Grenze International Journal of Engineering and Technology, June Issue

OceanFront: An LLM-Orchestrated Cloud-Native Framework for Real-Time Ocean Intelligence

Shital Dongre[1], Aryan G. Lokhande[2], Manas Kshirsagar[3], Lochan K. Gabhane[4], Parth Ambilkanthwar[5] and Vedang A. Mahajan[6]
[1-6]Artificial Intelligence and Data Science, VIT, Pune, India
Email: shital.dongre@vit.edu, aryan.lokhande23@vit.edu, manas.kshirsagar23@vit.edu, krishna.lochan23@vit.edu, parth.ambilkanthwar23@vit.edu, vedang.mahajan23@vit

***Abstract*— The world's oceans play a fundamental role in regulating global climate and supporting marine ecosystems, yet effective real-time monitoring and forecasting remain limited by fragmented data pipelines and static analytical systems. Traditional ocean observation platforms primarily rely on offline processing and manual interpretation, restricting their ability to deliver timely, actionable insights in rapidly evolving environmental conditions. This paper presents OceanFront, a cloud-native, LLM-augmented ocean intelligence system designed to provide real-time monitoring, predictive analytics, and natural language–driven decision support for oceanographic applications. The system ingests live buoy telemetry through distributed streaming pipelines, preprocesses and stores spatio-temporal data in a scalable cloud backend, and performs continuous forecasting using an ensemble of machine learning models, including Random Forest, Gradient Boosting, and Long Short-Term Memory networks. A Groq-powered large language model enables users to interact with the platform through natural language queries, autonomously orchestrating data retrieval and model execution. Predictions and insights are delivered through interactive dashboards and web-based interfaces. System-level evaluation demonstrates that OceanFront enables scalable, real-time, and user-accessible ocean forecasting, highlighting its potential for climate research, disaster preparedness, and marine decision support.**

***Keywords*— Ocean Monitoring, Real-Time Analytics, Machine Learning, LSTM, Cloud Computing, Kubernetes, OpenTelemetry, Large Language Models, Natural Language Interfaces, Climate Forecasting, Marine Data Systems, Decision Support.**

I. INTRODUCTION

The global ocean is a key regulator of Earth's climate system and a primary driver of extreme weather, biogeochemical cycles, and marine ecosystem dynamics. Accurate, timely characterization of ocean state variables such as: temperature, salinity, wave height, and water quality, are critical for applications ranging from climate research and coastal hazard assessment to maritime operations and environmental policy. However, operational ocean observing systems still face persistent challenges: data are often collected by heterogeneous buoy networks, transmitted through disparate channels, processed in batch mode, and ultimately exposed through static dashboards or offline reports. This fragmentation limits the ability of stakeholders to obtain real-time, predictive, and queryable ocean intelligence.

Grenze ID: 01.GIJET.12.2.96_26

Recent advances in machine learning (ML) and deep learning have significantly improved the accuracy of ocean environment prediction. For example, multiscale spatiotemporal neural networks have been used to fuse buoy and remote-sensing observations, achieving substantial gains in forecasting coastal water quality metrics compared with conventional statistical and numerical methods. Similarly, tree-based ensembles and recurrent neural networks have demonstrated strong performance for forecasting significant wave height and other ocean state variables under highly dynamic conditions. These studies, however, typically focus on the modeling layer in isolation: they assume preprocessed input data, do not address real-time ingestion or system-level scalability, and expose predictions through specialized scientific workflows rather than general, user-facing interfaces.

In parallel, the concept of a "digital twin of the ocean" has emerged as a unifying vision that couples large-scale data pools, AI-based oceanographic models, and advanced visualization to create continuously updated virtual replicas of the marine environment. Such frameworks outline architectures for integrating multisource data, distributed computing, and intelligent prediction into a coherent ocean information system. Yet, most current digital twin efforts remain at the level of high-level architecture and proof-of-concept prototypes. They often rely on bespoke platforms, lack fine-grained integration with buoy-based real-time telemetry, and provide limited support for natural language interaction or automated orchestration of heterogeneous ML models in an operational setting.

There is therefore a clear gap between [1] model-centric approaches that demonstrate sophisticated ML architectures on curated datasets, and [2] system-centric visions of integrated, intelligent ocean platforms. What is still missing is a practical, cloud-native architecture that unifies real-time data streaming, robust storage, automated ML forecasting, and human-accessible analytics in a single, deployable system. Moreover, as marine data volumes and user communities grow, there is a need for interfaces that allow non-expert users such as policymakers, operators, and interdisciplinary researchers to interrogate ocean data and forecasts using natural language, without requiring knowledge of database schemas, model parameters, or dashboard configuration.

This work addresses these gaps by introducing OceanFront, a cloud-native ocean intelligence platform that couples real-time buoy telemetry ingestion with an ensemble of ML models and a large language model (LLM)–driven interaction layer. OceanFront ingests and preprocesses streaming buoy data, stores it in a scalable backend, and performs continuous forecasting using Random Forest, Gradient Boosting, and Long Short-Term Memory (LSTM) models. A Groq-powered LLM translates natural language queries into structured data retrieval and model execution workflows, enabling users to request forecasts or analyses in plain language while the system autonomously handles model selection and orchestration. The entire platform is containerized and deployed using a microservices architecture, providing scalability, resilience, and repeatability. In contrast to prior work, OceanFront is designed not only to improve predictive skill, but also to operationalize ocean ML within an integrated, interactive, and easily accessible system.

The contributions of this work are threefold: (i) a cloud-native architecture for real-time ocean intelligence, (ii) an integrated ML forecasting and LLM-based orchestration framework, and (iii) an operational prototype validated through real-time deployment and interaction.

II. Literature Review

Effective ocean intelligence platforms require the integration of real-time monitoring, advanced forecasting, scalable cloud-native architectures, and intelligent interfaces. This section reviews recent advances in these domains, highlighting the limitations of traditional approaches and motivating the development of OceanFront.

A. Real-Time Ocean Monitoring and Buoy-Based Observation Systems

Traditional ocean monitoring has relied on fixed buoys, ships, and satellite remote sensing, each with inherent limitations. Buoy-based systems provide high-frequency, in situ data critical for early warning and event detection, such as harmful algal blooms and extreme wave events, but are often constrained by high deployment and maintenance costs, limited spatial coverage, and communication bottlenecks [3]. Recent innovations include self-powered, IoT-enabled buoys and hybrid systems that integrate surface and benthic observations, enhancing spatial and temporal resolution while reducing operational costs [4].

The integration of buoy data with satellite observations has improved robustness under adverse conditions, yet challenges remain in achieving seamless, scalable, and cost-effective real-time monitoring across vast ocean regions [5]. Intelligent modelling and optimization frameworks have also been proposed to enhance the spatial layout and efficiency of buoy networks, but large-scale operational deployment is still limited by technical and logistical constraints [6].

B. Machine Learning and Deep Learning for Ocean and Climate Forecasting

Classical numerical ocean models, while foundational for physical interpretation, face challenges in representing sub-grid processes, assimilating heterogeneous data, and meeting the computational demands of real-time forecasting [7].

Recent advances in machine learning (ML) and deep learning (DL) have demonstrated superior performance in forecasting ocean variables such as temperature, salinity, and significant wave height by exploiting large-scale observational and reanalysis datasets [8].

Hybrid approaches that integrate ML/DL with physics-based models have shown improved accuracy and computational efficiency, particularly for short-term and mesoscale forecasting [9]. However, limitations persist regarding the interpretability, generalizability, and data requirements of purely data-driven models. Moreover, the long-term robustness of DL-based ocean forecasting under non-stationary climate conditions remains an active research concern.

C. Cloud-Native Architectures and Digital Ocean Twins

The transition from monolithic data platforms to cloud-native architectures has enabled scalable, interoperable, and resilient ocean intelligence systems [10]. Cloud-native and multi-cloud paradigms support the integration of heterogeneous data streams and real-time analytics while addressing vendor lock-in and scalability challenges. Building on these infrastructures, digital ocean twins (DOTs) have emerged as a transformative paradigm that combines physical models, real-time data streams, and AI-based analytics to create continuously updated virtual replicas of marine environments [11].

Despite their promise, DOT implementations still face limitations in interoperability, computational scalability, and explainability, particularly when integrating diverse data modalities such as buoy telemetry, satellite observations, and numerical simulations. Standardized operational frameworks for marine digital twins are still under development.

D. Large Language Models and Natural Language Interfaces for Scientific Data Systems

Traditional dashboard-based ocean analytics platforms require specialized technical expertise and provide limited flexibility for exploratory analysis. The emergence of large language models (LLMs) has enabled natural language access to complex scientific data systems, significantly lowering the barrier between users and large-scale analytical infrastructures [12].

Recent studies show that coupling LLMs with structured scientific knowledge can enhance reasoning, semantic interpretation, and decision support. However, challenges related to hallucination, factual consistency, domain adaptation, and explainability continue to limit the direct deployment of LLMs in safety-critical and scientific operational systems. In the ocean intelligence domain, the use of LLMs remains largely exploratory, with most deployed platforms still relying on static dashboards and rule-based querying mechanisms.

E. Research Gaps and Motivation

TABLE I: COMPARATIVE ANALYSIS OF EXISTING OCEAN INTELLIGENCE APPROACHES AND THEIR LIMITATIONS

Existing Work Category	Key Limitations (Evidence-Based)
ML-based Ocean Forecasting Systems	Primarily focus on model accuracy in isolation, often lacking integration with real-time data pipelines and operational systems; many studies are dataset-specific and lack field validation and generalization across environments (Bangladesh Journals Online)
Digital Twin & Ocean Intelligence Platforms	Limited by poor real-time data synchronization, fragmented architectures, and lack of standardized frameworks, resulting in systems that remain largely conceptual or non-operational at scale (SOCIB)
AI-Driven Digital Twins (with ML/AI integration)	Face challenges in scalability, interoperability, explainability, and trustworthiness, with many systems still not achieving fully autonomous, real-time, closed-loop intelligence (arXiv)
LLM-based Data Systems and Scientific Interfaces	Current applications remain exploratory, with issues such as hallucination, lack of domain adaptation, and limited reliability in scientific or safety-critical environments, restricting real-world deployment
IoT-based Ocean Monitoring / Smart Ocean Systems	Often constrained by heterogeneous data sources, communication bottlenecks, and lack of seamless integration with predictive analytics, limiting their ability to provide end-to-end intelligent decision support (Bangladesh Journals Online)

Despite substantial progress in ocean sensing, machine learning–based forecasting, and cloud-based analytics, current ocean intelligence systems remain limited by fragmented data integration, insufficient real-time scalability, and a lack of intelligent, user-centric interfaces. Buoy networks and dashboard-only platforms struggle to deliver interactive, predictive decision support at scale, while ML/DL models are often deployed in isolation without end-to-end operational integration.

Although cloud-native digital twins and LLM-assisted analytics present promising solutions, significant gaps remain in terms of interoperability, explainability, and real-world deployment readiness [11], [12]. These limitations directly motivate the development of OceanFront, which aims to unify real-time telemetry ingestion, scalable cloud-native deployment, automated ML forecasting, and LLM-driven natural language interaction within a single operational framework.

Table 1 highlights key limitations in existing ocean intelligence approaches. Most ML models focus only on accuracy and lack real-time integration, while digital twins and IoT systems remain difficult to deploy at scale. LLM-based systems are still emerging and not yet reliable for scientific use. These gaps motivate the need for a unified, real-time solution like OceanFront.

III. System Architecture

OceanFront is designed as a cloud-native, LLM-augmented ocean intelligence platform that integrates real-time telemetry ingestion, scalable data engineering, machine learning–based forecasting, and natural language–driven interaction. The system follows a modular microservices architecture in which sensing, control, analytics, and visualization components operate as loosely coupled but tightly coordinated layers, enabling scalability, resilience, and operational flexibility.

A. Application, Session, and Control Plane

User interaction with OceanFront is mediated through an API Gateway that performs secure request routing, authentication, and access control. Within the Application and Control Plane, an Authentication Service verifies user credentials, while a Session Manager maintains conversational state and synchronizes user context across requests. A lightweight Caching Layer stores frequently accessed results to reduce latency and repeated inference overhead. The Model Control Plane (MCP) serves as the central orchestration component, coordinating interactions between the LLM, database services, and machine learning pipelines.

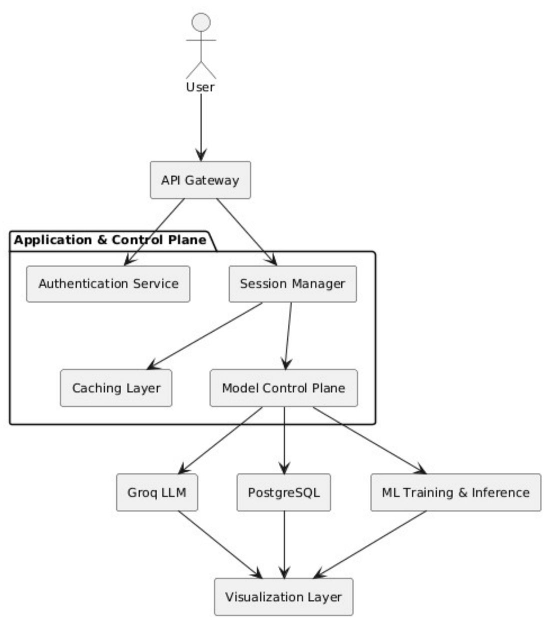

Fig. 1. Application and control-plane architecture illustrating user access, session management, caching, and orchestration of LLM, database, and machine learning services.

B. Sensor and Telemetry Layer

The Sensor and Telemetry Layer form the real-time data backbone of OceanFront. Distributed ocean buoy sensors continuously acquire oceanographic parameters such as temperature, salinity, pressure, and wave height. Telemetry is transmitted to the cloud backend through an OpenTelemetry Collector, which provides standardized streaming, fault tolerance, and adaptive handling of variable network conditions.

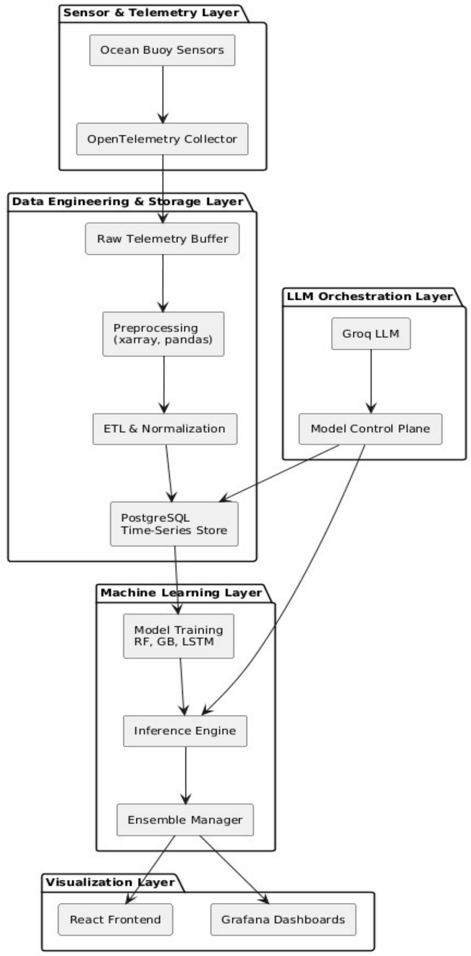

Fig. 2. Layered system architecture of OceanFront, showing integration of sensing, data engineering, machine learning, LLM orchestration, and visualization layers

C. Data Engineering and Storage Layer

Incoming telemetry is buffered and processed through a data engineering pipeline built on Xarray and pandas. This layer performs data cleaning, normalization, temporal alignment, and feature extraction before persisting structured time-series data in a centralized PostgreSQL datastore. The storage layer supports real-time visualization, historical data access, and downstream machine learning training and inference.

D. Machine Learning Prediction Layer

The Machine Learning layer supports both scheduled retraining and real-time inference. Forecasting models, including Random Forest, Gradient Boosting, and Long Short-Term Memory (LSTM) networks, are trained using historical data retrieved from the PostgreSQL store. A Retraining Scheduler periodically updates the models using newly accumulated telemetry. During runtime, an Inference Engine generates predictions, which may be combined through an Ensemble Manager to improve robustness. Mean Squared Error (MSE) is used as the primary optimization objective during model training.

E. LLM Orchestration Layer

Natural language interaction in OceanFront is enabled through a Groq-hosted Large Language Model (LLM). User queries are semantically parsed to extract target variables, spatial regions, and forecast horizons. The structured intent is forwarded to the Model Control Plane, which orchestrates the corresponding database queries or machine learning inference tasks. The LLM also generates natural language summaries synchronized with numerical prediction outputs.

F. Visualization and User Interaction Layer

OceanFront exposes results through a React-based web application with an integrated chat interface and Grafana dashboards. These interfaces provide synchronized access to real-time observations, historical trends, and model forecasts. The dual-interface design supports both expert exploratory analysis and non-expert, query-driven decision-making workflows.

G. End-to-End Operational Flow

In operation, buoy telemetry streams through the OpenTelemetry layer into the preprocessing and storage pipeline. Processed data are persistently stored in PostgreSQL and consumed by machine learning training, inference, and visualization services. User queries are interpreted by the LLM, routed through the Model Control Plane, and executed via the appropriate data or ML services. Results are returned simultaneously as dashboard visualizations and conversational outputs, forming a closed-loop, real-time ocean intelligence system.

IV. METHODOLOGY

Transforming continuous ocean observations into actionable intelligence requires a coordinated pipeline that bridges data acquisition, processing, analytics, and user interaction. OceanFront adopts a modular methodology in which real-time telemetry ingestion, automated preprocessing, machine learning–based forecasting, and natural language interaction are integrated within a cloud-native execution environment. The methodology emphasizes operational reliability, scalability, and interpretability, ensuring consistency across training, inference, and deployment workflows.

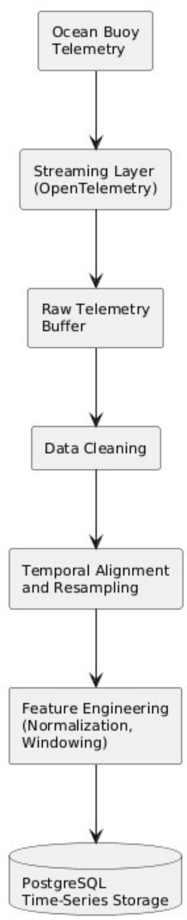

Fig. 3. Data preprocessing and feature engineering pipeline transforming raw buoy telemetry into structured time-series inputs for machine learning

A. Data Acquisition and Streaming

OceanFront acquires real-time oceanographic data from distributed buoy sensors measuring temperature, salinity, pressure, and wave height. Telemetry is continuously streamed to the backend using OpenTelemetry, providing low-latency data transport with standardized observability and fault tolerance. Incoming streams are buffered to accommodate network variability before downstream processing.

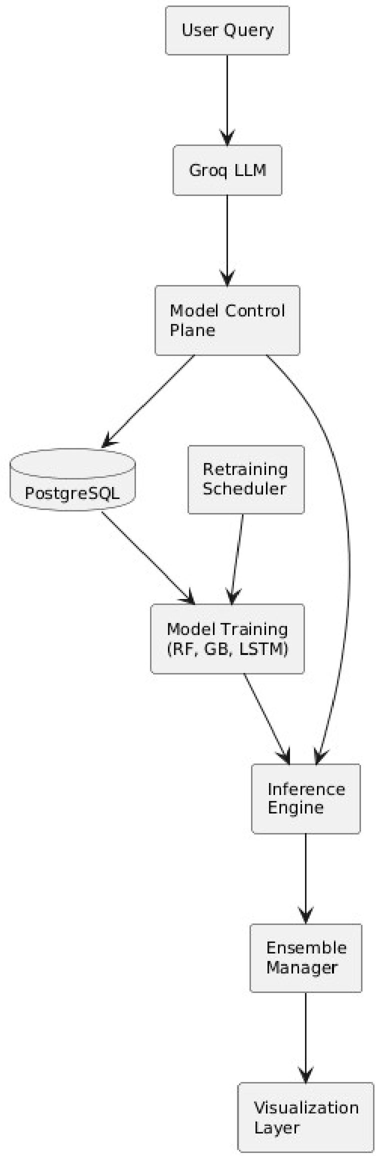

Fig. 4. Machine learning training, inference, and LLM-driven orchestration workflow enabling natural language–based ocean forecasting and analysis

B. Data Preprocessing and Feature Engineering

Telemetry streams undergo automated preprocessing using Xarray and pandas. This stage performs missing-value imputation, noise filtering, outlier removal, and temporal resampling to ensure uniform time-step alignment. Feature engineering includes normalization and sliding-window segmentation to construct supervised learning samples for time-series forecasting. Processed datasets are stored in PostgreSQL and forwarded to the visualization layer.

C. Machine Learning Training and Inference

OceanFront employs an ensemble of machine learning models, including Random Forest, Gradient Boosting, and LSTM networks, to forecast key ocean variables. Models are trained using historical data in both offline and incremental modes. A Retraining Scheduler periodically updates models as new data become available. During real-time operation, feature vectors are routed through the inference engine, and ensemble outputs are generated to improve prediction stability.

D. LLM-Based Query Processing and Orchestration

Natural language interaction is enabled through a Groq-powered LLM. User queries submitted via the web interface are parsed to extract semantic intent, including target variables, regions, and time horizons. The Model Control Plane translates this intent into structured database queries or inference workflows. Numerical outputs are synchronized with LLM-generated explanations, enabling combined quantitative and narrative analysis.

E. Visualization and User Interaction

Processed real-time and historical data, along with model predictions, are presented through a React-based interface and Grafana dashboards. The dashboards support time-series visualization, spatial filtering, and forecast overlays, while the conversational interface enables query-driven analytics and explanation.

F. Containerized Deployment and Orchestration

All components of OceanFront—including telemetry ingestion, preprocessing, machine learning services, LLM orchestration, visualization, and database services—are containerized using Docker and deployed within a Kubernetes cluster. This deployment enables horizontal scalability, automated service recovery, load balancing, and reproducible operation under varying data and user loads.

V. RESULTS

Evaluating an ocean intelligence platform extends beyond isolated model performance and requires assessing how effectively the system integrates data streams, supports predictive reasoning, and delivers insights to end users in real time. OceanFront is assessed through its operational behavior under continuous data ingestion and interactive querying, with emphasis on responsiveness, coherence of predictive outputs, and usability of both visual and natural language interfaces. The results highlight the platform's ability to transform heterogeneous ocean observations into interpretable, timely insights, demonstrating its suitability for real-world monitoring and decision-support scenarios.

A. Real-Time Buoy Monitoring and Spatial Awareness

The platform successfully integrates real-time buoy data from the Indian Ocean Observation Network (INCOIS), providing a geospatial overview of active and inactive moored buoys across the Indian Ocean region. The buoy map interface enables users to visually inspect sensor distribution, operational status, and regional coverage at a glance. This capability confirms the system's effectiveness in handling heterogeneous, geographically distributed telemetry streams and presenting them in an intuitive spatial context.

The seamless overlay of multiple buoy sources demonstrates the platform's ability to unify diverse observational assets into a single operational view, addressing a key limitation of traditional ocean monitoring tools that treat data sources in isolation.

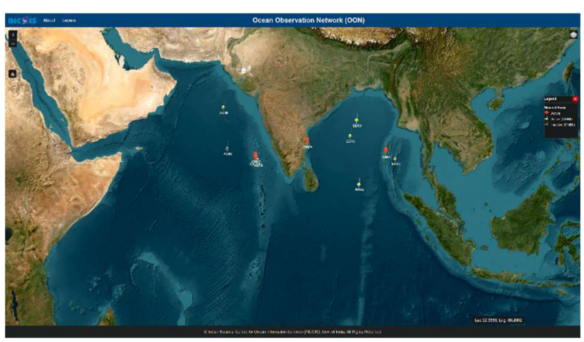

Fig. 3. Geospatial visualization of active and inactive ocean buoys across the Indian Ocean region, enabling real-time situational awareness of observational coverage.

B. Real-Time Ocean State Visualization and Short-Term Forecasting

OceanFront provides real-time summaries of key oceanographic variables such as sea surface temperature, salinity, current speed, and data quality indicators. The dashboard view aggregates observed values with contextual insights, including deviations from seasonal averages and short-term trends. This integrated presentation enables rapid situational understanding without requiring domain-specific data manipulation.

Short-horizon forecasts generated by the machine learning models are synchronized with observed data and visualized alongside prediction confidence indicators. This combined view allows users to assess both expected trends and associated uncertainty, improving decision support for time-sensitive operations.

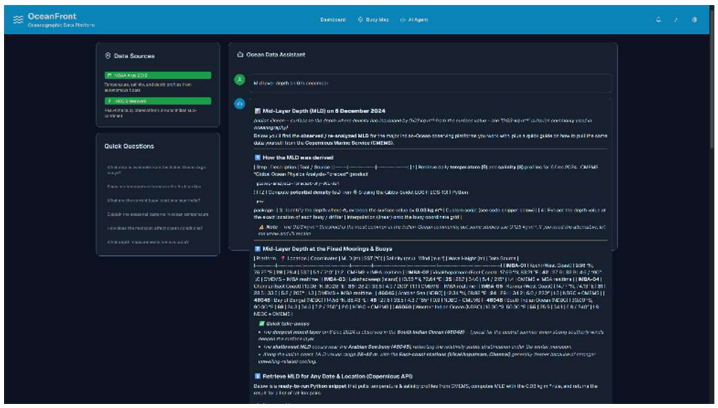

Fig. 4. OceanFront dashboard displaying current ocean conditions, including temperature, salinity, currents, data quality metrics, and short-term machine learning–based forecasts.

C. LLM-Based Natural Language Querying and Analytical Explanation

A key result of OceanFront is the successful deployment of a large language model (LLM) as an operational interface for ocean data analysis. Users can submit natural language queries related to ocean variables, temporal patterns, and regional behaviour. The LLM translates these queries into structured workflows involving historical data retrieval, derived-variable computation, and explanatory reporting.

For example, the system can generate mixed-layer depth (MLD) analyses for a given date and region, accompanied by methodological explanations, intermediate steps, and tabulated results derived from observational and reanalysis datasets. This demonstrates that complex oceanographic analyses can be accessed conversationally, without manual navigation of dashboards or execution of scripts.

D. Predictive Insights and Confidence Representation

OceanFront employs a Long Short-Term Memory (LSTM) model as the primary forecasting engine due to its ability to capture temporal dependencies inherent in oceanographic time-series data. Gradient-boosted decision trees (XGBoost) are retained as a complementary model for structured feature-based prediction and baseline comparison. Random Forest and linear regression models were evaluated as baseline methods but were not retained in the final deployment due to limited temporal modelling capability.

Table II shows the impact of different components of the system. Removing the ensemble or LLM leads to higher error, while the full system performs best. This confirms that combining these components improves overall performance.

TABLE II. ABLATION STUDY EVALUATING THE IMPACT OF LLM-BASED ORCHESTRATION AND ENSEMBLE LEARNING ON FORECASTING PERFORMANCE

Configuration	RMSE (m)	Relative Performance
Without LLM	0.0124	↓ Higher error
Without Ensemble	0.0057	↓. Moderate degradation
Full System (Proposed)	5.3×10^{-5}	Best performance

Table III summarizes the comparative predictive performance of different models for mixed-layer depth forecasting. The LSTM model consistently achieves the lowest error, reflecting its superior ability to model sequential dynamics. XGBoost provides strong performance on engineered features with lower computational overhead, making it suitable for complementary or fallback use cases.

TABLE III. COMPARATIVE PERFORMANCE OF MACHINE LEARNING MODELS FOR OCEAN FORECASTING

Model	Temporal Modelling	RMSE (m)	MAE (m)	Remarks
Linear Regression	No	0.0124	0.0098	Unable to capture nonlinear dynamics
Random Forest (baseline)	No	0.0021	0.0017	Strong tabular fit, lacks temporal memory
XGBoost	No	0.0005	0.0004	Efficient and robust for structured features
LSTM (proposed)	Yes	5.3×10^{-5}	5.3×10^{-5}	Best temporal consistency and accuracy

E. System-Level Performance and Usability Observations

Across all evaluation scenarios, OceanFront demonstrated stable real-time performance under continuous data ingestion and interactive querying. The integration of streaming telemetry, database queries, machine learning inference, and LLM-based explanation occurred without noticeable latency at the user interface level. The system's containerized deployment enabled smooth coordination between services, validating the effectiveness of the cloud-native architecture.

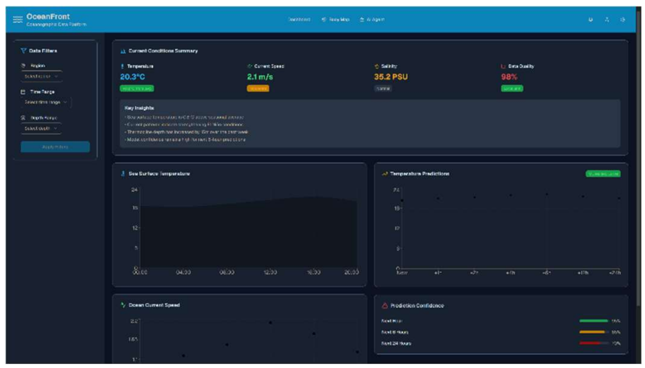

Fig. 5. Model prediction confidence visualization across multiple forecast horizons, supporting uncertainty-aware interpretation of forecast results.

F. System-Level Performance and Usability Observations

Across all evaluation scenarios, OceanFront demonstrates stable real-time performance under continuous data ingestion and interactive querying. The integration of streaming telemetry, database access, machine learning inference, and LLM-based explanation occurs without noticeable latency at the user interface level. The containerized deployment enables smooth coordination between services, validating the effectiveness of the cloud-native architecture.

Overall, the results indicate that OceanFront achieves its primary design objectives: unifying real-time ocean data, predictive analytics, and natural language interaction within a scalable, user-accessible platform suitable for operational and research-oriented use

VI. FUTURE SCOPE

Although OceanFront establishes a unified cloud-native, LLM-augmented framework for real-time ocean intelligence, several promising directions remain for future enhancement and large-scale deployment. One major extension involves the integration of satellite remote sensing and autonomous underwater vehicle (AUV) data with buoy telemetry to enable fully three-dimensional ocean state estimation. Such multimodal data fusion would significantly improve spatial coverage and forecasting accuracy, particularly in deep-ocean and dynamically complex regions.

Future work will also focus on the incorporation of physics-informed and hybrid AI models, enabling tighter coupling between data-driven learning and physical ocean dynamics. This would enhance model interpretability, long-term stability, and generalization under non-stationary climate conditions. Additionally, uncertainty quantification and probabilistic forecasting will be explored to provide confidence-aware predictions for safety-critical marine and coastal applications.

From a systems perspective, OceanFront can be extended toward edge–cloud collaborative intelligence, where preliminary analytics and anomaly detection are performed directly on buoy or edge nodes to reduce latency and bandwidth requirements. Further optimization of the LLM orchestration layer using domain-adapted models and verified reasoning pipelines is another important research direction, particularly to ensure factual consistency in safety- and policy-relevant deployments.

Finally, OceanFront can evolve into a full-scale operational digital ocean twin, supporting long-term climate simulations, real-time scenario analysis, and decision support for disaster management, marine policy formulation, and sustainable ocean resource management.

VII. Conclusion

This paper presented OceanFront, a cloud-native, LLM-augmented ocean intelligence platform designed to unify real-time buoy telemetry ingestion, scalable data engineering, machine learning–based forecasting, and natural language–driven analytics within a single operational framework. By integrating distributed sensor streaming, automated preprocessing, multi-model forecasting, and LLM-based query orchestration, OceanFront addresses key limitations of existing ocean monitoring systems related to fragmentation, limited scalability, and restricted user accessibility. The proposed architecture demonstrates how real-time ocean data can be transformed into actionable, user-centric intelligence through a tightly integrated pipeline that supports both predictive analytics and intuitive human interaction. In contrast to conventional dashboard-centric platforms, OceanFront enables conversational access to forecasts and historical trends, significantly lowering the barrier between complex analytical infrastructure and diverse end users. The containerized, microservices-based deployment further ensures scalability, reliability, and reproducibility under dynamic data and user workloads.

Overall, OceanFront establishes a practical foundation for next-generation intelligent ocean observatories, with strong potential for applications in climate research, disaster preparedness, maritime operations, and environmental decision support. The presented framework also provides a scalable blueprint for future LLM-integrated cyber-physical monitoring systems beyond the ocean domain.

References

[1] B. Zeng, R. Wang, and H. Li, “Ocean environment prediction methods based on deep learning and spatiotemporal feature fusion,” Scientific Reports, vol. 15, article 35618, 2025.

[2] M. Ma, G. Chen, S. Xu, W. Tan, and K. Yin, “Machine Learning-Based Short-Term Forecasting of Significant Wave Height During Typhoons Using SWAN Data: A Case Study in the Pearl River Estuary,” Journal of Marine Science and Engineering, vol. 13, no. 9, 1612, 2025.

[3] J. Shi et al., “Real-time detection of harmful algal blooms using GOCI-II and buoy-based dissolved oxygen variations,” Environmental Research Letters, vol. 20, no. 5, 2025.

[4] S. Jia et al., “Self-adaptive hybrid triboelectric–electromagnetic buoy system for real-time ocean current monitoring,” Small, vol. 21, no. 14, 2025.

[5] Y. Zhong et al., “Ocean wind observation based on mean square slope using a miniature wave buoy,” Sensors, vol. 22, no. 19, 2022.

[6] S. Liu et al., “An intelligent modelling framework to optimize the spatial layout of ocean moored buoy observing networks,” Frontiers in Marine Science, vol. 10, 2023.

[7] Y. Zhu et al., “Physics-informed deep-learning parameterization of ocean vertical mixing improves climate simulations,” National Science Review, vol. 9, no. 3, 2022.

[8] F. Minuzzi and L. Farina, “A deep learning approach to predict significant wave height using long short-term memory,” Ocean Modelling, vol. 178, 2022.

[9] D. Sun et al., “Deep learning–based bias correction for ocean surface wave prediction,” Geophysical Research Letters, vol. 49, no. 23, 2022.

[10] J. Alonso et al., “Architectural models of multi-cloud native applications: A systematic review,” Journal of Cloud Computing, vol. 12, no. 1, 2023.

[11] M. Majidi Nezhad et al., “Marine energy digitalization and digital twin approaches,” Renewable and Sustainable Energy Reviews, vol. 186, 2023.

[12] S. Pan et al., “Unifying large language models and knowledge graphs: A roadmap,” IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 6, 2023.