

VQLTI: Long-Term Tropical Cyclone Intensity Forecasting with Physical Constraints

Xinyu Wang¹, Lei Liu^{1*}, Kang Chen^{1,2}, Tao Han^{2,3}, Bin Li¹, Lei Bai^{2*}

¹ University of Science and Technology of China

² Shanghai Artificial Intelligence Laboratory

³ Hong Kong University of Science and Technology

xinyuwang@mail.ustc.edu.cn, liulei13@ustc.edu.cn, ck6@mail.ustc.edu.cn, hantao10200@gmail.com, binli@ustc.edu.cn, baisanshi@gmail.com

Abstract

Tropical cyclone (TC) intensity forecasting is crucial for early disaster warning and emergency decision-making. Numerous researchers have explored deep-learning methods to address computational and post-processing issues in operational forecasting. Regrettably, they exhibit subpar long-term forecasting capabilities. We use two strategies to enhance long-term forecasting. (1) By enhancing the matching between TC intensity and spatial information, we can improve long-term forecasting performance. (2) Incorporating physical knowledge and physical constraints can help mitigate the accumulation of forecasting errors. To achieve the above strategies, we propose the VQLTI framework. VQLTI transfers the TC intensity information to a discrete latent space while retaining the spatial information differences, using large-scale spatial meteorological data as conditions. Furthermore, we leverage the forecast from the weather prediction model FengWu to provide additional physical knowledge for VQLTI. Additionally, we calculate the potential intensity (PI) to impose physical constraints on the latent variables. In the global long-term TC intensity forecasting, VQLTI achieves state-of-the-art results for the 24h to 120h, with the MSW (Maximum Sustained Wind) forecast error reduced by **35.65%**-**42.51%** compared to ECMWF-IFS.

Code — <https://github.com/1457756434/VQLTI>

Introduction

Tropical cyclones (TCs), commonly known as typhoons or hurricanes, are cyclonic circulations that are stronger in the near-sea surface layer than in the upper layer (Emanuel 2003). These systems possess tremendous destructive potential, as evidenced by typhoon Mangkhut in 2018, which impacted 3.29 million people and resulted in USD 7.62 billion in direct economic losses (Yang et al. 2019). Consequently, accurate long-term forecasting of TC intensity is of critical importance for minimizing disaster-related damages and losses (Woodruff, Irish, and Camargo 2013; Zhang, Fang, and Yu 2023).

TC intensity forecasting typically refers to forecasting the maximum sustained wind (MSW) and minimum sea level

*Both are Corresponding Authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

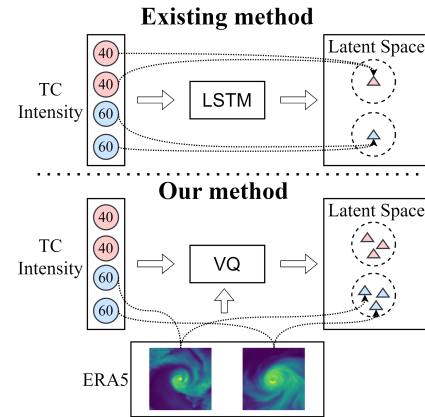


Figure 1: Past methods, such as LSTM, would map the same TC intensity to the same position in the latent space, thereby ignoring their spatial information. To preserve this spatial information, we use ERA5 data as a condition to map the TC intensity to a discrete latent space, such that the same TC intensity is distributed with a certain intra-class distance, thereby capturing their spatial differences.

pressure (MSLP) of a TC. In current operational forecasting, these forecasting methods primarily rely on Numerical Weather Prediction models (NWP) (Buizza et al. 2018), such as the Global Forecast System of the China Meteorological Administration (CMA-GFS) (Shen et al. 2023), the Global Forecast System of the National Centers for Environmental Prediction (NCEP-GFS) (Zhou and Juang 2023), and the Integrated Forecasting System of the European Centre for Medium-Range Weather Forecasts (ECMWF-IFS) (Haiden et al. 2018). However, these models usually require the use of high-performance computing resources to solve a large number of partial differential equations, resulting in high computational costs (Rüttgers et al. 2019). Therefore, the computationally less expensive deep learning (DL) methods have become a research hotspot. Deep learning has made pioneering progress in the field of meteorology, able to understand natural phenomena across temporal and spatial scales, thereby driving the advancement of natural sciences (Bi et al. 2023; Lam et al. 2023; Chen et al. 2023a,b). In recent

years, many researchers have begun to apply deep learning algorithms to the forecasting of TC intensity.

Researchers have conducted some work on deep learning methods for time series forecasting of TC intensity (Pan, Xu, and Shi 2019; Jiang et al. 2023; Huang et al. 2022). They use popular neural network structures like Transformer (Vaswani et al. 2017) and LSTM (Graves and Graves 2012), based on historical time series information to predict future intensity. However, these simple time series features are still not sufficient to fully represent the current state of the TC, so the forecasting performance is not ideal. More researchers (Zhang et al. 2022; Meng et al. 2023; Huang et al. 2023; Wang, Li, and Zheng 2024; Wang et al. 2024) have been trying to use multimodal data, such as the renowned fifth generation of ECMWF reanalysis data (ERA5) (Hersbach et al. 2020), to effectively extract the spatial features of the TC. This led to decent progress in short-term 24-hour (24h) forecasting. However, these methods generally suffer from significant error accumulation problems, resulting in poorer long-term forecasting performance.

The above methods have two problems. First, there is insufficient matching between TC intensity and spatial information. In the actual dataset, the same TC intensity can exhibit a variety of different spatial patterns. Some past methods (Vaswani et al. 2017; Huang et al. 2022) have neglected the impact of this spatial diversity, which is why relying solely on historical TC intensity performs poorly. Inspired by prior work demonstrating the benefits of data discretization for time series prediction (Rasul et al. 2022, 2024), we propose to represent the original intensity in a discrete space. The difference between our method and existing methods is illustrated in Figure 1. By representing the TC intensity information using discrete latent variables, and transforming the original TC intensity forecasting task into forecasting of the corresponding latent variables, this approach enables better matching between the TC intensity information and its spatial patterns.

Secondly, we believe that one important reason for the poor long-term forecasting performance of deep learning models is the insufficient learning of underlying physical knowledge. By harnessing the massive global ERA5 dataset, significant advancements have been achieved in machine learning-driven medium-range weather forecasting (Bi et al. 2023; Lam et al. 2023; Chen et al. 2023a,b; Han et al. 2024). Taking the FengWu model (Chen et al. 2023a) as an example, it has achieved a skillful 10.75-day forecast horizon for future atmospheric fields. Even more impressively, the TC track forecasting capabilities of FengWu rival or surpass those of the ECMWF-IFS, which is grounded in fluid dynamics equations (Hsu et al. 2024). Based on this observation, we believe that the forecast results of FengWu contain certain underlying physical knowledge. If we can effectively utilize this knowledge, it will significantly improve the performance of TC intensity forecasting. Additionally, since the ERA5 dataset itself underestimates TC intensity information (Bourdin et al. 2022; Malakar et al. 2020), the FengWu forecast field also underestimates this information. Supplementing this part of the information could also contribute to performance improvement. Potential intensity (PI)

is a theoretical model that characterizes the upper limit of TC intensity under given environmental conditions and constraints (Emanuel 1986). We calculate the PI value based on the forecast field, using physical constraints to enhance the forecasting performance.

To address the aforementioned issues, we propose the Vector Quantized Long-term Tropical cyclone Intensity forecasting model (VQLTI). The entire model training involves two stages. In the first stage, with the ERA5 data as a condition, we map the intensity information to a discrete latent space using the Conditional Vector Quantized-Variational AutoEncoder (CVQ-VAE) (Van Den Oord, Vinyals et al. 2017) framework. In the second stage, after mapping the historical intensity information to the latent space, similar to LSTM (Graves and Graves 2012), we transfer the forecasting task to the latent space. By utilizing the FengWu forecast field to calculate the PI, we apply physical constraints to the forecasting of the latent variables. Then, using the FengWu forecast field as a condition, we decode the forecast latent variables into TC intensity, thereby realizing the forecasting. By adjusting the number of iterations in the latent space, we can achieve forecasts of any time step. The key contributions of this work are:

- Proposing to map the TC intensity information to a discrete space, which enables better alignment with the TC spatial information.
- Incorporating the FengWu forecast field information to achieve high-precision real-time long-term TC intensity forecasting.
- By calculating the PI, the problem of the forecast field underestimating TC intensity is improved, and the long-term forecasting performance is enhanced.
- The VQLTI model achieves state-of-the-art (SOTA) performance in global long-term TC intensity forecasting.

Preliminaries

Problem Description and Data

VQLTI takes the past n -step TC intensity and ERA5 data, as well as the m -step FengWu forecast data as input, to predict the future m -step TC intensity. In this study, the “USA_WIND” (knots) and “USA_PRES” (hPa) variables from the IBTrACS dataset (Knapp et al. 2010) are employed as the ground truth for TC intensity, with a temporal resolution of 6 hours, denoted as I .

The ERA5 data can be interpreted as the fitted data of the actual atmospheric conditions. The same 69 variables as the FengWu model are selected, including 13 levels (50, 100, 150, 200, 250, 300, 400, 500, 600, 700, 850, 925, and 1000 hPa) of 5 variables (geopotential (z), specific humidity (q), zonal component of wind (u), meridional component of wind (v), and temperature (t)) and 4 surface variables (10-meter u wind component ($u10$), 10-meter v wind component ($v10$), 2-meter temperature ($t2m$), and mean sea level pressure (msl)) (Chen et al. 2023a), and the region with a 10° (longitude and latitude) (Zhang et al. 2022) around TC center is cropped. The spatial resolution of the data is 0.25° , and the temporal resolution is 6 hours, denoted as E .




Figure 2: For typhoon Hinnamnor at 2022-09-02 18:00, the ERA5 data and FengWu forecast data for 2-meter temperature are visualized, along with their differences. The figure shows FengWu’s 6-hour and 120-hour forecasts, with starting times of 2022-09-02 12:00 and 2022-08-28 18:00:00 respectively. As the forecast lead time increases, the tracked TC position from FengWu deviates more, resulting in spatial misalignment between the forecast and actual ERA5 data, as indicated by the red boxes.

The FengWu model performs autoregressive forecasting on the historical ERA5 data. Therefore, we utilize the ECMWF’s TC tracker (Magnusson et al. 2021) to search for the potential future position of the TC on the forecast field and then extract the data within a 10° diameter centered on this location, denoted as F . Figure 2 shows the difference between the FengWu forecast and the ERA5 data, using typhoon Hinnamnor in 2022 as an example.

In summary, this study aims to learn a function $f(\cdot)$ as follows:

$$\hat{I}_{1:m} = f(I_{-(n-1):0}, E_{-(n-1):0}, F_{1:m}, m), \quad (1)$$

where $\hat{I}_{1:m} \in \mathbb{R}^{m \times 2} = \{I_1, \dots, I_m\}$ denotes the forecast of the future m steps of TC intensity, $I_{-(n-1):0} \in \mathbb{R}^{n \times 2}$, $E_{-(n-1):0} \in \mathbb{R}^{n \times 69 \times 40 \times 40}$, and $F_{1:m} \in \mathbb{R}^{m \times 69 \times 40 \times 40}$.

Potential Intensity

The fundamental energy source for TCs is the release of latent heat from the condensation of water vapor. This process can be described using a Carnot heat engine model to characterize the energy cycle (Emanuel 1986, 1987). Specifically, the surrounding air enters the eyewall, absorbs heat from the ocean, then continuously rises and goes through the energy cycle, before finally flowing out to the outer walls and descending (Emanuel 2018). The incorporation of this theoretical model enables the formulation of TC’s theoretical maximum intensity, known as the Potential Intensity (PI), which is typically expressed in terms of wind speed as:

$$|V_{max}|^2 = \frac{C_k}{C_D} \frac{T_s - T_o}{T_o} (k_0^* - k), \quad (2)$$

where C_k represents the dimensionless surface exchange coefficient for enthalpy, C_D is momentum surface exchange coefficients, T_s denotes the sea surface temperature, T_o is the outflow temperature, and k is the enthalpy per unit mass, which is defined as:

$$k \equiv c_p T + L_v q, \quad (3)$$

Where c_p represents the specific heat capacity at constant pressure, L_v denotes the latent heat of vaporization, and q is the specific humidity. In this study, referring to Gilford’s calculation process (Gilford 2021), a more concise expression is utilized:

$$V_{max}, p_{min} = PI(t2m, msl, t, q). \quad (4)$$

The data of $t2m$, msl , t , and q have been mentioned in ERA5 data. Note that the original equation uses sea-surface temperature, but since FengWu does not perform forecasting, it chooses to use the relatively close $t2m$ as a substitute, and the resulting deviation is considered acceptable.

Methodology

The VQLTI model consists of two training stages, as Figure 3 illustrates. In the first stage, the model employs the CVQ-VAE training framework to learn the mapping from TC intensity information to a discrete latent space. The second phase involves learning within the latent space, forecasting future latent variables, and then decoding those into predicted TC intensity.

Pre-training

The first stage of training is inspired by CVAE (Sohn, Lee, and Yan 2015), CVAE-GAN (Bao et al. 2017), and VQ-TR (Rasul et al. 2024), forming an encoder-decoder CVQ-VAE framework. The Feature pyramid network (FPN) (Lin et al. 2017) is used to extract multi-scale feature maps of size $\{40 \times 40, 20 \times 20, 10 \times 10, 10 \times 10\}$ from the ERA5 data. Similar to the Vision Transformer (ViT) (Dosovitskiy et al. 2020), the feature maps are reshaped into a sequence of flattened 2D patches with 5×5 patch size, which are then concatenated to represent the spatial information of the TC as $E_p \in \mathbb{R}^{88 \times d_p}$, where d_p is the dimension of E_p . Conditioned on the TC spatial information E_p , the Encoder maps the TC intensity information to a latent representation. The TC intensity is then passed through a multi-layer perceptron (MLP) and performs cross-attention (Vaswani et al. 2017) computation with E_p . The overall encoding process can be expressed as:

$$\begin{aligned} h &= f_E(I | E) \\ &= \text{CrossAttn}(\text{MLP}(I)W_Q^E, E_p W_K^E, E_p W_V^E), \end{aligned} \quad (5)$$

$$\text{CrossAttn}(Q, K, V) = \text{softmax}\left(\frac{Q(K)^T}{\sqrt{d_k}}\right)V, \quad (6)$$

where W_Q^E , W_K^E , and W_V^E are learnable parameter matrices.

Define a learnable discrete latent space $\{e_1, \dots, e_J\}$ with a size of J and a dimension of d_z , commonly referred to as the codebook. The encoder output vector h is compared




Figure 3: VQLTI Framework. (a) The first stage involves CVQVAE pre-training, which uses ERA5 data as the condition to map TC intensity to a discrete latent space. (b) The second stage is the forecasting stage. Based on the pre-trained model, the blue module is frozen, and the red module is unfrozen. The TC intensity information sequence and the embedding vector in the codebook are matched to obtain the discrete latent variable sequence. When predicting the future latent variables, the Physical Information (PI) is calculated based on the FengWu forecast results to physically constrain the latent variables. Finally, the forecast latent variables are decoded into TC intensity using the FengWu forecast field as the condition.

against the discrete vectors in the codebook, and the nearest match is selected as the quantized representation, which can be expressed as:

$$Z = e_q = VQ(h), \text{ where } q = \operatorname{argmin}_j \|h - e_j\|_2. \quad (7)$$

The architectures of the decoder and encoder are analogous. Conditioning on the TC spatial information E_p , the decoder learns to reconstruct the quantized discrete latent variable Z back to the TC intensity, represented as:

$$\begin{aligned} \hat{I} &= f_D(Z | E) \\ &= \operatorname{MLP}(\operatorname{CrossAttn}(ZW_Q^D, E_p W_K^D, E_p W_V^D)). \end{aligned} \quad (8)$$

Besides the reconstruction loss L_{recon} , the overall network also aims to minimize the codebook loss $L_{codebook}$ and the commitment loss L_{commit} introduced by the VQ component. The reconstruction loss, using the mean absolute error (MAE) here, optimizes the encoder and decoder. The codebook loss encourages the encoder’s output to align with the codebook vectors. The commitment loss ensures the encoder commits to the embeddings, preventing unbounded output growth and facilitating training convergence. The overall aim is to minimize the following loss function:

$$\begin{aligned} L &= L_{recon} + L_{codebook} + \beta L_{commit} \\ &= \operatorname{MAE}(I, \hat{I}) + \|\operatorname{sg}[h] - e\|_2^2 + \beta \|h - \operatorname{sg}[e]\|_2^2. \end{aligned} \quad (9)$$

Note that Equation (7) does not have a gradient, so the stop gradient (sg) operator is used in $L_{codebook}$ and L_{commit} . During the forward pass, it is defined as the identity function and has zero partial derivatives. In the actual implementation, the “detach” operator blocks the gradient from propagating back. β is a hyperparameter.

Forecast

After the pre-training stage, we obtain discrete latent variables that capture the TC intensity information while preserving the associated spatial patterns. If we obtain the discrete latent variables corresponding to the future TC intensity, as well as the future ERA5 data, we can then realize accurate forecasting. Inspired by MSCAR (Wang et al. 2024) and LSTM (Graves and Graves 2012), forecasting in the latent space is feasible. The second stage of training will focus on forecasting the latent variables. Based on the pre-trained model, we freeze the encoding and codebook modules, unfreeze the decoding part, and use FengWu’s forecasting data to approximate the future ERA5 data, to achieve real-time forecasting. The second stage training process is shown in Figure 3(b).

The historical TC intensity information is quantified

Algorithm 1: Latent Variable Iteration Model

Input: $Z = \{Z_{-(n-1)}, \dots, Z_0\}$, $P = \{P_1, \dots, P_m\}$

Output: $\hat{Z} = \{\hat{Z}_1, \dots, \hat{Z}_m\}$

```
1: Initialize window:  $W = \{Z_{-(n-i)}, \dots, Z_{i-1}\}$ 
2: for  $i = 0; i < m; i ++$  do
3:    $score = CrossAttn(W[0], W[0], W[0])$ 
4:   for  $j = 1; j < n; j ++$  do
5:     for  $k = 0; k < (j + 1); k ++$  do
6:        $score_j = CrossAttn(W[k], W[j], W[j])$ 
7:        $score \leftarrow Concat\ score\ and\ score_j$ 
8:     end for
9:   end for
10:   $score_i^{future} = MLP(score)$ 
11:   $\hat{Z}[i] = CrossAttn(score_i^{future}, P[i], P[i])$ 
12:   $W \leftarrow Concat\ W\ and\ \hat{Z}[i]$ 
13:   $W = W[1:] \leftarrow Slide\ one\ step$ 
14: end for
15: return  $\{\hat{Z}_1, \dots, \hat{Z}_m\}$ 
```

into discrete latent variables window $\{Z_{-(n-1)}, \dots, Z_0\}$, which are then forecasted through a latent variable iteration model. As TC intensity is influenced by current and historical information, we perform cross-attention calculations on the relevant latent variables, concatenate all the results, and pass them through an MLP layer.

However, this latent space iteration approach suffers from the problem of error accumulation. Additionally, the ERA5 data itself tends to underestimate TC intensity. To address this, we calculate the future PI based on the FengWu forecast data and perform cross-attention between this PI and the MLP output. This provides an upper bound on the possible future TC intensity, which is used to constrain the iterative latent variables physically.

Using a residual connection, we add the last latent variable from the historical window to the one-step latent variable forecast result. We then concatenate this forecast result with the historical window, move the sliding window forward by one step, extract the new latent variables, and repeat the above process to complete the latent space iterative forecasting. The full algorithm is described in Algorithm 1.

For the forecasting task, it’s important to note that using future ERA5 data is impossible for real-time forecasting. Therefore, the final approach is to use the FengWu forecast as the conditioning input and decode all the forecast latent variables $\{\hat{Z}_1, \dots, \hat{Z}_m\}$ back to the TC intensity.

Experiments

We test the 5-day long-term forecast of global TC intensity and achieve SOTA performance. To assess its practical applicability, we further conduct real-time forecasting experiments, taking into account the actual operating conditions. Finally, we carry out ablation studies to validate the rationality of our model design.

Experimental Settings

Data Description. The experiment uses global TC-related data from the year 1980 to 2022. As mentioned earlier, the IBTrACS dataset (Knapp et al. 2010) is used as the ground truth for TC intensity, and 69 atmospheric variables are selected from the ERA5 data (Hersbach et al. 2020), with a 10° diameter area around the TC center extracted. The data from 1980 to 2017 is used for training, 2018 for validation, and 2019 to 2020 for testing. For real-time forecasting, we require FengWu forecast data. To address storage constraints, we use FengWu to generate 5-day forecasts for the ERA5 data in 2021 and 2022 and extract the corresponding TC-related variables. The PI is computed through future field calculations, including two-dimensional variables. All variables are normalized, as illustrated in:

$$\tilde{x} = \frac{(x - x_{mean})}{x_{std}}. \quad (10)$$

Metrics. Consistent with most research, we use mean absolute error (MAE) as the performance evaluation metric, calculating the absolute error of I and \hat{I} , with the formula:

$$MAE = \frac{1}{N} \sum_{i=1}^N |I_i - \hat{I}_i|. \quad (11)$$

Meanwhile, the unit of MSW is commonly expressed in m/s, and we convert the experimental results according to the conversion of 1 knot = 0.5144 m/s. Furthermore, when comparing the performance enhancement between different models, the forecast skill is often expressed as:

$$sf(\%) = 100 \times \frac{eb - ef}{eb}, \quad (12)$$

here, eb represents the error of the baseline model, while ef denotes the error of the model currently under evaluation. The value of sf can be easily understood as the percentage reduction in forecast error compared to the baseline model.

Implementation Details. VQLTI is trained using the PyTorch framework on an Nvidia A100 GPU. We employ the Adam optimization algorithm with a learning rate of 0.0001 and apply Exponential Moving Average (EMA) and L1 regularization with a coefficient of 0.00001. The batch size is set to 64, and the model is trained for 30 epochs.

In the first stage of training, the FPN extracts 256-dimensional feature maps with spatial scales of $\{40 \times 40, 20 \times 20, 10 \times 10, 10 \times 10\}$ through convolutional networks (Krizhevsky, Sutskever, and Hinton 2012). These feature maps are then split into patches of size 5 and concatenated, followed by a linear layer to output a dimension of $d_p = 128$. The codebook has a hidden space size $J = 1024$, a dimension $d_z = 128$, and the hyperparameter β is 0.25. The cross-attention computation uses 3 heads and a dropout of 0.1.

In the second stage of training, the input consists of historical information from the past (-18h, -12h, -6h, 0h), denoted by $n = 4$, and the forecast covers the next 2 days, denoted by $m = 8$. Thanks to the design of the VQLTI model, we can easily modify the value of m and the corresponding FengWu forecast data to achieve forecasting at any desired time step without additional training.

| Basin | Methods | | MSLP (hPa) | | | | | MSW (m/s) | | | | |
|-------------------------|----------|--------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | | 24h | 48h | 72h | 96h | 120h | 24h | 48h | 72h | 96h | 120h |
| Global 2019- 2020 | DL | TC-Pre | 7.22 | 11.78 | 14.42 | 15.71 | 15.69 | 5.27 | 8.52 | 10.40 | 11.11 | 11.08 |
| | | MSCAR | 7.13 | 11.52 | 14.33 | 15.68 | 15.37 | 5.56 | 8.76 | 10.72 | 11.58 | 11.31 |
| | | TCIF | 7.94 | 13.48 | 16.31 | 17.00 | 15.97 | 6.28 | 10.41 | 12.38 | 12.90 | 11.89 |
| | NWP | ECMWF | 7.65 | 9.12 | 9.80 | 10.35 | 11.07 | 7.48 | 8.22 | 8.32 | 8.05 | 8.15 |
| | | NCEP | 6.98 | 8.48 | 9.39 | 10.56 | 11.29 | 5.95 | 6.54 | 7.02 | 7.53 | 8.33 |
| | Base | ERA5 | 9.59 | 10.40 | 10.48 | 10.03 | 8.96 | 17.72 | 18.50 | 18.52 | 18.00 | 16.86 |
| | Ours | VQLTI | 5.87 | 6.81 | 7.09 | 7.15 | 6.64 | 4.30 | 4.86 | 5.06 | 5.18 | 5.05 |
| WP 2019 | DL | MMSTN | 9.68 | 19.18 | 26.64 | 29.08 | 28.62 | 5.73 | 11.34 | 16.23 | 18.92 | 19.94 |
| | | MGTCF | 10.20 | 15.07 | 14.79 | 13.32 | 14.26 | 5.55 | 8.25 | 8.07 | 8.01 | 8.24 |
| | Official | JTWC | — | — | — | — | — | 4.60 | 7.30 | 9.70 | 9.80 | 9.10 |
| | Ours | VQLTI | 6.62 | 8.05 | 8.96 | 8.99 | 8.09 | 4.24 | 4.97 | 5.43 | 5.56 | 5.23 |

Table 1: Comparison of long-term TC intensity forecast error with MAE as the metric. The comparison is shown for the global and Western North Pacific (WP) basin. Note that MMSTN and MGTCF produce ensemble forecast results, and for a fair comparison, we calculate the average error, so the results may differ considerably from the original paper.

Long-term Forecast

As shown in Table 1, we compare our global TC forecasts from 2019 to 2020 with some recent deep learning-based methods, including TC-Pre (Zhang et al. 2022), MSCAR (Wang et al. 2024), and TCIF (Wang, Li, and Zheng 2024). We also compare with the well-known NWP models ECMWF-IFS (Bougeault et al. 2010) and NCEP-GFS (Bougeault et al. 2010). To ensure the model achieves effective forecasting, we also present the error of the ERA5 data itself. Note that ERA5 does not have any forecast results. The error here is calculated based on the input of the future ERA5. For MSLP, the minimum value of the *msl* in ERA5 is selected, and for MSW, the maximum value of $\sqrt{(u10)^2 + (v10)^2}$ is used as an approximate substitute, although they are not entirely consistent. Meanwhile, in the Western North Pacific (WP) basin where TC activities are most active, we also make comparisons with the deep learning-based MMSTN (Huang et al. 2022) and MGTCF (Huang et al. 2023), as well as the operational forecasts of Joint Typhoon Warning Center (JTWC) (Chen et al. 2019). The deep learning-based methods mentioned mostly have fixed forecast horizons, so they are all set to a fixed 20-step forecast during training.

In the global long-term forecast comparison, VQLTI demonstrates SOTA performance in 24h to 120h forecasts. The forecast skill of VQLTI not only comprehensively outperforms other deep learning-based methods, but its error accumulation rate is also significantly lower than other deep learning-based methods. TC-Pre, MSCAR, and TCIF have very fast error accumulation rates in short and medium-term forecasts, with their 48h MSLP forecast errors increasing by 63.16%, 61.51%, and 69.77% respectively compared to the 24h errors, while VQLTI’s error only increases by 16.01%. Relative to the NWP models ECMWF-IFS and NCEP-GFS, VQLTI reduces the 24h to 120h MSW forecast error by 35.65%-42.51% and 25.69%-39.38%. The forecast errors of the VQLTI model are all lower than those of ERA5, which demonstrates the effectiveness of the model’s forecasts.

In the WP basin, we only perform comparisons in 2019 due to the unique data used by MMSTN and

MGTCF (Huang et al. 2022, 2023). MMSTN and MGTCF employ deep learning-based ensemble forecasting methods, with the mean of multiple forecasts used as the final result for a fair comparison. The 120-hour MSW forecast error of the VQLTI model is reduced by 71.73% and 36.53% compared to MMSTN and MGTCF, respectively. Furthermore, VQLTI’s forecast error is 42.53% lower than that of JTWC.

Real-time Forecast

Real-time forecasting needs to address two key issues. First, the ERA5 data requires data assimilation and correction, and cannot be obtained in real-time, so we use relatively available real-time analysis data to replace the historical ERA5 inputs. Second, future ERA5 fields are also unavailable, so we substitute them with FengWu’s forecast results. We use 2021 data to fine-tune VQLTI and evaluate it in 2022, as presented in Table 2.

VQLTI (real-time) continues to exhibit excellent performance in real-time forecasting. Its 24h forecast shows negligible changes compared to VQLTI and even experiences a slight improvement in 24h MSW forecasting. VQLTI (real-time)’s long-term forecasting does exhibit some performance degradation, with its 120h MSW forecast error increasing by 20.22% compared to VQLTI. We consider this error reasonable, as FengWu’s forecasts have inherent discrepancies compared to ERA5 (as shown in Figure 2), and the errors in FengWu’s long-term forecasts lead to deviations in the TC positions obtained by the ECMWF’s TC tracker, resulting in misalignment of the extracted spatial information. Although VQLTI (real-time) long-term forecasting has declined to some extent, its forecasting skill remains superior to other deep learning-based methods. Compared to TC-Pre, MSCAR, and TCIF, its 120h MSW forecast error decreases by 37.80%, 43.91%, and 54.60% respectively. Compared to the NWP models ECMWF-IFS and NCEP-GFS, VQLTI (real-time)’s 120h MSW forecast error decreases by 35.85% and 14.26%. We have also calculated the error of the FengWu forecast data, and the VQLTI’s forecast skill is significantly better than FengWu, which verifies the effectiveness of the fine-tuning.

| Methods | | MSLP (hPa) | | | | | MSW (m/s) | | | | |
|---------|-------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | 24h | 48h | 72h | 96h | 120h | 24h | 48h | 72h | 96h | 120h |
| DL | TC-Pre | 7.51 | 11.15 | 12.88 | 13.18 | 13.03 | 5.12 | 7.75 | 8.88 | 9.18 | 8.89 |
| | MSCAR | 7.64 | 11.22 | 13.18 | 13.99 | 14.63 | 5.72 | 8.22 | 9.25 | 9.95 | 9.86 |
| | TCIF | 8.11 | 12.98 | 15.74 | 16.83 | 17.43 | 6.22 | 9.97 | 11.52 | 12.33 | 12.18 |
| NWP | ECMWF | 7.49 | 8.28 | 8.54 | 9.05 | 9.66 | 8.99 | 9.19 | 9.10 | 8.63 | 8.62 |
| | NCEP | 6.31 | 7.63 | 8.77 | 9.29 | 10.18 | 5.72 | 5.87 | 6.10 | 6.07 | 6.45 |
| Base | FengWu | 11.25 | 12.65 | 12.85 | 12.80 | 12.90 | 20.04 | 20.51 | 19.93 | 19.28 | 18.70 |
| Ours | VQLTI | 6.01 | 6.74 | 6.54 | 6.22 | 6.37 | 4.51 | 4.80 | 4.69 | 4.56 | 4.60 |
| | VQLTI (real-time) | 6.42 | 7.98 | 8.29 | 7.97 | 7.86 | 4.47 | 5.58 | 5.78 | 5.69 | 5.53 |

Table 2: Comparison of real-time TC intensity forecast error with MAE as the metric. The comparison is shown for the global region in 2022.

| VQ | FengWu | PI | MSLP (hPa) | | | | | MSW (m/s) | | | | |
|----|--------|----|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | | 24h | 48h | 72h | 96h | 120h | 24h | 48h | 72h | 96h | 120h |
| | | | 22.68 | 25.17 | 25.45 | 25.34 | 25.10 | 15.47 | 17.99 | 18.25 | 18.17 | 17.84 |
| ✓ | | | 20.22 | 22.06 | 21.96 | 21.70 | 21.45 | 15.76 | 17.10 | 17.03 | 16.77 | 16.35 |
| ✓ | ✓ | | 6.60 | 8.56 | 9.44 | 10.24 | 11.35 | 4.64 | 5.92 | 6.34 | 6.96 | 7.31 |
| ✓ | ✓ | ✓ | 6.42 | 7.98 | 8.29 | 7.97 | 7.86 | 4.47 | 5.58 | 5.78 | 5.69 | 5.53 |

Table 3: Ablation experiments based on real-time forecasts in 2022 with MAE as the metric. VQ represents the discretization of latent variables by learning the codebook during training. FengWu represents the introduction of FengWu forecast field information. PI represents the calculation of PI data.

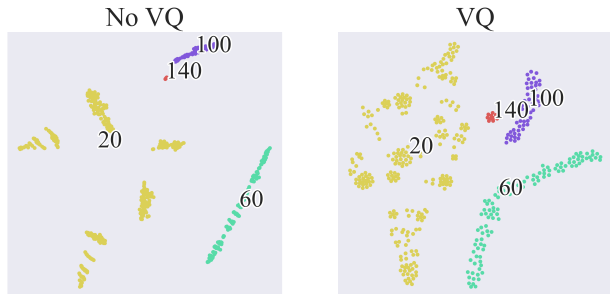


Figure 4: T-SNE visualization of latent variables. The same color represents the same intensity value, and here we use MSW (knot) as an example.

Ablation Studies

This ablation study explores the design rationale of the VQLTI model. As shown in Table 3, it explores whether to use a codebook to quantify latent variables (VQ), whether to introduce FengWu forecast data (FengWu), and whether to introduce PI data (PI). The experimental results show that using codebook for quantization of latent variables is beneficial for long-term forecasting, reducing the MSLP forecasting error by 14.54% for the 120h forecast. FengWu’s forecast field plays a huge role, and introducing the FengWu forecast field, the error in the 24h to 120h MSW forecast is reduced by 55.29%-70.56%. The constraint brought by PI also plays a significant role in improving the long-term forecast skill, reducing the MSLP and MSW forecast errors by 30.75% and 24.35% for the 120h forecast, respectively.

Furthermore, we visualize the latent variables Z of VQLTI with and without codebook quantization during

training using t-SNE (Van der Maaten and Hinton 2008), as shown in Figure 4. We use the same color to represent the same intensity, and we find that the latent variables quantized using codebook not only correctly represent the same intensity, but also maintain a certain intra-class distance. This is consistent with our original intention, where the codebook successfully transfers the TC intensity information to the discrete latent space while ensuring the difference caused by the difference in spatial information between the same intensities.

In summary, these results validate the rationality and effectiveness of the VQLTI model design. For additional validation experiments and further details, please refer to the Appendix.

Conclusion

This study presents VQLTI, a model that achieves highly accurate real-time global long-term tropical cyclone (TC) intensity forecasting. We transfer the TC intensity information into a discrete latent space and fully leverage the forecast fields of the machine learning-based weather prediction model FengWu as spatial information conditions to compensate for the lack of physical knowledge learning in the VQLTI model. Finally, we propose to compute the potential intensity (PI) to provide physical constraints, mitigating the underestimation of TC intensity by the forecast fields. To the best of our knowledge, this is the first real-time TC intensity forecasting model that utilizes the forecast fields of a machine learning-based weather prediction model, and it performs the forecasting in the discrete latent space, resulting in low inference cost and no need for post-processing by forecasters. The experimental results demonstrate the superior performance and the rationality of the design of VQLTI.

Acknowledgments

This research is supported by the National Key R&D Program of China (Grant No.2024YFB4613400), the Anhui Provincial Natural Science Foundation (Grant No.2408085QF214), the Fundamental Research Funds for the Central Universities (Grant No. WK2100000045), the Opening Project of the State Key Laboratory of General Artificial Intelligence(Grant No. SKLAGI2024OP10, Grant No.SKLAGI2024OP11), the National Natural Science Foundation of China under grand No.U19B2044 and the Shanghai Science and Technology Commission Project (23DZ1204704).

References

- Bao, J.; Chen, D.; Wen, F.; Li, H.; and Hua, G. 2017. CVAE-GAN: fine-grained image generation through asymmetric training. In *Proceedings of the IEEE international conference on computer vision*, 2745–2754.
- Bi, K.; Xie, L.; Zhang, H.; Chen, X.; Gu, X.; and Tian, Q. 2023. Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970): 533–538.
- Bougeault, P.; Toth, Z.; Bishop, C.; Brown, B.; Burridge, D.; Chen, D. H.; Ebert, B.; Fuentes, M.; Hamill, T. M.; Mylne, K.; et al. 2010. The THORPEX interactive grand global ensemble. *Bulletin of the American Meteorological Society*, 91(8): 1059–1072.
- Bourdin, S.; Fromang, S.; Dulac, W.; Cattiaux, J.; and Chauvin, F. 2022. Intercomparison of four tropical cyclones detection algorithms on ERA5.
- Buizza, R.; Balmaseda, M. A.; Brown, A.; English, S.; Forbes, R.; Geer, A.; Haiden, T.; Leutbecher, M.; Magnusson, L.; Rodwell, M.; et al. 2018. *The development and evaluation process followed at ECMWF to upgrade the Integrated Forecasting System (IFS)*. European Centre for Medium Range Weather Forecasts.
- Chen, G.; Zhang, X.; Bai, L.; and Wan, R. 2019. Verification of tropical cyclone operational forecast in 2018. *Proceedings of the ESCAP/WMO Typhoon Committee, Guangzhou, China*, 26.
- Chen, K.; Han, T.; Gong, J.; Bai, L.; Ling, F.; Luo, J.-J.; Chen, X.; Ma, L.; Zhang, T.; Su, R.; et al. 2023a. Fengwu: Pushing the skillful global medium-range weather forecast beyond 10 days lead. *arXiv preprint arXiv:2304.02948*.
- Chen, L.; Zhong, X.; Zhang, F.; Cheng, Y.; Xu, Y.; Qi, Y.; and Li, H. 2023b. FuXi: A cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and Atmospheric Science*, 6(1): 190.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissensborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Emanuel, K. 2003. Tropical cyclones. *Annual review of earth and planetary sciences*, 31(1): 75–104.
- Emanuel, K. 2018. 100 years of progress in tropical cyclone research. *Meteorological Monographs*, 59(1): 15–1.
- Emanuel, K. A. 1986. An air-sea interaction theory for tropical cyclones. Part I: Steady-state maintenance. *Journal of Atmospheric Sciences*, 43(6): 585–605.
- Emanuel, K. A. 1987. The dependence of hurricane intensity on climate. *Nature*, 326(6112): 483–485.
- Gilford, D. M. 2021. pyPI (v1. 3): Tropical cyclone potential intensity calculations in Python. *Geoscientific Model Development*, 14(5): 2351–2369.
- Graves, A.; and Graves, A. 2012. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37–45.
- Haiden, T.; Janousek, M.; Vitart, F.; Bouallègue, Z. B.; Fer-ranti, L.; Prates, F.; and Richardson, D. 2018. *Evaluation of ECMWF forecasts, including the 2018 upgrade*. European Centre for Medium Range Weather Forecasts Reading, UK.
- Han, T.; Guo, S.; Ling, F.; Chen, K.; Gong, J.; Luo, J.; Gu, J.; Dai, K.; Ouyang, W.; and Bai, L. 2024. Fengwu-ghr: Learning the kilometer-scale medium-range global weather forecasting. *arXiv preprint arXiv:2402.00059*.
- Hersbach, H.; Bell, B.; Berrisford, P.; Hirahara, S.; Horányi, A.; Muñoz-Sabater, J.; Nicolas, J.; Peubey, C.; Radu, R.; Schepers, D.; et al. 2020. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730): 1999–2049.
- Hsu, K.; Liu, C.-C.; Peng, M.; Chen, D.-S.; Chang, P.-L.; Hsiao, L.-F.; Fong, C.-T.; Hong, J.-S.; Cheng, C.-P.; Lu, K.-C.; et al. 2024. Performance of Five Machine Learning-based Global Weather Prediction Models in the East Asia Region.
- Huang, C.; Bai, C.; Chan, S.; and Zhang, J. 2022. MMSTN: A Multi-Modal Spatial-Temporal Network for Tropical Cyclone Short-Term Prediction. *Geophysical Research Letters*, 49(4): e2021GL096898.
- Huang, C.; Bai, C.; Chan, S.; Zhang, J.; and Wu, Y. 2023. MGTCF: multi-generator tropical cyclone forecasting with heterogeneous meteorological data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 5096–5104.
- Jiang, W.; Zhang, D.; Hu, G.; Wu, T.; Liu, L.; Xiao, Y.; and Duan, Z. 2023. Transformer-based tropical cyclone track and intensity forecasting. *Journal of Wind Engineering and Industrial Aerodynamics*, 238: 105440.
- Knapp, K. R.; Kruk, M. C.; Levinson, D. H.; Diamond, H. J.; and Neumann, C. J. 2010. The international best track archive for climate stewardship (IBTrACS) unifying tropical cyclone data. *Bulletin of the American Meteorological Society*, 91(3): 363–376.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Lam, R.; Sanchez-Gonzalez, A.; Willson, M.; Wirnsberger, P.; Fortunato, M.; Alet, F.; Ravuri, S.; Ewalds, T.; Eaton-Rosen, Z.; Hu, W.; et al. 2023. Learning skillful medium-range global weather forecasting. *Science*, 382(6677): 1416–1421.

- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Magnusson, L.; Majumdar, S.; Emerton, R.; Richardson, D.; Alonso-Balmaseda, M.; Baugh, C.; Bechtold, P.; Bidlot, J.-R.; Bonanni, A.; Bonavita, M.; Bormann, N.; Brown, A.; Browne, P.; Carr, H.; Dahoui, M.; Chiara, G. D.; Diamantakis, M.; Duncan, D.; English, S.; Forbes, R.; Geer, A.; Haiden, T.; Healy, S.; Hewson, T.; Ingleby, B.; Janousek, M.; Kuehnlein, C.; Lang, S.; Lock, S.-J.; McNally, T.; Mogenssen, K.; Pappenberger, F.; Polichtchouk, I.; Prates, F.; Prudhomme, C.; Rabier, F.; de Rosnay, P.; Quintino, T.; and Rennie, M. 2021. Tropical cyclone activities at ECMWF.
- Malakar, P.; Kesarkar, A.; Bhate, J.; Singh, V.; and Deshamukhya, A. 2020. Comparison of reanalysis data sets to comprehend the evolution of tropical cyclones over North Indian Ocean. *Earth and Space Science*, 7(2): e2019EA000978.
- Meng, F.; Yang, K.; Yao, Y.; Wang, Z.; and Song, T. 2023. Tropical cyclone intensity probabilistic forecasting system based on deep learning. *International Journal of Intelligent Systems*, 2023(1): 3569538.
- Pan, B.; Xu, X.; and Shi, Z. 2019. Tropical cyclone intensity prediction based on recurrent neural networks. *Electronics Letters*, 55(7): 413–415.
- Rasul, K.; Bennett, A.; Vicente, P.; Gupta, U.; Ghonia, H.; Schneider, A.; and Nevmyvaka, Y. 2024. VQ-TR: Vector Quantized Attention for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*.
- Rasul, K.; Park, Y.-J.; Ramström, M. N.; and Kim, K.-M. 2022. Vq-ar: Vector quantized autoregressive probabilistic time series forecasting. *arXiv preprint arXiv:2205.15894*.
- Rüttgers, M.; Lee, S.; Jeon, S.; and You, D. 2019. Prediction of a typhoon track using a generative adversarial network and satellite images. *Scientific reports*, 9(1): 6057.
- Shen, X.; Su, Y.; Zhang, H.; and Hu, J. 2023. New version of the CMA-GFS dynamical core based on the predictor–corrector time integration scheme. *Journal of Meteorological Research*, 37(3): 273–285.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, C.; Li, X.; and Zheng, G. 2024. Tropical cyclone intensity forecasting using model knowledge guided deep learning model. *Environmental Research Letters*, 19(2): 024006.
- Wang, X.; Chen, K.; Liu, L.; Han, T.; Li, B.; and Bai, L. 2024. Global Tropical Cyclone Intensity Forecasting with Multi-modal Multi-scale Causal Autoregressive Model. *arXiv preprint arXiv:2402.13270*.
- Woodruff, J. D.; Irish, J. L.; and Camargo, S. J. 2013. Coastal flooding by tropical cyclones and sea-level rise. *Nature*, 504(7478): 44–52.
- Yang, J.; Li, L.; Zhao, K.; Wang, P.; Wang, D.; Sou, I. M.; Yang, Z.; Hu, J.; Tang, X.; Mok, K. M.; et al. 2019. A comparative study of Typhoon Hato (2017) and Typhoon Mangkhut (2018)—Their impacts on coastal inundation in Macau. *Journal of Geophysical Research: Oceans*, 124(12): 9590–9619.
- Zhang, X.; Fang, J.; and Yu, Z. 2023. The forecast skill of tropical cyclone genesis in two global ensembles. *Weather and Forecasting*, 38(1): 83–97.
- Zhang, Z.; Yang, X.; Shi, L.; Wang, B.; Du, Z.; Zhang, F.; and Liu, R. 2022. A neural network framework for fine-grained tropical cyclone intensity prediction. *Knowledge-Based Systems*, 241: 108195.
- Zhou, X.; and Juang, H.-M. H. 2023. A Model Instability Issue in the NCEP Global Forecast System Version 16 and Potential Solutions. *EGU sphere*, 2023: 1–22.