Masters
Program
in **Geospatial**
Technologies

Leveraging Statistical Analysis and Interactive Dashboards for Diesel Consumption Monitoring in United Nations Peacekeeping Operations

Elaheh Torabidashti

Dissertation submitted in partial fulfilment of the requirements
for the Degree of *Master of Science in Geospatial Technologies*

NOVA IMS
INFORMATION MANAGEMENT SCHOOL
UNIVERSITAT JAUME·I
Universität
Münster

Leveraging Statistical Analysis and Interactive Dashboards for Diesel Consumption Monitoring in United Nations Peacekeeping Operations

Dissertation supervised by:

Prof. Dr. Sergi Trilles Oliver
Universitat Jaume I
Castellón, Spain

Co-supervised by:

Prof. Dr. Carlos Granell Canut
Universitat Jaume I
Castellón, Spain

Prof. Dr. Marco Octávio Trindade Painho
NOVA Information Management School
Lisbon, Portugal

STATEMENT OF INTEGRITY

I declare that the work described in this document is my own and not from someone else. All the assistance I have received from other people is duly acknowledged and all the sources (published or not published) are referenced. This work has not been previously evaluated or submitted to NOVA Information Management School or elsewhere. I further declare that I have fully acknowledged the Rules of Conduct and Code of Honor from the NOVA Information Management School.

Elaheh Torabidashti,
Castellón de la Plana, Spain,
20th February, 2026.

Table 1.: Use of Generative Artificial Intelligence

Tasks	**NO**	**YES**	**Generative Artificial Intelligence tools**
Better understand issues related to the research	×		
Summarizing text from bibliography / resources		×	Claude.ai
Summarizing the method(s) used	×		
Translating text	×		
Grammar check		×	Claude.ai
Paraphrase or rewriting text from other people / resources		×	Claude.ai
Coding in R, Python, etc.		×	ChatGPT
Get help on a software		×	ChatGPT
Creating and editing images, maps, videos, etc.	×		
Data analysis	×		
Specify below other tasks not mentioned above:			

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my beloved family for their unwavering support, endless encouragement, and constant belief in my abilities. Their patience, sacrifices, and confidence in me have been the foundation of my strength throughout this journey. Without their love and motivation, this achievement would not have been possible.

I am profoundly grateful to my supervisor, Prof. Sergi Trilles, for his exceptional guidance, insightful advice, and invaluable support throughout the development of this thesis. His constructive suggestions and continuous motivation have significantly shaped both this research and my growth as a scholar. I am sincerely thankful for his dedication and mentorship during every stage of this work.

I would also like to extend my sincere appreciation to my co-supervisors, Prof. Carlos Granell and Prof. Marco Painho, for their constructive feedback, thoughtful recommendations, and kind encouragement.

My heartfelt gratitude goes to the United Nations experts whose guidance and generous willingness to share their expertise have been truly invaluable. Their professional insights and extensive experience have significantly enriched the quality and depth of this research.

I am especially grateful to the Erasmus Mundus Master in Geospatial Technologies program for offering me this remarkable opportunity. This program has not only strengthened my academic and professional skills but has also contributed immensely to my personal development. The international environment, diverse perspectives, and academic excellence of the program have been truly transformative.

Finally, I would like to thank my friends and classmates who shared this journey with me. Their companionship, support, and shared experiences made this period both meaningful and memorable. The friendships formed and the moments shared will always remain among the most cherished aspects of this experience.

Abstract

Diesel fuel is a critical and costly resource in operations that rely on vehicle fleets and stationary generators, particularly in mission-based environments where efficiency and accountability are essential. Although IoT-based fuel monitoring systems enable continuous, high-frequency data collection, the resulting large and complex datasets are difficult to interpret using conventional reporting tools, limiting their practical value for operational decision-making. This thesis develops a structured analytical framework, guided by the Cross-Industry Standard Process for Data Mining (CRISP-DM), for IoT-based diesel fuel consumption monitoring using data collected from multiple storage tanks operated by the United Nations Interim Administration Mission in Kosovo (UNMIK). The proposed approach emphasizes robust preprocessing to transform raw sensor measurements into reliable consumption time series, followed by exploratory analysis using Seasonal–Trend decomposition by Loess (STL) to identify long-term trends, seasonal behavior, and irregular variations. Spatial autocorrelation analysis using Global Moran's I is applied to assess whether geographic proximity contributes additional explanatory information. Short-term fuel consumption forecasting is conducted over a 30-day horizon using a baseline statistical model and multiple machine learning and deep learning approaches, including LightGBM, XGBoost, Random Forest, and Long Short-Term Memory (LSTM) networks. The findings demonstrate that systematic preprocessing substantially enhances the reliability and interpretability of subsequent analyses. Spatial dependence across tanks was found to be weak, supporting independent per-tank modeling strategies. Forecasting evaluation reveals a clear performance hierarchy: the statistical Exponential Smoothing (ETS) baseline achieved an R^2 of 0.650, whereas tree-based ensemble models substantially improved predictive accuracy, achieving R^2 values above 0.91. The LSTM model delivered the strongest overall performance, reaching an R^2 of 0.929 and reducing RMSE by approximately 55% compared to the statistical baseline. The analytical outputs are deployed through an interactive Power BI dashboard that consolidates preprocessing outcomes, temporal consumption patterns, and forecasting results into an accessible decision-support tool. Overall, this thesis provides a practical and reusable framework for transforming high-volume IoT fuel data into actionable insights for operational monitoring and planning.

Keywords: IoT-based resource monitoring, diesel fuel consumption, time-series analysis, STL decomposition, feature engineering, machine learning forecasting, LSTM, spatial autocorrelation, CRISP-DM, decision support systems

Sustainable Development Goals (SDG):

Contents

List of Figures

List of Tables

Acronyms

Abbreviation	**Meaning**
API	Application Programming Interface
CRISP-DM	Cross-Industry Standard Process for Data Mining
ETS	Exponential Smoothing
GBDT	Gradient Boosting Decision Tree
GPS	Global Positioning System
ID	Identifier
IoT	Internet of Things
IQR	Interquartile Range
KNN	k-Nearest Neighbors
LightGBM	Light Gradient Boosting Machine
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MSE	Mean Squared Error
NRMSE	Normalized Root Mean Squared Error
RMSE	Root Mean Squared Error
STL	Seasonal–Trend decomposition using Loess
UNGSC	United Nations Global Service Centre
UNMIK	United Nations Interim Administration Mission in Kosovo
UTM	Universal Transverse Mercator
XGBoost	Extreme Gradient Boosting

1. Introduction

In operations heavily reliant on machinery and vehicle fleets, diesel fuel is a critical and costly resource. Effective management of its consumption is paramount for maintaining operational efficiency, controlling costs, and ensuring resource security (Ayebazibwe, 2023). The implementation of a formal fuel management program is recognised as an effective strategy for achieving long-term reductions in fuel expenditure. Through systematic monitoring of vehicle fuel consumption, organizations are able to establish control and management frameworks that facilitate the tracking, conservation, and optimization of fuel-related costs (Gitahi and Ogollah, 2014).

Fuel monitoring systems encounter several operational challenges, including undetected leaks, unauthorized usage or theft, and inefficient consumption patterns. Such issues not only result in direct financial losses but also elevate operational risks and exacerbate environmental impacts. The integration of advanced monitoring technologies and real-time data is therefore essential for enhancing fuel management efficiency, achieving cost savings, and minimizing the environmental footprint (Chaudhari et al., 2024).

The raw data generated by fuel sensors and transaction logs is often voluminous and complex, making it difficult for managers to extract meaningful insights through manual inspection or basic spreadsheets. This data overload can mask critical events, delaying the response to anomalies until substantial losses have already occurred (Vegesna, 2023). This demonstrates the pressing need for an analytical system that can process consumption data, identify significant deviations from the norm, and present findings in an accessible format for non-technical decision-makers.

1.1 Motivation

Fuel consumption represents one of the most significant operational expenses across transportation, logistics, construction, and industrial fleets, and can account for roughly a quarter to one-third of total operating costs in fuel-intensive sectors, rising higher when fuel prices spike (North American Council for Freight Efficiency, 2019; Vashisht, 2022). Effective management and optimization of fuel usage have become critical priorities for organizations seeking to reduce costs, enhance operational efficiency, and minimize environmental impact. Traditional fuel management approaches, which rely on manual tracking and periodic inspections, are increasingly inadequate for modern operations, suffering from inaccuracies, labor intensiveness, and limited real-time visibility into consumption patterns (Abediasl et al., 2023; Vashisht, 2022).

The emergence of Internet of Things (IoT) technology has revolutionized fuel monitoring by enabling continuous, automated data collection through sensor networks deployed across vehicles, equipment, and storage facilities. IoT-based fuel monitoring systems combine sensors for real-time fuel level detection with the Global Positioning System (GPS) for location tracking, providing comprehensive monitoring capabilities through cloud-based data processing and user-friendly interfaces. This data-driven approach improves service quality through analytics and provides accurate insights that enable managers to make effective business decisions. By providing real-time visibility into fuel consumption, detecting anomalies such as theft or leakage, and supporting data-driven decision-making, IoT-enabled fuel monitoring systems offer substantial advantages over conventional methods (Khatun et al., 2019; Anand et al., 2024).

However, the effectiveness of IoT-based fuel monitoring is fundamentally dependent on the quality and reliability of the underlying sensor data. Fuel-level measurements are particularly susceptible to various sources of error, including sensor noise, missing values, and systematic distortions caused by environmental factors and physical phenomena. Without appropriate preprocessing and data quality management, these issues can lead to misleading interpretations, false anomaly detection, and unreliable forecasts (Tech Mahindra, 2020; Teh et al., 2020).

Beyond the operational and financial implications, fuel consumption data can be viewed as a time series exhibiting complex temporal behavior, influenced by operational schedules, environmental conditions, and system-level constraints (Bhuiyan et al., 2024). Such data often display recurring daily and weekly patterns as well as long-term trends reflecting changes in usage intensity or efficiency (Inchauspe et al., 2020). These characteristics are common in many real-world time series, where trend and seasonality provide essential structure for understanding normal behavior and detecting anomalies (Zangrando et al., 2022). Time-series analysis techniques therefore play a crucial role in extracting meaningful patterns and supporting informed decision-making (Hyndman and Athanasopoulos, 2021).

Accurate fuel consumption prediction provides additional insights into how operational strategies impact efficiency and enables cost-effective management at scale. Time-series forecasting has emerged as a powerful tool for understanding and predicting fuel consumption behavior, with applications ranging from short-term operational planning to long-term strategic decision-making (Kabir et al., 2023). Despite its potential, energy forecasting faces significant challenges including data quality issues, model interpretability, and unpredictable external factors such as economic shifts and regulatory changes. The development of robust forecasting models requires careful consideration of data preprocessing, feature engineering, and model selection to balance accuracy, computational efficiency, and generalizability across diverse operational contexts (Hong et al., 2020; Zhou et al., 2023).

This thesis is motivated by the need to develop a comprehensive analytical framework that addresses the full pipeline of fuel consumption analysis—from raw IoT sensor data to actionable forecasts and insights. By integrating advanced data cleaning techniques, time-series decomposition methods, spatial analysis, and state-of-the-art forecasting models, this research aims to demonstrate how data-driven approaches can transform fuel management practices, enabling organizations to optimize consumption, reduce costs, and support sustainable operations.

1.2 Research Questions

This thesis aims to develop and evaluate an analytical framework for IoT-based diesel consumption monitoring that integrates data quality enhancement, advanced analytical methods, forecasting, and interactive visualization. To achieve this aim, the study addresses the following research questions:

1. How can IoT-based diesel consumption data be effectively integrated, cleaned, and validated to ensure reliable analysis? This question focuses on identifying data quality challenges in sensor-based fuel monitoring systems and evaluating preprocessing, correction, and validation techniques that enhance data reliability.

2. What statistical, temporal, spatial, and forecasting-relevant patterns emerge from the cleaned dataset, and how do these patterns reveal trends, seasonality, spatial clustering, and short-term fuel consumption dynamics? This question examines the underlying structure of diesel consumption behavior, including long-term trends, recurring temporal patterns, spatial dependencies, and characteristics relevant for short-term forecasting.

3. How can short-term diesel consumption be forecast using classical statistical, machine learning, and deep learning models, and how do these approaches compare in predictive performance? This question evaluates the effectiveness of different forecasting paradigms and quantifies their performance using appropriate accuracy metrics.

4. How can an interactive Microsoft Power BI dashboard be designed to integrate data quality indicators, analytical insights, and forecasting results, providing actionable information for mission decision-making? This question addresses the translation of analytical and forecasting outputs into an accessible decision-support system for non-technical stakeholders.

1.3 Research Objectives

The overall objective of this thesis is to develop a structured methodology for analyzing IoT-based diesel consumption data that integrates data cleaning, quality control, and advanced analytical techniques including forecasting to extract reliable and actionable insights. The proposed methodology aims to produce a reusable analytical framework, demonstrated through an interactive Power BI dashboard, to support effective decision-making in mission operations. To achieve this main objective, the study pursues the following sub-objectives:

1. To establish a robust data integration and cleaning process for IoT-based diesel consumption data collected within the United Nations Interim Administration Mission in Kosovo (UNMIK), ensuring data reliability and consistency for subsequent analysis.
2. To apply statistical, temporal, spatial, and forecasting analyses to the cleaned dataset in order to identify consumption patterns, long-term trends, seasonal behavior, spatial clustering, and short-term fuel consumption dynamics.
3. To develop and evaluate forecasting models for short-term diesel consumption prediction, comparing classical statistical, machine learning, and deep learning approaches.
4. To design and implement an interactive Power BI dashboard that integrates data quality indicators, analytical insights, and forecasting results to support operational monitoring and decision-making.

1.4 Study Area

The study area of this research is Kosovo, where the UNMIK operates throughout the territory. UNMIK was established in 1999 under United Nations Security Council Resolution 1244 with the mandate to support peacekeeping, civil administration, and institutional development in the region (United Nations Security Council, 1999). Since its establishment, the mission has relied on geographically distributed infrastructure and logistical assets to support its operations.

UNMIK's activities depend heavily on diesel fuel for transportation, backup power generation, and facility operations. Fuel is consumed across multiple operational sites, including administrative offices, vehicle depots, and mission facilities that are spatially dispersed across Kosovo. As a result, fuel consumption behavior varies by location and operational role, making the mission environment suitable for analyzing both temporal consumption patterns and cross-site variability (United Nations, 2023).

Fuel consumption data within UNMIK are collected using IoT-based monitoring systems installed on diesel storage tanks. These systems record time-stamped fuel-level measurements that enable continuous monitoring of consumption, anomaly detection, and performance evaluation. The geographically distributed nature of the monitored sites allows this study to examine not only temporal patterns such as trends and seasonality, but also potential spatial relationships between consumption behaviors across locations. Figure 1.1 shown above illustrates the geographic context of the study area, highlighting the territory of Kosovo.

Figure 1.1.: Study area showing the geographic location of Kosovo and the operational extent of UNMIK. Source: United Nations Geospatial.

By focusing on UNMIK operations in Kosovo, this thesis situates its analysis within a real-world mission environment where reliable fuel monitoring is critical for cost control, operational efficiency, and resource security. The findings of this study are therefore relevant not only to UNMIK, but also to other peacekeeping and large-scale operational missions that rely on diesel fuel and sensor-based monitoring systems.

1.5 Thesis Organization

This thesis is organized into ten chapters, each contributing to the development and evaluation of an analytical framework for IoT-based diesel consumption monitoring.

The first chapter introduces the research context, motivation, and objectives, and defines the scope of the study, including the data sources and study area. It also outlines the structure of the thesis.

The second chapter surveys related work in IoT-based fuel monitoring, consumption forecasting, and visualization-driven decision support, identifying key gaps that motivate the proposed approach.

The third chapter presents the overall methodology adopted in this study. It introduces the analytical workflow and describes the adaptation of the Cross-Industry Standard Process for Data Mining (CRISP-DM) framework to the context of IoT-based diesel consumption analysis, providing a structured foundation for the subsequent chapters.

The fourth chapter details the data cleaning and preprocessing framework applied to the raw sensor data. This chapter focuses on data quality assessment, fuel-level spike detection and classification, parameter calibration, and post-detection cleaning to generate reliable time series for analysis.

The fifth chapter explores the temporal characteristics of fuel consumption through time-series decomposition and exploratory analysis. Trends, seasonal patterns, residual behavior, and short-term consumption dynamics are examined to gain insight into operational fuel usage.

The sixth chapter evaluates whether geographic proximity influences fuel consumption behavior by applying spatial autocorrelation analysis and discussing its implications for modeling and forecasting.

The seventh chapter presents and compares multiple approaches for short-term fuel consumption forecasting, including statistical, machine learning, and deep learning models.

The eighth chapter demonstrates how analytical and forecasting results are integrated into an interactive Power BI dashboard, highlighting its role in supporting operational monitoring and decision-making.

The ninth chapter provides an overall discussion of the results. It synthesizes the findings from preprocessing, temporal analysis, spatial diagnostics, forecasting performance, and visualization, situating them within the broader literature. This chapter critically evaluates methodological choices and interprets the results in an integrated manner.

The final chapter summarizes the key findings of the study, discusses limitations, and suggests directions for future research and system enhancement.

2. Related Works

2.1 Energy Consumption Analysis Using Sensor and IoT Data

The proliferation of IoT technologies has fundamentally transformed fuel consumption monitoring and analysis across diverse operational contexts (Karunya et al., 2024). IoT-based fuel monitoring systems combine distributed sensing devices with wireless data transmission mechanisms and cloud-hosted data infrastructures to support real-time tracking, analysis, and optimization of fuel usage patterns (Hapsari et al., 2021). This section reviews the current state of research and practice in IoT-enabled fuel consumption analysis, with particular emphasis on sensor technologies, data quality challenges, and exploratory data analysis.

2.1.1 IoT-Based Energy Monitoring Systems and Sensor Technologies

This section reviews how recent work leverages heterogeneous sensor streams and IoT infrastructures to enable fine-grained monitoring and analysis of fuel consumption in operational settings. Early IoT-based vehicle fuel monitoring systems focused on integrating level or flow sensors with embedded controllers and wireless communication modules to record refuelling events and consumption patterns in real time, typically exposing this information through web or mobile interfaces for end users. For example, Khatun et al. (2019) design a vehicle fuel activities monitoring system in which an ultrasonic sensor measures tank level, a GPS module records vehicle position, and a Wi-Fi–enabled microcontroller sends data to a cloud backend where it can be inspected through a mobile application, thereby supporting theft detection, nearest fuel-station search, and centralized fleet supervision. Complementary work on smart fuel monitoring systems uses combinations of flow and load sensors connected to Raspberry Pi or similar edge-computing platforms, transmitting both intake and usage measurements to cloud services so that users can track refuelling quantities, infer consumption rates, and receive personalized efficiency recommendations based on historical driving behavior (P et al., 2019).

Beyond basic level sensing, several recent studies emphasise the role of IoT architectures in providing secure, scalable data acquisition pipelines for fuel management applications. Hapsari et al. (2021) proposes a real-time fuel consumption monitoring system that deploys flow sensors and ultrasonic level sensors on-board, with measurements relayed via IoT gateways to a central server that assists operational teams in accurately calculating fuel usage for maritime operations, addressing limitations of manual bookkeeping and intermittent meter readings. Other implementations demonstrate how Wi-Fi–enabled microcontrollers aggregate data from tank-mounted sensing devices to support high-precision, continuous observation of underground fuel storage conditions, with measurements such as fuel level, pressure, and temperature relayed to cloud-hosted data repositories for ongoing access and management (Edirisinghe et al., 2024). Architectures proposed in IoT-based fuel monitoring studies commonly employ lightweight communication protocols and cloud storage to preserve historical fuel and transaction data, which can enable subsequent analytics such as efficiency benchmarking and cost analysis at the vehicle or fleet level (Nita and Mihailescu, 2022).

IoT-based fuel monitoring has also been extended to incorporate multi-sensor context beyond the tank itself, enabling richer interpretations of fuel consumption behavior. Farrag et al. (2025) illustrate how vehicular telemetry, traffic conditions, and fuel consumption indicators, captured through networked on-board units, can be combined with IoT connectivity to support emission monitoring and predictive modeling, demonstrating that fused sensor data enables more accurate associations between operating regimes and fuel usage. Similarly, an IoT-based fuel monitoring and tracking system proposed by recent work integrates accelerometers for real-time fuel level estimation with GPS-based mileage calculation, and forwards these streams to cloud services and mobile applications so that fleet managers can jointly assess route efficiency, tank dynamics, and refuelling behavior (Anand et al., 2024). Such multi-modal sensing configurations underscore a shift from isolated gauge replacement toward comprehensive telematics solutions in which fuel consumption data are treated as part of a broader ecosystem of IoT-enabled operational metrics (Skar et al., 2023).

Across these contributions, researchers consistently report that IoT platforms improve the temporal resolution, reliability, and accessibility of fuel-related measurements compared with conventional mechanical gauges or manual logs (Vlachos et al., 2022). Continuous sensor-based monitoring allows systems to detect anomalous events such as sudden level drops, irregular refuelling patterns, or deviations from expected usage profiles, which can be surfaced through alerts on mobile or web dashboards for rapid investigation (Kadu et al., 2023). Moreover, by centralizing sensor data in cloud repositories, IoT fuel management solutions enable stakeholders to perform historical trend analysis, support auditing of fuel transactions, and prepare the data required for advanced tasks such as forecasting and optimization in subsequent stages of the fuel analytics pipeline (Kee and Simon, 2019).

2.1.2 Data Quality Challenges in IoT-Based Energy Monitoring

Despite the advantages of IoT-based sensing for fuel and energy monitoring, the quality of sensor-generated data remains a critical challenge. IoT sensor data are susceptible to data quality issues arising from multiple sources, including sensor malfunctions, communication disruptions, and environmental conditions (Goknil et al., 2023). Moreover, IoT applications commonly employ low-cost sensors with constrained battery capacity and processing capabilities, frequently deployed in challenging environments. These conditions can result in issues such as reduced measurement accuracy, calibration drift, device malfunction, improper sensor placement, limited sensing range, and data packet loss. Such sensor-related problems introduce errors into the collected data, thereby complicating subsequent analysis (Teh et al., 2020).

Another widely reported data quality issue in energy monitoring is the difficulty of distinguishing meaningful operational changes from spurious fluctuations. Prior studies emphasize that energy monitoring data can exhibit sudden spikes or drops that do not reflect actual consumption behavior, necessitating context-aware validation and domain-informed filtering strategies (Janczura et al., 2013). Without such validation, downstream analyses may misinterpret sensor artifacts as genuine behavioral signals.

The literature consistently highlights preprocessing as a necessary step for improving the reliability and interpretability of IoT-derived data. Techniques such as outlier detection via statistical methods like Interquartile Range and Z-Score, imputation using linear interpolation, and consistency checks are commonly employed to mitigate distortions and stabilize time series prior to analysis (Alsalehy and Bailey, 2025).

A traditional and widely employed technique for anomaly detection in smart grid and IoT-enabled energy monitoring systems is rule-based detection, which depends on predefined logical rules specifying acceptable operational ranges and behavioral patterns. Real-time observations that violate these rules, derived from temporal constraints, operational limits, or expert-defined thresholds, are identified as anomalous events (Peña et al., 2016).

Taken together, these studies highlight that Data quality issues may have a significant impact on the reliability of collected data as well as on the accuracy of downstream analytical results (Bertrand et al., 2023). Low-quality data can result in suboptimal decision-making, underscoring the importance of evaluating the quality of the data that informs such decisions. As a result, systematically identifying, characterizing, and communicating data quality issues across the data lifecycle, from acquisition and preprocessing to analysis and utilization, is essential for generating reliable insights (Byabazaire et al., 2020).

2.1.3 Exploratory Data Analysis of IoT-Based Energy Consumption Data

Exploratory data analysis plays a central role in understanding the structure, behavior, and variability of energy and fuel consumption data collected from IoT monitoring systems. Due to the high frequency and longitudinal nature of sensor-derived time series, exploratory analysis is commonly used as an initial step to reveal underlying patterns, assess data quality, and guide the selection of appropriate analytical and modeling techniques.

In the context of energy monitoring, exploratory analysis often focuses on identifying temporal consumption patterns across multiple time scales. Prior studies show that hourly, daily, and weekly aggregation of energy data can reveal recurring operational cycles, such as work-hour activity, maintenance schedules, and periodic demand fluctuations (Tyralis et al., 2017; Barkhordar et al., 2022). Visualization of these temporal patterns supports improved understanding of system behavior and provides context for subsequent forecasting and anomaly analysis (Tyralis et al., 2017).

Time-series decomposition techniques are frequently incorporated into exploratory analysis to separate long-term trends, seasonal components, and irregular variations. Research has demonstrated that decomposition-based exploration helps distinguish systematic consumption behavior from short-term fluctuations and measurement noise, particularly in energy and fuel monitoring applications (Bhuiyan et al., 2024; Lin et al., 2022).

The literature emphasizes that exploratory analysis should be tightly coupled with data preprocessing and modeling rather than treated as an isolated step. Integrating exploratory analysis into the analytical workflow enables iterative refinement of preprocessing decisions and supports more robust interpretation of underlying patterns (Hyndman and Athanasopoulos, 2021). This perspective motivates the exploratory analysis framework adopted in this study, which builds directly on validated and cleaned IoT-based fuel consumption data.

2.2 Forecasting Fuel and Energy Consumption

Forecasting energy and fuel consumption has been an active research area for several decades, motivated by the need to support operational planning, cost control, and resource management in energy-intensive systems (Eddaoudi et al., 2024). Early studies in energy consumption forecasting predominantly employed classical statistical forecasting techniques, including autoregressive integrated moving average and exponential smoothing methods, to capture temporal dependencies from historical data. These methods have been widely applied to electricity and fuel demand forecasting due to their interpretability and relatively low data requirements (Gardner, 2006; Dritsaki et al., 2021).

While classical statistical models perform well under stable conditions, several studies have shown that their predictive accuracy may be limited when consumption patterns exhibit nonlinear behavior or are influenced by multiple external factors (Vivas et al., 2020). As a result, machine learning techniques have increasingly been adopted for energy and fuel consumption forecasting, particularly in settings where large volumes of high-frequency data are available (Deb et al., 2017; Zhang et al., 2021).

Recent energy forecasting studies have demonstrated that data-driven learning approaches, including gradient boosting, random forests, and support vector regression, can surpass classical statistical methods by effectively modeling nonlinear behavior and complex feature interactions (Vivas et al., 2020; Dou et al., 2023).

Previous studies highlight key differences between statistical and machine learning approaches for energy and fuel consumption forecasting. Statistical models are commonly used for short-term prediction and can perform well with relatively small datasets, but they are typically limited by linear assumptions and strong dependence on historical patterns (Dou et al., 2023). In contrast, machine learning methods are better suited to multivariate forecasting problems, as they handle multivariate inputs and capture nonlinear relationships in high-dimensional data, when sufficient training data are available. Despite their strengths, both modeling paradigms face challenges in maintaining predictive accuracy over longer forecasting horizons, underscoring the importance of data quality and appropriate feature design (Dou et al., 2023).

In the context of energy demand prediction, these models have been successfully applied to hourly electricity consumption data, demonstrating improved short-term forecast accuracy compared to univariate models (Zafeiriou et al., 2023; Hu et al., 2024).

Feature engineering is a critical factor influencing the performance of machine learning–based forecasting models. Prior research consistently emphasizes that the inclusion of lagged consumption values, rolling statistics, and calendar-related features significantly improves predictive performance by encoding temporal dependencies and cyclical behavior (Wilfling, 2023; Karmaker, 2025). For energy systems affected by environmental conditions, exogenous variables such as ambient temperature have also been shown to be among the most informative predictors, particularly for short-term forecasting horizons (Zafeiriou et al., 2023).

Although research on energy and fuel consumption forecasting has expanded considerably, many studies focus primarily on predictive accuracy without fully integrating data quality assessment, preprocessing, and interpretability considerations into the forecasting workflow. (Shaikh et al., 2022). Moreover, relatively few works present end-to-end frameworks that combine robust data preparation, feature engineering, model comparison, and deployment within an operational decision-support context. These gaps motivate the forecasting framework proposed in this study, which emphasizes data reliability, transparent feature design, and comparative evaluation across statistical, machine learning, and deep learning models.

2.3 Visualization and Decision Support in Monitoring Systems

Visualization plays a critical role in transforming complex sensor-derived data into interpretable information that supports operational monitoring and decision-making. In monitoring systems that generate high-frequency and high-volume data, effective visualization helps reduce cognitive load and enables users to identify patterns, trends, and deviations that may not be apparent from raw data or tabular summaries (Keim et al., 2010).

Dashboards are commonly used as decision-support tools that integrate multiple data views, such as temporal trends, aggregated summaries, and contextual information, into a unified interface (Yigitbasioglu and Velcu, 2012). Several studies emphasize that visualization is most effective when combined with preprocessed and validated data, as raw sensor streams often contain noise and irregularities that can obscure meaningful signals (Zankl, 2009). Presenting cleaned and structured data enhances interpretability and increases user trust in analytical outputs, which is essential for operational decision support (Valdivia and Baca, 2025). As a result, visualization is increasingly positioned as the final stage of an analytical pipeline rather than a standalone component.

In the context of fuel consumption management, Vlachos et al. (2022) introduced an IoT-based vehicle monitoring system featuring a web-based dashboard for real-time visualization of parameters such as fuel level and acceleration. The dashboard provides end-users with intuitive access to live and historical sensor data through graphical interfaces, supporting situational awareness and operational monitoring. By enabling the identification of abnormal conditions and the exploration of long-term performance trends, the system contributes to maintenance planning and fleet-level operational optimization.

Similarly, Khatun et al. (2019) proposed a Vehicle Fuel Activities Monitoring System incorporating a dashboard interface for real-time visualization of fuel and vehicle data. While the system emphasizes usability and centralized monitoring through graphical representations, it primarily focuses on descriptive visualization and does not address automated anomaly detection or predictive analytics. As such, the approach highlights the role of dashboards in operational awareness but leaves opportunities for more advanced data-driven decision support.

2.4 Summary of Literature and Research Direction

The reviewed literature demonstrates that IoT-based fuel and energy monitoring systems enable high-resolution data collection and provide opportunities for improved consumption analysis and forecasting. Prior studies have shown that sensor-derived energy data can support temporal pattern discovery, short-term forecasting, and monitoring applications when appropriate analytical methods are applied. However, the literature also consistently highlights challenges related to data quality, data volume, and interpretability, which limit the direct usability of raw IoT data for analysis and decision support.

Research on fuel and energy consumption forecasting indicates that classical statistical models remain effective for short-term prediction under stable conditions, while machine learning and deep learning approaches offer improved flexibility and predictive accuracy in complex and nonlinear settings. At the same time, existing studies emphasize the critical role of feature engineering, preprocessing, and careful model evaluation in achieving reliable forecasting performance. Despite these advances, many forecasting-focused works treat data preparation as a secondary step and provide limited discussion of preprocessing decisions and their impact on downstream analysis.

The literature on visualization and decision-support systems underscores the importance of integrating analytical results into interactive dashboards to enhance interpretability and support operational monitoring. Nevertheless, many existing monitoring solutions emphasize data presentation, while offering limited support for integrating robust data preparation, systematic exploratory assessment, and forecasting capabilities within a cohesive analytical pipeline.

Overall, the reviewed studies reveal a gap in end-to-end methodologies that explicitly integrate data quality handling, exploratory analysis, forecasting, and visualization for IoT-based fuel consumption data. In response to this gap, the present study proposes a comprehensive analytical framework that emphasizes systematic preprocessing, transparent exploratory analysis, comparative forecasting, and dashboard-based deployment. By addressing these aspects within a single, coherent workflow, this research aims to improve the reliability, interpretability, and practical usability of IoT-based diesel fuel consumption data for operational decision-making.

3. Methodology

This chapter presents the overall methodology adopted in this study to analyze IoT-based diesel fuel consumption data. The methodology provides a structured and systematic framework that integrates data acquisition, preprocessing, analytical modeling, forecasting, and visualization to support reliable analysis and operational decision-making. The approach is designed to ensure data quality, extract meaningful patterns, and translate analytical results into actionable insights.

The analytical process is guided by the CRISP-DM methodology, which provides a widely accepted six-phase framework for data-driven projects. In this study, the CRISP-DM phases are adapted to the context of analyzing IoT-based diesel fuel consumption data from UNMIK.

3.1 CRISP-DM Framework Adaptation

3.1.1 Business Understanding

The primary objective of this study is to develop a robust methodology for analyzing IoT-based diesel fuel consumption data that supports reliable monitoring, pattern discovery, short-term forecasting, and decision-making in mission operations. Operational challenges such as sensor noise, refilling events, anomalous consumption, and limited interpretability of raw data motivate the need for a structured analytical approach. The study addresses these challenges by integrating data quality control, temporal and spatial analysis, forecasting models, and interactive visualization.

3.1.2 Data Understanding

The UNMIK mission fuel consumption dataset serves as the primary data source for this study. The dataset consists of high-frequency time-stamped fuel-level measurements collected from multiple diesel storage tanks, along with associated spatial information (e.g., geographic coordinates), tank capacities, and derived consumption metrics. Initial data exploration focused on understanding the structure and temporal coverage of the data, measurement resolution, and identifying potential quality issues such as missing values, unrealistic readings, noise spikes, and abrupt level changes.

Insights gained during this phase informed the design of the data cleaning, validation, and transformation procedures applied in subsequent stages.

3.1.3 Data Preparation

Data preparation constitutes a central component of the methodology. Raw fuel-level measurements were cleaned, validated, and transformed to ensure analytical reliability. This phase included data quality assessment, detection and classification of fuel-level spikes, empirical calibration of validation parameters, and post-detection cleaning through interpolation. Temporal alignment across data sources was ensured, and additional features were derived to support exploratory analysis and forecasting.

The detailed data cleaning and preprocessing framework is presented in Chapter 4.

3.1.4 Modeling

Following data preparation, multiple analytical models were applied to extract patterns and generate forecasts. Temporal decomposition using Seasonal-Trend Decomposition by LOESS (STL) was employed to separate long-term trends, seasonal components, and residuals. Spatial analysis was conducted using Global Moran's I to assess spatial dependence in fuel consumption behavior. For short-term forecasting, a range of models were implemented, including a baseline statistical approach, machine learning models, and a deep learning model, to predict hourly fuel consumption over a 30-day horizon.

The modeling techniques and their configurations are described in Chapters 5, 6, and 7.

3.1.5 Evaluation

Model evaluation was performed using both quantitative and qualitative criteria. Forecasting models were assessed using standardized error metrics to enable comparative evaluation of predictive performance. Analytical outputs were further evaluated through visual inspection of decomposition results, residual behavior, anomaly detection outcomes, and consistency across tanks. This evaluation phase ensured that the applied methods produced reliable, interpretable, and operationally meaningful results.

3.1.6 Deployment

The final phase of the CRISP-DM framework focuses on deploying analytical results in a form that supports practical use. In this study, deployment is realized through the development of an interactive Power BI dashboard that integrates data cleaning, Exploratory analysis results, and forecasting outputs. The dashboard enables users to explore fuel consumption behavior, validate analytical results, and support informed operational decision-making.

By embedding the analytical process within the CRISP-DM framework, this study ensures methodological rigor, coherence, and practical applicability, delivering actionable insights into IoT-based diesel fuel consumption patterns in mission environments.

3.2 Methodological Framework Overview

The proposed methodology follows a sequential and modular workflow, where each stage builds upon the outputs of the previous one. This structure ensures that data quality issues are addressed early in the analytical process and that all subsequent analyses are performed on validated and physically consistent data.

Figure 3.1 illustrates the overall analytical workflow adopted in this study. The workflow begins with data collection from IoT-based fuel monitoring systems and progresses through data cleaning and preprocessing, temporal and statistical analysis, spatial analysis, short-term forecasting, and interactive visualization. Each main block of the workflow is briefly described below, while detailed methodological explanations are provided in the subsequent chapters.

3.2.1 Data Acquisition and Initial Exploration

The workflow begins with the collection of high-frequency fuel-level data from IoT-based sensors installed in diesel storage tanks. The dataset includes time-stamped fuel-level measurements, spatial coordinates, and tank-specific attributes. Initial exploration is performed to understand the temporal coverage, measurement resolution, and general behavior of the sensor data, as well as to identify potential issues such as missing values, abrupt spikes, or unrealistic readings.

3.2.2 Data Preparation and Preprocessing

The data preparation stage focuses on cleaning and validating the raw sensor measurements. Initial filtering is applied to remove obvious inconsistencies, followed by spike detection and classification to distinguish between genuine refilling events and noise. Model parameters for spike detection are calibrated empirically to ensure robust performance across tanks. Detected noise points are then corrected through interpolation, resulting in a physically consistent fuel- level time series suitable for further analysis.

3.2.3 Temperature Correction and Time-Series Modeling

After cleaning, the fuel-level measurements are corrected for thermal expansion effects using ambient temperature data. This step ensures that variations caused by environmental temperature changes do not distort the actual consumption signal. The corrected series is then used for time-series modeling, including STL-based de-

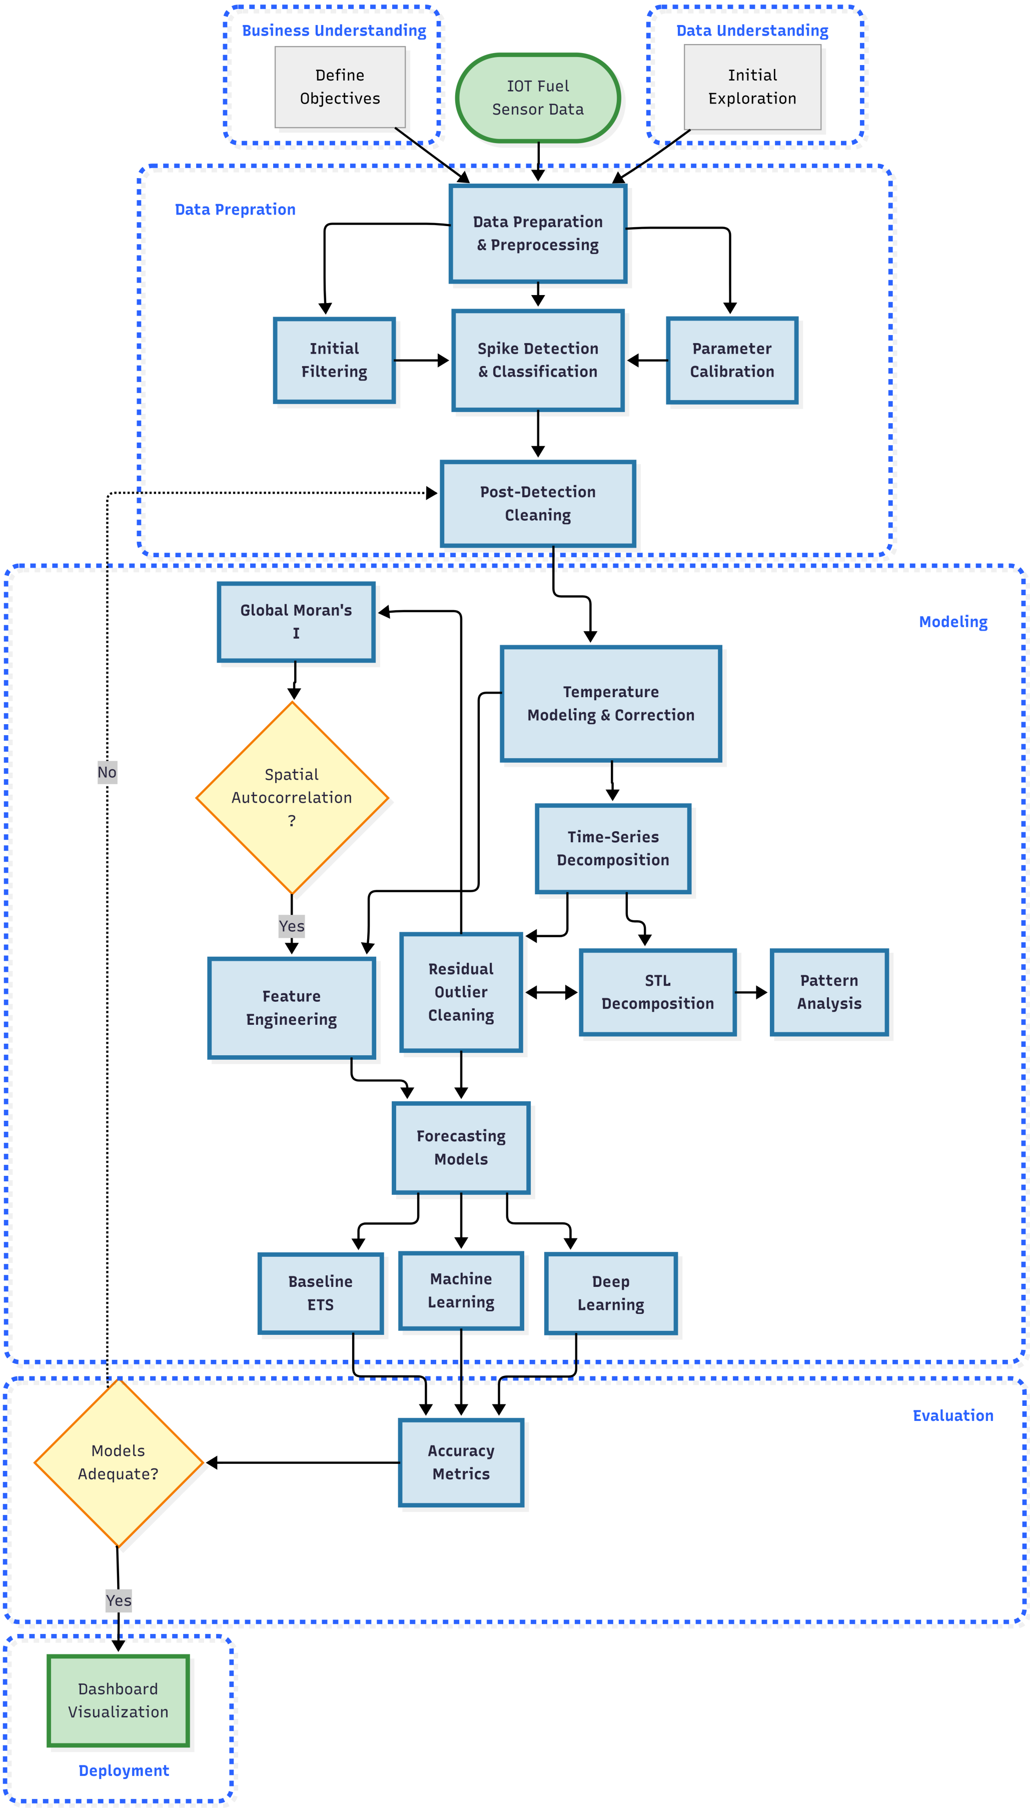

Figure 3.1.: Overall methodology workflow adopted in this study, structured according to the CRISP-DM framework.

composition to separate trend, seasonal, and residual components, as well as pattern analysis of consumption dynamics.

To improve model stability, extreme residual outliers are detected using an Interquartile Range (IQR)–based rule and replaced through interpolation. This step produces a smoothed and physically consistent consumption series. The resulting cleaned time series, obtained after residual-based outlier correction, serves as the final analytical input for subsequent modeling stages, including spatial autocorrelation analysis and short-term forecasting.

3.2.4 Spatial Analysis

To evaluate whether geographic proximity influences fuel consumption behavior, a spatial autocorrelation analysis is conducted using Global Moran's I. This step assesses whether nearby tanks exhibit similar consumption patterns. Depending on the results, spatial relationships may be incorporated into feature engineering for forecasting models.

3.2.5 Forecasting Models

Short-term forecasting models are developed using the cleaned and temperature-corrected consumption data. Three categories of models are considered: a baseline statistical model, machine learning models, and a deep learning model. All models are trained to predict hourly fuel consumption over a 30-day forecasting horizon using engineered temporal and environmental features. Feature engineering transforms the raw time series into a supervised learning format by creating lagged consumption variables, rolling statistics, and calendar-based attributes such as hour of day and day of week.

Ambient temperature variables are also included as exogenous predictors to account for environmental effects on fuel consumption. Together, these features capture recent consumption behavior, recurring operational cycles, and temperature-driven variations, improving short-term forecasting accuracy.

3.2.6 Evaluation and Deployment

Model performance is evaluated using standard forecasting accuracy metrics. The results are compared across modeling approaches to identify the most reliable method for short-term prediction. The final analytical outputs, including consumption trends, anomaly indicators, and forecasting results, are integrated into an interactive Power BI dashboard to support operational monitoring and decision-making.

3.3 Data Source and Description

The dataset employed in this study was obtained from the United Nations Global Service Centre (UNGSC) and was accompanied by a data dictionary to facilitate correct interpretation of the recorded variables. The dataset contains time-series measurements of fuel tank levels recorded within UNMIK operations.

The dataset comprises a total of 466,696 records collected from 10 fuel tanks distributed across 8 duty stations within the Kosovo mission. Measurements were recorded at an hourly interval, providing high temporal resolution suitable for detailed trend analysis and consumption assessment. The time span of the dataset ranges from 1 April 2020 to 30 September 2025.

Each record corresponds to a single timestamped measurement and includes metadata related to the site, device, mission, and geographical location. The dataset captures both raw sensor readings and derived values such as fuel level estimates. This comprehensive structure enables robust data cleaning, exploratory analysis, and subsequent modeling.

To support transparency and reproducibility, Table 3.1 presents the data dictionary describing each variable included in the dataset.

Table 3.1.: Data Dictionary for the UNMIK Fuel Tank Level Dataset

Column Name	**Description**
SITE_DEVICE_TAG	Unique identifier for each tag by site and device
TAG_NAME	Name of the tag
TAG_DISPLAY_UNIT	Unit in which the raw tag value is displayed
DEVICE_ID	Unique identifier for each device
DUTYSTATION_NAME	Name of the duty station
SITE_ID	Unique identifier for each site or location
SITE_NAME	Name of the site
TAG_VALUE_TIMESTAMP	Timestamp corresponding to the tag value
TAG_VALUE_RAW	Raw recorded value of the tag
MISSION_NAME	Name assigned to the mission
Date_Hour_Desc	Date and time value of the tag measurement
TAG_POWERBI_NAME	Measurement type used for reporting and visualization
AVG_MEAN	Average mean of the recorded values
STD_DEVIATION	Standard deviation of the recorded values

Column Name	**Description**
DATA_INTERVAL	Time interval between consecutive readings
LevelLiters	fuel level inside the tank in liters
TotalLiters	Maximum tank capacity in liters
UPDATED_AT_	Timestamp indicating when the data was updated in the framework
PREV_TAG_VALUE_RAW	Previous raw sensor reading
OGI_LONG	Longitude coordinate of the site
OGI_LAT	Latitude coordinate of the site

Figure 3.2 presents the spatial distribution of fuel tank locations across UNMIK duty stations in Kosovo. The figure serves as a contextual reference for the dataset, highlighting the geographic extent of the study area and the spatial arrangement of the monitored tanks used in subsequent analyses.

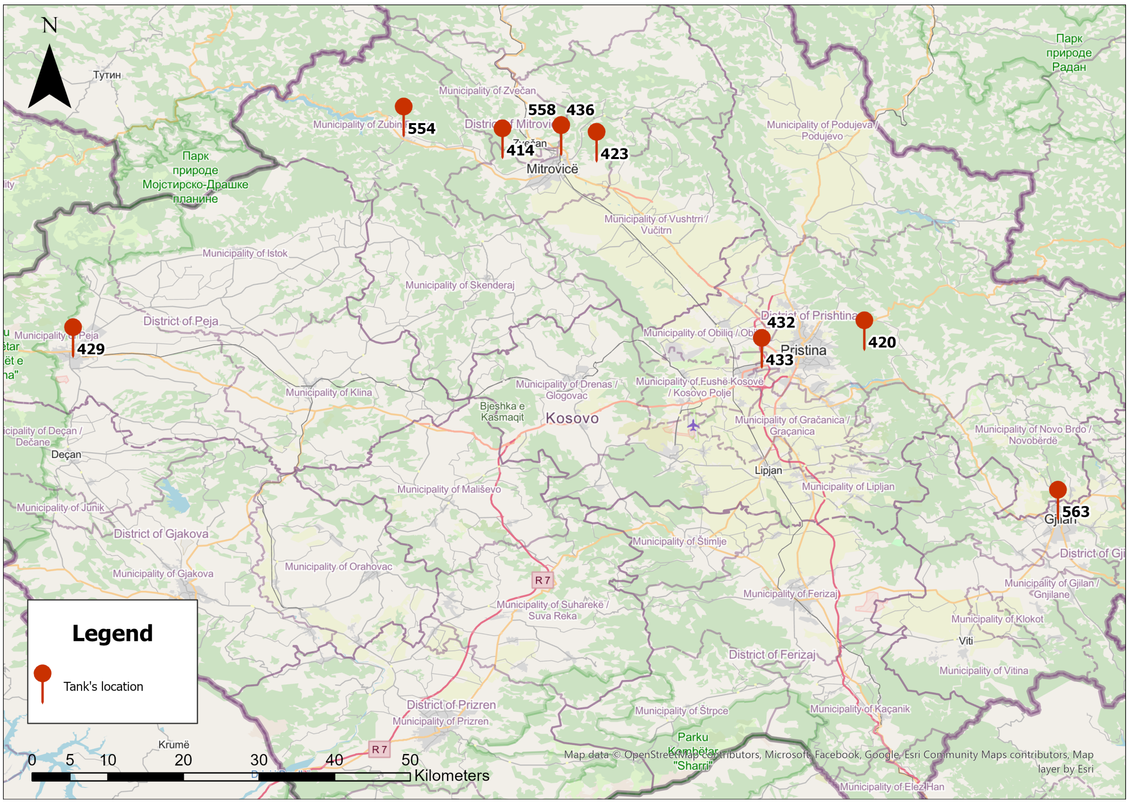

Figure 3.2.: Spatial distribution of fuel tank locations across UNMIK duty stations in Kosovo. Tank identification numbers are shown as labels to distinguish individual tank locations. Basemap source: OpenStreetMap contributors.

4. Data Cleaning and Preprocessing

4.1 Overview of the Data Cleaning Framework

The reliability of any analytical or forecasting task based on sensor data is fundamentally dependent on the quality and consistency of the underlying measurements. IoT-based fuel monitoring systems, while enabling high-frequency and continuous data collection, are particularly susceptible to noise, missing values, sensor drift, and physically implausible readings (Tech Mahindra, 2020; Teh et al., 2020). Consequently, a structured and systematic data cleaning framework is required to ensure that the fuel consumption data used in subsequent analyses accurately reflects real operational behavior rather than measurement artefacts.

The data cleaning framework adopted in this study is designed as a multi-stage pipeline that progressively transforms raw fuel-level sensor readings into a validated and analysis-ready time series. The framework integrates data quality assessment, anomaly detection, event classification, and corrective processing, ensuring that both random noise and systematic errors are addressed. Rather than applying generic filtering techniques, the framework incorporates domain knowledge about fuel tank behavior and operational processes, allowing anomalous observations to be interpreted within their physical and operational context.

The first stage of the framework focuses on an initial data quality assessment to identify missing values, invalid measurements, and obvious inconsistencies in the raw data. This includes detecting fuel-level readings that violate physical constraints, such as negative values or levels exceeding tank capacity. Records identified as invalid are removed or excluded based on clearly defined criteria, ensuring that subsequent processing steps are not influenced by gross measurement errors.

Following the initial filtering, the framework addresses sudden changes in fuel-level measurements through spike detection and classification. Abrupt changes may correspond to legitimate operational events, such as refueling, or may arise from sensor noise or transmission errors. To distinguish between these cases, candidate spikes are first detected using statistical methods applied to consecutive fuel-level differences. These candidates are then classified as either genuine refilling events or noise based on magnitude and persistence criteria, ensuring that real operational behavior is preserved while spurious fluctuations are removed.

Once potential refilling events and noise-induced spikes are identified, the framework applies empirical parameter calibration and stability validation to ensure robustness across different tanks and operating conditions. Key thresholds and temporal windows are calibrated using representative data, allowing the framework to adapt to variations in tank capacity and consumption behavior. Additional stability checks are incorporated to eliminate delayed false positives that may appear stable in the short term but revert to pre-event levels.

In the final stage, data points classified as noise are corrected through interpolation to restore continuity in the time series while preserving the underlying consumption trend. The resulting cleaned dataset represents a physically consistent and temporally stable fuel level signal suitable for exploratory analysis, time-series decomposition, forecasting, and visualization. Figure 4.1 illustrates the multi-stage data cleaning framework applied in this study, transforming raw IoT fuel-level measurements into a validated and analysis-ready time series.

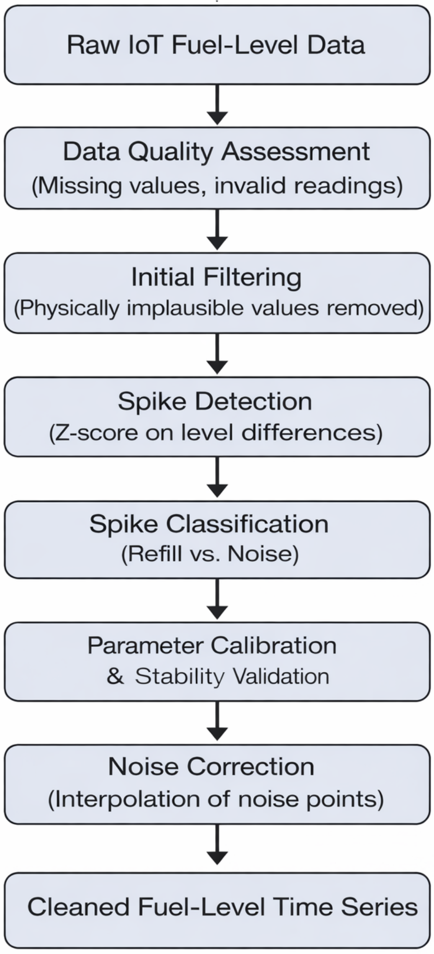

Figure 4.1.: Workflow of the data cleaning and preprocessing framework applied to IoT-based fuel-level data.

By structuring the cleaning process as a transparent and reproducible pipeline, the framework provides a reliable foundation for all subsequent analytical stages of this study.

4.2 Data Quality Assessment and Initial Filtering

Before applying advanced analytical and forecasting techniques, an initial data quality assessment was conducted to evaluate the integrity and reliability of the raw IoT-based fuel-level dataset. This step is essential for identifying issues that could compromise subsequent analyses, such as missing values, physically implausible measurements, and systematic sensor errors. The objective of this stage is to remove or exclude clearly invalid observations while preserving genuine operational behavior.

4.2.1 Missing Value Assessment

The first component of the data quality assessment involved examining the dataset for missing values across all recorded variables. The analysis revealed that the column PREV-TAG-VALUE-RAW contained nine missing entries. Since this variable is not required for the fuel-level or fuel consumption analysis conducted in this study, these missing values were not considered to pose a risk to the reliability of subsequent processing steps. Consequently, no imputation was applied for this variable, and it was excluded from further analysis. All other key variables required for fuel-level monitoring and consumption estimation were found to be complete.

4.2.2 Identification of Physically Implausible Readings

Following the missing value assessment, the dataset was examined for measurements that violate physical constraints of fuel tank operation. In particular, fuel-level values were checked against two fundamental conditions: non-negativity and tank capacity limits. No records were found with negative values in the LevelLiters variable, indicating that the sensor system did not produce invalid negative fuel-level measurements.

However, the assessment identified 186 records in which the recorded fuel level exceeded the corresponding TotalLiters tank capacity. These violations were traced to two specific tanks, with device identifiers 432 and 563. Further investigation revealed that these same records were associated with negative values in the raw sensor signal (TAG-VALUE-RAW). Since LevelLiters is directly derived from this raw signal, invalid raw measurements propagate into unrealistic fuel-level values, resulting in apparent capacity violations.

The extent of these invalid readings was quantified for each affected tank. Tank 563 contained 19 invalid observations out of 45,817 total records (approximately 0.04%), while tank 432 contained 167 invalid observations out of 47,501 records (approximately 0.35%). Given that these anomalies represent a very small fraction of the overall dataset, their removal was deemed unlikely to affect the underlying consumption patterns or introduce bias.

4.2.3 Initial Filtering and Data Exclusion

Based on the results of the data quality assessment, all records exhibiting physically implausible fuel-level values, specifically those exceeding tank capacity due to invalid raw sensor readings, were removed from the dataset prior to further analysis. This initial filtering step ensures that subsequent procedures, such as spike detection, refill classification, and time-series modeling, operate on data that is physically consistent and free from gross measurement errors.

By explicitly separating this initial filtering stage from later anomaly detection steps, the framework distinguishes between structural data quality issues and operational anomalies. The resulting dataset provides a clean and reliable foundation for detecting meaningful fuel-level changes and analyzing consumption behavior in later stages of the methodology.

4.3 Detection and Classification of Fuel Level Spikes

After initial data quality assessment and filtering, the next step in the preprocessing framework focuses on identifying abrupt changes in fuel-level measurements and distinguishing genuine operational events from sensor-induced noise. Fuel-level time series collected from IoT sensors often contain sudden jumps caused by refueling operations, sensor errors, or transmission artefacts. Treating all such spikes as anomalies would distort the analysis, as refueling events represent legitimate and essential components of normal tank behavior. Therefore, a two-stage approach is adopted: first detecting candidate spikes, and then classifying them as either refills or noise.

4.3.1 Spike Detection Using Level Differences

Spike detection is performed on the difference between consecutive fuel-level measurements rather than on the absolute fuel-level values. This transformation emphasizes sudden changes while reducing sensitivity to gradual consumption trends. For each tank, the difference series is computed as:

$$\Delta L_t = L_t - L_{t-1} \tag{4.1}$$

where Ł$_t$ represents the fuel level at time t. The resulting difference series is then analyzed using Z-score–based anomaly detection, applied independently for each tank to account for differences in consumption behavior and measurement variability.

To quantify how strongly each fuel-level change deviates from typical behavior, Z-score normalization is applied to the difference series. The Z-score for each observation is computed as:

$$Z_t = \frac{\Delta L_t - \mu_{\Delta L}}{\sigma_{\Delta L}} \tag{4.2}$$

where ΔL_t is the fuel-level difference at time t, and $\mu_{\Delta L}$ and $\sigma_{\Delta L}$ denote the mean and standard deviation of the difference series for a given tank. Observations with absolute Z-score values exceeding a predefined threshold were flagged as spike candidates for further evaluation.

This statistical screening step does not immediately classify events as anomalies but instead serves as a filtering mechanism to isolate sudden level changes that warrant further evaluation. By applying Z-score detection on a per-tank basis, the method adapts to tank-specific variability and avoids bias caused by differences in scale or operating conditions.

4.3.2 Conceptual Distinction Between Refills and Noise

Detected spikes can originate from two fundamentally different sources. Refilling events correspond to intentional operational actions in which fuel is added to the tank, producing a sudden increase in fuel level that persists over time and is followed by gradual consumption. In contrast, noise spikes are typically short-lived and arise from sensor malfunction, communication errors, or transient environmental effects. Such spikes often revert quickly to previous levels and do not reflect actual fuel usage.

Because refilling events are part of normal tank behavior, they must be preserved in the dataset. Noise spikes, however, must be corrected or removed to prevent distortion of consumption trends, seasonality analysis, and forecasting performance. Consequently, classification criteria must reliably separate these two cases.

4.3.3 Refill Identification Criteria

Each detected spike candidate is evaluated against two criteria to determine whether it represents a genuine refill event:

1. Jump magnitude criterion: The observed increase in fuel level must exceed a minimum threshold, defined relative to the tank's total capacity. This ensures that small fluctuations caused by sensor noise are not misclassified as refills.

2. Persistence criterion: Following the spike, the fuel level must remain elevated for a sustained period. A genuine refill is expected to result in a stable higher level that decreases gradually due to normal consumption, whereas noise spikes typically collapse back to pre-spike levels within a short time.

Only spikes that satisfy both conditions are classified as refilling events. Spikes that fail to meet either criterion are classified as noise and flagged for correction in subsequent processing stages.

4.4 Empirical Parameter Calibration and Stability Validation

The accurate classification of fuel-level spikes as either genuine refilling events or noise depends critically on the selection of appropriate threshold values and temporal parameters. Rather than relying on arbitrary or fixed thresholds, this study adopts an empirical calibration approach, grounding all key parameters in observed refill behavior from the dataset. This strategy ensures that the detection framework reflects real operational characteristics and remains robust across different tanks.

4.4.1 Capacity-Based Jump Threshold

To ensure consistency across tanks with different sizes, the refill jump threshold is defined as a fraction of the tank's total capacity:

$$\text{refill jump threshold} = \alpha \times C_{\text{tank}} \tag{4.3}$$

where C_{tank} denotes the tank capacity and α is a scaling factor. This formulation allows the same logic to be applied across all tanks while accounting for differences in absolute fuel volume. To empirically determine the value of α, a representative tank with clearly identifiable refill events was selected based on visual inspection of its fuel-level time series. For each manually identified refill event, a refill ratio was computed as:

$$\text{refill\_ratio} = \frac{\text{Jump in LevelLiters}}{\text{Tank capacity (TotalLiters)}} \tag{4.4}$$

The minimum refill ratio observed across all confirmed refilling events was selected as the value of α. This conservative choice ensures that genuine refills are retained while small, noise-induced fluctuations are excluded. The derived scaling factor was subsequently applied to all tanks in the dataset.

4.4.2 Persistence-Based Validation of Refill Events

While jump magnitude is a necessary condition for identifying refills, it is not sufficient on its own to distinguish refueling events from sensor noise. Genuine refills are characterized by persistence, meaning that the elevated fuel level remains stable for a sustained period before gradually decreasing due to normal consumption.

Two parameters define the persistence window:

- **window_after (in hours)**: Defines the length of the time window following a detected spike during which fuel-level stability is evaluated.
- **persist_points (number of readings)**: Specifies the minimum number of observations within the persistence window that must remain above a defined threshold for the event to qualify as a valid refill.

The future evaluation window is defined as:

$$[t_{\text{spike}},\ t_{\text{spike}} + \text{window\_after}] \tag{4.5}$$

A spike is classified as persistent if the following condition is satisfied:

$$\sum_{t \in \text{window\_after}} \Big[L_t > L_{\text{prev}} + \text{persistent\_threshold}\Big] \geq \text{persist\_points} \tag{4.6}$$

where L_{prev} denotes the fuel level immediately before the spike.

The persistence threshold is defined as:

$$\text{persistent\_threshold} = \max\big(0.01 \times C_{\text{tank}},\ 0.1 \times \text{jump}\big) \tag{4.7}$$

This formulation ensures adaptability to both tank capacity and refill magnitude.

4.4.3 Calibration of Persistence Parameters

The parameters *window_after* and *persist_points* were calibrated empirically using manually identified refill events from a representative tank. Refill timestamps were selected through visual inspection of the fuel-level time series, where each refill was characterized by a sudden increase followed by a stable high level and gradual consumption.

For each refill, the duration (in hours) during which the fuel level remained stable before decreasing was measured. The median of these durations was selected as the value of *window_after*. For example, when most refills remained stable for approximately 25–35 hours, the persistence window was set to:

$$\text{window\_after} = 30 \text{ hours}$$

Given the hourly sampling frequency, this window corresponds to approximately 30 observations. To allow for minor sensor noise, a refill was required to satisfy the persistence condition for at least 25% of the window. Empirical testing indicated that setting:

$$\text{persist\_points} = 8$$

provided a suitable balance between sensitivity and false detections.

4.4.4 Reverse Window Validation for False Spike Detection

Although the persistence check ensures short-term stability following a spike, certain false refills may appear stable initially but revert to pre-spike levels after a delay. To address this issue, a reverse window validation step was introduced to confirm long-term stability.

After a spike passes the persistence check, an additional reverse window is evaluated, typically spanning 24–48 hours beyond the persistence period. The condition is defined as:

$$\text{ReversedDrop} = \begin{cases} \text{True}, & \text{if } L(t) < L_{\text{prev}} + 0.1 \times C_{\text{tank}} \\ & \text{for any } t \in [t_{\text{spike}} + W_1,\ t_{\text{spike}} + W_1 + W_2] \\ \text{False}, & \text{otherwise} \end{cases} \tag{4.8}$$

where W_1 denotes the persistence window and W_2 the reverse window duration.

If a reversed drop is detected, the event is classified as noise and excluded from the list of valid refills.

4.4.5 Role of Stability Validation in the Framework

By combining jump magnitude thresholds, persistence-based validation, and reverse window checks, the framework effectively eliminates both short-lived and delayed false spikes. This multi-stage validation ensures that only genuine, sustained refilling events are retained, producing a physically consistent and reliable fuel-level time series suitable for downstream analysis.

Table 4.1 summarizes the final set of empirically calibrated parameters used for spike detection, refill classification, and stability validation.

Based on the calibrated parameters, Figure 4.2 illustrates the detected z-score anomaly candidates and their classification into refilling events and noise. Genuine refills are characterized by sustained increases in fuel level, while noise spikes exhibit short-lived behavior and are rejected by the persistence and reverse-window criteria.

Spike classification results for the remaining tanks are provided in Appendix A.2.

Table 4.1.: Final calibrated parameters for fuel-level spike detection and refill classification

Parameter	**Value**	**Description**
z_threshold	3	Z-score threshold used to identify candidate fuel-level spikes based on consecutive level differences.
capacity_factor_refill (α)	0.1	Scaling factor applied to tank capacity to define the minimum jump magnitude required for a refill candidate.
window_after	30 hours	Length of the persistence window following a detected spike used to evaluate short-term fuel-level stability.
persist_points	8	Minimum number of observations within the persistence window that must remain above the persistence threshold for an event to qualify as a valid refill.
reverse_window_hours	30 hours	Duration of the reverse window used to validate long-term stability and detect delayed false refilling events.

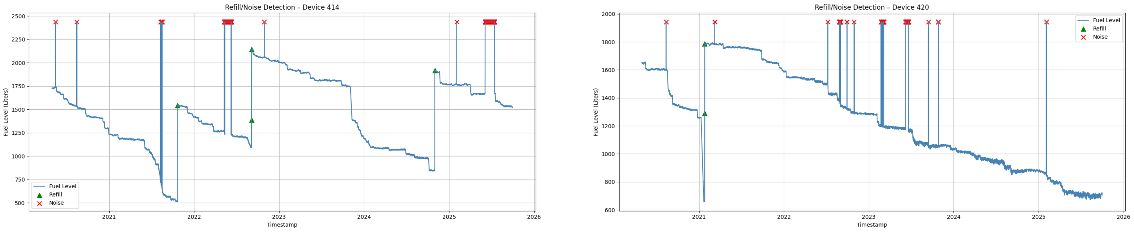

(a) Tank 414 (b) Tank 420

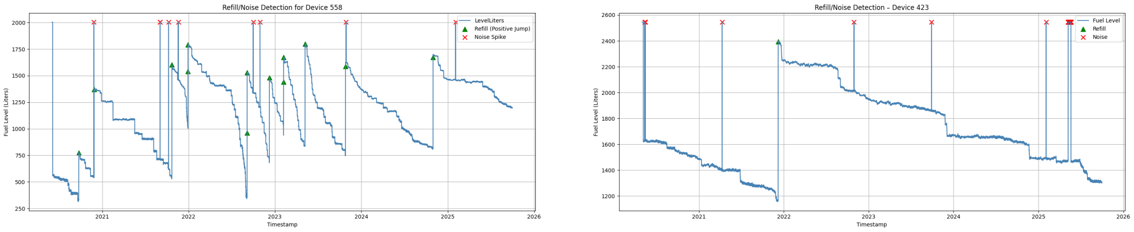

(c) Tank 558 (d) Tank 423

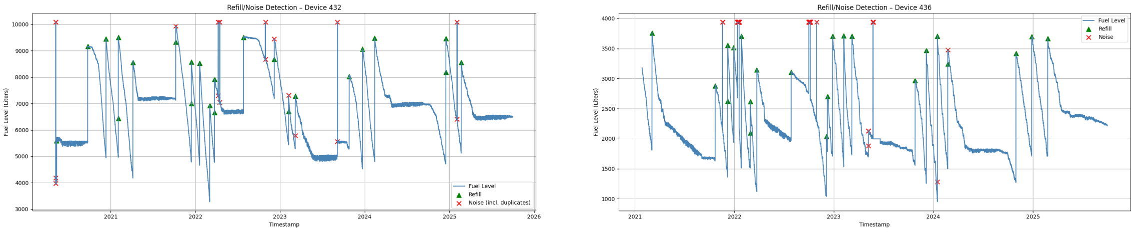

(e) Tank 432 (f) Tank 436

Figure 4.2.: Z-score anomaly candidates and refill versus noise classification for representative fuel tanks.

4.5 Post-Detection Cleaning and Preparation of Final Time Series

Following the detection and classification of fuel-level spikes, the final stage of the preprocessing framework focuses on removing noise-induced artefacts and preparing a clean, analysis-ready time series.While refilling events represent legitimate operational behavior and are retained, noise spikes must be corrected to prevent distortion of downstream analyses.

If noise-induced spikes remain in the dataset, they can significantly bias the interpretation of fuel consumption behavior. In particular, such artefacts may distort the underlying consumption trend, obscure seasonal patterns, increase forecasting error, and reduce the interpretability of statistical and machine learning models. Consequently, all data points classified as noise are explicitly addressed before applying time-series decomposition and forecasting methods such as STL and Exponential Smoothing (ETS) modeling.

4.5.1 Noise Identification and Correction

Based on the refill versus noise classification results obtained from the Z-score–based detection and validation framework, all timestamps corresponding to observations classified as noise were identified for each tank. These observations represent sensor-induced artefacts rather than genuine fuel-level changes and therefore require correction prior to further analysis.

For each identified noise timestamp, the corresponding LevelLiters value was replaced with a missing value (NaN), and time-based linear interpolation was applied between the nearest valid observations. This approach preserves the temporal continuity of the fuel-level series while removing unrealistic spikes, without altering genuine refilling events or normal consumption behavior.

4.5.2 Verification and Final Dataset Preparation

The interpolated fuel-level series was visually inspected to confirm the removal of abrupt, unrealistic jumps and to ensure that the overall consumption trend remained physically plausible. The resulting cleaned fuel-level time series provides a stable and reliable input for subsequent time-series decomposition, forecasting, and visualization.

Figure 4.3 shows the cleaned fuel-level time series after noise removal and interpolation. The resulting signal is continuous, free of spurious spikes, and suitable for subsequent time-series analysis.

Cleaned time-series results for the remaining tanks are provided in Appendix A.3.

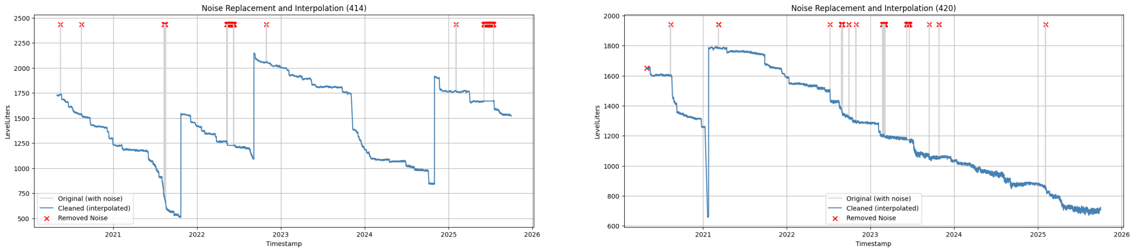

(a) Tank 414

(b) Tank 420

(c) Tank 558

(d) Tank 423

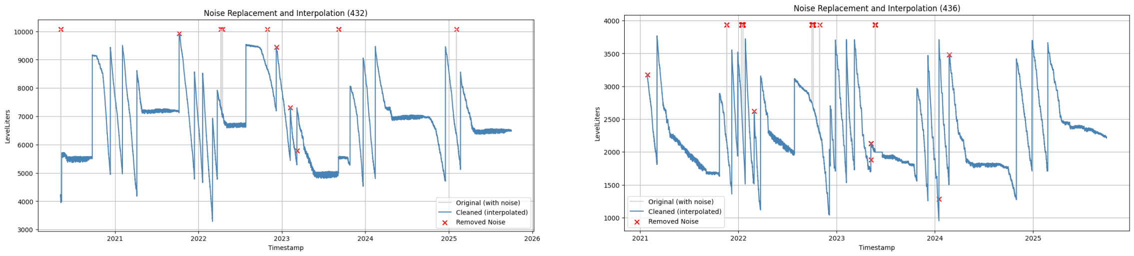

(e) Tank 432

(f) Tank 436

Figure 4.3.: Cleaned fuel-level time series after post-detection noise correction.

5. Time Series Decomposition and Exploratory Analysis

5.1 Overview of Time Series Decomposition Approach

After completing data cleaning, noise correction, and preparation of reliable fuel-level and consumption time series, the next step focuses on exploring the temporal structure of fuel consumption behavior. Time series decomposition is employed to separate the observed consumption signal into interpretable components, enabling a clearer understanding of long-term trends, recurring seasonal patterns, and irregular fluctuations.

Fuel consumption data typically exhibit multiple temporal characteristics driven by operational schedules, maintenance activities, and usage intensity. Short-term variations may reflect daily or weekly operational cycles, while longer-term trends can indicate changes in demand, efficiency, or mission activity levels. Decomposing the time series allows these effects to be analyzed independently, reducing confounding between systematic behavior and random noise.

In this study, STL is adopted as the primary decomposition framework. STL is well suited for fuel consumption analysis because it is flexible, robust to outliers, and capable of handling non-stationary trends and complex seasonal patterns (Cleveland, 1990). Prior aggregation to daily consumption further reduces high-frequency noise and emphasizes meaningful operational behavior.

The decomposition results form the foundation for subsequent exploratory analyses in this chapter, including the examination of trend evolution, seasonal dynamics, residual behavior, and anomaly detection. By isolating these components, the analysis provides insight into both typical fuel usage patterns and deviations that may indicate abnormal events or operational inefficiencies.

5.2 Temperature Correction and Consumption Derivation

During preliminary exploration of the IoT-based fuel-level data, irregular fluctuations were observed that did not correspond to actual fuel usage patterns. Following consultation with the data provider, these variations were attributed primarily to the thermal expansion and contraction of diesel fuel caused by ambient temperature changes. Because the monitored tanks are relatively large, even modest temperature variations can result in noticeable changes in measured fuel level. For example, a vertical change of approximately 1 mm in fuel height may correspond to nearly 2 L of volume difference. In regions such as Mitrovica, daily temperature variations can reach up to 10 °C, producing measurable expansion effects in the sensor readings.

5.2.1 Theoretical Background

The thermal expansion of liquid fuel can be described using the volumetric expansion relationship:

$$\Delta V = V_0 \times \beta \times \Delta T, \tag{5.1}$$

where V_0 denotes the fuel volume at a reference temperature, β is the volumetric expansion coefficient of diesel (approximately 0.00089 $°\mathrm{C}^{-1}$), and ΔT represents the temperature change in degrees Celsius. Accordingly, the measured volume at temperature T can be expressed as:

$$V_{\text{measured}} = V_0(1 + \beta \Delta T). \tag{5.2}$$

To normalize fuel volume measurements to a constant reference temperature ($T_{\text{ref}} = 15$ °C in this study), the corrected volume is obtained as:

$$V_0 = \frac{V_{\text{measured}}}{1 + \beta(T - T_{\text{ref}})}. \tag{5.3}$$

This correction removes artificial fluctuations induced by environmental heating or cooling and yields a more physically consistent representation of fuel volume.

5.2.2 Temperature Data Acquisition

Hourly ambient temperature data were retrieved from the Open–Meteo Application Programming Interface (API), which provides open-access historical weather observations for any latitude–longitude location. The API supports hourly temporal resolution, making it suitable for alignment with the fuel-level sensor timestamps. Each sensor's geographic coordinates were used to extract temperature data corresponding to the same time zone and hourly resolution as the fuel-level measurements.

5.2.3 Temperature Correction of Fuel-Level Data

After acquiring temperature data, it was merged with the hourly fuel-level dataset using timestamps as a common key. For each observation, the temperature-corrected fuel level was computed as:

$$\text{LevelLiters}_{\text{corrected}} = \frac{\text{LevelLiters}}{1 + \beta(T - 15)}. \tag{5.4}$$

This correction effectively compensates for temperature-induced expansion and contraction of the fuel. A comparative visualization is shown in Figure 5.1, illustrating the relationship between ambient temperature, raw fuel-level measurements, and the temperature-corrected fuel-level time series. The figure highlights how short-term fluctuations in the raw sensor readings coincide with temperature variations, while the corrected signal removes temperature-induced distortions and preserves the underlying fuel consumption behavior.

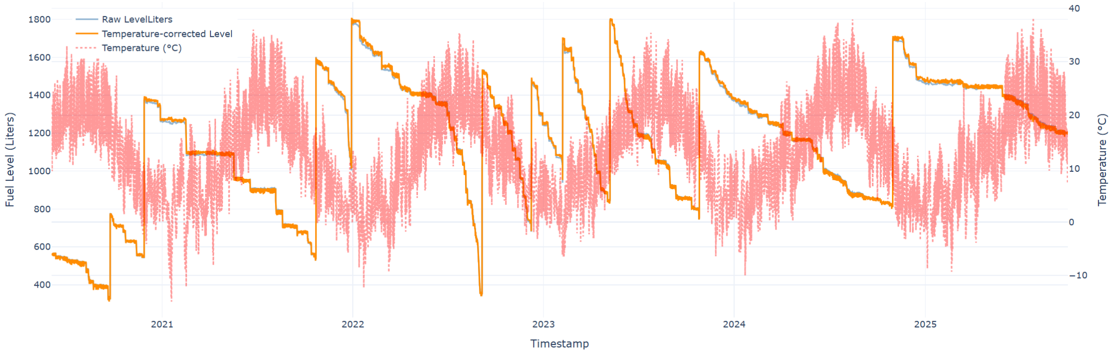

Figure 5.1.: Comparison of raw fuel-level measurements, temperature-corrected fuel levels, and ambient temperature for a representative tank.

5.2.4 Fuel Consumption Derivation

Hourly fuel consumption was derived from the temperature-corrected fuel-level series by computing the negative first difference:

$$\text{Consumption}_t = \max\left(0,\ \text{LevelLiters}_{t-1}^{\text{corrected}} - \text{LevelLiters}_t^{\text{corrected}}\right). \tag{5.5}$$

Negative differences, which correspond to refilling events, were set to zero to isolate true fuel usage. To reduce high-frequency sensor noise, a centered three-hour rolling mean was applied to the resulting consumption series. To ensure suitability for subsequent decomposition and analysis, additional preprocessing steps were applied to the derived consumption series, including outlier detection and correction. These steps are described in detail in Section 5.3.

5.3 Outlier Detection and Preprocessing for STL Stability

Before applying the STL decomposition, it is necessary to address extreme or erroneous values that could distort the extracted trend and seasonal components. Such outliers may arise from residual sensor noise, data transmission errors, or atypical operational events, and if left untreated, can adversely affect the stability and interpretability of the decomposition results.

Following the recommendations of Hyndman and Athanasopoulos (2021), an Interquartile Range (IQR)–based rule was employed to identify outliers in the daily fuel consumption series. Observations were classified as outliers if they lay outside the following bounds:

$$x_i < Q_1 - 1.5 \times \mathrm{IQR} \quad \text{or} \quad x_i > Q_3 + 1.5 \times \mathrm{IQR}, \tag{5.6}$$

where Q_1 and Q_3 denote the first and third quartiles of the series, respectively, and $\mathrm{IQR} = Q_3 - Q_1$.

This rule provides a robust, distribution-independent method for detecting extreme anomalies while avoiding the removal of legitimate variations in fuel consumption. Once identified, outlier values were replaced with missing values and subsequently reconstructed using time-based linear interpolation. This approach preserves the temporal continuity of the series while removing spurious spikes that could bias the decomposition.

The resulting preprocessed consumption series exhibits reduced irregular variability and improved stability, ensuring that the STL model captures meaningful trend and seasonal patterns rather than artefacts introduced by extreme observations.

5.4 STL Decomposition of Fuel Consumption

After completing temperature correction and outlier preprocessing, the daily fuel consumption time series was decomposed using STL. The objective of this decomposition is to separate the observed consumption signal into interpretable components that capture long-term behavior, recurring seasonal patterns, and irregular fluctuations.

Formally, STL decomposes a time series y_t into three additive components:

$$y_t = T_t + S_t + R_t, \tag{5.7}$$

where T_t represents the trend component, S_t denotes the seasonal component, and R_t corresponds to the residual or irregular component (Cleveland, 1990). In this study, the decomposition was applied to daily aggregated fuel consumption data, with a seasonal period of seven days to reflect weekly operational cycles.

The trend component captures gradual changes in average fuel usage over time and highlights longer-term variations associated with operational intensity, mission activity, or efficiency changes. The seasonal component reflects systematic and repeating patterns in consumption, particularly weekly cycles linked to regular work schedules and operational routines. The residual component contains short-term irregular fluctuations and unexplained variations, which may include random noise or anomalous consumption events.

Figure 5.2 illustrates the STL decomposition of daily fuel consumption for a representative tank (Tank 558). The top panel shows the observed daily consumption series, followed by the extracted trend, seasonal, and residual components. The trend reveals the overall evolution of fuel usage over time, while the seasonal component highlights consistent weekly consumption patterns. The residual component is largely centered around zero, indicating that the main systematic structures have been effectively captured by the decomposition.

To further emphasize the relationship between observed consumption and long-term behavior, Figure 5.3 overlays the STL trend component on the original daily consumption series. This visualization demonstrates how the trend smooths short-term variability while preserving the underlying evolution of fuel usage. Together, these results confirm that STL provides a clear and interpretable separation of fuel consumption dynamics, forming a robust foundation for subsequent exploratory analysis.

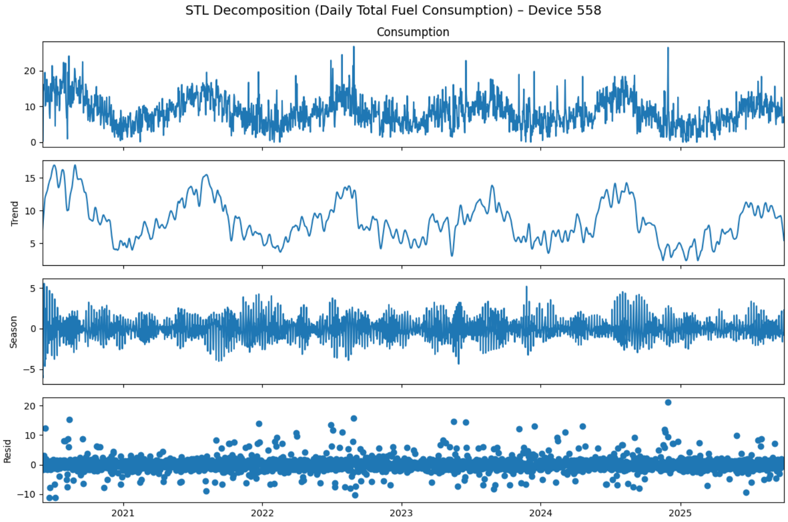

Figure 5.2.: STL decomposition of daily fuel consumption for a representative tank (Tank 558). The observed series is separated into trend, seasonal, and residual components.

5.5 Impact of Outlier Cleaning on Residual Behavior

To assess the effectiveness of the IQR-based outlier detection and interpolation procedure, the residuals obtained from the STL decomposition were examined before and after cleaning. The comparison focuses on evaluating whether the preprocessing steps successfully reduce irregular variability while preserving the underlying structure of the time series.

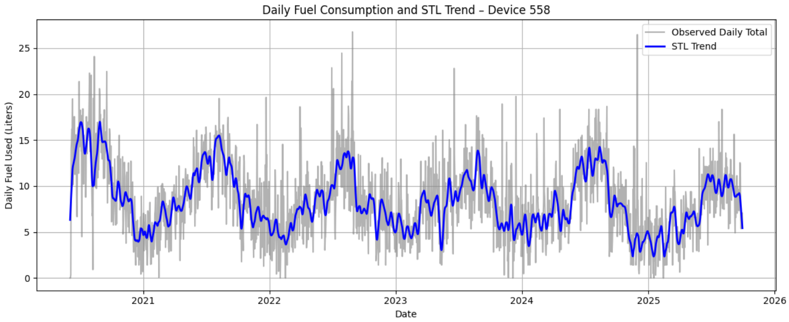

Figure 5.3.: Daily fuel consumption and extracted STL trend component for Tank 558.

Prior to cleaning, the residual series exhibited high-amplitude fluctuations and irregular variance, indicating the presence of extreme values and measurement noise. Such behavior suggests that a portion of the variability in the unprocessed data was not related to systematic consumption patterns, potentially distorting the interpretation of trend and seasonal components.

After applying the IQR-based outlier replacement and interpolation, the residuals became substantially smoother and more symmetrically distributed around zero. Figure 5.4 illustrates this improvement by comparing raw and cleaned residual time series for a representative tank. The cleaned residuals show reduced dispersion and fewer extreme deviations, indicating effective removal of spurious fluctuations.

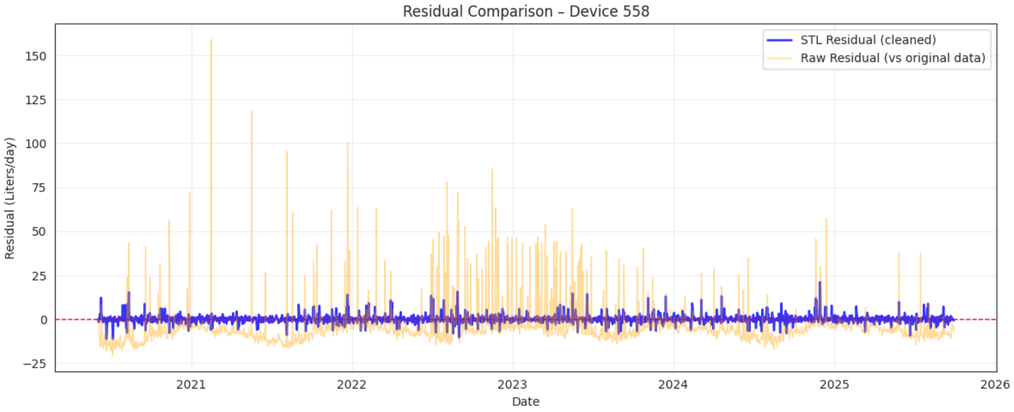

Figure 5.4.: Comparison of STL residuals before and after IQR-based outlier cleaning for Tank 558. The cleaned residuals exhibit reduced variability and are more closely centered around zero, indicating effective removal of extreme values and measurement noise.

Further evidence is provided by the residual distribution plots shown in Figures 5.5 and 5.6. The histogram comparison demonstrates a clear reduction in heavy tails after cleaning, while the boxplot highlights a narrower interquartile range and a significant decrease in the number and magnitude of outliers. Together, these results confirm that the cleaning procedure enhances the stability of the residual component.

Overall, the improved residual behavior indicates that the STL model more effectively captures the systematic structure of fuel consumption after preprocessing. This enhancement supports more reliable interpretation of trends and seasonality and provides a robust foundation for subsequent anomaly detection and forecasting analyses.

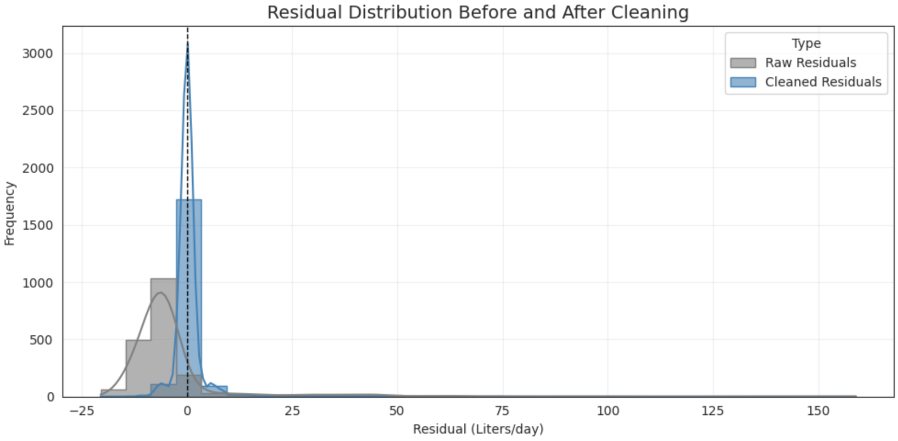

Figure 5.5.: Distribution of STL residuals before and after IQR-based outlier cleaning for Tank 558. The cleaned residuals show a narrower distribution and reduced heavy tails, suggesting improved model stability and reduced influence of extreme observations.

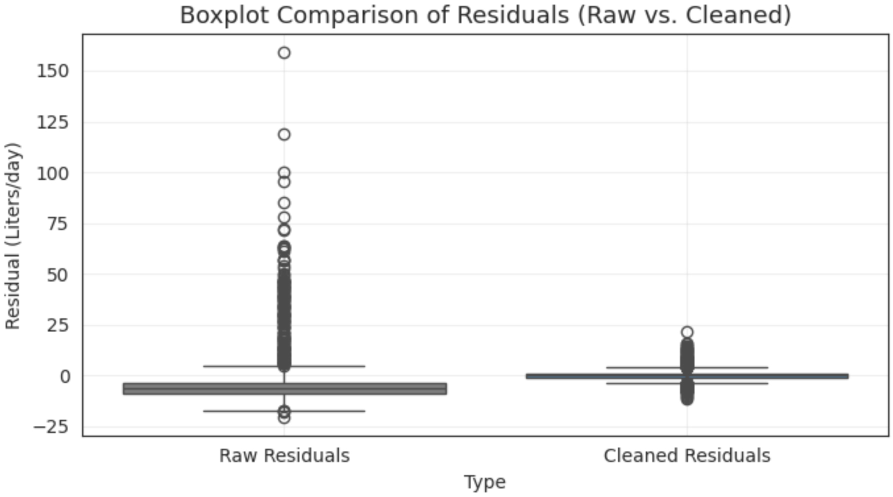

Figure 5.6.: Boxplot comparison of raw and cleaned STL residuals for Tank 558. The cleaned residuals display a smaller interquartile range and fewer extreme outliers, confirming the effectiveness of the outlier correction procedure.

5.6 Monthly and Seasonal Fuel Consumption Dynamics

To examine how fuel consumption varies throughout the year, the daily fuel consumption series was aggregated into monthly totals using the STL-derived trend component. By relying on the trend component rather than raw daily values, this approach emphasizes long-term consumption behavior while filtering out short-term fluctuations and irregular noise. Monthly totals were computed by summing the daily trend values within each calendar month, expressed in liters, yielding the total fuel consumed per month.

Figure 5.7 illustrates the resulting monthly fuel consumption trends across multiple years for a representative tank. Each line corresponds to one year, while the horizontal axis represents the months from January to December. This multi-line representation enables direct comparison of seasonal behavior and year-to-year differences in fuel usage. The results reveal recurring seasonal patterns, with fuel consumption increasing during specific months, potentially reflecting periods of heightened operational activity or environmental influences. Differences in the magnitude of monthly totals across years further indicate variations in operational intensity, efficiency changes, or mission-related factors. Months that consistently exhibit peaks across years suggest recurring high-demand periods, while troughs indicate intervals of reduced activity or improved efficiency.

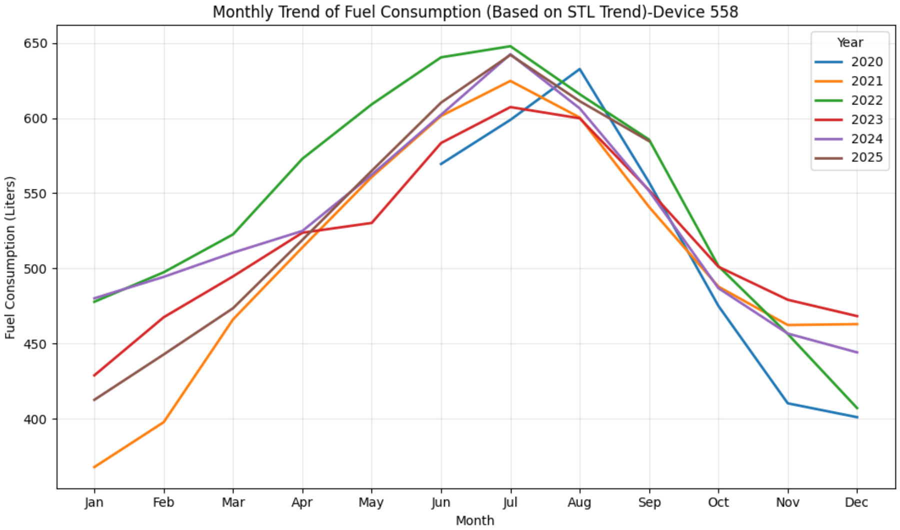

Figure 5.7.: Smoothed monthly fuel consumption trends for Tank 558 based on the STL trend component (2020–2025). Each line represents the total monthly fuel consumption for a given year, highlighting seasonal variations and year-to-year differences in overall fuel usage.

To further investigate recurring consumption patterns, the seasonal component extracted from the STL decomposition was analyzed independently. The seasonal component captures systematic, repeating deviations from the long-term trend and reflects short-term periodic behavior associated with operational routines or environmental conditions. Unlike the trend component, which evolves gradually, the seasonal component highlights consistent cyclical fluctuations.

Figure 5.8 presents a facet plot of the STL seasonal component, with one panel for each month from January to December. Within each panel, the black line represents the average seasonal deviation for that month across the analyzed years, while the blue horizontal line denotes the mean seasonal baseline. Positive deviations indicate periods when fuel consumption tends to exceed the long-term average, whereas negative deviations correspond to below-average usage. This representation highlights month-to-month differences in seasonal behavior and reveals which months consistently contribute to higher or lower fuel consumption relative to the baseline.

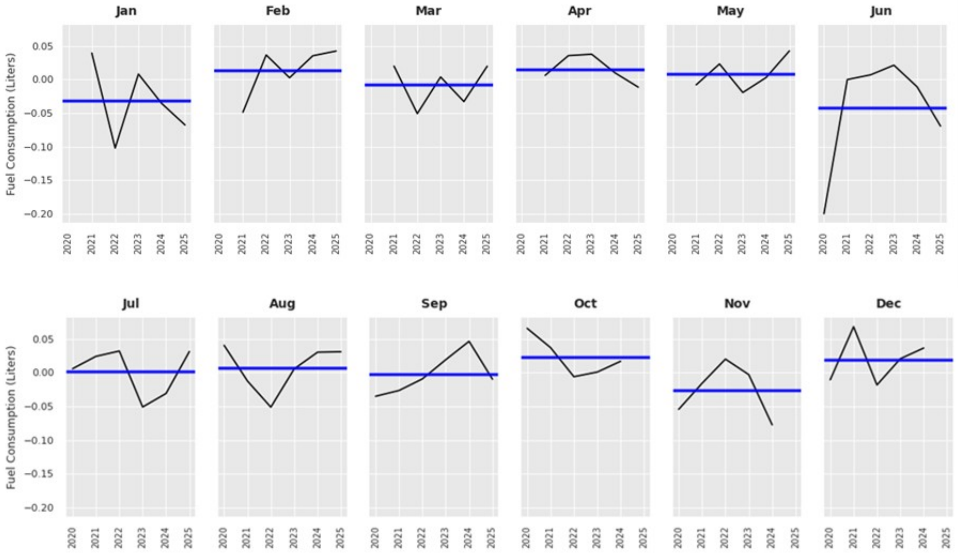

Figure 5.8.: Seasonality analysis of fuel consumption for Tank 558 based on annual averages of the STL seasonal component. Each panel corresponds to one month (January–December), showing year-to-year variations in seasonal deviations. Positive values indicate above-average consumption, while negative values represent below-average usage relative to the long-term trend.

By aggregating seasonal deviations across years, this analysis provides a clear, high-level view of how recurring consumption patterns evolve over time. Together, the monthly trend and seasonal component analyses offer complementary insights into both long-term consumption dynamics and systematic seasonal behavior.

5.7 Residual-Based Anomaly Detection in Fuel Consumption

To identify unusual fuel consumption behavior, anomalies were detected using the residuals obtained from the STL decomposition. Residuals represent the portion of the time series that is not explained by the long-term trend or the recurring seasonal patterns and therefore capture irregular fluctuations that may correspond to abnormal events.

Anomaly detection was performed on both the raw residuals (prior to outlier cleaning) and the cleaned residuals obtained after IQR-based smoothing and interpolation. In both cases, anomalies were defined as observations whose residual values exceeded $\pm 3\sigma$, where σ denotes the standard deviation of the residual series. This threshold-based approach assumes that, under normal operating conditions, residuals are approximately centered around zero and follow a stable distribution, making extreme deviations statistically unlikely.

Figure 5.9 illustrates the detected anomalies for a representative tank. The upper panel shows anomaly detection based on raw residuals, where a large number of extreme spikes exceed the $\pm 3\sigma$ confidence limits. Many of these spikes are attributable to measurement noise or uncorrected outliers, which can obscure truly meaningful abnormal events. The lower panel presents the results obtained using cleaned residuals, where the number of detected anomalies is reduced and the remaining events are more clearly distinguishable from normal variability.

Anomalies identified using cleaned residuals correspond to statistically significant deviations from expected consumption behavior. These events may indicate unusually high fuel usage, potentially associated with leaks or unauthorized consumption, or unexpectedly low usage, which may reflect sensor malfunctions, downtime, or data quality issues. By isolating such events, the residual-based anomaly detection framework provides a data-driven mechanism for monitoring fuel consumption and supporting targeted investigation of irregular operational behavior.

5.8 Daily and Weekly Fuel Consumption Patterns

To better understand short-term and operational fuel usage behavior, daily and weekly consumption patterns were analyzed using the hourly fuel consumption data. This analysis complements the STL-based decomposition by focusing on intra-day and intra-week variations that are directly linked to operational schedules and routine activities.

Figure 5.10 summarizes the daily and weekly fuel consumption patterns for a

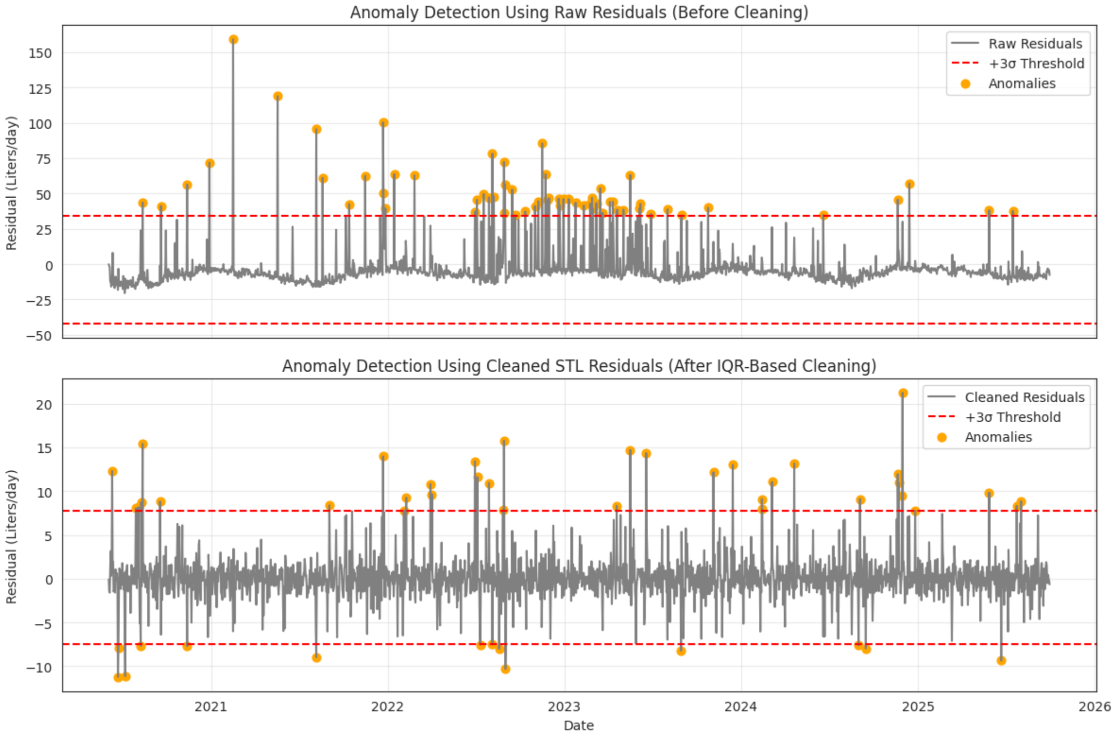

Figure 5.9.: Residual-based anomaly detection in daily fuel consumption for Tank 558. The upper panel shows anomalies detected using raw STL residuals prior to outlier cleaning, while the lower panel presents anomalies identified using cleaned residuals after IQR-based preprocessing. Points outside the $\pm 3\sigma$ confidence limits indicate statistically significant deviations from expected consumption behavior.

representative tank (Tank 436). The left panel presents the average hourly fuel consumption over a 24-hour cycle, illustrating how fuel usage varies throughout the day. Consumption generally increases during active operational hours and decreases during off-peak or idle periods, reflecting typical daily work routines. The right panel shows total daily fuel consumption aggregated by weekday, highlighting systematic differences in fuel usage across the week. Higher consumption levels are typically observed on workdays, while reduced usage during weekends or maintenance periods indicates lower operational activity.

Together, these visualizations provide insight into short-term operational behavior and reveal recurring daily and weekly rhythms that influence overall fuel demand. Such patterns are important for distinguishing expected operational variability from potentially abnormal consumption events.

To further examine the joint influence of hour-of-day and day-of-week effects, a heatmap of average hourly fuel consumption by weekday was generated, as shown in Figure 5.11. Each cell represents the mean fuel consumption for a specific hour and day combination, with warmer colors indicating higher usage levels and cooler tones corresponding to lower activity. The heatmap reveals consistent high-consumption periods during working hours on specific weekdays, as well as reduced activity during nights and weekends.

This combined daily and weekly visualization highlights recurring operational patterns and identifies periods of peak fuel demand. Such insights are valuable for operational planning, resource allocation, and supporting the interpretation of anomalies and forecasting results presented in subsequent sections.

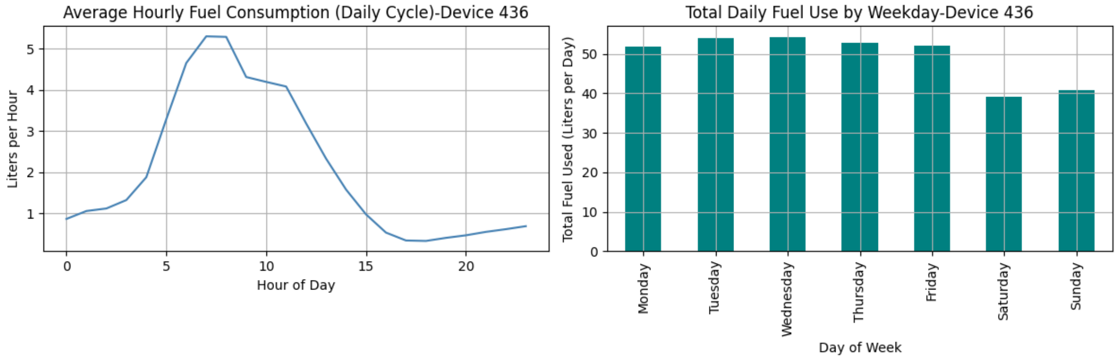

Figure 5.10.: Average hourly fuel consumption (left) and total daily fuel consumption by weekday (right) for Tank 436. The plots illustrate typical intra-day usage cycles and systematic weekday variations, reflecting operational schedules and activity levels.

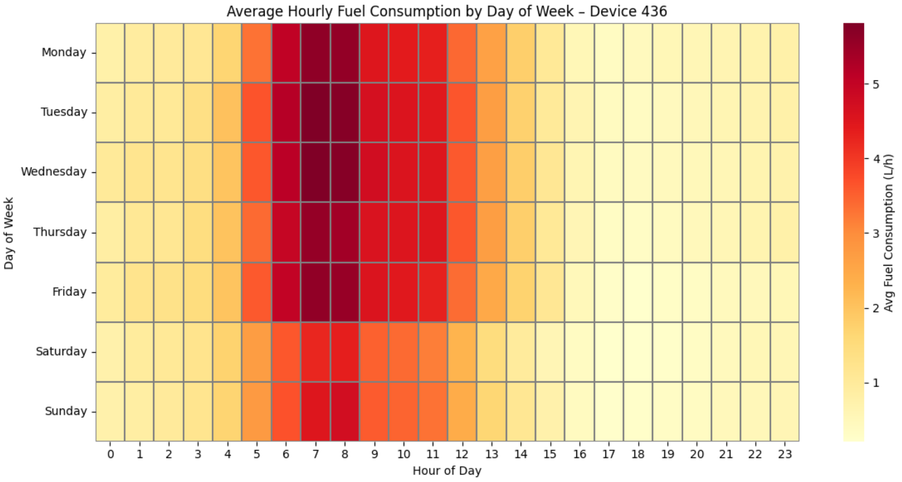

Figure 5.11.: Heatmap of average hourly fuel consumption by day of the week for Tank 436. Brighter colors indicate higher consumption levels, highlighting recurring daily and weekly operational patterns in fuel usage.

6. Spatial Autocorrelation Analysis of Fuel Consumption

6.1 Overview and Motivation for Spatial Analysis

While fuel consumption patterns are primarily driven by site-specific operational behavior and temporal dynamics, geographic location may also influence consumption through shared environmental conditions, logistical constraints, or coordinated operational activities. Spatial autocorrelation analysis is therefore introduced to assess whether spatial proximity between fuel tanks provides additional information beyond what is captured by independent time-series analysis.

In the context of fuel consumption monitoring and forecasting, the presence of spatial dependence would imply that nearby tanks exhibit similar consumption behavior. Such similarity could arise from shared infrastructure, common usage schedules, or regional environmental effects, suggesting that spatial information might be leveraged to enhance predictive models. Conversely, weak or absent spatial autocorrelation would indicate that fuel consumption behavior is largely independent of geographic location, supporting the use of per-tank modeling approaches.

Accordingly, spatial autocorrelation analysis is employed in this study as a diagnostic tool rather than a primary modeling framework. Its purpose is to evaluate whether spatial relationships contribute meaningful structure to fuel consumption patterns across sites and to determine whether more complex spatial or spatial–temporal forecasting models are justified. The analysis focuses on identifying global spatial dependence in fuel consumption values and examining whether temperature correction alters the observed spatial structure.

By establishing whether spatial proximity influences fuel consumption behavior, this chapter provides an empirical basis for selecting appropriate forecasting strategies in subsequent analyses and supports methodological decisions regarding model complexity.

6.2 Global Spatial Autocorrelation Methodology

To quantify the degree of spatial dependence in fuel consumption across tanks, Global Moran's I was employed as the primary measure of spatial autocorrelation. Moran's I is a widely used statistic that evaluates whether similar values of a variable tend to occur near one another in geographic space, indicating spatial clustering, or whether dissimilar values are spatially proximate, indicating spatial dispersion. Values close to zero suggest a random spatial pattern.

Positive values of Moran's I indicate positive spatial autocorrelation, where nearby tanks exhibit similar fuel consumption levels, while negative values indicate spatial dispersion, where neighboring tanks tend to display dissimilar consumption behavior. Moran's I values that are not statistically significant imply the absence of meaningful spatial structure.

The Global Moran's I statistic is defined as:

$$I = \frac{N}{W} \frac{\sum_{i=1}^{N} \sum_{j=1}^{N} w_{ij}(x_i - \bar{x})(x_j - \bar{x})}{\sum_{i=1}^{N} (x_i - \bar{x})^2} \tag{6.1}$$

where N is the number of spatial units (tanks), x_i and x_j represent the fuel consumption values at locations i and j, $\bar{x}$ is the mean fuel consumption, w_{ij} denotes the spatial weight defining the neighborhood relationship between locations i and j, and W is the sum of all spatial weights, given by $W = \sum_{i=1}^{N} \sum_{j=1}^{N} w_{ij}$.

Statistical significance of the observed Moran's I values was assessed using a permutation-based test. In this procedure, fuel consumption values are randomly reassigned across spatial locations multiple times to generate a reference distribution under the null hypothesis of spatial randomness. The empirical Moran's I statistic is then compared against this distribution to obtain a p-value, indicating whether the observed spatial pattern is unlikely to have arisen by chance.

6.3 Spatial Weights Specification and Data Aggregation for Spatial Analysis

The computation of Global Moran's I requires the specification of a spatial weights matrix that defines the neighborhood structure among fuel tanks. The spatial weights matrix encodes which locations are considered neighbors and determines how spatial relationships influence the autocorrelation measure.

6.3.1 Spatial Weights Specification

In this study, spatial relationships between tanks were defined using a k-nearest neighbor (KNN) spatial weights matrix with $k = 2$. Under the KNN approach, each tank is assigned exactly two nearest neighboring tanks based on geographic distance. This choice was motivated by the relatively small number of tanks and their clustered spatial configuration, which resulted in disconnected components when distance-based threshold methods were tested.

The KNN specification ensures that each spatial unit has the same number of neighbors, producing a fully connected spatial weights matrix and avoiding the need to impose arbitrary distance cutoffs. As a result, local spatial relationships are preserved while maintaining consistency across all sites. The spatial weights were row-standardized so that the influence of neighboring tanks sums to one for each location.

All spatial distance calculations were performed using projected coordinates in Universal Transverse Mercator (UTM) Zone 34N (EPSG:32634) to ensure accurate and comparable distance measurements.

6.3.2 Data Aggregation for Spatial Analysis

To support spatial autocorrelation analysis, fuel consumption data were aggregated to the site–year level, yielding one observation per tank per year. The mean annual fuel consumption was used as the aggregation metric, representing typical consumption behavior while reducing the influence of short-term fluctuations.

This aggregation strategy ensures comparability across tanks with different operating durations, data availability, and consumption scales. By summarizing consumption at the annual level, the analysis avoids conflating spatial dependence with temporal variability, tank capacity differences, or irregular monitoring periods. The resulting aggregated dataset provides a stable and interpretable basis for evaluating global spatial autocorrelation across sites.

6.4 Results of Spatial Autocorrelation Analysis

Global Moran's I was computed separately for each year to evaluate spatial dependence in fuel consumption across tanks. The analysis was conducted for both observed fuel consumption and temperature-corrected fuel consumption in order to assess whether temperature normalization influences the spatial structure of the data.

6.4.1 Spatial Autocorrelation of Observed Fuel Consumption

Across all analyzed years, Moran's I values for observed fuel consumption were consistently negative, with magnitudes ranging approximately from -0.13 to -0.36. None of the computed Moran's I statistics were statistically significant at the 5% significance level based on permutation testing.

These results indicate a weak tendency toward spatial dispersion rather than spatial clustering, suggesting that nearby tanks do not exhibit similar fuel consumption behavior. Instead, tanks in close geographic proximity tend to display dissimilar consumption patterns, and the observed spatial arrangement does not deviate significantly from what would be expected under spatial randomness.

6.4.2 Spatial Autocorrelation of Temperature-Corrected Fuel Consumption

After correcting fuel consumption values for temperature effects, Moran's I values remained negative and of similar magnitude across all years. As with observed fuel consumption, none of the temperature-corrected Moran's I statistics were statistically significant.

The close correspondence between observed and temperature-corrected Moran's I values indicates that temperature effects do not introduce, mask, or alter the spatial structure of fuel consumption across sites. This finding suggests that the absence of spatial clustering is not attributable to thermal expansion effects but reflects the underlying operational characteristics of individual tanks.

Table 6.1 summarizes the Global Moran's I statistics and associated *p*-values for both observed and temperature-corrected fuel consumption across all analyzed years.

Table 6.1.: Global Moran's I results for observed and temperature-corrected fuel consumption

	Observed Consumption		**Temperature-Corrected Consumption**	
Year	Moran's I	*p*-value	Moran's I	*p*-value
2020	−0.142	0.189	−0.125	0.160
2021	−0.359	0.242	−0.250	0.428
2022	−0.340	0.364	−0.240	0.408
2023	−0.242	0.389	−0.242	0.421
2024	−0.248	0.415	−0.213	0.472
2025	−0.132	0.324	−0.184	0.426

6.4.3 Summary of Spatial Autocorrelation Results

Overall, the results provide no evidence of positive global spatial autocorrelation in fuel consumption across tanks for any analyzed year. Both observed and temperature-corrected consumption exhibit weak spatial dispersion and statistically non-significant Moran's I values, indicating that geographic proximity alone does not explain similarities in fuel usage patterns.

6.5 Interpretation and Implications for Forecasting

The spatial autocorrelation analysis conducted in this chapter provides no evidence of positive spatial dependence in fuel consumption across tanks. For both observed and temperature-corrected fuel consumption, Global Moran's I values were consistently negative and statistically non-significant across all analyzed years. These results indicate that fuel consumption behavior does not exhibit spatial clustering and that geographic proximity alone does not lead to similar consumption patterns.

The weak tendency toward spatial dispersion suggests that fuel usage is primarily governed by site-specific operational factors rather than shared locational or environmental influences. Factors such as equipment usage intensity, operational schedules, maintenance practices, and local demand appear to dominate fuel consumption behavior, outweighing any potential spatial effects between nearby tanks.

Importantly, the similarity between observed and temperature-corrected Moran's I results indicates that temperature effects do not introduce or obscure spatial structure in the data. This finding confirms that the absence of spatial autocorrelation reflects intrinsic operational characteristics rather than artefacts arising from thermal expansion or measurement bias.

From a methodological perspective, these findings have direct implications for fuel consumption modeling and forecasting. The lack of significant spatial dependence suggests that incorporating spatial or spatial–temporal dependencies is unlikely to yield substantial predictive improvements. Consequently, independent per-tank time-series modeling approaches are justified, supporting the forecasting framework adopted in subsequent chapters.

Overall, the spatial autocorrelation analysis serves as a diagnostic step that informs model selection and complexity. By demonstrating that fuel consumption patterns are largely independent across locations, this chapter provides empirical support for focusing on temporal dynamics and site-specific behavior in downstream analysis and decision-support applications.

7. Time-Series Forecasting of Fuel Consumption

7.1 Overview of Forecasting Framework and Feature Engineering

After establishing reliable, cleaned, and temperature-corrected fuel consumption time series and analyzing their temporal structure, the next stage of the study focuses on short-term forecasting of fuel consumption. The objective of this chapter is to evaluate and compare different forecasting approaches in their ability to predict hourly fuel usage over the final 30-day period, supporting operational planning and decision-making.

All forecasting models were trained using temperature-corrected fuel consumption data measured at an hourly resolution as the target variable. The dataset was divided into training and testing periods, with the final 30 days reserved for out-of-sample evaluation at an hourly resolution. This evaluation strategy ensures an unbiased assessment of predictive performance under realistic operational conditions.

To enable both statistical and machine learning models to capture relevant temporal dependencies and contextual influences, a structured feature engineering framework was applied. Temporal features were derived from the historical consumption series, including lagged consumption values (up to 24 hours and seven days) to represent short-term and daily autocorrelation. Rolling statistics such as moving averages and standard deviations were computed over multiple windows to capture local trends and variability at the hourly scale.

In addition to consumption history, environmental and calendar-based features were incorporated to reflect known drivers of fuel usage. Temperature-related features included hourly ambient temperature, one-hour temperature lag, and rolling temperature statistics to account for residual thermal effects. Calendar features such as hour of day, day of week, month, and weekend indicators were included to capture recurring intra-day and weekly operational cycles identified in the exploratory analysis.

This unified forecasting framework allows direct comparison between classical univariate statistical models and multivariate machine learning and deep learning approaches. By applying a consistent feature set and evaluation protocol, the analysis isolates the contribution of model complexity and feature integration to forecasting accuracy.

7.2 Baseline Statistical Forecasting Model

As a benchmark for evaluating forecasting performance, a classical Exponential Smoothing model was employed as the baseline statistical approach. Exponential Smoothing methods model time series data by decomposing it into error, level, trend, and seasonal components, updating these components recursively over time based on recent observations using exponential smoothing (Hyndman and Athanasopoulos, 2021). Due to their simplicity and interpretability, such models are commonly used as reference points in forecasting studies.

In this work, the baseline model was applied to the temperature-corrected fuel consumption series measured at an hourly resolution. The model relies solely on the internal temporal structure of the series and does not incorporate external regressors or engineered features. This univariate formulation allows the baseline performance to reflect how well the intrinsic trend and seasonal patterns of fuel consumption can be captured without additional contextual information.

Seasonality was specified to reflect recurring daily operational cycles, with a seasonal period corresponding to 24 hours. This configuration enables the model to account for systematic intra-day variations in fuel usage identified during exploratory analysis. Model parameters were estimated automatically by minimizing the in-sample error, ensuring an objective and data-driven fitting process.

The Exponential Smoothing model serves as a reference against which the performance of more complex machine learning and deep learning models can be assessed. By establishing a baseline based on a purely statistical and univariate approach, the subsequent comparison quantifies the predictive gains achieved through multivariate feature integration and nonlinear modeling.

Figure 7.1 illustrates the ETS forecast for hourly fuel consumption over the final 30-day test period. The figure compares the observed consumption values with the corresponding ETS forecasts, highlighting the model's ability to capture recurring daily patterns and short-term consumption dynamics. Deviations between observed and predicted values indicate the limitations of a purely univariate approach when external drivers and nonlinear effects are not explicitly modelled.

7.3 Machine Learning–Based Forecasting Models

To capture nonlinear relationships and incorporate multiple contextual drivers of fuel consumption, several machine learning models were employed for hourly forecasting. Unlike the baseline Exponential Smoothing model, which relies solely on historical consumption values, machine learning approaches allow fuel usage to be modeled as a function of both temporal dependencies and external factors.

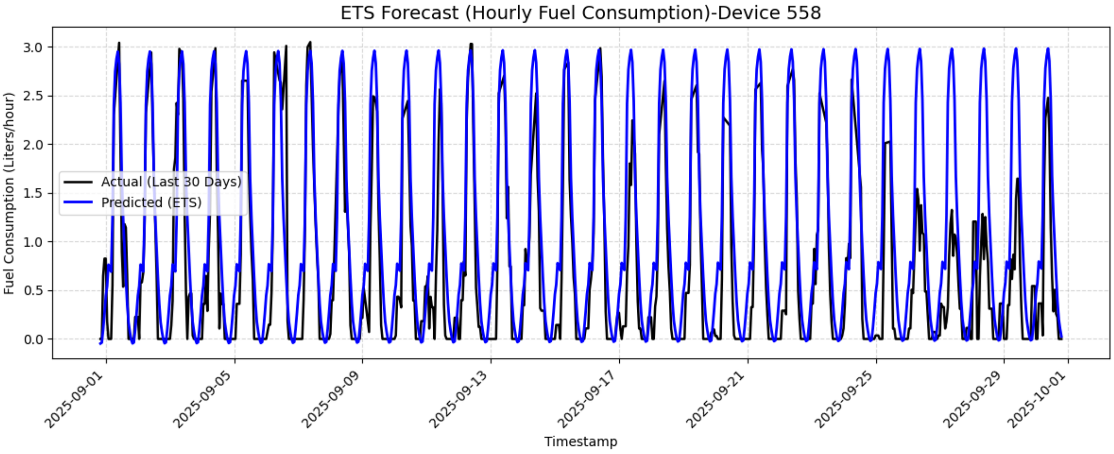

Figure 7.1.: Hourly fuel consumption forecast generated using the ETS model over the final 30-day test period. The plot compares observed values with model predictions, illustrating the baseline forecasting performance based solely on historical consumption patterns.

All machine learning models were trained using temperature-corrected fuel consumption data measured at an hourly resolution as the target variable. A supervised learning formulation was adopted, where historical observations and engineered features were used to predict future consumption values. The same training and testing split described in Section 7.1 was applied to ensure a fair comparison across models.

Feature Set for Machine Learning Models

The following features were engineered and used consistently across all machine learning forecasting models:

- **Lagged consumption features:** Hourly fuel consumption values from recent time steps, including short-term lags and daily lags, to capture temporal autocorrelation and short-term memory effects.
- **Rolling statistics:** Moving averages and standard deviations computed over multiple rolling windows to represent recent consumption trends and local variability.
- **Temperature-related features:** Hourly ambient temperature, lagged temperature values, and rolling temperature statistics to account for residual thermal effects influencing fuel behavior.
- **Calendar features:** Hour of day, day of week, month, and weekend indicators to capture recurring intra-day and weekly operational cycles.

This common feature engineering framework was applied consistently across all machine learning models to isolate the effect of model architecture on forecasting performance. The following subsections describe the individual machine learning algorithms used in this study and discuss their suitability for hourly fuel consumption forecasting.

7.3.1 Light Gradient Boosting Machine (LightGBM)

LightGBM was implemented as one of the primary machine learning models for hourly fuel consumption forecasting. LightGBM is a gradient boosting decision tree (GBDT) framework designed for high computational efficiency and strong predictive performance on large-scale datasets; it builds an ensemble of decision trees sequentially, where each new tree corrects errors of its predecessors (Ke et al., 2017).

In this study, LightGBM was trained to predict temperature-corrected fuel consumption at an hourly resolution using the feature set described in Section 7.3. The model leverages both historical consumption patterns and contextual variables, enabling it to capture complex and nonlinear relationships between fuel usage, environmental conditions, and operational timing.

Several characteristics make LightGBM particularly suitable for this application. First, its tree-based structure allows automatic feature selection by prioritizing the most informative predictors, such as recent consumption lags or rolling statistics, while reducing the influence of redundant features. Second, LightGBM efficiently handles nonlinear interactions between variables, which are difficult to represent using classical statistical models. Finally, its gradient boosting framework provides strong generalization performance, making it well suited for short-term forecasting tasks (Ke et al., 2017).

Model training was performed using a standard regression objective, and hyperparameters were selected to balance predictive accuracy and model stability. The resulting forecasts produced by LightGBM are evaluated and compared against other forecasting approaches in Section 7.5 to assess the benefits of multivariate and nonlinear modeling in fuel consumption prediction.

Figure 7.2 presents the hourly fuel consumption forecast generated using the LightGBM model for the final 30-day test period. The figure compares the observed consumption values with the corresponding model predictions, illustrating LightGBM's ability to capture both short-term variability and recurring operational patterns. Compared to the baseline statistical model, the LightGBM forecast more closely follows fluctuations in fuel usage, reflecting the benefits of multivariate feature integration and nonlinear modeling.

7.3.2 Extreme Gradient Boosting (XGBoost)

XGBoost model was implemented as an additional tree-based ensemble learning approach for hourly fuel consumption forecasting. XGBoost is a highly optimized

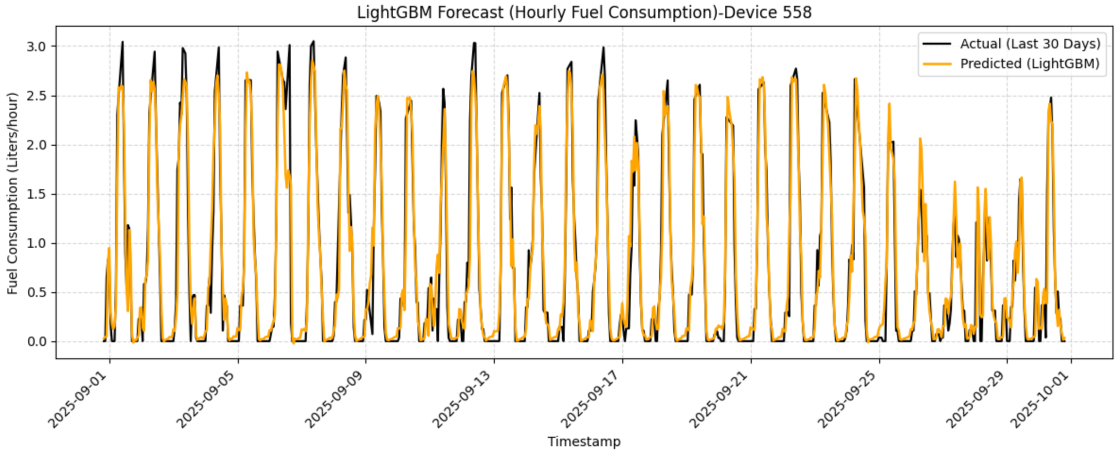

Figure 7.2.: Hourly fuel consumption forecast generated using the LightGBM model over the final 30-day test period. The plot compares observed values with model predictions, illustrating the model's ability to capture nonlinear consumption patterns and short-term temporal dynamics.

gradient boosting framework that extends classical boosting algorithms through regularization, efficient tree construction, and parallel processing, resulting in strong predictive performance and robust generalization capabilities (Chen and Guestrin, 2016).

In this study, XGBoost was trained using the same temperature-corrected hourly fuel consumption data and feature set described in Section 7.3. By adopting a consistent feature engineering framework, differences in forecasting performance can be attributed primarily to model architecture rather than input representation. The model learns complex nonlinear relationships between historical consumption, temperature effects, and operational timing variables.

A key advantage of XGBoost lies in its regularization mechanisms, including both L1 and L2 penalties, which help control model complexity and reduce overfitting, making the model well suited for high-frequency time series data characterized by noise and short-term variability. In addition, its tree-based boosting structure enables efficient modeling of nonlinear feature interactions and combined effects between predictors (Chen and Guestrin, 2016).

Figure 7.3 illustrates the hourly fuel consumption forecast produced by the XGBoost model over the final 30-day test period. The comparison between observed and predicted values demonstrates the model's ability to follow short-term fluctuations and recurring operational cycles. The forecasting performance of XGBoost is quantitatively evaluated and compared with other models in Section 7.5.

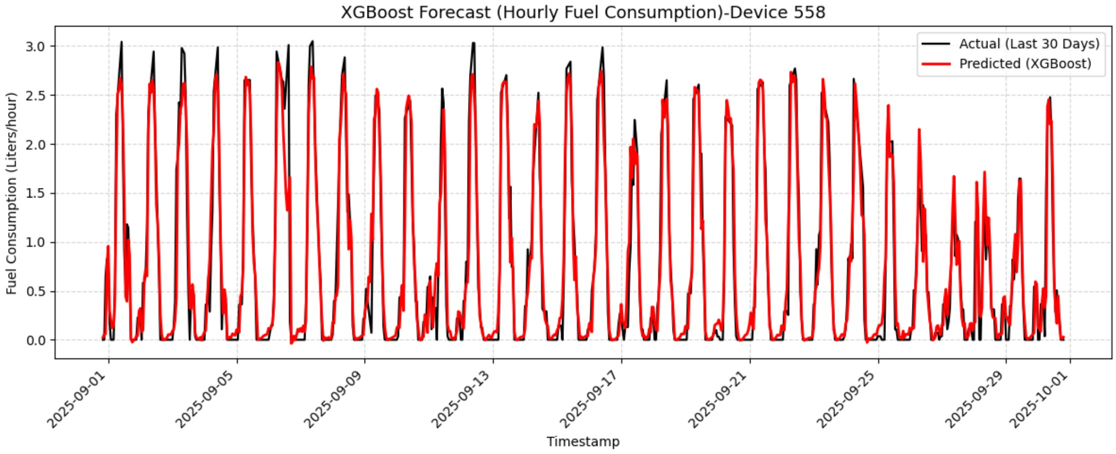

Figure 7.3.: Hourly fuel consumption forecast generated using the XGBoost model over the final 30-day test period. The plot compares observed values with model predictions, illustrating the model's capability to capture nonlinear temporal patterns in fuel consumption.

7.3.3 Random Forest Regressor

The Random Forest Regressor was implemented as an ensemble learning benchmark for hourly fuel consumption forecasting. Random Forest is an ensemble learning approach based on bootstrap aggregation, in which multiple decision trees are trained on different resampled versions of the training data and on randomly selected subsets of input features. The final prediction is obtained by averaging the outputs of all trees, which improves stability and reduces variance (Breiman, 2001).

In this study, the Random Forest model was trained using the same temperature-corrected hourly fuel consumption data and feature set. By maintaining a consistent input representation, differences in forecasting performance reflect the characteristics of the learning algorithms rather than variations in feature engineering. The model leverages historical consumption patterns, temperature-related variables, and calendar features to capture nonlinear relationships and recurring operational cycles.

A key strength of Random Forest lies in its robustness to noise and its ability to model nonlinear dependencies without strong parametric assumptions. Because individual trees are trained independently and combined through averaging, the model is less sensitive to short-term fluctuations and outliers, making it suitable for high-frequency time series data. Additionally, the random feature selection mechanism provides an implicit form of feature importance assessment, offering insight into the drivers of fuel consumption (Breiman, 2001).

Figure 7.4 presents the hourly fuel consumption forecast produced by the Random Forest model over the final 30-day test period. The comparison between observed and predicted values demonstrates the model's ability to capture general consumption trends and recurring operational patterns, while smoothing short-term irregular variability. The predictive performance of XGBoost is assessed and compared with that of the other forecasting models in Section 7.5.

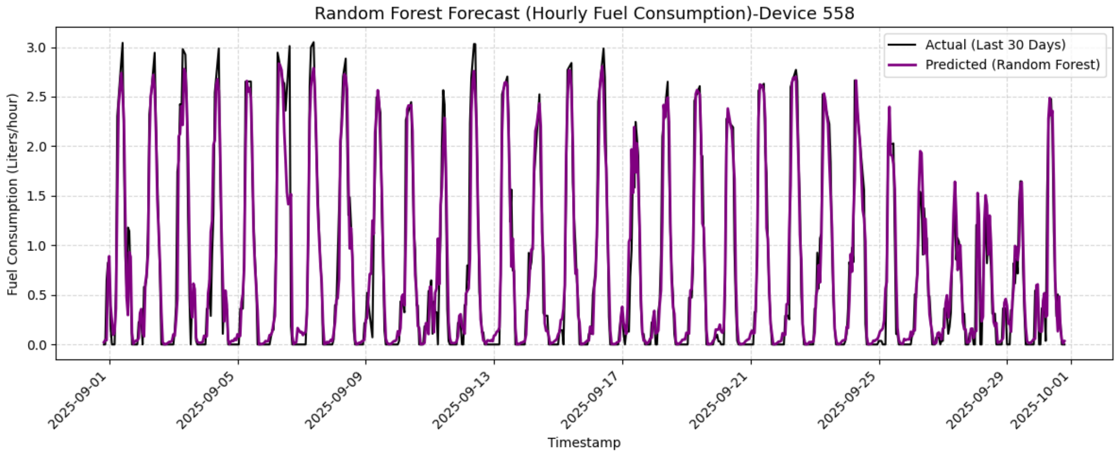

Figure 7.4.: Hourly fuel consumption forecast generated using the Random Forest Regressor over the final 30-day test period. The plot compares observed values with model predictions, illustrating the ensemble model's ability to capture nonlinear consumption patterns while providing stable and robust forecasts.

7.4 Deep Learning–Based Forecasting Model

In addition to statistical and machine learning approaches, a deep learning model was implemented to evaluate the benefits of sequence-based representation learning for hourly fuel consumption forecasting. Deep learning methods are particularly well suited for time series data due to their ability to learn complex temporal dependencies directly from sequential inputs, without requiring strong assumptions about data stationarity or linearity (Lim and Zohren, 2021).

The deep learning approach adopted in this study follows a supervised learning formulation, where historical sequences of temperature-corrected fuel consumption and associated contextual features are used to predict future consumption values. Unlike tree-based models, which rely on explicitly engineered lag features, deep learning models can internally learn temporal representations that capture both short-term and longer-term dependencies in the data.

To ensure comparability with the machine learning models, the same feature set described in Section 7.1 was used as input to the deep learning model. However, the input data were reshaped into a three-dimensional structure of samples, time steps, and features, allowing the model to process temporal information across multiple parallel feature streams. Prior to training, all input features were normalized to improve numerical stability and convergence during optimization.

The Long Short-Term Memory (LSTM) network was selected as the deep learning architecture for this study due to its ability to mitigate the vanishing gradient problem and retain information over extended temporal horizons (Hochreiter and Schmidhuber, 1997). The following subsection details the LSTM architecture, training configuration, and its application to hourly fuel consumption forecasting.

7.4.1 Long Short-Term Memory (LSTM) Network

The LSTM architecture used in this study consisted of a single recurrent layer with 64 hidden units, followed by fully connected layers that produced the hourly fuel consumption forecast. This configuration was selected to balance model expressiveness and computational efficiency for short-term forecasting.

Model training was performed using the Adam optimizer with a learning rate of 0.001, and Mean Squared Error (MSE) was adopted as the loss function. The model was trained on the historical portion of the dataset and evaluated on the final 30-day test period, ensuring a consistent evaluation protocol across all forecasting approaches.

Figure 7.5 illustrates the hourly fuel consumption forecast generated by the LSTM model over the test period. The comparison between observed and predicted values highlights the model's ability to track temporal dynamics in fuel consumption. A quantitative comparison of LSTM performance with statistical and machine learning models is presented in Section 7.5.

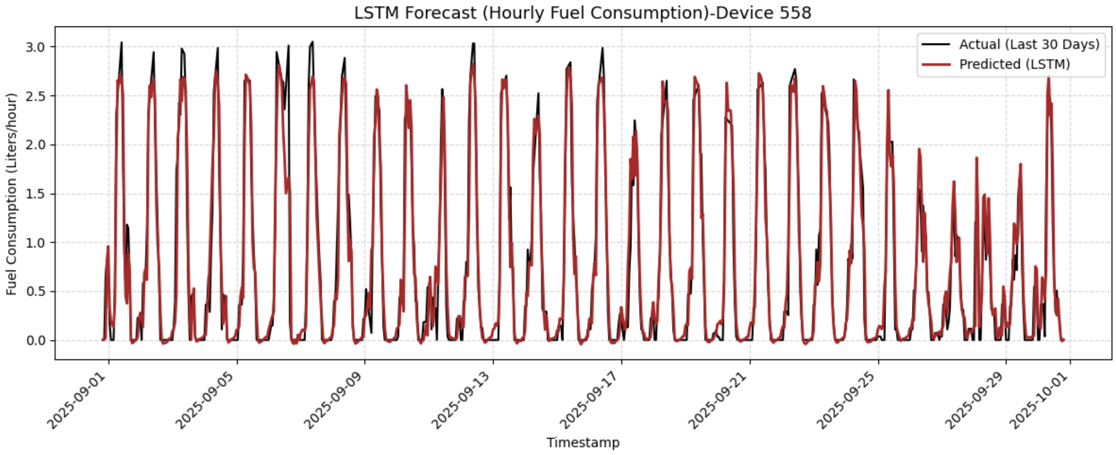

Figure 7.5.: Hourly fuel consumption forecast generated using the LSTM network over the final 30-day test period. The plot compares observed values with model predictions, illustrating the LSTM's ability to learn temporal dependencies from sequential input data.

7.5 Forecast Evaluation and Model Comparison

The forecasting models were evaluated using multiple complementary performance metrics to provide a robust comparison of predictive accuracy over the final 30-day hourly test period. The selected metrics include Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Normalized Root Mean Squared Error (NRMSE), and the coefficient of determination (R^2). Together, these measures capture absolute error magnitude, sensitivity to large deviations, and overall goodness of fit.

Table 7.1 summarizes the forecasting performance of the baseline statistical model (ETS), the machine learning models (LightGBM, XGBoost, and Random Forest), and the deep learning model (LSTM).

Table 7.1.: Forecasting performance comparison for hourly fuel consumption over the 30-day test period. Lower values indicate better performance for MAE, RMSE, and NRMSE, while higher values indicate better performance for R^2.

Model	**MAE**	**RMSE**	**NRMSE (range)**	$\mathbf{R^2}$
ETS	0.398	0.569	18.66%	0.650
LightGBM	0.172	0.275	9.01%	0.918
XGBoost	0.171	0.278	9.13%	0.916
Random Forest	0.172	0.282	9.26%	0.914
LSTM	0.165	0.255	8.38%	0.929

-

The ETS model exhibits substantially higher error values across all metrics, indicating limited ability to capture the complex and nonlinear dynamics present in the hourly fuel consumption series. In contrast, all machine learning and deep learning models achieve markedly lower errors and higher R^2 values, demonstrating their effectiveness in modeling short-term consumption patterns.

Among the machine learning approaches, LightGBM, XGBoost, and Random Forest show very similar performance, with only marginal differences in error metrics. This suggests that tree-based ensemble methods are consistently effective for this forecasting task when provided with appropriately engineered features.

The LSTM model achieves the best overall performance, yielding the lowest MAE, RMSE, and NRMSE values, as well as the highest R^2. This indicates a superior ability to track temporal dynamics in hourly fuel consumption over the test period. However, the performance gains over the best-performing machine learning models are relatively modest, suggesting that ensemble-based approaches already capture a large portion of the predictive structure in the data.

Overall, these results demonstrate that advanced machine learning and deep learning models significantly outperform the baseline statistical approach for short-term hourly fuel consumption forecasting, while differences among the top-performing models remain comparatively small.

8. Visualization and Dashboard-Based Analysis

8.1 Purpose and Role of Visualization

Visualization plays a central role in this study by transforming complex analytical outputs into interpretable and actionable insights for operational decision-making. While the preceding chapters focus on data cleaning, statistical analysis, and forecasting, visualization provides the interface through which these results can be effectively explored, validated, and communicated to non-technical stakeholders.

In the context of IoT-based fuel monitoring, large volumes of high-frequency sensor data and derived analytical outputs can be difficult to interpret through numerical summaries alone. Interactive visual representations enable users to quickly identify consumption trends, recurring patterns, refilling events, anomalies, and forecasted behavior across time and devices. This supports both retrospective analysis and near-real-time situational awareness.

Visualization also serves an important validation function within the analytical workflow. By visually comparing raw and cleaned data, temperature-corrected series, decomposition components, and forecasting outputs, potential modeling artifacts or data quality issues can be more readily detected. This visual inspection complements quantitative evaluation metrics and increases confidence in the analytical results.

Finally, the visualization layer bridges the gap between advanced analytical methods and practical mission operations. By integrating results into a unified dashboard, the system enables decision-makers to monitor fuel usage, investigate irregular behavior, and support planning and resource allocation based on evidence derived from the data.

8.2 Dashboard Design and Data Integration

The Power BI dashboard was designed as an integrated decision-support interface that consolidates the analytical outputs developed throughout this thesis. Its primary objective is to transform cleaned, validated, and modeled IoT-based fuel data into an interactive environment that supports monitoring, exploration, and short-term forecasting of fuel consumption across mission sites.

The dashboard integrates multiple data layers, including raw and cleaned fuel-level measurements, derived consumption metrics, anomaly detection outputs, spatial information, and forecasting results. These datasets were harmonized using consistent temporal indexing and device identifiers, enabling seamless interaction between spatial, temporal, and analytical views.

An Overview and Map page provides a high-level summary of the monitoring system. This page displays key indicators, such as the number of monitored tanks, alongside interactive slicers for device identifiers and time periods. A geographic map visualizes tank locations using latitude and longitude coordinates, with marker attributes representing aggregated fuel consumption. This spatial overview supports rapid identification of high-usage sites and facilitates geographic comparison across the mission area. Figure 8.1 presents the Overview and Map page of the Power BI dashboard.

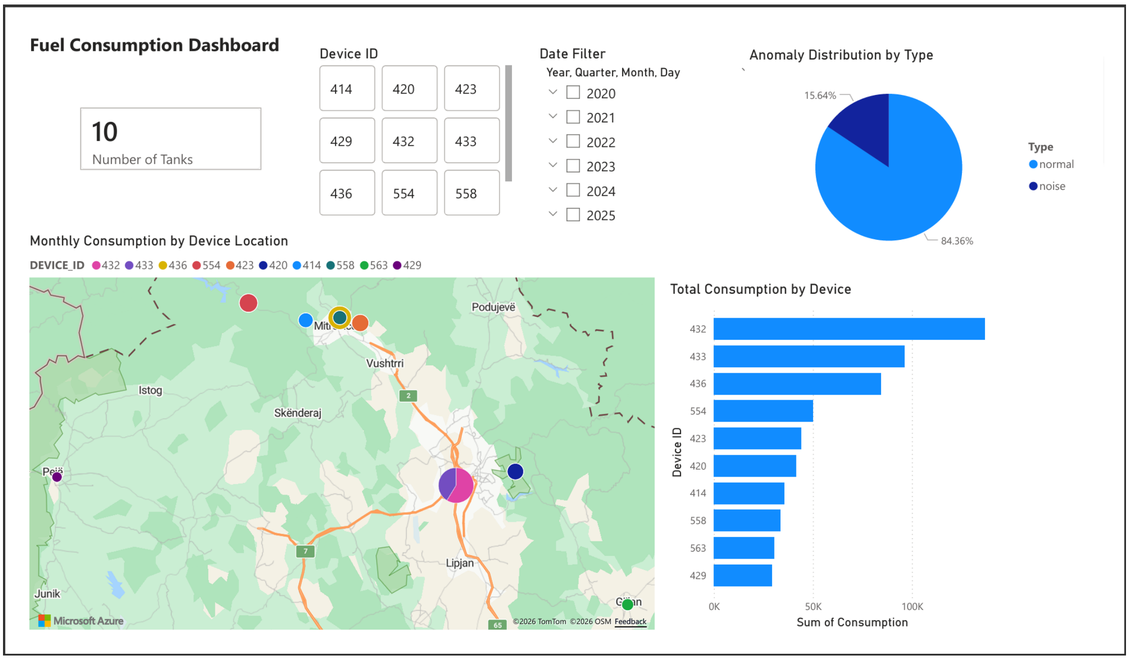

Figure 8.1.: Overview and Map page of the Power BI fuel consumption dashboard. The dashboard displays the spatial distribution of monitored fuel tanks together with aggregated fuel consumption indicators and interactive filters for device selection and temporal exploration.

In addition to the overview page, the dashboard is organized into thematic pages corresponding to the main analytical components of this study. A dedicated page focuses on anomaly detection and data cleaning, enabling visual inspection of raw fuel-level signals, detected refilling events, noise spikes, and the resulting cleaned time series. Another page supports exploratory analysis, presenting monthly, daily, and hourly fuel consumption patterns to reveal seasonal behavior and operational cycles. A final page is dedicated to forecasting, where short-term hourly predictions are compared against observed consumption to support evaluation of predictive performance.

Across all pages, interactive filters, based on device Identifier (ID), date range, and temporal hierarchy, allow users to dynamically link spatial context with temporal behavior and analytical results. By integrating data quality assessment, exploratory analysis, and forecasting outputs within a unified dashboard, the visualization framework operationalizes the analytical methodology developed in this thesis and provides actionable insights for fuel management decision-making. Detailed descriptions of individual visualization components are provided in Sections 8.2.1– 8.2.3.

8.2.1 Visualization of Anomalies and Events

The visualization of anomalies and refilling events plays a critical role in validating data quality and supporting operational monitoring. To this end, a dedicated dashboard page was designed to present detected anomalies alongside the original and cleaned fuel-level time series.

Figure 8.2 illustrates the anomaly detection and data cleaning page of the Power BI dashboard. The upper panel displays the raw fuel-level measurements over time, with detected refilling events and noise spikes explicitly marked using distinct visual indicators. This representation allows users to visually assess the classification results produced by the spike detection and validation framework described in Chapter 4.

The lower panel presents the corresponding cleaned fuel-level time series after the removal of noise spikes and interpolation. By juxtaposing raw and cleaned signals, the dashboard enables direct verification that genuine refilling events are preserved while spurious sensor-induced fluctuations are effectively removed. Interactive filters for device selection and time range allow users to examine individual tanks in detail and inspect specific events of interest.

Overall, this visualization supports transparency and trust in the preprocessing pipeline by making anomaly detection outcomes directly observable. It provides an intuitive mechanism for identifying irregular events, validating cleaning decisions, and ensuring that subsequent analyses are based on reliable and physically consistent fuel-level data.

8.2.2 Visualization of Fuel Consumption Patterns

To support exploratory analysis and operational interpretation, a dedicated dashboard page was designed to visualize temporal patterns in fuel consumption across multiple time scales. This page focuses on summarizing long-term trends, seasonal behavior, and short-term operational cycles derived from the cleaned fuel consumption data.

Figure 8.3 presents the fuel consumption patterns visualization page of the Power

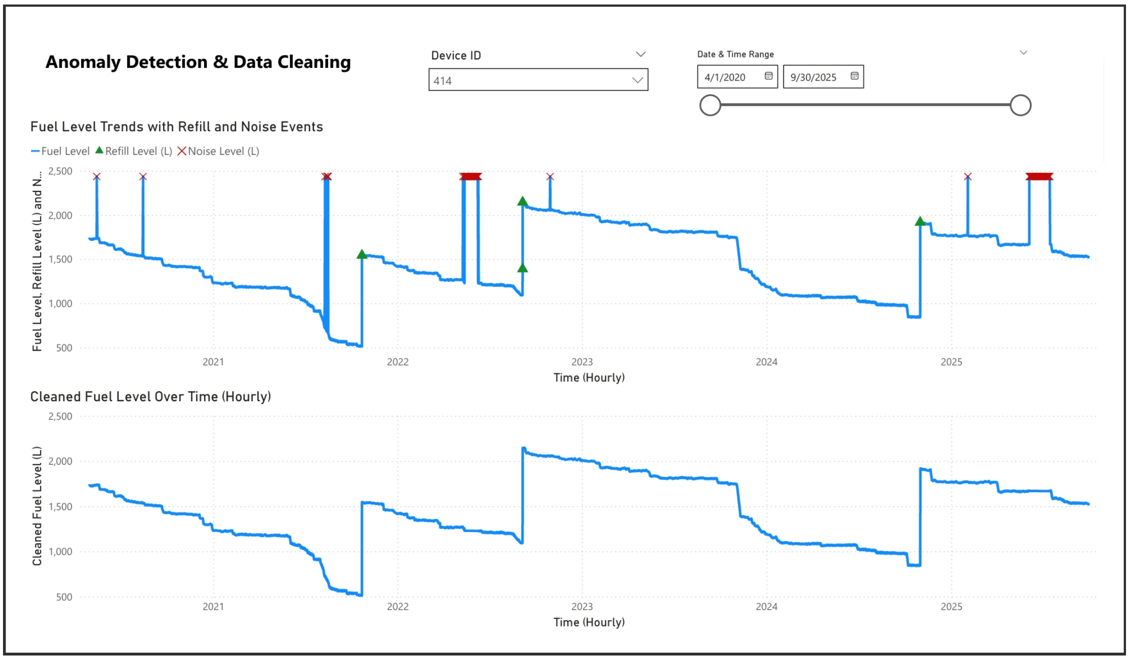

Figure 8.2.: Anomaly detection and data cleaning page of the Power BI dashboard. The upper panel shows raw fuel-level measurements with detected refilling events and noise spikes highlighted, while the lower panel displays the corresponding cleaned fuel-level time series after noise removal and interpolation.

BI dashboard. Monthly consumption trends are displayed using multi-year line charts, enabling comparison of seasonal behavior and year-to-year variability in fuel usage. This representation highlights recurring peak and low consumption periods and supports identification of long-term changes in operational intensity.

Complementary visualizations summarize short-term consumption behavior. Bar charts illustrate average daily fuel consumption by weekday, revealing weekly operational cycles. Line charts of average hourly fuel consumption depict intra-day usage patterns, highlighting typical operating hours and idle periods.

Interactive filters allow users to select individual tanks and time ranges, dynamically updating all visual elements. This enables both comparison and detailed inspection of individual tanks. By combining monthly, weekly, and hourly perspectives, the dashboard provides a comprehensive view of fuel consumption dynamics and supports data-driven interpretation of operational behavior.

8.2.3 Visualization of Forecasting Results

To support interpretation and evaluation of predictive models, a dedicated dashboard page was designed to visualize short-term fuel consumption forecasts. This page focuses on communicating forecasting behavior in an intuitive manner, enabling users to assess expected consumption patterns and identify potential deviations from

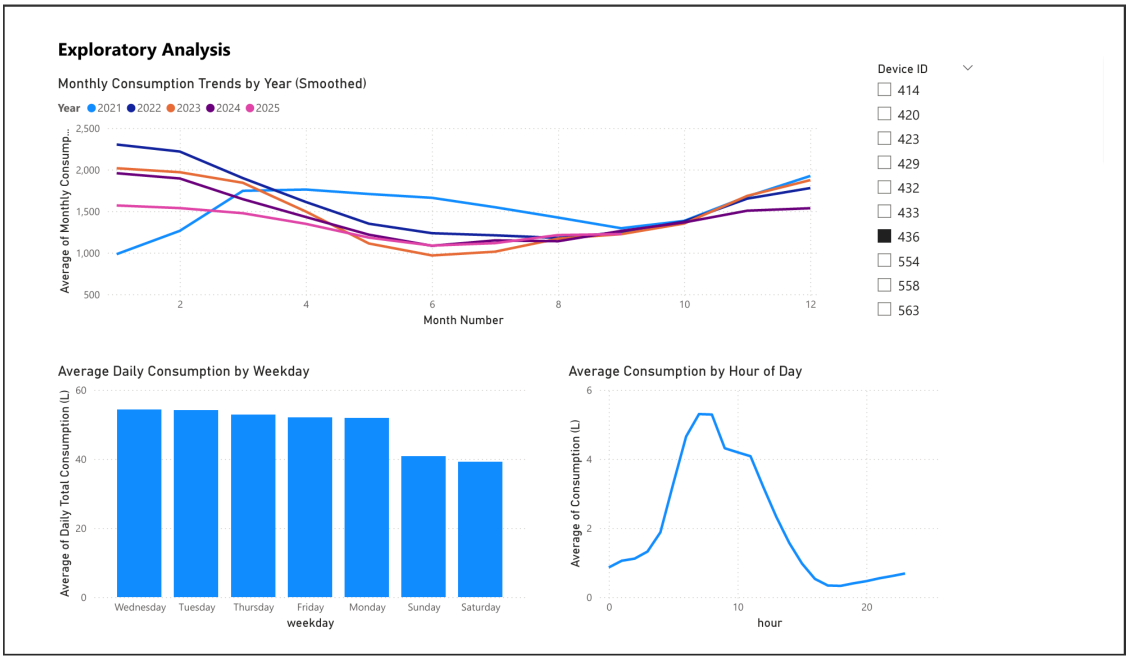

Device ID

☐ 414
☐ 420
☐ 423
☐ 429
☐ 432
☐ 433
■ 436
☐ 554
☐ 558
☐ 563

Figure 8.3.: Exploratory analysis visualization page of the Power BI dashboard. The page displays monthly consumption trends across multiple years, average daily fuel usage by weekday, and average hourly consumption profiles, supporting analysis of seasonal, weekly, and intra-day fuel usage behavior.

predicted values.

Figure 8.4 presents the forecasting results visualization page of the Power BI dashboard. Time-series plots compare observed hourly fuel consumption with model-generated forecasts over the final 30-day horizon. This representation allows users to assess forecast accuracy across different models, temporal alignment, and the model's ability to capture short-term fluctuations in fuel usage. Interactive filters enable selection of individual tanks and methods, allowing users to examine forecasting performance at different levels.

Overall, the forecasting visualization complements the quantitative evaluation presented in Chapter 7 by providing an accessible and interpretable view of predictive model outputs within the broader analytical workflow.

8.3 Summary of Visualization Insights

The Power BI dashboard developed in this study demonstrates how interactive visualization can effectively translate complex analytical results into actionable operational insights. By integrating data quality indicators, temporal analysis, spatial context, and forecasting outputs within a unified interface, the dashboard supports comprehensive monitoring of fuel consumption behavior across mission sites.

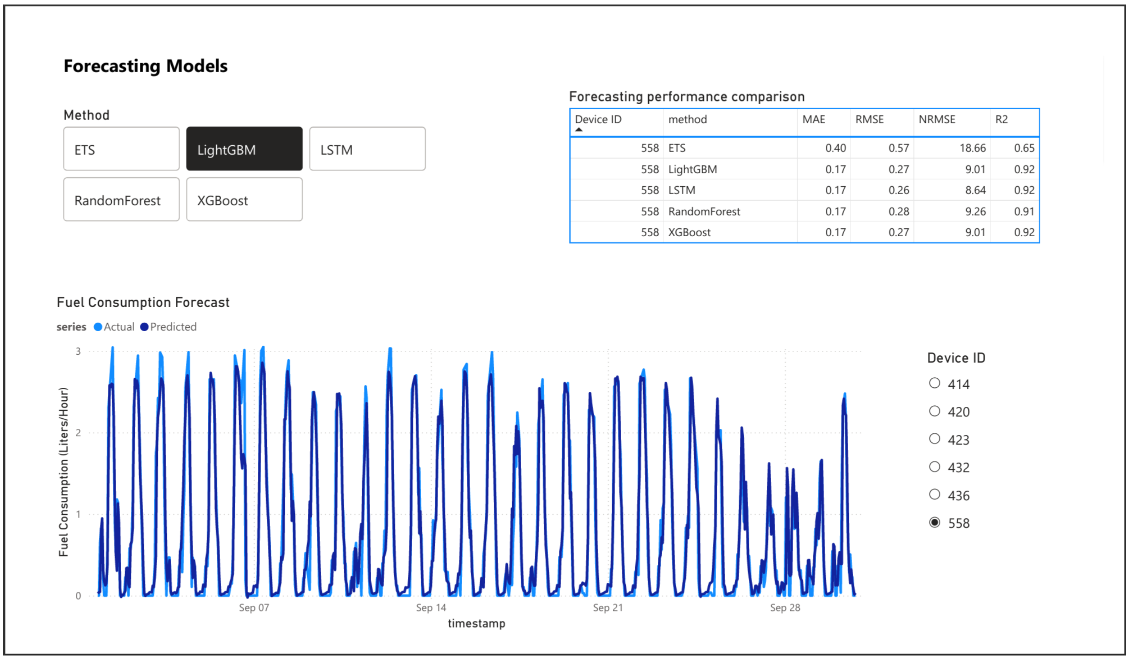

Forecasting performance comparison

Device ID	method	MAE	RMSE	NRMSE	R2
558	ETS	0.40	0.57	18.66	0.65
558	LightGBM	0.17	0.27	9.01	0.92
558	LSTM	0.17	0.26	8.64	0.92
558	RandomForest	0.17	0.28	9.26	0.91
558	XGBoost	0.17	0.27	9.01	0.92

Device ID
○ 414
○ 420
○ 423
○ 432
○ 436
● 558

Figure 8.4.: Forecasting results visualization page of the Power BI dashboard. The figure compares observed and predicted hourly fuel consumption over a 30-day horizon, enabling visual assessment of short-term forecasting performance.

Visualization of anomalies and events enables transparent validation of data cleaning and spike classification procedures, allowing users to directly inspect refilling events and noise removal outcomes. Temporal consumption visualizations reveal long-term trends, seasonal behavior, and recurring daily and weekly operational patterns, providing contextual understanding of fuel usage dynamics. Forecasting visualizations further extend this capability by presenting short-term predictions alongside historical observations, supporting planning and early detection of potential deviations from expected consumption.

Overall, the visualization framework complements the analytical methods presented in earlier chapters by enhancing interpretability, facilitating exploratory analysis, and bridging the gap between advanced data analysis and practical decision-making. The dashboard operationalizes the proposed methodology and provides a reusable foundation for monitoring, analysis, and forecasting of IoT-based fuel consumption data in operational environments.

9. Overall Discussion

This chapter provides an integrated interpretation of the results presented in the previous chapters. While methodological steps and corresponding findings were discussed individually, this section synthesizes the outcomes of data preprocessing, exploratory analysis, spatial assessment, forecasting performance, and dashboard implementation. The findings are interpreted in light of the research objectives and existing literature.

9.1 Impact of Systematic Preprocessing on Analytical Reliability

A central objective of this study was to transform raw IoT-based fuel-level measurements into a reliable and physically consistent analytical time series. The results demonstrate that systematic preprocessing substantially altered the statistical properties and interpretability of the signal.

The data cleaning and preprocessing stage followed a structured multi-stage pipeline. It began with an initial data quality assessment that identified and removed a small number of physically implausible observations caused by sensor malfunctions, where recorded fuel levels violated known tank capacity constraints. These invalid measurements were excluded prior to further processing to prevent contamination of downstream analyses.

The core of the framework focused on detecting and classifying abrupt fuel-level changes. Rather than treating all spikes as anomalies, the pipeline distinguished between genuine refilling events and sensor-induced noise by evaluating both the magnitude of each abrupt change and the stability of the fuel level in the subsequent period. A reverse validation step was introduced to detect cases in which apparent refilling events reverted to pre-spike levels, indicating measurement artefacts rather than true operational events. Importantly, all thresholds and validation parameters were empirically calibrated from the dataset itself, ensuring that the procedure reflected actual operational behavior instead of relying on arbitrary assumptions.

Following spike classification, observations identified as noise were corrected through interpolation, restoring temporal continuity without modifying legitimate refilling events or the underlying consumption trend. In addition, temperature correction was applied to mitigate the effect of thermal expansion on fuel-level measurements. This step reduced artificial volatility caused by environmental temperature fluctuations and produced a more physically meaningful representation of actual consumption dynamics. Visual comparison between raw and temperature-corrected series confirmed that while the long-term trend remained consistent, high-frequency fluctuations were significantly attenuated.

Subsequently, STL-based decomposition was applied, and the residual component was analyzed using the IQR method to detect and remove remaining extreme deviations. This additional refinement further stabilized the signal while preserving genuine consumption behavior. The resulting time series exhibited reduced noise, improved continuity, and clearer temporal structure, making it suitable for spatial assessment and short-term forecasting.

These findings align with prior research emphasizing that data quality management is critical in IoT-based energy monitoring systems, where sensor noise and environmental effects can distort analytical outputs (Goknil et al., 2023; Bertrand et al., 2023). Without systematic data cleaning, calibration, and correction, measurement errors would likely propagate into modeling stages and degrade predictive accuracy.

Overall, the preprocessing framework was not merely a technical necessity but a foundational component that enhanced analytical reliability and enabled robust downstream temporal, spatial, and forecasting analyses.

9.2 Temporal Consumption Patterns

To investigate the temporal structure of diesel consumption, STL decomposition was applied to the cleaned daily fuel consumption time series. The method separated the observed signal into trend, seasonal, and residual components, enabling a structured interpretation of consumption dynamics.

Building on the decomposition results, monthly and seasonal analyses further revealed recurring consumption patterns across years, with certain periods consistently exhibiting elevated or reduced usage relative to the long-term average. Analysis of daily and weekly consumption patterns confirmed clear intra-day and intra-week rhythms driven by operational schedules, with peak usage concentrated during working hours on weekdays and noticeably lower activity during nights and weekends. Similar periodic dynamics have been documented in energy demand forecasting studies, where operational and human activity patterns strongly influence temporal energy usage (Tyralis et al., 2017; Barkhordar et al., 2022).

The residual component, following systematic preprocessing and outlier correction, was largely centered around zero, indicating that the dominant trend and seasonal structures had been effectively captured. Comparing residuals before and after outlier cleaning demonstrated a clear improvement in stability, with the corrected series showing substantially reduced irregular variability. This confirms that preprocessing enhanced the clarity and interpretability of the decomposed components.

Residual-based analysis also enabled the identification of statistically significant deviations from expected consumption behavior.

Importantly, the cleaned residual component provided a stable foundation for subsequent modeling tasks. This supports the argument that decomposition-based preprocessing enhances both interpretability and predictive robustness, as discussed in time-series literature.

These findings justify the use of time-series-based forecasting approaches, as the underlying process exhibits predictable temporal dynamics suitable for short-term prediction.

9.3 Spatial Analysis Considerations

The spatial autocorrelation assessment using Global Moran's I did not reveal strong statistically stable spatial clustering. The computed indices were consistently negative and statistically non-significant for both raw and temperature-adjusted series, indicating no evidence of positive spatial clustering.

These results suggest that fuel consumption behavior is largely driven by local operational conditions rather than geographic proximity. In other words, neighboring tanks do not exhibit systematically similar usage patterns. The comparable outcomes obtained before and after temperature correction further indicate that measurement adjustments did not influence the spatial structure of the data.

It should be noted, however, that the strength of spatial inference is constrained by the limited number of monitored tanks and their geographic configuration, which is approximately linear. Previous research has shown that Moran's I is sensitive to sample size and spatial configuration (Bivand et al., 2009), and small or unevenly distributed datasets can reduce its statistical power. Therefore, spatial findings in this study should be interpreted as exploratory rather than definitive.

Nevertheless, incorporating spatial diagnostics represents a methodological contribution by extending fuel consumption analysis beyond purely temporal modeling. From a modeling perspective, the findings support the decision to adopt independent per-tank time-series forecasting models rather than spatial or spatial–temporal approaches.

9.4 Forecasting Performance in Context

The forecasting experiments compared a classical statistical model (ETS) with machine learning and deep learning approaches for short-term hourly diesel consumption prediction over a 30-day horizon. All models were trained on temperature-corrected consumption data, and a consistent feature engineering framework was applied to the machine learning and deep learning models, incorporating lagged consumption values, rolling statistics, temperature variables, and calendar-based operational indicators.

The ETS model, which relies solely on the internal temporal structure of the series, achieved an R^2 of 0.650 and served as a baseline reference. In contrast, the machine learning models (LightGBM, XGBoost, and Random Forest) achieved R^2 values above 0.91, with RMSE values reduced by more than 50% compared to ETS. These results highlight the advantage of incorporating multivariate contextual information and nonlinear relationships into the forecasting framework.

Among all models, the LSTM achieved the highest predictive performance ($R^2 = 0.929$; RMSE reduced by approximately 55% relative to ETS), demonstrating its capability to capture sequential temporal dependencies in hourly consumption data. However, the performance gap between LSTM and the tree-based ensemble models was relatively small, suggesting that well-engineered features combined with robust ensemble methods are already sufficient to capture most of the predictive structure present in the dataset.

These findings are consistent with prior research indicating that machine learning and deep learning approaches tend to outperform traditional statistical models in energy forecasting tasks characterized by nonlinearity and multivariate interactions (Vivas et al., 2020; Dou et al., 2023). At the same time, the modest differences among the advanced models underscore the critical role of preprocessing quality and feature design, which may be as influential as model complexity in determining forecasting accuracy.

9.5 Operational Implications and Decision Support

Beyond predictive accuracy, the practical value of this study lies in the translation of analytical outputs into an operational monitoring framework. Prior research has shown that visualization-driven decision support systems improve interpretability and facilitate the adoption of advanced analytics in operational environments (Valdivia and Baca, 2025). In this study, analytical results are embedded within an interactive Power BI dashboard that enables structured exploration of fuel consumption behavior.

The dashboard integrates validated and temperature-corrected data, time-series decomposition outputs, and forecasting results within a unified interface. Users can examine long-term trends, weekly and daily operational cycles, and short-term predictions alongside historical observations. This integrated view supports contextual understanding of consumption patterns and enhances transparency in the analytical process.

By presenting forecasted consumption trajectories together with recent observations, the system supports short-term planning and informed decision-making. Rather than replacing operational expertise, the dashboard acts as a structured analytical support tool that enhances situational awareness and facilitates data-driven fuel management.

Overall, the visualization component operationalizes the methodological pipeline developed in this thesis, bridging the gap between advanced time-series modeling and practical monitoring applications in mission environments.

9.6 Overall Contribution

Taken together, the findings demonstrate that the structured application of the CRISP-DM methodology, encompassing systematic preprocessing, decomposition-based temporal analysis, comparative forecasting using statistical, machine learning, and deep learning models, and interactive visualization, provides a coherent and operationally relevant framework for IoT-based diesel consumption analytics.

While individual components such as time-series forecasting or exploratory analysis are well established in the literature, this study contributes by integrating these elements into a unified analytical pipeline tailored to mission-scale fuel monitoring. The work demonstrates how rigorous data preparation, decomposition-informed modeling, and feature engineering collectively enhance predictive reliability and interpretability.

By linking methodological rigor with practical deployment through an interactive dashboard, the thesis advances an end-to-end framework that supports both analytical robustness and operational usability in resource-constrained mission environments.

10. Conclusion

This thesis addressed the challenges of analyzing IoT-based diesel fuel consumption data by developing an integrated framework that combines data cleaning, analytical modeling, forecasting, and visualization to support operational decision-making. The study was guided by four research questions, each targeting a key aspect of reliable fuel monitoring and analysis.

Regarding Research Question 1 (RQ1), which investigated how IoT-based diesel consumption data can be effectively cleaned, integrated, and validated, a systematic data preprocessing framework was developed. This framework included data quality assessment, detection and classification of fuel-level spikes, empirical calibration of validation parameters, and post-detection cleaning through interpolation. The proposed approach successfully removed sensor-induced noise while preserving genuine refilling events, resulting in physically consistent and analytically reliable fuel-level and consumption time series. These steps established a robust foundation for subsequent temporal, spatial, and forecasting analyses.

In relation to Research Question 2 (RQ2), which focused on identifying statistical, temporal, and spatial patterns in the cleaned dataset, multiple analytical techniques were applied. STL decomposition revealed clear long-term trends, recurring seasonal behavior, and residual components capturing irregular consumption patterns. Exploratory analyses further highlighted consistent daily and weekly operational cycles, as well as monthly and interannual consumption dynamics. Spatial autocorrelation analysis using Moran's I consistently indicated weak and statistically insignificant spatial dependence across tanks, both before and after temperature correction. These findings suggest that fuel consumption behavior is primarily driven by site-specific operational factors rather than geographic proximity. However, spatial inference is constrained by the limited number of monitored tanks and their geographic configuration, which reduces the statistical power of global spatial indicators and limits the strength of generalizable spatial conclusions.

With respect to Research Question 3 (RQ3), which investigated how short-term diesel consumption can be forecast using classical statistical, machine learning, and deep learning models, multiple forecasting approaches were evaluated using hourly data over a 30-day horizon. The baseline statistical model (ETS) demonstrated limited predictive capability for high-frequency consumption data. In contrast, machine learning and deep learning models achieved substantially improved performance, indicating their effectiveness in capturing nonlinear relationships and short-term temporal dynamics. While the LSTM model achieved the highest overall accuracy, the performance differences among the best-performing models were relatively modest, suggesting that ensemble-based machine learning approaches already capture much of the predictive structure present in the data.

Addressing Research Question 4 (RQ4), which examined how analytical results can be effectively communicated through interactive visualization, a Power BI dashboard was developed as an integrated decision-support tool. The dashboard combines spatial context, anomaly detection outputs, temporal consumption patterns, and forecasting results within a unified interface. By enabling interactive exploration across tanks, time periods, and analytical perspectives, the dashboard enhances transparency, interpretability, and accessibility of complex analytical outputs for decision-makers.

Overall, this study demonstrates that integrating robust data cleaning, temporal and spatial analysis, forecasting, and interactive visualization provides a reusable and effective methodology for analyzing IoT-based diesel fuel consumption data. The proposed framework enhances data reliability, reveals meaningful consumption patterns, and translates analytical results into actionable operational insights.

10.1 Limitations

Despite the contributions of this study, several limitations should be acknowledged.

First, Spatial correlation results are limited by the very small number of tanks. Moran's I is sensitive to sample size and spatial configuration, and statistical significance is difficult to establish with few spatial units. Therefore, the lack of significant spatial autocorrelation should be interpreted cautiously, as it may reflect limited statistical power.

Second, forecasting performance was evaluated over a fixed 30-day horizon, and longer-term predictive behavior was not assessed. Additionally, while multiple models were compared, hyperparameter tuning was intentionally limited to maintain methodological consistency and interpretability. More extensive tuning or model-specific optimization could potentially improve predictive performance.

10.2 Future Work

Future research could extend this work in several directions.

The residual component extracted from STL decomposition provides a promising basis for detecting abnormal fuel consumption behavior, such as potential leaks, theft, or atypical usage patterns, and could be further explored for automated anomaly detection and alerting.

Forecasting models could be evaluated over longer horizons and enhanced through more extensive hyperparameter optimization or hybrid modeling approaches. Incorporating additional contextual information, such as operational schedules or asset utilization data, may further improve predictive performance.

From a visualization perspective, integrating real-time data streaming and automated alerts into the dashboard would significantly enhance its applicability for continuous monitoring and proactive fuel management in operational environments.

Besides, the proposed analytical framework could be extended beyond diesel fuel to other tank-stored liquids, such as gasoline or chemical supplies. Adapting the methodology to different liquid types and operational contexts would help assess its generalizability and support broader resource monitoring applications.

Beyond monitoring and forecasting, future research could also explore how the analytical outputs can be directly integrated into operational optimization. Forecasting results could support optimized refilling schedules by predicting when tank levels are expected to reach critical thresholds, thereby reducing the risk of stockouts while minimizing unnecessary refilling trips. Additionally, combining consumption forecasts with logistical constraints (e.g., delivery routes, vehicle availability, or refilling costs) could enable the development of decision-support tools for cost-efficient fuel supply planning. Such extensions would move the system from descriptive and predictive analytics toward prescriptive analytics, enhancing its value for operational resource management.

A. Appendix

A.1 Accessing the Code Repository

The complete code repository for this thesis is publicly available at `https://github.com/ElaheTorabi/Masters-Thesis.git`.

A.2 Z-score Spike Detection and Classification for Remaining Tanks

This appendix presents additional Z-score–based spike detection and refill classification results for the remaining analyzed fuel tanks not shown in the main chapters. The figures included here provide supplementary visual evidence of the methodology described in Chapter 4 and demonstrate the consistency of the proposed detection and validation framework across different tanks. All parameters and decision criteria used in these figures are identical to those defined and calibrated in Section 4.4.

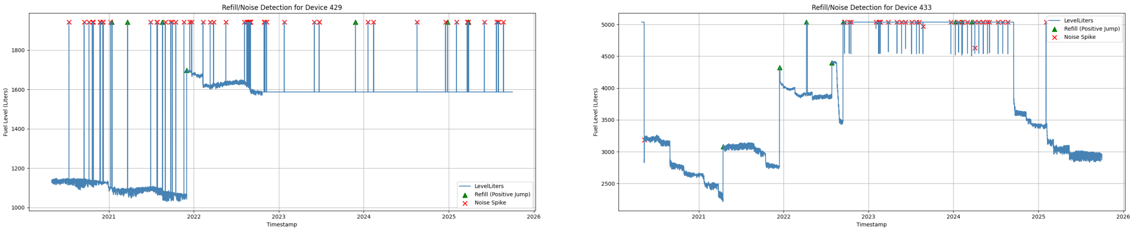

(a) Tank 429

(b) Tank 433

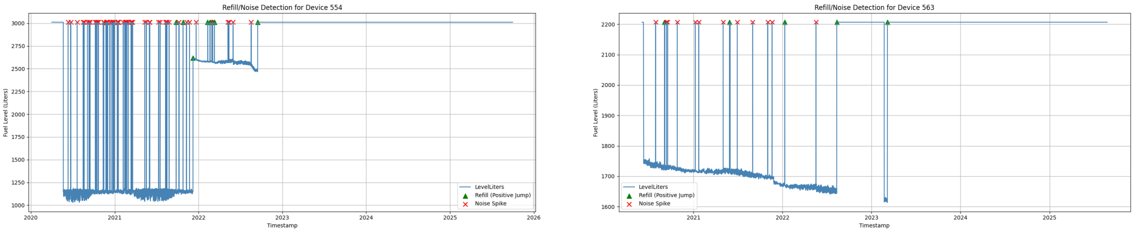

(c) Tank 554

(d) Tank 563

Figure A.1.: Z-score anomaly candidates and refill versus noise classification for additional fuel tanks. The same empirically calibrated parameters and validation criteria described in Chapter 4 were applied.

A.3 Post-Detection Cleaned Fuel-Level Time Series for Remaining Tanks

This appendix presents the cleaned fuel-level time series for the remaining analyzed tanks that are not shown in the main body of the thesis. These figures complement the example results presented in Section 4.5 and demonstrate the consistent application of the post-detection noise correction and interpolation procedure across all tanks. All cleaning steps and parameters applied here are identical to those described in Chapter 4.

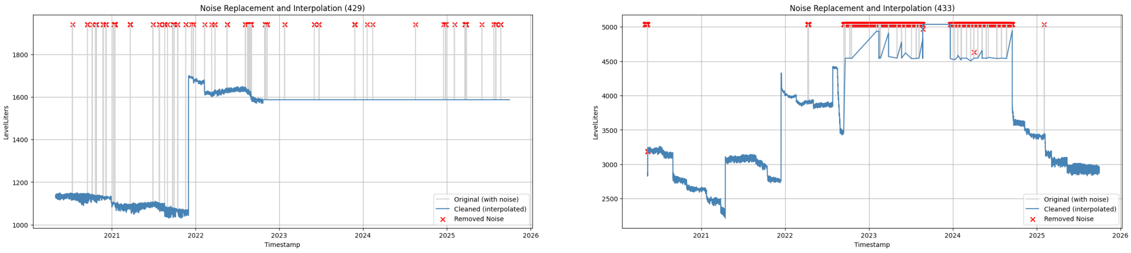

(a) Tank 429 (b) Tank 433

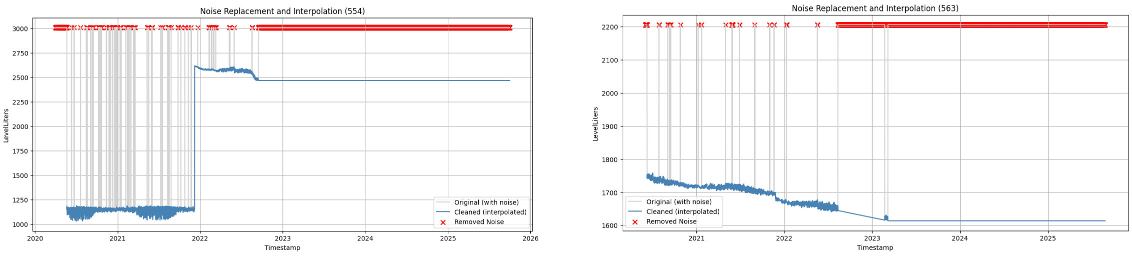

(c) Tank 554 (d) Tank 563

Figure A.2.: Cleaned fuel-level time series for additional fuel tanks following post-detection noise correction.

Bibliography

Abediasl, H., Ansari, A., Hosseini, V., Koch, C. R., and Shahbakhti, M. (2023). Real-time vehicular fuel consumption estimation using machine learning and on-board diagnostics data. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 238(12):3779–3793.

Alsalehy, A. and Bailey, M. (2025). Improving time series data quality: Identifying outliers and handling missing values in a multilocation gas and weather dataset. *Smart Cities*, 8:82.

Anand, B. A., Kathyayini, H. S., and Kumar, A. (2024). An iot-based fuel monitoring and tracking system with real-time fuel level detection using accelerometer. *Journal of Scientific Research and Reports*, 30(11):1064–1070.

Ayebazibwe, S. (2023). The impact of fuel management on supplier performance in ministry of energy and mineral development: A case study of ministry of energy and mineral development.

Barkhordar, Z. A., Habibzadeh, S., and Alizadeh, N. (2022). Deriving electricity consumption patterns using a decomposition approach. *Results in Engineering*, 16:100628.

Bertrand, Y., Van Belle, R., De Weerdt, J., and Serral, E. (2023). Defining data quality issues in process mining with iot data. In Montali, M., Senderovich, A., and Weidlich, M., editors, *Process Mining Workshops*, pages 422–434, Cham. Springer Nature Switzerland.

Bhuiyan, M. M. H., Sakib, A., Alawee, S., and Razzaghi, T. (2024). Fueling the future: A comprehensive analysis and forecast of fuel consumption trends in u.s. electricity generation. *Sustainability*, 16:2388.

Bivand, R., Müller, W. G., and Reder, M. (2009). Power calculations for global and local moran's i. *Computational Statistics Data Analysis*, 53(8):2859–2872.

Breiman, L. (2001). Random forests. *Machine Learning*, 45:5–32.

Byabazaire, J., O'Hare, G., and Delaney, D. (2020). Data quality and trust: Review of challenges and opportunities for data sharing in iot. *Electronics*, 9(12).

Chaudhari, D. S., Talmale, A. P., Phulwale, G. R., Bangal, N. R., and Karande, N. U. (2024). Fuel monitoring system using iot device. *International Research Journal of Modernization in Engineering Technology and Science*, 6(5):519–528.

Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. pages 785–794.

Cleveland, R. B. (1990). Stl : A seasonal-trend decomposition procedure based on loess.

Deb, C., Zhang, F., Yang, J., Lee, S. E., and Shah, K. W. (2017). A review on time series forecasting techniques for building energy consumption. *Renewable and Sustainable Energy Reviews*, 74:902–924.

Dou, Y., Tan, S., and Xie, D. (2023). Comparison of machine learning and statistical methods in the field of renewable energy power generation forecasting: a mini review. *Frontiers in Energy Research*, Volume 11 - 2023.

Dritsaki, C., Niklis, D., and Stamatiou, P. (2021). Oil consumption forecasting using arima models: An empirical study for greece. *International Journal of Energy Economics and Policy*, 11(4):214–224.

Eddaoudi, Z., Aarab, Z., Boudmen, K., Elghazi, A., and Rahmani, M. D. (2024). A brief review of energy consumption forecasting using machine learning models. *Procedia Computer Science*, 236:33–40. International Symposium on Green Technologies and Applications (ISGTA'2023).

Edirisinghe, K. G. J. C., Amarasinghe, A. S. M., and Heelibathdeniya, H. A. N. B. (2024). Real-time monitoring of gas station underground fuel tanks using iot technology. In *17th International Research Conference (KDUIRC-2024)*, pages 148–153.

Farrag, O., Mansour, A., Abed, B., Abu Alhaj, A., Landolsi, T., and Al-Ali, A. R. (2025). Ivemps: Iot-based vehicle emission monitoring and prediction system. *IEEE Access*, 13:95628–95646.

Gardner, E. S. (2006). Exponential smoothing: The state of the art—part ii. *International Journal of Forecasting*, 22(4):637–666.

Gitahi, M. P. and Ogollah, K. (2014). Influence of fleet management practices on service delivery performance to refugees in united nations high commissioner for refugees kenya programme. *European Journal of Business Management*, 2:336–341.

Goknil, A., Nguyen, P., Sen, S., Politaki, D., Niavis, H., Pedersen, K. J., Suyuthi, A., Anand, A., and Ziegenbein, A. (2023). A systematic review of data quality in cps and iot for industry 4.0. *ACM Comput. Surv.*, 55(14s).

Hapsari, A., Prastowo, H., and Pitana, T. (2021). Real-time fuel consumption monitoring system integrated with internet of things (iot). *Kapal: Jurnal Ilmu Pengetahuan dan Teknologi Kelautan*, 18:88–100.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9:1735–1780.

Hong, T., Pinson, P., Wang, Y., Weron, R., Yang, D., and Zareipour, H. (2020). Energy forecasting: A review and outlook. *IEEE Open Access Journal of Power and Energy*, 7.

Hu, H., Gong, S., and Taheri, B. (2024). Energy demand forecasting using convolutional neural network and modified war strategy optimization algorithm. *Heliyon*, 10(6).

Hyndman, R. J. and Athanasopoulos, G. (2021). *Forecasting: Principles and Practice*. OTexts, Melbourne, Australia, 3 edition.

Inchauspe, J., Li, J., and Park, J. (2020). Seasonal patterns of global oil consumption: Implications for long term energy policy. *Journal of Policy Modeling*, 42(3):536–556.

Janczura, J., Trück, S., Weron, R., and Wolff, R. C. (2013). Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling. *Energy Economics*, 38:96–110.

Kabir, R., Remias, S. M., Waddell, J., and Zhu, D. (2023). Time-series fuel consumption prediction assessing delay impacts on energy using vehicular trajectory. *Transportation Research Part D: Transport and Environment*, 117:103678.

Kadu, A., Patil, A., Patil, A., Pardhi, R., and Patil, B. (2023). Iot based fuel theft monitoring system. *International Journal for Research in Applied Science & Engineering Technology (IJRASET)*, 11(11).

Karmaker, S. (2025). *Feature Engineering in Time Series Forecasting.*

Karunya, L., Sugapriya, S. A., Balu, B. P., and Thiagarajan, G. (2024). Fuel smart: Real-time iot surveillance for vehicle. *International Journal of Advanced Research in Computer and Communication Engineering*, 13(2).

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree.

Kee, K.-K. and Simon, B.-Y. L. (2019). Cloud-based iot solution for predictive modeling of ship fuel consumption. In *Proceedings of the 2019 8th International*

Conference on Software and Computer Applications, ICSCA '19, page 44–49, New York, NY, USA. Association for Computing Machinery.

Keim, D., Kohlhammer, J., Ellis, G., and Mansmann, F. (2010). *Mastering The Information Age – Solving Problems with Visual Analytics.*

Khatun, R., Antor, S., Ullah, A., and Hossain, A. (2019). Vehicle fuel activities monitoring system using iot. *Advances in Internet of Things*, 09:63–71.

Lim, B. and Zohren, S. (2021). Time-series forecasting with deep learning: a survey. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194):20200209.

Lin, C.-H., Nuha, U., Lin, G.-Z., and Lee, T.-F. (2022). Hourly power consumption forecasting using robuststl and tcn. *Applied Sciences*, 12(9).

Nita, S. L. and Mihailescu, M. I. (2022). Fuel monitoring system based on iot: Overview and device authentication. In *2022 14th International Conference on Communications (COMM)*, page 1–4. IEEE Press.

North American Council for Freight Efficiency (2019). 2019 annual fleet fuel study. Technical report, North American Council for Freight Efficiency (NACFE), Fort Wayne, IN. Accessed: 2026-01-21.

P, V., V, G., Patwari, P., Singh, R., A, S., Tayade, P., Rajashree, R., and Tamilselvi, M. (2019). Iot based smart fuel monitoring system. *International Journal of Recent Technology and Engineering (IJRTE)*, 8:14–20.

Peña, M., Biscarri, F., Guerrero, J. I., Monedero, I., and León, C. (2016). Rule-based system to detect energy efficiency anomalies in smart buildings, a data mining approach. *Expert Systems with Applications*, 56:242–255.

Shaikh, A., Nazir, A., and Shah, A. S. (2022). Short term energy consumption forecasting using neural basis expansion analysis for interpretable time series. *Scientific Reports*, 12.

Skar, A., Vestergaard, A., Pour, S. M., and Pettinari, M. (2023). Internet-of-things (iot) platform for road energy efficiency monitoring. *Sensors*, 23(5).

Tech Mahindra (2020). Iot analytics: Data quality challenges. Technical report, Tech Mahindra. Whitepaper accessed January 21, 2026.

Teh, H. Y., Kempa-Liehr, A. W., and Wang, K. I. K. (2020). Sensor data quality: A systematic review. *Journal of Big Data*, 7:11.

Tyralis, H., Karakatsanis, G., Tzouka, K., and Mamassis, N. (2017). Exploratory data analysis of the electrical energy demand in the time domain in greece. *Energy*, 134:902–918.

United Nations (2023). United Nations Mission in Kosovo (UNMIK): Mandate and operations. `https://unmik.unmissions.org`. Accessed 2026-01-21.

United Nations Security Council (1999). Resolution 1244 (1999) adopted by the Security Council at its 4011th meeting, on 10 June 1999. `https://undocs.org/S/RES/1244(1999)`.

Valdivia, F. d. L. P. and Baca, H. A. H. (2025). Guidance for interactive visual analysis in multivariate time series preprocessing. *Sensors*, 25(18).

Vashisht, B. (2022). Reducing operational costs through fleet optimization and maintenance strategies. *International Journal of Scientific Advances in Technology*, 13(1). Summarizes ATRI's *An Analysis of the Operational Costs of Trucking* data, including fuel at about 24% of total operational expenses in 2021.

Vegesna, R. (2023). Automated reporting software for fuel usage and dispenser activity. *International Journal of Multidisciplinary Research and Growth Evaluation*, 4:1132–1134.

Vivas, E., Allende-Cid, H., and Salas, R. (2020). A systematic review of statistical and machine learning methods for electrical power forecasting with reported mape score. *Entropy*, 22(12):1412.

Vlachos, M., Lopardo, R., and Amditis, A. (2022). Real time vehicle status monitoring under moving conditions using a low power iot system. *Journal on Internet of Things*, 4(4):235–261.

Wilfling, S. (2023). Augmenting data-driven models for energy systems through feature engineering: A python framework for feature engineering.

Yigitbasioglu, O. M. and Velcu, O. (2012). A review of dashboards in performance management: Implications for design and research. *International Journal of Accounting Information Systems*, 13(1):41–59.

Zafeiriou, A., Chantzis, G., Jonkaitis, T., Fokaides, P., and Papadopoulos, A. (2023). Smart energy strategy - a comparative study of energy consumption forecasting machine learning models. *Chemical Engineering Transactions*, 103:691–696.

Zangrando, N., Fraternali, P., Petri, M., Pinciroli Vago, N., and Herrera, S. (2022). Anomaly detection in quasi-periodic energy consumption data series: a comparison of algorithms. *Energy Informatics*, 5:62.

Zankl, S. (2009). Visualizing sensor data. Seminar paper, Ludwig-Maximilians-Universität München. Media Informatics Advanced Seminar on Information Visualization (2008–2009).

Zhang, L., Wen, J., Li, Y., Chen, J., Ye, Y., Fu, Y., and Livingood, W. (2021). A review of machine learning in building load prediction. *Applied Energy*, 285:116452.

Zhou, E., Gadzanku, S., Hodge, C., Campton, M., de la Rue du Can, S., and Zhang, J. (2023). Best practices in electricity load modeling and forecasting for long-term power system planning. Technical Report OSTI ID: 1973409, Lawrence Berkeley National Laboratory.

2026

Leveraging Statistical Analysis and Interactive Dashboards for Diesel Consumption Monitoring in United Nations Peacekeeping Operations

Elaheh Torabidashti

Masters Program in **Geospatial Technologies**

Supported by:
Education and Culture
ERASMUS MUNDUS