BEYOND BLURRINESS AND ARTIFACTS: A SYNERGISTIC DETERMINISTIC-PROBABILISTIC APPROACH FOR RADAR RECONSTRUCTION

Jun Chen[†], *Lin Chen*[†], *Yan Fang*[†], *Minghui Qiu*[†], *Shuxin Zhong*[†], *Yu Zhang*[‡], *Binghong Chen*[‡], *Kaishun Wu*[†*]

[†] IoT & DSA Thrust, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China
[‡] Guangzhou Meteorological Observatory, Guangzhou, China

ABSTRACT

Synthesizing high-fidelity radar reflectivity from satellite observations is crucial for weather forecasting and climate hazard tracking, particularly in regions with sparse radar coverage such as mountainous and oceanic areas. Existing methods either ensure global consistency but over-smooth details, or capture details while deviating from the true distribution. To address this challenge, we propose `CasS2R` a new cascaded framework consisting of deterministic and probabilistic components. The framework operates through the synergy of two meticulously designed modules: a *Physics-Constrained Structural Extractor*, which captures the macro-scale structure of the radar field that conforms to physical statistical distributions; and a *Flow-based Distribution Adaptor*, precisely maps the initial macroscopic result to the fine-grained distribution of radar observations. Comprehensive experiments on two large-scale, real-world datasets demonstrate that `CasS2R` significantly outperforms existing baseline models, not only improve the CSI_{50} by over 16% but also achieving a breakthrough in key structural fidelity metrics.

***Index Terms*—** Radar reconstruction, flow matching, meteorological observation

1. INTRODUCTION

Meteorological radar observation serves as a critical source of information for weather condition, playing an essential role in the observation and early warning of extreme weather events. However, the deployment of meteorological radars is highly uneven, with significant gaps over open oceans and complex mountainous regions. These coverage deficiencies create observable blind zones, which severely hinder maritime navigation safety and impair the capacity for early warning of sudden-onset hazards in mountainous areas. In contrast, satellites are not constrained by terrain or environmental limitations, offering continuous large-scale observation capabilities[1]. Training reconstruction models using satellite data thus presents a promising approach to mitigating the absence of radar coverage.

* is the corresponding author.

The success of deep learning has recently spurred significant advances across meteorological science. In the field of radar reconstruction, existing methods have diverged into deterministic and probabilistic. Deterministic models, primarily based on CNN[2, 3, 4] and Transformer[5] architectures, excel at capturing the large-scale structure. However, their optimization via pixel-wise losses (e.g., MSE) will inevitably causing the blurring of high-frequency details. In contrast, probabilistic models, particularly diffusion models, excel at generating small-scale weather phenomena by modeling the stochasticity of atmospheric processes through sampling from a learned distribution[6, 7, 8]. However, the unconstrained stochastic sampling can cause the generated fields to deviate from the real large-scale patterns, compromising structural accuracy[9]. **Consequently, previous methods struggling to reconcile the accurate reconstruction of global structures with the realistic synthesis of local stochastic details.**

To tackle this challenge, we propose `CasS2R` a novel radar reconstruction framework.`CasS2R` cascades the task into the reconstruction of the overall global observation and the refinement of small-scale systems. Specifically, the model first employs *Physics-Constrained Structural Extractor*, a deterministic module, to establish the overall distribution of the radar observations. Then, building upon this captured global distribution, a probabilistic module, *Flow-based Distribution Adaptor*, is used to generate small-scale weather phenomena, through solving Ordinary Differential Equation (ODE). Our contributions can be summarized as follows:

- We propose a cascaded framework `CasS2R` which decouples radar reconstruction into deterministic global distribution capture and stochastic small-scale detail modeling, enabling high-fidelity radar reconstruction.
- `CasS2R` contains two components: i) *Physics-Constrained Structural Extractor* provides a large-scale consistent estimation of the radar field; and ii) *Flow-based Distribution Adaptor* generates high-frequency detailed, capturing the stochastic nature of convective weather.
- Extensive evaluations on two large-scale dataset shows that `CasS2R` achieves a significant improvement of over 16% in CSI_{50} over strong baselines.

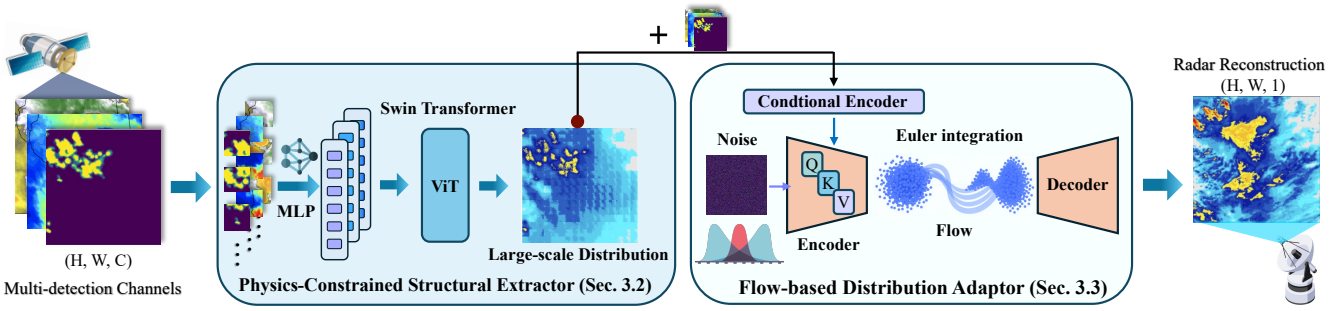

Fig. 1. The Framework of `CasS2R`.

2. RELATED WORK

Geostationary meteorological satellites overcome the terrain and environmental limitations of radars, enabling large-scale and continuous Earth observation, providing multi-spectral information[10]. Existing radar reconstruction approaches generally fall into two paradigms. Deterministic models, often based on UNets for spatial feature extraction and transformers for global context modeling[2, 3, 5], have become the dominant choice for radar reconstruction. However, by optimizing pixel-wise errors, they inevitably produce over-smoothed results and fail to capture high-frequency details of small-scale systems[11]. In contrast, probabilistic models, such as diffusion and flow-matching models, are particularly effective at representing the stochasticity of weather[6, 12]. Nevertheless, this stochastic advantage often comes at the expense of accurately reconstructing deterministic large-scale observations structures[9]. Thus, existing methods struggle to reconcile the accurate recovery of macroscopic structures with the realistic synthesis of microscopic stochastic details, leaving a critical gap for further exploration.

3. METHODOLOGY

To reconstruct high-fidelity radar observations, we introduce `CasS2R` a cascaded framework composed of deterministic and probabilistic componentsas shown in Figure 1.

3.1. Problem Definition

We formulate the radar reconstruction task as **cross-modal image-to-image translation**. Our goal is to learn a reconstruction model $\mathcal{F}_\theta : \mathbb{R}^{H\times W\times C} \to \mathbb{R}^{H\times W\times 1}$ that transforms a multi-channel satellite observation $\mathcal{X}$ into a corresponding radar reflectivity field $\hat{\mathcal{Y}}$. Here, H and W are spatial dimensions, and C is the number of satellite channels. The model is trained to minimize the divergence between the generated distribution $\hat{\mathcal{Y}}$ and the real observations $\mathcal{Y}$:

$$\hat{\mathcal{Y}} = \mathcal{F}_\theta(\mathcal{X}). \tag{1}$$

3.2. Physics-Constrained Structural Extractor

To capture the large-scale spatial organization of radar observations, we design *Physics-Constrained Structural Extractor* (PCSE), adopting a ViT module as the backbone[13]. The multi-channel satellite input $\mathcal{X} \in \mathbb{R}^{H\times W\times C}$ is first divided into N non-overlapping patches x_i, each of size $P \times P \times C$. Each patch is projected to an embedding $z \in R^D$ by adopting a linear layer.

$$z_i = \mathcal{G}(x_i) + p_i, \quad i = 1, \ldots, N, \tag{2}$$

where $\mathcal{G}$ denotes the embedding function and p_i is the positional encoding. The sequence $\{z_i\}_{i=1}^N$ is then fed into a stacked transformer block:

$$Z_{attn} = \text{MHSA}(\text{LN}(Z)) + Z, \quad Z' = \text{FFN}(\text{LN}(Z) + Z_{attn}, \tag{3}$$

where MHSA denotes multi-head self-attention[14], and FFN and LN represent a feedforward network and layer normalization, respectively. In this circumstance, $z \in R^D$ is first projected into query, key, and value matrices, i.e., $Q, K, V \in \mathbb{R}^{N\times D}$. The attention weights $\mathbf{A} \in \mathbb{R}^{N\times N}$ are then obtained via the softmax function:

$$\mathbf{A} = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V, \tag{4}$$

where d_k denotes the dimension of each attention head. By leveraging self-attention mechanism across the entire spatial domain, the ViT ensures that the reconstructed radar reflectivity preserves coherent storm structures and large-scale precipitation envelopes.

To balance over-prediction and under-prediction of reflectivity values, we design a physical- constrained loss. Specifically, it is defined as

$$\mathcal{L}_D = \frac{1}{n}\sum_{i=1}^{n} \exp\left(at_i^b\right) \cdot (\hat{y'}_i - y_i)^2, \tag{5}$$

where y and $\hat{y'}$ are the ground truth and predicted values, n is the number of pixels, and a and b are determined by calculating categorical bias of reflectivity threshold[4].

Model	CONUS3					SEVIR				
	MAE (↓)	SSIM (↑)	LPIPS (↓)	CSI_{15} (↑)	CSI_{50} (↑)	MAE (↓)	SSIM (↑)	LPIPS (↓)	CSI_{15} (↑)	CSI_{50} (↑)
SU3plus	0.0705	0.7624	0.5263	0.3319	0.1121	0.0405	0.4376	0.2034	0.6912	0.2989
SRViT	0.0689	0.7916	0.3635	0.4579	0.1332	0.0407	0.4353	0.2034	0.6762	0.3145
BBDM	0.0785	0.7199	0.3316	0.4039	0.1839	0.0458	0.6792	0.1984	0.6712	0.3627
CasS2R	**0.0662**	**0.8245**	**0.2717**	**0.4632**	**0.2573**	**0.0396**	**0.6997**	**0.1862**	**0.7012**	**0.4213**

Table 1. Quantitative comparison of different models on **CONUS3** and **SEVIR** datasets. Best results are bolded.

3.3. Flow-based Distribution Adaptor

The radar reconstruction results $\hat{y}'$ from multi-modal prediction tend to blur in small-scale, reducing their effectiveness. To capture the strong randomness, exhibited in the small-scale systems, we design *Flow-based Distribution Adaptor* (FDA), which employ flow matching technology. Let $\mathbf{y}$ denote the ground-truth radar observation, $\mathbf{c}$ the corresponding initial reconstruction from the deterministic module $\hat{y}'$ and satellite observations $\mathcal{X}$ (used as the condition), and $\mathbf{x}_0 \sim \mathcal{N}(0, I)$ a noise image. We define a time-dependent interpolation:

$$\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{y}, \quad \text{for } t \in (0,1), \tag{6}$$

and train a neural network $\mathbf{v}_\theta(\mathbf{x}, t, \mathbf{c})$ to approximate the optimal transport vector field that maps $\mathbf{x}_0$ to $\mathbf{y}$ under the guidance of $\mathbf{c}$. Specifically, we minimize the following flow matching loss:

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{\mathbf{x}_0,\mathbf{y},t}\left[\left\|\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - \frac{\mathbf{y}-\mathbf{x}_0}{t(1-t)}\right\|^2\right]. \tag{7}$$

Here, the velocity $\frac{\mathbf{y}-\mathbf{x}_0}{t(1-t)}$ corresponds to the vector that would transport the sample $\mathbf{x}_t$ along the shortest path to the target.

At inference time, we adopt Euler's method for ODE integration using a fixed number of steps to obtain the refined observation $\mathbf{x}(1) \sim \mathcal{Y}$. Specifically, we initialize $\mathbf{x}(0) \sim \mathcal{N}(0, I)$ and numerically solve the following ODE problem:

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{v}_\theta(\mathbf{x}(t), t, \mathbf{c}), \quad \mathbf{x}(0) = \mathbf{x}_0, \tag{8}$$

In essence, the procedure operates as an atmospheric simulator, with blurry observations defining synoptic-scale boundary conditions that steer random perturbations toward dynamically consistent fine-scale convective structures.

4. EVALUATION

4.1. Evaluation Settings

- **Datasets**: We evaluate `CasS2R` on two real world datasets: (i) **CONUS3**[4], containing 82,449 observations from the GOES-R satellites (ABI infrared channels 7, 9, 13, and GLM) with corresponding MRMS composite radar reflectivity over the United States (2020-2022); and (ii) **SEVIR**[15], comprising 11,640 strom events with Vertically Integrated Liquid (VIL) radar data, GOES-16 infrared channels (6.9 m and 10.7 m), and GLM.
- **Baselines**: We compare our `CasS2R` with strong baselines that belong to 2 categories: (i) **Deterministic models**: **SU3plus**[2], which leverages the Swin-UNet backbone, and **SRViT**[5], which adopts a ViT-based architecture. (ii) **Probabilistic models**: **BBDM** [16] models image-to-image translation as a stochastic Brownian Bridge process.
- **Metrics**: To evaluate both the accuracy and physical realism of `CasS2R`, we employ a comprehensive set of metrics including, Mean Absolute Error (MAE) measures per-pixel numerical deviation. Critical Success Index (CSI) to assess model's ability to detect storms of varying severity, Structural Similarity Index (SSIM) and Learned PerceptualImagePatch Similarity (LPIPS) evaluate the preservation of storm morphology and textural sharpness.
- **Implementation**: `CasS2R` is implmented using PyTorch and trained on 4 NVIDIA A800 GPUs and a Intel(R) Xeon(R) Platinum 8358P CPU @ 2.60GHz 100 epochs with batch size 32. For optimization, we use AdamW with an initial learning rate of 5×10^{-4}, scheduled by a OneCycle policy. Each dataset is split chronologically, with the first 80% for training and the last 20% for testing.

4.2. Overall Performance

Table 1 compares `CasS2R` with representative models from two categories in **CONUS3** and **SEVIR**.

- `CasS2R` consistently outperforms all baselines on both **CONUS3** and **SEVIR**, reduced LPIPS by 22.05% and 6.55%, while boosting SSIM by 14.58% and 3.02%, and CSI_{50} by 39.91% and 16.16%, respectively. In meteorology, high reflectivity values serve as crucial indicators of severe weather phenomena such as intense convective storms. The substantial improvement in high-threshold CSI highlight `CasS2R`'s ability to accurately reconstruct radar observations by modeling the interplay between deterministic large-scale forcing and small-scale stochasticity.
- While deterministic models achieve low MAE by robustly capturing large-scale precipitation structures, they fail to recover high-value regions tied to severe convection, yielding

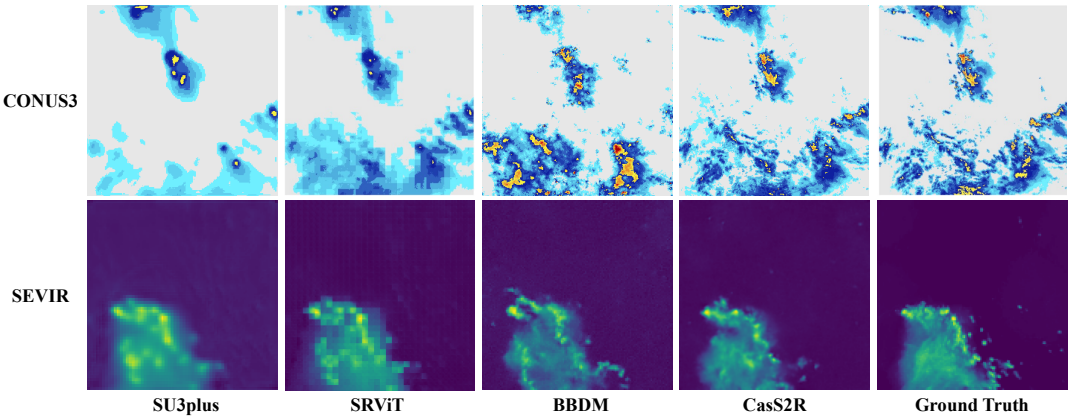

Fig. 2. Radar reconstruction of different model.

poor CSI_{50}. In contrast, probabilistic models generate realistic details but compromise structural fidelity, leading to inferior CSI_{15} in reproducing overall precipitation patterns.

CONUS3					
Model	MAE (↓)	SSIM (↑)	LPIPS (↓)	CSI_{15} (↑)	CSI_{50} (↑)
w/o PCSE	0.0756	0.7392	0.3341	0.3516	0.2112
w/o FDA	**0.0658**	0.7892	0.4579	0.3919	0.1542
CasS2R	0.0662	**0.8245**	**0.2717**	**0.4632**	**0.2573**
SEVIR					
Model	MAE (↓)	SSIM (↑)	LPIPS (↓)	CSI_{15} (↑)	CSI_{50} (↑)
w/o PCSE	0.0462	0.6382	0.2073	0.6791	0.3982
w/o FDA	0.0417	0.5782	0.2137	**0.7019**	0.3841
CasS2R	**0.0396**	**0.6997**	**0.1862**	0.7012	**0.4213**

Table 2. Ablation study results on **CONUS3** and **SEVIR**.

4.3. Key Component Assessment

To assess the impact of each design in `CasS2R` we perform ablation experiments cross two datasets (Table 2). The results yield two key insights:

- Removing the *Physics-Constrained Structural Extractor* (w/o PCSE) increases MAE and decreases CSI15, indicating that although the deterministic module struggles with fine-scale convective features, it plays a crucial role in calibrating overall observations and mitigating overestimation.
- Excluding the *Flow-based Distribution Adaptor* (w/o FDA) leads to a substantial performance drop in reconstructing small convective regions-for instance, CSI_{50} decreases by 21.83% and 9.68% on **CONUS3** and **SEVIR**, highlighting its importance in capturing fine-scale convective patterns.

4.4. Case Study

Figure 2 presents a cross-model visual comparison of radar reconstructions on both datasets. In all circumstances, `CasS2R` preserves remarkable spatial fidelity and intensity consistency. Even under highly complex convective patterns, it accurately captures both global organization and local details, aligning closely with ground-truth observations. By contrast, SU3plus and SRViT tend to produce locally blurred rainfall areas and misplaced storm cores, leading to underestimated peak intensities. SRViT outputs further exhibit noticeable tiling artifacts, likely caused by token-wise decoding, which undermines natural transitions and coherence. BBDM maintains visual plausibility but suffers from spatial drift, particularly in the temporal evolution of storm centers. These results highlight the superior capability of `CasS2R` to reconstruct radar fields from satellite observations.

4.5. Real-Time Inference Efficiency

The Flow-based Distribution Adaptor uses 100 ODE steps during inference. On a single NVIDIA A800 GPU, the model reconstructs a 384×384 radar observation in 12.43 ms, with a memory footprint of 283 MB. These results demonstrate that CasS2R is well suited for near-real-time operational settings.

5. CONCLUSION

We propose a cascaded deterministic-probabilistic framework, named `CasS2R`, for reconstructing radar reflectivity from satellite observations. Unlike prior approaches that rely solely on deterministic or probabilistic paradigms, `CasS2R` explicitly decouples the task into two complementary stages: global structure reconstruction and local detail refinement. Specifically, a deterministic module, establishes the large-scale radar distribution, while a probabilistic module, generates fine-scale convective structures by solving ODEs. Extensive experiments demonstrate that `CasS2R` outperforms strong baselines across multiple metrics, achieving up to 16% improvement in CSI_{50}. Beyond quantitative gains, `CasS2R` delivers reconstructions with both high-value weather features (e.g., storm cores) and realistic fine-grained structures. Overall, `CasS2R` establishes a powerful foundation for bridging satellite and radar observations.

6. ACKNOWLEDGEMENTS

This work is supported partly by China NSFC Grant (62472366), the Project of DEGP (No.2023KCXTD042, 2024GCZX003), Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things (No.2023B1212010007), 111 Center (No.D25008), Shenzhen Science and Technology Foundation (ZDSYS20190 902092853047) Guangdong-Hong Kong-Macao Greater Bay Area Academy of Meteorological Research (GHMA2025Y03). Guangdong Basic and Applied Basic Research Foundation (2025A1515510015).

7. REFERENCES

[1] Joel R Norris, Robert J Allen, Amato T Evan, Mark D Zelinka, Christopher W ODell, and Stephen A Klein, "Evidence for climate change in the satellite cloud record," *Nature*, vol. 536, no. 7614, pp. 72–75, 2016.

[2] Jingtao Li, Zhixuan Zhou, Xintong Zhao, Mounir Kaaniche, Congcong Wang, and Ding Liu, "Su3plus: An enhanced swin-unet for synthesizing vertically integrated liquid from multiple meteorological satellite data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 63, pp. 1–13, 2025.

[3] Jianwei Si, Haonan Chen, and Lei Han, "Enhancing weather radar reflectivity emulation from geostationary satellite data using dynamic residual convolutional network," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 63, pp. 1–11, 2025.

[4] Kyle A Hilburn, Imme Ebert-Uphoff, and Steven D Miller, "Development and interpretation of a neural-network-based synthetic radar reflectivity estimator using goes-r satellite observations," *Journal of Applied Meteorology and Climatology*, vol. 60, no. 1, pp. 3–21, 2020.

[5] Jason Stock, Kyle Hilburn, Imme Ebert-Uphoff, and Charles Anderson, "Srvit: Vision transformers for estimating radar reflectivity from satellite observations at scale," in *ICML 2024 Workshop on Machine Learning for Earth System Modeling*, July 2024.

[6] Minseok Seo, YoungTack Oh, Doyi Kim, Dongmin Kang, and Yeji Choi, "Improving flood insights: Diffusion-based sar to eo image translation," in *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*, 2023.

[7] Wei Zhang, Xinyu Zhang, Zhuyu Jin, Youqi Wen, and Jie Liu, "Inpainting radar missing data via denoising iterative restoration," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 10715–10725, 2024.

[8] Xuming He, Zhiwang Zhou, Wenlong Zhang, Xiangyu Zhao, Hao Chen, Shiqi Chen, and Lei Bai, "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)*, 2025, pp. 1–5.

[9] Junchao Gong, Lei Bai, Peng Ye, Wanghan Xu, Na Liu, Jianhua Dai, Xiaokang Yang, and Wanli Ouyang, "Cascast: Skillful high-resolution precipitation nowcasting via cascaded modelling," in *ICML*, 2024.

[10] Pan Xia, Lu Zhang, Min Min, Jun Li, Yun Wang, Yu Yu, and Shengjie Jia, "Accurate nowcasting of cloud cover at solar photovoltaic plants using geostationary satellite images," *Nature Communications*, vol. 15, no. 1, pp. 510, 2024.

[11] Demin Yu, Xutao Li, Yunming Ye, Baoquan Zhang, Chuyao Luo, Kuai Dai, Rui Wang, and Xunlai Chen, "Diffcast: A unified framework via residual diffusion for precipitation nowcasting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 27758–27767.

[12] Siwei Tu, Ben Fei, Weidong Yang, Fenghua Ling, Hao Chen, Zili Liu, Kun Chen, Hang Fan, Wanli Ouyang, and Lei Bai, "Satellite observations guided diffusion model for accurate meteorological states at arbitrary resolution," in *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, June 2025, pp. 28071–28080.

[13] Kun Wang, Hao Wu, Yifan Duan, Guibin Zhang, Kai Wang, Xiaojiang Peng, Yu Zheng, Yuxuan Liang, and Yang Wang, "Nuwadynamics: Discovering and updating in causal spatio-temporal modeling," in *The Twelfth International Conference on Learning Representations*, 2024.

[14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[15] Mark Veillette, Siddharth Samsi, and Chris Mattioli, "Sevir: A storm event imagery dataset for deep learning applications in radar and satellite meteorology," *Advances in Neural Information Processing Systems*, vol. 33, pp. 22009–22019, 2020.

[16] Bo Li, Kaitao Xue, Bin Liu, and Yu-Kun Lai, "Bbdm: Image-to-image translation with brownian bridge diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, 2023, pp. 1952–1961.