research

Need for the Study

The research responds to three compounding deficiencies identified through a systematic review of published literature (Section 4) and confirmed through direct empirical investigation of SMC's data environment: limited multi model, feature rich forecasting for Indian municipal contexts; reliance on synthetic routing and untested geocoding feasibility; and the absence of GPS free, operationally validated decision support frameworks for Indian ULBs.

First, no prior published study had conducted a systematic comparison of model families as wide as the one undertaken here spanning linear, regularised, kernel, gradient-boosting, bagging, deep learning, and classical time-series methods within a single, uniformly applied evaluation protocol on Indian municipal MSW data. Existing Indian studies typically test between one and five models, and the specific question of whether feature-engineering quality can substitute for model architectural complexity in data-scarce monthly panel settings had not been empirically resolved.

Second, published route-optimisation research for Indian MSW systems is dominated by synthetic or simplified infrastructure representations. The feasibility of open-source geocoding, the foundational spatial-data requirement for any GPS-based routing formulation had never been systematically evaluated on Indian municipal address data at operational scale; nor had contractor compliance, subsequently identified in this research as the dominant empirical cause of missed collection stops, been incorporated into the theoretical formulation of any published CVRP or inventory-routing model for Indian MSW.

Third, decision-support system research for Indian municipal waste management has consistently assumed GPS telemetry and structured GIS infrastructure that is absent for approximately ninety per cent of India's 4,041-plus ULBs. No published GPS-free DSS combining real operational data, verified stakeholder interfaces, and a structured deployment roadmap had been developed prior to this research.

The consequence, evidenced directly in SMC's own operational data, is chronic collection underperformance - an on-time collection rate of 22.0% by the utility's own operational classification (6.1% under independently verified timestamp analysis), a missed-stop rate of 2.01%, and 147 routes simultaneously exhibiting both high demand and high delay - that has conventionally been attributed to fleet-capacity shortfalls, an attribution this research subsequently shows to be empirically incorrect. Beneath this integrated problem statement, twenty-six specific sub-gaps are identified during the literature review across forecasting methodology, feature engineering, routing and infrastructure representation, geographic-information-system data quality, and decision-support-system design (Section 4), which are consolidated into five major research gaps that collectively motivate and structure the three research objectives stated below. The need for the study is therefore not merely the absence of any one technique for Surat, but the absence, in the published literature generally, of a framework that treats forecasting accuracy, routing feasibility given actual Indian data-infrastructure conditions, and decision-support usability as a single connected design problem rather than as three independently optimisable components.

Research Objectives

The research is organised around three formally approved research objectives.

RO-1 (Forecasting). To develop and comparatively evaluate a machine learning-based forecasting framework for monthly MSW generation in Surat, incorporating a systematic feature engineering pipeline, a seventeen-model comparative analysis across statistical, ensemble, and deep learning families, a multidimensional eight-metric evaluation protocol, and a COVID-19 interrupted-time-series impact analysis.

RO-2 (Route Optimisation). To design and evaluate a multi-stage vehicle route optimisation and operational efficiency framework for SMC waste collection, progressing from a published LSTM-CVRP proof-of-concept through real transfer-station efficiency assessment to a GPS-free IMPACT framework validated against October 2025 operational data.

RO-3 (Decision Support System). To design and validate an integrated Decision Support System that operationalises the forecasting and optimisation outputs for SMC deployment, encompassing a four-layer DSS architecture, forecasting-to-operations integration via a demand scalar, stakeholder-specific interfaces, and a phased municipal deployment strategy aligned with the Smart Cities Mission, Swachh Bharat Mission 2.0, and Sustainable Development Goals 11, 12, and 13.

Research Questions

RQ-1. Which machine learning model family achieves the highest forecasting accuracy for monthly MSW generation in Surat when evaluated across statistical, ensemble, and deep learning families using a 419-feature engineering pipeline, and does systematic feature-engineering quality supersede model architectural complexity in data-scarce municipal forecasting contexts?

RQ-2. How can a multi-stage vehicle route optimisation and operational efficiency framework be designed and progressively validated - from a published LSTM-CVRP proof-of-concept through real transfer-station efficiency assessment to GPS-free operational analysis - to characterise and address the structural causes of missed waste-collection stops in Surat?

RQ-3. How can a GPS-free, operationally deployable Decision Support System be designed to integrate monthly forecasting and near-real-time optimisation outputs for SMC stakeholders across three personnel tiers, and how can such a system be aligned with the Smart Cities Mission, Swachh Bharat Mission 2.0, and Sustainable Development Goals 11, 12, and 13?

Scope of the Study

The research is conducted exclusively within the SMC municipal boundary, nine administrative zones and thirty wards, subdivided into approximately 130 administrative sub-wards together with the eight verified transfer stations and the Khajod disposal site. The primary forecasting dataset spans 186 months (June 2009–December 2024); a separate, independent 196-month raw dataset is used exclusively for the COVID-19 interrupted-time-series analysis and is not otherwise referenced. Route-optimisation Stage 3 and the IMPACT framework are validated on 662,803 real SMC operational collection records from October 2025. The DSS is developed and validated as a computational proof-of-concept; full

production deployment within SMC's institutional IT infrastructure lies outside the scope of the present research and is identified as a future-work item with an explicit phased pathway. The study is single-city in scope; multi-city validation is identified as a priority direction for future research (Section 8).

4. Literature Review and Research Gaps

The literature review synthesises published research across nine thematic domains: MSW generation forecasting; feature engineering and explainable artificial intelligence; comparative machine learning and deep learning model families; evaluation metrics and validation protocols; the impact of COVID 19 on MSW systems; route optimisation for waste collection; GIS and spatial applications; decision support systems and smart city waste management; and integrated forecasting–optimisation–DSS frameworks.

Evolution of forecasting methodology. MSW forecasting research has progressed through five broadly identifiable eras: classical statistical methods (dominant before 2010, principally ARIMA/SARIMA and regression based approaches); early artificial intelligence approaches (2010–2016, typically using three to eight input variables); machine learning ensemble methods (2016–2022, led by Random Forest and gradient boosting families); deep learning approaches (2018–2024, principally LSTM, GRU, and hybrid architectures); and an emerging integrated analytics era (post 2022) combining SHAP based explainability with COVID aware forecasting pipelines. A structured review of eighteen representative MSW forecasting studies across these five eras shows a consistent pattern: reported accuracy improves with each successive era, but reported feature counts, evaluation metric rigour, and validation protocol sophistication do not improve at the same rate – many later eras studies report state of the art architectures evaluated against only one or two accuracy metrics and a small, often opaque, feature set. No published study identified in the review benchmarks more than nine model families under one uniform evaluation protocol, and Indian tier II municipal contexts remain almost entirely unrepresented in the comparative literature.

Feature engineering and explainability. The literature establishes feature set quality as at least as consequential as model architecture. Socio economic and demographic variables are the most extensively studied predictors, followed by climatic variables and systematically engineered temporal (lag and rolling) features. SHAP has become the dominant explainability framework since around 2021. The largest reported candidate feature pool

identified in the reviewed studies comprises approximately fifty variables [5] – an order of magnitude smaller than the 419-feature pipeline developed in this research – and no reviewed study applies SHAP based attribution consistently across more than a small number of compared model families [6].

Comparative model evaluation. Random Forest, support vector regression, and gradient boosting methods dominate published shallow machine learning comparisons, while LSTM, GRU, and CNN LSTM hybrids dominate deep learning comparisons. The widest single study benchmark identified in the review spans nine model families [7]; regularised linear models (Ridge, Lasso, ElasticNet) and Gaussian Process Regression are essentially absent from the comparative MSW literature, despite this research subsequently finding both regularised linear and Gaussian Process approaches among the top three performers.

Evaluation protocols. RMSE and MAE dominate reported metrics; R^2 is common but can mislead on trended series; and MAPE is scale independent but numerically unstable near zero. Theil’s U statistic and Mean Directional Accuracy are almost never reported in the reviewed MSW literature, and rolling origin (walk forward) validation is markedly less common than simple holdout validation. As a result, many published forecasting accuracy claims cannot be interpreted as demonstrating genuine skill above a naïve baseline.

COVID-19 impact studies. Evidence on pandemic-period waste generation is markedly context-dependent [8], [9]: aggregate and commercial waste volumes fell under lockdown restrictions in some settings [9], residential per-capita waste rose in others, and Chile showed an initial lockdown-phase reduction followed by a post-stimulus economic surge [10]. Interrupted time-series and difference-in-differences designs provide the appropriate quasi-experimental toolkit for quantifying such disruptions [11], but no reviewed study integrates a formal interrupted-time-series design with a comprehensive multi model machine learning forecasting pipeline, and no COVID 19–MSW study in the reviewed literature covers an Indian tier II city.

Route optimisation. Classical and capacitated vehicle-routing formulations, typically solved via metaheuristics (genetic algorithms, particle-swarm optimisation, ant-colony optimisation, tabu search, and hybrid combinations thereof) [12], deliver documented efficiency gains generally in the range of fifteen to thirty-five per cent over manual scheduling, evidenced across a reviewed set of seven representative VRP/CVRP/inventory-routing studies [13]. The closest existing approach to the forecast-driven routing developed

in this research is the demand-pull inventory-routing paradigm [14], extended in one notable case to real municipal deployment in Geneva using IoT fill-level sensors [15]. These implementations rely on IoT sensing or statistical demand decomposition rather than machine-learning forecasts, and in every reviewed instance treat transfer stations as dimensionless depot nodes without throughput, compliance, or capacity modelling, a gap of direct relevance to Surat's transfer-station-centred collection network, where the eight transfer stations function as the operational bottleneck rather than as simple pass-through points.

GIS and spatial data. GIS applications in the reviewed literature have matured from basic network routing to multi-criteria facility siting and, in a small number of advanced cases, IoT-integrated real-time systems. Indian peer-reviewed GIS-MSW precedent is confined to three cities - Chennai [16], Lucknow [17], and Erode [18] - and does not include Surat. Every reviewed framework assumes GPS-level spatial data availability; a condition this research's own geocoding pilot (Section 6) demonstrates does not hold for Surat's operational address data.

Decision support and integrated frameworks. Published DSS approaches range from classical scenario-planning tools through IoT-enabled smart systems [19] to digital-twin platforms, with the most architecturally advanced reviewed implementation being a municipal digital-twin platform [20] that reaches the second-highest level on the integration-maturity taxonomy applied during the review. The five-level integration taxonomy (Level 1: single-component analytical tool; Level 5: forecasting, optimisation, and DSS fully integrated, empirically validated on real data, and explicitly policy-aligned) shows that fewer than ten reviewed studies globally reach forecasting-plus-optimisation integration, fewer than seven include a dedicated decision-support layer, and only one Indian implementation i.e., a multi-modal reporting and estimation system for a different Indian city [21], reaches the highest integration level identified, while still requiring IoT infrastructure incompatible with the informal addressing constraint documented for Surat. No published framework in the reviewed literature achieves GPS free, machine learning driven, Indian policy aligned integration across all three components (forecasting, optimisation, and decision support) simultaneously – the specific combination this research is positioned to deliver.

Research Gaps. The critical analysis consolidates the review's findings into five major research gaps, each mapped explicitly to one or more of the three research objectives:

1. **Insufficient multi-model benchmarking and COVID-19 structural-break methodology for Indian tier-II MSW forecasting** - no study benchmarks the full model-family spectrum on a single Indian tier-II dataset under a uniform multi-metric protocol, and formal interrupted-time-series methodology has never been integrated with a comprehensive machine-learning pipeline for this context. ($\rightarrow RO-1$)
2. **Narrow feature engineering and limited cross-model explainability attribution** - published feature pools rarely exceed fifty variables, and cross-model SHAP attribution across a full comparative model suite is absent from the literature. ($\rightarrow RO-1$)
3. **Disconnection between machine-learning forecasting and demand-adaptive routing, with transfer-station analytics absent** - no study couples a comprehensive forecasting benchmark with formal route- or inventory-optimisation, and transfer-station-level operational analytics are absent from all reviewed routing formulations. ($\rightarrow RO-1, RO-2$)
4. **Pervasive GPS-dependence of routing and GIS frameworks, inadequate for Indian informal-settlement addressing** - every reviewed routing and GIS framework assumes GPS-quality spatial data unavailable in Indian informal urban environments, and Indian GIS-routing precedent is confined to three non-Surat cities. ($\rightarrow RO-2$)
5. **Absence of an integrated, GPS-free, policy-aligned decision-support framework for Indian tier-II cities** - no reviewed framework combines machine-learning forecasting, GPS-free routing, and multi-stakeholder decision support for an Indian tier-II city aligned with Swachh Bharat Mission, Smart Cities Mission, and SDG targets. ($\rightarrow RO-3$)

5. Study Area and Data Collection

Study Area

The research is conducted for the city of Surat, Gujarat's second largest city and a Smart Cities Mission beneficiary, with an estimated population of approximately 6.1 million. Solid waste collection and disposal are administered by the Surat Municipal Corporation (SMC) across nine administrative zones and thirty wards, further subdivided into approximately 130 administrative sub wards. Collection infrastructure comprises eight transfer stations with verified geographic coordinates – Anjana, Bhatar, Bhestan, Dindoli, Katargam, Kosad, Pal, and Varachha – feeding into the Khajod regional disposal facility, located about 16 kilometres from the city centre and operational since 2014. Total municipal waste handled

has increased from approximately 3,800 tonnes per day in June 2009 to roughly 6,200 tonnes per day in December 2024.




Figure: 1. Study Area Map of Surat Municipal Corporation

Zones are colour-coded by administrative boundary. Transfer station locations are indicated by verified GPS coordinates. Source: SMC administrative data; map generated using Python matplotlib and Surat shapefiles.

Data Collection

Three distinct, purpose-specific datasets underpin the research, each explicitly scoped to a specific analytical objective and never interchanged.

Forecasting dataset (RO-1, canonical). A 186-month monthly panel (June 2009–December 2024) of SMC waste generation records, obtained via a formal data sharing arrangement and disaggregated by waste category and zone, forms the canonical forecasting dataset used throughout the thesis. It is supplemented by four categories of secondary data: demographic and economic indicators (Census based population, Gujarat state GDP, national Consumer Price Index, monthly diesel prices); climatic variables (India

Meteorological Department monthly mean temperature and rainfall for the Surat district); mobility data (Google Community Mobility Reports, available from February 2020); and an event/festival calendar (major Indian festivals and SMC civic events). This 186-month span is used consistently for every general statement about the forecasting dataset throughout the research.

COVID-19 analysis dataset (RO-1, Section 4.5.3 only). A separate, independently sourced 196-month raw dataset (June 2009–October 2025), associated with the JSWTM publication and carrying only minimal lag-based feature engineering, is used exclusively for the interrupted time-series COVID 19 impact analysis. This dataset is explicitly labelled "JSWTM raw dataset (196 months)" wherever it appears and does not inform any other section of the research.

Route-optimisation dataset (RO-2, Stage 3). A large scale real operational dataset comprising 662,803 individual collection event records from October 2025 – covering nine zones, 131 wards, approximately 936 active collection vehicles, 3,341 named collection routes, and 455,827 individual point of interest stop records – forms the empirical basis for Stage 3 of the route optimisation programme and for the IMPACT framework. This dataset contains no GPS coordinates for individual collection points – a structural absence that is itself a central empirical finding of the research (Section 6.3).

A fourth, smaller dataset - raw SMC monthly records spanning approximately 192 months, distinct from both the 186-month canonical dataset and the 196-month COVID dataset - underlies the Stage 1 LSTM forecasting model published in the IJAM paper, reflecting the earlier data-availability conditions at the time that analysis was conducted.

Data preprocessing. All datasets pass through a common four stage preprocessing pipeline prior to feature engineering: (i) datetime decomposition and monthly index construction, with unparseable date entries excluded; (ii) coercion of object typed fields to numeric form, with non-convertible entries treated as missing; (iii) missing value handling, in which columns exceeding fifty per cent missingness are dropped entirely and remaining missing values are filled using column wise mean imputation (the primary SMC waste generation series itself has a missing rate of under two per cent for the target variable); and (iv) a two stage sparsity and near zero variance filter that reduces the 419 feature engineered set first to 413 features (removal of features exceeding fifty per cent missingness) and then to 396

active features (removal of near zero variance features), the latter forming the final model input set for all seventeen forecasting models.

6. Research Methodology, Data Analysis and Results

This section presents the research methodology and the associated results together, chapter by chapter, in keeping with the completed thesis's structure (Chapters 3 through 6), since the methodological design and the empirical findings of each component are most clearly understood in relation to one another rather than as separated theoretical and empirical treatments. Section 6.1 describes the overall four phase research design; Section 6.2 describes the feature engineering framework common to the forecasting component; Section 6.3 presents the forecasting methodology and results (RO 1); Section 6.4 presents the three stage route optimisation methodology and results, including the IMPACT framework (RO 2); Section 6.5 presents the Decision Support System design (RO 3); and Section 6.6 synthesises the three components, positions the integrated framework against prior published work, and assesses its robustness.

6.1 Overall Research Design

The research follows a positivist and empirically grounded analytical approach, organised as a deductive, four-phase integrated methodology: (1) data collection, preprocessing, and feature engineering; (2) development and comparative evaluation of forecasting models; (3) vehicle routing optimisation, further structured into three progressive sub-stages; and (4) design of the decision support system. The forecasting component adopts a cross-sectional design, based on 186 months of historical panel data, while the route optimisation component is quasi-experimental, where interventions are applied to real operational data and evaluated analytically. Throughout the study, a pragmatic approach guides the selection of models and methods, with emphasis placed on proven performance, computational practicality, and interpretability, rather than unnecessary architectural complexity.




FIGURE 2. Integrated four-phase research methodology framework for the Surat MSW management study.

6.2 Feature Engineering Framework

A systematic 419-feature engineering pipeline was constructed from the 186-month canonical dataset across five feature domains: temporal-lag features (twelve monthly lags); rolling statistical features (three window widths - three, six, and twelve months - for mean and standard deviation); socio-economic features (population, Consumer Price Index, diesel price, and Gujarat state GDP, each as both level and year-on-year percentage change); climatic features (monthly mean temperature and rainfall, plus derived temperature-deviation and seasonal indicators); and mobility/event features (six Google Mobility change variables, a COVID-19 indicator, and seven binary festival flags), together with a set of interaction and year-on-year change terms. Two sequential filtering stages were applied for model input: a sparsity filter removing features with more than fifty per cent missing observations (419 $\rightarrow$ 413 features), followed by a near-zero-variance filter (413 $\rightarrow$ 396 active features used for all model training). The 419-feature figure is used to describe the engineered feature set; the 396-feature figure describes the filtered model-input set - both figures are used consistently and are explicitly distinguished throughout the research.

6.3 Forecasting Framework (RO-1)

Seventeen forecasting models drawn from six algorithmic families – linear and regularised regression, kernel methods, gradient-boosting and bagging ensembles, deep learning architectures, and classical time-series baselines – were trained on the 419/396-feature dataset. All models were evaluated under a uniform eight-metric protocol (R^2 , MAE, RMSE, MAPE, coefficient of variation, Theil’s U, Mean Directional Accuracy, and residual standard deviation) using a 70/15/15 train/validation/test split for cross-sectional models and a chronological split for time-series models. Hyperparameters were tuned using grid or random search within family-appropriate ranges, and Diebold–Mariano tests were applied to assess statistical significance of forecast-accuracy differences between leading models. Full model specifications and tuning grids are documented in the thesis.

Results. The seventeen models are organised into four tiers according to their test-set R^2 . Tier 1 ($R^2 > 0.996$) contains OLS, Gaussian Process Regression (GPR), and Ridge regression, all delivering very high accuracy (for OLS, $R^2 = 0.9972$, MAPE = 1.43%). Tier 2 ($R^2 \geq 0.97$) includes ElasticNet, LightGBM, Lasso, CatBoost, and XGBoost, which perform well but still clearly below the Tier 1 models. Tier 3 ($R^2 \geq 0.95$) comprises the tuned

Multilayer Perceptron and Random Forest. Beyond these top-performing models, seven additional candidates – Holt–Winters, Support Vector Regression, SARIMA, ConvLSTM, LSTM, GRU with attention, and Prophet – perform poorly under the engineered-feature protocol, with several showing negative R^2 values on the test set.




FIGURE 3. Comprehensive feature engineering pipeline, illustrating the sequential transformation of 186 months of raw SMC MSW records through five domain-specific feature categories — temporal, socio-economic, climatic, mobility, and event — with derived interaction terms and model-specific scaling to yield a final set of 419 engineered features for the comparative forecasting analysis.

These models are therefore not considered suitable for operational deployment in the Surat context, and their detailed results are omitted from this synopsis but documented in full in

the thesis. A computational-efficiency comparison further shows that OLS trains in 0.010 seconds, whereas CatBoost requires 28.517 seconds for materially worse accuracy.

Cross-model feature-importance analysis (LightGBM, XGBoost, and Random Forest, aggregated with SHAP confirmation) identifies short-horizon lag features (one- and three-month) and twelve-month rolling statistics as the dominant predictors, with Consumer Price Index and diesel price contributing consistent secondary importance and Google Mobility features showing marked importance specifically during the 2020–2022 pandemic window. An ablation study, systematically excluding each feature domain from the OLS model, shows only minimal R^2 degradation under any single-domain removal, indicating that predictive information is distributed redundantly across feature domains rather than concentrated in any one category.

Why feature engineering matters more than architecture. The central interpretive finding is that the 419/396 feature pipeline pre encodes the causal structure of monthly municipal waste generation – seasonal, economic, climatic, and event driven – sufficiently well that a linear model can exploit this structure directly without having to learn it from raw sequences. Gradient boosting ensembles partially exploit the same engineered structure and achieve respectable, but consistently inferior, performance (Tier 2). Deep sequence models (LSTM, GRU with attention, ConvLSTM), by contrast, are data hungry; with only 186 monthly observations they lack sufficient training signal to outperform models that already receive the temporal structure as explicit inputs. This is a context specific result: in data scarce, monthly granularity municipal panels typical of Indian ULBs, investment in systematic feature engineering yields a substantially higher return than investment in model architectural complexity. A municipality that implements the feature pipeline and deploys OLS will, in this setting, outperform one that deploys an LSTM on raw monthly data ($R^2 = 0.9972$ versus $R^2 = -3.04$), a finding with direct relevance for how resource constrained ULBs prioritise analytical capacity. The feature construction methodology itself is designed to be transferable across comparable tier II cities, with coefficients and weights re estimated on local socio economic, climatic, and event calendar data.

TABLE 1 Eight-metric performance evaluation of 10 successful forecasting models on the held-out test set.

Model	R²	MAE (t)	RMSE (t)	MAPE (%)	CV (%)	Theil's U	MDA (%)	Res. Std (t)	Family
OLS	0.9972	198.3	275.6	1.43	1.84	0.0642	92.86	268.0	Linear
GPR	0.9970	214.4	288.4	1.51	1.92	0.0672	85.71	278.7	Kernel
Ridge	0.9962	249.9	325.3	2.03	2.17	0.0758	85.71	317.6	Linear
ElasticNet	0.9940	323.3	406.6	2.67	2.71	0.0948	85.71	404.3	Linear
LightGBM	0.9830	580.7	684.0	4.07	4.56	0.1594	75.00	677.2	Boosting
Lasso	0.9787	553.9	765.8	4.04	5.10	0.1785	82.14	756.5	Linear
CatBoost	0.9766	629.3	803.5	5.22	5.35	0.1873	85.71	802.2	Boosting
XGBoost	0.9697	625.5	913.9	4.32	6.09	0.2130	82.14	900.8	Boosting
Tuned MLP	0.9655	808.6	974.5	6.24	6.49	0.2271	71.43	974.3	NeuralNet
Random Forest	0.9596	803.2	1054.4	5.79	7.02	0.2457	71.43	1053.5	Bagging




FIGURE 4. Model performance landscape - R^2 values for all 17 models arranged by tier, with model family colour coding and tier boundary annotations.

COVID-19 interrupted-time-series analysis (Section 4.5.3, JSWTM 196-month dataset only). Using the independent 196-month raw dataset, six models (Linear Regression, Random Forest, XGBoost, Support Vector Regression, a stacking ensemble, and a Keras Multilayer Perceptron) were each evaluated with and without a binary COVID-19 indicator (March 2020–March 2022). Random Forest achieved the best overall performance in this sub-analysis ($R^2 = 0.8506$, RMSE = 0.4318, normalised scale), followed by a stacking ensemble using Linear Regression as meta-learner ($R^2 = 0.8106$), with the plain Linear-Regression baseline ($R^2 = 0.5255$) and Support Vector Regression ($R^2 = 0.1692$) performing markedly worse on this smaller, less richly engineered dataset. Including the COVID-19 indicator reduced RMSE by approximately sixteen per cent in Linear Regression and approximately ten per cent in XGBoost, demonstrating that explicit pandemic-period encoding materially improves forecast accuracy during structural-break conditions [3]. Benchmarked against five comparable published studies applying similar model families to municipal waste or related utility-demand forecasting under COVID-19 or structural-break conditions, the present sub-analysis's Random Forest result ($R^2 = 0.8506$) compares

favourably, while the primary JETS-based canonical result (OLS $R^2 = 0.9972$) substantially exceeds all reviewed benchmarks - a difference attributed principally to the richer 396-feature input space rather than to algorithmic choice, since Linear Regression, Random Forest, and XGBoost are common to both the canonical and the COVID-specific analyses.

6.4 Route Optimisation Framework (RO-2)

The route-optimisation research progresses through three empirically distinct, sequentially motivated stages, each responding directly to the data-availability constraints encountered in the preceding stage.

Stage 1 - Published proof-of-concept (IJAM, 2025). A purpose-built LSTM network, trained chronologically on approximately 192 months of raw SMC waste data, achieved a test RMSE of 0.1095 and $R^2 = 0.93$. [2]. The resulting monthly demand estimates were disaggregated and fed into a Capacitated Vehicle Routing Problem solved with Google OR-Tools on a synthetic, standards-validated Surat-scale network of forty-three collection points and a fleet of twenty-five vehicles, with vehicle mix and capacity parameters aligned to MoHUA and CPCB standards. Optimisation reduced total daily collection distance from 392.04 km to 194.78 km (50.31% reduction) and fleet requirement from twenty-five to nineteen vehicles (24% reduction), with associated environmental benefits estimated for the representative synthetic network. These results are consistently qualified as proof-of-concept outcomes rather than validated operational results for SMC's actual collection fleet.

Stage 2 - Transfer-station efficiency assessment. Applying the forecasting output to eight real, GPS-verified transfer stations (Anjana, Bhatar, Bhestan, Dindoli, Katargam, Kosad, Pal, and Varachha) and the Khajod disposal site, a Haversine-distance-based transportation-efficiency analysis established a baseline of 926.8 km/day, computed from forty-five vehicle-trips per day at CPHEEO-standard ten-tonne vehicle capacity, against an analytical lower bound of 166.8 km/day (assuming one direct round trip per transfer station under maximum load consolidation) - an 82% theoretical optimisation potential, explicitly labelled as an idealised planning target rather than an operationally achievable outcome. Demand is distributed unevenly across the eight stations, with Kosad handling the largest share (in excess of 100 tonnes per day) and distances to the Khajod disposal site ranging from just over one kilometre (Kosad) to nearly ten kilometres (Pal), underscoring that the

transfer-station network itself - not merely the final collection round - is a material driver of total system distance and fuel consumption.

Stage 3 - Full operational-data investigation and the geocoding constraint. Acquisition of 662,803 real SMC operational collection records from October 2025 enabled an attempt at beat-level vehicle routing using the actual collection network. A geocoding feasibility pilot, conducted on the West Zone (99 routes, 162 unique collection-point names), using the open-source Nominatim geocoder, successfully resolved usable coordinates for only 8.6% of unique name - far below the seventy-to-eighty per cent success rate typically expected for structured addresses - rendering beat-level, GPS-based vehicle-routing formulation infeasible on this dataset. This finding, documented as a valid and original negative result, is interpreted as evidence of a structural data-infrastructure gap affecting the large majority of Indian Urban Local Bodies, and directly motivated a research pivot toward a GPS-free operational-optimisation architecture.

The IMPACT Framework (GPS-free operational optimisation). IMPACT - Integrated Multi-level Planning for Adaptive Collection Timing - comprises five sequentially interdependent components, validated on the full 662,803-record operational dataset: Time Window Recalibration (TWR); Vehicle Allocation Optimisation (VAO); the Missed Stop Root Cause System (MSRS); Ward Demand Balancing and Route Balancing (WDB+RB); and Forecasting Integration Optimisation (FIO). A sixth analytical layer, Multi-Objective Optimisation across the four IMPACT objectives, was fully formulated but not implemented within the present research scope and is identified as the principal direction for future work (Section 8.4).

The baseline diagnostic established an on-time collection rate of 22.0% under SMC's own operational status classification (6.1% under independently parsed timestamp analysis, a discrepancy explicitly reconciled as reflecting two distinct, equally valid measurement definitions), a system-wide missed-stop rate of 2.01%, and 147 routes simultaneously exhibiting both high demand and high delay, with the Central Zone recording the highest missed-stop rate (9.31%) despite having the lowest demand-per-vehicle ratio in the system.

Time Window Recalibration (TWR) operates in two parts. TWR-A tested thirty-seven candidate global time-shift scenarios and identified an optimal shift of minus seventy-five minutes, improving the strict-definition on-time rate from 6.10% to 8.19% - a system-wide gain of 2.09 percentage points. TWR-B then develops individually calibrated window

recommendations for all nine zones, 131 wards, and 3,260 qualifying routes: the majority of wards (just over half) require a major schedule adjustment exceeding sixty minutes, roughly a third require a moderate adjustment, and a small remainder require little or no change - with a phased rollout sequencing the zones showing the strongest early-phase gains first, followed by the remaining zones including the Central Zone.




FIGURE 5 Three-Stage Route Optimization Research Evolution Diagram

Vehicle Allocation Optimisation (VAO), solved as an integer linear program using Google OR-Tools' CP-SAT solver to optimality, computed a minimum operationally required fleet of approximately 689 vehicles against the deployed fleet, indicating a surplus of approximately 260 vehicles concentrated overwhelmingly in a small number of zones. The Central Zone in particular is shown to be substantially over-resourced in vehicle count relative to its demand - it has the lowest demand-per-vehicle ratio of any zone in the system - while nevertheless recording the system's highest missed-stop rate, and only two of the nine zones (South East Zone and South Zone B) are found to be under-resourced relative to the computed minimum. Reallocating the identified surplus toward these two under-resourced zones produced a change in the system-wide missed-stop rate of only -0.0008 percentage points - a result interpreted as strong evidence that fleet capacity is not the binding constraint on collection performance anywhere in the system.

The Missed Stop Root Cause System (MSRS), an XGBoost classifier (Area Under Curve = 0.940, Recall = 88.1%) with SHAP-based attribution across five causal categories, identified contractor compliance as responsible for 75.72% of the model's predictive attribution for missed stops, against 1.65% for vehicle capacity, 5.42% for scheduling factors, and smaller contributions from demand and operational factors. This finding replicates without exception across all nine zones and is corroborated by a striking natural test: the Central Zone, with the lowest demand-per-vehicle ratio in the system, records the *highest* missed-stop rate - the opposite of what a capacity-driven explanation would predict.

Ward Demand Balancing and Route Balancing (WDB+RB), using a First-Fit-Decreasing heuristic supplemented by linear-programming minimax balancing, eliminated all 232 statistically overloaded routes (defined as more than 1.5 standard deviations above mean zonal demand) and reduced the coefficient of variation of route-level demand from 1.42 to 0.60 - a 57.5% improvement in load distribution equity. Rebalancing performance varies by zone: some zones (including South East Zone) show near-complete elimination of demand-load variability, while others (including the Central Zone) show more modest gains, reflecting the underlying differences in each zone's route-density and demand-distribution structure.

Forecasting Integration Optimisation (FIO) formalises the link between the monthly OLS forecast and daily operational planning through a demand scalar (the ratio of the forecast daily demand to a baseline reference value of 1,195,756 kilograms per day), with three

defined operating scenarios - routine, elevated-demand alert, and contingency - that translate forecast deviations into pre-positioning, alerting, and surplus-mobilisation actions respectively.

The central IMPACT finding. In Surat, missed municipal waste collection stops are driven primarily by contractor compliance failure rather than fleet capacity, with SHAP-based attribution assigning 75.72% of predictive importance to compliance and only 1.65% to capacity. This result has direct implications for capital allocation policy, shifting the main question from “how many additional vehicles are required?” to “how should contractor performance be measured, monitored, and incentivised?”

IMPACT Framework: System Architecture Overview <i>Surat MSW Route Optimization</i>				
Component	Method	Key Result	Status	Lever (pp ONTIME / miss)
MSW Forecast (Layer 0)	Ordinary Least Squares (OLS)	$R^2 = 0.9972$ $MAPE = 1.43\%$ Daily scalar $s_d = F_d / 1,195,756$	DONE	+Anchors all downstream
KPI Baseline (Stage 1)	Descriptive stats 455,827 stops	ONTIME = 22.0% MISS = 2.0% 147 overloaded routes CV = 1.48	DONE	Baseline reference
TWR - Time Window Recal. (Level 3)	Timestamp-parsed shift simulation	Global shift: ~75 min ONTIME: 6.1% → 8.2% (+2.09 pp) South A: +7.97 pp	DONE	+2.09 pp ONTIME
VAO - Vehicle Allocation (Level 1)	ILP (OR-Tools CP-SAT) + kappa	Surplus: 260 vehicles Miss delta: -0.0008 pp (negligible) Compliance floor CZ: 93.1%	DONE	-0.0008 pp MISS
WDB+RB - Demand Balancing (Level 2)	FFD heuristic + LP minimax	Mean CV: 1.42 → 0.60 (-57.5%) 232 overloaded routes → 0 Max miss delta: -0.12 pp	DONE	-0.12 pp MISS
MSRS - Root Cause System (Cross-cutting)	XGBoost + SHAP 24 features	AUC = 0.940 Recall = 88.1% Compliance: 75.72% Capacity: 1.65%	DONE	×46 leverage vs capacity
MOO - Pareto Frontier (Unified)	e-constraint multi-objective	Trade-off: ONTIME vs MISS vs CV vs COST Pending implementation	PENDING	Synthesis layer
IMPACT = 6 validated modules 455,827 stop records GPS-free Deployable with existing SMC data				

FIGURE 6 IMPACT System Architecture Diagram

6.5 Decision Support System Framework (RO-3)

The Decision Support System operationalises the forecasting and optimisation outputs within a four-layer architecture. The Data Ingestion Layer reads from verified operational and model-output data sources and is designed for scheduled extraction from SMC’s operational database. The Model Execution Layer hosts the OLS monthly forecasting model, a Random Forest daily-disaggregation model, and all five IMPACT components. The

Analytics Layer aggregates zone-, ward-, and route-level key-performance indicators, computes SHAP-based root-cause attribution, and evaluates the FIO demand-scalar logic. The User Interface Layer, implemented as a Streamlit proof-of-concept, presents stakeholder-oriented dashboards with role-differentiated views spanning strategic, management, and operational tiers.




Figure: 7 DSS Four-Layer Architecture Diagram




FIGURE 8 DSS Stakeholder Dashboard Wireframes

IMPACT DSS Wireframe 01 - Executive Dashboard (left panel) and Forecast Dashboard (right panel). IMPACT DSS Wireframe 02 - KPI Baseline Dashboard (left panel) and Optimisation Dashboard (right panel). IMPACT DSS Wireframe 03 - Compliance Dashboard (left panel) and IMPACT Framework Dashboard (right panel).

6.6 Integrated Framework Synthesis, Comparative Positioning, and Robustness

The three components are not merely reported alongside one another; their interaction is itself analysed. The forecasting-to-operations link is quantified through the FIO demand scalar: because OLS forecasting accuracy is high ($\text{MAPE} = 1.43\%$), the forecast-induced uncertainty injected into daily operational planning is small relative to the ten-per-cent and twenty-per-cent thresholds that trigger FIO's alert and contingency scenarios respectively - meaning the routing and allocation layers built on top of the forecast are not undermined by forecast noise. This is presented as a structural argument for why forecasting accuracy is not merely a reporting statistic but an operationally load-bearing property of the integrated system: a materially less accurate forecasting model - of the kind more typical in the reviewed literature, where double-digit MAPE values are common - would be expected to generate spurious alerts and erode confidence in the downstream optimisation layers.

The integrated framework was benchmarked against eight prior forecasting–optimisation–decision-support frameworks identified in the literature, including several digital-twin and IoT-based architectures. [20], [21], [22], [23], [24]. Applying a five-level integration-maturity taxonomy (from single-component analytical tools to fully integrated, policy-aligned systems validated on real operational data), all reviewed frameworks fell short of full integration on at least one dimension – typically the absence of a validated decision-support layer, reliance on synthetic rather than real operational data, or dependence on GPS/IoT infrastructure. The present framework uniquely combines large-scale operational validation, a GPS-free design, and explicit alignment with Indian policy missions, placing it at the highest integration level in this taxonomy.

The framework's generalisability is assessed against ten transferability criteria - including population scale, monthly-data availability, GPS-infrastructure status, contractor-operated collection, and zone/ward administrative structure - for application to comparable Indian tier-II cities and, more broadly, to developing-economy urban contexts sharing Surat's informal-addressing and limited-telemetry conditions. The 419-feature engineering pipeline is assessed as generalisable in structure but requiring local socio-economic, meteorological, and event-calendar recalibration for each new city; the MSRS compliance-attribution model is assessed as requiring independent re-estimation per city, since the 75.72% compliance-attribution figure is explicitly Surat-specific rather than a universal constant; and the GPS-free architectural principle itself is assessed as the most broadly transferable element, since the underlying data-infrastructure constraint it responds to - most Indian ULBs' lack of GPS-tracked collection fleets - is a structural condition shared across the large majority of India's more than four thousand Urban Local Bodies, and, more broadly, across many developing-economy urban contexts characterised by similarly informal, landmark-based addressing systems.

Robustness and internal consistency. Three potentially contradictory results, each identified during the research's own internal consistency review, are explicitly reconciled rather than left unaddressed. First, the LSTM model produces markedly different results depending on evaluation context ($R^2 = -3.04$ under the tabular, random-split protocol used for the seventeen-model comparison; $R^2 = 0.93$ under the chronological, purpose-built sequence-modelling protocol used in the published IJAM study) - both values are correct within their respective evaluation designs, and the divergence is treated as a substantive methodological finding rather than an error to be resolved away. Second, the on-time

collection rate is reported under two distinct and equally valid definitions - 22.0% under SMC's own operational status classification and 6.1% under independently parsed timestamp analysis - with the TWR simulation using the more conservative, independently verifiable 6.1% baseline throughout. Third, the dataset-span distinction between the 186-month canonical forecasting dataset and the 196-month dataset used exclusively for COVID-19 analysis is maintained without exception throughout the research; no analysis in the thesis uses an intermediate or unexplained dataset span.

Sensitivity of the integrated conclusions. Three central claims of the research are additionally assessed for sensitivity to their underlying assumptions. The compliance-dominance claim is tested against the possibility that the vehicle-level miss-rate feature used by MSRS is itself partly circular (since it is derived from the same missed-stop outcome the model predicts); this is acknowledged as a genuine methodological caveat for prospective, forward-looking deployment, but does not undermine the claim's validity as a retrospective causal-attribution finding, given its independent corroboration through nine-zone replication and the Central Zone's counter-intuitive result. The forecasting-accuracy claim is tested against the counterfactual of a materially less accurate model (of the kind more typical in the reviewed literature): under such conditions, the FIO scalar's alert and contingency thresholds would be expected to trigger spuriously, eroding operational trust in the DSS - an argument that treats forecast accuracy as operationally load-bearing rather than as a standalone reporting statistic. The GPS-free deployability claim is tested against the possibility that commercial geocoding services (rather than the open-source Nominatim geocoder used in the Stage 3 pilot) might achieve materially higher success rates; this is acknowledged as plausible and is explicitly identified as the primary pathway toward eventual GPS-based routing, but does not weaken the finding that open-source geocoding - the only option realistically available to most resource-constrained Indian ULBs - fails at operational scale on Surat's address data, which is the specific empirical claim on which the GPS-free architectural decision rests.

TABLE 2 Transferability Assessment of the Surat MSW Framework to Comparable Tier-II Indian Cities and Developing-Economy Urban Contexts

Criterion	Surat Value	Applicability Condition	Generalizable?	Adaptation Required
Population scale	~6.9 million	>1 million ULB with monthly MSW data	Yes - all MoHUA Class I cities	None - pipeline scales

Monthly data availability	186 months (Jun 2009–Dec 2024)	Minimum ~5 years continuous records	Yes for 5+ year records	Feature set reduction for shorter series
GPS infrastructure	Absent (geocoding 8.6% success)	GPS absent or partial coverage	Yes - primary design condition	No adaptation needed
Contractor-operated collection	Yes - all 9 zones	Contractor-operated vs direct ULB operation	Yes for contractor systems	MSRS must be re-estimated for direct ULB operation
Zone/ward structure	9 zones, 131 wards	Any zoned administrative structure	Yes	Zone-level parameterisation of IMPACT
MSW composition	Organic-dominant (~52%)	Broadly applicable	Largely yes	Feature weights for industrial waste proxies
419-feature pipeline	396 active from 419	Public data: IMD, CPCB, Census, Google Mobility	Yes - all sources available across India	Localisation of CPI, population, event calendars
IMPACT components	All 5 operational on Oct 2025 data	Stop-level visit records with timestamps and vehicle IDs	Yes - standard contractor reporting	Local vehicle capacity, zone demand data
MSRS compliance model	XGBoost AUC=0.940 (Surat)	Contractor systems with vehicle-level records	Methodologically yes; empirically city-specific	75.72% attribution is Surat-specific; re-estimate per city
DSS deployment	Streamlit PoC, GPS-free	Basic office-grade computing	Yes - no specialised hardware	Zone structure, RBAC tier, language localisation

7. Major Findings, Contributions and Validation

7.1 Objective-wise Findings

RO 1 (Forecasting and Feature Engineering). The seventeen-model benchmark on Surat’s 186-month panel indicates that Ordinary Least Squares (OLS), Gaussian Process Regression (GPR), and Ridge regression, when used with the 419/396-feature engineering pipeline, deliver the highest accuracy, with OLS achieving $R^2 = 0.9972$ and $MAPE = 1.43\%$. Deep sequence models perform poorly under the same protocol, reinforcing that, in data-scarce monthly municipal settings, systematic feature engineering yields larger gains than added architectural complexity.

RO 2 (Route Optimisation and IMPACT). The three-stage routing programme shows that synthetic proof-of-concept optimisation can cut collection distance by half and reduce fleet

size, that the existing transfer-station infrastructure has substantial untapped efficiency potential, and that a GPS-free IMPACT framework can be rolled out on 662,803 real operational records. Analysis of missed collections further indicates that contractor compliance, rather than fleet capacity, is the main structural cause of service failure across Surat's nine zones.

RO 3 (Decision Support System and Integrated Framework). A four-layer, GPS- and IoT-free Decision Support System (DSS) architecture was designed and implemented as a Streamlit proof-of-concept, integrating monthly forecasting and operational optimisation outputs for stakeholders across three institutional tiers. The integrated framework is aligned with Smart Cities Mission, Swachh Bharat Mission 2.0, and SDG 11/12/13 priorities and is positioned for phased deployment within SMC's institutional context.

7.2 Contributions of the Research

The research makes eighteen contributions, which are grouped here for brevity.

Forecasting and feature engineering (C1–C5). The study introduces a 419/396-feature engineering pipeline for municipal MSW data; systematically benchmarks seventeen models across six algorithmic families on a 186-month Surat panel; demonstrates that linear and regularised models with rich features outperform deep architectures in data-scarce settings; integrates SHAP-based explainability across multiple model families; and provides a reusable feature-construction methodology for other tier-II cities.

Routing and operational analytics (C6–C11). It develops a three-stage route-optimisation programme from synthetic proof-of-concept to GPS-free operational validation; quantifies transfer-station efficiency potential; documents the failure of open-source geocoding at operational scale for informal Indian addresses; designs and validates the IMPACT framework on 662,803 real collection records; and establishes contractor compliance – rather than fleet capacity – as the dominant structural driver of missed collection.

Decision support and integrated framework (C12–C18). It designs and implements a four-layer, GPS- and IoT-free DSS architecture; integrates forecasting and routing outputs into stakeholder-specific dashboards; defines an integration-maturity taxonomy for municipal waste frameworks and positions the proposed system at the highest level using real-data evidence; aligns the framework with Smart Cities Mission, Swachh Bharat Mission

2.0, and SDG 11/12/13; and outlines a phased deployment roadmap for SMC with clear policy and governance implications.

7.3 Validation

Validation is carried out along four dimensions. Forecasting models are evaluated using multi-metric performance, statistical significance testing, and residual diagnostics. Routing and IMPACT results are validated through large-scale application to real October 2025 operational data, sensitivity checks on key parameters, and consistency with observed field constraints. The DSS is validated as a working prototype through internal consistency checks and structured walkthroughs with intended user groups. Together, these validations provide sufficient evidence that the proposed framework is technically sound and operationally plausible within SMC's context.

8. Conclusions, Limitations and Future Recommendations

8.1 Conclusions

This research set out to address, as a single integrated problem, three deficiencies in Indian municipal solid waste management research and practice: the absence of systematically benchmarked, feature rich forecasting for Indian municipal contexts; the reliance of route optimisation research on synthetic infrastructure and unavailable GPS data; and the absence of an integrated, GPS free decision support framework grounded in real operational validation. All three research objectives are met on the evidence summarised in Section 7.1, and the three research questions in Section 3 are answered directly by the results obtained. In response to RQ 1, OLS regression built on the 419/396 feature engineering pipeline is identified as the highest accuracy model among the seventeen tested, with systematic feature engineering quality shown empirically to supersede model architectural complexity under data scarce, monthly granularity municipal forecasting conditions. In response to RQ 2, the three-stage route optimisation evolution – synthetic proof of concept, real infrastructure efficiency assessment, and GPS free operational validation – characterises the structural causes of missed collection stops in Surat, identifying contractor compliance failure, rather than fleet capacity, as the dominant cause across all nine zones. In response to RQ 3, a GPS free, four-layer Decision Support System is designed, implemented as a working proof of

concept, and validated across technical, internal consistency, interface, and policy alignment dimensions, with a phased pathway toward institutional deployment aligned with Smart Cities Mission, Swachh Bharat Mission 2.0, and SDG 11/12/13 targets.

The single most policy-consequential finding of the research is that contractor compliance failure, not fleet capacity, is the dominant structural cause of missed waste-collection stops in Surat - a finding that directly inverts the conventional policy response (fleet expansion) and redirects the highest-leverage recommended investment toward contractor-accountability governance: vehicle-level performance monitoring through the MSRS system, time-window recalibration (a zero-capital-cost intervention improving system-wide on-time performance by 2.09 percentage points), route-load rebalancing (eliminating all 232 statistically overloaded routes), and deferral of fleet expansion pending revision of vehicle allocation. This finding carries a direct capital-reallocation implication: at an indicative procurement cost in the range of INR 25–40 lakh per collection vehicle, the identified surplus of approximately 260 vehicles represents capital that current evidence suggests would deliver negligible service-quality improvement if spent on further fleet expansion, and that is more productively redirected toward contractor-governance infrastructure - vehicle-level compliance monitoring, performance-linked contracting, and the proof-of-concept Decision Support System itself, whose Phase 1 institutional deployment cost is estimated as a small fraction of the cost of a single additional collection vehicle.

The research additionally demonstrates that a rigorously engineered feature pipeline, combined with a comparatively simple linear model, can outperform substantially more complex deep learning architectures in data-scarce monthly municipal forecasting - a finding with direct practical relevance for resource-constrained municipalities considering where to invest limited technical capacity. Finally, the documented geocoding failure and the resulting GPS-free IMPACT and DSS design demonstrate that Indian municipal data-infrastructure constraints, rather than being obstacles to rigorous optimisation research, can themselves motivate methodologically original, practically deployable alternatives applicable to the large majority of India's Urban Local Bodies that share the same infrastructure gap.

The research is intended to be read not only as a record of what was found for Surat specifically, but as a demonstration of a transferable research and deployment methodology for Indian municipal solid waste management more broadly - one in which forecasting, routing, and decision support are designed together from the outset, in which data-

infrastructure constraints are investigated empirically rather than assumed away, and in which policy recommendations are derived from, and traceable to, quantified operational evidence rather than from conventional assumption. A four-step replication framework is proposed for other Indian Urban Local Bodies seeking to adopt the integrated approach: (i) a data-infrastructure and geocoding-feasibility audit, including a structured sample-based geocoding test analogous to the West Zone pilot conducted in this research, to determine at the outset whether a GPS-dependent or GPS-free architecture is appropriate; (ii) deployment of the 419-feature forecasting framework with city-specific localisation of socio-economic, climatic, and event-calendar inputs; (iii) deployment of the IMPACT framework's five components, with the compliance-attribution model independently re-estimated on local operational data rather than the Surat-specific 75.72% figure being assumed to transfer directly; and (iv) phased Decision Support System deployment beginning with the same GPS-free, low-capital-cost proof-of-concept architecture validated in this research.

8.2 Recommendations for Municipal Authorities

On the evidence developed in this research, the following prioritised actions are recommended for SMC and comparable Indian ULBs: (i) establish contractor-compliance monitoring infrastructure using vehicle-level historical miss-rate data as the primary signal; (ii) implement system-wide time-window recalibration, an administrative, near-zero-capital-cost intervention; (iii) implement route- and workload-balancing before any route-network expansion; (iv) defer fleet-expansion capital expenditure pending revision of vehicle allocation, given the identified fleet surplus; and (v) deploy the proof-of-concept Decision Support System's data-infrastructure layer on existing municipal administrative systems, without requiring GPS hardware or proprietary software, as the foundation for the subsequent deployment phases.

8.3 Limitations

1. The forecasting models are calibrated on a single city's monthly series with 186 observations; although unusually long by Indian standards, this still limits deep learning performance and constrains generalisability.

2. The IMPACT framework is validated on a single month (October 2025) of operational data, so seasonal and inter-annual variations in missed-collection patterns are not yet captured.
3. Geocoding conclusions are based on one large-scale Nominatim pilot on raw operational address strings; alternative address-standardisation methods or commercial geocoding services might yield different success rates.
4. The DSS is validated as a technical proof-of-concept without a full production deployment or longitudinal user studies within SMC.
5. The research is single-city in scope; transferability to other Indian or international cities will require adaptation of feature sets, institutional arrangements, and policy contexts.

8.4 Future Scope

1. Extend the forecasting framework to multi-city and multi-stream panels, incorporating segregation quality and special-waste streams, and explore additional probabilistic and Bayesian modelling approaches.
2. Generalise IMPACT to longer multi-month and multi-season datasets, embed explicit multi-objective optimisation, and use simulation to evaluate alternative contractor-incentive schemes.
3. Integrate richer spatial data sources (for example, standardised addressing, selective GPS pilots, or high-resolution base maps) to improve routing granularity while preserving GPS-free fallback options.
4. Undertake participatory DSS design and full production deployment within SMC, including longitudinal tracking of decision cycles, user adoption, and service outcomes.
5. Adapt the framework for other Indian and developing-country cities, documenting the adjustments required in institutional design, data infrastructure, and policy alignment.

9. List of Papers Published / Presented

Sr. No.	Title	Journal	Indexing	Status
1	A Comparative Study of Advanced Machine Learning and Deep Learning Models for Municipal Solid Waste Forecasting: A Case Study of Surat, India	Journal of Engineering and Technological Sciences (JETS) - ISSN:2337-5779E-ISSN:2338-5502	Scopus Indexed Q3	Published (2026) , DOI: https://doi.org/10.5614/j.eng.tech.nol.sci.2026.58.5.1
2	Data-Driven Automation for Solid Waste Forecasting Using Machine Learning, with COVID-19 Impact Analysis: A Case Study of Surat, Gujarat	Journal of Solid Waste Technology and Management (JSWTM), Vol. 52, Issue 2 - ISSN 1088-1697	Scopus Indexed Q4	Published (2026) , DOI: https://doi.org/10.5276/jswtm/iswmaw/522/2026.515
3	AI-Driven Decision Support System for Optimized Municipal Solid Waste Management: A Machine Learning and Route Optimization Framework for Smart Cities	International Journal of Applied Mathematics (IJAM), Vol. 38, No. 12s, pp. 2465–2488 - ISSN: 1314-8060	Scopus Indexed Q3	Published (2025) , DOI: https://doi.org/10.12732/ijam.v38i12s.1565
4	A Systematic Review of Models for Predicting Municipal Solid Waste Generation for Efficient Waste Management	International Journal of Scientific Development and Research – ISSN 2455-2631	Peer Reviewed Journal	Published (2025)
5	Comparative Study of Machine Learning Techniques for Predictive Modeling: Approaches, Applications, and Performance Evaluation	27 th ISTE Conference on “Revolutionizing Multidisciplinary Technical Education with Innovative Instructional Strategies: Vision 2047 for New Bharat” - ISBN: 978-81-987520-2-4, 312-316.	National Conference	Published (2025)

10. References

- [1] B. H. Abbasi and A. El Hanandeh, “Forecasting municipal solid waste generation using artificial intelligence modelling approaches,” *Waste Management*, vol. 56, pp. 13–22, Oct. 2016.
- [2] A. Rathod and others, “AI-Driven Decision Support System for Optimized Municipal Solid Waste Management: A Machine Learning and Route Optimization Framework for Smart Cities,” *International Journal of Applied Mathematics*, vol. 38, no. 12s, 2025.
- [3] A. Rathod and others, “Data-Driven Automation for Solid Waste Forecasting Using Machine Learning, with COVID-19 Impact Analysis: A Case Study of Surat, Gujarat,” *Journal of Solid*

- Waste Technology and Management*, vol. 52, no. 2, p. 515, 2026, doi: 10.5276/jswtm/iswmaw/522/2026.515.
- [4] A. Rathod and others, "A Comparative Study of Advanced Machine Learning and Deep Learning Models for Municipal Solid Waste Forecasting: A Case Study of Surat, India," *Journal of Engineering and Technological Sciences*, 2026.
 - [5] Y. Zhao *et al.*, "Prediction of municipal solid waste generation and analysis of dominant variables in rapidly developing cities based on machine learning—a case study of China," *Waste Management and Research*, 2024, doi: 10.1177/0734242X231192766.
 - [6] V. Nourani and others, "Predicting municipal solid waste generation using artificial intelligence: A hybrid approach of entropy analysis and SHAP for optimal feature selection," *Waste Management*, vol. 205, p. 115012, 2025, doi: 10.1016/j.wasman.2025.115012.
 - [7] O. Mudannayake, D. Rathnayake, J. D. Herath, D. K. Fernando, and M. Fernando, "Exploring Machine Learning and Deep Learning Approaches for Multi-Step Forecasting in Municipal Solid Waste Generation," *IEEE Access*, vol. 10, pp. 122570–122585, 2022, doi: 10.1109/ACCESS.2022.3221941.
 - [8] S. Wan, T. Mahmud, and others, "Prophet time series modeling of waste disposal rates in the context of COVID-19," *Sustainable Cities and Society*, vol. 105, p. 105205, 2024, doi: 10.1016/j.scs.2024.105205.
 - [9] Maksymilian Cieśla and Anna Kosior, "IMPACT OF COVID-19 PANDEMIC ON AMOUNT OF MUNICIPAL WASTE GENERATION – A CASE STUDY," *JCEEA*, vol. 69, pp. 5–16, Oct. 2022, doi: 10.7862/rb.2022.1.
 - [10] V. D. Muñoz-Soto, S. Dávila-Gálvez, A. E. Pérez, L. Rojo-González, N. Valenzuela-Levi, and Ó. C. Vásquez, "COVID-19, waste production and municipal recycling programs: Insights from Chile to the global south," *Science of The Total Environment*, vol. 899, p. 165388, Nov. 2023, doi: 10.1016/j.scitotenv.2023.165388.
 - [11] D. Kong, C. Liu, and L. Han, "COVID-19 lockdown and municipal solid waste: evidence from the discarding records of 252 communities in China," *Envir. Dev. Econ.*, vol. 28, no. 4, pp. 307–331, Aug. 2023, doi: 10.1017/S1355770X22000353.
 - [12] Q. Qiao, F. Tao, H. Wu, X. Yu, and M. Zhang, "Optimization of a Capacitated Vehicle Routing Problem for Sustainable Municipal Solid Waste Collection Management Using the PSO-TS Algorithm," *IJERPH*, vol. 17, no. 6, p. 2163, Mar. 2020, doi: 10.3390/ijerph17062163.
 - [13] A. M. Restrepo-Franco, O. Valencia-Rodríguez, and E. M. Toro-Ocampo, "The vehicle routing problem as applied to residential solid waste collection operations: Systematic literature review," *10.5267/j.ijiec*, vol. 16, no. 1, pp. 197–220, 2025, doi: 10.5267/j.ijiec.2024.10.002.
 - [14] M. Mes, M. Schutten, and A. Pérez Rivera, "Inventory routing for dynamic waste collection," *Waste Management*, vol. 34, no. 9, pp. 1564–1576, Sep. 2014, doi: 10.1016/j.wasman.2014.05.011.
 - [15] I. Markov, M. Bierlaire, J.-F. Cordeau, Y. Maknoon, and S. Varone, "Waste collection inventory routing with non-stationary stochastic demands," *Computers and Operations Research*, vol. 113, p. 104798, 2020, doi: 10.1016/j.cor.2019.104798.
 - [16] V. Sanjeevi and P. Shahabudeen, "Optimal routing for efficient municipal solid waste transportation by using ArcGIS application in Chennai, India," *Waste Management and Research*, vol. 34, no. 1, pp. 11–17, 2016, doi: 10.1177/0734242X15607430.
 - [17] P. Kumar, R. Singh, S. SHUKLA, and S. Behera, "A GIS-based Route Optimization Approach for Municipal Solid Waste Management," *Indian Journal of Environmental Protection*, vol. 42, pp. 1616–1623, Oct. 2022.
 - [18] V. Navin Ganesh, K. Vishnuvardhan, and K. Sakthiprasanth, "Neural-network-driven approach in optimization of municipal solid waste collection integrated with geo-spatial techniques," *Global NEST Journal*, Oct. 2024, doi: 10.30955/gnj.06187.
 - [19] I. Sosunova and J. Porras, "IoT-enabled smart waste management systems for smart cities: A systematic literature review," *IEEE Access*, vol. 10, pp. 73326–73363, 2022, doi: 10.1109/ACCESS.2022.3188308.
 - [20] I. Cárdenas-León, M. Koeva, P. Nourian, and C. Davey, "Urban digital twin-based solution using geospatial information for solid waste management," *Sustainable Cities and Society*, vol. 115, p. 105798, Nov. 2024, doi: 10.1016/j.scs.2024.105798.

- [21]Sanjeev Kumar, “IoT-Based Management, Monitoring, Reporting, and Evaluation (MMRE) Implementation Strategy for Dhanbad Municipal Solid Waste Management and Disposal,” *jisem*, vol. 10, no. 51s, pp. 1064–1075, May 2025, doi: 10.52783/jisem.v10i51s.10625.
- [22]E. B. Tirkolaei, V. Simic, M. Ghobakhloo, B. Foroughi, S. Asadi, and M. Iranmanesh, “Integrated design of a sustainable waste management system with co-modal transportation network: A robust bi-level decision support system,” *Journal of Cleaner Production*, vol. 449, p. 141760, Apr. 2024, doi: 10.1016/j.jclepro.2024.141760.
- [23]I. Sosunova, A. Wolff, J. Porras, and A. Happonen, “Smarter Waste: A decision support framework for smart waste management systems,” *International Journal of Social Ecology and Sustainable Development*, vol. 15, no. 1, pp. 1–18, 2024, doi: 10.4018/IJSESD.361770.
- [24]L. Barth, L. Schweiger, R. Benedech, and M. Ehrat, “From data to value in smart waste management: Optimizing solid waste collection with a digital twin-based decision support system,” *Decision Analytics Journal*, vol. 9, p. 100347, Dec. 2023, doi: 10.1016/j.dajour.2023.100347.