

Hybrid Transformer AI enables operational forecasting of turbine-scale Kuroshio power generation at optimal sites

Chau-Ron Wu^{a,b,*}, Yu-Chia Chang^a

^a Research Center for Critical Issues, Academia Sinica, Taiwan

^b Department of Earth Sciences, National Taiwan Normal University, Taiwan

ARTICLE INFO

Keywords:

Ocean current energy
Energy conversion
Site-scale forecasting
Observation-constrained machine learning
Transformer-based deep learning
Power system integration

ABSTRACT

Western boundary currents offer a persistent source of marine renewable energy, yet their deployment remains limited by the lack of accurate, site-scale forecasts that can be translated into electricity. Transformer-based artificial intelligence has transformed prediction, but its translation into site-specific ocean-current power forecasting for real-world energy applications remains strikingly limited. Here we develop an observation-constrained hybrid deep-learning (HyDL) framework that advances Kuroshio prediction toward grid-relevant, ocean-current energy applications. This study delivers four key contributions. First, we develop a hybrid Transformer-based AI framework that improves site-scale Kuroshio current prediction across multiple spatio-temporal scales. Second, by synergistically integrating global self-attention, statistical correction, and observation-driven real-time learning, HyDL achieves an average $\sim 30\%$ reduction in RMSE relative to a state-of-the-art operational system. Third, predicted currents are translated into turbine-scale power output under realistic operational constraints, and providing a direct pathway to reduce generation uncertainty and system-level cost. Finally, this work provides a scalable blueprint that integrates Kuroshio power generation, Transformer-based AI forecasting, and Taiwan's energy management system, accelerating the operational deployment of ocean-current energy in Taiwan and across major western boundary current regions worldwide.

1. Introduction

The Kuroshio, originating from the North Equatorial Current, flows northward along Luzon and Taiwan before entering the East China Sea and extending eastward south of Japan (see Fig. 1a). As one of the world's major western boundary currents, it transports large volumes of tropical heat toward mid-latitudes and plays a central role in regional climate and ocean-atmosphere coupling. Its velocity, however, exhibits pronounced variability across synoptic to multidecadal timescales, driven by tides, storms [1–4], mesoscale eddies [5–8], monsoonal forcing [9–12], climate modes [13,14], and long-term climate change [15,16]. This multiscale forcing fundamentally limits the predictability of Kuroshio currents.

A recent study [17] shows that the state-of-the-art US Navy HYCOM + NCODA system systematically underestimates Kuroshio velocities—and thus kinetic energy—by $\sim 30\%$ ($\sim 70\%$ in power) at optimal sites off southeast Taiwan. This bias is consistent with independent observational assessments [18], which demonstrate that HYCOM-derived current fields systematically underestimate ocean-current

energy densities across other major western boundary current systems, including regions off South Africa and Florida, when benchmarked against Acoustic Doppler Current Profiler (ADCP) measurements. Such persistent underestimation obscures the true magnitude of the resource and introduces a structural bias in resource assessment, thereby constraining the reliability of energy estimates and undermining the economic viability of ocean-current power. Beyond the inability of existing models to accurately assess ocean-current power potential, several fundamental obstacles continue to hinder the transition of ocean-current energy from resource evaluation to practical operational deployment. In particular, a persistent disconnect remains between ocean-dynamics forecasting and dispatch-oriented power-system applications, limiting the ability to convert ocean predictability into actionable energy value.

Motivated by recent advances in Transformer-based algorithms [19–24] and site-scale feasibility-driven Kuroshio resource screening [17,25], this study is guided by three key research questions. First, to what extent can an AI-driven framework improve the predictive skill of Kuroshio current dynamics relative to state-of-the-art operational models? Second, can Transformer-based learning systematically

* Corresponding author at: Research Center for Critical Issues, Academia Sinica, Taiwan.

E-mail address: cwu@ntnu.edu.tw (C.-R. Wu).




Fig. 1. Study region, site prioritization, and Transformer-based forecasting framework for Kuroshio power. (a) Basin-scale circulation of the western North Pacific, highlighting the Kuroshio as a persistent western boundary current linking the Pacific interior to marginal seas. The red box marks the Taiwan domain analyzed in this study. (b) Spatial distribution of a multi-criteria feasibility index around Taiwan, integrating oceanographic, bathymetric, and engineering constraints. Elevated index values delineate narrow, site-resolved Kuroshio energy corridors along the southeast and northeast coasts, identifying candidate power-plant sites including CG1 (near Chenggong) and SA2 (near Suao). (c) Schematic diagram of the Transformer-based hybrid deep learning (HyDL) framework for ocean-current power generation. The HyDL framework employs a Transformer architecture to assimilate continuous satellite-derived current fields within the black dashed domain, together with long-term drifter measurements, and supervised machine learning, to resolve the spatiotemporal structure of the Kuroshio Current. The trained model enables forward prediction of flow fields and site-scale power generation. Independent validation against moored ADCP observations at the CG1 and SA2 sites links basin-scale circulation dynamics with site-specific energy assessment.

Transformer-based Hybrid Deep Learning (HyDL) Development Pipeline




Fig. 2. Transformer-based hybrid deep learning (HyDL) development pipeline for ocean-current energy forecasting. The workflow illustrates the stepwise construction of the HyDL framework. A baseline Transformer-hybrid model is first established using minimal features and default hyperparameters to provide a reference benchmark. Model skill is subsequently enhanced through systematic hyperparameter optimization, followed by a spatiotemporal domain search to identify optimal spatial coverage and temporal windows that balance predictive accuracy and computational efficiency. Feature engineering is then applied to isolate the most informative predictors from multi-source observations. These elements are integrated during advanced fine-tuning to train high-resolution HyDL models. The final model enables time-resolved current predictions suitable for turbine-scale power estimation.

outperform state-of-the-art operational models under realistic ocean conditions? If so, how much can prediction skill surpass operational models? Third, can forecast currents be translated into turbine-scale power generation and operationally predicted at optimally selected deployment sites?

To address these questions, we develop a hybrid Transformer framework that forecasts Kuroshio velocities and converts them into turbine-scale power output at Taiwan's highest-value extraction sites (Fig. 1b), identified using a multi-criteria feasibility index [17,26]. Model skill was further benchmarked against state-of-the-art operational forecasting systems, establishing a robust performance baseline under real-world ocean conditions. The remainder of this paper is organized as follows. Section 2 describes the methodology and data sources. Section 3 presents forecasting performance and power prediction results. Section 4 discusses limitations and future research directions. Section 5 concludes the paper.

2. Methodology

Typical turbine operating depths east of Taiwan are concentrated

within $\sim 30\text{--}50$ m [17], coinciding with the oceanic mixed layer (OML). Climatological analyses across the northwestern Pacific constrain OML depths to $\sim 20\text{--}60$ m, with a representative mean of $\sim 30\text{--}50$ m [27], placing turbine-relevant layers firmly within a dynamically coherent upper-ocean regime. Within the OML, wind-driven turbulence and buoyancy forcing act to vertically redistribute momentum, suppress shear, and produce a quasi-depth-uniform velocity structure [28]. This physical expectation is borne out by extensive shipboard ADCP observations in the Kuroshio east of Taiwan [17], which show closely comparable velocity magnitude and variability between $0\text{--}30$ m (1.13 ± 0.36 m s⁻¹; $>4,000$ observations) and $30\text{--}60$ m (1.05 ± 0.35 m s⁻¹; $>7,000$ observations) within the current core. When evaluated relative to intrinsic variability, both layers converge to $\sim 1.1 \pm 0.4$ m s⁻¹, indicating no statistically significant vertical differentiation across turbine-operating depths. The convergence of independent climatological constraints and in situ observations demonstrates that upper-ocean currents exhibit strong vertical coherence within the OML. This provides a robust, observation-constrained basis for treating near-surface currents as a physically consistent proxy for flow conditions at turbine depth, thereby enabling reliable translation of ocean-current predictions into energy-relevant metrics.

We developed a Transformer-based hybrid deep-learning (HyDL) framework to predict ocean current dynamics within the core Kuroshio power corridor east of Taiwan ($120^\circ\text{--}124^\circ\text{E}$, $21.5^\circ\text{--}25.5^\circ\text{N}$). The framework couples Transformer-based deep learning with supervised machine learning to fuse satellite-derived geostrophic currents and near-surface drifter observations, enabling site-specific current velocity and turbine-scale power output forecasts at 1–7 day lead times. Model generalization is rigorously assessed using independent in situ observations from two Acoustic Doppler Current Profiler (ADCP) moorings deployed at the highest-value energy sites, CG1 and SA2 (Fig. 1b). The overall architecture and end-to-end workflow are illustrated in Fig. 1c, with methodological details and the operational implementation described in the subsections below.

2.1. Data

Geostrophic velocity fields were obtained from the Copernicus Marine Environment Monitoring Service (CMEMS), derived from multi-mission satellite altimetry using optimal interpolation of sea-level anomalies relative to a 1993–2012 climatology [29]. These products provide dynamically consistent estimates of large-scale circulation suitable for forecasting applications.

Near-surface absolute velocities were obtained from the NOAA Global Drifter Program (GDP), which forms a core component of the Global Ocean Observing System [30]. Drifter trajectories provide direct observations of ageostrophic and mesoscale flow components unresolved by satellite altimetry.

To ensure consistency across heterogeneous datasets, satellite and drifter observations were harmonized through ensemble averaging onto a common spatial grid [31–35]. This preprocessing step preserves dominant current structures while enabling direct co-location of geostrophic and in situ velocity fields for model training and validation.

2.2. Transformer-based HyDL

The HyDL framework couples a Transformer forecaster [19–24] with an observation-constrained supervised learning stage to reconstruct absolute surface velocities.

Stage 1: Transformer geostrophic forecasting.

Daily near-real-time geostrophic velocity snapshots at times $t_1 \dots t_n$ are encoded using a sequence-input layer with positional embeddings and processed through two stacked multi-head self-attention blocks followed by a regression head. This stage produces 1–7-day forecasts of geostrophic velocity fields, capturing large-scale and mesoscale circulation patterns. In this formulation, bathymetric effects are not included

Table 1

Training configuration of the Transformer-based HyDL framework.

Category	Component	Setting	Description
Model architecture	Backbone	Transformer HyDL	Hybrid framework combining attention-based forecasting and supervised correction
	Attention heads (H)	8	Optimized to balance accuracy and computational efficiency
	Embedding dimension (d_{model})	512	Ensures stable representation ($d_{\text{model}} / H \approx 64$)
	Key/query dimension (d_k)	128	Selected based on hyperparameter optimization
	Number of layers	2	Stacked multi-head self-attention blocks
Loss & optimization	Output layer	Regression head	Predicts continuous current velocity
	Loss function	Mean Squared Error (MSE)	Consistent with RMSE evaluation
	Optimizer	Adam	Adaptive gradient optimization
	Learning rate	1×10^{-3}	Tuned for stable convergence
	Batch size	32	Mini-batch training
Data preprocessing	Normalization	Z-score standardization	Mean = 0, Std = 1 (training set statistics)
	Input data	CMEMS & NOAA GDP	Geostrophic & drifter velocities
	Feature alignment	Spatial-temporal co-location	Ensures consistency across datasets
Scalability	Spatial domain	96 sites	Kuroshio corridor east of Taiwan
Training strategy	Framework	Rolling forecast training	Sequential training windows
	Forecast horizon	1–7 days	Operationally relevant timescale
	Number of forecasts	1,008	Independent rolling forecasts
Validation	Ground truth	Moored ADCP	Independent in situ observations
	Validation sites	CG1, SA2	High-value energy sites
	Metrics	RMSE, MAE, MAPE, MSE	Multi-metric evaluation

as input variables in our framework. Bathymetric effects are assumed to be implicitly captured through the observational velocity datasets used for training, which inherently reflect the integrated response of ocean circulation to underlying topographic constraints.

Stage 2: Observation-constrained reconstruction of absolute velocity.

Observed drifter velocity U is decomposed into geostrophic and ageostrophic components [32,33],

$$U = U_g + U_a \quad (2)$$

where U_g is derived from altimetric sea-surface height gradients and U_a represents wind-driven and other non-geostrophic contributions. Velocities are expressed in complex form,

$$U = u + iv, U_g = u_g + iv_g, U_a = u_a + iv_a, i = \sqrt{-1} \quad (3)$$

The ageostrophic component is defined as the residual between altimetry-derived geostrophic velocity and drifter-observed absolute flow, representing unresolved processes that cannot be captured by geostrophic balance alone, including wind-driven and submesoscale, and other non-geostrophic contributions [32,33]. To restore this missing signal, we trained a lightweight supervised machine-learning (ML) model using co-located CMEMS geostrophic currents and NOAA GDP drifter observations, enabling predicted geostrophic velocity fields to be mapped into full-field absolute currents. The ML model adopts an observation-constrained reconstruction framework based on the motivated decomposition of ocean velocity into geostrophic and ageostrophic components. The model can be expressed in the unified form $U = aU_g + b$, yielding an annual scaling of $a = 1.3$ and $b = 0.07$. Importantly, the fitted parameters exhibit a physically consistent seasonal modulation, with $a = 1.2$ and $b = 0.06$ in spring, $a = 1.0$ and $b = 0.14$ in summer, $a = 0.9$ and $b = 0.12$ in fall, and $a = 0.8$ and $b = 0.17$ in winter. Across all cases, the regressions achieve high explanatory power ($R^2 = 0.876\text{--}0.984$), indicating that geostrophic variability explains most of the seasonal current signal, whereas the supervised ML model systematically recovers the unresolved ageostrophic contribution. This representation captures the dominant first-order relationship between geostrophic and absolute flow. This hybrid design enhances the fidelity of absolute current reconstruction and improves forecast skill, effectively integrating observation-constrained statistical correction with data-driven flexibility.

2.3. Validation with moored ADCP observations

Independent validation is provided by two moored ADCP moorings

located at CG1 (southeast Taiwan) and SA2 (northeast Taiwan). Each mooring consists of an anchored ADCP with subsurface flotation and acoustic release (see Fig. 1c), recording continuous velocity profiles through the water column. ADCP time series underwent standard quality control, including removal of outliers and application of a moving-average filter to suppress high-frequency noise while preserving dominant flow variability. Model forecasts were evaluated against these in situ observations using root mean squared error (RMSE), mean absolute error (MAE), mean absolute percentage error (MAPE), and mean squared error (MSE). Validation against previously unseen ADCP observations provides a stringent test of generalization under real ocean conditions.

To ensure data independence, strict temporal separation is enforced between training and validation. For each forecast cycle, the model is trained using only historical data prior to the forecast initialization time and evaluated on future periods, preventing information leakage. The validation dataset consists of moored ADCP observations, which are not used during model training. These measurements therefore provide an independent in situ benchmark for assessing predictive performance. Drifter and ADCP observations differ fundamentally in sampling strategy (Lagrangian vs. Eulerian). As such, model evaluation reflects generalization across observational modalities rather than implicit data overlap.

2.4. HyDL optimization pipeline

The Transformer-based HyDL framework was developed through a structured, multi-stage pipeline designed to progressively enhance predictive skill while maintaining computational tractability (Fig. 2). We first established a baseline Transformer-hybrid architecture using a minimal feature set and default hyperparameters, providing a transparent benchmark against which all subsequent refinements could be objectively quantified. Model performance was then systematically improved via targeted hyperparameter optimization.

To explicitly resolve the scale dependence of ocean-current dynamics, we performed a spatiotemporal domain search in which both spatial coverage and temporal window length were varied, identifying configurations that optimally balance forecast accuracy against computational cost. Feature engineering was subsequently conducted to evaluate the predictive contribution of candidate variables drawn from multi-source observations and to isolate the most informative feature subset. The resulting optimized architecture, spatiotemporal configuration, and feature set were integrated during advanced fine-tuning to train the final HyDL models. This end-to-end workflow ultimately




Fig. 3. Field validation of operational ocean models against moored ADCP observations in the Kuroshio east of Taiwan. Upper panels show near-surface current speed (color shading; m s^{-1}) and velocity vectors from HYCOM + NCODA (left) and GOPAF (right), illustrating the structure of the Kuroshio jet, its interaction with Green Island, and its downstream evolution along the shelf break from CG1 toward SA2. Lower panels compare time series of current speed at CG1 (upper) and SA2 (lower) from June to September 2023, derived from HYCOM + NCODA (blue), GOPAF (orange), and moored ADCP observations (black).

delivers time-resolved ocean-current forecasts that can be directly translated into turbine-scale power output and energy production estimates.

To ensure reproducibility, the training configuration of the Transformer-based HyDL framework is explicitly defined as follows. Model optimization is performed using the mean squared error (MSE) loss function, consistent with RMSE-based evaluation metrics. All input features are standardized using z-score normalization computed from the training dataset, and the same scaling is applied to validation and test data. Training is conducted using the Adam optimizer with a learning rate of 1×10^{-3} and a mini-batch size of 32. The model is trained under a rolling forecast framework, in which each training cycle uses only historical data prior to the forecast initialization time, ensuring strict temporal causality and preventing information leakage. Hyperparameters, including attention head number and embedding dimension, are selected through systematic optimization as described in Section 3.1. A summary of all training parameters is provided in Table 1.

3. Results and discussion

3.1. Skill-cost frontier design

To determine an attention configuration that maximizes predictive skill without unnecessary complexity, we conducted systematic hyperparameter optimization across combinations of head count (H) and key/query dimension (d_k). The results identify a clear optimum at $H = 8$ and $d_k = 128$, which consistently yields the lowest RMSE. This configuration achieves the best balance between expressive capacity and computational efficiency, indicating that predictive improvements arise from resolving the essential dynamical degrees of freedom of the Kuroshio system rather than from indiscriminate architectural scaling.

This optimal regime is consistent with established insights from Transformer scaling studies. Multi-head self-attention enhances representational capacity by enabling parallel learning of heterogeneous dependencies [19], but systematic pruning experiments demonstrate that performance gains saturate beyond moderate head counts [36]. Although increasing both embedding dimension and head multiplicity




Fig. 4. Lead-time-dependent forecast skill and inter-model comparison for Kuroshio current prediction. (a) Forecast errors of the Transformer-based HyDL model as a function of lead time (1–7 days) at the CG1 and SA2 moorings, evaluated against ADCP observations. Panels show root mean squared error (RMSE), mean absolute error (MAE), mean absolute percentage error (MAPE), and mean squared error (MSE), demonstrating stable error growth with lead time and consistently lower errors at SA2 than at the more energetic CG1 site. (b) Comparison of RMSE, MAE, MAPE, and MSE among GOPAF, HYCOM + NCODA, LSTM-based HyDL, and Transformer-based HyDL at CG1 and SA2.

can improve accuracy in large-language and vision models [20,37], such gains come at sharply increasing computational and memory costs, and stable performance typically emerges when the ratio d_{model}/H remains near ~ 64 . Together, these results suggest that Transformer performance is governed by an optimal balance between head multiplicity and per-head dimensionality rather than by scale alone.

Beyond the identified optimum, our results reveal a clear saturation

regime in which forecast accuracy converges once the dominant spatiotemporal modes of the current are sufficiently represented. Additional model capacity yields diminishing returns, indicating an upper bound on achievable skill. This behaviour reflects a fundamental distinction between oceanic and atmospheric prediction: unlike atmospheric systems, where dense observations and strong external forcing can sustain continued gains from model scaling, ocean-current

Table 2

Performance comparison of operational and AI-based models for site-scale Kuroshio current forecasting. Model skill is evaluated using 1,008 independent 7-day forecasts against moored ADCP observations at CG1 and SA2, quantified by RMSE, MAE, MAPE and MSE. Operational prediction systems (HYCOM + NCODA and GOPAF) are benchmarked against LSTM-based and Transformer-based hybrid deep-learning (HyDL) frameworks. The Transformer-based HyDL consistently delivers the lowest errors (bold) across all metrics and both sites, approaching the observational uncertainty floor.

Station	System	RMSE	Imp.	MAE	Imp.	MAPE	Imp.	MSE	Imp.
CG1	GOPAF	0.43 m/s	−34%	0.36 m/s	−38%	34%	−3%	0.185 m ² /s ²	−81%
	HYCOM + NCODA	0.32 m/s	0%	0.26 m/s	0%	33%	0%	0.102 m ² /s ²	0%
	LSTM-based HyDL	0.24 m/s	+25%	0.21 m/s	+19%	30%	+9%	0.057 m ² /s ²	+44%
	Transformer-based HyDL	0.18 m/s	+42%	0.15 m/s	+42%	22%	+33%	0.035 m²/s²	+64%
SA2	GOPAF	0.30 m/s	−30%	0.23 m/s	−28%	29%	−71%	0.090 m ² /s ²	−70%
	HYCOM + NCODA	0.23 m/s	0%	0.18 m/s	0%	17%	0%	0.053 m ² /s ²	0%
	LSTM-based HyDL	0.22 m/s	+4%	0.19 m/s	−5%	18%	−6%	0.048 m ² /s ²	+9%
	Transformer-based HyDL	0.17 m/s	+26%	0.14 m/s	+22%	13%	+23%	0.028 m²/s²	+47%

forecasting is intrinsically constrained by unresolved mesoscale dynamics, sparse in situ measurements, and limited observational bandwidth to constrain ageostrophic variability. As a result, model performance is ultimately limited by the intrinsic information content of the ocean system rather than by network expressiveness.

Feature-ablation experiments further corroborate this interpretation. Satellite-derived geostrophic currents capture the majority of the Kuroshio's kinetic energy and large-scale structure, while drifter-derived velocities provide the ageostrophic and mesoscale corrections essential for site-scale prediction. Together, these two inputs explain nearly all physically learnable variance in the target signal. In contrast, adding wind forcing or sea surface height yields only marginal improvement owing to strong collinearity with existing predictors, indicating a feature-saturation regime in which additional inputs increase redundancy rather than predictive insight.

Taken together, these results demonstrate that information efficiency, rather than feature proliferation or aggressive architectural scaling, defines the practical upper bound of robust ocean-current prediction. This principle has direct implications for AI-enabled marine energy systems, where optimal forecasting performance depends on aligning model complexity with the informational constraints imposed by multiscale ocean dynamics and sparse in situ validation, rather than pursuing scale for its own sake.

3.2. Field-validated hybrid transformer forecasts

Although data-assimilative numerical models capture the basin-scale Kuroshio pathway, they exhibit systematic and station-dependent biases in the energetically active region near Green Island (see Fig. 3). Previous studies [13,38] have shown that coarse-resolution bathymetry and island topography impose a fundamental bottleneck on predictive skill in operational numerical modeling frameworks. As a result, HYCOM and GOPAF show pronounced discrepancies relative to moored ADCP measurements at both CG1 and SA2, with RMSE reaching 0.32 m s^{−1} and 0.23 m s^{−1} for HYCOM, and 0.43 m s^{−1} and 0.30 m s^{−1} for GOPAF, respectively. By contrast, the Transformer-based HyDL maintained robust skill across more than one thousand rolling forecasts (Fig. 4a). A total of 1,008 independent 7-day forecasts were generated at two stations using 72 progressively extended training windows spanning 10 June to 20 August 2023, with each model trained on observations from 1 January 2022 up to the forecast initialization date. For each forecast cycle, predicted velocities were directly compared against observed moored ADCP velocities. At CG1, where the Kuroshio jet reaches its maximum intensity, the Transformer-based HyDL reduced RMSE from 0.32 m s^{−1} (HYCOM + NCODA) to 0.18 m s^{−1}, corresponding to a 44% improvement, and lowered MAE by 42%. Even at SA2, where currents are weaker and more variable, RMSE decreased by 26% relative to HYCOM, accompanied by a 22% reduction in MAE (Table 2). Across all sites and evaluation metrics, the Transformer-based HyDL framework consistently achieved the lowest prediction errors among all models considered, demonstrating the advantage of attention-based learning for

resolving the Kuroshio's intrinsically multiscale and nonlinear dynamics.

Beyond operational comparisons, we further implemented an LSTM-based HyDL benchmark to evaluate whether these gains primarily arise from hybrid learning design or from the Transformer architecture itself. Building on the early demonstration by Bayindir [39] that LSTM networks can capture temporal variability in ocean currents, our LSTM-based HyDL provides a rigorous recurrent baseline for modern deep-learning ocean forecasting. Across all sites and metrics, the Transformer-based HyDL consistently outperformed the LSTM benchmark (Fig. 4b), indicating that attention-based learning substantially enhances generalization beyond what is achievable with conventional recurrent architectures.

To assess seasonal transferability, we conducted 73 rolling 7-day forecasts during the weakened winter Kuroshio at CG1, training on data windows spanning January 1, 2022 to end dates between December 10, 2023 and February 20, 2024. Wintertime observations at SA2 were unavailable, precluding parallel validation. Across all winter windows, the Transformer-based HyDL maintained stable performance and achieved an average 25% reduction in RMSE relative to HYCOM + NCODA. In contrast, during the summer regime—characterized by a strong, coherent jet and larger-amplitude mesoscale anomalies—the same framework reduced RMSE by 44%, revealing that AI-driven gains are strongly regime dependent. In winter, reduced signal amplitudes and elevated relative noise in both altimetric and in situ measurements constrain baseline errors, leaving less margin for improvement, whereas in summer intensified jets and enhanced eddy activity amplify resolution-driven deficiencies in numerical models.

Importantly, HyDL retained pronounced skill under event-scale disturbances, including cyclonic-eddy-induced weakening of the Kuroshio near Green Island (Fig. 5), conditions under which operational models frequently fail to reproduce localized jet attenuation and splitting. Mechanistically, the Transformer framework leverages global self-attention to learn long-range spatiotemporal dependencies that connect upstream eddy evolution to downstream velocity anomalies, thereby capturing nonlinear modulation pathways that are difficult to resolve in coarse-grid dynamical models. In dynamically complex regimes, forecast errors can approach the combined magnitude of observational uncertainty and unresolved variability, implying that residual discrepancies may partly reflect limitations in the observations rather than solely deficiencies in the model. Moored ADCP velocity measurements typically exhibit uncertainties on the order of ~ 0.01 – 0.05 m s^{−1}, with a nominal accuracy of ~ 0.3 – 0.5% of water velocity, depending on instrument configuration and environmental conditions [40,41]. Against this benchmark, the Transformer-based HyDL framework consistently achieves the lowest prediction errors across all metrics and both validation sites (Table 2). At CG1, RMSE is reduced from 0.32 to 0.18 m s^{−1} (−42%), with comparable reductions in MAE and MSE. At SA2, RMSE decreases from 0.23 to 0.17 m s^{−1} (−26%), again with consistent improvements across all error metrics. These gains are robust across more than 1,000 independent rolling forecasts, indicating




Fig. 5. Mesoscale eddy modulation of the Kuroshio and impacts on site-scale predictability. (a)–(d) Sea level anomaly (SLA; m) on 28 June, 12 July, 26 July and 9 August 2023, showing the evolution of a mesoscale eddy (E) interacting with the Kuroshio jet and modulating conditions near the CG1 mooring (triangle). (e) Current-speed time series at CG1 in August 2023 from HYCOM + NCODA and GOPAF (dashed), Transformer-based HyDL forecasts (1-, 4- and 7-day lead times; solid), and moored ADCP observations (black), highlighting skill under eddy-perturbed regimes.

Table 3

Velocity-dependent capacity factor and turbine-scale energy output for Kuroshio ocean-current power. Instantaneous capacity factor (CF) and corresponding hourly energy production (EP) are estimated for representative installed capacities (IC = 100 MW, 1 GW, and 4 GW) using the semi-empirical velocity–capacity-factor (U–CF) relationship of Chang et al. [17]. Values illustrate the nonlinear amplification of power output with increasing current speed. EP is an idealized estimate, as actual turbine-specific characteristics (e.g., cut-in/out velocities, efficiency, losses, and flow interactions) are not explicitly resolved.

Current speed, U	0.7 m/s	0.9 m/s	1.1 m/s	1.3 m/s	1.5 m/s
Capacity factor, CF	8%	19%	34%	52%	100%
EP (IC = 100 MW)	8 MWh	19 MWh	34 MWh	52 MWh	100 MWh
EP (IC = 1 GW)	80 MWh	190 MWh	340 MWh	520 MWh	1 GWh
EP (IC = 4 GW)	320 MWh	760 MWh	1.36 GWh	2.08 GWh	4 GWh




Fig. 6. Forecast evolution of Kuroshio current speed and energy metrics east of Taiwan. Spatial distributions of predicted absolute current speed (U_F ; m s^{-1}), theoretically available power (TAP; W m^{-2} , per unit cross-sectional area), and corresponding daily energy production (EP; MWh day^{-1}) for a notional 1-GW installed capacity in the Kuroshio region east of Taiwan are shown for lead times of 1, 4, and 7 days, corresponding to 16, 19, and 22 July 2023 (panels a–i). Maps highlight the persistence and downstream evolution of the Kuroshio jet relative to the Taiwanese coastline, with markers indicating the CG and SA sites. Panels j–l summarize the associated time series of area-averaged U_F , TAP, and EP as a function of forecast lead time (1–7 days) at both sites.

systematic enhancement in predictive skill rather than case-specific performance. Crucially, the magnitude of RMSE reduction ($\sim 0.06\text{--}0.14 \text{ m s}^{-1}$) substantially exceeds the typical observational uncertainty ($\sim 0.01\text{--}0.05 \text{ m s}^{-1}$), demonstrating that the improvements cannot be attributed to measurement noise. Instead, they reflect genuine gains in resolving physically relevant flow variability.

The comparison between HyDL and basin-scale operational models (e.g., HYCOM, GOPAF) should be interpreted in the context of differing model objectives. Operational systems are designed to maintain large-scale dynamical consistency and basin-wide continuity, whereas HyDL is optimized for site-specific prediction at energy-relevant locations. Accordingly, the reported performance improvements reflect enhanced local predictive skill under site-specific constraints, rather than a general

replacement of basin-scale models. HyDL is therefore best viewed as a complementary, site-focused forecasting layer that can be integrated with existing operational systems to improve prediction accuracy at targeted locations. While the present evaluation focuses on two ADCP-validated sites, these locations provide stringent independent validation in energetically critical regions. In energy-critical regions, model predictions exhibit improved agreement with independent moored ADCP observations, indicating that the learned mapping effectively captures topographic controls on the flow. Nevertheless, broader spatial validation is required to assess transferability beyond the Kuroshio corridor.

A Framework for Sustainable Energy Management




Fig. 7. Integrated AI forecasting and power engineering framework for a sustainable, net-zero energy system. The schematic illustrates a unified architecture in which Kuroshio ocean-current generation and the proposed AI framework are embedded within Taiwan's existing Energy Management System (EMS). Distributed renewable resources—including Kuroshio ocean-current, solar, and offshore wind power—are coordinated through the Distributed Renewable Energy Advanced Management System (DREAMS) and the Green Energy Estimate and Monitoring System (GEMS), and interfaced with the Regional Dispatch Center (RDC) and Distribution Dispatch Center (DDC) via SCADA-based supervisory control. In high-renewable power systems, such forecasting capability is critical, translating predictive skill into tangible system-level value through reduced reserve requirements, lower operating costs, and enhanced operational reliability—thereby accelerating Taiwan's transition toward a resilient net-zero energy future.

3.3. Bridging ocean and grid

The significance of Transformer-based HyDL extends beyond improved hydrodynamic prediction to the operational core of energy systems. The speed–capacity factor (U–CF) relationship validated in previous studies [17,42–44] directly maps forecast currents to turbine-scale energy output, yielding time-resolved and generation estimates that can be readily integrated into grid operations (Table 3). By resolving synoptic-to-mesoscale variability, the proposed framework enables more accurate reserve allocation, generation scheduling, and system cost optimization—realizing, for the first time, operationally predictive Kuroshio power generation (see Fig. 6). Power generation is presented as an idealized estimate. In practice, output depends on actual turbine-specific characteristics, including cut-in and cut-out velocities, efficiency curves, system-level losses, and device-scale flow interactions [45,46], which are not explicitly resolved in the present framework.

The theoretical available power (TAP) defines the upper bound of kinetic energy incident on a marine current turbine and scales with the cube of current speed [17,38],

$$TAP = \frac{1}{2} \rho A U^3 \quad (1)$$

where ρ is the seawater density (kg/m^3), A is the turbine's swept area (m^2), and U is the current speed (m/s). The TAP represents an upper bound on the kinetic energy available in the flow and does not account for turbine-specific conversion efficiency or operational constraints. In this study, TAP is therefore used as an idealized upper bound rather than an operationally realistic estimate. Representative forecasts (Fig. 6) demonstrate that HyDL resolves both synoptic and mesoscale structures during seasonal transitions, yielding spatially and temporally coherent estimates of TAP density and energy output. When embedded within Taiwan's advanced distribution management system (see Fig. 7), such forecasts support dynamic reallocation of generation resources, adaptive

voltage control, and demand-response coordination, thereby enhancing grid resilience. The previous study [47] has shown that accurate renewable-generation forecasting can reduce curtailment and redispatch losses while lowering reserve requirements, thereby translating predictive skill directly into economic value. In this context, forecasting evolves from a passive diagnostic tool into an active instrument of system optimization. Renewable energy systems are typically integrated with energy storage and conventional generation within grid-based energy systems. Here, we address this gap by introducing an AI forecasting framework, enabling dispatch-oriented predictability and dynamic coordination between future variable renewables, energy storage, and conventional power systems.

More broadly, our study delivers a timely advance at the intersection of artificial intelligence, ocean dynamics, and power-system operation, addressing a long-standing barrier in marine renewable energy. By rendering the Kuroshio predictable at turbine and grid timescales, HyDL elevates ocean-current energy from a persistent but uncertain resource to a system-relevant component of renewable portfolios, offering a scalable blueprint for integrating western boundary currents into future low-carbon energy systems worldwide.

4. Limitations and future work

The improved predictive capability of the proposed Transformer-based HyDL framework stems from the combined strengths of global self-attention, observation-constrained statistical correction, and real-time learning data. By exploiting long-range spatiotemporal dependencies, the Transformer learns dynamically coherent patterns across the Kuroshio system, enabling it to capture variability spanning mesoscale eddy interactions, current-path shifts, and synoptic-scale forcing—processes that are typically under-resolved by local or recurrent architectures. In parallel, the hybrid design integrates supervised learning components trained on drifter observations to capture

observation-constrained ageostrophic variability, supporting improved generalization relative to purely data-driven approaches.

Despite these advances, several limitations remain. First, this constraint is particularly evident in coastal and shallow-water regions, where satellite altimetry suffers from reduced accuracy. Second, the daily sampling frequency and coarse effective resolution of satellite-derived currents impose an intrinsic ceiling on the model's ability to resolve sub-daily and submesoscale processes, including near-inertial oscillations, internal-wave-driven variability, and rapid accelerations induced by typhoons. These unresolved dynamics may lead to underestimation of short-lived extremes that are critical for both turbine loading assessment and operational forecasting.

Future work will advance HyDL along two key directions. First, extending the framework to other major western boundary current systems—including the Gulf Stream and Agulhas Current—will rigorously assess cross-basin generalizability and establish a scalable global platform for site-specific ocean-current energy forecasting. Second, enhancing computational efficiency through model pruning, quantization, and reduced-order learning will enable real-time, large-scale deployment and accelerate integration into global decision-support and dispatch frameworks.

5. Conclusions

1. Our work establishes a hybrid Transformer AI framework that improves site-scale Kuroshio current prediction, enabling high-fidelity forecasting of multi-scale current variability that has long remained challenging for conventional approaches. This provides a new technical foundation for translating ocean observations into actionable energy intelligence.
2. The proposed Transformer-based HyDL framework attains its predictive advantage through a synergistic coupling of global self-attention, statistical correction, and observation-driven real-time learning, yielding an $\sim 30\%$ reduction in RMSE relative to the leading operational model. This performance gain underscores how observation-constrained AI can systematically overcome persistent model biases.
3. Our framework bridges the long-standing divide between oceanographic prediction and engineering-scale energy conversion. By enabling site-scale forecasts of ocean-current power and energy production, the proposed AI framework directly reduces system-level uncertainty and operational costs.
4. Ultimately, this study provides a scalable and operational blueprint to accelerate the deployment of ocean-current energy across major western boundary current regions worldwide, supporting the transition toward low-carbon power systems.

CRedit authorship contribution statement

Chau-Ron Wu: Writing – review & editing, Supervision, Software, Resources, Project administration, Data curation, Conceptualization.
Yu-Chia Chang: Writing – original draft, Visualization, Validation, Methodology.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to acknowledge the Grants from the Academia Sinica, Taiwan (AS-IAIA-114-M04; AS-IAIA-115-M04). We are grateful for the comments of editor and reviewers.

Data availability

The various data sources used in this study can be accessed and downloaded from the following locations: Argo float data and metadata from Global Data Assembly Centre (<https://doi.org/10.17882/42182>) [48], NOAA Global Drifter Program quality-controlled 6-hour interpolated data from ocean surface drifting buoys (<https://doi.org/10.25921/7ntx-z961>) [32], Ocean Surface Current Analyses Real-time (<https://doi.org/10.5067/OSCAR-25N20>), Global Ocean Gridded L4 Sea Surface Heights And Derived Variables (<https://doi.org/10.48670/moi-00149>), Global Ocean Physics Analysis and Forecast (<https://doi.org/10.48670/moi-00016>) [49], HYCOM+NCODA 1/12° Reanalysis (GLBu0.08/expt_19.1, available at: https://tds.hycom.org/thredds/catalogs/GLBu0.08/expt_19.1.html [50], and HYCOM+N-CODA Global 1/12° Analysis (GOFS 3.1, GLBy0.08, available at: https://tds.hycom.org/thredds/catalogs/GLBy0.08/expt_93.0_uv3z.html [51]).

References

- [1] Wu C-R, Chang Y-L, Oey L-Y, Chang C-W, Hsin Y-C. Air-sea interaction between tropical cyclone Nari and the Kuroshio. *Geophys Res Lett* 2008;35:L12605. <https://doi.org/10.1029/2008GL033942>.
- [2] Chang Y-C, Tseng R-S, Centurioni LR. Typhoon-induced strong surface flows in the Taiwan Strait and Pacific. *J Oceanogr* 2010;66:175–82. <https://doi.org/10.1007/s10872-010-0015-4>.
- [3] Zheng Z-W, Zheng Q, Lee C-Y, Gopalakrishnan G. Transient modulation of Kuroshio upper layer flow by directly impinging typhoon Morakot in east of Taiwan in 2009. *J Geophys Res Oceans* 2014;119:4462–78.
- [4] Kuo Y-C, Zheng Z-W, Zheng Q, Gopalakrishnan G, Lee H-Y. Typhoon-Kuroshio interaction in an air-sea coupled system: case study of typhoon Nanmadol (2011). *Ocean Model* 2018;132:130–46.
- [5] Nitani H. Beginning of the Kuroshio. In: Stommel H, Yoshida K, editors. *Kuroshio: Physical Aspects*. Univ. Tokyo Press; 1972. p. 129–163.
- [6] Zhang D, Lee TN, John WE, Liu C, Zantopp R. The Kuroshio east of Taiwan: modes of variability and relationship to interior ocean mesoscale eddies. *J Phys Oceanogr* 2001;31:1054–74.
- [7] Wang Y-L, Wu C-R. Discordant multi-decadal trend in the intensity of the Kuroshio along its path during 1993–2013. *Sci Rep* 2018;8:14633. <https://doi.org/10.1038/s41598-018-32843-y>.
- [8] Yan X, Kang D, Pang C, Zhang L, Liu H. Energetics analysis of the eddy-Kuroshio interaction east of Taiwan. *J Phys Oceanogr* 2022;52(4):857–76.
- [9] Hsin Y-C, Qiu B, Chiang T-L, Wu C-R. Seasonal to interannual variations in the intensity and central position of the surface Kuroshio east of Taiwan. *J Geophys Res Oceans* 2013;118:4305–16. <https://doi.org/10.1002/jgrc.20323>.
- [10] Wu C-R, Wang Y-L, Lin Y-F, Chao S-Y. Intrusion of the Kuroshio into the South and East China Sea. *Sci Rep* 2017;7:7895. <https://doi.org/10.1038/s41598-017-08206-4>.
- [11] Wang Y-L, Jin F-F, Wu C-R, Qiu B. Northwestern Pacific oceanic circulation shaped by ENSO. *Sci Rep* 2024;14:11684. <https://doi.org/10.1038/s41598-024-62361-z>.
- [12] Lin K, Han T, Zhang Y, Shen C-C, Lee S-Y, Wang J, et al. Influences of East Asian winter monsoon and El Niño–Southern Oscillation variability on the Kuroshio intrusion to the South China Sea over the past 60 years. *Geophys Res Lett* 2024;51:e2023GL104155.
- [13] Wu C-R. Interannual modulation of the Pacific Decadal Oscillation on the low-latitude western North Pacific. *Prog Oceanogr* 2013;110:49–58. <https://doi.org/10.1016/j.pocean.2012.12.001>.
- [14] Wu C-R, Wang Y-L, Chao S-Y. Disassociation of the Kuroshio current with the Pacific Decadal Oscillation since 1999. *Remote Sens* 2019;11:276. <https://doi.org/10.3390/rs11030276>.
- [15] Wu C-R, Hsin Y-C, Chiang T-L, Lin Y-F, Tsui I-F. Seasonal and interannual changes of the Kuroshio intrusion onto the East China Sea shelf. *J Geophys Res Oceans* 2014;119:5039–51. <https://doi.org/10.1002/2013JC009748>.
- [16] Wang Y-L, Wu C-R, Chao S-Y. Warming and weakening trends of the Kuroshio during 1993–2013. *Geophys Res Lett* 2016;43:9200–7. <https://doi.org/10.1002/2016GL069432>.
- [17] Chang Y-C, Wu C-R, Chu PC, Wang Y-L, Centurioni LR, Chen G-Y, et al. Underestimated Kuroshio power and its potential sites off Southeast Taiwan. *Appl Energy* 2025;396:126316. <https://doi.org/10.1016/j.apenergy.2025.126316>.
- [18] VanZwieten J, Meyer I, Alsenas G. Evaluation of HYCOM as a tool for ocean current energy assessment, proceedings of the 2nd marine energy technology symposium, METS14. 2014.
- [19] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need, 2017. arXiv preprint arXiv:1706.03762. <https://doi.org/10.48550/arXiv.1706.03762>.
- [20] Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, et al. An image is worth 16×16 words: transformers for image recognition at scale. 2020. arXiv preprint arXiv:2010.11929. <https://doi.org/10.48550/arXiv.2010.11929>.

- [21] Liu Z, Lin Y, Cao Y, Hu H, Wei Y, Zhang Z, et al. Swin Transformer: Hierarchical vision transformer using shifted windows. 2021. arXiv preprint arXiv:2103.14030. <https://doi.org/10.48550/arXiv.2103.14030>.
- [22] Liu Z, Hu H, Lin Y, Yao Z, Xie Z, Wei Y, et al. Swin Transformer V2: Scaling up capacity and resolution, 2022. arXiv preprint arXiv:2111.09883. <https://doi.org/10.48550/arXiv.2111.09883>.
- [23] Bi K, Xie L, Zhang H, Chen X, Gu X, Tian Q. Accurate medium-range global weather forecasting with 3D neural networks. *Nature* 2023;619:533–8. <https://doi.org/10.1038/s41586-023-06221-7>.
- [24] Bodnar C, Bruinsma WP, Lucic A, Stanley M, Brandstetter J, Garvan P, et al. A foundation model for the Earth system. *Nature* 2025;641:1180–7.
- [25] Sadoughipour E, VanZwieten J, Tang Y. Drifter-based global ocean current energy resource assessment. *Renew Energy* 2025;244:122576.
- [26] Tseng R-S, Chang Y-C, Chu P. Use of global satellite altimeter and drifter data for ocean current resource characterization, marine renewable energy. *Resour Characterization Phys Effect* 2017;159–77.
- [27] Boyer MC, Madec G, Fischer A, Lazar A, Iudicone D. Mixed layer depth over the global ocean: an examination of profile data and a profile-based climatology. *J Geophys Res Oceans* 2004;109(C12):C12003. <https://doi.org/10.1029/2004JC002378>.
- [28] Price J, Weller R, Pinkel R. Diurnal cycling: Observations and models of the upper ocean response to diurnal heating, cooling, and wind mixing. *J Geophys Res Oceans* 1986;91(C7):8411–27.
- [29] Copernicus Marine Service (CMEMS). GOGAD global ocean gridded L4 sea surface heights and derived variables NRT. CMEMS Marine Data Store; 2025. <https://doi.org/10.48670/moi-00149>.
- [30] Centurioni L. NOAA global drifter program quality-controlled 6-hour interpolated data. NOAA NCEI 2010. <https://doi.org/10.25921/7ntx-z961>.
- [31] Centurioni LR, Niiler PP. On the surface currents of the Caribbean Sea. *Geophys Res Lett* 2003;30:1279. <https://doi.org/10.1029/2002GL016231>.
- [32] Chang Y-C, Chen G-Y, Tseng R-S, Centurioni LR, Chu PC. Observed near-surface currents under high wind speeds. *J Geophys Res Oceans* 2012;117:C11026. <https://doi.org/10.1029/2012JC007996>.
- [33] Chang Y-C, Chen G-Y, Tseng R-S, Centurioni LR, Chu PC. Observed near-surface flows under all tropical cyclone intensity levels using drifters in the northwestern Pacific. *J Geophys Res Oceans* 2013;118:2367–77. <https://doi.org/10.1002/jgrc.20187>.
- [34] Chang Y-C, Chu PC, Tseng R-S. Site selection of ocean current power generation from drifter measurements. *Renew Energy* 2015;80:737–45. <https://doi.org/10.1016/j.renene.2015.02.046>.
- [35] Wang G, Wu L, Mei W. Ocean currents show global intensification of weak tropical cyclones. *Nature* 2022;611:496–500. <https://doi.org/10.1038/s41586-022-05455-7>.
- [36] Michel P, Levy O, Neubig G. Are sixteen heads really better than one? arXiv preprint arXiv:1905.10650; 2019. <https://doi.org/10.48550/arXiv.1905.10650>.
- [37] Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805; 2019. <https://doi.org/10.48550/arXiv.1810.04805>.
- [38] Barnier B, et al. Modelling the impact of flow-driven turbine power plants on great wind-driven ocean currents. *Nat Energy* 2020;5:240–9. <https://doi.org/10.1038/s41560-020-0562-8>.
- [39] Bayindir C. Predicting the ocean currents using deep learning. *TWMS J Appl Eng Math* 2023;13(1):373–85.
- [40] Instruments RD. Acoustic Doppler current profiler principles of operation: a practical primer. Teledyne RD Instruments 2001.
- [41] Teledyne Marine, Workhorse II 300, 600, 1200 kHz direct-reading and self-contained ADCPs– Product Datasheet. 2023.
- [42] Shirasawa K, Tokunaga K, Iwashita H, Shintake T. Experimental verification of a floating ocean-current turbine with a single rotor for use in Kuroshio currents, *renew. Energy* 2016;189–95.
- [43] Yeh P, Tsai S, Chen W, Wu S, Hsieh M, Chen B. A simple nozzle-diffuser duct used as a Kuroshio energy harvester. *Processes* 2021;9:1552. <https://doi.org/10.3390/pr9091552>.
- [44] Corporation IHI. IHI demonstrated the world's Largest Ocean current turbine for the first time in the world. *IHI Engr Rev* 2019;52(1):6–9.
- [45] Bahaj A, Myers L, Thomson M. Fundamentals applicable to the utilisation of marine current turbines for energy production. *Renew Energy* 2003;28:2205–11.
- [46] Myers L, Bahaj A. Experimental analysis of the flow field around horizontal axis tidal turbines by use of scale mesh disk rotor simulators. *Ocean Eng* 2010;37: 218–27. <https://doi.org/10.1016/j.renene.2010.01.009>.
- [47] Aguilar C, Hernández M, Muñoz F. Optimal scheduling of renewable energy microgrids. *Appl Energy* 2024;353:122137. <https://doi.org/10.1016/j.apenergy.2023.122137>.
- [48] Argo. Argo float data and metadata from the Global Data Assembly Centre. SEANOE 2025. <https://doi.org/10.17882/42182>.
- [49] Copernicus Marine Service (CMEMS). GOPAF Global Ocean Physics analysis and forecast. CMEMS Marine Data Store 2025. <https://doi.org/10.48670/moi-00016>.
- [50] Cummings JA. Operational multivariate ocean data assimilation. *Q J R Meteorol Soc* 2005;131:3583–604. <https://doi.org/10.1256/qj.05.105>.
- [51] Cummings JA, Smedstad OM. Variational data assimilation for the global ocean. In: *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications II*. Springer; 2013. p. 303–43. https://doi.org/10.1007/978-3-642-35088-7_13.