

SwinRDM: Integrate SwinRNN with Diffusion Model towards High-Resolution and High-Quality Weather Forecasting

Lei Chen*, Fei Du*, Yuan Hu, Zhibin Wang, Fan Wang

Alibaba Group

{fanjiang.cl, dufei.df, lavender.hy, zhibin.waz, fan.w}@alibaba-inc.com

Abstract

Data-driven medium-range weather forecasting has attracted much attention in recent years. However, the forecasting accuracy at high resolution is unsatisfactory currently. Pursuing high-resolution and high-quality weather forecasting, we develop a data-driven model SwinRDM which integrates an improved version of SwinRNN with a diffusion model. SwinRDM performs predictions at 0.25-degree resolution and achieves superior forecasting accuracy to IFS (Integrated Forecast System), the state-of-the-art operational NWP model, on representative atmospheric variables including 500 hPa geopotential (Z500), 850 hPa temperature (T850), 2-m temperature (T2M), and total precipitation (TP), at lead times of up to 5 days. We propose to leverage a two-step strategy to achieve high-resolution predictions at 0.25-degree considering the trade-off between computation memory and forecasting accuracy. Recurrent predictions for future atmospheric fields are firstly performed at 1.40625-degree resolution, and then a diffusion-based super-resolution model is leveraged to recover the high spatial resolution and finer-scale atmospheric details. SwinRDM pushes forward the performance and potential of data-driven models for a large margin towards operational applications.

Introduction

Accurate weather forecasting is beneficial to human beings in several areas such as agriculture, energy, and public transportation. Numerical Weather Prediction (NWP) has long been adopted for weather forecasting. It has been improved by better physics parameterization techniques and high-quality atmospheric observations in the past few decades. However, this approach requires huge amounts of computing power, which may limit its application in industry.

With the development of machine learning (ML), especially deep learning (DL) techniques, many studies are employing data-driven DL methods to forecast atmospheric variables. The purely data-driven DL models are often orders of magnitude faster than the NWP model. However, the performance of current DL models is unsatisfactory currently. To facilitate the development of data-driven weather forecasting, some benchmarks (Rasp et al. 2020; Garg,

Rasp, and Thuerey 2022; de Witt et al. 2020) are constructed to enable a thorough comparison of different methods. Among them, WeatherBench (Rasp et al. 2020) is one of the widely used benchmarks, which is constructed by re-gridding the ERA5 reanalysis dataset (Hersbach et al. 2020) from 0.25° resolution to three different resolutions (i.e., 5.625°, 2.8125° and 1.40625°). It focuses on the medium-range global prediction of a few key variables at lead times of up to 5 days. Several works have tried to improve the prediction performance on WeatherBench (Rasp and Thuerey 2021; Weyn, Durran, and Caruana 2020; Hu et al. 2022). Among them, SwinVRNN (Hu et al. 2022) achieves the best performance by integrating a variational recurrent neural network (SwinRNN) with a feature perturbation module.

Although these works have achieved great success in global weather forecasting, their methods are built on low-resolution (usually lower than 1°) data. The largest resolution of WeatherBench is 1.40625°, corresponding to a 128×256 pixels grid. And the distance between every two pixels is larger than 100km, which is too coarse for a forecasting model to capture the fine-scale dynamics (Pathak et al. 2022). (Keisler 2022) builds a graph neural network (GNN) to forecast global weather on the 1-degree scale. It shows comparable performance on wind and relative humidity to the IFS. However, its resolution is still relatively low. FourCastNet (Pathak et al. 2022) trains an Adaptive Fourier Neural Operator (AFNO) model directly at 0.25° resolution ERA5 dataset, which achieves comparable performance to the IFS at short-range lead times and can resolve many important small-scale phenomena. However, there still exists a performance gap between data-driven models and the IFS at lead times of up to 5 days, especially on representative variables such as Z500 and T850.

In this paper, we focus on building a global weather forecasting model at 0.25° resolution and propose a SwinRDM model by integrating an improved SwinRNN (Hu et al. 2022) with a diffusion-based super-resolution model. Since SwinRNN achieves superior performance at low resolution, we employ it as our base model. We experimentally analyze the SwinRNN model and build an improved version named SwinRNN+ by replacing the multi-scale network with a single-scale design and adding a feature aggregation layer. Our SwinRNN+ achieves higher performance than IFS on all key variables at lead times of up to 5 days at 1.40625°

*These authors contributed equally.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

resolution. Note that this is a considerable improvement compared to SwinRNN that can only compete with the IFS model on surface-level variables at 5.625° resolution. Different from FourCastNet, to generate high-resolution global weather prediction, we resort to the super-resolution (SR) technique rather than directly train the SwinRNN+ model on 0.25° resolution data due to the prohibitive computational cost. The super-resolution task is implemented using a conditional diffusion model (Saharia et al. 2021b), which trains a U-Net model (Ronneberger, Fischer, and Brox 2015) to iteratively refine the outputs starting from pure Gaussian noises. This model is shown to be able to generate photo-realistic outputs compared to traditional super-resolution models (Saharia et al. 2021b). We show in this work that the diffusion model-based super-resolution conditioned on low-resolution predictions can capture small-scale variations and generate high-quality weather forecasting at high resolution.

Our contribution can be summarized as follows:

- We conduct experimental studies on the SwinRNN model and propose an improved version — SwinRNN+. It achieves superior performance than the state-of-the-art IFS model on all representative variables at the resolution of 1.40625° and lead times of up to 5 days.
- We employ a conditional diffusion model for super-resolution conditioned on SwinRNN+ outputs, which achieves high-quality weather forecasting at the resolution of 0.25° with an optimal trade-off between computation cost and forecast accuracy.
- Experimental results on the ERA5 dataset show that our SwinRDM model not only outperforms IFS but also achieves high-quality forecasts with finer-scale details, which sets a solid baseline for data-driven DL weather forecasting models.

Related Works

In this section, we briefly review some deep learning-based weather forecasting methods and super-resolution methods.

Deep Learning-based Weather Forecasting

Deep learning has been investigated to perform data-driven weather forecasting in recent years, and the goal is to fully replace the NWP model. Some works focus on a particular local area (Shi et al. 2017; Sønderby et al. 2020; Bihlo 2021). However, using data in a local spatial domain may result in uncertainty around the boundary regions (Bihlo 2021). For global weather forecasting, some widely used networks in computer vision has been applied, including ResNet (Rasp and Thuerey 2021), U-Net (Weyn, Durran, and Caruana 2020), VAE (Hu et al. 2022), and GNNs (Keisler 2022). Among global weather forecasting methods, SwinVRNN (Hu et al. 2022) is the first work that can compete with the IFS on representative surface-level variables (T2M and TP) at lead times of up to 5 days. It constructs a deterministic SwinRNN model based on the Swin Transformer block (Liu et al. 2021) and builds a perturbation module to perform ensemble forecasting. However,

the resolution of SwinVRNN is relatively low, and it cannot achieve comparable performance on pressure-level variables such as Z500 and T850. FourCastNet (Pathak et al. 2022) is the first work that directly builds the network on the 0.25° resolution data. Although it achieves high performance on short timescales, it cannot compete with the IFS at a 5-day lead time. In this paper, we also intend to predict the atmospheric variables at 0.25° resolution. We improve the SwinRNN network at low resolution and employ the diffusion-based super-resolution model to generate high-resolution and high-quality results.

Super Resolution

The goal of super-resolution is to reconstruct a high-resolution image from a low-resolution counterpart. The simplest way to realize super-resolution is interpolation. It is computationally efficient but often suffers from detail loss in regions with complex textures (Chen et al. 2022). SR-CNN (Dong et al. 2015) is a pioneering work that exploits CNNs to perform super-resolution. Later, many deep learning methods (Kim, Lee, and Lee 2016; Lim et al. 2017; Zhang et al. 2018; Liang et al. 2021) have been proposed to improve super-resolution performance. These methods usually employ a reconstruction loss such as MSE loss to train the network. However, the network simply trained with a reconstruction loss can hardly capture high texture details and generate perceptually satisfactory results (Bihlo 2021). Some works (Ledig et al. 2017; Wang et al. 2018, 2021) employ generative adversarial networks (GANs) (Goodfellow et al. 2014) to encourage the generator to output high-resolution images that are hard to distinguish from the real high-resolution images. Although these methods can generate high-quality images, GAN-based methods are difficult to optimize (Arjovsky, Chintala, and Bottou 2017).

Recently, diffusion models (Sohl-Dickstein et al. 2015) have attracted much attention in image generation since it is able to generate high-quality images (Song and Ermon 2019; Ho, Jain, and Abbeel 2020; Song and Ermon 2020) that are comparable to GANs. A U-Net architecture is trained with a denoising objective to iteratively refine the outputs starting from pure Gaussian noise. Diffusion models have also been successfully applied to image super-resolution, such as SR3 (Saharia et al. 2021b) which employs a diffusion model to generate realistic high-resolution images conditioned on low-resolution images. Since the diffusion model is proven to generate high-quality images, we exploit this model in weather forecasting to generate high-quality and high-resolution predictions.

Methodology

We formulate the high-resolution weather forecasting problem as a combination of weather forecasting at low resolution and super-resolution to high resolution. The proposed model first recurrently forecasts the future atmospheric variables via a recurrent neural network and then reconstructs the high-resolution results from the predicted low-resolution counterparts via a super-resolution network. The recurrent neural network is built based on SwinRNN

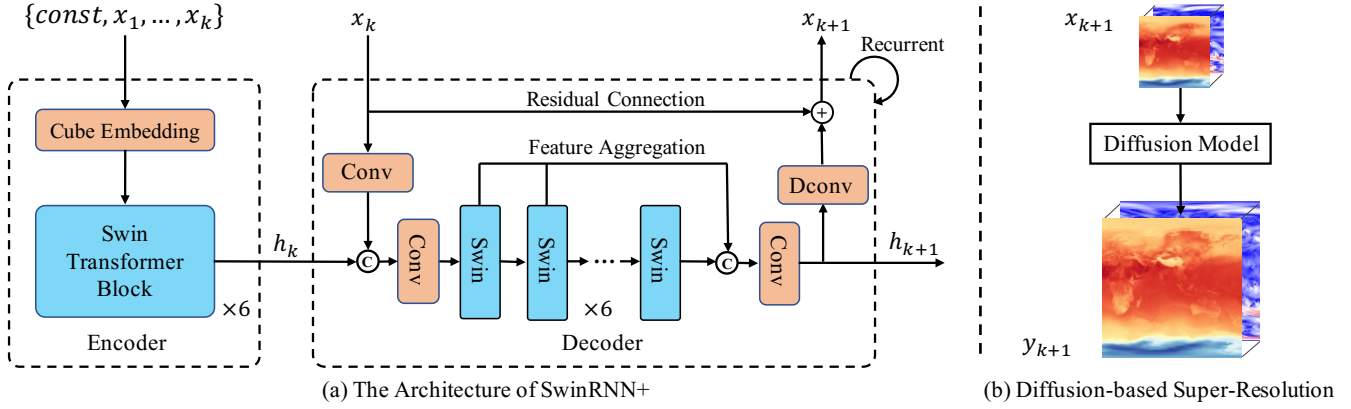


Figure 1: The SwinRDM consists of two parts: (a) the low-resolution forecasting model SwinRNN+ is an improved version of SwinRNN, which adopts a single-scale architecture and adds a multi-layer feature aggregation component, and (b) the diffusion-based super-resolution model conditions on the prediction x_{k+1} from SwinRNN+.

(Hu et al. 2022). We experimentally analyze the architecture of SwinRNN and design an improved version named SwinRNN+ which considerably improves the performance of SwinRNN at 1.40625° resolution. However, SwinRNN+ can hardly be directly trained at 0.25° because the training recurrent steps have to be reduced due to limited computational resources, which will lead to inferior performance. Thus, we adopt the super-resolution technique to achieve high-resolution prediction. To generate high-resolution and high-quality results at 0.25° resolution, we build our super-resolution model based on the diffusion model, which can help capture fine-grained scales compared to traditional super-resolution methods. The architecture of our method is demonstrated in Figure 1.

Background on SwinRNN

The SwinRNN mainly consists of a multi-scale encoder for historical context information extraction and a multi-scale decoder for hidden state propagation and variable prediction at each recurrent step. The atmospheric variables at each time step are stacked together with a shape $C_{in} \times H \times W$, where C_{in} denotes the number of atmospheric variables, and $H \times W$ denotes the global grid size. The stacked result can be regarded as a multi-channel frame similar to many computer vision tasks. The encoder takes k consecutive historical frames as input. It first employs a 3D convolutional network-based cube embedding block to project all frames to features with a size $C \times H \times W$ and then extracts four-scale features $(h_k^1, h_k^2, h_k^3, h_k^4)$ via a hierarchical Swin Transformer. The historical context information is embedded in the features, and they are used to initialize the hidden states of the decoder at time step k . Then in each future time step, the decoder takes as input the combination of hidden states $(h_k^1, h_k^2, h_k^3, h_k^4)$ and current frame x_k , and outputs the updated hidden states $(h_{k+1}^1, h_{k+1}^2, h_{k+1}^3, h_{k+1}^4)$ for next time step and the predicted frame x_{k+1} . The results of SwinRNN show that it achieves higher performance than the IFS on T2M and TP variables.

Although SwinRNN achieves high performance on surface-level variables, it is only trained at 5.625° (32×64) resolution and cannot compete with the IFS model on pressure-level variables such as Z500 and T850. Based on its Swin Transformer-based recurrent structure, we propose SwinRNN+ that is able to achieve superior performance on all key surface-level and pressure-level variables than the IFS at 1.40625° resolution.

SwinRNN+

It is non-trivial to transfer the SwinRNN to high-resolution data even at 1.40625° (128×256) since the memory consumption in the training stage is quadratic to resolution (e.g., 1 vs. 16 for 5.625° vs. 1.40625°) for the Swin Transformer architecture. Thus, it is important to balance the capacity of the network and the computational cost. The improvements of our SwinRNN+ over SwinRNN are two folds. First, we replace the multi-scale network with a single-scale network with higher feature dimensions. Second, we fuse the output features of all Swin Transformer layers in the decoder to generate the hidden states and the output predictions.

Trade Multi-Scale Design for Higher Feature Dimensions. To increase the capacity of the network, a straightforward way is to increase the dimension of the feature. However, the memory cost increases dramatically with the increase of the feature dimension since there are several recurrent steps during training. We observed from (Hu et al. 2022) that SwinRNN benefits little from the multi-scale architecture, whereas the structure significantly increases the number of parameters and memory consumption. Thus, we conduct an ablation experiment to compare the performance of multi-scale and single-scale structures on different feature dimensions. The experimental results show that the multi-scale network generally achieves better performance compared to the single-scale network with the same feature dimensions. However, with the increase of the feature dimension, the performance gap between the two different structures is narrowed rapidly, whereas the memory and the pa-

rameters of the multi-scale structure increase dramatically. Notably, the single-scale network with a high feature dimension shows better performance and higher efficiency compared to the multi-scale network with a low feature dimension. Thus, we draw a conclusion that the multi-scale architecture limits the potential of SwinRNN, and it is more effective to increase the feature dimension of the single-scale network rather than use a multi-scale structure.

Multi-Layer Feature Aggregation. Our second improvement is to aggregate features from multiple layers to learn the hidden states. As shown in Figure 1, the decoder fuses features from the 6 layers to update the hidden states via a convolutional layer, while the original SwinRNN only treats the features from the final layer as the hidden states. Our multi-layer feature aggregation network has two advantages compared to SwinRNN. First, the representation power of the hidden states is improved, which is beneficial to the prediction in the current time step and feature propagation to the next time step. Second, the gradient backward propagation path is reduced, and the information can be more easily propagated back to former time steps, which eases the optimization of the recurrent network.

As shown in Figure 1, the proposed network consists of 6 Swin Transformer blocks with the same scale in both the encoder and decoder. To enable training on 1.40625° resolution data, we first use a patch size of 2×2 to split the image into non-overlapping patches. This is achieved by a convolutional layer with a kernel size of 2 and a stride of 2 in the cube embedding block. Thus, we have a hidden state h_k with a size of $C \times H/2 \times W/2$. In the decoder, the frame x_k is also embedded with a convolutional layer with the same settings, and x_{k+1} is predicted by a transposed convolutional layer. Features from all layers in the decoder are aggregated to increase the representation power.

Diffusion Model for Super Resolution

In order to achieve high-resolution weather forecasting, we integrate the SwinRNN+ with a super-resolution component based on the diffusion model since it is able to generate realistic images with rich details (Saharia et al. 2021b), which can help resolve fine-grained features and generate high-quality and high-resolution forecasting results.

Diffusion models (Ho, Jain, and Abbeel 2020) are a class of generative models consisting of a forward process (or diffusion process) that destroys the training data by successive addition of Gaussian noise and a reverse process that learns to recover the data by reversing this noising process. More specifically, given a sample from data distribution $y_0 \sim q(y_0)$, the diffusion process is a Markov chain that gradually adds Gaussian noise to the data according to a fixed variance schedule β_1, \dots, β_T :

$$q(y_t|y_{t-1}) = \mathcal{N}(y_t; \sqrt{1 - \beta_t}y_{t-1}, (\beta_t)\mathcal{I}). \quad (1)$$

If the magnitude β_t of the noise added at each step is small enough, and the total step T is large enough, then y_T is equivalent to an isotropic Gaussian distribution. It is convenient to produce samples from a Gaussian noise input $y_T \sim \mathcal{N}(0, \mathcal{I})$ by reversing the above forward process.

| Dim. | Fusion | Z500($m^2 s^{-2}$) | T850(T) | T2M(T) | TP(mm) |
|------|--------|----------------------|-------------|------------|------------|
| 128 | | 456 | 2.354 | 2.196 | 2.265 |
| 128 | ✓ | 386 | 2.050 | 1.926 | 2.182 |
| 256 | | 394 | 2.092 | 1.957 | 2.207 |
| 256 | ✓ | 371 | 1.971 | 1.843 | 2.148 |

Table 1: The RMSE results of the SwinRNN+ with and without multi-layer feature aggregation on different feature dimensions.

However, the posterior $q(y_{t-1}|y_t)$ need for sampling is hard to compute, and we need to learn a model parameterized by θ to approximate these conditional probabilities:

$$p_\theta(y_{t-1}|y_T) = \mathcal{N}(\mu_\theta(y_t), \Sigma_\theta(y_t)). \quad (2)$$

While there exists a tractable variational lower-bound on $\log p_\theta(y_0)$, better results arise from optimizing a surrogate denoising objective:

$$E_{\epsilon \sim \mathcal{N}(0, \mathcal{I}), t \sim [0, T]} [\|\epsilon - \epsilon_\theta(y_t, t)\|^2], \quad (3)$$

where $y_t \sim q(y_t|y_0)$ is obtained by applying Gaussian noise ϵ to y_0 , and ϵ_θ is the model to predict the added noise.

Diffusion models can be conditioned on class labels, text, or low-resolution images (Dhariwal and Nichol 2021; Ho et al. 2022; Saharia et al. 2021b,a; Whang et al. 2022; Nichol et al. 2021; Ramesh et al. 2022). In our case, we make it condition on the low-resolution output of SwinRNN+ for the super-resolution task. During training, the low-resolution data x_k is generated on the fly, and forecasting quality is decreased with the lead time. To account for such variation, the model additionally conditions on the time step k of SwinRNN+. Thus, we have a posterior

$$p_\theta(y_k^{(t-1)}|y_k^{(t)}, x_k, t, k), \quad (4)$$

where y_k is the corresponding high-resolution target of x_k . Instead of using the ϵ -prediction formulation, we predict the original targets y_k directly, following (Ramesh et al. 2022). The model acts like a denoising function, trained using a mean squared error loss:

$$E_{t \sim [0, T], y_k^{(t)} \sim q_t} [\|y_k - f_\theta(y_k^{(t)}, x_k, t, k)\|^2]. \quad (5)$$

For the sampling process, given the same low-resolution input x_k , the diffusion model can produce diverse outputs y_k starting from different Gaussian noise samples. Such property makes it possible to perform ensemble forecasting, which is an effective way to improve forecast skills. Thus, we achieve super-resolution and ensemble forecasting at the same time with the diffusion model.

Experiments

Experimental Setup

Dataset We evaluate our proposed SwinRDM method on the ERA5 dataset (Hersbach et al. 2020) provided by the ECMWF. ERA5 dataset is an atmospheric reanalysis dataset, which consists of atmospheric variables at a 0.25° spatial resolution from 1979 to the present day. Data from

| Feat. Dim. | Multi-Scale | Mem. | Params | Z500 | T850 |
|------------|-------------|-------|--------|------|-------|
| 128 | | 12.0G | 4.0M | 417 | 2.182 |
| 128 | ✓ | 13.1G | 62.5M | 386 | 2.050 |
| 256 | | 14.1G | 15.5M | 374 | 1.998 |
| 256 | ✓ | 23.7G | 248.6M | 371 | 1.971 |
| 384 | | 19.4G | 34.8M | 359 | 1.929 |
| 512 | | 25.2G | 61.6M | 354 | 1.912 |

Table 2: Comparison (RMSE) between the single-scale and multi-scale architectures. The units are the same as Table 1.

| Methods | FID | Z500 | T850 | T2M | TP |
|----------|-----|------|------|------|------|
| Bilinear | 235 | 374 | 2.00 | 1.89 | 2.18 |
| SwinIR | 227 | 376 | 1.99 | 1.68 | 2.16 |
| SwinRDM | 60 | 378 | 2.11 | 1.73 | 2.79 |
| SwinRDM* | 145 | 374 | 2.02 | 1.65 | 2.30 |

Table 3: Quantitative comparison (FID and RMSE) of different methods at a lead time of 120 hours. The units are the same as Table 1.

1979 to 2016 are chosen as the training set, and data from 2017 and 2018 are used for evaluation, following (Rasp et al. 2020). We sub-sample the dataset at 6-hour intervals to train our model as in (Pathak et al. 2022). The input signal of the SwinRDM contains 71 variables, including geopotential, temperature, relative humidity, longitude-direction wind, and latitude-direction wind at 13 vertical layers, four single-layer fields (2m temperature, 10m wind, and total precipitation), and two constant fields (land-sea mask and orography).

Evaluation Metrics We follow (Rasp et al. 2020) to evaluate forecast quality using latitude-weighted RMSE (root-mean-square error). In addition, we adopt Fréchet inception distance (FID) (Heusel et al. 2017) to quantitatively assess the sample fidelity of the super-resolution outputs.

Implementation Details We end-to-end train two models (i.e., SwinRNN+ and SwinRDM) for 50 epochs with a batch size of 16. SwinRNN+ is trained on 1.40625° data, while SwinRDM takes 1.40625° data as input and outputs 0.25° predictions. During training, our models take 6 historical frames as input and recurrently predict 20 frames at 6-hour intervals. For SwinRDM, we randomly select one of them to train the diffusion-based super-resolution model. The dimensions of the feature for the encoder and the decoder are set to 768 and 512, respectively. The cosine learning rate policy is used with initial learning rates of 0.0003 for SwinRNN+ and 0.0002 for the diffusion model. The models are optimized by AdamW using PyTorch on 8 NVIDIA A100 GPUs. For the diffusion model, we adopt the implementation in (Nichol and Dhariwal 2021) and use 10 sampling steps to get decent results during inference.

Ablation Study on SwinRNN+

For efficient ablation studies, we train all networks for only 10 epochs.

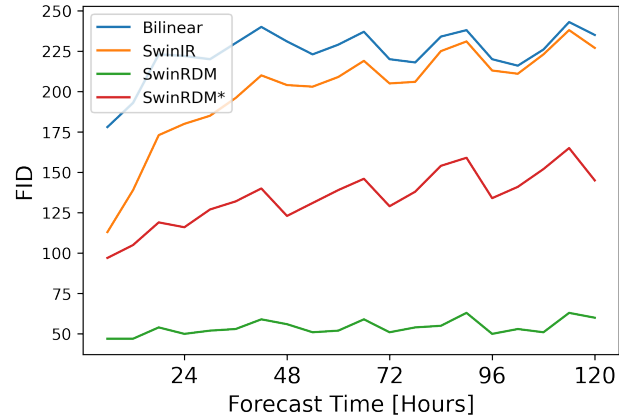


Figure 2: FID in different forecasting times.

Effectiveness of multi-layer feature aggregation The multi-layer feature aggregation component fuses features from all Swin Transformer layers of the decoder to generate the hidden states, which can help improve the representation power and ease the optimization process. Table 1 shows the comparison of SwinRNN+ with and without the aggregation layer. The aggregation layer improves the performance of the networks with different feature dimensions, which verifies the effectiveness of the multi-layer feature aggregation.

Trade-off between multi-scale design and higher feature dimensions The number of parameters and the training memory cost increases significantly with the feature dimensions for the multi-scale architecture of the original SwinRNN, which limits the improvement of the model capacity since the computational cost is prohibitive if we use a large feature dimension. We conduct ablation experiments to compare the performance and potential of the multi-scale structure and the single-scale structure. Table 2 shows the comparison results. For both the encoder and decoder, the multi-scale network contains 4 scales, and each scale consists of 6 Swin Transformer blocks, while the single-scale network only has one scale with 6 Swin Transformer blocks. As can be seen from the table, when the feature dimension is set to 128, there exists a large performance gap between the multi-scale structure and the single-scale structure. However, the performance gap is narrowed significantly when the feature dimension increases from 128 to 256. Meanwhile, the training memory and the parameters of the multi-scale structure increase dramatically with the feature dimensions. It is noticeable that the single-scale network with 384-dimensional features shows better performance and less memory cost compared to the multi-scale network with 256-dimensional features. The dimensions can be set to 512 for the single-scale network, while it is unaffordable for the multi-scale one. Thus, instead of using a multi-scale structure to improve the model capacity, it is more effective to increase the feature dimensions.

Ablation Study on Diffusion-based SR

In this subsection, we compare three variants of the super-resolution methods: bilinear interpolation, SwinIR (Liang

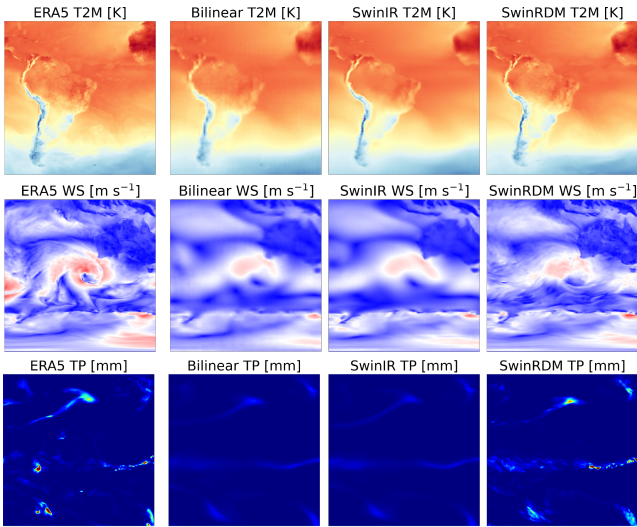


Figure 3: Visual comparison of example fields at the initialization time of June 4, 2018, 00:00 UTC. The first column shows the ERA5 fields for Z500, WS10, and TP at a lead time of 120 hours. The second to fourth columns represent the corresponding forecasting of different SR methods.

| Methods | CSI2 | CSI5 | CSI10 | CSI20 | CSI50 |
|----------|-------|-------|-------|-------|-------|
| Bilinear | 0.171 | 0.051 | 0.006 | 0.000 | 0.000 |
| SwinIR | 0.186 | 0.069 | 0.013 | 0.001 | 0.000 |
| SwinRDM | 0.246 | 0.179 | 0.111 | 0.046 | 0.010 |
| SwinRDM* | 0.262 | 0.190 | 0.123 | 0.049 | 0.006 |

Table 4: Critical success index (CSI) of six-hour accumulate total precipitation (TP) with different thresholds, i.e., 2 mm, 5 mm, 10 mm, 20 mm, and 50 mm.

et al. 2021), and our diffusion model. Bilinear interpolation simply upsamples the low-resolution output. SwinIR is a state-of-the-art SR method constructed by several residual Swin Transformer blocks. All methods are jointly trained with the same low-resolution forecasting model SwinRNN+. In addition to the RMSE metric, FID is used here to evaluate the ability to generate realistic results. FID is computed over six variables by modifying the input channels and weights of the Inception V3 model (Szegedy et al. 2016).

Diffusion-based SR achieves high visual quality As can be seen from Table 3, there is little difference between different super-resolution methods in terms of the RMSE metrics of all variables. However, the FID metrics are significantly different. The bilinear interpolation obtains the worst FID score, and SwinIR slightly improves it by 8. Equipped with the diffusion-based super-resolution model, our SwinRDM considerably improves the FID score by 175, indicating that the diffusion-based super-resolution model can help generate high-quality and realistic results. SwinRDM* is a 10-member ensemble version of our SwinRDM, and it makes a good trade-off between the RMSE and FID. Figure 2 shows the FID scores for different lead times. For a data-driven recurrent forecasting model, the predictions may be

| Methods | Z500 | T850 | T2M | TP |
|----------|----------------|------------------|------------------|------------------|
| IFS | 154/334 | 1.36/2.03 | 1.35/1.77 | 2.36/2.59 |
| SwinRNN | 207/392 | 1.39/2.05 | 1.18/1.63 | 2.01/2.14 |
| SwinRNN+ | 152/316 | 1.12/1.75 | 0.99/1.42 | 1.88/2.07 |
| SwinRDM | 156/316 | 1.23/1.83 | 1.07/1.49 | 2.02/2.24 |
| SwinRDM* | 153/313 | 1.15/1.76 | 1.01/1.43 | 1.87/2.06 |

Table 5: Comparison with state-of-the-art IFS and SwinRNN. The RMSE scores for 3 and 5 days forecast times in the years 2017 and 2018 are shown. SwinRDM* is the 10-member ensemble version of SwinRDM. The units are the same as Table 1.

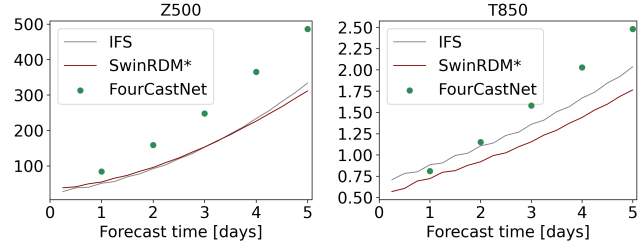


Figure 4: Comparison with IFS and FourCastNet on the test data of 2018. The RMSE of Z500 and T850 is shown.

come smoother with the increase in the lead time. The FID scores for bilinear interpolation and SwinIR increase with the forecast time, while our SwinRDM keeps low FID scores for lead times of up to 5 days. This again shows the good property of the diffusion-based super-resolution model. Although our SwinRDM* increases the FID scores of SwinRDM, it still maintains relatively stable FID scores compared with bilinear interpolation and SwinIR.

High visual quality means high forecasting quality The qualitative results of different SR methods are shown in Figure 3. SwinRDM successfully captures small-scale structures and generates high-quality results at a 5-day lead time. By contrast, both Bilinear and SwinIR produce blurry results and fail to generate rich details. These details are essential for weather forecasting, especially for variables that have complex structures and variations, such as wind speed (WS) and TP. To further verify the superiority of our method, we calculate the critical success index (CSI) (Jolliffe and Stephenson 2012) for TP. The CSI can indicate the prediction performance under different thresholds. As shown in Table 4, SwinRDM surpasses the strong baseline SwinIR by a large margin, which increases the CSI2, CSI5, and CSI10 by 6%, 11%, 10%, respectively. For higher CSI thresholds (20mm and 50mm), only our SwinRDM and SwinRDM* can get decent results, indicating that our methods can forecast extreme precipitation more effectively. Thus, our diffusion-based super-resolution model can generate high-resolution and high-quality forecasting results.

Comparison with State-of-the-art Methods

We compare our methods with state-of-the-art operational IFS and data-driven methods. Since the high-resolution IFS

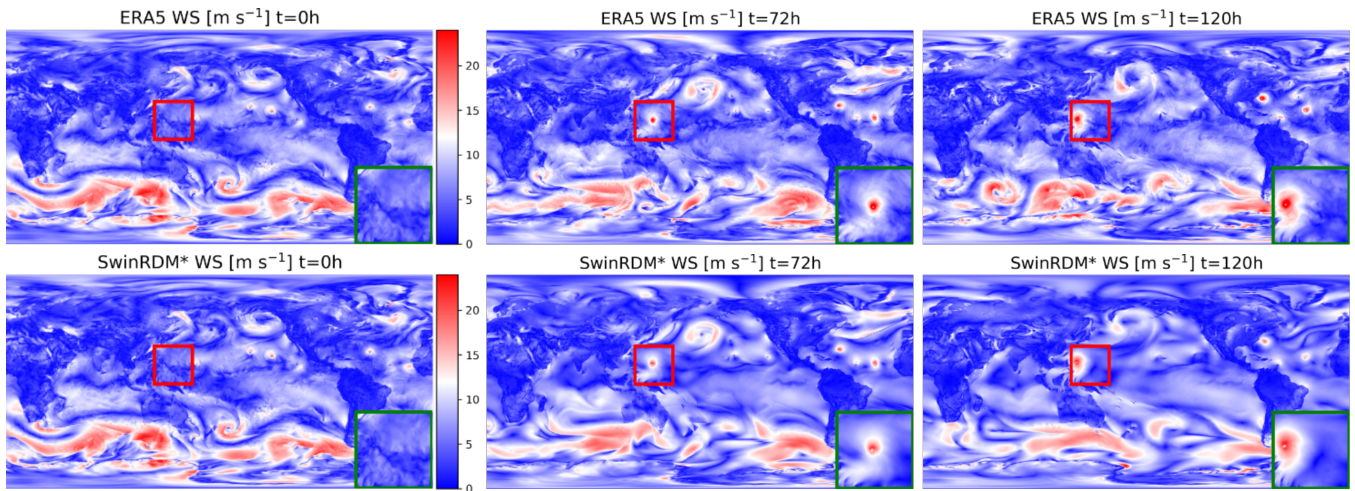


Figure 5: Qualitative illustration of a global near-surface wind forecast generated by our SwinRDM*. The prediction starts at the initial time of September 8, 2018, 06:00 UTC. The zoom-in area shows the beginning of Super Typhoon Mangkhut. Our method successfully forecasts Super Typhoon Mangkhut with high accuracy and rich fine-scale features.

results are not available online due to the data center migration currently, we regrid our results to 5.625° so that we can compare them to the IFS results provided by the WeatherBench (Rasp et al. 2020). Different methods can be directly compared due to the equivalence of evaluation on different resolutions. As stated in (Rasp et al. 2020), there is nearly no difference for evaluation at different resolutions. We have also verified this statement. For data-driven methods, we choose the SwinRNN (Hu et al. 2022) and FourCastNet (Pathak et al. 2022) for comparison. To the best of our knowledge, SwinRNN is the best method at 5.625° resolution, and FourCastNet is the best method at 0.25° resolution. Table 5 shows the comparison results for the years 2017 and 2018. Our SwinRNN+ method is trained on 1.40625° resolution data, and it achieves significantly better performance compared to SwinRNN, showing the effectiveness of our improvement. The SwinRNN+ is also the first method that can surpass the IFS on both the surface-level and pressure-level variables. Our SwinRDM* trained for forecasting at 0.25° resolution also shows better performance compared to the IFS at lead times of 3 days and 5 days. Specifically, at a lead time of 5 days, it outperforms the IFS by 21, 0.27, 0.34, and 0.53 in terms of Z500, T850, T2M, and TP, respectively. Since FourCastNet is only evaluated for the year 2018, we show the comparison results in terms of Z500 and T850 for the year 2018 in Figure 4. The results of FourCastNet are obtained from the original paper. As shown in the figure, our SwinRDM* method shows better performance compared with FourCastNet, indicating that our method shows state-of-the-art performance at 0.25° resolution. Note that our SwinRDM* shows slightly lower performance at lead times less than 3 days compared to IFS. This may attribute to the lower representational power of the encoder since we find that the capacity of the encoder is important to the short-range forecasting performance. How to better extract the historical context information is essential for the encoder, and

we leave this for future research.

Qualitative Illustration

Figure 5 shows qualitative results of the proposed SwinRDM*. Our model is tested on 8 September 2018, 06:00 UTC to forecast the near-surface wind speeds (WS) at lead times of 3 days and 5 days. The wind speeds are computed from the predicted zonal and meridional components of the wind velocity i.e., $WS = \sqrt{U_{10}^2 + V_{10}^2}$. The results of ERA5 are the ground truth. As shown in the figure, our method can forecast wind speeds for up to 5 days with high resolution and high quality. Specifically, we can see from the zoom-in area in the figure that our method can successfully forecast and track Super Typhoon Mangkhut. The ability to forecast this kind of extreme event is really beneficial for the mitigation of loss of life and property damage. Our method shows high forecasting accuracy and the ability to capture fine-scale dynamics at high resolution.

Conclusion

We propose a high-resolution data-driven medium-range weather forecasting model named SwinRDM by integrating SwinRNN+ with a diffusion-based super-resolution model. Our SwinRNN+ is improved upon SwinRNN by trading the multi-scale design for higher feature dimensions and adding a feature aggregation layer, which achieves superior performance compared to the operational NWP model on all key variables at 1.40625° resolution and lead times of up to 5 days. Our SwinRDM uses a diffusion-based super-resolution model conditioned on the forecasting results of SwinRNN+ to achieve high-resolution forecasting. The diffusion model helps generate high-resolution and high-quality forecasting results and can also perform ensemble forecasting. The experimental results show that our method achieves SOTA performance at 0.25° resolution.

References

- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*, 214–223. PMLR.
- Bihlo, A. 2021. A generative adversarial network approach to (ensemble) weather prediction. *Neural Networks*, 139: 1–16.
- Chen, B.; Lin, M.; Sheng, K.; Zhang, M.; Chen, P.; Li, K.; Cao, L.; and Ji, R. 2022. ARM: Any-Time Super-Resolution Method. *arXiv preprint arXiv:2203.10812*.
- de Witt, C. S.; Tong, C.; Zantedeschi, V.; De Martini, D.; Kalaitzis, F.; Chantry, M.; Watson-Parris, D.; and Bilinski, P. 2020. RainBench: towards global precipitation forecasting from satellite imagery. *arXiv preprint arXiv:2012.09670*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34: 8780–8794.
- Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2015. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2): 295–307.
- Garg, S.; Rasp, S.; and Thuerey, N. 2022. WeatherBench Probability: A benchmark dataset for probabilistic medium-range weather forecasting along with deep learning baseline models. *arXiv preprint arXiv:2205.00865*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- 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.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Ho, J.; Saharia, C.; Chan, W.; Fleet, D. J.; Norouzi, M.; and Salimans, T. 2022. Cascaded Diffusion Models for High Fidelity Image Generation. *J. Mach. Learn. Res.*, 23: 47–1.
- Hu, Y.; Chen, L.; Wang, Z.; and Li, H. 2022. SwinVRNN: A Data-Driven Ensemble Forecasting Model via Learned Distribution Perturbation. *arXiv preprint arXiv:2205.13158*.
- Jolliffe, I. T.; and Stephenson, D. B. 2012. *Forecast verification: a practitioner's guide in atmospheric science*. John Wiley & Sons.
- Keisler, R. 2022. Forecasting Global Weather with Graph Neural Networks. *arXiv preprint arXiv:2202.07575*.
- Kim, J.; Lee, J. K.; and Lee, K. M. 2016. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1646–1654.
- Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4681–4690.
- Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Van Gool, L.; and Timofte, R. 2021. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1833–1844.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Mu Lee, K. 2017. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 136–144.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, 8162–8171. PMLR.
- Pathak, J.; Subramanian, S.; Harrington, P.; Raja, S.; Chatopadhyay, A.; Mardani, M.; Kurth, T.; Hall, D.; Li, Z.; Azizzadenesheli, K.; et al. 2022. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Rasp, S.; Dueben, P. D.; Scher, S.; Weyn, J. A.; Mouatadid, S.; and Thuerey, N. 2020. WeatherBench: a benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11): e2020MS002203.
- Rasp, S.; and Thuerey, N. 2021. Data-driven medium-range weather prediction with a resnet pretrained on climate simulations: A new model for weatherbench. *Journal of Advances in Modeling Earth Systems*, 13(2): e2020MS002405.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Saharia, C.; Chan, W.; Chang, H.; Lee, C. A.; Ho, J.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2021a. Palette: Image-to-image diffusion models. *arXiv preprint arXiv:2111.05826*.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2021b. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*.
- Shi, X.; Gao, Z.; Lausen, L.; Wang, H.; Yeung, D.-Y.; Wong, W.-k.; and Woo, W.-c. 2017. Deep learning for precipitation nowcasting: A benchmark and a new model. *Advances in neural information processing systems*, 30.

- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, 2256–2265. PMLR.
- Sønderby, C. K.; Espenholt, L.; Heek, J.; Dehghani, M.; Oliver, A.; Salimans, T.; Agrawal, S.; Hickey, J.; and Kalchbrenner, N. 2020. Metnet: A neural weather model for precipitation forecasting. *arXiv preprint arXiv:2003.12140*.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32.
- Song, Y.; and Ermon, S. 2020. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33: 12438–12448.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2818–2826.
- Wang, X.; Xie, L.; Dong, C.; and Shan, Y. 2021. Real-rgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1905–1914.
- Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; and Change Loy, C. 2018. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, 0–0.
- Weyn, J. A.; Durran, D. R.; and Caruana, R. 2020. Improving data-driven global weather prediction using deep convolutional neural networks on a cubed sphere. *Journal of Advances in Modeling Earth Systems*, 12(9): e2020MS002109.
- Whang, J.; Delbracio, M.; Talebi, H.; Saharia, C.; Dimakis, A. G.; and Milanfar, P. 2022. Deblurring via stochastic refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16293–16303.
- Zhang, Y.; Li, K.; Li, K.; Wang, L.; Zhong, B.; and Fu, Y. 2018. Image super-resolution using very deep residual channel attention networks. In *Proceedings of the European conference on computer vision (ECCV)*, 286–301.