Human Interaction and Emerging Technologies (IHIET-FS 2026), Vol. 235, 2026, 227–237
https://doi.org/10.54941/ahfe1007277

Design of an Intelligent Agent for Offshore Cage Aquaculture Based on a Cold-Damage Early-Warning Model

Hsun-Yu Lan[1], Chi-Lin Tsai[2], and Shyi-Chyi Cheng[2]

[1]Department of Aquaculture, National Taiwan Ocean University, Keelung 202301, Taiwan
[2]Department of Computer Science and Engineering, National Taiwan Ocean University, Keelung 202301, Taiwan

ABSTRACT

Climate-change-driven anomalies in seawater temperature frequently trigger cold-damage events, leading to significant fish mortality and economic losses in marine cage aquaculture. Although Internet of Things (IoT) systems and marine meteorological observation (MMO) data have been widely used for environmental monitoring, challenges remain in addressing microclimate data scarcity, domain shift, and the prediction of rare extreme events, limiting their effectiveness for actionable farm management. In this study, long-term (2000–2025) MMO datasets from 14 stations of the Taiwan Central Weather Administration, including seawater temperature, air pressure, and wave dynamics, were integrated with real-time offshore sensor data. Based on these inputs, a time-series deep-learning-based intelligent agent, TPP-CASTformer, is proposed. To address data scarcity and domain shift, Test-Time Training (TTT) enables model adaptation during inference using unlabeled target-site data. In addition, a classification mechanism combining Temporal Point Processes and Prototypical Networks is designed to improve interpretability and handle extreme-event imbalance through few-shot learning. The proposed framework integrates heterogeneous MMO and local sensing data, incorporates domain knowledge to define temperature thresholds and exposure durations, and provides interpretable early-warning signals with corresponding management actions. Experimental results based on documented cold-damage events in Penghu, Taiwan indicate that the proposed approach can reduce potential losses by at least 30%, demonstrating improved accuracy and timeliness in cold-damage risk assessment.

Keywords: Intelligent agent, Offshore cage aquaculture, Marine meteorological observation data, Internet of Things, Time-series deep learning model, Cold-damage early-warning model

INTRODUCTION

Data-driven cold-damage early warning in marine aquaculture requires both oceanographic knowledge and long-term environmental observations. In recent years, the development of the Artificial Intelligence of Things (AIoT) has made real-time environmental monitoring more feasible for aquaculture operations (Errachdi, 2023). By continuously collecting environmental variables such as air temperature, seawater temperature, and air pressure,

Received March 13, 2026; Revised March 31, 2026; Accepted April 20, 2026; Available online June 22, 2026

AIoT systems can detect abnormal patterns that may precede cold-damage events (Wang, 2024; Zhang, 2019).

However, deep-learning-based marine meteorological prediction usually depends on long and continuous observation records (Jiang, 2025). In offshore aquaculture farms, sensor systems are often deployed only recently, so the available historical data are limited. In addition, sensor configurations may differ across monitoring locations. Some stations may lack certain measurements, such as wave direction or air pressure. These differences in feature availability make it difficult to combine data from multiple sites for model training (Pérez, 2011; Ha, 2025).

To address these challenges, this study uses long-term marine meteorological observation (MMO) data provided by the Taiwan Central Weather Administration (CWA). The dataset includes records from 14 stations covering the period from 2000 to 2025. These observations form a large historical dataset for training a global prediction model. A sliding-window strategy is used to segment the time-series data into training samples. Stochastic feature masking and patching are further applied to convert heterogeneous environmental observations into unified tensor representations suitable for deep-learning models.

Figure 1 shows the overall architecture of the proposed system. The framework integrates two main data sources: long-term MMO datasets and real-time offshore IoT sensor observations. These environmental data are processed by the proposed AI model, TPP-CASTformer, which performs seawater temperature forecasting and cold-damage risk assessment. The prediction results are then presented through an explainable human–machine interface (HMI), allowing aquaculture operators to monitor risk levels and respond in time.

When the pretrained model is applied to offshore cage aquaculture farms where sensors have been deployed only recently, prediction performance may decline because of domain shift between the MMO dataset and the local farm environment (Sun, 2020). To address this problem, the proposed framework incorporates Test-Time Training (TTT), which allows the model to adapt during inference using unlabeled data from the target site (Zhang, 2026).

Figure 2 illustrates the core workflow of the proposed TPP-CASTformer model. The model first extracts spatiotemporal features through a multi-layer CASTformer encoder. It then performs seawater temperature regression and cold-damage risk classification. The classification module combines Temporal Point Processes (TPP) with Prototypical Networks to model the probability of extreme events and perform interpretable few-shot classification. This design helps address the scarcity of extreme cold-damage samples while improving the interpretability of the prediction results.

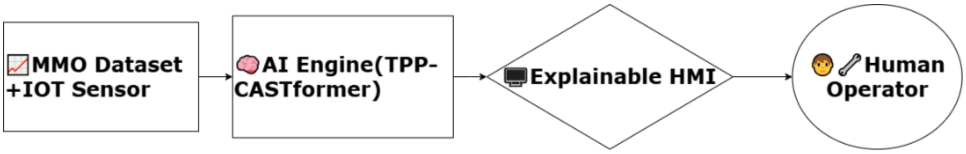

Figure 1: Overall architecture of the proposed cold-damage early-warning system.

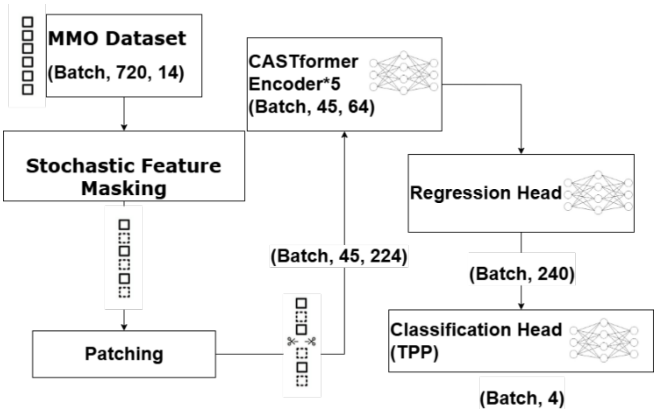

Figure 2: Workflow of the proposed TPP-CASTformer model.

METHODOLOGY

In this section, we describe the proposed TPP-CASTformer early-warning model and its data-processing pipeline. The model is trained using historical marine meteorological observations provided by the Taiwan Central Weather Administration (CWA), which include long-term sensor time-series records from 14 stations between 2000 and 2025 (Chen, 2020). The overall workflow of the model is illustrated in Figure 2.

The first step in building the prediction model is data preparation, including feature selection, missing-value imputation, feature engineering, and data normalization. After preprocessing, the environmental observations are converted into spatiotemporal tensors using a sliding-window strategy and labeled with the corresponding cold-damage risk classes. Table 1 summarizes the datasets used for model training, validation, and evaluation. The following subsections describe each data-processing step and the model design in detail.

Table 1: Description of training and test datasets.

Attribute	CWA Training Set	CWA Validation Set	CWA Test Set	Target Site Test Set
Observation Period	2000/01–2025/12	2000/01–2025/12	2000/01–2025/12	2025/06–2026/02
Stride	6 hr	12 hr	12 hr	6 hr
Samples	101,399	20,854	20,855	861
Class 0	19,172	–	19,172	825
Class 1	856	–	856	36
Class 2	32	–	32	0
Class 3	225	–	225	0

This study uses marine meteorological observation data from the Taiwan Central Weather Administration (CWA) as the primary dataset. The dataset contains hourly observation records from 2000–2025. The raw data include

12 environmental variables, such as air pressure, seawater temperature, wave direction, peak wave period, and gust wind speed. These variables capture important environmental dynamics and are widely used to analyze and predict extreme weather events, including cold surges affecting coastal waters (Chen, 2020).

Feature selection is performed to identify the environmental variables that are most relevant to cold-damage prediction. To evaluate the relative importance of these variables, this study applies the **XGBoost Gain metric** as a feature importance measure. The Gain metric estimates the contribution of each feature to the improvement of model performance during the training process (Chen and Guestrin, 2016).

Based on the resulting feature importance scores, variables with low contributions are removed to reduce noise and improve model efficiency. The remaining variables are used as input features for the proposed **TPP-CASTformer model,** allowing the model to focus on variables that are most relevant to seawater temperature dynamics and cold-damage risk.

Data preprocessing aims to reduce noise and large variations in sensor data and to construct an appropriate input format for the deep learning model. The preprocessing pipeline consists of several key steps, including missing value imputation, feature engineering, data normalization, and tensor construction.

Missing Value Imputation: Marine sensors may produce data gaps due to communication interruptions or harsh environmental conditions. To preserve the temporal continuity of the observations, bidirectional linear interpolation is applied to estimate missing values (Moritz, 2017).

Feature Engineering: To help the model capture multi-scale temporal dynamics, additional features are derived from the target time series (seawater temperature) (Wu, 2021). The following feature transformations are applied:

- Rolling Statistics: Moving averages (MA) with 24-hour and 168-hour windows are computed to smooth short-term fluctuations and capture weekly climate trends.
- Exponential Moving Average (EMA): By assigning a larger decay weight α to recent observations, the model becomes more responsive to sudden environmental changes.
- Cyclical Encoding: Time variables (e.g., the transition from 23:00 to 00:00) contain numerical discontinuities. Sine and cosine transformations are used to represent the cyclical nature of time (Oreshkin, 2020).

Data Normalization and Tensor Construction: To ensure that different environmental variables have comparable scales, all input features are normalized before model training. Normalization helps stabilize the training process and prevents variables with large numerical ranges from dominating the learning process. After normalization, the time-series data are segmented using a sliding-window strategy to construct model inputs and prediction targets. For each sample, a fixed historical window is used as the input sequence, while a future time window is used as the prediction target.

Following this procedure, the processed dataset is converted into input and target tensors, generating the input tensor X (101399 × 720 × 14) and the prediction target tensor Y_{Reg} (101399 × 240 × 1).

Cold-Damage Risk Classification Definition: Cold-damage risk levels are defined according to seawater temperature thresholds and exposure duration derived from aquaculture operational experience. Based on these criteria, the prediction targets are categorized into four risk classes.

- Class 0: Seawater temperature ≥ 18 °C.
- Class 1: Seawater temperature between 15 °C and 18 °C.
- Class 2: Seawater temperature between 14 °C and 15 °C for at least three consecutive days.
- Class 3: Seawater temperature < 14 °C.

Under this definition, the target experimental site (total sample size: 861) exhibits an extreme class imbalance. Class 0 accounts for 95.82% of the samples, while Class 2 and Class 3 samples are absent. This long-tail distribution poses a significant challenge for model generalization in real-world environments. In marine environments, extreme cold-damage events (Class 2 and Class 3) that cause severe fishery losses occur with very low probability, resulting in a highly imbalanced training dataset. When such unbalanced data are used directly for model training, the model may fall into the majority-class trap, losing sensitivity to extreme climate conditions (Jackson, 2018). To improve the model's ability to learn decision boundaries for minority classes, this study introduces data augmentation techniques. Specifically, Mixup (Zhang, 2018) and Jitter (Um, 2017) strategies are applied to generate synthetic samples. These synthetic samples are used to expand the training data for Class 2 and Class 3, enabling the model to better capture the characteristics of extreme low-temperature events and reducing the false negative rate in practical applications.

As shown in Figure 3, before the input tensor is fed into the encoder, the data first passes through two preprocessing modules: Stochastic Feature Masking and Patching.

First, in practical deployments, the types and feature dimensions of sensor data may vary across locations (e.g., some sites may lack wave direction or air pressure measurements). To improve the model's ability to handle heterogeneous data and missing features, this study introduces a Stochastic Feature Masking mechanism (He et al., 2022). During training, specific input feature channels are randomly masked, either by setting them to zero or replacing them with the mean. This encourages the model to learn latent cross-correlations among variables and maintain stable prediction performance when certain features are unavailable.

Second, to reduce the quadratic computational complexity of Transformers and mitigate local semantic loss in long sequences, Patching is adopted (Nie et al., 2023). The input sequence is divided into non-overlapping blocks (length = 16), which reduces memory overhead and improves the model's ability to capture long-term meteorological dependencies. High-dimensional marine meteorological data often suffer from limited temporal context and high-frequency noise. Traditional RNNs are prone to gradient vanishing, while standard Transformers lack inductive bias for temporal periodicity (Vaswani et al., 2017). Therefore, this study adopts **CASTformer**, based on an orthogonal decomposition mechanism, as the core encoder. The detailed architecture is shown in Figure 3. The key idea of CASTformer is to address spatiotemporal entanglement and non-causal information leakage in high-dimensional time-series prediction.

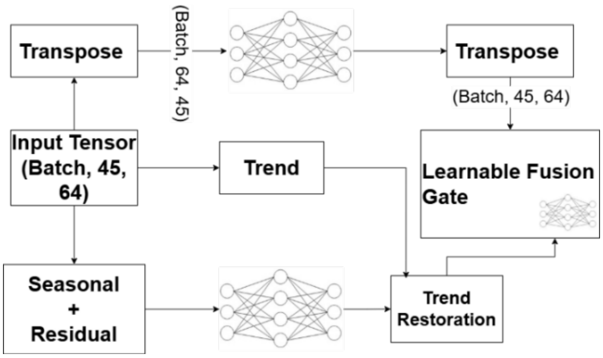

Figure 3: Architecture of the CASTformer encoder for seawater temperature prediction.

Series Decomposition and Causal Decoupling: Referencing the design of Autoformer (Wu et al., 2021), the CASTformer encoder introduces a series decomposition mechanism to separate the input time series into trend and seasonal components. This decomposition reduces temporal interference and allows the model to capture long-term meteorological patterns more effectively.

Formally, given the input sequence X, the seasonal component is obtained by subtracting the smoothed trend component estimated through moving-average pooling:

$$X_{seasonal} = X - \text{AvgPool}(\text{Padding}(X)) \tag{1}$$

where AvgPool(Padding(X)) represents the trend component derived from moving-average smoothing. The decomposed components are subsequently processed by the encoder to model their temporal dependencies.

To prevent future information leakage during prediction, a causal decoupling mechanism is further introduced in the attention computation. This mechanism constrains the model to attend only to historical observations, ensuring that the prediction strictly follows temporal causality. Through this decomposition–decoupling design, the encoder can effectively isolate slow-changing climate trends from short-term fluctuations, thereby improving robustness when handling sudden meteorological mutations.

Dual-Branch Attention Architecture: As shown in Figure 3, to decouple spatiotemporal dependencies, this study designs a dual-branch structure. The lower branch is responsible for modeling the temporal dynamic evolution of

$X_{seasonal}$, while the upper branch captures the interaction relationships among feature channels through a transpose-dimension mechanism.

After transposing the input tensor to $X^T \in \mathrm{R}^{B\times D\times T}$, channel-wise self-attention is applied to model interactions among the meteorological variables:

$$\text{Attention}_c(X^T) = \text{Softmax}\left(\frac{(X^T W_Q)(X^T W_K)^{\mathrm{T}}}{\sqrt{d_k}}\right)(X^T W_V) \tag{2}$$

where W_Q, W_K, and W_V are learnable projection matrices for the query, key, and value transformations, and d_k denotes the dimensionality of the key vectors.

This mechanism captures global dependencies among meteorological variables and assigns higher weights to physically meaningful features, thereby suppressing high-dimensional noise interference. After processing, the tensor dimensions are transposed back to the original shape.

Learnable Fusion Gate: Considering that marine systems may be governed by different mechanisms during stable periods and extreme weather events, a simple feature addition cannot adapt to dynamic environmental changes. Therefore, we introduce a learnable fusion gate to combine the outputs of the dual-branch attention structure.

Let $H_{Channel}$ denote the feature representation extracted from the channel-attention branch and $H_{temporel}$ denote the representation from the temporal branch. The final fused representation is defined as:

$$Z_{fused} = g \cdot H_{Channel} + (1 - g) \cdot (X_{input} + H_{temporel}) \tag{3}$$

where $g \in [0,1]$ is a learnable gating parameter. The term $X_{input} + H_{temporel}$ forms a residual connection that preserves the temporal dynamics of the input sequence.

This adaptive fusion mechanism allows the model to dynamically balance channel interactions and temporal evolution, thereby improving generalization under changing meteorological conditions.

Prototype-based TPP Classification Head: To address the scarcity of extreme disaster samples at the target site and improve prediction interpretability, this study combines Temporal Point Processes (TPP) (Shchur, 2021) with Prototypical Networks (Snell, 2017). We design a multi-feature fusion mechanism in the classification head. Specifically, the classifier receives and concatenates three groups of heterogeneous features:

- **Spatiotemporal & Static Features ($\mathbf{f}_{\text{static}}$):** High-dimensional spatiotemporal features ($\mathbf{x}_{\text{feat}}$) extracted by the CASTformer encoder, combined with station embeddings to incorporate site-specific static information.
- **Predicted Future Trend ($\mathbf{f}_{\text{pred\_trend}}$):** The predicted seawater temperature sequence $\hat{Y}_{reg}$ for the next 240 hours produced by the regression head is processed through a 1D convolution layer (Conv1D) to extract temporal features.

- **Raw Historical Trend ($\mathbf{f}_{\text{hist\_trend}}$):** The raw historical input data (after stochastic feature masking) is processed through an independent convolution branch followed by Global Average Pooling, extracting low-frequency background climate features without interference from the Transformer attention mechanism.

Assume that each cold-damage risk class k ∈ {0,1,2,3} has a learnable prototype center $\mathbf{c}_k$ in the embedding space. For a feature vector $\mathbf{h}_i$ produced by the CASTformer encoder, the squared Euclidean distance to each prototype is computed as

$$d(\mathrm{h}_j, \mathrm{c}_k) = \left\| \mathrm{h}_j, \mathrm{c}_k \right\|_2^2 \tag{4}$$

Euclidean distance is selected instead of cosine similarity because it belongs to the family of Bregman divergences. Minimizing this distance is equivalent to maximizing the likelihood of an exponential-family distribution, which helps form tighter intra-class compactness.

Based on the TPP architecture, the probability of a cold-damage event is modeled as inversely related to the distance between the feature representation and the corresponding class prototype. To address the extreme sparsity of disaster samples, the classifier adopts Focal Loss (Lin et al., 2017) instead of the standard cross-entropy loss, enabling the model to focus on hard samples and improving sensitivity to rare extreme-weather events.

As shown in Table 1, the offshore cage aquaculture farm used as the target test site contains only 861 short sequence samples. When a zero-shot model pre-trained on the large-scale CWA dataset is directly applied, prediction performance often degrades due to domain shift caused by geographical microclimate differences and heterogeneous sensor configurations (Sun et al., 2020).

When only a small amount of calibration data $\mathbf{X}_{calib}$ is available, updating all model parameters may lead to overfitting and catastrophic forgetting, thereby destroying the general spatiotemporal knowledge learned from the large historical dataset (Houlsby et al., 2019). To address this issue, this study adopts a local unfreezing strategy during the Test-Time Training (TTT) phase. Specifically, the regression branch is extracted from the original TPP-CASTformer architecture. During adaptation, the parameters of the CASTformer backbone θ_{frozen} remain fixed, while only the regression head parameters θ_{head} are updated. The optimization objective is defined as

$$\theta_{head}^{*} = \arg\min_{\theta_{head}} \frac{1}{N} \sum_{i=1}^{N} \left\| Y_{calib}^{(i)} - f_{MLP}\left(f_{backbone}(X_{calib}^{(i)}; \theta_{frozen}); \theta_{head} \right) \right\|_2^2. \tag{5}$$

This local adaptation strategy allows the model to correct systematic bias caused by domain shift while preserving the global spatiotemporal knowledge learned during pretraining. As a result, the model avoids overfitting to noisy calibration data and improves prediction accuracy compared with the original zero-shot model, significantly reducing the final Root Mean Square Error (RMSE).

CONCLUSION

This study proposes an intelligent forecasting and decision-support framework for cold-damage early warning in offshore cage aquaculture. By integrating long-term marine meteorological observation (MMO) datasets with farm-level sensing systems, the proposed TPP-CASTformer model effectively captures seawater temperature dynamics and detects potential cold-risk conditions.

The forecasting results are translated into interpretable risk levels through a structured classification mechanism, enabling aquaculture operators to make timely management decisions. As summarized in Table 2, the ablation study confirms that the proposed architecture improves prediction stability and extreme-event detection. A retrospective analysis of historical cold-damage events in Penghu further suggests that the proposed early-warning framework could potentially reduce economic losses by up to 30%, demonstrating its practical value for risk management in offshore aquaculture systems.

Table 2: Ablation study and performance comparison of the proposed model.

Model	Changed	CWA RMSE	CWA Class 2 Recall	CWA Class 2 F1	CWA Class 3 Recall	CWA Class 3 F1	CWA Macro Avg F1	Target Zero-Shot RMSE	Target TTT RMSE
Ours	Full Model	**0.99**	0.66	0.27	**0.77**	**0.86**	0.7	3.74	**1.02**
1	Standard Encoder	1.04	0.22	0.15	0.75	0.85	0.67	4.53	1.25
2	No Masking	1.07	0.31	0.15	0.6	0.74	0.62	3.48	1.21
3	No Patching	3.67	**1**	0.18	0.17	0.3	0.52	4.15	1.63
4	No TPP	**0.99**	0.56	**0.31**	0.75	0.85	**0.72**	4.05	1.07
5	No Static	1.03	0.81	0.27	0.51	0.68	0.64	2.69	1.3
6	No Pred Trend	1.03	0.78	0.25	0.72	0.84	0.69	2.86	1.2
7	No Hist Trend	1.03	0.50	0.20	0.69	0.79	0.66	**2.53**	1.24

ACKNOWLEDGMENT

This work was supported in part by the National Science and Technology Council, Taiwan, under grant numbers NSTC 114-2221-E-019-022-MY3 and NSTC 114-2637-8-011-004-.

REFERENCES

Chen, T. and Guestrin, C. (2016) “XGBoost: A scalable tree boosting system”, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 785–794.

Errachdi, H., El Malki, M., Sefiani, N. and Hachimi, H. (2023) “Bridging the gap: Challenges and opportunities of IoT and wireless sensor networks in marine environmental monitoring”, Engineering Proceedings.

Ha, J., Kim, S. and Lee, D. (2025) “Handling heterogeneous marine sensor data for environmental prediction using deep learning”, Sensors, 25(2), 814.

He, K., Chen, X., Xie, S., Li, Y., Dollár, P. and Girshick, R. (2022) “Masked autoencoders are scalable vision learners”, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16000–16009.

Houlsby, N., Giurgiu, A., Jastrzębski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M. and Gelly, S. (2019) “Parameter-efficient transfer learning for NLP”, in Proceedings of the International Conference on Machine Learning. pp. 2790–2799.

Jackson, D.A. and Somers, K.M. (2018) “Class imbalance and rare-event prediction in environmental data analysis”, Ecological Informatics, 47, pp. 161–170.

Jiang, T., Liu, Q., Zhang, L. and Sun, X. (2025) “Deep-learning-based ocean environmental forecasting using multi-source meteorological observations”, Applied Ocean Research, 146, 104012.

Lin, T.Y., Goyal, P., Girshick, R., He, K. and Dollár, P. (2017) “Focal loss for dense object detection”, in Proceedings of the IEEE International Conference on Computer Vision. pp. 2980–2988.

Moritz, S. and Bartz-Beielstein, T. (2017) “imputeTS: Time series missing value imputation in R”, The R Journal, 9(1), pp. 207–218.

Nie, Y., Nguyen, N.H., Sinthong, P. and Kalagnanam, J. (2023) “A time series is worth 64 words: Long-term forecasting with transformers”, in Proceedings of the International Conference on Learning Representations.

Oreshkin, B.N., Carpov, D., Chapados, N. and Bengio, Y. (2020) “N-BEATS: Neural basis expansion analysis for interpretable time series forecasting”, in Proceedings of the International Conference on Learning Representations.

Pérez, C.A., Rodríguez, J., Rodríguez, D. and Gutiérrez, J. (2011) “A system for monitoring marine environments based on wireless sensor networks”, in Proceedings of OCEANS 2011 IEEE.

Shchur, O., Biloš, M. and Günnemann, S. (2021) “Intensity-free learning of temporal point processes”, in Proceedings of the International Conference on Learning Representations.

Snell, J., Swersky, K. and Zemel, R. (2017) “Prototypical networks for few-shot learning”, in Proceedings of the Advances in Neural Information Processing Systems. pp. 4077–4087.

Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A.A. and Hardt, M. (2020) “Test-time training with self-supervision for generalization under distribution shifts”, in Proceedings of the 37th International Conference on Machine Learning.

Um, T.T., Pfister, F.M.J., Pichler, D., Endo, S., Lang, M., Hirche, S., Fietzek, U. and Kulić, D. (2017) “Data augmentation of wearable sensor data for Parkinson’s disease monitoring using convolutional neural networks”, in Proceedings of the ACM International Conference on Multimodal Interaction. pp. 216–220.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I. (2017) “Attention is all you need”, in Advances in Neural Information Processing Systems.

Wang, Y., Li, H., Chen, Z. and Zhao, X. (2024) “Artificial intelligence-based environmental monitoring for marine aquaculture: A data-driven approach”, Ocean Engineering, 295, 116985.

Wu, H., Xu, J., Wang, J. and Long, M. (2021) “Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting”, in *Advances in Neural Information Processing Systems*.

Zhang, H., Cisse, M., Dauphin, Y.N. and Lopez-Paz, D. (2018) “mixup: Beyond empirical risk minimization”, in *Proceedings of the International Conference on Learning Representations*.

Zhang, M., Levine, S. and Finn, C. (2026) “Adaptive test-time training for domain-shifted time-series prediction”, in *Proceedings of the International Conference on Learning Representations*.

Zhang, X., Liu, H., Li, Y. and Liu, S. (2019) “Internet of things in marine environment monitoring: A review”, *Sensors,* 19(7).