DM

Data-Driven Methodology for Estimation of Cruise Ships and Port Energy Profiles in Small Data Settings

MASTER DISSERTATION

Luis Alberto Gonçalves Nunes
MASTER IN INFORMATICS ENGINEERING

UNIVERSIDADE da MADEIRA
A Nossa Universidade
www.uma.pt

November | 2025

Data-Driven Methodology for Estimation of Cruise Ships and Port Energy Profiles in Small Data Settings

MASTER DISSERTATION

Luis Alberto Gonçalves Nunes
MASTER IN INFORMATICS ENGINEERING

SUPERVISOR
Fábio Ruben Silva Mendonça

CO-SUPERVISOR
Amâncio Lucas de Sousa Pereira

Faculdade de Ciências Exatas e da Engenharia

Master of Science Degree in Informatics Engineering

Data-Driven Methodology for Estimation of Cruise Ships and Port Energy Profiles in Small Data Settings

Luis Alberto Gonçalves Nunes

Supervised by:

Dr. Fábio Ruben Silva Mendonça

Dr. Amâncio Lucas de Sousa Pereira

Constituição do júri de provas públicas:

Dra. Karolina Baras, Presidente

Dr. Hugo Gabriel Valente Morais, Vogal

Dr. Fábio Ruben Silva Mendonça, Vogal

November 2025

Resumo

Os portos são essenciais para o comércio global, mas enfrentam uma pressão crescente para cumprir objetivos de sustentabilidade, em particular para reduzir as emissões de gases com efeito de estufa. Uma parte importante destas emissões provém dos navios de cruzeiro, que necessitam de grandes quantidades de energia enquanto atracados para manter serviços como HVAC, iluminação e instalações para os passageiros. Esta necessidade de energia não propulsiva, conhecida como carga hoteleira, depende tradicionalmente de geradores a bordo a combustíveis fósseis. Consequentemente, estimar com precisão esta necessidade energética é um passo crítico para a descarbonização portuária e a transição para a eletrificação em larga escala.

Esta tese apresenta uma ferramenta baseada em modelos de machine learning para gerar estimativas de consumo energético, incluindo intervalos de confiança, para navios de cruzeiro enquanto atracados. Foram aplicadas e avaliadas várias técnicas de regressão, desde modelos lineares e não lineares até abordagens probabilísticas, comparadas através de métricas de desempenho como o erro quadrático médio e o coeficiente de determinação.

Os resultados mostram a viabilidade do uso de machine learning para prever a carga hoteleira, apoiando o planeamento energético e estratégias de eletrificação como a shore side electricity. O Porto do Funchal, na Ilha da Madeira, é o caso de estudo utilizado para validar esta metodologia.

Keywords: Shore Side Electricity, Navios de Cruzeiro, Machine Learning, Previsão de Consumo de Energia, Modelos de Regressão

Abstract

Ports are essential to global trade but face increasing pressure to meet sustainability goals, particularly to reduce greenhouse gas emissions. An important portion of these emissions comes from cruise ships, which require substantial amounts of energy while at berth to sustain services such as HVAC, lighting, and guest facilities. This non-propulsive power requirement, known as the hotel load, traditionally relies on onboard fossil-fuel generators. Consequently, accurately estimating this demand is a critical step for port decarbonization and the transition toward large-scale electrification.

This thesis presents a tool based on machine learning models designed to generate energy consumption estimates, including confidence intervals, for cruise ships while docked. Specifically, various regression techniques were applied and evaluated, ranging from linear and nonlinear models to probabilistic approaches. These models were compared using performance metrics, such as mean squared error and the coefficient of determination, to identify the most accurate methods for predicting consumption.

The results show the feasibility of using machine learning to forecast hotel load demand, providing a basis for energy planning and supporting electrification strategies like shore side electricity. The Port of Funchal in Madeira Island serves as the case study for validating this methodology.

Keywords: Shore Side Electricity, Cruise Ships, Machine Learning, Energy Forecasting, Regression Models

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my thesis supervisors, Dr. Lucas Pereira and Dr. Fábio Mendonça, for their patience, constant guidance, and invaluable support throughout this process. Their mentorship was essential not only for the development of this work but also for my professional and academic growth. I would also like to thank Dr. Sheikh Mostafa, who actively participated in several meetings during the preparation of this research, offering valuable advices that contributed to the quality of the work.

My special thanks go to my parents, João and Marisol, whose unconditional support was crucial in difficult moments, and to my family, loved ones, and friends.

I am grateful to the Universidade da Madeira and the Faculdade de Ciências Exatas e da Engenharia, as well as to all the professors who took part in my academic journey, for providing the foundation of my scientific and professional training.

This research is supported by the EU's Horizon Europe R&I programme under grant agreement no. 101160665, project AHEAD, and grant agreement no. 101136131, project Shift2DC. The views and opinions expressed in this thesis are my sole responsibility and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

Table of Contents

List of Figures

List of Tables

List of Acronyms

AIS	Automatic Identification System
ANN	Artificial Neural Network
APRAM	Administração dos Portos da Região Autónoma da Madeira
BESS	Battery Energy Storage System
CO_2	Carbon dioxide
CWC	Coverage Width-based Criterion
GT	Gross Tonnage
HL	Hotel Load
HVAC	Heating, Ventilation, and Air Conditioning
HVSC	High Voltage Shore Connection
IMO	International Maritime Organization
KL	Kullback-Leibler
LOA	Length Overall
MAE	Mean Absolute Error
MAPE	Mean Average Percentage Error
ML	Machine Learning
MLR	Multiple Linear Regression
MPC	Maximum Passenger Capacity
NMPIW	Normalized Mean Prediction Interval Width
NO_X	Nitrogen oxides
OLS	Ordinary Least Squared
PICP	Prediction Interval Coverage Probability
PM	Particulate Matter

PV	Photovoltaic
R^2	Coefficient of Determination
RE	Renewable Energy
RMSE	Root Mean Square Error
SO_X	Sulfur oxides
SSE	Shore Side Electricity
UNFCCC	United Nations Framework Convention on Climate Change

1 Introduction

Ports are essential parts of the global logistics network, as they serve as strategic points for the movement of goods. In addition to their logistical role, they are also key to local and regional economic development by promoting employment, tourism, and industrial growth. It is estimated that more than 80% of the volume of international trade in goods is transported by sea, with this percentage being even higher in many developing countries [1].

Alongside global trade, the cruise tourism industry has experienced explosive growth, becoming a cornerstone of the tourism economy for many coastal and island regions like Madeira. Before the global pandemic, the number of ocean cruise passengers in Europe grew from over 6 million in 2012 to more than 8 million by 2019, with projections indicating a continued upward trend. [2]. This growth has led to the construction of larger 'mega ships', which, while bringing undeniable economic benefits, also present significant environmental challenges, as projections indicate that overall Carbon dioxide (CO_2) emissions from shipping could increase substantially by 2050 if no effective mitigation measures are implemented [3].

To address this issue, the Paris Agreement, adopted at the 21st Conference of the Parties to the United Nations Framework Convention on Climate Change (UNFCCC) in Paris in 2015, set the overarching goal of limiting the increase in global average temperature well below 2,°C above pre industrial levels, and pursuing efforts to limit it to 1.5,°C. This highlighted the urgent need to accelerate the transition towards cleaner technologies across all sectors, including maritime transport [4]. Following this, the International Maritime Organization (IMO) has established specific measures to reduce the carbon intensity of international shipping by at least 40% by 2030 compared to 2008 levels and to reach net-zero greenhouse gas emissions around 2050. Currently, emissions from maritime transport account for approximately 2.9% of global CO_2 emissions, which corresponds to a substantial volume of greenhouse gas emissions [5].

While the entire maritime sector is under review, cruise ships pose a unique challenge at the local level. Due to their massive hotel operations, their energy consumption and the resulting emissions in the berth can be equivalent to that of a small town, severely affecting the air quality of the cities they visit [6].

One of the most promising strategies to reduce local air pollution and greenhouse gas emissions from ships at berth is the implementation of Shore Side Electricity (SSE), also known as cold ironing [7]. This set of technologies allows vessels to shut down their auxiliary engines while docked and connect to the onshore power grid to meet their energy needs. This significantly reduces emissions of CO_2, Nitrogen oxides (NO_X), Sulfur oxides (SO_X), and Particulate Matter (PM) in the port area, leading to improved air quality and contributing to broader climate goals [8]. Additionally, SSE helps reduce noise pollution and supports environmental sustainability in densely populated coastal regions.

The architecture of this type os systems can vary depending on the port and the specifications of the vessels berthed, but a common configuration is illustrated in Figure 1. In this setup, electricity from the onshore utility grid is first stepped down in voltage and then passed through a frequency converter to match the ship's requirements (from 50 Hz to 60 Hz). The voltage is then stepped up again for efficient transmission to the berth, where a specialized high voltage cable connects to the ship's shore connection system. Once onboard, the power is adjusted to the vessel's internal voltage level and distributed to its systems, allowing the diesel generators to be safely turned off while at berth.

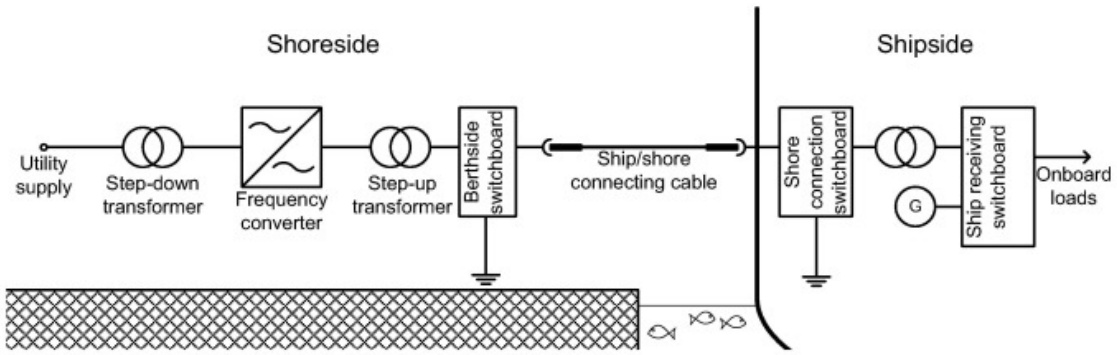

Figure. 1: Typical SSE system from shore side to the seaside [9].

To fully benefit from the advantages of this technology, ports must ensure that the shore side power infrastructure is properly dimensioned and capable of responding to the energy demand of vessels [10]. This requires an accurate estimation of the energy needs of ships during docking periods. This is especially true for cruise ships, where the vast majority of energy consumption at berth is dedicated to their Hotel Power, the energy required to keep all onboard services running, such as lighting, Heating, Ventilation, and Air Conditioning (HVAC), kitchens, elevators, entertainment systems, and others [11]. Accurately forecasting this highly variable load is the primary challenge

to avoid inaccurate sizing of the electrical systems. By anticipating the energy requirements of incoming ships, port authorities can plan energy distribution more efficiently, ensure grid stability, and reduce unnecessary greenhouse gas emissions [12].

However, a considerable obstacle in this process is the limited availability of real consumption data, which constrains the capacity to design, validate, and optimize SSE solutions. In this context, data-driven predictive models provide a promising alternative, using ship specific physical specifications to simulate realistic energy consumption profiles. These models enable scalable and automated forecasting of shore side power demand, adjusted to the operational dynamics of each port [13,14].

1.1 Motivation

This thesis is developed within the scope of a research initiative aimed at improving the sustainability and energy efficiency of port operations through predictive modeling. Focusing specifically on the Port of Funchal, on Madeira Island, it proposes the development of Machine Learning (ML) models to estimate the energy demand of cruise ships during docking periods, a key step in supporting the transition to SSE. Providing a concrete case study that directly addresses the technological challenges and opportunities identified by Nunes et al. [15].

In this context, the motivation of the work is directly aligned with the European projects Shift2DC [16], and AHEAD [17] , both of which aim to advance energy efficiency through innovative electrification solutions. While Shift2DC accelerates the adoption of DC technologies by testing their technical, economic, and environmental feasibility, AHEAD develops AI-based tools to plan and optimize charging infrastructure for electric vehicles and vessels, introducing smart charging algorithms to reduce grid impact and improve operational efficiency.

In this way, both projects are closely related and directly involved with the case of the Port of Funchal, which serves as a strategic demonstrator for sustainable port electrification, benefiting from the gap that this thesis seeks to address: the lack of reliable data on vessels' energy consumption while docked. By developing predictive models capable of estimating the Hotel Load (HL) demand of cruise ships, this work contributes to the generation of realistic data that can serve as a basis for the planning, validation, and optimization of electrification scenarios in the port.

1.2 Objectives

The main goal of this thesis is to develop methods for generating realistic energy consumption profiles of cruise ships while docked, using ML techniques. These profiles are then employed to conduct a comprehensive study of the SSE demand at the Port of Funchal, providing valuable insights to support the planning and sustainable management of port operations. The specific objectives of this study are:

O1 Review the state-of-the-art in SSE systems and existing approaches for estimating hotel power consumption in cruise ships, with particular attention to data-driven and ML techniques that could be applied to the case of the Port of Funchal.

O2 Develop and evaluate ML models to estimate the HL demand of cruise ships during docking periods, using both historical consumption profiles and physical specifications of the vessels, and compare their performance in terms of accuracy and uncertainty.

O3 Simulate the Port of Funchal over one year of docking operations, combining generated consumption profiles with docking records to estimate SSE demand under spatial and capacity constraints.

1.3 Research Questions

Based on the objectives defined for this work, the following research questions are proposed:

RQ1 To what extent do the different ML models provide consistent estimates of the Port of Funchal's SSE demand?

RQ2 Which model demonstrates to be most suitable according to the evaluation metrics defined?

RQ3 How can the results obtained be used to improve infrastructure planning and the port's energy efficiency?

1.4 Thesis Outline

This thesis is structured into six chapters, which are organized as follows: The current Chapter 1 introduces the study by presenting its motivation, objectives, and research questions, while framing the importance of port electrification and the role of ML in energy forecasting. Chapter 2 reviews the literature on port electrification and SSE, addressing its concept and historical development,

its environmental and socio-economic effects, as well as recent applications of ML in this field. It also explores works focused on ships' energy consumption, with particular attention to cruise ships and the forecasting of HL demand.

Chapter 3 describes the methodological development of the study, detailing the workflow followed and explaining the different stages carried out to obtain energy consumption estimates at the port level, while also including the mathematical formulations and theoretical foundations of the regression and curve fitting models applied. Chapter 4 presents the case study at the Port of Funchal, describing the local context and the characteristics of the port, as well as the process of applying the proposed methodology. In addition, it analyzes port energy consumption based on the collected data, integrating the estimation of cruise ship HL with the total port energy demand.

Chapter 5 presents the results obtained by the implemented models and analyses their behaviour based on the defined evaluation metrics. It also analyses the outcomes of the port-wide SSE demand simulation in energy terms. The chapter concludes by proposing a set of techniques and operational measures grounded in these results to improve the port's energy efficiency. Finally, Chapter 6 concludes the thesis by examining the overall work, revisiting the objectives and research questions, discussing the main findings in light of the study's limitations, and outlining directions for future work.

2 Literature Review

This chapter presents a review of the current state of the literature, covering the main concepts and advances related to port electrification, with a particular focus on the cruise industry. It also examines various studies on the prediction of energy consumption in ships, with special emphasis on the Hotel Power demand of cruise ships and the usage of ML techniques. In this way, it provides the context and the scientific background for the development of this thesis.

2.1 Electrification in Ports

Before exploring the concept of SSE, it is important to recognize the broader efforts made to improve energy efficiency in ports. In recent years, ports around the world have adopted a range of technologies (electrified equipment, renewable energy sources, and smart energy management systems) to reduce fuel consumption and lower emissions [18–21]. These advancements not only support environmental goals but also help optimize daily operations. Among them, SSE has emerged as one of the key solutions for making port activities cleaner and more sustainable.

2.1.1 SSE Concept and History

The concept of supplying SSE to vessels while docked has been explored since the 1980s, initially through low voltage systems designed for ferries [22]. As maritime traffic and ship energy requirements grew, high voltage systems became necessary to power large commercial ships, including cruise ships, tankers, and cargo ships.

One of the main pioneers in the implementation of SSE systems was the Port of Los Angeles. In 2004, it became the first port in the world to install a SSE supply for container ships through its West Basin Container Terminal. That same year, it welcomed the NYK Atlas, the first container ship specifically designed to connect to this type of infrastructure while docked [23]. Despite the various technical and operational challenges at the time, including compatibility issues, high investment costs, and the lack of standardized procedures, the port led the way in demonstrating that such systems were both feasible and beneficial [24].

These early difficulties highlighted the need for harmonized technical guidelines, which led to the development of international standards for shore power systems. The Port of Los Angeles played a key role in this process, actively contributing to the creation of the IEC/ISO/IEEE 80005-1

standard, first published in 2012 [25]. Its facilities were among the first to comply with this international reference for High Voltage Shore Connection (HVSC) systems [26]. This standard was later revised and consolidated as IEC/IEEE 80005-1:2019, the current version in use, which establishes a comprehensive framework for ensuring equipment compatibility, earthing configurations, safety measures, and standardized operational procedures in HVSC system implementations [27].

Building on this foundation, Prenc et al. [28] provide a detailed technical analysis of various HVSC network configurations and their integration into national power grids. Focusing on Croatian ports, their study explores centralized, distributed, and single connection models, each evaluated according to the IEC/ISO/IEEE 80005-1 standard. Their contribution serves as a practical guide for port authorities and planners aiming to develop compliant and efficient shore power infrastructure.

2.1.2 SSE Environmental Effects

One of the main motivations for the shift to SSE was the need to reduce greenhouse gas and pollutant emissions, especially in densely populated coastal regions [29]. Several studies have demonstrated the environmental benefits of SSE in mitigating air pollution generated from berthed ships. Among them, Kotrikla et al. [30] explore a small port on a Greek island during the peak of the tourist season, a time when passenger ships are constantly arriving and local air quality is especially vulnerable. Their findings showed that these passenger vessels were the main source of emissions, given their high energy use while at berth. By implementing SSE alongside renewable energy sources such as wind turbines and solar panels, the port was able to reduce CO_2 emissions by up to 96%, while also cutting harmful pollutants such as NO_X, SO_X, and PM almost entirely. Beyond the clear environmental gains, the study illustrates how combining SSE with clean energy can make port operations not only more sustainable but also more compatible with sensitive tourism-based economies.

In parallel, Stolz et al. [31] present an analysis of the emission reduction potential of SSE across European ports. Using Automatic Identification System (AIS) and Monitoring, Reporting and Verification data from over 10,000 ships and 714 ports, they estimated that up to 3 Mt of CO_2 emissions could be avoided annually if SSE were supplied by national grids, and up to 5 Mt if powered by CO_2 neutral sources. The study also highlights that passenger ships account for the highest auxiliary energy demand, offering the greatest potential for CO_2 reduction up to 10.2% of their total emissions. They also found out that the implementation of SSE could eliminate

significant amounts of other pollutants, particularly in major ports like Rotterdam, Hamburg, and Antwerp.

Expanding the environmental case for SSE, recent studies are now investigating sophisticated strategies for specialized vessels. The work of Pérez et al. [32], for instance, assesses an innovative bidirectional SSE framework in a Chilean port focusing tanker ships. Departing from the traditional unidirectional power flow, their approach enables a two-way energy exchange, allowing the vessel's auxiliary engines to maintain operation within their optimal efficiency range, thereby curtailing pollutant generation. When coupled with cleaner fuels like liquefied natural gas, this advanced control system achieves a marked reduction in emissions. This demonstrates that the potential of SSE is evolving beyond simple power substitution towards intelligent energy management to maximize environmental benefits.

2.1.3 SSE Socio-Economics Effects

In addition to its environmental advantages, the implementation of SSE also brings important economic implications for both ports and ship operators [33]. Despite the considerable initial investment in infrastructure, various studies have explored how these costs can be balanced by long-term savings, operational efficiency, and regulatory compliance. The economic dimension of SSE adoption is therefore a main factor in shaping its feasibility and wider deployment, especially in regions with growing environmental requirements.

For instance, Ballini et al. [34] highlights that implementing SSE infrastructure in the Port of Copenhagen is not only an environmental solution but also a financially viable investment. Through a cost-benefit analysis based on real cruise ship traffic data during the 2012 high season, the authors estimate that supplying SSE to 60% of the vessels could lead to annual external health cost savings of approximately €2.8 million. Although the total investment required for the infrastructure is around €37 million, the study concludes that these costs could be recovered within 12 to 13 years due to the significant reduction in harmful emissions. These findings highlight the strategic value of adopting sustainable technologies that simultaneously address environmental concerns and reduce societal costs linked to port-related air pollution.

Similarly, Gore et al. [35] perform a cost-benefit analysis of SSE in Dublin and Belfast. Their findings suggest that targeting the top 10 most frequent passenger callers yields a 10-year Net Present Value of €34.06 million and €15.44 million, respectively. Despite high upfront costs, approximately

€1.5 million per berth and €337,000 per vessel, these expenses are outweighed by fuel savings and a substantial reduction in external socio-environmental costs. Furthermore, Ireland's 2030 renewable targets are projected to increase project viability by 50%, requiring only a negligible ticket supplement of €0.03–€0.04 to ensure financial sustainability.

2.1.4 ML Applications in SSE

In recent years, ML techniques have increasingly played a key role in supporting the logistical and operational activities of port operations [36–39]. While traditionally used for tasks like berth allocation and traffic forecasting, these techniques also contribute significantly to the implementation of SSE by forecasting key variables that determine energy demand.

Some studies focus on predicting intermediate variables, such as berthing duration. A relevant example is the work by Abu Bakar et al. [40], which proposed an ML-based approach to forecast ship berthing duration using data from the Port of Aalborg. The authors tested five models, namely, Artificial Neural Network (ANN), Multiple Linear Regression (MLR), Decision Tree, Random Forest, and XGBoost, finding that ANN reached the best performance with a Root Mean Square Error (RMSE) of 3.13 hours. While valuable for berth allocation, this approach requires a subsequent step to translate duration into energy consumption.

In this context, the study by [41] presents one of the most relevant frameworks for forecasting the total electricity demand of a port microgrid, directly supporting the integration of SSE systems. Using real consumption data collected from the Valletta Cruise Port in Malta, the authors compared several ML models, including Random Forest, Least-Squares Boosting, and Gaussian Process Regression, against simpler baselines such as Linear Regression and Naive models. Their analysis showed that Random Forest achieved the highest accuracy, with an RMSE of 1.18 MW and an Mean Average Percentage Error (MAPE) of 7.25%, demonstrating the model's ability to capture the non-linear consumption patterns associated with cruise ship arrivals and departures.

The study also incorporates real-time meteorological and operational variables, such as ambient temperature, wind speed, and ship berthing schedules, in order to account for environmental and operational factors influencing the demand. This level of integration makes it one of the first works to combine data-driven forecasting with SSE energy planning within a real port environment. However, the authors highlight an important limitation. Namely, the dataset is restricted to a

single port and is not publicly available, which limits the reproducibility and generalization of the results. To partially address this limitation, they complemented the missing cruise ship consumption data with estimations obtained from their earlier study [14], which employed a constraint-based methodology considering engine ratings, propulsion type, and installed power to estimate the hotel load under maximum operational conditions. This integration between both studies strengthened the representation of the cruise ship component within the port's total energy demand, although it still relied on confidential ship information and the assumption of a constant hoteling load during berthing.

2.2 Ships Energy Consumption

Forecasting ship energy consumption is one of the main components in the planning and management of port operations, fuel logistics, and the implementation of SSE. The energy demand of a ship is influenced by multiple dynamic factors, including vessel type, Gross Tonnage (GT), Length Overall (LOA), main and auxiliary engine configurations, cargo load, speed, and weather conditions [42–44].

Traditional forecasting methods have relied on empirical formulas and statistical models based on historical consumption data, often provided by ship logs or AIS-based monitoring [45]. These approaches aim to estimate the power demand during both navigation and port stays, which helps in optimizing port resource allocation, grid management, and emissions control.

Reinforcing this trend, recent research has shifted towards applying various ML models to predict energy consumption. The comprehensive work by Uyanık et al. [46] exemplifies this advance, evaluating nine distinct ML algorithms on a container vessel using real voyage data in addition to AIS Data. Their findings revealed that model performance varies by task: Deep Neural Networks proved most accurate for power prediction, while MLR was superior for fuel consumption. This study is a strong benchmark for forecasting propulsion loads, but it also highlights a critical distinction for cruise ships, where the primary challenge is forecasting the highly variable HL.

2.2.1 Cruises Ships Energy Consumption

Cruise ships stand out among all types of vessels due to their exceptionally high auxiliary energy demand. Unlike cargo ships, where propulsion typically dominates overall consumption, cruise ships

maintain a substantial and continuous energy load even when docked. This is primarily driven by hotel services, necessary to ensure passenger comfort onboard [11].

The main components that constitute these extensive hotel systems are represented in Figure 2. They encompass a wide range of onboard services, from energy intensive operations like Chiller & HVAC Systems, Food Storage & Preparation, and Galleys & Entertainment Facilities, to essential utilities such as Lighting and Elevator Systems. The combined and continuous operation of these diverse systems is what drives the uniquely high and variable HL of a cruise ship.

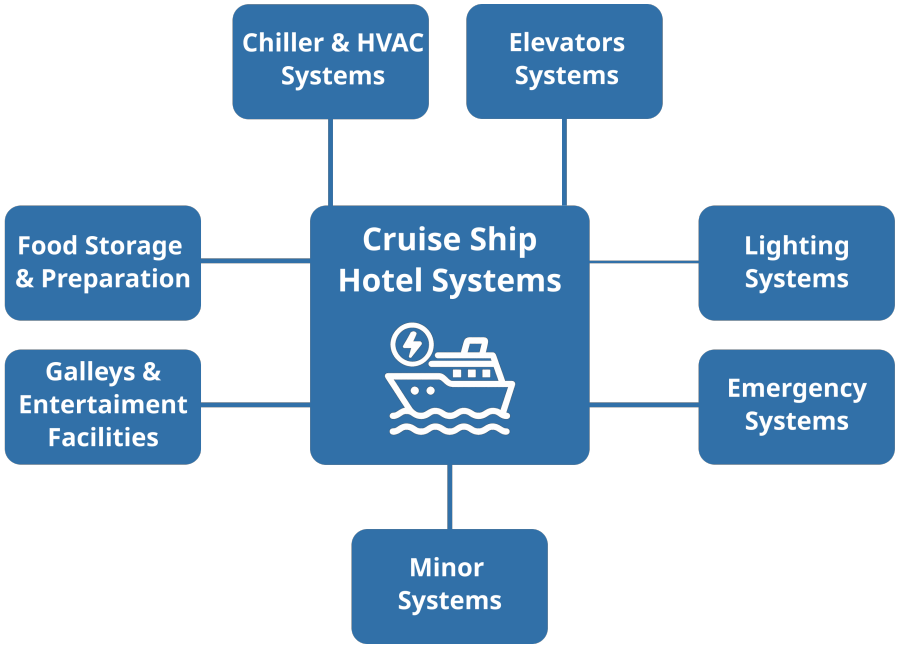

Figure. 2: Main components of cruise ship hotel systems.

Among these, the HVAC systems represent the single largest consumer of electrical power within the HL. Early dynamic modeling studies highlighted this, estimating that HVAC systems alone can be responsible for roughly one third of a cruise liner's total power consumption, an amount comparable to that of the main propulsion system [47].

Given their complexity, these systems are responsible for maintaining a comfortable climate across vast interior spaces, and their energy consumption is highly dynamic. It is sensitive to a range of external factors, such as ambient temperature and solar radiation, as well as internal operational factors like the number of passengers onboard and time of day. This sensitivity makes the HVAC system not only a primary driver of the high baseline load but also a major source of its variability.

The challenge of modeling these complex hotel systems is a current focus of research. For instance, the recent work by Brækken et al. [48] developed a detailed dynamic model to study how hotel systems on a cruise ship consume energy under the real operating conditions of a Nordic climate. Their study underscores the fluctuating nature of the demands for HVAC, hot water, and lighting by accounting for factors like passenger behavior and operational mode (at sea or in port), reinforcing the need for dynamic, rather than static, analysis.

2.2.2 Forecasting Cruise Ships Hotel Load

Accurate forecasting of the HL of cruise ships is, as a result, highly relevance for both on board energy system design and land-based infrastructure planning, especially in ports implementing shore power solutions. Unlike propulsion related demand, HL is highly sensitive to operational schedules, passenger activities, climate conditions, and ship design specifications. Among the most influential factors are the ship's physical attributes, particularly GT, LOA, and Maximum Passenger Capacity (MPC), which are now widely used in forecasting models due to their strong correlation with on board energy demand [13, 14].

Early attempts to model HL were often secondary to broader environmental impact studies. A notable example is the work of Battistelli et al. [12], who, as part of an air quality study in the Port of Naples, attempted a simple linear regression to correlate a ship's LOA with its winter load. The model's poor performance, represented by an Coefficient of Determination (R^2) of 0.20, highlighted a main challenge, as predicting HL is a complex problem, and simple physical characteristics alone are insufficient, thus establishing the need for more robust, data-driven approaches.

This challenge was directly addressed by Espinosa et al. [13], in what became a foundational step for data-driven estimation of cruise ship HL. They systematically developed and compared regression-based models (linear, logarithmic, and potential) to estimate the maximum hotelling power demand of cruise ships using the key physical specifications (GT, LOA, and MPC). Their analysis was based on empirical scheduling data from the Port of Barcelona combined with fixed hotel load values reported by Ericsson and Fazlagić [49]. While this approach enabled a practical estimation of port-wide SSE demand profiles, it relied entirely on static assumptions of onboard power usage rather than on real time or time series measurements. Consequently, the HL was treated as a constant representative value, typically the average or peak load during hotelling,

ignoring short term operational variations caused by passenger activity, climatic conditions, or onboard systems. This constraint simplified model development and allowed the authors to assess SSE infrastructure requirements for the port, but limited the model's applicability for dynamic analyses or time varying demand estimation, highlighting the need for better modeling frameworks in subsequent research.

On the other hand, as mentioned in Micallef et al. (2025) [41], the pursuit of higher accuracy in static load estimations was previously advanced by Micallef et al. (2024) [14], who introduced a more sophisticated hybrid ML methodology specifically designed to estimate the hotel load of cruise ships under maximum operational conditions. Although their approach remained focused on static estimation, it refined the process through a two-stage framework that combined a constraint-based algorithm, incorporating both physical ship specifications and parameters related to propulsion and engine engineering, with Gaussian Process Regression capable of capturing non-linear relationships. This integration enabled a more realistic representation of the physical dependencies between ship design characteristics and energy demand, overcoming the simplifications and limited variable scope that characterized earlier regression-based methods such as those by Espinosa et al. [13].

The performance of the proposed models was assessed using standard regression metrics (R^2 and RMSE). The linear regression model achieved an R^2 of 0.6949 and an RMSE of 3.06 MW, while the Gaussian Process Regression model improved these results to an R^2 of 0.9767 and an RMSE of 0.99 MW, confirming the enhanced predictive capability of the hybrid approach. As a result, the model achieved a substantial improvement in the accuracy of static peak load estimation and provided a stronger foundation for subsequent port energy demand forecasting studies, including its complementary use in Micallef et al. (2025) [41].

However, this literature review also highlights two major limitations that remain unresolved in the field. The first limitation concerns the lack of publicly available and detailed datasets on ship energy consumption. Most of the studies reviewed, including those by Espinosa et al. [13] and Micallef et al. [14,41], rely on proprietary or aggregated data collected from individual ports, often complemented by fixed assumptions or indirect estimations of hotel load. As a result, reproducibility and cross validation of models remain highly constrained. Furthermore, existing datasets typically provide only static or averaged consumption values rather than high resolution time series, which hinders the validation of approaches that aim to capture the dynamic variability of onboard power demand.

The second limitation is the absence of established methodologies capable of estimating dynamic HL profiles for cruise ships. Although several studies have made notable progress in static peak load estimation, none have yet addressed the temporal behavior of energy consumption during port stays. This represents a substantial research gap, as cruise ships exhibit highly variable power demand patterns influenced by passenger activity, onboard systems, and environmental conditions. Moreover, across the studies reviewed, there is no consensus on the methodological framework or performance evaluation standards. Metrics such as R^2, RMSE, and MAPE are not always reported or applied consistently, which limits the comparability, validation, and reproducibility of published results.

To address these limitations, this thesis proposes the development of a two-stage methodology for generating regression-based energy consumption profiles. The proposed approach first estimates the HL based on ship specifications and then applies regression-based curve fitting to real world time series consumption data, resulting in a model that functions as a dynamic generator of energy consumption profiles. Unlike previous studies, which rely on fixed or static values of hotel load, this methodology incorporates the temporal variability of actual hotel power consumption, enabling more realistic and adaptive estimations of onboard energy demand. This tool is designed to support the planning of shore side electricity infrastructure and strategic decision making in ports with high volumes of cruise traffic.

3 Methodological Development

This chapter details the methodological framework designed specifically to address the gaps identified in the literature review. The primary objective is to move beyond simple static predictions and develop a robust pipeline for the generation of complete energy consumption profiles.

This chapter will be organized as follows. First, a methodological pipeline is presented, outlining the systematic workflow from data collection and preprocessing to the two-stage model development and its final evaluation. Following this, the Regression and Modeling section explores the technical specifics of this approach, justifying the selection of the specific regression models used. Finally, the Port Estimation section explains how the validated model is used to conduct practical analysis, including scenario simulation and the interpretation of confidence intervals for strategic decision making at the Port of Funchal.

3.1 Methodological Pipeline

The methodological pipeline developed in this thesis was designed to estimate port energy consumption using different classes of data. These include the scheduling of ship arrivals and departures, the physical specifications of the vessels, and time series data collected from control ships.

This pipeline adopts a dynamic and adaptable approach, making it applicable to different ports, as long as the three mentioned data sources are available. Its modular structure also facilitates the reuse and progressive updating of the model, allowing for easy adjustments as new records are added or operational conditions at the port evolve. Furthermore, its performance can continuously improve with the inclusion of additional historical data, improving its predictive power and its usefulness for long term decision making.

The entire process was divided into five sequential stages, where each stage serves as the input for the next. The defined stages are Data Preparation, Data Preprocessing and Resampling, Two-Step Regression Model Component, Model Evaluation, and Port level Energy Demand Estimation. Figure 3 presents a schematic representation of the interconnection between the different stages of the pipeline.

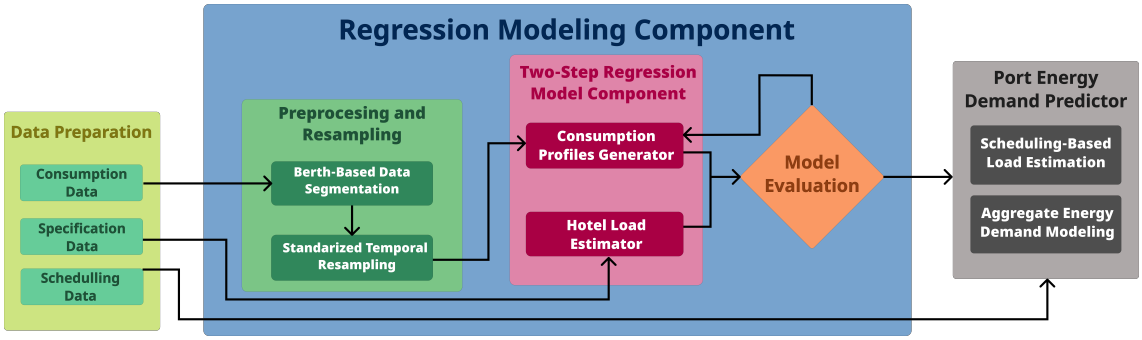

Figure. 3: Overview of the Methodological Pipeline.

3.1.1 Data Preparation

The first stage of this pipeline, as its name suggests, involves the preparation and gathering of all necessary data required for its proper functioning. The main types of data used were Port Call Schedules, Vessel Specifications, and Energy Consumption Time Series.

The port call schedules are essential to obtain both historical data and future scheduling information for the various vessels calling at the port. The most important data points to collect are the ship name, along with the exact date and time (HH:MM:SS) of arrival and departure.

Once the vessels calling at the port have been identified, it is necessary to collect their physical and technical specifications. Among the different parameters that can be considered, the most relevant and readily available in practice are the GT, LOA, and MPC. These variables are widely recognized as strong predictors of HL, as discussed in the literature review (Section 2) and demonstrated in previous works such as [13] and [14], where they were successfully employed to model the energy requirements of cruise vessels with high predictive accuracy.

Energy consumption data, expressed as time series in kW, were used for a subset of reference vessels. These vessels were not selected based on specific criteria but correspond to those for which energy consumption data were available. This information is critical for constructing hotel consumption profiles and developing the regression curve fitting models developed in this thesis.

3.1.2 Data Preprocessing and Resampling

Once collected, the raw data required a methodical preprocessing phase to be structured and cleaned for use in the modeling pipeline. The initial steps focused on organizing the port scheduling information and the technical specifications of the vessels.

The most important and complex aspect of this preprocessing stage involved handling the time series energy consumption data. The primary objective was to accurately isolate the periods during which the reference vessels were berthed at the port. This step was essential to ensure that the model was trained exclusively on relevant HL data, excluding any energy consumption associated with maneuvering or navigation, but also to align the scope of the data directly with the practical application of SSE, which is only utilized during the berthing phase.

To facilitate further processing, the data corresponding to these berthing periods were extracted and saved in separate files, referred to as *"Slices"*. The extracted variables were limited to Hotel Power and Chillers Power, as these constitute the primary components of the vessel's HL profile and were the only consistently available measurements across all ships in the dataset, ensuring direct comparability between vessels and strong relevance for the modeling stage.

A challenge identified during the preprocessing phase was that the original datasets for each vessel were recorded at different sampling rates. To address this, a resampling procedure was applied to standardize all records to a common resolution of 1 minute. This step ensured that, while each slice preserved its actual temporal duration, all data shared the same temporal resolution, allowing for direct comparison in the subsequent modeling stages.

3.1.3 Two-Stage Model Development

The main contribution of this thesis lies in the development of a two-stage model, divided into two core components. The first stage estimates the HL of cruise ships based on their physical and technical specifications, while the second generates energy consumption profiles using regression-based curve fitting. The integration of both stages allows the prediction of realistic energy consumption patterns for ships without historical data and the reconstruction of their temporal behavior during berthing periods.

The first stage of the model (HL estimation through Regression) is designed to predict a static baseline HL value (in kW) for any given cruise ship, using only its physical and technical specifications. As mentioned earlier, the input features selected for this stage are the GT, LOA, and MPC, which have been consistently employed in previous studies as reliable predictors of HL. The rational is that by relying on these key attributes, the model ensures interpretability and applicability to vessels without historical consumption data available.

The second stage of the model, which is the curve fitting for consumption profiles that focuses on generating dynamic energy consumption profiles for new vessels by leveraging the subset of reference ships. This is achieved through a regression-based curve fitting approach combined with a best-fit prototype comparison strategy. The process is structured as follows.

First, each available consumption slice for the reference ships is individually processed, and the curve fitting procedure is applied separately to each one. This approach allows capturing the specific variability of different berthing periods while preserving their original characteristics.

Next, different curve fitting regression models are trained on each slice, aiming to approximate its energy profile with the most suitable functional form. During this step, various regression techniques and validation strategies are tested to evaluate the suitability of the method and to select the most robust configuration for each case. Additionally, uncertainty intervals are derived to quantify the variability around the fitted curves, providing statistically robust upper and lower bounds for every time step. Depending on the model, these take the form of prediction intervals or confidence intervals. This analysis is discussed in further detail in Section 3.2.2.

3.1.4 Model Evaluation

This stage was designed to assess the models developed across the different phases. The evaluation relies on a set of metrics, computed under an 80/20 random train–test split per slice, that cover three perspectives, namely, point prediction quality, uncertainty quality, and the consistency of the port's aggregated consumption.

First, the quality of point predictions is considered as the model's ability to approximate observed values without resorting to uncertainty intervals. RMSE and Mean Absolute Error (MAE) are used as error magnitudes, where the former is more sensitive to large deviations and the latter is more robust to average errors, and R^2 is used as a measure of explained variance. In HL estimation, these metrics are reported with descriptive intent to illustrate each model's effectiveness and are not used as a selection criterion. For the curve fitting stage, the empirical Kullback-Leibler (KL) Divergence is additionally included to capture potential shape mismatches between the distributions of predictions and observations (for example, tails or skewness) that may not be fully reflected in average errors.

Next, for curve fitting, the uncertainty quality is evaluated. Prediction Interval Coverage Probability (PICP) measures their calibration, that is, the proportion of observations that fall within the interval, meaning that values close to the nominal level indicate bands consistent with the data. In addition, Normalized Mean Prediction Interval Width (NMPIW) quantifies "simplicity" via the mean width normalized by the dispersion of the target, enabling fair comparisons across scales, as lower values imply more informative bands. Finally, Coverage Width-based Criterion (CWC) combines both dimensions by penalizing lack of coverage relative to width, favoring intervals that are simultaneously well calibrated and narrow.

Lastly, the consistency of the port's aggregated consumption metric was defined to aggregate the shore-side power demand over the entire analysis period. For each model, the predicted power time series is integrated to obtain a single annual energy value, from these annual values, the mean across models is computed, and the deviation of each model from that mean is quantified. This formulation enables an operational comparison between models, highlighting those that imply a higher or lower total SSE demand and those that are closest to the central estimate.

3.1.5 Port Energy Demand Estimation

The final stage of this pipeline uses the validated two-stage model to estimate the total energy demand at the port level. To generate the consumption profile of each target vessel, its estimated HL is first compared against the set of reference ships. Then, the generalized curve of the closest-matching reference ship is selected, and a scaling factor, defined as

$$\text{Scaling Factor} = \frac{\text{Hotel Load Target Ship}}{\text{Hotel Load Closest Reference Ship}}, \tag{1}$$

is applied to every point in the chosen profile, ensuring that the resulting series reflects the specific energy demand of the target vessel.

Once each vessel has been assigned its scaled consumption profile, these profiles are combined with the port call schedule data, which specifies the berthing periods of each ship. This integration produces a continuous representation of total energy demand at the port, explicitly considering overlapping stays and the operational capacity constraints of the Port of Madeira.

The model's outputs can be analyzed across multiple time scales, enabling scenario simulation and infrastructure planning. At the daily level, it highlights peak demand hours that are critical for operational and energy management. At monthly and yearly scales, it reveals recurring traffic patterns and seasonal variations, while also providing a robust basis for estimating annual energy consumption, essential for financial planning and sustainability assessments.

3.2 Regression and Modeling

The development of this two-stage model relies on the implementation of several regression techniques, each carefully selected to meet the specific requirements of its corresponding phase. Every regression technique was chosen based on a criterion that balances predictive performance, robustness to data variability, and interpretability of the results. The specific techniques employed and the rationale for their selection are detailed below.

3.2.1 Baseline Hotel Load Estimation via Regression

This stage builds upon the description provided in Section 3.1.3, where the objective of estimating the baseline HL from vessel specifications was introduced. Here, the focus is on detailing the regression techniques implemented to accomplish this task.

For the implementation of this estimation stage, data collected from the study conducted by Ericsson et al. [49] were used. This study is based on an extensive compilation of peak HL values recorded during shore power connections at the Port of San Francisco for various types of vessels, with the present work focusing exclusively on cruise ships. The dataset used in this work contains GT, LOA, MPC, and HL, which serve as the basis for the regression modeling described in this section.

In addition to this dataset, the specifications of a subset of reference cruise ships with known real HL values were incorporated to provide additional reference points and improve model performance. Once the dataset to be used was established, two different approaches were implemented for estimating HL. The first approach follows the methodology described by Espinosa et al. [13], in which, after comparing different types of regression, a potential regression model using GT as the sole explanatory variable was selected due to its high correlation and wide range of applicability.

In this study, the potential model was estimated by applying a log-log transformation to both the dependent variable (ower) and the independent variable (GT), allowing the model to be of the form

$$y = A \cdot x^B, \tag{2}$$

and to be expressed in its linearized form

$$log(y) = \log(A) + B \cdot \log(x), \tag{3}$$

which was then fitted using Ordinary Least Squared (OLS) regression [50].

The second approach, developed and proposed in the present study, is based on Multiple Linear Regression (MLR) using GT, LOA, and MPC as independent variables. This choice is motivated by the strong pairwise Pearson correlations observed between these vessel specifications and the HL values, as illustrated in Figure 4. All three explanatory variables show strong positive correlations with the target variable (HL), with coefficients ranging from 0.78 to 0.84, indicating that each carries relevant predictive information and justifying the adoption of a multivariate approach rather than relying on a single predictor.

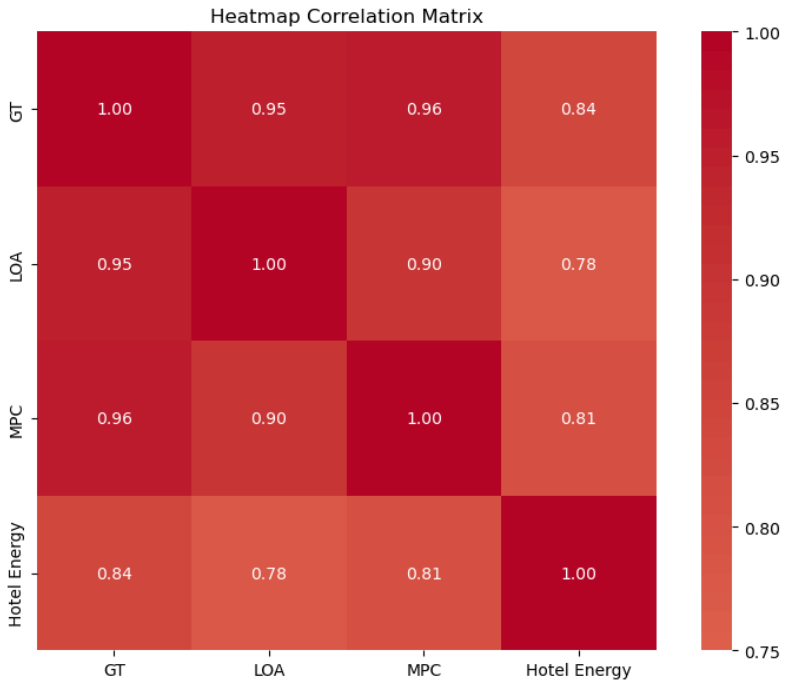

Figure. 4: Correlation Matrix for the Regression dataset.

The general form of the MLR model can be expressed as

$$\hat{y} = \beta_0 + \beta_1 \cdot \mathrm{GT} + \beta_2 \cdot \mathrm{LOA} + \beta_3 \cdot \mathrm{MPC}. \tag{4}$$

In order to estimate the coefficients of the MLR model, the OLS method was again employed. OLS determines the values of the regression coefficients by minimizing the sum of the squared differences between the observed and predicted values, thus providing the best linear unbiased estimators under the classical regression assumptions. In this formulation, $\hat{y}$ denotes the estimated HL (kW), β_0 is the intercept, and β_i are the regression coefficients associated with each explanatory variable.

Despite the proposal to implement multiple linear regression, the potential regression model remains a validated reference due to its previous application and demonstrated robustness. Due to that fact, in this work, only the results obtained from this model are considered, serving as a reliable basis to support complementary analyses and the development of subsequent models.

3.2.2 Curve Fitting for Consumption Profiles

In line with the objective of the second stage of the model, which seeks to generate accurate representations of cruise ship hotel energy consumption patterns, a wide range of regression-based curve fitting techniques was applied.

In more detail, curve fitting, as depicted in Figure 5, is a data analysis technique used to model the underlying relationship between a set of observed data points and a mathematical function. By minimizing the discrepancy between measured and predicted values, it constructs a continuous approximation of the process under study. As a result, it is widely used in engineering and energy modeling, since it can reveal global trends and local variations within a dataset—even when observations are irregular or affected by noise. In this context, the red curve $\hat{y}(t)$ represents the model's mean prediction over time, while the shaded region around it illustrates an uncertainty band conveying the plausible variability of the predictions.

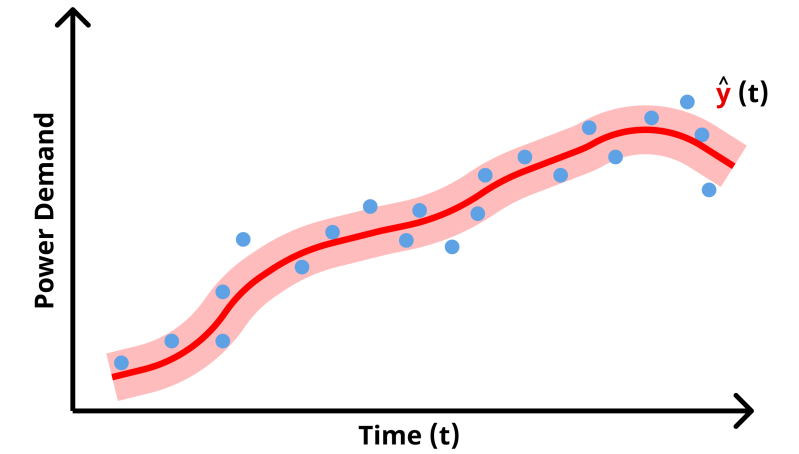

Figure. 5: Schematic of power demand forecasting with confidence intervals. Blue points are real measurements, the red line ($\hat{y}(t)$) is the mean model prediction over time and the shaded band indicates the prediction confidence interval.

In the context of cruise ship hotel energy consumption analysis, curve fitting can then be used to reconstruct realistic consumption trajectories from discrete or incomplete records. Unlike static estimation methods that rely on fixed average or peak values, curve fitting captures the temporal variability of energy use, reflecting operational dynamics during berthing periods, such as fluctuations in HVAC systems, lighting, or passenger activity.

To ensure that predictions are accurate and reliable, each fitted curve is complemented with an estimation of its associated uncertainty, allowing the quantification of prediction confidence and the identification of time intervals where the model exhibits higher or lower reliability. Representing this uncertainty through confidence or prediction bands around the fitted curve provides an intuitive visualization of model robustness, which is particularly valuable for operational decision making in port energy management.

Table 1 summarizes the regression methods used for curve fitting, their model categories, and the corresponding uncertainty estimation approaches. These groups differ in their underlying assumptions, learning mechanisms, and ability to capture the variability of cruise ship energy consumption.

Regression Method	**Model Category**	**Uncertainty Estimation**
Linear Regression	Linear Models	Residual-based Interval
Ridge Regression	Regularized Models	Residual-based Interval
Lasso Regression	Regularized Models	Residual-based Interval
ElasticNet Regression	Regularized Models	Residual-based Interval
Bayesian Ridge Regression	Probabilistic Models	Bayesian Predictive Interval
Support Vector Regression (SVR)	Kernel-based Models	Residual-based Interval
Kernel Regression	Kernel-based Models	Residual-based Interval
Decision Tree	Tree-based Models	Jackknife Empirical Coverage
Random Forest	Tree-based Models	Jackknife Empirical Coverage
K-Nearest Neighbors (KNN)	Lazy Learning Models	Jackknife Empirical Coverage
Quantile Regression	Distribution-aware Models	Quantile Estimates
Spline Regression	Other Models	Residual-based Interval
Isotonic Regression	Other Models	Residual-based Interval

Table 1: Summary of regression methods, model categories, and uncertainty estimation techniques used for curve fitting.

For a subset of the models considered, the uncertainty estimation is explicitly derived from the residual analysis. In these cases, the prediction bands are obtained by treating the residuals as an estimate of the irreducible noise and assuming Gaussian errors. Let $\hat{y}(t)$ denote the fitted curve and $\hat{\sigma}_{\text{res}}$ the standard deviation of the training residuals. The associated prediction bands are then defined as

$$\hat{y}(t) \pm z\,\hat{\sigma}_{\text{res}}, \tag{5}$$

with $z = 1.96$ for a nominal 95% prediction level. This construction provides a simple, model-agnostic way to quantify uncertainty around the fitted curve and to highlight time intervals where deviations are more likely.

Residual analysis was conducted to verify the assumptions underlying this construction. The histogram in Figure 6 shows that the residuals are symmetrically distributed around zero, while the Q–Q plot in Figure 7 indicates that they align closely with the theoretical normal distribution. These diagnostics confirm that the errors behave approximately as Gaussian noise, supporting the use of the standard normal quantile when defining the prediction bands.

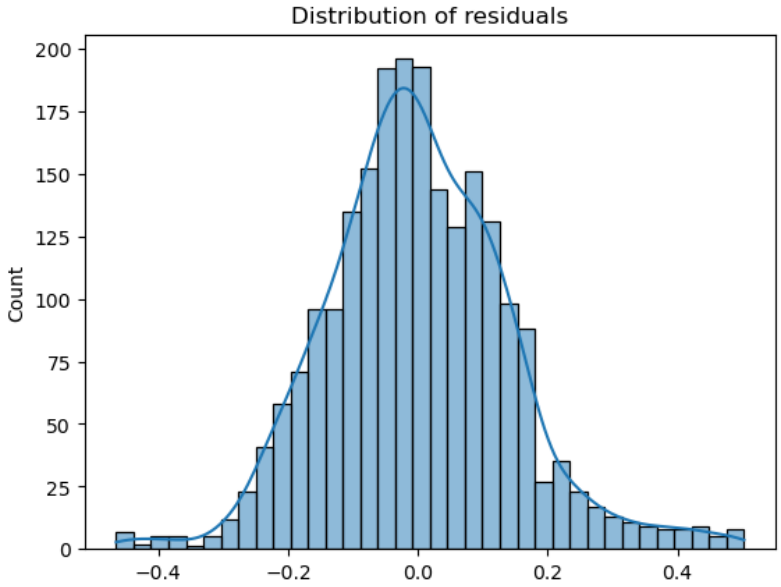

Figure. 6: Histogram and kernel density of residuals showing approximate normality.

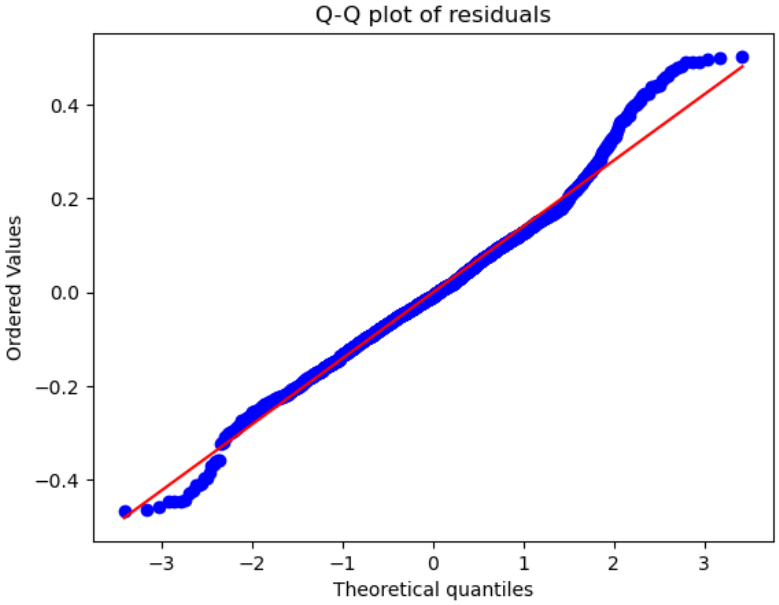

Figure. 7: Q–Q plot of residuals against the theoretical normal distribution.

3.2.2.1 Linear Models

The Linear Regression Model was used as the baseline for assessing the performance of more advanced approaches, given its simplicity, interpretability, and analytical transparency. Despite its restrictive assumptions, it provides a valuable reference for understanding the underlying struc-

ture of cruise ship energy consumption and serves as the foundation for subsequent extensions incorporating regularization, nonlinearity, or probabilistic inference.

Mathematically, linear regression expresses the response variable y as a linear combination of p explanatory variables x_j plus an additive error term

$$\hat{y} = \beta_0 + \beta_1 x_1 + \cdots + \beta_p x_p + \varepsilon, \tag{6}$$

where β_0 is the intercept, β_j are the regression coefficients, and ε represents the residual error. The coefficients were estimated by minimizing the sum of squared residuals

$$\min_{\beta} \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2, \tag{7}$$

which leads to a closed-form solution under the OLS criterion [50].

For this model, the associated uncertainty is quantified using the residual-based approach previously described, where the empirical distribution of the training residuals is used to construct prediction bands around the fitted curve under the assumption of approximately Gaussian errors.

3.2.2.2 Regularized Models

The family of Regularized Linear Models represents a natural extension of standard linear regression, maintaining its structure while introducing additional control over model complexity. The goal is to improve generalization, stabilize coefficient estimates in the presence of correlated predictors, and avoid overfitting when data are limited or noisy.

Formally, regularization augments the least-squares criterion with a penalty $\Omega(\boldsymbol{\beta})$ scaled by a tuning parameter λ as

$$\min_{\beta} \left(\sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2 + \lambda\, \Omega(\boldsymbol{\beta}) \right). \tag{8}$$

Depending on the choice of $\Omega(\boldsymbol{\beta})$, we obtain the three standard estimators. The first one is the Ridge Regression that uses an L_2 penalty that shrinks coefficients proportionally and mitigates multicollinearity, given by

[51]

$$\min_{\beta} \left(\sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \; + \; \lambda \sum_{j=1}^{p} \beta_j^2 \right), \tag{9}$$

where $\lambda > 0$ controls the strength of the shrinkage. The second one is the Lasso Regression that employs an L_1 penalty that promotes sparsity and enables automatic variable selection, using [52]

$$\min_{\beta} \left(\sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \; + \; \lambda \sum_{j=1}^{p} |\beta_j| \right), \tag{10}$$

The last one is the ElasticNet Regression, which combines L_1 and L_2 penalties to balance sparsity and stability, thus [53]

$$\min_{\beta} \left(\sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \; + \; \lambda_1 \sum_{j=1}^{p} |\beta_j| \; + \; \lambda_2 \sum_{j=1}^{p} \beta_j^2 \right), \tag{11}$$

As with the linear model, the uncertainty associated with predictions is quantified from the residual variability. The empirical standard deviation of the residuals represents the average deviation of observed values from the fitted curve. Assuming the residuals follow an approximately normal distribution, confidence and prediction intervals can also be computed as in Equation 5.

3.2.2.3 Probabilistic Models

The Probabilistic Regression Models represent a methodological shift in this study, as they introduced the concept of uncertainty bands directly into the regression process. In contrast to earlier approaches, where uncertainty was estimated from residual analysis, these models inherently quantify the confidence associated with each prediction. This capability proved important for cruise ship energy consumption forecasting, since power demand is highly variable and subject to operational and environmental fluctuations. For SSE planning, having access to both the expected value and its plausible range of variation provides a more robust basis for decision-making and infrastructure sizing.

Within this category, the Bayesian Ridge Regression model was adopted as the representative approach. Instead of producing a single deterministic fitted curve, it estimates a full predictive

distribution for each time step, providing a more comprehensive depiction of model confidence and variability. Mathematically, Bayesian Ridge regression assumes a Gaussian likelihood for the observed targets as [54]

$$p(\mathbf{y} \mid \mathbf{X}, \mathbf{w}, \alpha) = \mathcal{N}(\mathbf{X}\mathbf{w}, \alpha^{-1}\mathbf{I}), \tag{12}$$

and a Gaussian prior over the regression weights, given by

$$p(\mathbf{w} \mid \lambda) = \mathcal{N}(\mathbf{0}, \lambda^{-1}\mathbf{I}), \tag{13}$$

where α represents the precision (inverse variance) of the observation noise, and λ controls the precision of the weight prior. The Bayesian formulation integrates through Bayes' theorem, resulting in a posterior distribution over the model parameters that reflects the degree of certainty in each estimated coefficient.

For a new observation $\mathbf{x}_*$, the predictive distribution is also Gaussian, defined by a mean $\hat{y}$ and a standard deviation $\sigma_{\hat{y}}$. From this distribution, prediction intervals can be derived as [55]

$$\hat{y} \pm z_\gamma \cdot \sigma_{\hat{y}}, \tag{14}$$

where z_γ corresponds to the standard normal quantile ($z_{0.975} = 1.96$ for a 95% confidence level). Thus, this formulation naturally incorporates parameter uncertainty and data noise into the predictions, producing smooth fitted curves surrounded by credible intervals that represent the model's confidence.

3.2.2.4 Kernel-based Models

The Kernel-based Models were introduced to overcome the limitations of linear and regularized regressions, whose inherent linear structure constrains their ability to represent the nonlinear and cyclical dynamics of cruise ship HL demand. Kernel methods address this limitation by implicitly projecting the original input space $\mathbf{x}$ into a higher dimensional feature space $\phi(\mathbf{x})$, where complex nonlinear relationships can be captured through linear estimators. This flexibility makes them particularly suitable for modeling the irregular consumption patterns associated with passenger activity, weather variations, and onboard system operation.

Two kernel-based approaches were applied in this work. The first one was the Support Vector Regression (SVR), which introduces the ϵ-insensitive loss function, ignoring small deviations within a tolerance margin ϵ, preventing minor fluctuations from influencing the fitted curve. The optimization problem is formulated as

$$\min_{\mathbf{w},b,\xi,\xi^*} \frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i=1}^{n}(\xi_i + \xi_i^*), \tag{15}$$

subject to

$$\begin{cases} y_i - (\mathbf{w}^\top \phi(\mathbf{x}_i) + b) \leq \epsilon + \xi_i, \\ (\mathbf{w}^\top \phi(\mathbf{x}_i) + b) - y_i \leq \epsilon + \xi_i^*, \\ \xi_i, \xi_i^* \geq 0, \end{cases} \tag{16}$$

where C controls the trade off between model smoothness and tolerance for large deviations, and γ in the RBF kernel defines the influence radius of each support vector [56].

The second one was Kernel Regression, which is a fully non-parametric approach in which predictions are obtained as weighted averages of neighboring observations. In the Nadaraya–Watson formulation, the estimator is defined as [57, 58]

$$\hat{y}(\mathbf{x}) = \frac{\sum_{i=1}^{n} K\left(\frac{\mathbf{x}-\mathbf{x}_i}{h}\right) y_i}{\sum_{i=1}^{n} K\left(\frac{\mathbf{x}-\mathbf{x}_i}{h}\right)}, \tag{17}$$

where $K(\cdot)$ is the kernel function and h the bandwidth controlling smoothness. It is important to note that, unlike SVR, Kernel Regression does not assume any parametric structure, relying entirely on local similarity to reconstruct smooth consumption trajectories.

Regarding uncertainty quantification, the two kernel-based approaches differ in their methodological treatment. For Kernel Regression, uncertainty estimation follows the same residual-based strategy as in previous models. In contrast, the SVR model employs a non-parametric resampling strategy based on the block Jackknife method [59]. In this approach, the time series is divided into B contiguous segments, and the model is iteratively retrained, each time leaving out one block from the training set. The ensemble of predictions across these resampled models provides an empirical estimate of the predictive dispersion. From this variability, the uncertainty bands are defined as $\hat{y} \pm z\sigma_{\text{jack}}$, where σ_{jack} denotes the standard deviation of predictions across Jackknife replicates. To

ensure methodological consistency and comparability with the uncertainty estimates of the other model families, the same critical threshold $z = 1.96$ (corresponding to a 95% confidence level when the sample is large enough) is adopted.

Unlike residual-based Gaussian methods, this approach does not rely on any assumption about the distribution of errors. Instead, it derives uncertainty directly from the variability observed across the resampled predictions, making it a fully data-driven and distribution free technique. This is particularly useful in cases where the data exhibit irregular fluctuations and changing levels of variability over time, as often occurs in cruise ship energy consumption patterns.

3.2.2.5 Tree-Based Models

The Tree-Based Models, represented in this study by Decision Trees and Random Forest, were selected for their capacity to capture nonlinear interactions between temporal and operational variables that characterize cruise ship HL energy consumption behavior. These models are particularly effective when energy consumption depends simultaneously on multiple interacting factors, such as the operational state of onboard systems (Chillers Power) and time-related patterns like hour of the day or day of the week. Furthermore, unlike parametric regressions, tree-based approaches make no assumptions about the data distribution, relying instead on recursive partitioning of the feature space to minimize prediction error.

A Decision Tree divides the input space into homogeneous regions by successively selecting the splits that maximize the reduction of node impurity. With setting the considered criterion to be the squared error, the impurity in a node t is given by the mean squared error

$$I_t; =; \frac{1}{N_t} \sum_{i \in t} \left(y_i - \bar{y}_t\right)^2, \tag{18}$$

and the corresponding reduction in impurity after a split is expressed as [60]

$$\frac{N_t}{N} \left(I_t - \frac{N_{t_L}}{N_t} I_{t_L} - \frac{N_{t_R}}{N_t} I_{t_R} \right), \tag{19}$$

where N_t is the number of samples in node t, and t_L, t_R denote its left and right child nodes.

The Random Forest Regressor extends this principle by combining multiple trees, each trained on bootstrap samples of the data, and averaging their predictions to improve generalization and reduce variance as

$$\hat{y}(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^{T} f_t(\mathbf{x}), \tag{20}$$

where T is the total number of trees in the ensemble. This aggregation mechanism provides robustness against noise and overfitting, which is especially useful when modeling high variability energy consumption patterns in cruise ships. For uncertainty quantification, both tree-based models employed the same block Jackknife resampling approach previously introduced for the SVR model.

3.2.2.6 Lazy Learning Models

One of the main challenges in this research was the limited amount of available data, which makes Lazy Learning models particularly suitable. Specifically, unlike parametric models that require a training phase to estimate coefficients, lazy learners store all training samples and generate predictions by comparing new inputs to previously observed instances. This property allows them to adapt flexibly to small datasets and local consumption variations without assuming any specific functional form. Within this category, the K-Nearest Neighbors (KNN) algorithm was selected for its simplicity, interpretability, and ability to capture local nonlinear dependencies.

The KNN model predicts the target value of a new observation by averaging the outputs of its k nearest samples in the feature space. Similarity between instances $\mathbf{x}_i$ and $\mathbf{x}_j$ is typically measured using the Euclidean distance, given by [61]

$$d(\mathbf{x}_i, \mathbf{x}_j) = \sqrt{\sum_{m=1}^{M} (x_{im} - x_{jm})^2}, \tag{21}$$

where x_{im} is the value of the m-th feature for instance i. Given a new input $\mathbf{x}$, the model estimates its prediction as

$$\hat{y}(\mathbf{x}) = \frac{1}{k} \sum_{i \in \mathcal{N}_k(\mathbf{x})} y_i, \tag{22}$$

where $\mathcal{N}_k(\mathbf{x})$ represents the set of the k nearest neighbors, and y_i are their corresponding outputs.

Smaller values of k make the estimator more sensitive to local variations, while larger ones yield smoother, more stable predictions. To further improve adaptability, distance-weighted averaging was employed, assigning higher influence to closer observations and improving the smoothness of the resulting energy consumption curve. For uncertainty quantification, the KNN model followed the same block Jackknife resampling procedure previously introduced.

3.2.2.7 Distribution-Aware Models

In contrast to traditional regression methods that focus exclusively on estimating the conditional mean of the target variable, the Distribution-Aware models aim to characterize the entire conditional distribution. Within this study, Quantile Regression was adopted as the representative of this category, as this approach extends the concept of linear regression by estimating conditional quantiles, instead of the mean. As a result, it allows the model to capture asymmetric error behavior and varying levels of variability across time, providing a more complete view of the possible range of energy consumption values.

Mathematically, quantile regression estimates the parameter vector $\boldsymbol{\beta}$ by minimizing the check loss function $\rho_\tau(\cdot)$ [62]

$$\min_{\boldsymbol{\beta}\in\mathbb{R}^p} \sum_{i=1}^{n} \rho_\tau \left(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}\right), \tag{23}$$

where

$$\rho_\tau(u) = \begin{cases} \tau u, & \text{if } u \geq 0, \\ (\tau - 1)u, & \text{if } u < 0, \end{cases} \tag{24}$$

and $\tau \in (0, 1)$ denotes the quantile level.

In this work, three quantile levels were estimated, namely, $\tau = 0.025$, $\tau = 0.5$, and $\tau = 0.975$. The model trained at $\tau = 0.5$ represents the conditional median of the energy consumption distribution, while the lower and upper quantiles define an empirical uncertainty band that encompasses approximately 95% of the predicted distribution.

Unlike methodologies that could also be used, such as the Gaussian-based approaches, this technique does not rely on any assumption about the error distribution, deriving its prediction intervals directly from the observed quantiles. Thus, by using quantiles instead of residuals or normal as-

sumptions, the method produces asymmetric confidence bands that more accurately reflect the actual risk profile of energy consumption during berthing periods.

3.2.2.8 Other Models

To incorporate aspects not yet addressed by the previous regression families, such as smoothness and monotonicity, two additional models were considered in this study. Spline Regression was the first one and represents the regression function as a sequence of low degree polynomials joined smoothly at specific points called knots. This construction allows the model to capture gradual variations in the power signal while maintaining a controlled level of smoothness, avoiding the oscillations commonly introduced by high degree global polynomials. For a spline of degree p with K knots $\xi_1, \ldots, \xi_K$, the general formulation is [63]

$$f(x) = \beta_0 + \beta_1 x + \cdots + \beta_p x^p + \sum_{k=1}^{K} \gamma_k (x - \xi_k)_+^p, \tag{25}$$

where $(x - \xi_k)_+^p$ denotes the truncated power function that activates only when $x > \xi_k$. Adapting the number and placement of knots allows the model to balance flexibility and smoothness.

Isotonic Regression was the second method. It is a non-parametric technique that enforces a monotonic relationship between the predictor and the target variable. This makes it particularly suitable in contexts where domain knowledge implies a one directional trend. For instance, when the hotel load is expected to increase steadily during passenger boarding or gradually decrease as onboard activity declines.

The isotonic fit is obtained by solving the following constrained optimization problem [64]

$$\min_{\hat{y}_1 \leq \cdots \leq \hat{y}_n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2, \tag{26}$$

which enforces the monotonicity condition across all fitted points. This ensures that the resulting consumption profile follows a consistent directional trend, aligning with realistic operational behavior.

For both Spline and Isotonic Regression models, predictive uncertainty was quantified using residual-based confidence bands consistent with the rest of the framework ($z = 1.96$ for 95% confidence), allowing direct comparison of uncertainty ranges across all regression families.

4 Case Study Design

This chapter focuses on explaining how the methodology described in the previous chapter is adopted in the case study of the Port of Funchal. The goal is to integrate the developed regression models with real docking records and vessel specifications to simulate the port's SSE demand. The case study assesses in realistic terms the electrification needs of the port and its main constraints by applying the general framework to the conditions of Funchal.

4.1 Case Context: The Port of Funchal

Due to its strong reputation as a leading European tourist destination, the island of Madeira receives a large number of cruise ships every year at the Port of Funchal. This prominence led to the island being awarded *"Europe's Best Cruise Destination 2022"* by the World Cruise Awards [65], further highlighting the significant traffic handled by the port. However, the port is not yet equipped with facilities to provide SSE to vessels, forcing them to keep their auxiliary engines running during berthing periods.

The continuous use of these onboard generators results in high fossil fuel consumption and the emission of atmospheric pollutants SO_x, NO_x, and PM. These emissions deteriorate local air quality and pose negative health impacts on the residents of Funchal. This concern is amplified by the city's high population density, which in 2023 reached 1,411.2 inhabitants/km^2, far above the regional average of 320.3 inhabitants/km^2, making it the most densely populated municipality in the archipelago [66].

In this context, the adoption of SSE in the Port of Funchal emerges as a key solution to improve air quality and reduce noise within the city. By enabling cruise ships to connect to the onshore power grid during their stay, the most harmful emissions to residents' health could be significantly reduced, while at the same time aligning with the decarbonization targets established by the IMO.

For this reason, the application of the methodology developed in this thesis represents both a significant advancement and a valuable opportunity for the study and future implementation of port electrification in Madeira, providing a realistic and data-driven foundation for assessing the potential impact of SSE in the Port of Funchal.

4.2 Methodology Application

For the application and implementation of the previously defined methodology, **Python** (3.12.4) was the programming language selected, due to its prominence in the scientific community for ML applications, its extensive documentation, clear syntax, and prior familiarity with the language. Additionaly, Python offers a broad ecosystem of libraries that simplify and structure the different stages of the workflow.

In data processing, **pandas** (2.2.2) was used to read `.xlsx` files, clean and standardize columns, combine dates and times into consistent timestamps, and merge the scheduling information with Madeira's hotel power data, enabling reliable handling of the different `slices`. **NumPy** (1.26.4) provided vectorized computation to generate temporal grids, build masks for some of the resampling procedures implemented (jackknife), compute statistics (such as interval widths), and perform interpolation when "stretching" profiles to the required number of points. To speed up batch processing, **joblib** (1.4.2) parallelized the generation and loading of profiles by ship and `slice`. Modeling and evaluation were implemented using **scikit-learn** (1.4.2), which supported the implementation of the defined models, as well as the use of `scalers` and evaluation metrics. For obtaining plots and visual representations that allowed a better understanding of the port's energy consumption, **Plotly** (5.22.0) was used, facilitating the exploration of time series, comparison of scenarios, and visualization of prediction bands.

Some libraries do not have a specific version, only the standard one defined by the version of **Python** used (3.12.4). The **os** and **pathlib** libraries were employed to orchestrate paths and directories (including the cache of profiles and the slices index) in a cross platform and readable way; **hashlib** generated unique MD5 identifiers for each configuration/`slice` to ensure reproducible cache keys; **pickle** enabled persistence (serialization) of profiles, metrics, and metadata to avoid recomputation; **functools** (`lru_cache`) memorized recurrent results such as the selection of the best `slice` and the "stretched" profiles; **dataclasses** was used to structure and organize the configuration of the models (parameters and tags of each configuration), thus simplifying class definition; and **typing** was used for the typed definition of structures (such as `Dict` and `Tuple`), improving readability, enabling static verification, and facilitating code maintenance.

4.2.1 Data Preparation

Building on the methodology itself, the first step for implementing the methodological pipeline is the Data Preparation stage. The port scheduling data, which provides details on the different calls made by vessels at the Port of Funchal, as well as the duration of each stay, was supplied by Administração dos Portos da Região Autónoma da Madeira (APRAM) in a `.xlsx` file. This dataset contained the arrival and departure dates and times of each ship during the year 2024, as shown in Figure 8. It should be noted that the same information is also publicly available through the official APRAM platform in the berth reservation section [67].

	A	B	C	D	E
1	Arrival Date	Arrival Time	Departure Date	Departure Time	Ship Name
2	01/01/2024	06:00:00	01/01/2024	23:00:00	AIDA NOVA
3	01/01/2024	04:00:00	01/01/2024	23:58:00	AIDA COSMA
4	02/01/2024	06:00:00	02/01/2024	22:30:00	AZURA
5	02/01/2024	06:00:00	02/01/2024	23:59:00	AIDA COSMA
6	02/01/2024	06:00:00	02/01/2024	22:00:00	MS AMADEA
7	02/01/2024	07:00:00	03/01/2024	06:00:00	AIDA STELLA
8	07/01/2024	07:00:00	07/01/2024	18:00:00	AURORA
9	07/01/2024	07:00:00	07/01/2024	18:00:00	VENTURA
10	07/01/2024	07:00:00	07/01/2024	18:00:00	MEIN SCHIFF 3
11	07/01/2024	07:00:00	07/01/2024	16:00:00	NORWEGIAN EPIC
12	08/01/2024	06:00:00	08/01/2024	22:30:00	AZURA
13	08/01/2024	05:00:00	08/01/2024	23:58:00	AIDA NOVA
14	08/01/2024	06:00:00	08/01/2024	23:58:00	AIDA COSMA
15	08/01/2024	23:00:00	10/01/2024	14:00:00	MEIN SCHIFF 1
16	08/01/2024	09:30:00	08/01/2024	20:00:00	BALMORAL

Figure. 8: Sample of the port scheduling dataset provided by APRAM, containing the arrival and departure times of cruise ships in Funchal during 2024.

Once the different vessels scheduled to berth at the Port of Funchal were identified, the next step consisted of collecting their corresponding physical specifications, namely GT, LOA, and MPC for each registered ship. These data were retrieved from freely accessible online maritime databases [68, 69] and subsequently stored in a `.xlsx` file to facilitate structured integration into the modeling pipeline, as illustrated in Figure 9.

	A	B	C	D
1	Ship Name	Gross Tonnage	Length	Passengers
2	Sea Dream 2	4228	138	112
3	Corinthian	4077	88	100
4	National Geographic Explorer	6471	112	162
5	Exploris One	6130	108	173
6	Ocean Explorer	8228	104	162
7	World Explorer	9930	126	200
8	World Navigatorr	9930	126	200
9	World Traveller	9930	126	200
10	World Voyager	9930	126	200
11	MS Spitsbergen	7344	101	335

Figure. 9: Sample of the cruise ship physical specifications dataset, showing a small portion of the dataset containing GT, LOA and MPC.

Finally, the most relevant dataset consisted of the time series of energy consumption, which formed the core input for generating hotel energy profiles. These correspond to real measurements from four reference cruise ships, referred to as `Ship 1`, `Ship 2`, `Ship 3`, and `Ship 4`, provided by an anonymous supplier in a single `.xlsx` file divided into four sheets, each corresponding to one vessel. The dataset included both the fixed HL value of each vessel and the timestamped records of total hotel power demand (in kW), as depicted in Figure 10.

	A	B	C
1	Timestamp	Chillers Power	Hotel Power
2	2021-08-16 04:20:00	684,1666698	9186,409116
3	2021-08-16 04:21:00	684,1968117	9161,769557
4	2021-08-16 04:22:00	684,2269535	9137,129999
5	2021-08-16 04:23:00	684,2570953	9112,490441
6	2021-08-16 04:24:00	684,2872371	9087,850882
7	2021-08-16 04:25:00	684,3173789	9063,211324
8	2021-08-16 04:26:00	684,3475207	9038,571766

Figure. 10: Sample of the time series data used for modeling, containing timestamped records of Chillers Power and Hotel Power for each ship.

4.2.2 Data Preprocessing and Resampling

Following the data acquisition stage, the next step of the methodological pipeline focused on the Data Preprocessing and Resampling phase. This process aimed to transform the raw cruise ship energy data into standardized and temporally consistent time series suitable for subsequent curve fitting and modeling.

The procedure began with the raw `CruiseData.xlsx` file, which contained the operational power measurements of several vessels. Each cruise ship had its own dedicated sheet, so the first step consisted of selecting the reference ship sheet to be processed. Once selected, the data were loaded and parsed, ensuring proper interpretation of the raw structure.

After loading, only the relevant columns (`Timestamp`, `Hotel Power`, and `Chillers Power`) were retained. This selection step isolated the key variables required for the construction of hotel energy consumption profiles, removing unrelated information and reducing data dimensionality.

Next, the dataset was filtered to retain only the records corresponding to periods when the ship was docked (`State = "In_Port"`), thus isolating the energy consumption associated with hotelling operations. After filtering, the timestamps were cleaned, ordered chronologically, and any invalid or missing entries were removed to guarantee temporal coherence.

To identify interruptions in the data, consecutive time differences were calculated, and any gap longer than 60 minutes was flagged as a discontinuity. These gaps defined the boundaries for splitting the time series into individual segments, or slices, each representing a continuous stay of the vessel at berth. Finally, the slices were exported as independent Excel files following the standardized naming convention `ShipX_slice_K.xlsx`, where `X` indicates the ship number and `K` the chronological order of its port calls. The used workflow is depicted in Figure 11.

Once the creation of the slices was completed, all time series underwent a resampling and standardization process to ensure that the data from different vessels shared the same temporal resolution. Each slice was resampled to a uniform 1-minute interval and interpolated using a time-based method to fill minor gaps, guaranteeing temporal consistency across all slices datasets.

After discarding anomalous or excessively short slices, we retained 19 slices for Ship 1, 37 for Ship 2, 71 for Ship 3, and 82 for Ship 4 (total: 209). These preprocessed and resampled datasets served as the core input for the curve fitting framework.

4.2.3 Two-Stage Model Development

At the end of the second phase, all slices were obtained and standardized, providing the necessary input for the generation of hotel energy consumption profiles of the vessels calling at the Port of Funchal, a task carried out during the Two-Stage Model Development stage, which is structured

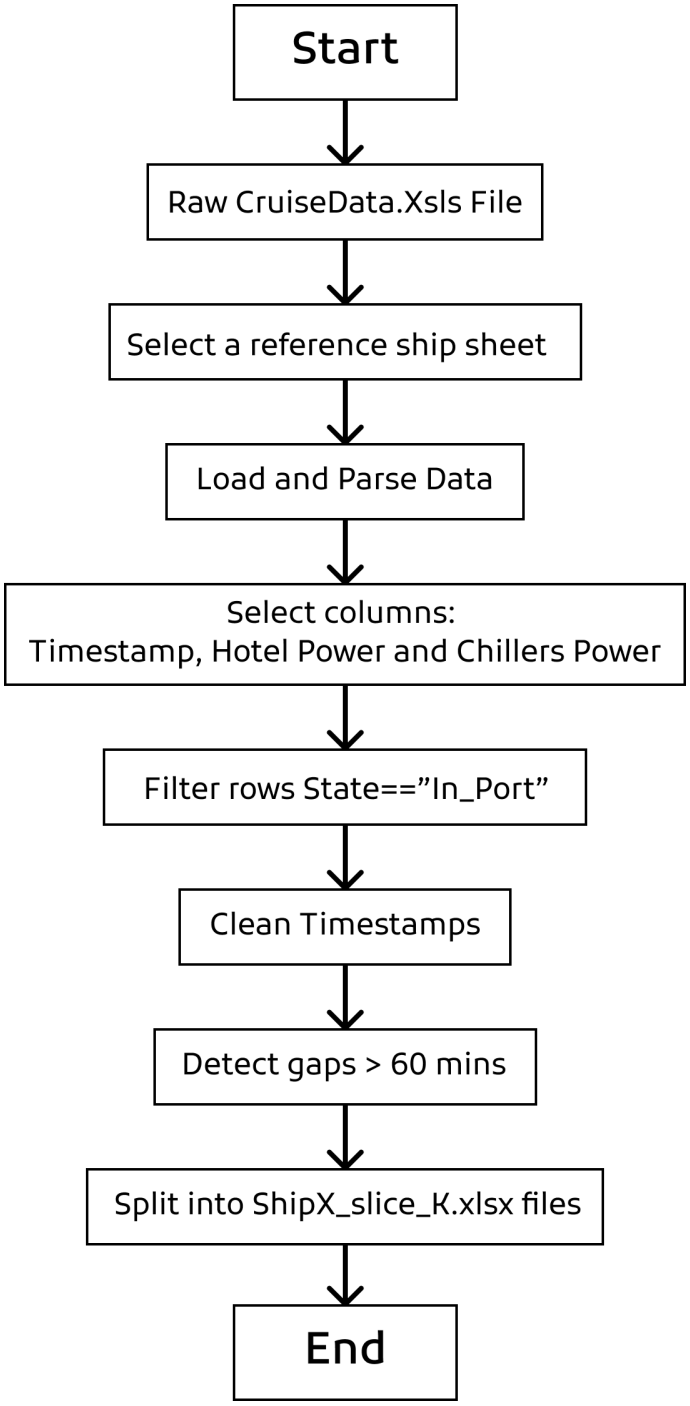

Figure. 11: Workflow for generating time series slices for each reference vessel.

into two main parts: HL Estimation through Regression, and Curve Fitting for Consumption Profiles.

For the first main part, as explained in Subsection 3.2.1, the process integrates the physical specifications of a set of real cruise vessels, along with four reference ships used for model calibration. The variables considered include GT,LOA, MPC, and estimated HL.

Figure 12 and Table 2 together illustrate an extract from the regression dataset used in this study (`RegressionData.xlsx`), which combines the data obtained by Ericsson et al. [49] with the physical specifications of the reference vessels. This unified dataset served as the foundation for establishing the baseline hotel load estimation.

Name	IMO	GT	Length m	Passengers	Hotelling kW
NATIONAL GEOGRAPHIC ENDEAVOUR	6611863	3132	88	110	1280
ISLAND SKY	8802894	4280	91	114	3600
SEA CLOUD II	9171292	3849	117	96	1290
VISTAMAR	8701193	7498	121	330	2360
HANSEATIC	9000168	8378	123	184	4400
OCEAN MAJESTY	6602898	10417	131	535	4260
SEABOURN PRIDE	8707343	9975	134	204	3200
SPIRIT OF ADVENTURE	7904889	9570	139	470	2112
SILVER CLOUD	8903923	16927	156	296	6912
SILVER WIND	8903935	16927	156	315	6912

Figure. 12: Extract of real vessel specifications included in `RegressionData.xlsx`, containing GT, length, passengers, and estimated hotelling power (kW) [49].

Reference Ship	Gross Tonnage (GT)	Length (m)	MPC	HL (kW)
Ship 1	150,000	320	4,650	15,000
Ship 2	55,000	220	1,506	6,300
Ship 3	140,000	290	4,378	13,000
Ship 4	41,000	210	1,029	4,500

Table 2: Specifications of the four reference ships included in `RegressionData.xlsx`.

With this information defined, it was possible to apply the two approaches for HL estimation, as described in Subsection 3.2.1. The first approach follows the procedure of Espinosa et al. [13], using a potential regression framework where GT acts as the single explanatory variable. Both GT and HL were transformed into their logarithmic forms to fit a linear regression under a potential (power-law) formulation, as expressed in Equation 3. The resulting equation derived from the regression dataset was

$$y = 67.90 \cdot x^{0.4187}. \tag{27}$$

Applying this model to the regression dataset yielded an in-sample R^2 of 0.6703, with a RMSE of 1,590.31 kW and a MAE of 1,209.05 kW. Figure 13 shows the fitted potential model, where the non-linear relationship between GT and HL is captured through the power-law function.

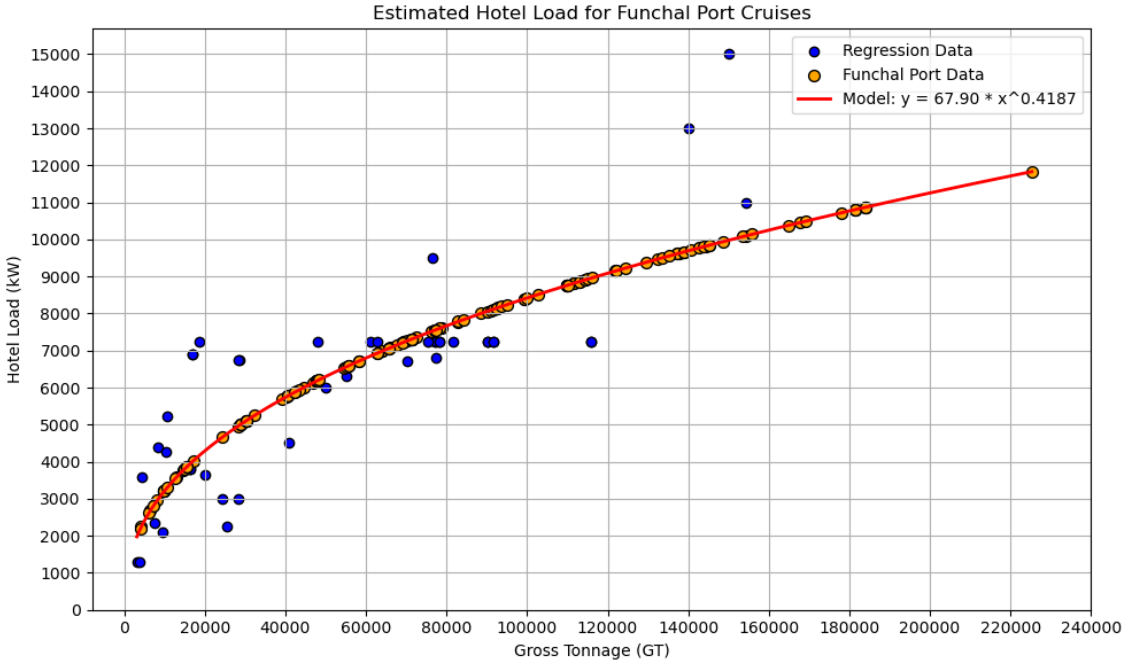

Figure. 13: Potential regression fit (power law) between GT and HL. Blue points represent the regression dataset cruises, orange points represent Funchal Port cruises, and the red line is the fitted power-law model.

Although this model effectively represents the non-linear dependency between GT and HL, the dispersion observed around the fitted curve highlights the inherent variability across real world vessels due to differences in onboard systems, operational conditions, and passenger services. Nonetheless,

the potential regression serves as a validated and interpretable baseline for estimating hotel load in cruise ships.

The second approach for HL estimation relied on an MLR model, which incorporated all available ship characteristics (GT, LOA, and MPC) as explanatory variables, while the HL was used as the dependent variable. An important methodological step in this approach was the estimation of missing MPC values for certain reference ships. To ensure a complete dataset, MPC was estimated using an auxiliary regression model with GT and LOA as predictors, leveraging their strong correlations (0.96 and 0.90 respectively, as shown in Figure 4). This ensured consistency and completeness for the subsequent modeling stage. The final MLR model obtained from the consolidated dataset (`RegressionData.xlsx`) was expressed as:

$$\hat{y}\ (\text{kW}) = 3734.23 + 0.055034 \cdot GT - 4.519549 \cdot LOA + 0.265916 \cdot MPC. \tag{28}$$

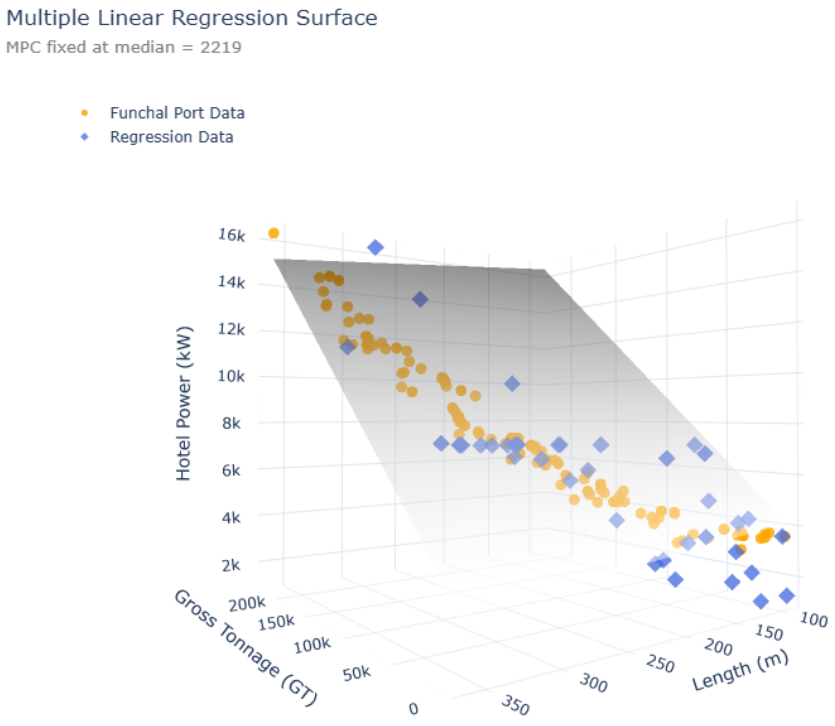

Figure. 14: Multiple Linear Regression (MLR) surface relating GT, LOA, and HL (with MPC fixed at the median value). Blue points represent the regression dataset cruises, orange points correspond to Funchal Port cruises, the grey surface represents the fitted MLR plane.

Figure 14 presents the three-dimensional visualization of this regression model. The grey plane represents the fitted regression surface (with MPC fixed at its median value), while the blue and red markers correspond respectively to the predicted HL of the Madeira cruise dataset and the original regression data used for model calibration.

As observed, both datasets exhibit a consistent alignment with the fitted plane, confirming the linear relationship between vessel size parameters and predicted hotel power demand. The inclusion of both MPC and LOA introduces an additional dimension that refines the model's predictive ability, accounting for variability associated with onboard services and occupancy. The MLR model, therefore, provides a more interpretable framework for estimating HL across vessels of varying size and capacity.

When compared to the potential regression baseline, the MLR model achieved higher explanatory power and lower error: $R^2 = 0.7035$ vs 0.6703, RMSE = 1,508.16 kW vs 1,590.31 kW, and MAE = 1,198.37 kW vs 1,209.05 kW. These results show that incorporating multiple physical and operational parameters yields measurable accuracy gains while preserving interpretability and computational simplicity.

Although both models perform well for HL estimation, the potential regression remains aligned with the previous literature and provides a robust baseline for cruise ship power prediction. In contrast, the MLR model leverages multiple correlated predictors to offer a more refined, yet still transparent, representation of cruise energy-demand behavior. Consequently, we apply the potential regression formula (Equation 27) to the dataset of vessels calling at the Port of Funchal to estimate HL for each ship, and a sample of the resulting values is illustrated in Table 3.

Name	**Gross Tonnage**	**Length (m)**	**Passengers**	**Hotel Load (kW)**
Corinthian	4,077	88	100	2,205.31
Sea Dream 2	4,228	138	112	2,239.15
Exploris One	6,130	108	173	2,615.91
National Geographic Explorer	6,471	112	162	2,675.90
AIDA Cosma	183,900	337	6,600	10,866.08
AIDA Nova	183,900	337	6,600	10,866.08
Iona	184,089	345	6,264	10,870.75
Oasis of the Seas	225,282	362	6,431	11,829.8

Table 3: Sample of estimated HL for vessels calling at Funchal based on the potential regression model.

For the second main part of this stage, the focus is placed on the generation of hotel power consumption profiles through curve fitting techniques. While the regression models provide the static HL for each vessel, to represent the dynamic behavior of energy demand during a port call, these values need to be translated into continuous temporal values. To achieve this goal, the preprocessed and standardized slices obtained in the earlier data preparation phase were used as the basis for applying the different curve fitting methods described in Subsection 3.2.2.

Since the original slices were provided in `.xlsx` format, an additional preprocessing step was introduced to enhance efficiency. Each file was converted into `Parquet` format and stored in a dedicated cache directory. This conversion was motivated by the performance limitations of working directly with Excel files, which tend to be slower and heavier in large scale modeling. The Parquet format, on the other hand, integrates well with the `pandas` library, allowing for faster read/write operations, selective column access (avoiding loading the entire dataset into memory), and efficient compression that significantly reduces storage requirements. It also preserves data types more accurately, preventing unintended type conversions. As a result, the adoption of Parquet ensured greater efficiency, consistency, and scalability in handling all slice files throughout the fitting pipeline.

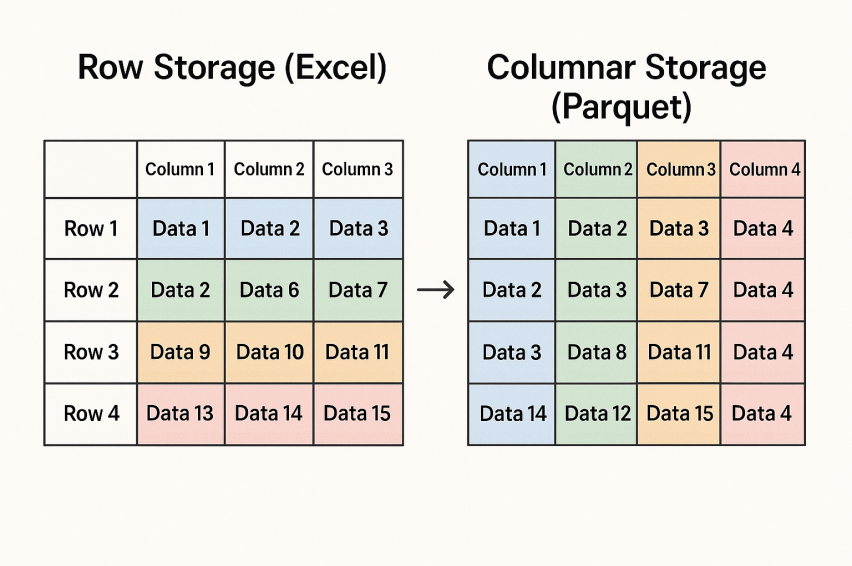

Figure. 15: Comparison between Row Storage (Excel) and Columnar Storage (Parquet).

As illustrated in Figure 15, the shift from the traditional row-based storage used by Excel to the columnar structure of Parquet directly contributes to performance improvements.

An important aspect to consider is that, beyond using Parquet, another alternative to improve data processing performance would have been to implement a relational database (such as an `SQL` system), which can offer optimized read operations through indexing and query planning. However, given that this is a single user, batch-oriented system, where neither concurrent queries nor real time service are required, setting up and maintaining an SQL server would have introduced operational complexity that is not justified by the marginal performance benefits. Therefore, the simpler approach based on Parquet files was maintained, providing an appropriate balance between simplicity and performance for the workflow considered.

Once the conversion to Parquet was implemented, a set of helper functions was designed to standardize the fitting process. These utilities manage the construction of feature matrices, the partitioning of data into training and testing subsets, the estimation of uncertainty, and the packaging of results in a consistent format across all model backends.

As explained in the subsection 3.1.4, the data partitioning for standard model evaluation was handled by a function that performs a reproducible random split corresponding to 80% for training and 20% for testing. This approach was used to assess predictive performance under typical conditions.

For models requiring uncertainty quantification, two complementary methods were implemented. The first computes confidence intervals based on the standard deviation of the training residuals, assuming approximately normal residual distributions and producing symmetric confidence bands around the mean predictions.

The second employs the Jackknife resampling method, which replaces the fixed 80/20 partitioning with a systematic resampling strategy. As previously discussed, this approach divides the dataset into contiguous blocks, and one block is excluded at each iteration while the model is trained on the remaining data. This process is repeated until every block has been left out once. As illustrated in Figure 16, this leave-one-block-out procedure generates multiple models, each producing slightly different predictions due to the varying training subsets.

The ensemble of resulting predictions is then aggregated to compute both the mean estimate and the variability bands (confidence intervals) across iterations. These bands reflect how sensitive the model's predictions are to changes in the training data, providing a non-parametric estimate of uncertainty that does not rely on distributional assumptions (normality of residuals, for example).

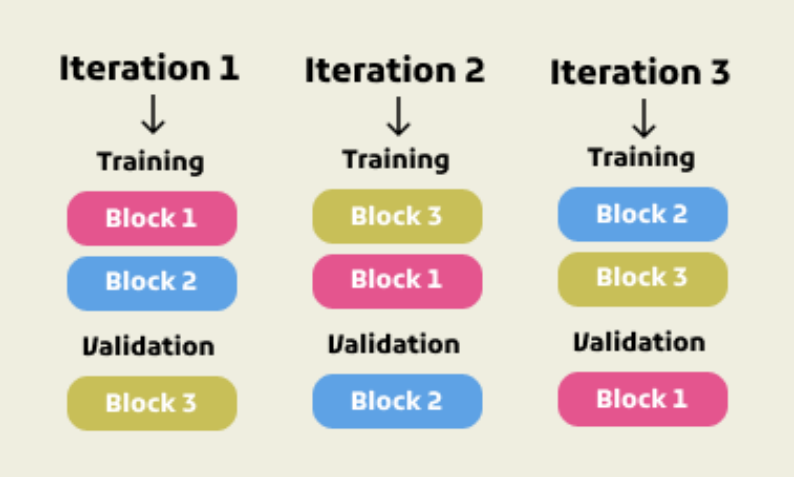

Figure. 16: Jackknife by blocks with three partitions. Each iteration excludes one block while training on the remaining data to estimate prediction variability.

Compared to a single random split, this approach offers a more robust and data efficient evaluation of model stability.

Another aspect addressed in the set of helpers was feature normalization. To ensure consistent model performance, especially given that many regression algorithms are sensitive to differences in magnitude among predictors, a normalization was applied. The used method transforms all features to a range of $[0, 1]$, allowing each variable to contribute comparably to the model fitting process. This approach maintains the shape of the original distributions (which is important when working with data that, in our case, comes from different reference ships and presents distinct value ranges for the same features). This measure was consistently applied to the training, testing, and full datasets, ensuring methodological coherence across all experiments. Figure 17 illustrates an example of the normalization process.

Another aspect considered was the robustness of the pipeline against a lack of data. Although the data had already undergone the preprocessing stage, situations could arise where a slice contained an insufficient number of data points to allow for meaningful and statistically valid modeling. To prevent pipeline interruptions in such scenarios, a safeguard mechanism was introduced. Specifically, in these cases, the observed values of the segment were directly passed through as predictions. At the same time, the performance metrics associated with that segment were assigned as null values (`NaN`).

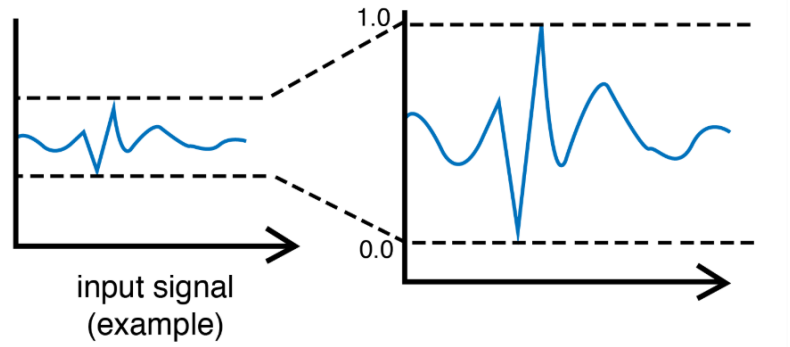

Figure. 17: Visual Representation of the Normalization Process using a Min-Max Scaler [70].

This set of helper functions ensured the integrity and continuity of the data processing stage, even under imperfect conditions. Building on this foundation, the next part of the framework focused on how models are executed and results are managed, introducing an architecture designed to keep the system scalable and efficient across different regression models. This architecture can be separated into two main processes. First is the Model Execution Process, which handles model configuration and computation, and then the Caching and Management Process, which manages result storage and retrieval to ensure computational efficiency and reproducibility.

The first process begins with the configuration class (`ModelConfig`), which encapsulates the model name and parameters while generating a unique `cache tag` that links every stored result to its specific setup. This mechanism automatically invalidates cached files whenever model parameters change, preventing inconsistencies between stored and current configurations.

Each model is implemented as an independent backend function (`ridgecv_profile_stats`, `lassocv_profile_stats`, etc.), which follows a standardized interface, receiving the input data and shared parameters, executing the model, and returning both a prediction `DataFrame` and a dictionary of metrics. The resulting `DataFrame` is generated internally by the `_pack` function, which formats the model outputs into a consistent structure including the timestamp, predicted mean, and the lower and upper confidence bounds, thus ensuring that all models generate comparable outputs ready for evaluation and visualization.

These backends are registered centrally in the `BACKENDS` dictionary, which acts as a dynamic dispatcher that selects the appropriate backend at runtime. The function `compute_profile_df` then serves as the unified entry point for model execution: given a `ModelConfig`, it retrieves the correct

backend, applies the configured parameters, and returns standardized outputs for evaluation. The main structural components and their interaction are shown in Listing 1.

```python
# Model configuration + unique cache tag
@dataclass(frozen=True)
class ModelConfig:
    name: str = "linear"
    params: Dict[str, Any] = None
    @property
    def cache_tag(self) -> str:
        # Generates a unique identifier from model name and parameters.
        h = hashlib.md5(str(self.params).encode()).hexdigest() if self.params else "none"
        return f"model={self.name}|params={h}"
# Unified output format for predictions
def _pack(df, mean, low, up) -> pd.DataFrame:
    return pd.DataFrame({
        "Timestamp": df["Timestamp"].values, "Predicted": mean, "Lower": low, "Upper": up
    })
# Common metrics
def _metrics(*args, **kwargs) -> dict:
    # ... Full implementation in source code ...
    return {"R2": np.nan, "RMSE": np.nan, "MAE": np.nan, "PICP": np.nan, "CWC": np.nan, "KL": np.nan}
# Backends (one per model): return (pred_df, metrics)
def ridgecv_profile_stats(df, degree=3, **params):
    # ... Full implementation in source code ...
    return _pack(df, y_full, lower, upper), metrics
# Other models follow the same structure
# Central backend registry
BACKENDS: Dict[str, Callable[..., Tuple[pd.DataFrame, dict]]] = {
    "ridge":      ridgecv_profile_stats,
    "lasso":      lassocv_profile_stats,
    "elasticnet": elasticnetcv_profile_stats,
    # ... more models ...
}
# Runs the selected backend model and returns standardized predictions and metrics
def compute_profile_df(df: pd.DataFrame, cfg: ModelConfig, seed=42, test_frac=0.2, z=1.96):
    fn = BACKENDS.get(cfg.name.lower())
    if fn is None:
        raise ValueError(f"Unsupported model: {cfg.name}. Available: {list(BACKENDS.keys())}")
    params = dict(cfg.params or {}); params.update({"seed": seed, "test_frac": test_frac, "z": z})
    return fn(df, **params)  # -> standardized (pred_df, metrics)
```

Listing 1: Core components: configuration, backend registry, and standardized output format

The second process handles the execution of each model per data slice, ensuring efficient reuse of results and avoiding redundant computations. This layer checks whether a cached result already exists for a given combination of ship, slice, and model configuration. If a cached result is available, the precomputed predictions and metrics are retrieved directly. Otherwise, the data slice is loaded, processed, and passed through the corresponding backend.

Caching is implemented using Python's native `pickle` serialization format, which stores both the predictions and evaluation metrics in binary form. The cache key incorporates the ship identifier, slice name, random seed, confidence level, last modification time of the source file, and model tag, ensuring that each cached entry corresponds uniquely to a specific configuration and dataset version. The main caching logic and orchestration flow is illustrated in Listing 2.

```
CACHE_DIR = Path("./_cache_ports"); CACHE_DIR.mkdir(exist_ok=True)
MODEL_CACHE = True
# Builds cache key and handles load/save
def _model_cache_key(ship_ref, slice_name, seed, z, path_xlsx, model_tag):
    meta = f"{ship_ref}|{slice_name}|seed={seed}|z={z}|{model_tag}"
    return CACHE_DIR / f"profile_{hashlib.md5(meta.encode()).hexdigest()}.pkl"
# Loads profile from cache
def _load_profile_from_cache(path):
    try: return pickle.load(open(path, "rb"))
    except: return None
# Computes and stores a profile in cache
def _save_profile_to_cache(path, df_pred, metrics):
    pickle.dump((df_pred, metrics), open(path, "wb"), protocol=pickle.HIGHEST_PROTOCOL)
# Loads from cache if available, otherwise computes and stores
def _compute_or_load_profile(ship_ref, slice_name, slices_dir, MODEL_CFG, seed, z, test_frac):
    path = os.path.join(slices_dir, ship_ref, slice_name)
    cache = _model_cache_key(ship_ref, slice_name, seed, z, path, MODEL_CFG.cache_tag)
    if MODEL_CACHE and (cached := _load_profile_from_cache(cache)):
        return ship_ref, slice_name, *cached
    df = load_slice_df(path)
    pred_df, metrics = compute_profile_df(df, cfg=MODEL_CFG, seed=seed, test_frac=test_frac, z=z)
    if MODEL_CACHE: _save_profile_to_cache(cache, pred_df, metrics)
    return ship_ref, slice_name, pred_df, metrics
```

Listing 2: Caching layer and management logic for per-slice execution

Figure 18 summarizes the defined architecture. The process begins by checking the cache for existing results; if they are available, they are retrieved directly. Otherwise, the pipeline loads the slice (from its Parquet version if possible), applies the appropriate preprocessing, runs the selected backend model, and stores the resulting predictions and metrics back into the cache.

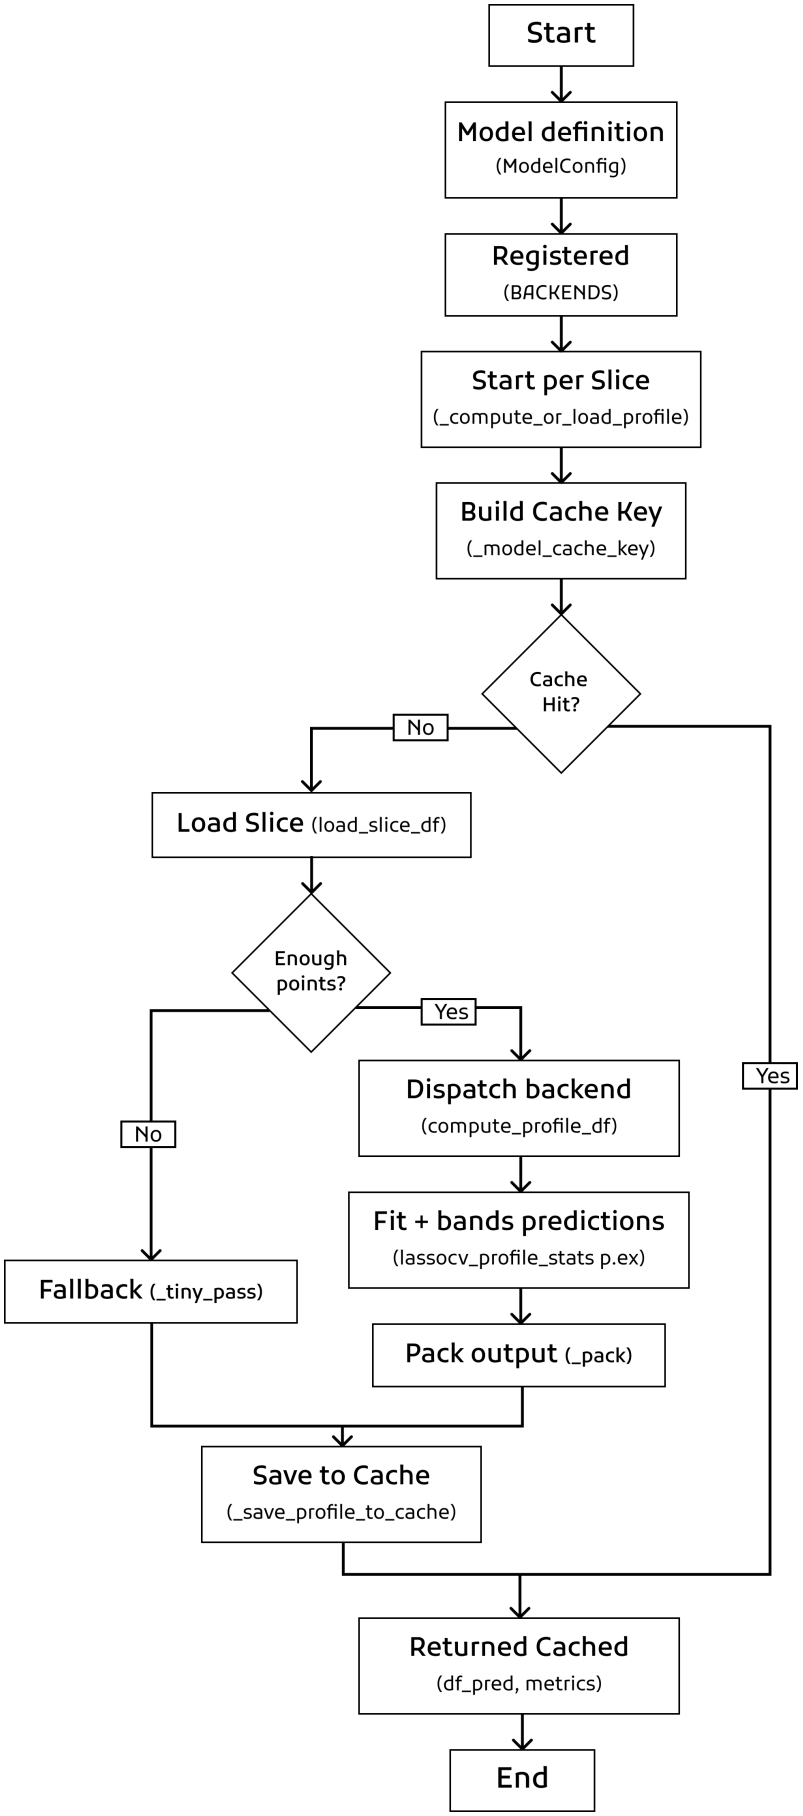

Figure. 18: Flow diagram of the fitting management layer, showing cache verification, slice loading, backend dispatch, and persistence of predictions.

4.2.4 Model Evaluation

As introduced in subsection 3.1.4 and illustrated in Figure 19, all metrics are computed under an 80/20 random train–test split per slice. Concretely, 80% of each slice is used for training and the remaining 20% for testing, so that the reported performance reflects generalization on unseen samples from the same temporal context. Importantly, metrics are aggregated using a macro average. This means that each metric is first computed on the test portion of every slice independently, and the mean of those per slice results is then reported across all slices. This procedure gives equal weight to each slice regardless of its length or number of test samples.

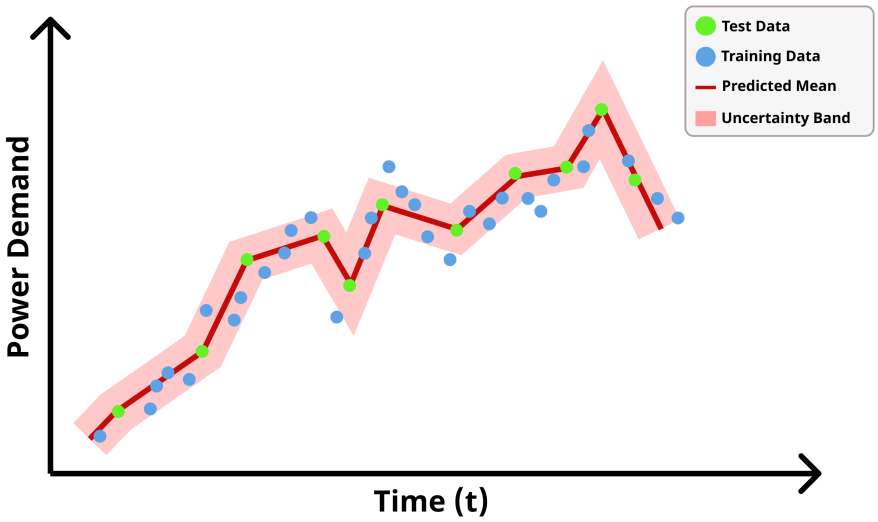

Figure. 19: Illustration of the 80/20 train–test data split used. Blue points represent training data, green points correspond to test data, the red line denotes the predicted mean, and the shaded area represents the associated uncertainty band.

It is important to note that evaluation covers three complementary perspectives, indicated as point–prediction quality, uncertainty quality, and the consistency of the port's aggregated consumption. The first one assesses how closely point forecasts replicate the observed trajectory at each timestamp, independently of uncertainty bands. It focuses on the magnitude of errors, their robustness to outliers, and the fraction of variance in the target explained by the model.

Regarding the used metrics [71], the RMSE indicates the typical error in the same units as y (kW) and *magnifies* large deviations due to the square, according to

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2}, \tag{29}$$

where y_i are the observed values, $\hat{y}_i$ are the model predictions, and n is the number of samples.

The MAE then gives the average absolute difference (kW) and is more robust than RMSE to occasional large misses.

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}\left|y_i - \hat{y}_i\right|. \tag{30}$$

R^2 is particularly relevant in this context because it measures how much of the variability in the observed target values is explained by the model. Formally, it compares the model's residual error to the variability of the data around its mean:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}i)^2}{\sum i = 1^n (y_i - \bar{y})^2}, \tag{31}$$

where $\bar{y}$ denotes the average of the observed values.

Lastly, KL Divergence was used between the distributions of y and $\hat{y}$, using shared range histograms with B bins, according to [72]

$$\mathrm{KL}(P \parallel Q) = \sum_{j=1}^{B} p_j \, \log\!\left(\frac{p_j}{q_j}\right), \tag{32}$$

where p_j and q_j are normalized frequencies of the observed and predicted data in bin j.

A model shows good point–prediction quality when RMSE and MAE are small (errors in kW are low) and R^2 is close to 1. In curve fitting, a small KL additionally indicates that predicted and observed distributions agree.

For the coverage analysis, the PICP was used, indicating the empirical coverage on the test set. Hence, it shows the fraction of targets falling within the predicted band $[L_i, U_i]$ as [73]

$$\mathrm{PICP} = \frac{1}{N_{\mathrm{test}}}\sum_{i=1}^{N_{\mathrm{test}}} \mathbf{1}\{\, L_i \le y_i \le U_i \,\}, \tag{33}$$

where L_i and U_i denote the lower and upper bounds of the i-th prediction interval, y_i is the true value, and N_{test} is the number of test samples.

The NMPIW then averages the bandwidth and scales it by a robust dispersion S of the target (here $S = P_{97.5}(y) - P_{2.5}(y)$) to allow fair comparisons between ranges and periods as [74]

$$\text{NMPIW} = \frac{1}{S N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} (U_i - L_i), \tag{34}$$

where $U_i - L_i$ is the width of the i-th prediction interval, S is a global robust scale computed over all accepted slices as $P_{97.5}(y) - P_{2.5}(y)$ (the central 95% empirical range), ensuring robustness to outliers and comparability across periods and models.

Lastly, CWC combines the coverage probability and the average width of prediction intervals into a single score. It "rewards" narrow intervals with sufficient coverage and penalizes those that fail to reach the expected coverage level as [75]

$$\text{CWC} = \text{NMPIW} \left(1 + \gamma(\text{PICP})\, e^{-\eta\,(\text{PICP}-\mu)}\right), \tag{35}$$

where μ is the nominal confidence level (expected to be 0.95), η is a control parameter (typically $\eta = 50$), and $\gamma(\text{PICP})$ is defined by

$$\gamma(\text{PICP}) = \begin{cases} 0, & \text{if PICP} \geq \mu, \\ 1, & \text{if PICP} < \mu, \end{cases} \tag{36}$$

meaning that the exponential penalty activates only when the empirical coverage PICP falls below the target μ. This criterion penalizes under coverage exponentially.

A model exhibits good uncertainty quality or well calibrated uncertainty when PICP is close to or higher than 0.95% while NMPIW is small; a small CWC concisely reflects this balance.

However, the consistency of the port's aggregated consumption needs to be estimated, as it translates predictions into an operational magnitude by integrating the predicted power over the analysis window, yielding the total demand attributable to SSE. To implement this metric, the procedure is as follows. First, the total port energy consumption in MWh is calculated for each model. Then, the mean of this consumption across all models is computed. Lastly, the deviation of each model with respect to the mean is determined.

As a restuls, the total port consumption per model indicates predicted power accumulated over time using power in MW and Δt in hours, resulting in the port's total annual energy consumption (MWh)

$$C_{\text{port},m}\,[\text{MWh}] \;=\; \sum_{t=t_0}^{t_1} \hat{P}_m^{(\text{MW})}(t)\,\Delta t \;=\; \int_{t_0}^{t_1} \hat{P}_m^{(\text{MW})}(t)\,dt, \tag{37}$$

where m indexes the model, allowing an operational comparison of how much each model would require from the system over the same period.

The multi-model consumption mean then, given M models, estimates the cross-model mean as

$$\bar{C}_{\text{port}}\,[\text{MWh}] \;=\; \frac{1}{M}\sum_{m=1}^{M} C_{\text{port},m}. \tag{38}$$

The deviation from the mean is then assessed by

$$\Delta_m\,[\text{MWh}] \;=\; C_{\text{port},m} - \bar{C}_{\text{port}}. \tag{39}$$

These values represent each model's deviation from the multi-model mean. Positive values are above the mean, negative values are below it. A model exhibits good consistency of aggregated consumption when its annual total $C_{\text{port},m}$ is close to the multi-model mean $\bar{C}_{\text{port}}$. In practice, this corresponds to $\Delta_m \approx 0$.

4.2.5 Port Energy Demand Estimation

After evaluating the alternative models for profile generation, the final stage assigns reference consumption profiles to vessels calling at the Port of Funchal and aggregates them into a port-wide demand series. The procedure operates at the level of scheduled stays (arrival–departure intervals) and maps each stay to a reference profile that is then temporally aligned and scaled in magnitude before being stacked into the aggregate.

For each stay, the estimated HL (in kW) is compared against the set of reference ships (Tables 3 and 2). The reference vessel whose HL is closest to the target is selected, and a scaling factor is computed as the ratio between the target and the chosen reference load to adapt the magnitude of the reference profile to the scheduled ship. To avoid undue influence of extreme cases, the scaling factor can be constrained to a plausible operating range.

Prior to the assignment stage, the scheduling dataset is standardized (column names and types), timestamps are constructed by combining dates and times, invalid records (departures earlier than arrivals or missing values) are removed, and HL is merged from the Madeira registry via a standardized vessel-name key. This produces a consistent schedule in which each stay has a reliable HL and valid temporal bounds.

Once reference ships have been mapped to Funchal calls, and the schedule has been cleaned, timestamped, and enriched with HL from the Madeira registry, the next step is to ensure temporal consistency. For the selected reference vessel, we choose the consumption slice whose start and end times best align with the stay's actual arrival–departure window. This matching uses a compact, precomputed, and cached metadata index (start/end per slice) stored as a pickle file, `slices_meta.pkl`, so the best-aligned slice can be retrieved in memory without reopening the original spreadsheets. The result, for each stay, is an assignment of `AssignedShip`, `BestSlice`, and the corresponding scaling factor. Figure 20 illustrates how this index is organized in `slices_meta.pkl`: top level keys are reference ships; within each, every slice file maps to a small record with `start_min`, `end_min`, and `mtime` (the file's last modification timestamp, in seconds, since corresponds to epoch) used to invalidate/update the cache.

```
{
 'Ship1':
   {'Resampled_Ship1_slice_01_resampled_1min.xlsx':
                                          {'end_min': 930,
                                           'mtime': 1759845707.0,
                                           'start_min': 260},

     'Resampled_Ship1_slice_02_resampled_1min.xlsx':
                                          {'end_min': 980,
                                           'mtime': 1755863460.17,
                                           'start_min': 320},

     'Resampled_Ship1_slice_03_resampled_1min.xlsx':
                                          {'end_min': 790,
                                           'mtime': 1755863460.22,
                                    :      'start_min': 500},

 'Ship2':{ ... }
 'Ship3':{ ... }
 'Ship4':{ ... }
}
```

Figure. 20: Structure of the `slices_meta.pkl` metadata index.

Finally, after temporal alignment, the scaling factor is applied to the predicted time series for the stay, multiplying every value of the `Mean`, `Lower`, and `Upper` profiles, so that magnitude adjustments are consistently propagated to the port-wide aggregate. However, it is also necessary to adjust the length of the predicted segment to the actual stay duration. To this end, the profile is linearly scaled to match exactly the number of time steps between arrival and departure before scaling and placement on the uniform grid.

Once the profile generation stage is complete, we move to the port level simulation. To reflect operational realities, we enforce the maximum number of simultaneous berths and a berth priority rule. The Port of Funchal accommodates at most three cruise ships at the same time. In berth conflicts, vessels of the **TUI Group** receive priority: if a TUI vessel arrives when three berths are already occupied, the longest staying non-TUI vessel is preempted to free capacity for the TUI call. Ties at identical arrival times are resolved randomly under a fixed seed to ensure reproducibility. The TUI vessels identified in the scheduling dataset are, under Marella Cruises: Marella Discovery, Marella Discovery 2, Marella Explorer, Marella Explorer 2, and Marella Voyager; and under Mein Schiff Cruises: Mein Schiff 1, Mein Schiff 2, Mein Schiff 3, Mein Schiff 5, Mein Schiff 6, and Mein Schiff 7.

The port-wide series is built on a uniform grid by stacking all accepted, aligned, and scaled segments. For each timestamp, the method sums the `Mean` and the `Lower/Upper` predictions of every active segment and records the list of active ships. In parallel, it accumulates per-vessel energy contributions (kWmin and kWh) to produce a concise breakdown of participation. The outputs are, therefore, the aggregate series with mean and uncertainty bands, annotated with active ships per timestamp; a contributions table with the number of stays and the energy contributed by each vessel; and an updated profile/metrics cache that ensures consistency for subsequent analyses and visualizations. The function that implements this final step is shown in Listing 3.

```
def build_port_series_for_range(df_sched, start, end, agg_freq="1min"):
    # Time grid
    start = pd.to_datetime(start); end = pd.to_datetime(end)
    step  = pd.to_timedelta(agg_freq)
    n     = int((end - start) // step)
    idx   = pd.date_range(start=start, periods=n, freq=agg_freq)
    if n <= 0:
        empty = pd.DataFrame(columns=["Mean","Lower","Upper","ActiveShips"
    ], index=idx)
```

```
        return empty, pd.DataFrame(columns=["Boat","NumStays","kWmin","kWh"]), {}
    # Window filter
    stays = df_sched[(df_sched["Arrival"] < end) & (df_sched["Departure"] > start)].copy()
    # Capacity + TUI priority
    stays["IsTUI"] = stays["Boat"].astype(str).apply(_is_tui)
    accepted = enforce_capacity_with_tui_preemption(stays, start, end, cap=PORT_CAP, seed=random_state)
    accepted = accepted[accepted["BestSlice"].astype(str) != ""].reset_index(drop=True)
    # Ensure profiles (cache) and expose to stretcher
    prof_cache, met_cache = ensure_profiles_for(accepted)
    global _GLOBAL_PROFILE_CACHE
    _GLOBAL_PROFILE_CACHE = prof_cache
    get_stretched_arrays.cache_clear()
    # Accumulators
    mean = np.zeros(n); low = np.zeros(n); up = np.zeros(n)
    active = [set() for _ in range(n)]
    contrib = {}
    # Stretch, apply Scaling factor and stack
    for _, r in accepted.iterrows():
        left, right = max(r["Arrival"], start), min(r["Departure"], end)
        m = int((right - left) // step)
        if m <= 0: continue
        i0 = int((left - start) // step); i1 = min(i0 + m, n)
        keep = i1 - i0
        if keep <= 0: continue
        ship_ref, slc = r["AssignedShip"], r["BestSlice"]
        sf = float(r.get("SF", 1.0)) # scaling factor
        # Stretch profiles to exactly m points (temporal alignment)
        seg_mean = get_stretched_arrays(ship_ref, slc, m, "Predicted")[:keep] * sf
        seg_low  = get_stretched_arrays(ship_ref, slc, m, "Lower")[:keep] * sf
        seg_up   = get_stretched_arrays(ship_ref, slc, m, "Upper")[:keep] * sf
        # Pointwise stacking of Mean / Lower / Upper
        mean[i0:i1] += seg_mean; low[i0:i1] += seg_low; up[i0:i1] += seg_up
        # Active ships and per-vessel energy (kWmin/kWh)
        b = str(r["Boat"])
        for j in range(i0, i1): active[j].add(b)
        d = contrib.setdefault(b, {"NumStays":0,"kWmin":0.0})
        d["NumStays"] += 1; d["kWmin"] += float(seg_mean.sum())
    # Outputs: aggregate series + per-vessel contributions + metrics/profile cache
    df_series = pd.DataFrame({
        "Timestamp": idx, "Mean": mean, "Lower": low, "Upper": up,
        "ActiveShips": [_format_ship_list(list(s)) for s in active]
```

```
    }).set_index("Timestamp")
    df_contrib = (pd.DataFrame(
        [{"Boat":b,"NumStays":v["NumStays"],"kWmin":v["kWmin"],"kWh":v["
    kWmin"]/60.0}
         for b,v in contrib.items()])
        .sort_values(["kWh","Boat"], ascending=[False,True]).reset_index(
    drop=True)
        if contrib else pd.DataFrame(columns=["Boat","NumStays","kWmin","
    kWh"])
    )
    return df_series, df_contrib, met_cache
```

Listing 3: Compact core of port-wide aggregation.

Once each ship is aligned with its `AssignedShip`, `BestSlice`, and their respective scaling factor `SF`, `build_port_series_for_range` performs the final aggregation as follows. It filters the schedule to the target window (`[end, start]`), tags vessels with `_is_tui`, and enforces berth capacity with priority using `enforce_capacity_with_tui_preemption`. Next, it ensures the required prediction profiles with `ensure_profiles_for` and exposes them to stretching; for each admitted stay, it clips the interval to the analysis window, and stretches the selected profile to the exact number of time steps and `get_stretched_arrays`, simultaneously scales the resulting segment (by multiplying the `SF` on `Mean`, `Lower`, and `Upper`), and places it on the uniform grid where all active segments are summed pointwise. Finally, it records the list of active ships per timestamp and builds the per-vessel energy summary, returning the aggregate series (`df_series`), the contributions table (`df_contrib`), and the updated metrics/profile cache (`met_cache`). In more detail, Figure 21 illustrates the execution flow of this function.

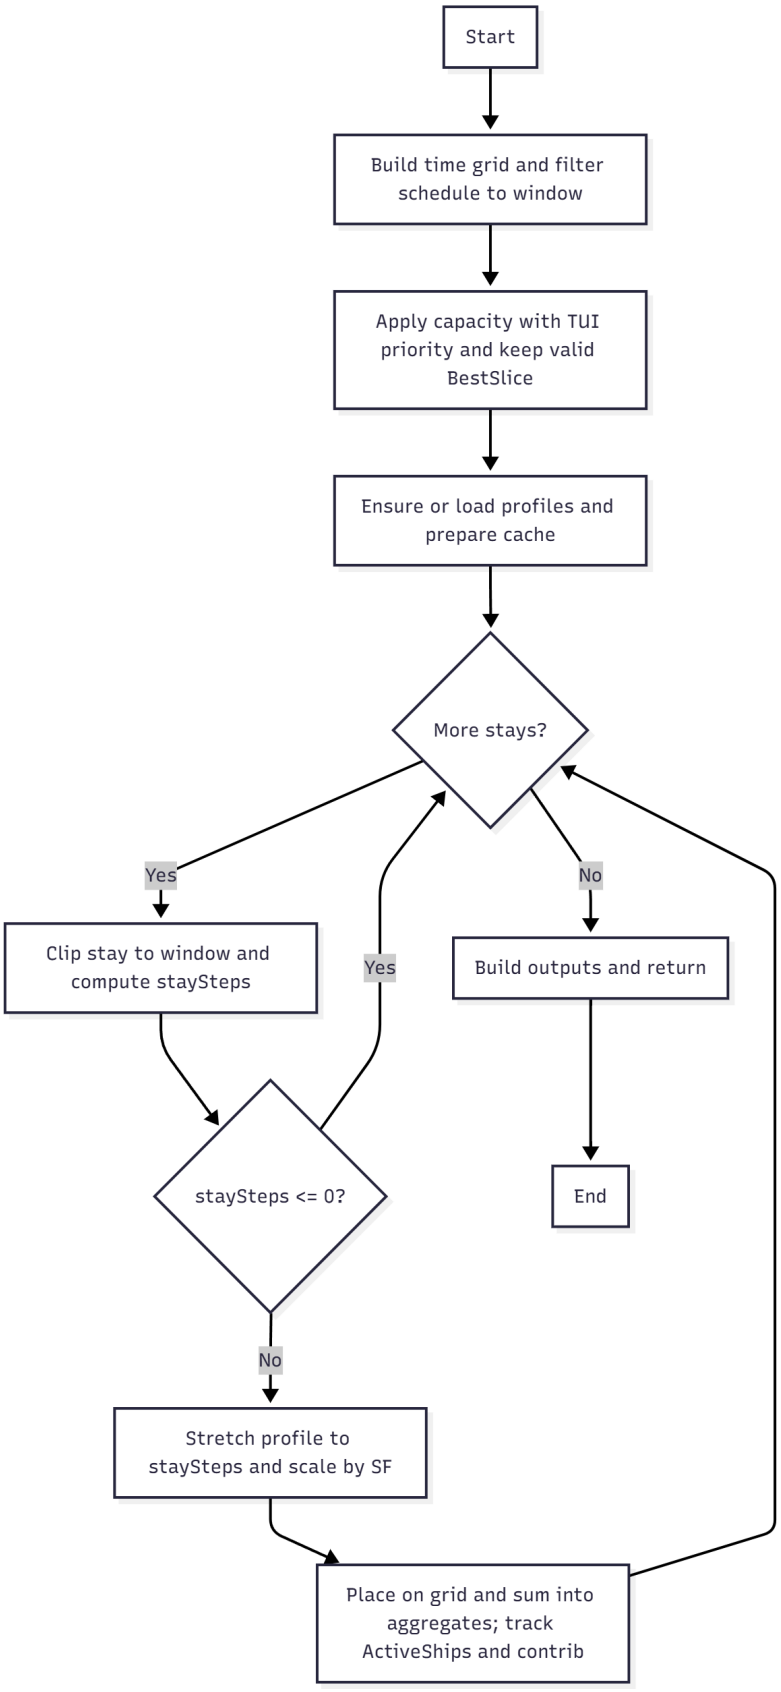

Figure. 21: Flow diagram of the port-wide aggregation layer, involving filtering, capacity/TUI, stretch–scale, and aggregation.

5 Results and Discussion

This chapter presents the results obtained from the different techniques and models defined in the previous sections, to compare their performance and analyze the energy behavior observed in the simulation of the Port of Funchal. The discussion addresses the results from both a visual and a numerical standpoint, examining how each model captures the dynamics of energy consumption and relating these outcomes to their corresponding performance metrics and measures of uncertainty. In addition, the discussion extends to the general energy demand of the port, where visual components provide further insight into consumption patterns.

5.1 Evaluation Metrics and Model Comparison

To compare how the different models capture the dynamics of port energy demand, we show a zoomed segment of the annual port-wide series for November 16–18. This is a reference profile (not raw measurements): it is constructed by concatenating and scaling the reference cruise ship slices according to the Port of Funchal schedule. The highlighted intervals mark days with three simultaneous ships at berth, the operating point where capacity constraints and scaling effects are most visible. The ship names shown inside each highlighted box correspond to the vessels actually alongside at that moment (for instance, Nov 16 at 15:59 — Caribbean Princess, Mein Schiff 5, and Voyager of the Seas). By contrast, the list shown in the bottom right panel enumerates the full set of vessels considered in the simulation, not only those present during the highlighted window. Figure 22 provides this baseline against which the models are compared.

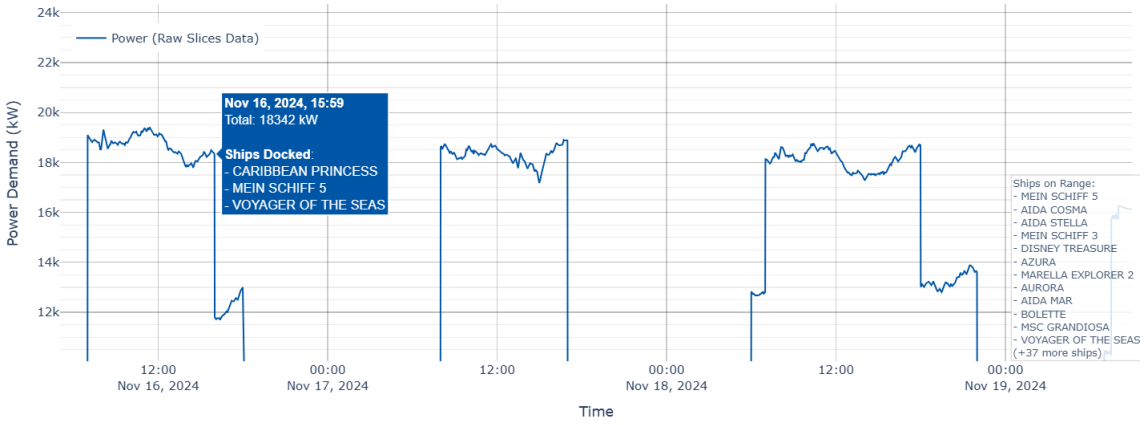

Figure. 22: Zoomed sample (Nov 16–18) of the annual port-wide reference profile (slices aligned to schedule).

The comparison of models for profile generation considers, on the one hand, their proximity to the observed values, namely, the ability to reproduce the level and trend of the profile shown in Fig. 22, and, the quality with which uncertainty is captured, assessed through the coverage and width of the intervals. These aspects are quantified by the previously defined metrics, whose consolidated values for the annual simulation can be observed in Figure 23, which represents a table where conditional color formatting was applied to each column. In this example, green means better performance and red worse, with colors linearly scaled between the column specific minimum and maximum values.

Model	RMSE	MAE	R2	KL	PICP (%)	NMPIW	CWC
Linear Regression	106,599	84,028	0,244	14,796	95,5	0,053	0,161
Ridge Regression	71,416	56,083	0,582	10,228	94,6	0,036	0,189
Lasso Regression	83,265	64,543	0,483	13,665	94,8	0,041	0,204
ElasticNet Regression	74,777	59,852	0,530	11,908	94,8	0,037	0,168
Bayesian Ridge	48,974	38,878	0,805	3,057	95,2	0,024	0,111
SVR	29,780	21,464	0,929	1,842	94,4	0,031	1,384
Kernel Regression	43,323	33,415	0,857	6,067	93,3	0,021	0,254
Decision Tree	7,803	5,396	0,996	0,829	100,0	0,024	0,024
Random Forest	8,072	5,680	0,995	0,877	99,8	0,023	0,023
KNN	7,686	5,402	0,996	0,665	100,0	0,024	0,024
Quantile Regression	87,568	65,336	0,452	6,910	93,0	0,036	0,127
Spline Regression	44,338	34,809	0,838	2,900	94,0	0,022	0,275
Isotonic Regression	86,075	62,089	0,445	16,448	93,9	0,043	0,249

Figure. 23: Comparison of regression models based on performance metrics.

As defined in subsection 3.2.2, the linear regression model is used as a baseline for comparison. The linear model is a standard OLS regression using the `LinearRegression` class from the `sklearn.linear_model` module, with default parameters (no regularization, intercept included).

Figure 24 illustrates its performance over the selected data interval, where it partially captures the upward trend but fails to adequately reproduce either the peaks or the drops in consumption.

These limitations are confirmed by the metrics. In terms of point prediction quality, the model exhibits a high RMSE (106.599) and an equally high MAE (84.028), reflecting significant prediction errors. The R^2 is very low (0.244), indicating a poor ability to explain the variability of the series. In addition, the KL reaches a high value (14.796), showing that the estimated distribution deviates substantially from the reference curve.

In terms of uncertainty quality, the model achieves the average coverage within the defined interval, with a PICP slightly above 95% (95.5%). However, it does so using rather wide bands, as reflected

in the high NMPIW value (0.053). This behavior results in a slightly elevated CWC (0.161), indicating that while the model adequately captures the variability of the data, it does so at the cost of producing broad and less informative intervals.

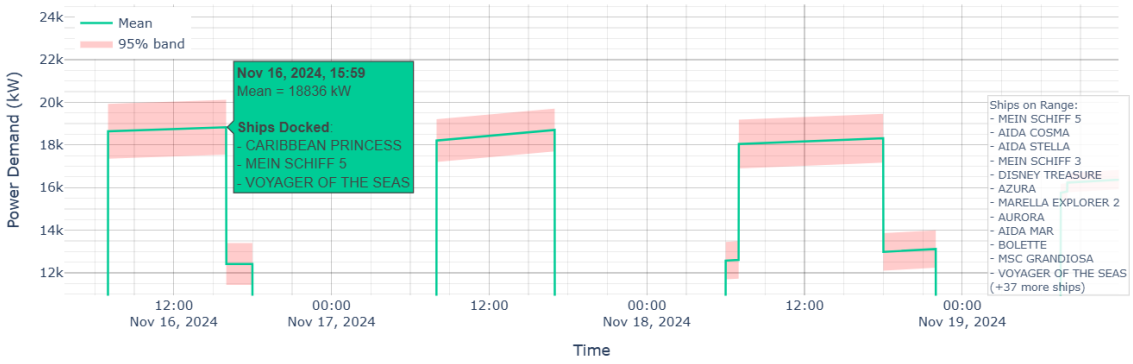

Figure. 24: Predicted shore side power demand profile using linear regression.

In the next category, the regularized models improve over plain linear regression by introducing penalties that control model complexity and help capture both the trend and the fluctuations, an effect that is most evident in the point prediction quality metrics. For consistency, all regularized pipelines use the same degree-3 polynomial feature expansion (using the `PolynomialFeatures` class from the `sklearn.preprocessing` module) and a common scaler fitted on the training split, and each model is trained via its cross validated variant to select hyperparameters on the training data (`RidgeCV` tuning α, `LassoCV` tuning α, and `ElasticNetCV` jointly tuning α and $l1\_ratio$). Within this group, Ridge, illustrated in Figure 25, offers the best balance, with RMSE = 71.416, MAE = 56.083, R^2 = 0.582, and KL = 10.228, clearly outperforming Lasso and ElasticNet. With that in mind, the results still show room for improvement.

In terms of coverage, it attains a PICP slightly below the 95% target (94.6%), which explains a CWC (0.189) higher than the interval width NMPIW (0.036). Overall, Ridge more reliably represents the stable segment of the demand curve and maintains smoother transitions than the other regularized models.

However, ElasticNet yields slightly better results regarding uncertainty quality, reaching a PICP of 94.8%, a smaller NMPIW (0.037), and a lower CWC (0.168), indicating tighter and more efficient intervals without significantly compromising coverage.

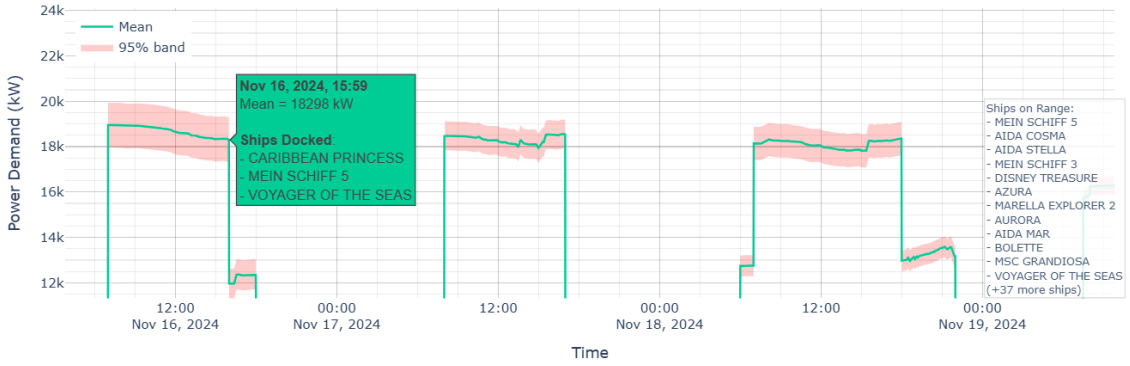

Figure. 25: Predicted shore side power demand profile using ridge regression.

The Bayesian Ridge model, illustrated in Figure 26, extends Ridge with a probabilistic treatment of coefficients, improving point accuracy and yielding well calibrated uncertainty. It was implemented using the `BayesianRidge` class from the `sklearn.preprocessing` module, with targeted parameters: `fit_intercept=False` because inputs and target are scaled; `tol=1e-6` for stricter convergence; `compute_score=True` to track the marginal likelihood during optimization; `alpha_init = 1/Var(y_train)` to start the noise precision on the correct scale (important for the initial width of the uncertainty band); and `lambda_init=1.0` as a neutral prior precision for the coefficients. With RMSE = 48.974, MAE = 38.878, and R^2 = 0.805, it outperforms standard Ridge, and its KL (3.057) shows a closer match to the reference distribution. In terms of uncertainty, the results are very positive, achieving the nominal coverage (PICP = 95.2%) with relatively thin bands (NMPIW = 0.024), leading to a low CWC (0.111) and indicating accurate, efficient intervals.

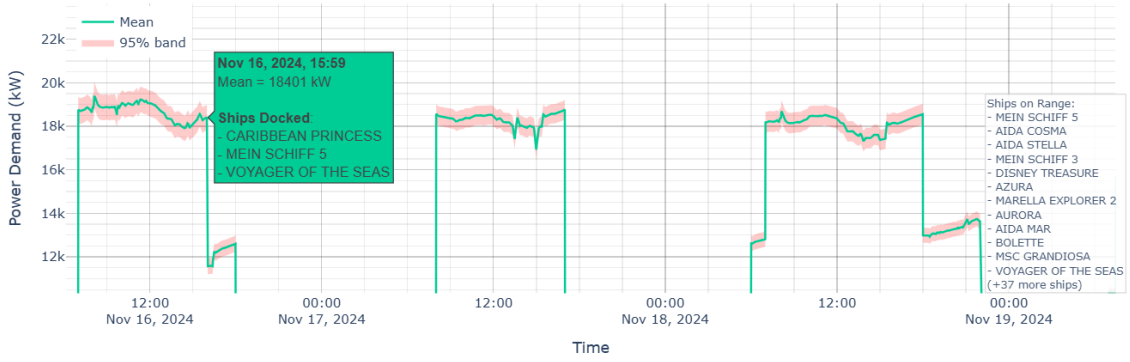

Figure. 26: Predicted shore side power demand profile using Bayesian ridge regression.

Regarding the Kernel-based models, the SVR model, observed in Figure 27, is implemented with an RBF kernel using the `SVR` class from and the following parameters: `C=10` allowing the fit to

follow sharp spikes without excessive underfitting; `gamma="scale"` adapting kernel width to feature variance for stable behavior across scales; `epsilon=0.005`, tightening the ϵ-tube and improving point accuracy at the risk of local undercoverage.

This model presents a good point prediction performance, with RMSE = 29.780, MAE = 21.464, R^2 = 0.929, and KL = 1.842. However, its uncertainty quality is less promising. On average, coverage is slightly below the 95% target (PICP = 94.4%), and despite relatively narrow intervals (NMPIW = 0.031), the combined score is high (CWC = 1.384), indicating episodes of undercoverage even with thin bands. This is consistent with jackknife-based intervals, which are locally adaptive, wider in volatile periods, and tighter in stable ones. As Figure 27 shows, a high spike in the uncertainty band indicates that occasional very wide bands (and large penalties) can occur even when the global average is moderate.

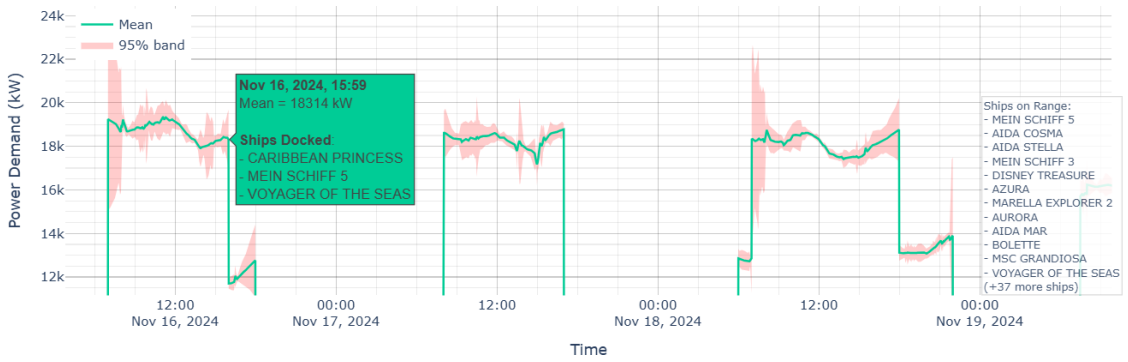

Figure. 27: Predicted shore side power demand profile using SVR.

On the other hand, the Kernel Regression model also provides good point accuracy, although slightly below that of SVR, but it clearly improves in terms of uncertainty, reaching a PICP of 93.3% (still below the nominal 95%), but result in tighter bands (NMPIW = 0.021) and a much lower CWC (0.254) compared to SVR, thus producing much sharper and more consistent intervals. It is implemented as a Nadaraya–Watson estimator with a Gaussian kernel using `bandwidth = 0.08`, a moderate smoothing level that reduces noise while preserving local structure.

Among the tree-based models, both the Decision Tree and the Random Forest models reached good point prediction results. Both models are implemented with the `DecisionTreeRegressor` class from the`sklearn.tree` module and the `RandomForestRegressor` class`sklearn.ensemble`

module, respectively, using default parameters; `max_depth=None` for both, and `n_estimators=100` for Random Forest.

Decision Tree records RMSE = 7.803, MAE = 5.396, and R^2 = 0.996, while Random Forest, observed in Figure 28 achieves RMSE = 8.070, MAE = 5.680, and R^2 = 0.995, both also show very low KL divergence (0.829 and 0.877, respectively), confirming distributions close to the reference. In terms of uncertainty, both models exhibit near ideal coverage with PICP = 100% for Decision Tree and 99.8% for Random Forest, both with thin and sharp bands (NMPIW = 0.024 and 0.023) and minimal penalty (CWC = 0.024 and 0.023). Although Decision Tree posts slightly better error metrics, it introduces abrupt transitions. By contrast, Random Forest reconstructs the series more smoothly and realistically, making it the more robust and reliable option overall.

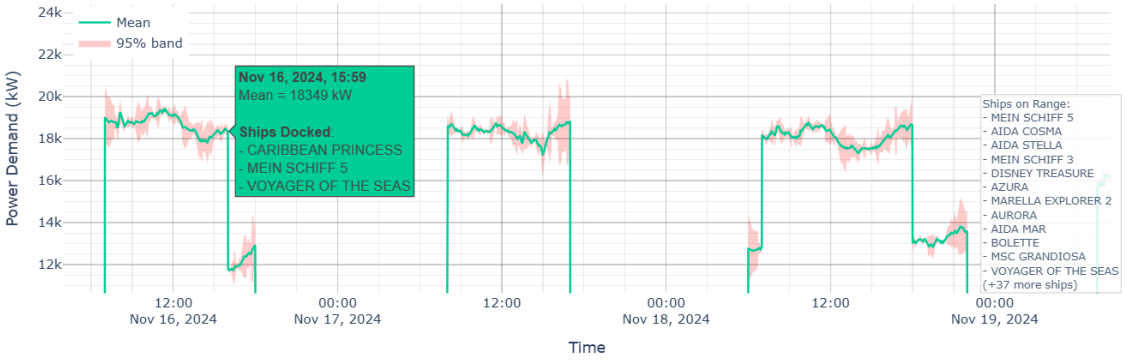

Figure. 28: Predicted shore side power demand profile using random forest.

The KNN model, observed in Figure 29, achieves the best point prediction quality among all evaluated models, with RMSE = 7.686, MAE = 5.402, R^2 = 0.996, and the lowest KL (0.665), indicating a predictive distribution that closely matches the reference. In terms of uncertainty, it attains ideal coverage (PICP = 100%) with thin and very sharp bands (NMPIW = 0.024) and no effective penalty (CWC = 0.024), yielding reliable and well calibrated intervals on par with the best tree-based approaches. It is implemented using the `KNeighborsRegressor` class from the `sklearn.neighbors` module, configured with `n_neighbors=5` (the default number of neighbors) and `weights="distance"` (giving higher weight to closer neighbors, allowing the model to capture local variations in the profile).

The Quantile Regression model, shown in Figure 30, stands out for providing explicit prediction intervals instead of a single trajectory. The model is implemented using the `QuantileRegressor`

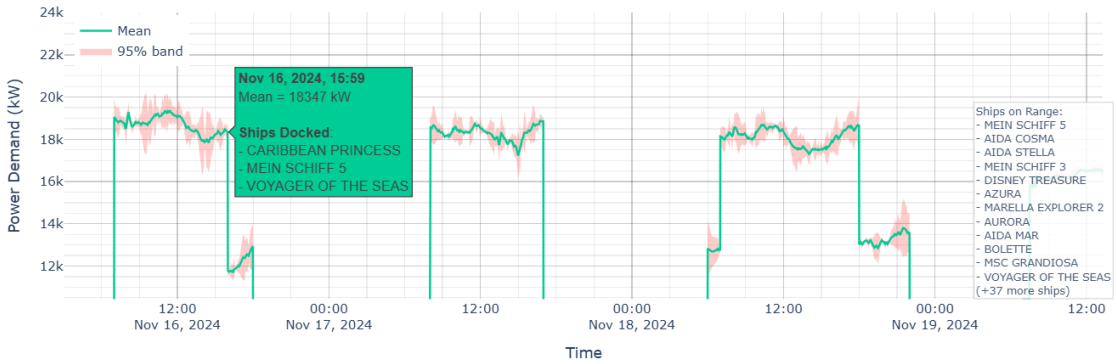

Figure. 29: Predicted shore side power demand profile using KNN.

class from the `sklearn.linear_model` module. The parameters follows the estimator's defaults (including `quantile = 0.5` and `fit_intercept = True`), except for the regularization parameter, set to `alpha = 0`, which removes the default L2 penalty and allows the model to estimate the conditional quantile directly without smoothing the coefficients.

In terms of point prediction quality, its metrics incline to worse compared to the previously analyzed models, achieving a relatively high RMSE = 87.568 and MAE = 65.336, with a slightly low R^2 = 0.452, although with a moderate distributional mismatch (KL = 6.910). However, from an uncertainty perspective, it behaves as expected and tends to the positive side, providing empirical coverage slightly below the nominal level (PICP = 93.0%), but reasonably thin bands (NMPIW = 0.036), and a favorable CWC (0.127). Overall, although its point predictions are less accurate than models such as Random Forest or KNN, its uncertainty quantification is comparatively reliable and well calibrated.

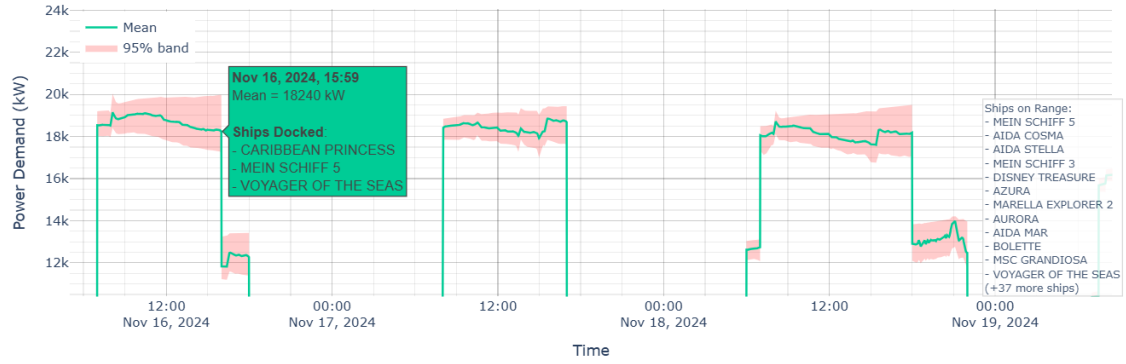

Figure. 30: Predicted shore side power demand profile using quantile regression.

Finally, the Splines model, shown in Figure 31, provides smooth and stable reconstructions of the demand curve, achieving good point prediction quality with RMSE = 44.338, MAE = 34.809, R^2 = 0.838, and a low KL = 2.900, indicating a distribution close to the reference. In terms of uncertainty, the model achieves near nominal coverage (PICP = 94.0%) with very thin bands (NMPIW = 0.022), resulting in a moderate CWC (0.275) that could reflect slight under coverage given the tightness of the intervals.

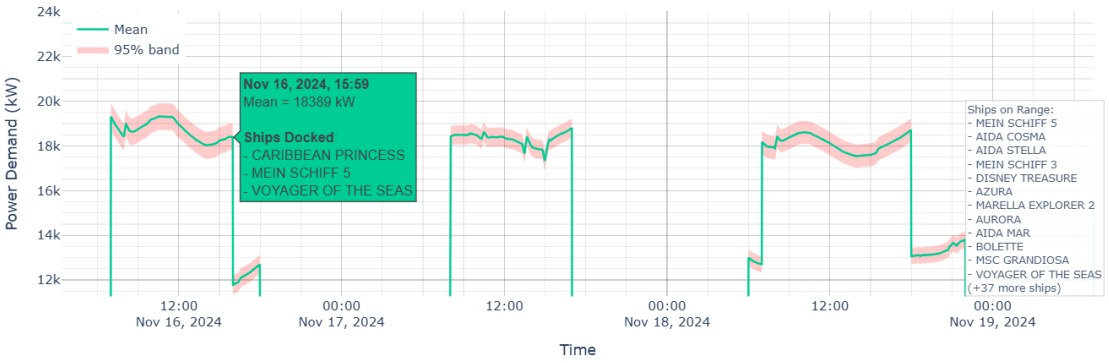

Figure. 31: Predicted shore side power demand profile using splines.

On the other hand, the isotonic regression model, shown in Fig. 32, presents weak point prediction performance, with RMSE = 86.075, MAE = 62.089, and R^2 = 0.445, and a very high KL = 16.448 indicating substantial distributional mismatch. Regarding uncertainty, it attains coverage close to (but below) nominal (PICP = 93.9%) with wider bands than Splines (NMPIW = 0.043), yet it still achieves a moderateCWC (0.249), benefiting from the broader intervals.

The modelis implemented as a spline-basis linear model using the `SplineTransformer` from the `sklearn.preprocessing` module inside a `ColumnTransformer` (from `sklearn.compose`) followed by `LinearRegression`, with `degree=3` and `n_knots=5`, cubic splines provide smooth curvature while five knots offer moderate flexibility without overfitting.

The model is implemented with the `IsotonicRegression` class from the`sklearn.isotonic` module using `increasing=True`, which enforces global monotonicity, suited to non decreasing signals, but misaligned with the alternating rises and drops in port demand, yielding less realistic fits in decreasing segments.

It is also important to examine the aggregated energy outcome predicted by each method over the entire year. Thus, as established in previous sections, the consistency of the port's aggregated

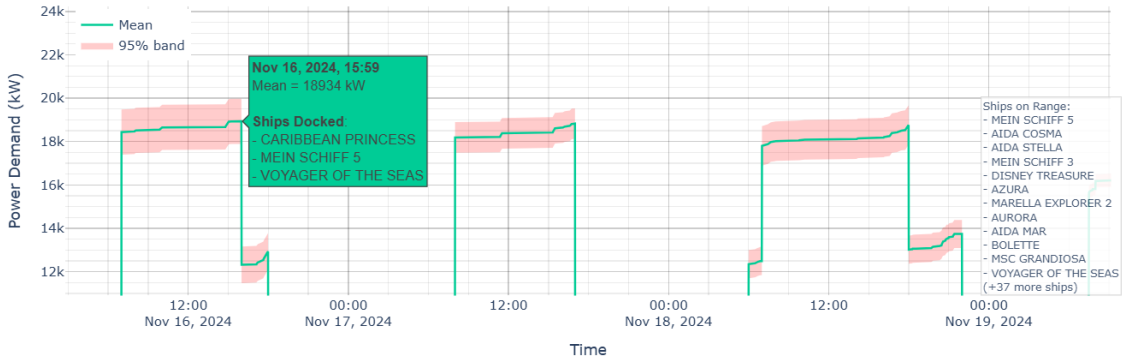

Figure. 32: Predicted shore side power demand profile using isotonic regression.

consumption is analyzed by comparing the annual totals estimated by each model and their relative deviations with respect to the ensemble mean.

Figure 33 reports the total annual energy estimated by each model (MWh), computed according to Eq. 37. All methods produce very similar values (between 37,630.9 MWh and 37,652.9 MWh), indicating that, despite methodological differences, the overall energy balance remains consistent within the modeling framework.

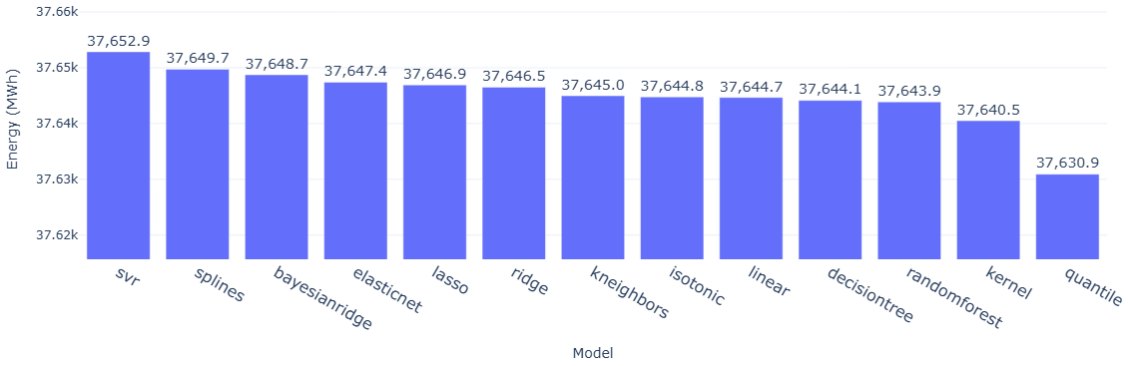

Figure. 33: Total annual shore side energy consumption estimated by each model (MWh).

Figure 34 deepens this result by showing each model's deviation from the multi-model mean (Δ MWh), computed via Eq. 39 with the reference $\bar{C}$ defined in Eq. 38. Positive values correspond to slight overestimation, and negative values to underestimation. Deviations lie within a small band (approximately ± 15 MWh), confirming that the ensemble is well centered and that no model exhibits a material annual scale bias.

In terms of performance, the KNN model stands out as the one closest to zero deviation, making it the most consistent in reproducing the ensemble mean. Other models, such as Linear Regression

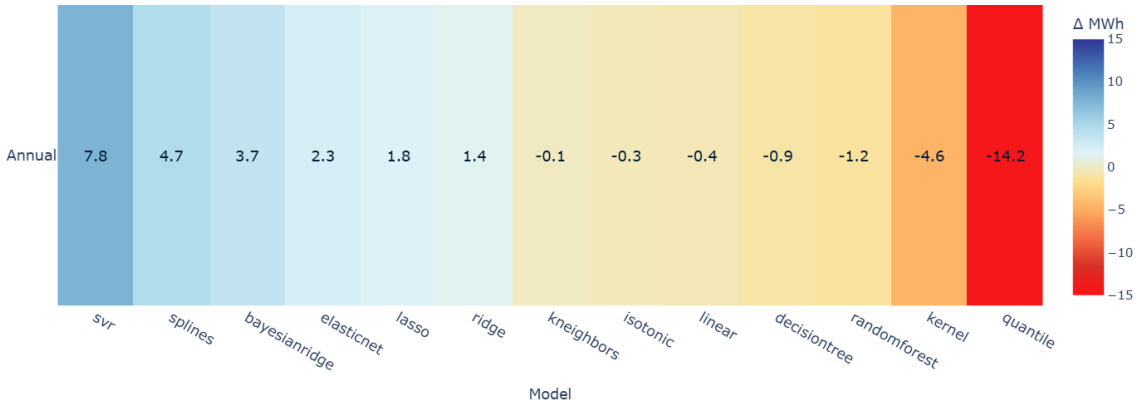

Figure. 34: Model deviation from the multi-model mean (Δ MWh).

and Isotonic Regression, also remain close to the central value, indicating balanced annual energy estimates, despite not exhibiting particularly strong point prediction or uncertainty metrics. In the other hand, the SVR and Quantile Regression models, although not far from the ensemble mean, represent the most distant cases within the group, slightly overestimating and underestimating, respectively, the total annual consumption. Overall, these results confirm that deviations across models are minimal, reinforcing the stability and robustness of the proposed modeling framework.

5.2 Port Energy Consumption Analysis

Building on the results of the previous subsection, the KNN model is adopted as the reference for the consumption analysis presented here. This choice reflects its overall superiority across the evaluated metrics, achieving low point prediction errors, well calibrated uncertainty, and a near zero annual deviation relative to the multi-model mean. The analyses that follow use KNN-based aggregated profiles, first mapped to the vessels scheduled for the Port of Funchal and adjusted by their corresponding scaling factor, to reveal consumption pattern findings in a way that can be readily interpreted by port authorities.

A key aspect is to identify seasonality patterns, since once characterized, they provide several advantages in both energy and operational terms, enabling the anticipation of high demand periods, the adjustment of energy supply planning, and the scheduling of maintenance during low activity intervals. In the case of the Port of Funchal, this seasonality is especially clear as the summer months correspond to the low season for cruise ship arrivals. This pattern is visible in Figure 35, which shows sustained activity between January and May, with multiple power peaks around 20,000 kW that reflect an intense traffic of cruise ships.

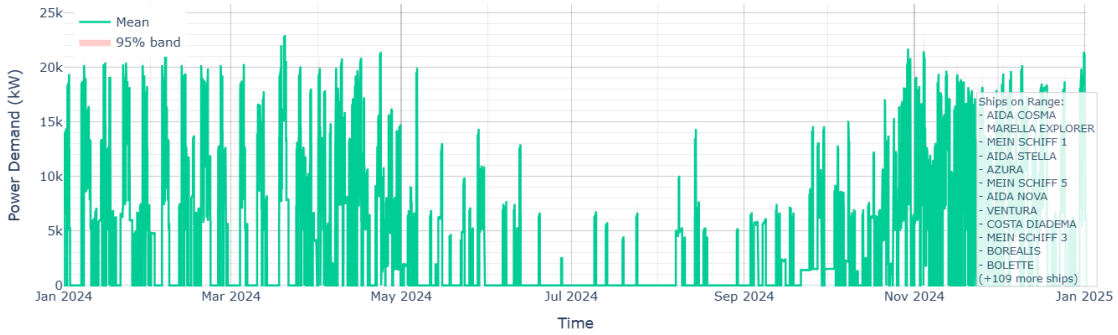

Figure. 35: Annual estimated power profile of the Port of Funchal (KNN-based).

Subsequently, between June and September, activity drops sharply, coinciding with the absence of cruise calls in the region during summer, which may be due to most cruise lines redeploying their ships to Mediterranean or Northern European itineraries.

The period of high energy demand resumes in October as autumn–winter operations restart, with power peaks returning and a denser concentration of demand events, indicating an increase in ships arrivals and overall energy consumption.

This rise in energy consumption is clearly visible in Figure 36, which shows the monthly totals. Demand is high during the first five months, approximately 3,900 to 5,100 MWh, with the first major peak in April (approximately 5,123 MWh). A sharp drop follows in summer, reaching the annual minimum in July (around 281 MWh). After September, consumption climbs again, marking the restart of autumn–winter operations, and the port reaches its highest monthly demand in November (approximately 6,062 MWh), more than 20% above the early year average.

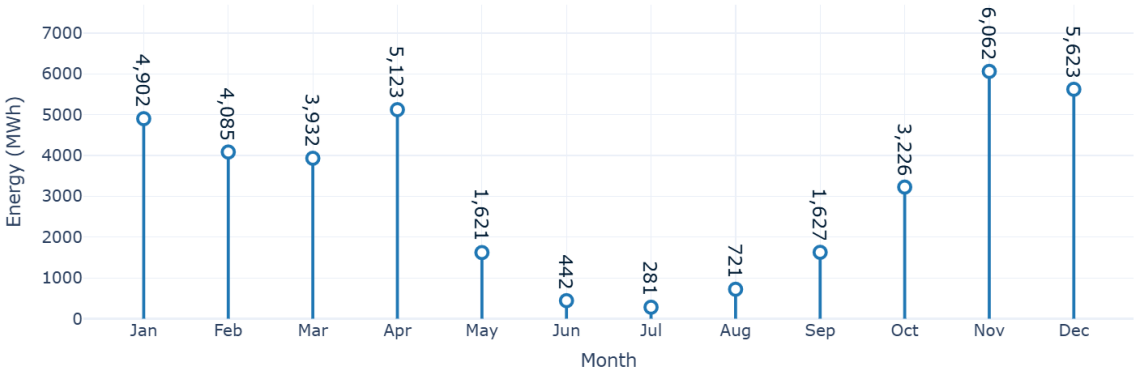

Figure. 36: Monthly estimated shore side energy consumption (KNN-based), 2024.

This pattern reveals a pronounced seasonality in the port's electrical demand profile. The concentration of high demand months between October and May suggests that the supply system should be adequately dimensioned to accommodate recurrent power peaks without compromising operational performance. On the other hand, the low activity summer period represents an optimal window for maintenance scheduling, operational adjustments, and the implementation of energy efficiency measures aimed at stabilizing the system and reducing costs before the start of the autumn–winter season.

In addition to seasonality, another important dimension is the port's intra-day behavior (or hourly load pattern), which is of utmost relevance for operational planning, as it enables port authorities to anticipate the hours with the highest demand peaks and to adjust energy efficiency and supply strategies accordingly. To explore this aspect, Figure 37 shows the average power demand by hour of day and month.

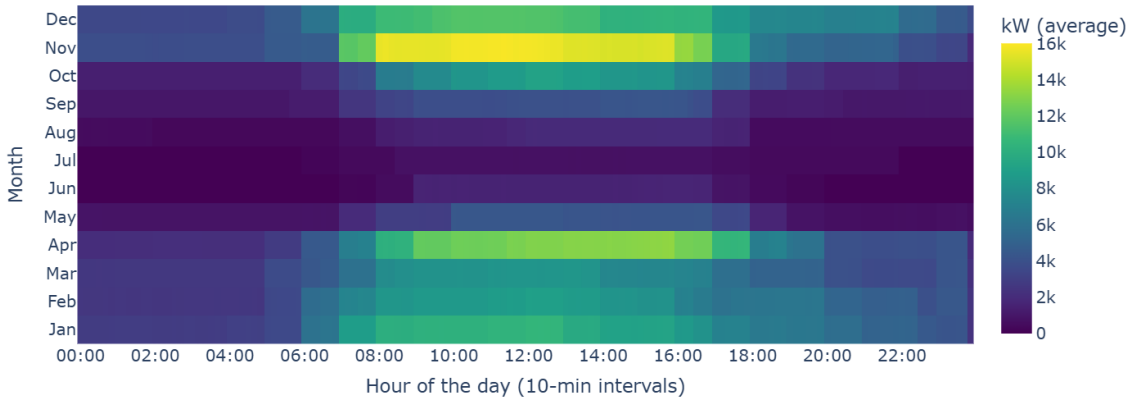

Figure. 37: Average shore side energy consumption by month and hour of day (10-min intervals aggregated to kW) of the Port of Funchal, 2024.

The heatmap reveals a clear daytime concentration of demand (roughly 08:00–16:00), with higher intensities during the October–May period and a minimal profile throughout summer, showing higher average consumption in April, November, and December, as also observed in Figure 36. This hourly pattern suggests that typical connection windows occur during daylight hours, with morning and early afternoon showing the strongest activity, while nighttime demand remains limited. From an operational perspective, this hourly pattern enables the planning of time specific actions so that they align with the hours of highest demand observed during the high demand season.

To provide a more detailed view of intra-day dynamics during the peak demand season, Figure 38 presents a zoomed view for November (days 1–22), which corresponds to the month with the highest total energy consumption, showing the distribution of shore side power demand across the day (in 10-minute bins).

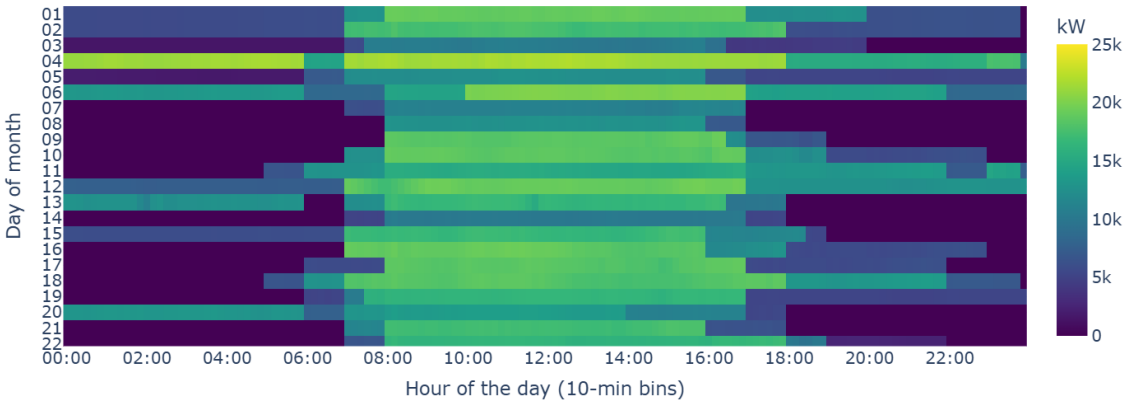

Figure. 38: Daily distribution of shore side energy consumption for November, between days 1-22 (10-min intervals aggregated to kW), 2024.

This heatmap reinforces that daily demand is predominantly concentrated between approximately 08:00 and 16:00. However, specific exceptions can occur outside this daytime window. For example, on day 4, the figure shows a pronounced high demand interval between 00:00 and 06:00, indicating that occasional overnight operations or atypical connection schedules can generate substantial loads. These deviations highlight the importance of complementing average intra-day planning with short-term monitoring to accommodate off hour peaks when they arise.

Overall, the observed seasonal and hourly dynamics provide valuable inputs for optimizing the port's electrical infrastructure. By identifying when and how demand peaks happen, the port can better align its supply side resources and explore advanced solutions in order to improve their energy management strategies.

5.3 Potential Implementations to Improve Port Energy Management

Building on the temporal demand characterization from the previous simulation, this subsection proposes strategies to reinforce the Port of Funchal's energy management. The pronounced seasonality and intra-day variability observed in the results underscore the need for adaptive solutions that can mitigate recurrent demand peaks across multiple time scales. This is both an operational

and an economic challenge, since demand charges are determined by the highest monthly peak and often represent 25 to 30% of the electricity bill [18].

To address this, a smart grid-oriented approach is proposed, focused on controlling power peaks and balancing energy flows between the grid and local resources. In this approach, peak shaving is implemented, understood as temporarily limiting the power drawn from the grid during high demand periods by supplying part of the load with local resources.

Within this approach, two complementary techniques are prioritized. First one is the deployment of a Battery Energy Storage System (BESS) that charges during low demand periods and discharges during peak hours, shifting energy in time to reduce grid draw when it is most expensive. The second one is the integration of Renewable Energy (RE) to supply part of the load during high consumption periods, thereby lowering dependence on external power and improving the port's overall economic performance.

For the first approach, the idea behind this technique is to apply peak shaving through a load sharing concept, whereby a local asset temporarily supplies part of the demand so that the grid does not bear the full load during critical periods. As illustrated in Figure 39, the BESS stores energy in low demand hours and releases it when demand rises to keep grid power within target limits.

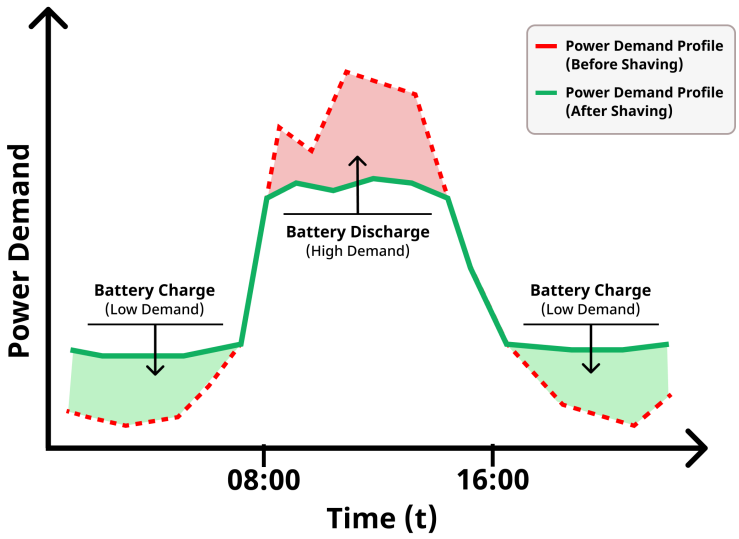

Figure. 39: Peak shaving via load sharing.

In the Port of Funchal, the operational objective is to accumulate energy before 08:00 and use it between 08:00 and 16:00 to limit the maximum power drawn from the grid. The energy management system schedules charging ahead of the peak window to reach a target state of charge while respecting power, efficiency, and safety limits.

During the 08:00–16:00 interval, it modulates discharge to cover the margin between the actual demand and the grid cap, prioritizing the most intense portions of the peak and preserving sufficient state of charge for short lived spikes. Once the peak period subsides, the system resumes charging in the next low-demand window to restore the battery for the following day, smoothing the demand profile and reducing exposure to demand charges without altering port operations.

The idea behind the second approach is to integrate local renewable generation and use the BESS to store the energy generated, so that part of the demand can be supplied locally during high consumption hours. In Madeira Island, multiple renewable sources are already established, including hydro, wind, solid waste incineration (SWI), and Photovoltaic (PV), which together represented roughly 30% of the island's electricity in 2017 [76]. This context strengthens the rationale for a renewables with BESS strategy that stores locally generated energy for later use in peak mitigation.

As illustrated in Figure 40, renewables cover concurrent demand when available, and any excess production is stored in the battery to be released later for peak mitigation.

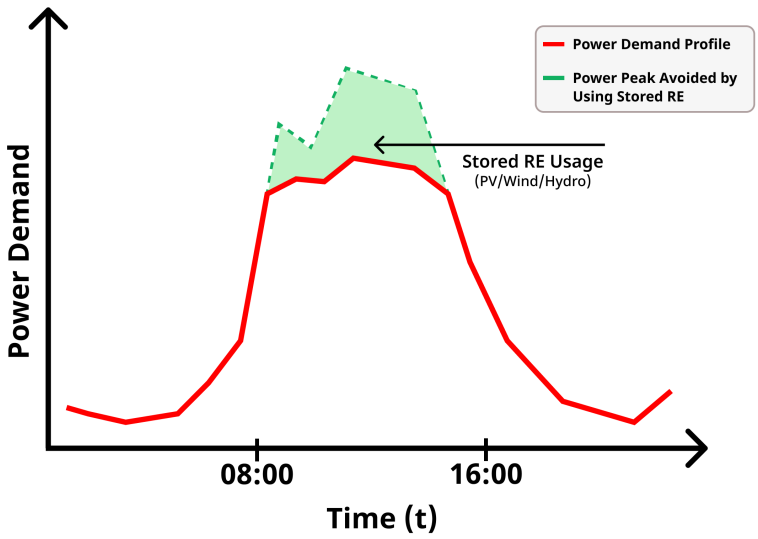

Figure. 40: Renewable supported peak avoidance.

In the Port of Funchal, the operational context favors the use of accumulated renewable energy to mitigate the peaks observed between 08:00 and 16:00. The control objective is to charge the BESS with the renewable energy available outside the peak period, so that it can be applied during that interval to keep the grid power within the desired limits.

The use of renewable energy also helps to address the seasonality observed in the simulation, mainly during the summer months when overall demand is very low and environmental conditions favor PV generation. Under these conditions, a large portion, and on some days even the entirety, of the port's daytime consumption could be supplied through solar panels together with the other renewable sources available, while the BESS shifts energy in time to cover the most demanding hours.

Based on the previously proposed peak shaving techniques, the smart grid architecture for the Port of Funchal can be represented as shown in Figure 41. At the center of the system lies an Energy Management System (EMS), responsible for coordinating power flows between the electrical grid connection, the BESS, and the local RE generation (including PV, hydro, and wind power).

An important aspect to consider is that the Port of Funchal is primarily a tourist port, so it does not involve container handling operations or other heavy industrial loads. Consequently, based on the simulation results, the system would supply three main endpoints, namely, SSE to berthed cruise ships, port buildings, and a few EV charging stations.

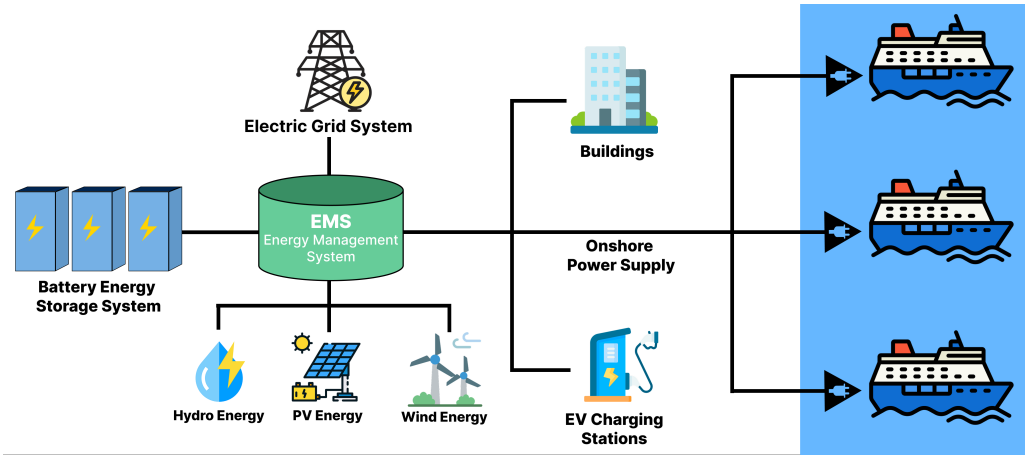

Figure. 41: Proposed smart grid architecture for the Port of Funchal.

In this way, the proposed architecture translates the defined energy strategies into an operational solution for the daytime peak from 08:00 to 16:00, improving energy efficiency and supporting the sustainable electrification of the port.

6 Conclusion

This thesis develops a data-driven methodology to estimate the hotel load demand of cruise ships calling at the Port of Funchal and to assess its implications for port energy management. The workflow begins with a data preparation stage, distinguishing ship physical specifications, real hotel energy consumption from reference ships, and the port's 2024 arrivals and departures schedule. A resampling and temporal alignment procedure is then applied to the consumption data to unify resolutions, creating per-call slices for the reference ships so that each port stay is treated separately and consistent consumption profiles can be constructed.

Following this, a two-stage regression model is implemented. First, static hotel load values are estimated from ship's physical characteristics. Then, using the reference consumption data, dynamic profiles representative of hotel energy use during port stays are created through a curve fitting process, capturing hourly demand variations, where a theoretical 95% uncertainty band is incorporated into the profiles as a safety margin for port planning and decision making, providing operationally reliable estimations.

During this stage, different types of models were compared and evaluated in terms of point prediction accuracy, uncertainty coverage, and deviation from the mean aggregated port consumption generated by each model.

Finally, the resulting dynamic profiles are assigned to ships calling at Funchal according to their estimated static hotel load, enabling the simulation of port-wide energy demand over time. The simulated results are then analyzed to identify seasonality and intra-day patterns, and these findings are briefly discussed to outline potential improvements to the port's electrical operation, with emphasis on smart grid strategies such as battery energy storage and renewable supported peak shaving.

Regarding Objective O1, the literature review focused on studies applying ML to estimate cruise ships HL, identified works such as [13] and [14], which provided a useful basis for static HL estimation from vessel specifications. However, they also exposed important gaps, including the scarcity of public high resolution datasets, the lack of consolidated methodologies for dynamic modeling, and inconsistencies in evaluation criteria. These insights helped establish the scientific foundation of this thesis and confirmed that dynamic consumption prediction remains underexplored, thereby motivating the methodological framework proposed for the Port of Funchal.

For Objective O2 a the two-stage methodology was successfully developed and implemented. In the first stage, regression models were used to estimate HL from the vessel's physical specifications, such as GT, LOA, and MPC. Two approaches were compared, potential regression and MLR, and it was shown that incorporating multiple variables significantly improves estimation accuracy over univariate models. In the second stage, curve fitting regression models were applied to real consumption profiles to generate dynamic profiles that capture the temporal variability of demand. Finally, after evaluating each model according to the defined metrics, the KNN, Random Forest, and Decision Tree approaches were considered among the best, exhibiting low prediction errors and well calibrated uncertainty.

Lastly, for Objective O3, using the KNN model was possible to simulate the SSE demand of the Port of Funchal for the entire year 2024. The simulation combined the dynamic consumption profiles generated for each vessel with the port's actual berthing schedule, considering the operational constraint of a maximum of three cruise ships berthed simultaneously and applying priority to TUI Group ships. The results revealed clear seasonal and intraday patterns.

Now addressing the research questions. For Research Question RQ1, in general terms, the models provide highly consistent annual estimates. Despite their differing approaches (linear, regularized, probabilistic, tree-based, or non-parametric), the aggregate demand predicted by the ensemble of models falls within a narrow range of 37.63-37.65 GWh, with typical deviations of about ± 15 MWh from the mean. This indicates that, for computing the port's annual aggregated SSE estimate, the choice of model is practically insensitive.

For Research Question RQ2, according to the defined evaluation metrics, the KNN model showed the best overall performance. It achieved the highest point prediction quality, with notably low RMSE and MAE values (7,686kW and 5,402kW, respectively). The model also obtained the highest explanatory capacity ($R^2 = 0.996$) and the lowest distributional divergence (KL = 0.665), indicating a very precise fit to the reference data. In terms of uncertainty, KNN also outperformed the other models, achieving perfect coverage (PICP = 100%) with thin (NMPIW = 0.024) and efficient (CWC = 0.024) uncertainty bands. Furthermore, the absolute deviation of KNN from the multi-model annual SSE estimate was only ≈ 0.1 MWh, which further supports its selection as the best option (and reference model) for the port's annual estimation.

Finally, for Research Question RQ3, the attained models allowed to produce simulation results revealed several aspects that can inform strategic decision making at the Port of Funchal. The main findings were the seasonality, as demand is higher between October and May, showing a notable decline during the summer months, and the intraday patterns, as daily demand is concentrated mainly between 08:00-16:00. Based on these results, measures can be considered to improve infrastructure planning throughout scheduled maintenance during the low demand season and reinforce operational staffing during peak demand hours to ensure the continuity and safety of the electrical supply.

In parallel, for the port's energy efficiency, two main lines are proposed. The first one use BESS for load sharing throughout the day, supplying energy during high demand hours and recharging during lower demand periods, and to integrate renewable energy sources to enable peak shaving, thereby reducing power peaks and associated costs.

6.1 Limitations

The main limitations identified during the elaboration of this study were the limited data availability and the single port scope.

First, data availability was limited, as the dataset included only the consumption records of the four ships available, and it contained a very small number of variables for each vessel. In practice, the usable features were restricted to two load categories, Hotel Power and Chillers Power, with no additional operational, environmental, or equipment level information. Second, the scope was limited to a single port (Funchal). The specific operational, and infrastructural characteristics of this setting may not be representative of other ports, meaning that the findings may not transfer directly to different contexts.

6.2 Future Work Direction

Building on the identified limitations, the most important future direction is to expand the consumption dataset. It is important to increase the number of ships and broaden the feature set, incorporating, for example, sub metered HL components (HVAC/Chillers, Galleys, Entertainment facilities, Lighting, Laundry, Elevators) alongside operational context and environmental variables.

It would also be important to assess the economic feasibility and long term impact. Future research could focus on quantifying the economic implications of the proposed energy strategies. This in-

cludes estimating the port's current annual energy expenditure and comparing it with projected costs under different scenarios involving SSE, BESS, and renewable energy integration. Such analysis would help determine whether these implementations are financially viable in the long term by identifying potential savings, payback periods, and overall economic efficiency improvements.

References

[1] United Nations Conference on Trade and Development, "Review of maritime transport 2024: Navigating maritime chokepoints," *UNCTAD*, 2024. [Online]. Available: https://unctad.org/publication/review-maritime-transport-2024

[2] Statista Research Department, "Number of cruise passengers sourced from europe from 2012 to 2024," *Statista*, 2025. [Online]. Available: https://www.statista.com/statistics/386688/number-of-cruise-passengers-in-europe/

[3] Ø. Buhaug, J. J. Corbett, Ø. Endresen, V. Eyring, J. Faber, S. Hanayama, D. S. Lee, D. Lee, H. Lindstad, A. Z. Markowska, A. Mjelde, D. Nelissen, J. Nilsen, C. Pålsson, J. J. Winebrake, W. Wu, and K. Yoshida, "Second IMO GHG study 2009," International Maritime Organization (IMO), London, UK, Tech. Rep., April 2009. [Online]. Available: https://www.imo.org

[4] United Nations, "Paris agreement," Adopted at the 21st Conference of the Parties (COP21) to the United Nations Framework Convention on Climate Change, Paris, 12 December 2015, 2015. [Online]. Available: https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement

[5] International Maritime Organization, "Fourth IMO GHG study 2020," International Maritime Organization (IMO), London, UK, Tech. Rep., 2020. [Online]. Available: https://wwwcdn.imo.org/localresources/en/OurWork/Environment/Documents/Fourth%20IMO%20GHG%20Study%202020%20-%20Full%20report%20and%20annexes.pdf

[6] Q. Chen, Y. yip Lau, M. Kanrak, X. Sun, P. Zhang, and Y.-M. Tang, "Using a bottom-up method to assess cruise ship activity impacts on emissions during 2019–2020 in china," *Heliyon*, vol. 10, p. e27101, 2024. [Online]. Available: https://doi.org/10.1016/j.heliyon.2024.e27101

[7] N. N. A. Bakar, J. M. Guerrero, J. C. Vasquez, and N. Bazmohammadi, "Electrification of onshore power systems in maritime transportation towards decarbonization of ports: A review of the cold ironing technology," *Renewable and Sustainable Energy Reviews*, vol. 178, p. 113243, 2023. [Online]. Available: https://doi.org/10.1016/j.rser.2023.113243

[8] T. Zis, R. J. North, P. Angeloudis, W. Y. Ochieng, and M. G. H. Bell, "Evaluation of cold ironing and speed reduction policies to reduce ship emissions near and at ports," *Maritime Economics & Logistics*, vol. 16, pp. 371–398, 2014. [Online]. Available: https://doi.org/10.1057/mel.2014.6

[9] E. A. Sciberras, B. Zahawi, and D. J. Atkinson, "Electrical characteristics of cold ironing energy supply for berthed ships," *Transportation Research Part D: Transport and Environment*, vol. 39, pp. 31–43, 2015. [Online]. Available: https://dx.doi.org/10.1016/j.trd.2015.05.007

[10] J. Kumar, L. Kumpulainen, and K. Kauhaniemi, "Technical design aspects of harbour area grid for shore to ship power: State of the art and future solutions," *Electrical Power and Energy Systems*, vol. 104, pp. 840–852, 2019. [Online]. Available: https://doi.org/10.1016/j.ijepes.2018.07.051

[11] C. Boertz, "Forecasting energy consumption of cruise ships using machine learning methods," Master's thesis, Norwegian University of Science and Technology (NTNU), Trondheim, Norway, 2023.

[12] L. Battistelli, T. Coppola, M. Fantauzzi, and F. Quaranta, "The environmental impact of cruise ships in the port of naples: Analysis of the pollution level and possible solutions," *Maritime Environment Research II*, pp. 945–958, 2012.

[13] S. Espinosa, P. Casals-Torrens, and M. Castells, "Hoteling cruise ship's power requirements for high voltage shore connection installations," *Journal of Maritime Research*, vol. 13, pp. 19–28, 2016.

[14] A. Micallef, M. Apap, J. Licari, and C. Caruana, "A two-stage approach combining constraint-based algorithms and gaussian process regression for estimation of cruise ship hotel loads," in *2024 IEEE 22nd Mediterranean Electrotechnical Conference (MELECON)*. IEEE, 2024, pp. 750–755. [Online]. Available: https://doi.org/10.1109/MELECON56669.2024.10608691

[15] L. Nunes, F. Mendonça, and L. Pereira, "Machine-learning applications in ports electrification – a rapid review," in *2025 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe)*. IEEE, 2025, pp. 1–5. [Online]. Available: https://doi.org/10.1109/ISGTEurope64741.2025.11305556

[16] Shift2DC Project. (2025) Shift2dc. Accessed: 6 November 2025. [Online]. Available: https://shift2dc.eu/

[17] AHEAD Project. (2024) Ahead — ai-informed holistic electric vehicles integration approaches for distribution grids. Accessed: 6 November 2025. [Online]. Available: https://horizon-ahead.eu/

[18] Ç. Iris and J. S. L. Lam, "A review of energy efficiency in ports: Operational strategies, technologies and energy management systems," *Renewable and Sustainable Energy Reviews*, vol. 112, pp. 170–182, 2019. [Online]. Available: https://doi.org/10.1016/j.rser.2019.04.069

[19] S. Vakili and A. I. Ölçer, "Techno-economic-environmental feasibility of photovoltaic, wind and hybrid electrification systems for stand-alone and grid-connected port electrification in the philippines," *Sustainable Cities and Society*, vol. 96, p. 104618, 2023. [Online]. Available: https://doi.org/10.1016/j.scs.2023.104618

[20] S. Zhu, M. M. Kinnon, J. Soukup, A. Paradise, D. Dabdub, and S. Samuelsen, "Assessment of the greenhouse gas, episodic air quality and public health benefits of fuel cell electrification of a major port complex," *Atmospheric Environment*, vol. 275, p. 118996, 2022. [Online]. Available: https://doi.org/10.1016/j.atmosenv.2022.118996

[21] T. Notteboom and H. Haralambides, "Seaports as green hydrogen hubs: advances, opportunities and challenges in europe," *Maritime Economics & Logistics*, vol. 25, pp. 1–27, 2023. [Online]. Available: https://doi.org/10.1057/s41278-023-00253-1

[22] D. Paul, V. Haddadian, B. Chavdarian, and K. Peterson, "Low-voltage shore connection power systems: Optional designs and a safety loop circuit," *IEEE Industry Applications Magazine*, vol. 24, pp. 62–68, 2018. [Online]. Available: https://doi.org/10.1109/MIAS.2017.2740448

[23] Port of Los Angeles. (2024) Alternative maritime power (amp). Accessed: 2025-07-23. [Online]. Available: https://www.portoflosangeles.org/environment/air-quality/alternative-maritime-power-(amp)

[24] Y. Khersonsky, M. Islam, and K. Peterson, "Challenges of connecting shipboard marine systems to medium voltage shoreside electrical power," in *IEEE Industry Applications Society Petroleum and Chemical Industry Conference (PCIC)*, no. PCIC-2005-16, 2005.

[25] "Iec/iso/ieee utility connections in port–part 1: High voltage shore connection (hvsc) systems–general requirements," pp. 1–68, 2012. [Online]. Available: https://doi.org/10.1109/IEEESTD.2012.6243150

[26] A. Cuculić, D. Vučetić, and V. Tomas, "High voltage shore connection," in *53rd International Symposium ELMAR-2011*. University of Rijeka, Faculty of Maritime, 2011, pp. 257–259.

[27] "Iec/ieee international standard - utility connections in port – part 1: High voltage shore connection (hvsc) systems – general requirements," *IEC/IEEE 80005-1:2019*, pp. 1–78, 2019. [Online]. Available: https://doi.org/10.1109/IEEESTD.2019.8666180

[28] R. Prenc, D. Vucetić, and A. Cuculić, "High voltage shore connection in croatia: Network configurations and formation of the connection point to the utility power grid," *Electric Power Systems Research*, vol. 157, pp. 106–117, 2018. [Online]. Available: https://doi.org/10.1016/j.epsr.2017.12.011

[29] J. J. Corbett, J. J. Winebrake, E. H. Green, P. Kasibhatla, V. Eyring, and A. Lauer, "Mortality from ship emissions: A global assessment," *Environmental Science & Technology*, vol. 41, pp. 8512–8518, 2007.

[30] A. M. Kotrikla, T. Lilas, and N. Nikitakos, "Abatement of air pollution at an aegean island port utilizing shore side electricity and renewable energy," *Marine Policy*, vol. 75, pp. 238–248, 2016. [Online]. Available: https://doi.org/10.1016/j.marpol.2016.01.026

[31] B. Stolz, M. Held, K. Meyer, and M. Schmied, "The co_2 reduction potential of shore-side electricity in europe," *Transportation Research Part D: Transport and Environment*, vol. 97, p. 102918, 2021. [Online]. Available: https://doi.org/10.1016/j.apenergy.2020.116425

[32] J. R. Pérez Osses, V. M. Palma, and C. A. Reusser, "Emissions assessment of a tanker in a chilean port using bi-directional cold ironing integrated to lng," *Sustainable Energy Technologies and Assessments*, vol. 52, p. 102135, 2022. [Online]. Available: https://doi.org/10.1016/j.seta.2022.102135

[33] J. Qi, S. Wang, and C. Peng, "Shore power management for maritime transportation: Status and perspectives," *Maritime Transport Research*, vol. 1, p. 100004, 2020. [Online]. Available: https://doi.org/10.1016/j.martra.2020.100004

[34] F. Ballini and R. Bozzo, "Air pollution from ships in ports: The socio-economic benefit of cold-ironing technology," *Research in Transportation Business & Management*, vol. 17, pp. 92–98, 2015. [Online]. Available: https://doi.org/10.1016/j.rtbm.2015.10.007

[35] K. Gore, P. Rigot-Müller, and J. Coughlan, "Cost-benefit assessment of shore side electricity: An irish perspective," *Journal of Environmental Management*, vol. 326, p. 116755, 2023. [Online]. Available: https://doi.org/10.1016/j.jenvman.2022.116755

[36] S. Filom, A. M. Amiri, and S. Razavi, "Applications of machine learning methods in port operations – a systematic literature review," *Transportation Research Part E: Logistics and Transportation Review*, vol. 161, p. 102722, 2022. [Online]. Available: https://doi.org/10.1016/j.tre.2022.102722

[37] Z. Zhang, W. Rong, Y. Liu, and Y. Yang, "Port carbon emission estimation: Principles, practices, and machine learning applications," *Transportation Research Part E: Logistics and Transportation Review*, vol. 199, p. 104159, 2025. [Online]. Available: https://doi.org/10.1016/j.tre.2025.104159

[38] F. D'Agostino, G. P. Schiapparelli, S. Dallas, D. Spathis, V. Georgiou, and J. Prousalidis, "On Estimating the Port Power Demands for Cold Ironing Applications," in *2021 IEEE Electric Ship Technologies Symposium (ESTS)*, 2021.

[39] M. Mansoursamaei, A. Asadi, M. Moradi, M. R. Mes, and K. Yakideh, "A novel hybrid machine learning model for electricity demand forecasting in offshore cold ironing," 2024, sSRN preprint, 21 pages, posted on July 17, 2024.

[40] N. N. Abu Bakar, N. Bazmohammadi, H. Çimen, T. Uyanik, J. C. Vasquez, and J. M. Guerrero, "Data-driven ship berthing forecasting for cold ironing in maritime transportation," *Applied Energy*, vol. 326, p. 119947, 2022. [Online]. Available: https://doi.org/10.1016/j.apenergy.2022.119947

[41] A. Micallef, M. Apap, J. Licari, C. Spiteri Staines, and X. Zhaoxia, "A comparative framework for evaluating machine learning models in forecasting electricity demand for port microgrids," *Energy and AI*, vol. 20, p. 100494, 2025. [Online]. Available: https://doi.org/10.1016/j.egyai.2025.100494

[42] X. Lai, X. Jin, and X. Gu, "An integrated physical-based and parameter learning method for ship energy consumption prediction under varying operating conditions," in *2017 13th IEEE Conference on Automation Science and Engineering (CASE)*, 2017, pp. 1180–1183. [Online]. Available: https://doi.org/10.1109/COASE.2017.8256263

[43] J. Yuan and V. Nian, "Ship energy consumption prediction with gaussian process metamodel," in *Energy Procedia*, vol. 152, 2018, pp. 655–660. [Online]. Available: https://doi.org/10.1016/j.egypro.2018.09.226

[44] Y. Peng, H. Liu, X. Li, J. Huang, and W. Wang, "Machine learning method for energy consumption prediction of ships in port considering green ports," *Journal of Cleaner Production*, vol. 264, p. 121564, 2020. [Online]. Available: https://doi.org/10.1016/j.jclepro.2020.121564

[45] J. Chen, H. Chen, J. Shi, M. Shi, J. Xu, H. Jiang, Y. Xiang, Y. Liu, and H. Chen, "Ais data-driven assessment of shore side electricity's emission reduction potential in china," *Transport Policy*, vol. 167, pp. 130–144, 2025. [Online]. Available: https://doi.org/10.1016/j.tranpol.2025.03.027

[46] T. Uyanık, Y. Yalman, Özcan Kalenderli, Y. Arslanoğlu, Y. Terriche, C.-L. Su, and J. M. Guerrero, "Data-driven approach for estimating power and fuel consumption of ship: A case of container vessel," *Mathematics*, vol. 10, no. 22, p. 4167, 2022. [Online]. Available: https://doi.org/10.3390/math10224167

[47] D. Sfakianakis and D. Vassalos, "Dynamic modelling of thermal energy flows in ships," in *Low Carbon Shipping Conference*. London, UK: University of Strathclyde, September 2013.

[48] A. Brækken, C. Gabrielii, and N. Nord, "Energy use and energy efficiency in cruise ship hotel systems in a nordic climate," *Energy Conversion and Management*, vol. 288, p. 117121, 2023. [Online]. Available: https://doi.org/10.1016/j.enconman.2023.117121

[49] P. Ericsson and I. Fazlagic, "Shore-side power supply: A feasibility study and a technical solution for an on-shore electrical infrastructure to supply vessels with electric power while in port," Master's thesis, Chalmers University of Technology, 2008.

[50] K. Imai, "Linear regression," 2016, pOL572 Quantitative Analysis II, Princeton University.

[51] IBM. What is ridge regression? Accessed: 2025-05-14. [Online]. Available: https://www.ibm.com/think/topics/ridge-regression

[52] GeeksforGeeks. Lasso regression. [Online]. Available: https://www.geeksforgeeks.org/machine-learning/what-is-lasso-regression/

[53] A. Jain. (2023) Elastic net regression: Combined features of l1 and l2 regularization. Accessed: 2025-08-19. [Online]. Available: https://medium.com/@abhishekjainindore24/elastic-net-regression-combined-features-of-l1-and-l2-regularization-6181a660c3a5

[54] E. Egorov. Relevance vector machine: Alpha and beta update. Accessed: 2025-08-19. [Online]. Available: https://evgenii-egorov.github.io/sk-bdl/seminar_2/slides/rvm_alfa_beta_upd.pdf

[55] Scikit-learn developers. Bayesian ridge regression. Accessed: 2025-08-19. [Online]. Available: https://scikit-learn.org/stable/auto_examples/linear_model/plot_bayesian_ridge_curvefit.html

[56] C. Gambella, B. Ghaddar, and J. Naoum-Sawaya, "Optimization problems for machine learning: A survey," *European Journal of Operational Research*, vol. 290, no. 3, pp. 807–828, 2021. [Online]. Available: https://doi.org/10.1016/j.ejor.2020.08.045

[57] E. García-Portugués, *Notes for Nonparametric Statistics*. Self-published, 2025, version 6.12.1. ISBN 978-84-09-29537-1. [Online]. Available: https://bookdown.org/egarpor/NP-UC3M/

[58] R. Tibshirani and L. Wasserman, "Nonparametric regression," 2015, course notes for Statistical Machine Learning.

[59] S. Hewage and Y. Sang, "Jackknife empirical likelihood confidence intervals for the categorical gini correlation," *Journal of Statistical Planning and Inference*, vol. 231, p. 106123, 2024. [Online]. Available: https://doi.org/10.1016/j.jspi.2023.106123

[60] Scikit-learn developers. sklearn.tree.decisiontreeregressor. Accessed: 2025-08-19. [Online]. Available: https://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeRegressor.html

[61] GeeksforGeeks. K-nearest neighbours(knn). Accessed: 2025-08-19. [Online]. Available: https://www.geeksforgeeks.org/machine-learning/k-nearest-neighbours/

[62] M. Xu, "Quantile regression model and its application research," *Academic Journal of Science and Technology*, vol. 8, pp. 172–176, 2023. [Online]. Available: https://doi.org/10.54097/vt1qpm59

[63] S. Jackson, *Splines*. Self-published via Bookdown, 2023, ch. 9, accessed: 2025-08-19. [Online]. Available: https://bookdown.org/ssjackson300/Machine-Learning-Lecture-Notes/splines.html

[64] Scikit-learn developers. Isotonic regression. Accessed: 2025-08-19. [Online]. Available: https://scikit-learn.org/stable/modules/isotonic.html

[65] (2022) World Cruise Awards 2022 – Europe's Best Cruise Destination: Madeira (Portugal). World Cruise Awards. Accessed: 2025-08-27. [Online]. Available: https://worldcruiseawards.com/winners/2022

[66] (2024) Resident population of the autonomous region of madeira in 2023 exceeded 256,000 persons, recording the largest increase in the last five years. Instituto de Estatística da Madeira. Accessed: 2025-08-27. [Online]. Available: https://estatistica.madeira.gov.pt/en/download-now-3/social-gb/popcondsoc-gb/demografia-gb/demografia-noticias-gb/4607-18-06-2024-in-2023-the-resident-population-of-the-autonomous-region-of-madeira-exceeded-256-000-persons-recording-the-largest-increase-in-the-last-five-years.html

[67] S. APRAM Administração dos Portos da Região Autónoma da Madeira. (2024) Berth reservation portal. Accessed: 2025-08-30. [Online]. Available: https://apram.pt/reserva-cais

[68] (2025) Cruisemapper – cruise ship tracker and information database. Accessed: 2025-08-30. [Online]. Available: https://www.cruisemapper.com/

[69] (2025) Myshiptracking – real time ship tracking. Accessed: 2025-08-30. [Online]. Available: https://www.myshiptracking.com/

[70] Plaquette. MinMaxScaler. [Online]. Available: https://plaquette.org/MinMaxScaler.html

[71] C. Miller, T. Portlock, D. M. Nyaga, and J. M. O'Sullivan, "A review of model evaluation metrics for machine learning in genetics and genomics," *Frontiers in Bioinformatics*, vol. 4, p. 1457619, 2024. [Online]. Available: https://doi.org/10.3389/fbinf.2024.1457619

[72] J. M. Joyce, "Kullback-leibler divergence," in *International encyclopedia of statistical science*. Springer, 2011, pp. 720–722.

[73] I. K. Bazionis and P. S. Georgilakis, "Review of deterministic and probabilistic wind power forecasting: Models, methods, and future research," *Electricity*, vol. 2, pp. 13–47, 2021. [Online]. Available: https://doi.org/10.3390/electricity2010002

[74] B. Płaczek, "A multi-agent prediction method for data sampling and transmission reduction in internet of things sensor networks," *Sensors*, vol. 23, p. 8478, 2023. [Online]. Available: https://doi.org/10.3390/s23208478

[75] Y. Shen, X. Wang, and J. Chen, "Wind power forecasting using multi-objective evolutionary algorithms for wavelet neural network-optimized prediction intervals," *Applied Sciences*, vol. 8, p. 185, 2018. [Online]. Available: https://doi.org/10.3390/app8020185

[76] M. U. Hashmi, L. Pereira, and A. Bušić, "Energy storage in madeira, portugal: co-optimizing for arbitrage, self-sufficiency, peak shaving and energy backup," in *2019 IEEE Milan PowerTech*, 2019, pp. 1–6.