

SynWeather: Weather Observation Data Synthesis Across Multiple Regions and Variables via a General Diffusion Transformer

Kaiyi Xu^{1,2*}, Junchao Gong^{2,3*}, Zhiwang Zhou^{2,4}, Zhangrui Li^{2,5}, Yuandong Pu^{2,3}, Yihao Liu², Ben Fei^{2,6}, Fenghua Ling², Wenlong Zhang^{2†}, Lei Bai^{2†}

¹University of Science and Technology of China

²Shanghai Artificial Intelligence Laboratory

³Shanghai Jiao Tong University

⁴Tongji University

⁵Nanjing University

⁶The Chinese University of Hong Kong

xukaiyi@pjlab.org.cn, zhangwenlong@pjlab.org.cn, bailei@pjlab.org.cn

Abstract

With the advancement of meteorological instruments, abundant data has become available. However, due to instruments' intrinsic limitations such as environmental sensitivity and orbital constraints, raw data often suffer from temporal or spatial gaps, making it urgent to leverage data synthesis techniques to fill in missing information. Current approaches are typically focus on single-variable, single-region tasks and primarily rely on deterministic modeling. This limits unified synthesis across variables and regions, overlooks cross-variable complementarity and often leads to over-smoothed results. To address above challenges, we introduce SynWeather, the first dataset designed for **Unified Multi-region and Multi-variable Weather Observation Data Synthesis**. SynWeather covers four representative regions: the Continental United States, Europe, East Asia, and Tropical Cyclone regions, as well as provides high-resolution observations of key weather variables, including Composite Radar Reflectivity, Hourly Precipitation, Visible Light, and Microwave Brightness Temperature. In addition, we introduce SynWeatherDiff, a general and probabilistic weather synthesis model built upon the Diffusion Transformer framework to address the over-smoothed problem. Experiments on the SynWeather dataset demonstrate the effectiveness of our network compared with both task-specific and general models. Moreover, SynWeatherDiff is able to generate results that are both fine-grained and accurate in high-value regions. Through the dataset and baseline model, we aim to advance meteorological downstream tasks and promote the development of general models for weather variable synthesis.

Website — <https://dtdtxuky.github.io/SynWeather-Proj/>

Code — <https://github.com/Dtdtxuky/SynWeather>

Introduction

With the development of various meteorological observation instruments, including geostationary satellites, polar-orbiting

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

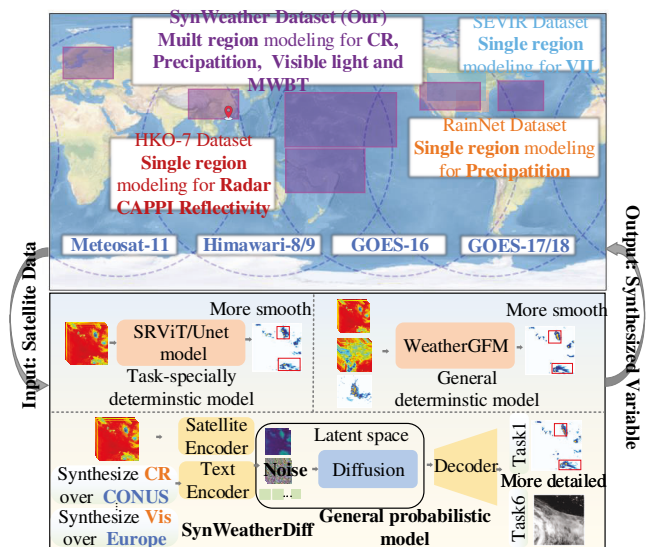


Figure 1: Overview of datasets and pipelines in weather variable synthesis. Compared to existing single-region, single-variable and deterministic modeling, SynWeather enables general multi-region, multi-variable probabilistic modeling.

satellites and radar systems, a wide range of weather data has become available (Schmit et al. 2017; Bessho et al. 2016). Such data have contributed to advancements in weather forecasting, disaster monitoring, and climate research (Bauer, Thorpe, and Brunet 2015; Joyce et al. 2009; Stephens et al. 2002) and have been used to form various datasets (Zhou et al. 2025; Zhao et al. 2025; Wang et al. 2025). However, due to the intrinsic characteristics and deployment constraints of the aforementioned instruments, original weather data face limitations in both temporal and spatial coverage. For instance, radar observations often suffer from sparse coverage in regions with complex terrain or limited economic development (Germann et al. 2022; Ovchynnykova, Svazas, and Navickas 2025). Similarly, visible satellite images are also

unavailable at night (Harder et al. 2020; Pasillas et al. 2024).

To bridge these gaps in spatial and temporal coverage, weather variable synthesis has emerged as a significant area of research (Oliver and Webster 1990; Liu et al. 2023; Hayawi, Shahriar, and Hacid 2025). Recent methods have used satellite infrared and microwave observations to reconstruct Composite Reflectivity (Stock et al. 2024; He et al. 2025b), estimate precipitation fields (Goroooh et al. 2022; Cannon et al. 2024), and generate visible light on night (Harder et al. 2020; Chirokova et al. 2023). Li, Tan, and Bai first used diffusion models conditioned on high-resolution geostationary infrared data to extend the spatial coverage of microwave observations. Apart from these task-specific models, WeatherGFM (Zhao et al. 2024) as a general model has also been applied to a variety of weather understanding tasks, including radar and visible light synthesis.

However, weather variable synthesis is still limited in the following aspects: (1) *Lack of Global-Scale and Multi-Variable Observations Dataset*: Most existing datasets for weather variable synthesis are limited to single-region or/and single-variable. For example, HKO-7 (Shi et al. 2017) and SEVIR (Veillette, Samsi, and Mattioli 2020) focus solely on radar-related variables over Hong Kong and the Continental United States (CONUS), respectively. While some global datasets like RainBench (Schroeder de Witt et al. 2020) include multiple weather variables, these are mostly derived from reanalysis rather than observation. Moreover, due to the limited geographical coverage, existing datasets rely on a single satellite source. For instance, SEVIR (Veillette, Samsi, and Mattioli 2020) only uses three GOES satellites, and Digital Typhoon (Kitamoto et al. 2023) only uses one Himawari channel. (2) *Customized Deterministic Modeling*: Current weather variable synthesis methods use a special network for a particular weather variable synthesis task. Besides, these methods are often deterministic models trained with pixel-wise squared loss. As a result, the outputs are often overly smooth and unable to capture high-intensity areas during severe weather events. Moreover, different regions and observation variables can be viewed as multi-modalities reflecting the atmospheric state, but existing datasets and modeling overlook their cross-modal complementarities. In summary, weather variable synthesis faces an important challenge: ***How to perform general-purpose generative modeling of multiple weather observation variables across multiple regions?***

To address the above challenge, we first introduce SynWeather, a dataset that supports unified multi-region and multi-variable weather variable synthesis. The dataset spans four key regions: the CONUS, Europe, East Asia, and Tropical Cyclone regions (TC regions), and integrates 10 channels infrared observations from corresponding global geostationary satellites (GOES-16/17/18, Meteosat-11, and Himawari-8/9) to synthesize a variety of weather variables including composite reflectivity (CR), precipitation, visible light, and microwave brightness temperature (MWBT). Building upon SynWeather, we design six standard tasks and one out-of-distribution (OOD) task to comprehensively evaluate weather variable synthesis models. We further propose SynWeatherDiff, a general-purpose weather synthesis model built on a diffusion transformer. It leverages text prompts to guide

the generation of diverse weather variables across different regions and satellite sources. By adopting a probabilistic modeling framework, SynWeatherDiff is capable of generating fine-grained spatial structures and recovering high-intensity regions, which are often missed by deterministic models. Moreover, the synthesized results can serve as valuable inputs for both nowcasting (Gong et al. 2024b; He et al. 2025a; Xu et al. 2025) and medium-range forecasting systems (Chen et al. 2025a), as well as data assimilation (Sun et al. 2025). Our main contributions can be summarized as:

- We construct SynWeather, the first standardized dataset supporting unified multi-region and multi-variable weather observation data synthesis tasks.
- We propose SynWeatherDiff, the first generative and probabilistic framework that unifies many variables and regions under a single text-prompt-driven interface.
- Extensive experiments demonstrate the effectiveness of SynWeatherDiff over both task-specific and general models, particularly in generating fine-grained details.

SynWeather Dataset

Dataset Collection

The input data are collected from all infrared channels of six geostationary satellites (i.e., GOES-16/17/18, Himawari-8/9, and Meteosat-11) and could seamlessly cover the four target regions as shown in Fig. 2. These satellites capture full-disk images every 10–15 minutes at spatial resolutions from 0.5 to 4 km. The target weather variables are sourced from authoritative datasets in each respective region: composite reflectivity (CR) from the GREMLIN CONUS3 dataset (Hilburn 2023) and 1-hour quantitative precipitation estimates (QPE) from Multi-Radar Multi-Sensor (MRMS) in the CONUS; radar gauge fused precipitation from EURADCLIM (Overeem et al. 2023) in Europe; satellite-based precipitation from Global Precipitation Measurement (GPM) mission in East Asia; and microwave brightness temperatures (MWBT) from AMSR-2 and GMI provided by the Tropical Cyclone Precipitation, Infrared, Microwave, and Environmental Dataset (TC PRIMED) (Razin et al. 2023) for TC regions. Details of the data coverage and composition are summarized in Table 1.

Dataset Processing

Since the original inputs and targets vary temporally and spatially, all data are first standardized to a unified resolution of 1-hour and 4-kilometer and aligned by the closest timestamps. To ensure consistent input size for training, all regional data are cropped to 256×256 patches using a sliding window with a 128-pixel stride. However, not all patches contain meaningful signals; therefore, we set two thresholds γ_1 , γ_2 and keep patches that contain connected components formed by pixels that exceed γ_1 in value with sizes larger than γ_2 . Finally, to align variable ranges, we log-transform precipitation for its long-tailed distribution, followed by min-max normalization. For other variables, we directly apply min-max normalization.

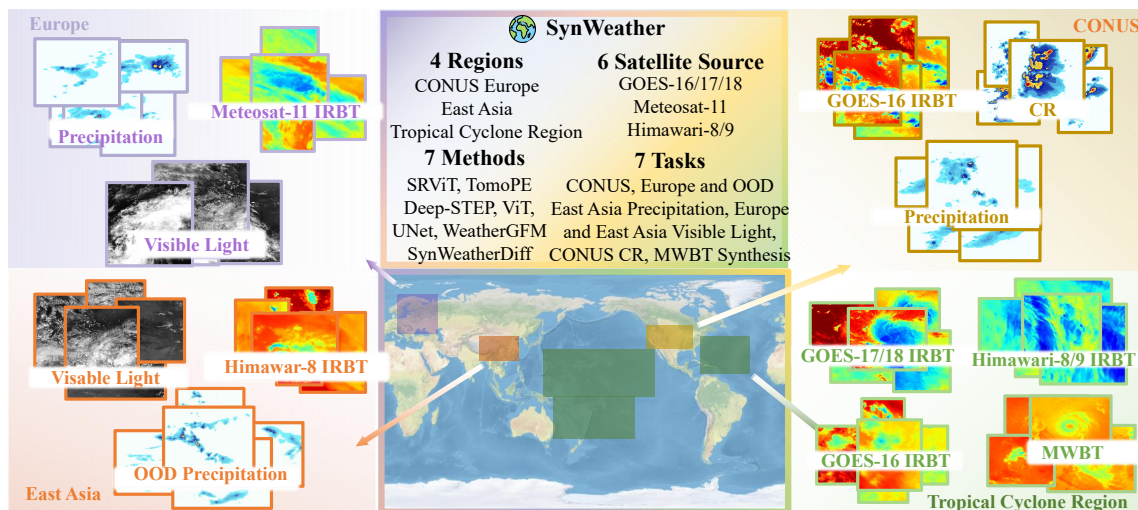


Figure 2: Overview of SynWeather. SynWeather is a comprehensive dataset that covers four distinct regions and four key weather observation variables, integrating data from six satellite sources as a condition to support seven synthesis tasks. Extensive evaluations are conducted on seven models, comprising both task-specific and general synthesis models.

Region	Data	Input			Target		Sample numbers	Year
	Satellite	Band	Spatial Res. (km)	Variable	Source	Spatial Res. (km)		
CONUS	GOES-16	C07-16	2	CR Precipitation	GREMLIN CONUS3 Dataset MRMS	3 1	142k 20k	2020–2022
Europe	Meteosat-11	IR_016-134 WV_062-073	3	Visible light Precipitation	Meteosat-11 EURADCLIM	3 1	372k 25k	2019–2021
East Asia	Himawari-8	C07-16	2	Visible light Precipitation	Himawari-8 GPM	2 10	503k 15k	2019–2021 2021.7
TC Region	GOES-16/17/18 Himawari-8/9	C07-16 C07-16	2	MWBT	AMS2-2 GMI	7×12, 3×5 8.6×14, 4.4×7.2	9k	2015–2023

Table 1: Detailed information of SynWeather. Spatial resolution is denoted as “Spatial Res”

Dataset Statistics

Apart from the MWBT collected event-wise during tropical cyclones in the TC PRIMED dataset (Razin et al. 2023), all other SynWeather variables are sampled continuously at fixed intervals, resulting in a relatively balanced number of raw samples across variables. However, because precipitation is inherently sparse in space and time, many samples contain little or no rainfall. After removing these non-precipitating cases, the number of valid precipitation patches becomes substantially smaller than for other variables. Table 1 summarizes the sample counts, showing that visible imagery accounts for the largest share, followed by CR, while precipitation and MWBT remain comparatively limited.

Proposed Baseline: SynWeatherDiff

One major challenge in weather data synthesis is the diversity of satellite sources and observed variables across different regions and types. Under such conditions, a specialized modeling strategy requires training separate models for each region-variable pair, resulting in considerable computational

and operational overhead. However, many weather variables are often physically correlated. For instance, precipitation can be inferred from CR using a Z-R relationship (Wu et al. 2018; Peng et al. 2022), highlighting the potential for joint modeling across variables. Inspired by these observations, we propose SynWeatherDiff, a text-prompt-based general model that enables unified modeling and conditional generation across multiple regions and variables similar to unified models in natural images (Chen et al. 2025b; Pu et al. 2025).

Problem Formulation

Based on the SynWeather, we define seven weather observation variable synthesis tasks, including six standard tasks: (i) CR synthesis over CONUS, (ii, iii) precipitation synthesis over CONUS and Europe, (iv, v) visible light synthesis over East Asia and Europe, and (vi) MWBT synthesis over tropical cyclone regions, as well as a out-of-distribution task: precipitation synthesis over East Asia. For each task, we compare specialized models (trained on region- and variable-specific data) with a general model (trained on all regions

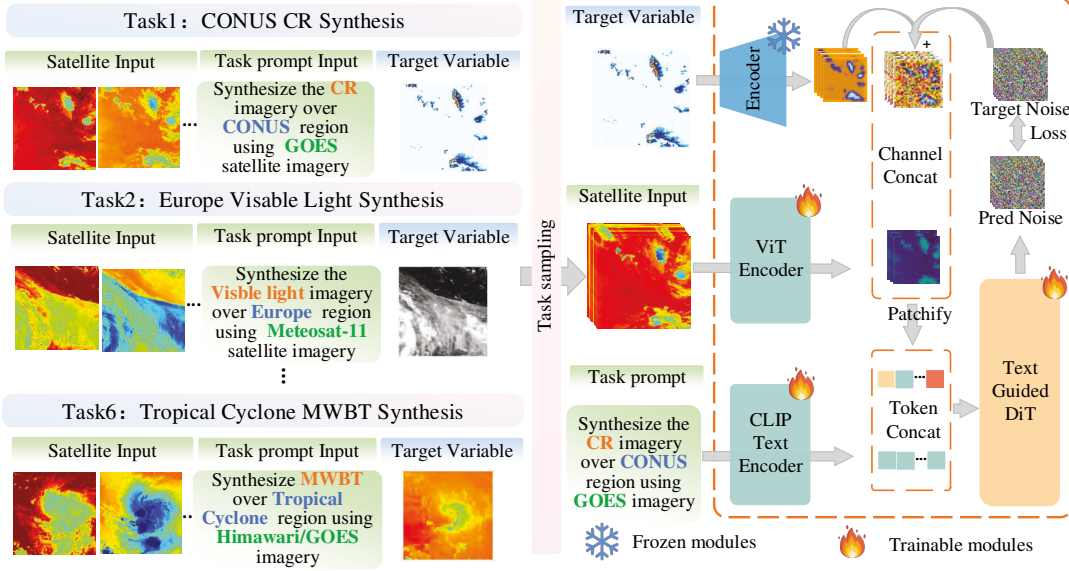


Figure 3: An overview of our SynWeatherDiff. The target variables are projected into a unified latent space using a general autoencoder. The satellite inputs are processed through a ViT-based encoder to extract features. A task-specific text prompt is encoded using a fine-tuned CLIP text encoder. The text tokens serve as conditional information to guide Text-Guided DiT for different weather synthesis tasks.

and variables). For single tasks, the specialized model $f_{r,b}$ maps satellite observations $X_r \in \mathbb{R}^{C \times H \times W}$ to weather observation variables $Y_{r,b} \in \mathbb{R}^{1 \times H \times W}$:

$$Y_{r,b} = f_{r,b}(X_r), \quad (1)$$

where r denotes the target region, including CONUS, Europe (Eur), East Asia (EA), and TC regions. And b represents the type of weather observation variable, containing CR, Precipitation (Prec), Visible Light (Vis) and MWBT; for OOD tasks, the specialized model $f_{r,b}$ generalizes as:

$$Y_{EA, Prec}^{OOD} = f_{r_1, Prec}(X_{EA}), \quad (2)$$

where $r_1 \in \{\text{CONUS, Eur}\}$.

In addition, we introduce a general model g to perform all types of weather variable synthesis tasks, which takes both the satellite observation X_r and text prompts $P_{r,b}$: “Synthesize the b variable over the r region using corresponding satellite imagery.” as inputs to guide the generation of $Y_{r,b}$:

$$Y_{r,b} = g(X_r, P_{r,b}), \quad (3)$$

Similarly, the OOD tasks are represented as:

$$Y_{EA, Prec}^{OOD} = g(X_{EA}, P_{EA, Prec}). \quad (4)$$

Network Architecture

Fig. 3 outlines the architecture of our general weather variable synthesis model, SynWeatherDiff. Firstly, it uses a general autoencoder that encodes all weather variables into a shared latent space. And then, a diffusion transformer-based denoising network is trained to perform conditional generation in the latent space, guided by both the encoded satellite inputs and a task-specific text prompt.

General AutoEncoder. Unlike natural images, meteorological images often contain redundancy (Luo, Xu, and Ji 2015). It is rare to observe precipitation or typhoons across areas spanning hundreds or thousands of kilometers at the same time. Therefore, even after careful filtering, meteorological images still include background areas that offer limited meteorological value. In addition, some weather variables also exhibit physical similarity and thus have the potential to complement each other. The above characteristics motivate us to compress various weather variables into a shared latent space. Following (Rombach et al. 2022), we train a general autoencoder using a combination of pixel-wise reconstruction loss, KL divergence loss, and adversarial loss. Specifically, the encoder encodes each weather observation variable $Y_{r,b} \in \mathbb{R}^{1 \times H \times W}$ into latent representation $z_{r,b} \in \mathbb{R}^{C_z \times H_z \times W_z}$ and the decoder reconstructs it as $\hat{Y}_{r,b}$.

Text-Guided Diffusion Transformer. During this stage, the diffusion transformer learns to recover clean latent representations of different weather variables from their noisy versions, under the guidance of satellite inputs and task-specific prompts. For each region-variable pair (r, b) , we define a prompt $P_{r,b}$ following the format: “Synthesize the b variable over the r region using corresponding satellite imagery.” This prompt is embedded using a pretrained CLIP text encoder (Radford et al. 2021), with only the final transformer block fine-tuned to adapt to the weather synthesis domain. Similarly, the satellite observation input X_r is encoded by a ViT-based encoder that is jointly trained with the diffusion transformer to extract features. In contrast to SD3 (Esser et al. 2024), SynWeatherDiff adopts an early fusion strategy:

Task name	CR Synthesis				Precipitation Synthesis						
	CONUS				CONUS			Europe			
	RMSE↓	CSI/25↑	CSI/35↑	CSI/40↑	RMSE↓	CSI/2↑	CSI/5↑	CSI/15↑	RMSE↓	CSI/2↑	CSI/5↑
SRViT [#]	3.561	0.277	0.120	0.069	-	-	-	-	-	-	-
Deep-STEP [#]	-	-	-	-	0.916	0.262	0.111	0.007	<u>0.415</u>	<u>0.083</u>	0.016
TomoPE [#]	-	-	-	-	0.986	0.247	0.149	0.036	0.413	0.060	0.009
UNet [#]	3.395	0.299	0.069	0.023	0.976	0.231	0.166	0.059	0.641	0.035	0.016
ViT [#]	3.487	0.309	0.141	<u>0.089</u>	0.981	0.250	0.157	0.038	0.497	<u>0.083</u>	<u>0.044</u>
WeatherGFM [†]	<u>3.124</u>	<u>0.366</u>	0.166	0.086	1.049	<u>0.288</u>	<u>0.198</u>	<u>0.090</u>	0.714	0.018	0.013
SynWeatherDiff [†]	2.820	0.382	<u>0.158</u>	0.101	<u>0.976</u>	0.312	0.223	0.113	0.569	0.084	0.079

Task name	Visible Light Synthesis						MWBT Synthesis				
	East Asia			Europe			Tropical Cyclone Region				
	SSIM↑	PSNR↑	CSI/50↑	SSIM↑	PSNR↑	CSI/50↑	RMSE↓	SSIM↑	PSNR↑	LPIPS↓	CSI/300↑
ViT [#]	<u>0.870</u>	<u>20.87</u>	0.672	0.860	<u>24.03</u>	0.496	<u>4.768</u>	0.783	21.56	0.324	<u>0.792</u>
UNet [#]	0.917	21.67	0.711	0.878	24.82	0.556	5.803	0.816	20.6	<u>0.329</u>	0.741
WeatherGFM [†]	0.822	18.43	0.465	0.836	22.26	0.396	4.979	<u>0.828</u>	<u>21.86</u>	0.325	0.777
SynWeatherDiff [†]	0.868	19.79	<u>0.690</u>	<u>0.864</u>	23.65	<u>0.508</u>	4.456	0.837	22.33	0.254	0.795

Table 2: Quantitative results on standard weather synthesis tasks. [#] specialized model. [†]: general model trained with all six weather variable synthesis. Best results are **bold**, second-best are underlined.

the noisy latent $z_{r,b}$ is first concatenated with satellite encoder features and then patchified. These patches are further concatenated with the embedded prompt tokens and passed through the self-attention layers of the diffusion transformer to perform conditional denoising. The training objective in this stage is to predict noise loss as follows:

$$\mathcal{L} = \mathbb{E}_{z_{r,b}, \epsilon, t} \left[\left\| \epsilon_{\theta}(z_{r,b}^t, t, X_{r,b}, P_{r,b}) - \epsilon \right\|_2^2 \right], \quad (5)$$

$$z_{r,b}^t = \sqrt{\alpha_t} \cdot z_{r,b} + \sqrt{1 - \alpha_t} \cdot \epsilon, \quad (6)$$

where $\epsilon_{\theta}(\cdot)$ is the noise predicted by the DiT. $\epsilon \sim \mathcal{N}(0, I)$, and α_t is the cumulative noise schedule product.

Experiments

Evaluation Protocol

We evaluate performance using three metric categories:

Event detection accuracy: We first use the Critical Success Index (CSI), a common metric in weather variable synthesis (Schaefer 1990; Gong et al. 2024a). It is defined as:

$$\text{CSI} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}}, \quad (7)$$

where TP, FP, and FN denote true positives, false positives, and false negatives. Due to varying variable ranges, we apply task-specific CSI thresholds: 25, 35, 40 for CR synthesis and 2, 5, 15 for precipitation. For visible light and MWBT, we focus more on perceptual quality. Therefore, we adopt coarse thresholds of 50 and 300.

Regression quality: In addition to thresholds CSI, we report the Root Mean Squared Error (RMSE) to measure pixel-wise differences between predictions and ground truth.

Perceptual similarity: Finally, to assess perceptual similarity, we employ the Structural Similarity Index (SSIM), Peak Signal-to-Noise Ratio (PSNR), and Learned Perceptual Image Patch Similarity (LPIPS).

Benchmark Setup

We benchmark against three categories of models: (i) general models: WeatherGFM (Zhao et al. 2024) that focus on multi-variable synthesis but lack regional generalization; (ii) task-specific models: SRViT (Stock et al. 2024) for vertically integrated liquid (VIL) and Deep-STEP (Gorooh et al. 2022), TomoPE (Cannon et al. 2024) for precipitation; (iii) popular deep learning models: UNet and ViT adapted for satellite-to-variable mapping. All models are trained using a unified protocol on SynWeather.

Implementation Details

We train SynWeatherDiff for 600K steps with a batch size of 16 on 4×80GB NVIDIA A100 GPUs. The model is optimized using AdamW with a cosine learning rate scheduler, decaying the learning rate from 5e-4 to 1e-5.

Experiment Results

SynWeatherDiff achieves strong universal capabilities. As shown in Table 2 and Fig 4, SynWeatherDiff, guided by task-specific text prompts, successfully synthesizes multiple weather variables across different regions using heterogeneous satellite sources. It outperforms the existing generalist model, WeatherGFM, across most tasks. Furthermore, SynWeatherDiff is capable of distinguishing and generating different variables from the same satellite input, such as synthesizing both CR and precipitation over CONUS, visible

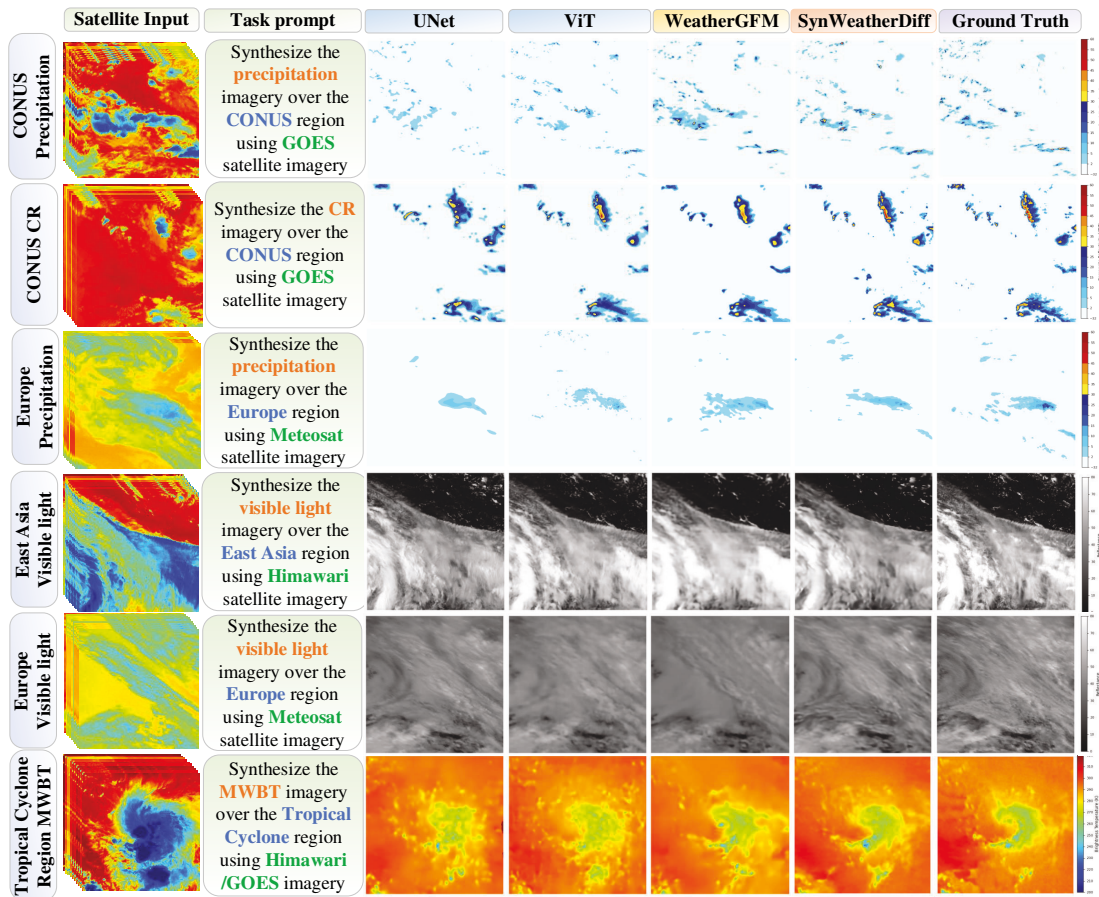


Figure 4: Visual results of the weather synthesis standard tasks by our SynWeatherDiff and other models.

light and precipitation over Europe. This highlights our SynWeatherDiff’s strong universality and flexibility, attributed to the task-specific text prompts.

SynWeatherDiff outperforms specialized models across most synthesis tasks. As shown in Table 2, SynWeatherDiff achieves state-of-the-art CSI performance with comparable RMSE performance for precipitation synthesis, which is more challenging and critical in meteorological applications. In the visible light synthesis task, UNet achieves better results than SynWeatherDiff as it operates directly in pixel space. Visible light data contains abundant high-frequency details, which are difficult to reconstruct through the autoencoder used by SynWeatherDiff. Nevertheless, with improvements to the autoencoder, SynWeatherDiff holds promising potential to surpass specialized models across all tasks.

SynWeatherDiff demonstrates a strong ability to generate results with both fine-grained and high-value details. In Fig 4, SynWeatherDiff clearly outperforms UNet, ViT and WeatherGFM in generating fine-grained details. The latter three models produce fewer individual cells, which can be attributed to two factors: (1) weak signals being overlooked; (2) multiple small weather cells tend to merge into larger

areas, which smooths out boundaries and removes important local details. For high-value regions, ViT and WeatherGFM can barely recover some intensity centers in CR but fail to reconstruct the surrounding convective structure. UNet performs worse, often missing the intensity centers entirely in both CR and precipitation tasks. In contrast, SynWeatherDiff successfully restores the number and distribution of individual scattered cells. It also accurately captures the location and shape of intensity centers, closely matching the ground truth. These observations are consistent with the quantitative results in Table 4. SynWeatherDiff achieves higher CSI scores and demonstrates superior performance in generating fine-grained patterns and high-intensity events.

Ablation Studies and Exploration

Exploration of different task sampling ratios. In this section, CONUS CR, CONUS precipitation, Europe visible light and MWBT synthesis are set as the main task with a 0.5 sampling ratio in turn, with the other five tasks each at 0.1. Table 3 reveals three findings. (1) CR synthesis helps CONUS precipitation (as the Z-R relationship makes CR a strong prior, and its large sample size and convective pattern capturing ability benefit precipitation modeling) and MWBT synthesis

Sampling Setting	CONUS Precipitation			CONUS CR			Europe Visible Light			MWBT		
	RMSE	CSI/2	CSI/15	RMSE	CSI/25	CSI/35	SSIM	PSNR	CSI/70	SSIM	LPIPS	PSNR
Uniform ($\frac{1}{6}$)	0.976	<u>0.312</u>	0.113	<u>2.82</u>	<u>0.382</u>	<u>0.158</u>	0.864	23.65	0.172	0.837	0.254	22.33
CONUS CR ($\frac{1}{2}$)	0.961	0.320	0.137	2.69	0.403	0.187	0.857	23.33	0.168	0.843	0.252	<u>22.50</u>
CONUS Prec($\frac{1}{2}$)	<u>0.974</u>	0.292	0.087	2.88	0.343	0.110	0.855	23.34	0.166	0.841	0.258	22.15
Europe Vis ($\frac{1}{2}$)	1.014	0.298	<u>0.115</u>	2.77	0.377	0.145	<u>0.877</u>	24.16	0.185	<u>0.842</u>	<u>0.254</u>	22.51
MWBT ($\frac{1}{2}$)	1.001	0.295	0.112	2.77	0.374	0.141	0.879	<u>24.13</u>	<u>0.184</u>	<u>0.842</u>	<u>0.254</u>	22.38

Table 3: Ablation study on task sampling ratios. In each setting, the primary task is assigned a sampling ratio of 0.5, while the remaining five tasks are set to 0.1. We denote each experiment by the name of the primary task along with its sampling ratio. Best results are **bold**, second-best are underlined.

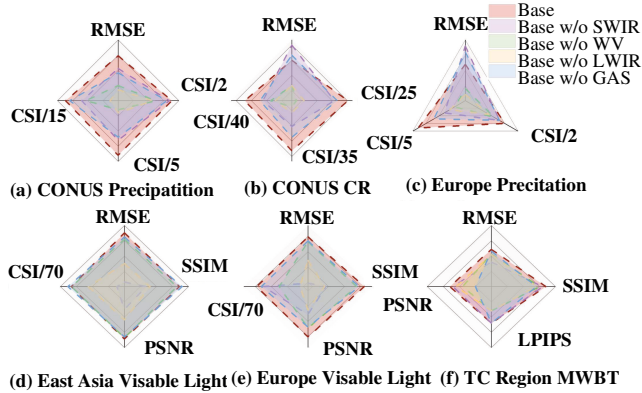


Figure 5: The effect of input channel is analyzed across six tasks, focusing on four groups: (1) SWIR (shortwave infrared), (2) WV (water vapor channels), (3) LWIR (longwave infrared channels), and (4) GAS (gas absorption channels).

(since CR from ground-based S-band radars shares a similar physical domain with satellite-based sensors and offers higher spatial resolution). Thus, a higher CR task sampling ratio than uniform sampling boosts both. (2) Visible light and MWBT tasks benefit each other. When Europe’s visible light or tropical cyclone MWBT synthesis is the main task, they surpass uniform sampling. This is because the visible band and MWBT have spatial correspondence in strong convective systems, aiding cross-task knowledge transfer. (3) CR task performance gains mainly come from the increased number of training samples. It does not benefit from interactions with other tasks. Notably, the CR task only improves when CR synthesis is set as the main task. Also, there are conflicting tasks like CR and visible light synthesis. Therefore, choosing a proper task ratio to balance mutually beneficial and conflicting tasks is key for overall performance gains.

Exploration of different input channels. We choose all 10 infrared channels instead of the usual 3-channel setup, and divided them into four groups based on spectral range and atmospheric sensitivity (SWIR, WV, LWIR, GAS). Ablation experiments involved removing each group to assess its contribution to synthesis performance. Fig. 5 shows that removing any group worsens all synthesis tasks. Detailed analysis indicates that water vapor and longwave infrared

channels are vital for precipitation and radar reflectivity synthesis, while shortwave infrared channels are more important for visible light synthesis. For example, in the East Asia visible light task, removing the SWIR channel alone causes a bigger performance drop than removing all LWIR channels. Gas absorption channels have less impact on visible light synthesis. For MWBT synthesis, removing any single group reduces performance, but the overall effect is small.

Metric	RMSE	SSIM	PSNR	CSI/2	CSI/5
UNet [#]	1.383	0.671	30.02	0.161	0.076
ViT [#]	<u>1.233</u>	<u>0.749</u>	29.49	<u>0.232</u>	<u>0.103</u>
SynWeatherDiff [†]	1.150	0.771	<u>29.80</u>	0.235	0.108

Table 4: Quantitative results on OOD weather synthesis tasks. [#] specialized model. [†]: general model. Best results are **bold**, and second-best are underlined.

Exploration on out-of-distribution tasks. To evaluate the generalization ability of SynWeatherDiff, we conduct OOD experiments on East Asia precipitation synthesis. Table 4 shows that SynWeatherDiff surpasses the specialized model trained exclusively on CONUS precipitation, demonstrating superior generalization across both region and variable. This not only highlights the advantages of using text prompts for flexible task control but also shows that different variables can complement each other in a general model, compared to region-specific and single-variable models. Nevertheless, inference under the current prompt-based framework is also constrained to the specific regions and variables that were individually present in the training set.

Conclusion

We construct the first multi-region and multi-variable weather observation data synthesis dataset (i.e., SynWeather) and introduce a general diffusion-based model as a baseline (i.e., SynWeatherDiff). By introducing text prompts to guide the synthesis process, SynWeatherDiff flexibly generates diverse weather variables across multiple regions within a unified framework. We conduct a comprehensive analysis of task sampling ratios and the input channel. Through the proposed dataset and baseline model, we aim to facilitate research in weather synthesis and inspire the development of future generalist models for weather-related downstream tasks.

Acknowledgements

This work is supported by Shanghai Artificial Intelligence Laboratory. This work was done during her internship at Shanghai Artificial Intelligence Laboratory.

References

- Bauer, P.; Thorpe, A.; and Brunet, G. 2015. The quiet revolution of numerical weather prediction. *Nature*, 525(7567): 47–55.
- Bessho, K.; Date, K.; Hayashi, M.; Ikeda, A.; Imai, T.; Inoue, H.; Kumagai, Y.; Miyakawa, T.; Murata, H.; Ohno, T.; et al. 2016. An introduction to Himawari-8/9—Japan’s new-generation geostationary meteorological satellites. *Journal of the Meteorological Society of Japan. Ser. II*, 94(2): 151–183.
- Cannon, F.; Pfreundschuh, S.; Taylor, B.; Munchak, S. J.; Nelson, S.; L’heureux, J.; Owens, C.; Conibear, L.; Flam-pouris, S.; and Chawla, A. 2024. Deep Learning for Multi-Satellite Precipitation Retrievals: Impact of Tomorrow. io’s Microwave Sounders. *Authorea Preprints*.
- Chen, K.; Han, T.; Ling, F.; Gong, J.; Bai, L.; Wang, X.; Luo, J.-J.; Fei, B.; Zhang, W.; Chen, X.; et al. 2025a. The operational medium-range deterministic weather forecasting can be extended beyond a 10-day lead time. *Communications Earth & Environment*, 6(1): 518.
- Chen, X.; Zhu, K.; Pu, Y.; Cao, S.; Li, X.; Zhang, W.; Liu, Y.; Qiao, Y.; Zhou, J.; and Dong, C. 2025b. Exploring Scalable Unified Modeling for General Low-Level Vision. *arXiv preprint arXiv:2507.14801*.
- Chirokova, G.; Knaff, J. A.; Brennan, M. J.; DeMaria, R. T.; Bozeman, M.; Stevenson, S. N.; Beven, J. L.; Blake, E. S.; Brammer, A.; Darlow, J. W.; et al. 2023. ProxyVis—A proxy for nighttime visible imagery applicable to geostationary satellite observations. *Weather and Forecasting*, 38(12): 2527–2550.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Germann, U.; Boscacci, M.; Clementi, L.; Gabella, M.; Hering, A.; Sartori, M.; Sideris, I. V.; and Calpini, B. 2022. Weather radar in complex orography. *Remote Sensing*, 14(3): 503.
- Gong, J.; Bai, L.; Ye, P.; Xu, W.; Liu, N.; Dai, J.; Yang, X.; and Ouyang, W. 2024a. Cascast: Skillful high-resolution precipitation nowcasting via cascaded modelling. *arXiv preprint arXiv:2402.04290*.
- Gong, J.; Tu, S.; Yang, W.; Fei, B.; Chen, K.; Zhang, W.; Yang, X.; Ouyang, W.; and Bai, L. 2024b. Postcast: Generalizable postprocessing for precipitation nowcasting via unsupervised blurriness modeling. *arXiv preprint arXiv:2410.05805*.
- Gorooh, V. A.; Asanjan, A. A.; Nguyen, P.; Hsu, K.; and Sorooshian, S. 2022. Deep neural network high SpatioTemporal resolution precipitation estimation (Deep-STEP) using passive microwave and infrared data. *Journal of Hydrometeorology*, 23(4): 597–617.
- Harder, P.; Jones, W.; Lguensat, R.; Bouabid, S.; Fulton, J.; Quesada-Chacón, D.; Marcolongo, A.; Stefanović, S.; Rao, Y.; Manshausen, P.; et al. 2020. NightVision: generating nighttime satellite imagery from infra-Red observations. *arXiv preprint arXiv:2011.07017*.
- Hayawi, K.; Shahriar, S.; and Hacid, H. 2025. Climate data imputation and quality improvement using satellite data. *Journal of Data Science and Intelligent Systems*, 3(2): 87–97.
- He, X.; You, Z.; Gong, J.; Liu, C.; Yue, X.; Zhuang, P.; Zhang, W.; and Bai, L. 2025a. RadarQA: Multi-modal Quality Analysis of Weather Radar Forecasts. *arXiv preprint arXiv:2508.12291*.
- He, X.; Zhou, Z.; Zhang, W.; Zhao, X.; Chen, H.; Chen, S.; and Bai, L. 2025b. DiffSR: Learning Radar Reflectivity Synthesis via Diffusion Model from Satellite Observations. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Hilburn, K. 2023. GREMLIN CONUS3 Dataset for 2022.
- Joyce, K. E.; Belliss, S. E.; Samsonov, S. V.; McNeill, S. J.; and Glassey, P. J. 2009. A review of the status of satellite remote sensing and image processing techniques for mapping natural hazards and disasters. *Progress in physical geography*, 33(2): 183–207.
- Kitamoto, A.; Hwang, J.; Vuillod, B.; Gautier, L.; Tian, Y.; and Clanuwat, T. 2023. Digital typhoon: Long-term satellite image dataset for the spatio-temporal modeling of tropical cyclones. *Advances in Neural Information Processing Systems*, 36: 40623–40636.
- Li, Z.; Tan, Z.-M.; and Bai, L. 2025. Generative Deep Learning Reconstructs Tropical Cyclone Microwave Data from Geostationary Infrared Radiometers. *Authorea Preprints*.
- Liu, M.; Huang, H.; Feng, H.; Sun, L.; Du, B.; and Fu, Y. 2023. Pristi: A conditional diffusion framework for spatiotemporal imputation. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 1927–1939. IEEE.
- Luo, Y.; Xu, Y.; and Ji, H. 2015. Removing rain from a single image via discriminative sparse coding. In *Proceedings of the IEEE international conference on computer vision*, 3397–3405.
- Oliver, M. A.; and Webster, R. 1990. Kriging: a method of interpolation for geographical information systems. *International Journal of Geographical Information System*, 4(3): 313–332.
- Ovchynnykova, O.; Svazas, M.; and Navickas, V. 2025. Assessing economic profiles of coastal regions in the blue economy: a radar chart approach. *Challenges in sustainability*, 13(2): 177–192.
- Overeem, A.; van den Besselaar, E.; van der Schrier, G.; Meirink, J. F.; van der Plas, E.; and Leijnse, H. 2023. EURADCLIM: The European climatological high-resolution gauge-adjusted radar precipitation dataset. *Earth System Science Data*, 15(3): 1441–1464.
- Pasillas, C. M.; Kummerow, C.; Bell, M.; and Miller, S. D. 2024. Turning Night into Day: The Creation and Validation of Synthetic Nighttime Visible Imagery Using the Visible Infrared Imaging Radiometer Suite (VIIRS) Day–Night Band

- (DNB) and Machine Learning. *Artificial Intelligence for the Earth Systems*, 3(3): e230002.
- Peng, W.; Bao, S.; Yang, K.; Wei, J.; Zhu, X.; Qiao, Z.; Wang, Y.; and Li, Q. 2022. Radar Quantitative Precipitation Estimation Algorithm Based on Precipitation Classification and Dynamical ZR Relationship. *Water*, 14(21): 3436.
- Pu, Y.; Zhuo, L.; Zhu, K.; Xie, L.; Zhang, W.; Chen, X.; Gao, P.; Qiao, Y.; Dong, C.; and Liu, Y. 2025. Lumina-omnilv: A unified multimodal framework for general low-level vision. *arXiv preprint arXiv:2504.04903*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Razin, M. N.; Slocum, C. J.; Knaff, J. A.; Brown, P. J.; and Bell, M. M. 2023. Tropical cyclone precipitation, infrared, microwave, and environmental dataset (TC PRIMED). *Bulletin of the American Meteorological Society*, 104(11): E1980–E1998.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Schaefer, J. T. 1990. The critical success index as an indicator of warning skill. *Weather and forecasting*, 5(4): 570–575.
- Schmit, T. J.; Griffith, P.; Gunshor, M. M.; Daniels, J. M.; Goodman, S. J.; and Lebar, W. J. 2017. A closer look at the ABI on the GOES-R series. *Bulletin of the American Meteorological Society*, 98(4): 681–698.
- Schroeder de Witt, C.; 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 e-prints*, arXiv–2012.
- 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.
- Stephens, G. L.; Vane, D. G.; Boain, R. J.; Mace, G. G.; Sassen, K.; Wang, Z.; Illingworth, A. J.; O’connor, E. J.; Rossow, W. B.; Durden, S. L.; et al. 2002. The CloudSat mission and the A-Train: A new dimension of space-based observations of clouds and precipitation. *Bulletin of the American Meteorological Society*, 83(12): 1771–1790.
- Stock, J.; Hilburn, K.; Ebert-Uphoff, I.; and Anderson, C. 2024. Srvit: Vision transformers for estimating radar reflectivity from satellite observations at scale. *arXiv preprint arXiv:2406.16955*.
- Sun, J.-A.; Fan, H.; Gong, J.; Fei, B.; Chen, K.; Ling, F.; Zhang, W.; Xu, W.; Yan, L.; Gentine, P.; et al. 2025. Align-DA: Align Score-based Atmospheric Data Assimilation with Multiple Preferences. *arXiv preprint arXiv:2505.22008*.
- Veillette, M.; Samsi, S.; and Mattioli, C. 2020. Sevir: A storm event imagery dataset for deep learning applications in radar and satellite meteorology. *Advances in Neural Information Processing Systems*, 33: 22009–22019.
- Wang, F.; Chen, M.; He, X.; Zhang, Y.; Liu, F.; Guo, Z.; Hu, Z.; Wang, J.; Xu, J.; Li, Z.; et al. 2025. OmniEarth-Bench: Towards Holistic Evaluation of Earth’s Six Spheres and Cross-Spheres Interactions with Multimodal Observational Earth Data. *arXiv preprint arXiv:2505.23522*.
- Wu, W.; Zou, H.; Shan, J.; and Wu, S. 2018. A Dynamical Z-R Relationship for Precipitation Estimation Based on Radar Echo-Top Height Classification. *Advances in Meteorology*, 2018(1): 8202031.
- Xu, K.; Gong, J.; Zhang, W.; Fei, B.; Bai, L.; and Ouyang, W. 2025. SynCast: Synergizing Contradictions in Precipitation Nowcasting via Diffusion Sequential Preference Optimization. *arXiv preprint arXiv:2510.21847*.
- Zhao, X.; Xu, W.; Liu, B.; Zhou, Y.; Ling, F.; Fei, B.; Yue, X.; Bai, L.; Zhang, W.; and Wu, X.-M. 2025. MSEarth: A Benchmark for Multimodal Scientific Comprehension of Earth Science. *arXiv preprint arXiv:2505.20740*.
- Zhao, X.; Zhou, Z.; Zhang, W.; Liu, Y.; Chen, X.; Gong, J.; Chen, H.; Fei, B.; Chen, S.; Ouyang, W.; et al. 2024. Weathergfm: Learning a weather generalist foundation model via in-context learning. *arXiv preprint arXiv:2411.05420*.
- Zhou, Y.; Wang, Y.; He, X.; Xiao, R.; Li, Z.; Feng, Q.; Guo, Z.; Yang, Y.; Wu, H.; Huang, W.; et al. 2025. Scientists’ First Exam: Probing Cognitive Abilities of MLLM via Perception, Understanding, and Reasoning. *arXiv preprint arXiv:2506.10521*.