

## Precipitation Nowcasting Using Physics Informed Discriminator Generative Models

Yin, Junzhe; Meo, Cristian; Roy, Ankush; Cher, Zeineh Bou; Lică, Mircea; Wang, Yanbo; Imhoff, Ruben; Uijlenhoet, Remko; Dauwels, Justin

**DOI**

[10.23919/EUSIPCO63174.2024.10715141](https://doi.org/10.23919/EUSIPCO63174.2024.10715141)

**Publication date**

2024

**Document Version**

Final published version

**Published in**

32nd European Signal Processing Conference, EUSIPCO 2024 - Proceedings

**Citation (APA)**

Yin, J., Meo, C., Roy, A., Cher, Z. B., Lică, M., Wang, Y., Imhoff, R., Uijlenhoet, R., & Dauwels, J. (2024). Precipitation Nowcasting Using Physics Informed Discriminator Generative Models. In *32nd European Signal Processing Conference, EUSIPCO 2024 - Proceedings* (pp. 1967-1971). (European Signal Processing Conference). European Signal Processing Conference, EUSIPCO. <https://doi.org/10.23919/EUSIPCO63174.2024.10715141>

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

***Green Open Access added to TU Delft Institutional Repository***

***'You share, we take care!' - Taverne project***

***<https://www.openaccess.nl/en/you-share-we-take-care>***

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

# Precipitation Nowcasting Using Physics Informed Discriminator Generative Models

Junzhe Yin<sup>1\*</sup>, Cristian Meo<sup>1\*</sup>, Ankush Roy<sup>1</sup>, Zeineh Bou Cher<sup>1</sup>, Mircea Lică<sup>1</sup>, Yanbo Wang<sup>1</sup>, Ruben Imhoff<sup>2</sup>  
Remko Uijlenhoet<sup>1</sup>, Justin Dauwels<sup>1</sup>

<sup>1</sup>Delft University of Technology, Netherlands

<sup>2</sup>Deltares, Netherlands

\*Equal Contribution

Corresponding author: Junzhe.Yin1999@gmail.com

**Abstract**—Nowcasting leverages real-time atmospheric conditions to forecast weather over short periods. State-of-the-art models, including PySTEPS, encounter difficulties in accurately forecasting extreme weather events because of their unpredictable distribution patterns. In this study, we design a physics-informed neural network to perform precipitation nowcasting using the precipitation and meteorological data from the Royal Netherlands Meteorological Institute (KNMI). This model draws inspiration from the novel Physics-Informed Discriminator GAN (PID-GAN) formulation, directly integrating physics-based supervision within the adversarial learning framework. The proposed model adopts a GAN structure, featuring a Vector Quantization Generative Adversarial Network (VQ-GAN) and a Transformer as the generator, with a temporal discriminator serving as the discriminator. Our findings demonstrate that the PID-GAN model outperforms numerical and SOTA deep generative models in terms of precipitation nowcasting downstream metrics.

## I. INTRODUCTION

More frequent global extreme precipitation has led to severe flooding, soil erosion, loss of agricultural productivity, and heightened health risks [1]. Traditional Numerical Weather Prediction (NWP) models, though comprehensive, face challenges in short-term forecasting of rainfall due to their high computational requirements and too low spatial-temporal resolution [2], [3]. Precipitation nowcasting techniques aim to provide accurate forecasts of upcoming precipitation events within six hours for local regions, reducing response time and efforts to handle extreme weather events [4]–[6]. Radar extrapolation methods, such as PySteps, leverage real-time data and use optical flow and statistical analysis techniques to improve weather prediction accuracy [7]. However, while statistical nowcasting methods eliminate the need for historical data, they do not explicitly account for the growth and decay processes of convective rainfall, limiting their usefulness for severe rainfall events. Due to limitations in radar extrapolation methods, there has been a shift toward deep learning models [4], [8], which have shown promise in weather prediction by capturing complex spatio-temporal patterns without relying on traditional data assimilation techniques.

Deep generative models like Generative Adversarial Networks (GAN) [9] and Variational Autoencoder (VAE) [10] have excelled in precipitation nowcasting, offering more accurate and realistic forecasts by modelling the distribution underlying the precipitation dynamics [4], [8], [11]. Nevertheless,

these models have many limitations, for instance, they are not able to produce consistent prediction over medium- to long-term horizons and struggle with extreme event modeling and prediction. Moreover, a notable limitation of these models is their tendency to overlook fundamental physical laws, resulting in predictions that may not always align with them, particularly in scenarios not encountered during training. Physics-informed machine learning (PIML) [12] has emerged as a promising solution, enhancing weather and climate forecasting by improving physical consistency [13]. In this paper, we propose a novel PIML model that integrates deep generative models with physical priors of precipitation, aiming to produce accurate and physically consistent precipitation forecasts.

## A. Dataset

This paper explores precipitation nowcasting in the Netherlands, using weather radar data from the Royal Netherlands Meteorological Institute (KNMI) for 2008–2021 [8], along with hourly meteorological data from Automatic Weather Stations (AWS) [14] and ERA5 reanalyses [15]. The meteorological data from AWS contains hourly observations of wind speed, wind direction, air temperature, dew point temperature and global radiation. ERA5 offers hourly estimates for a wide variety of quantities related to the atmosphere, ocean waves, and the land surface [15]. We use radar precipitation maps with a 5-minute temporal resolution and a 1 km spatial resolution. Our study focuses on a  $256 \times 256$  pixel area covering most of the country and 12 Dutch catchment regions, as detailed in [3, Fig. 1]. Following [8], an event is classified as extreme if its average rainfall over three hours ranks in the top 1% of all recorded events between 2008 and 2021. This threshold is defined as exceeding 5 mm per 3 hours, as shown in [8, Table. 1]. To match the temporal resolution of AWS and ERA5 meteorological data with the used radar dataset, we apply cubic interpolation to estimate half-hour intervals of the latter (e.g.,  $T - 60$ ,  $T - 30$ , to  $T + 180$  minutes) from the original hourly data points (e.g.,  $T - 60$ ,  $T$ ,  $T + 60$ ,  $T + 120$ ,  $T + 180$  minutes). To address the limited spatial coverage of the AWS and ERA5 reanalyses datasets, we applied kriging interpolation using the PyKrig package [16], achieving a spatial resolution compatible with our radar data. Kriging uses spatial autocorrelation and variance to accurately estimate

values in unmeasured locations, creating maps that align with the resolutions of the radar map [17].

### B. Problem Formulation

In the last few years, precipitation nowcasting using deep learning models have been cast as a video prediction problem [8], where given an input spatio-temporal sequence of  $N$  frames  $\mathbf{x}_{\text{in}} \in \mathbb{R}^{N \times H \times W \times C}$ , where  $H, W$  denote the spatial resolution and  $C$  represents the image channels or the different types of measurements (e.g., radar maps, humidity maps, etc). The goal is to predict the next  $M$  precipitation maps  $\mathbf{x}_{\text{out}} \in \mathbb{R}^{M \times H \times W \times 1}$  and classify extreme events based on the defined threshold. Throughout this work we use  $N = 3$  and  $M = 6$ .

## II. RELATED WORKS

Generative Adversarial Networks, as the current state-of-the-art in various fields, have seen many adaptable variations in recent years, demonstrating their flexibility. They serve as key architectures in precipitation nowcasting tasks. For instance, the Adversarial Extrapolation Neural Net (AENN) [11], a variation of GAN with 2 discriminators, has shown superior performance in weather radar echo extrapolation than NWP approaches. In [4], a GAN utilizing temporal and spatial discriminators was proposed. This advancement underscores the potential of GANs in enhancing the accuracy of precipitation nowcasting through sophisticated data generation techniques.

Despite their advancements, these models face persistent challenges, such as delivering consistently clear and precise forecasts and adequately generalizing across diverse weather conditions, particularly those underrepresented in training datasets. A significant obstacle faced by sophisticated neural networks such as AENN is their lack of interpretability. A critical hurdle with advanced neural networks lies in their interpretability. The uninterpretable reasoning behind their predictions complicates the understanding of their decision-making processes, making it difficult for experts to fully trust and utilize these models. Moreover, these systems must adhere to the fundamental physical principles that govern meteorological phenomena, ensuring that their predictive capabilities do not compromise physical accuracy. Maintaining a balance between physical accuracy and predictive performance represents a significant challenge, underscoring the demand for innovation in meteorological forecasting. This balance is crucial for developing reliable models that can accurately predict weather patterns while adhering to the fundamental principles of meteorology. Recently, PID-GAN, a physics-informed discriminator GAN formulation that does not suffer from an imbalance of gradient flow compared to physics-informed neural networks (PINN) [18], was introduced in [19]. In this work, we adopt the PID-GAN framework to integrate physical constraints into our model, ensuring the generation of physically coherent precipitation forecasts. This includes the incorporation of moisture conservation equations, as outlined in [20], to accurately capture the dynamics of precipitation.

## III. METHODOLOGY

In this section, we define and describe the proposed PID-GAN architecture used in this work. Moreover, we describe the physics-based loss function used to inject the moisture conservation equation information into the PID-GAN architecture.

### A. PID-GAN: Model Architecture

This section describes the backbone architecture used for the proposed PID-GAN model, which is inspired by the successful implementation detailed in [4]. Our design adopts this proven structure, incorporating a generator alongside spatial and temporal discriminators, due to its demonstrated effectiveness in capturing complex spatio-temporal relationships in data. This approach ensures a robust framework capable of handling the intricate dynamics of precipitation nowcasting by efficiently learning from both the spatial patterns and temporal sequences inherent in meteorological data.

1) *Generator*: The generator in our model incorporates a deep generative model with two components: Vector Quantization Generative Adversarial Network (VQ-GAN) [21], that learns the mapping between precipitation maps and discrete tokens, and an Autoregressive Transformer (AT) [22] that models the dynamics between tokens of consecutive timesteps. **VQ-GAN**: The architecture of the VQ-GAN model consists of a CNN encoder ( $E$ ) and decoder ( $G$ ), a codebook ( $\mathcal{Z}$ ), which define a VQ-VAE [23] and a patch-based discriminator [24], indicated as spatial discriminator ( $D$ ). The VQ-VAE is trained by optimizing:

$$\mathcal{L}_{\text{vq-vae}} = \|x - \hat{x}\|_1 + \|sg[E(x)] - z_q\|_2^2 + \|sg[z_q] - E(x)\|_2^2 + \mathcal{L}_{\text{perceptual}}(x, \hat{x}). \quad (1)$$

Here,  $sg[\cdot]$  represents the stop-gradient operation, which prevents back-propagating gradients. The loss function comprises four terms: the reconstruction loss  $\mathcal{L}_{\text{rec}} = \|x - \hat{x}\|_1$ , comparing the original input  $x$  (radar maps) with its reconstruction  $\hat{x}$  (reconstructed maps). The commitment loss, covered by the second and third terms, penalizes discrepancies between the encoded representations and codebook entries and optimizes the codes within the codebook. The fourth term, perceptual loss, assesses high-level semantic differences between  $x$  and  $\hat{x}$  [21]. The VQ-GAN is optimized training adversarially VQ-VAE, which acts as a generator, and a patch-based discriminator (i.e., spatial discriminator), using the following loss function:

$$\text{argmin}_{E, G, \mathcal{Z}} \max_D \mathbb{E}_{x \sim p(x)} [\mathcal{L}_{\text{vq-vae}}(E, G, \mathcal{Z}) + \lambda \mathcal{L}_{\text{GAN}}(\{E, G, \mathcal{Z}\}, D)], \quad (2)$$

$$\mathcal{L}_{\text{GAN}}(\{E, G, \mathcal{Z}\}, D) = [\log D(x) + \log(1 - D(\hat{x}))] \quad (3)$$

$$\lambda_{\text{GAN}} = \frac{\nabla_G [\mathcal{L}_{\text{rec}}]}{\nabla_G [\mathcal{L}_{\text{GAN}}] + \delta}. \quad (4)$$

Here,  $\mathcal{L}_{\text{GAN}}(\{E, G, \mathcal{Z}\}, D)$  represents the discriminator loss and  $\lambda_{\text{GAN}}$  is the adaptive weight determined by  $\nabla_G[\cdot]$ , which represents the gradient of the input concerning the final layer of decoder.  $\delta = 10^{-6}$  is a scalar for numerical stability.

**Autoregressive Transformer**: The AT architecture aims to model the dynamics between consecutive precipitation maps [22]. The ground truth precipitation maps are quantized into

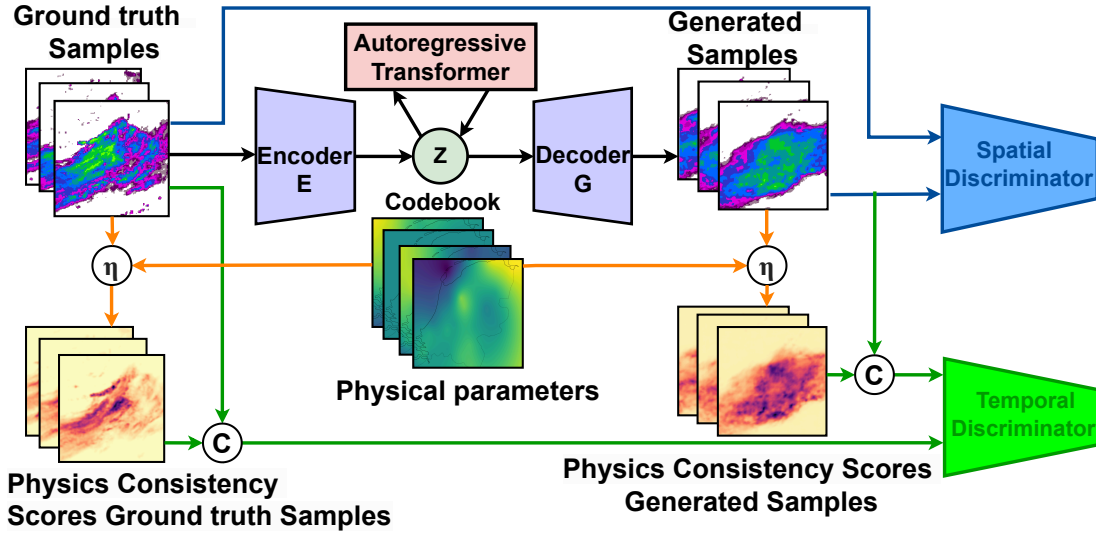


Fig. 1: The proposed PID-GAN model structure. C means concatenation and  $\eta$  represent  $\eta_k = e^{-\lambda \mathcal{R}^{(k)}(x, \hat{x})}$ , referring to the equation of the physics consistency score.

$\mathbf{z}_q = q(\mathbf{E}(x))$ , producing a sequence  $\mathbf{s} \in \{0, \dots, |\mathbf{Z}|-1\}^{h \times w}$ , representing VQ-GAN codebook indices. These indices are transformed into continuous vectors by an embedder and augmented with positional embeddings to provide order information. The transformer then processes these vectors, with the head module refining the output into logits, which represent the probability of using a specific token. These logits are used to compute a cross-entropy loss that compares predicted token probabilities with the actual tokens:

$$\mathcal{L}_{\text{Transformer}} = \mathbb{E}_{x \sim p(x)} \left[ -\log \prod_{i=1}^N p(\mathbf{s}_i | \mathbf{s}_{<i}) \right]. \quad (5)$$

Given a sequence of indices  $\mathbf{s}_{<i}$ , the transformer is trained to predict the distribution of the consecutive indices  $\mathbf{s}_i$ . The AT employs a causal attention mechanism which accesses only previously seen and current tokens when predicting the next one in a sequence, enabling efficient and context-sensitive output production.

### B. Physics-Informed Discriminator: PID-GAN

Following the PID-GAN [19] framework, we use physics residuals to compute a physics consistency score ( $\eta$ ) for each prediction, indicating the likelihood of the prediction being physically consistent. These physics consistency scores are fed into a temporal discriminator [11] as additional inputs, such that it distinguishes between real and fake sequences of precipitation maps by learning from the underlying distribution of labelled points and using the additional physics supervision.

**Estimating Physics Consistency Scores:** Formally, we compute the physics consistency score of a prediction  $\hat{x}$  for the  $k$ -th physical constraints as  $\eta_k = e^{-\lambda \mathcal{R}^{(k)}(x, \hat{x})}$ , where  $\mathcal{R}^{(k)}$  represents the physical equation used to describe the phenomena. The larger  $\eta_k$ , the more prediction  $\hat{x}$  obeys the

$k$ -th physical constraint. Following [19], the temporal physics-informed discriminator is trained by optimizing:

$$\mathcal{L}_D(\phi) = -\frac{1}{N} \sum_{i=1}^N \log(D(x_i, \eta_i)) - \frac{1}{N} \sum_{i=1}^N \log(1 - D(\hat{x}_i, \hat{\eta}_i)). \quad (6)$$

Here,  $\eta_i$  and  $\hat{\eta}_i$  represent the physics consistency score for ground truth and generated data, ensuring alignment with physical laws.

**Moisture Conservation Equation:** The Moisture Conservation Equation [25], [26] in NWP, which describes the relationship between atmospheric moisture content, evaporation, and precipitation (P), is given by:

$$\frac{\partial q}{\partial t} = -u \frac{\partial q}{\partial x} - v \frac{\partial q}{\partial y} - \omega \frac{\partial q}{\partial z} + ET - P \quad (7)$$

$$ET = 0.65 \frac{\Delta}{\Delta + \gamma} \frac{R_s}{\lambda}. \quad (8)$$

Here,  $u$  is the west-to-east wind component ( $m/s$ ),  $v$  is the south-to-north wind component ( $m/s$ ),  $\omega$  is the vertical wind component ( $m/s$ ),  $q$  is the specific humidity ( $g \ g^{-1}$ ),  $ET$  is evapotranspiration rate ( $mm/h$ ),  $\Delta$  is the derivative with respect to temperature of the saturation vapour pressure ( $Pa/^\circ C$ ),  $\gamma$  is the psychrometric constant ( $J/m^2$ ),  $\lambda$  is the latent heat of vaporization ( $J/g$ ), and  $R_s$  is global radiation ( $Pa/^\circ C$ ). The evapotranspiration can be estimated using the Makkink Equation (8) [27], which relies on temperature and solar radiation data, yielding relatively accurate results in cold and temperate humid climates.

Measuring vertical wind speed  $\omega$  poses challenges due to its low intensity and high spatial variability, the requirement for sensitive equipment, atmospheric stability, and the high costs and complexity of accurate measurements, resulting in a lack of datasets with these measurements. Omitting the term  $\omega \frac{\partial q}{\partial z}$  simplifies the process but risks overlooking crucial

moisture transport between atmospheric layers, potentially impacting the balance of moisture within the atmosphere's three-dimensional dynamics. To compensate for the lack of vertical wind speed data, the model shifts focus to the horizontal wind components (U and V) at different elevation levels, using the ERA5 dataset for wind measurements at 100 meters ( $u_{100}$  and  $v_{100}$ ) and AWS data at 10 meters ( $u_{10}$  and  $v_{10}$ ). It's important to note that atmospheric interactions, which significantly influence weather patterns, extend up to the end of the troposphere, approximately 10 km in altitude. Therefore, by focusing on wind speed within the lower 100 meters, the approach inevitably entails a degree of uncertainty, given the comprehensive atmospheric interactions occurring beyond this range. By integrating these horizontal components from both 10-meter and 100-meter altitudes, the model adapts to variations in altitude for atmospheric moisture analysis. This approach simplifies equation (7) to approximate moisture transport across altitudes without needing direct vertical wind speed measurements, offering a more detailed view of the atmosphere's moisture dynamics. As a result, we can define the physical constraint for the proposed PID-GAN as:

$$\mathcal{R}_q = -\frac{\partial q}{\partial t} - u_{10} \frac{\partial q}{\partial x} - v_{10} \frac{\partial q}{\partial y} - u_{100} \frac{\partial q}{\partial x} - v_{100} \frac{\partial q}{\partial y} + ET - P. \quad (9)$$

where the equation is computed at the pixel level.

#### IV. EXPERIMENTS

Our investigation addresses two main research questions: (a) Can physical supervision improve the accuracy of precipitation nowcasting? (b) Does incorporating physical data into the model's design enhance its ability to detect extreme precipitation events? We compare the proposed model to a classic benchmark, namely Pysteps [7], a temporally consistent video prediction benchmark, TECO [28], and two benchmarks that use Extreme Value loss regularization, namely Nuwä-EVL [8] and NowcastingGPT+EVL [29]. The nowcasting task is configured with parameters set to  $N = 3$  conditioning timesteps and  $M = 6$  predicted maps, with each timestep representing a realistic scenario of 30 minutes. Outputs from all models are the average ensembles of five sample predictions. The PySteps model employed in this study adheres to the probabilistic configuration parameters outlined in [3].

To quantitatively assess the predictions we calculate visual fidelity metrics including Mean Squared and Absolute Errors (MSE, MAE) and Pearson Correlation Score (PCC), and nowcasting metrics, such as Critical Success Index (CSI), False Alarm Ratio (FAR) and Fractional Skill Score (FSS). Precipitation thresholds of 1 and 8 mm are set for CSI and FAR, and FSS is evaluated at spatial scales of 1, 10, and 20 km with a 1 mm threshold. Since fidelity metrics cannot capture extreme event classification, we plot a precision-recall curve of the extremes to assess the considered baselines in terms of extreme classification capabilities. Table I presents a comparison of nowcasting downstream performances for the validated models across the entire study area. The PID-GAN model's advancements are particularly notable in Mean Squared Error (MSE) and Pearson Correlation Score (PCC), indicating a strong correlation with actual precipitation events and a high degree of prediction accuracy. Additionally, the model excelled in the Fractional Skill Score (FSS) across different spatial




Fig. 2: The precision-recall curves for detecting extreme events over 3 hours at 12 Dutch catchments. Every point from right to left represents a different precipitation threshold (0.5 to 10mm/3h) for prediction and a fixed threshold for ground truth by definition of extreme events [8].

scales and significantly outperformed all benchmarks in the Critical Success Index (CSI) at both light (1mm) and heavy (8mm) precipitation thresholds. This performance suggests an enhanced ability to detect and accurately predict a wide range of precipitation events, from light showers to severe storms.

Figure 2 highlights the PID-GAN model's outstanding performance in detecting extreme precipitation events within 12 Dutch catchments [3], as illustrated by its precision-recall curves. These curves not only underscore the model's superior capability in forecasting extreme events when compared to all benchmarks but also reveal PID-GAN's exceptional balance between precision and recall, evidenced by achieving the highest area under the Precision-Recall curve (AUC). This is particularly noteworthy in the context of effective nowcasting, where identifying extreme weather patterns is crucial. Additionally, the analysis reveals a significant insight: removing physical constraints from the PID-GAN model (PID-GAN(-P)) leads to a decrease of 6.17% in AUC. This drop highlights the critical role that physical constraints play in enhancing the accuracy of precipitation map predictions, further validating the importance of integrating physical data into the model's design for improved performance.

#### V. CONCLUSION

In conclusion, the proposed PID-GAN model has demonstrated significant effectiveness in nowcasting precipitation and accurately predicting extreme precipitation events. Addressing a notable challenge in nowcasting, the model not only outperforms existing benchmarks in terms of accuracy but also significantly enhances forecast precision. This study innovates by integrating the moisture conservation equation as a physical constraint, offering a novel approach to improving forecast ac-

TABLE I: 3-hour averaged precipitation nowcasting skill of different models (Pixel-level evaluation). Top and second-best performances are highlighted in bold and underlined, respectively. PID-GAN(-P) represents the proposed model without physical constraints, and PID-GAN(-PT) represents the proposed model without physical constraints and temporal discriminator.

	PySTEPS	TECO	Nuwā-EVL	NowcastingGPT+EVL	NowcastingGPT	PID-GAN	PID-GAN(-P)	PID-GAN(-PT)
PCC (↑)	0.219	0.149	0.202	0.253	0.241	<b>0.313</b>	0.288	0.250
MAE (↓)	0.798	<b>0.664</b>	0.938	0.714	0.725	0.686	0.706	0.692
MSE (↓)	4.210	3.335	3.592	3.298	3.293	<b>3.117</b>	3.162	3.271
CSI(1mm) (↑)	0.250	0.097	0.262	0.267	0.21	<b>0.313</b>	0.296	0.234
CSI(8mm) (↑)	0.008	0.001	0.006	0.009	0.005	<b>0.011</b>	0.008	0.004
FAR(1mm) (↓)	0.617	0.662	0.623	0.587	0.579	0.583	0.601	<b>0.549</b>
FAR(8mm) (↓)	0.592	<b>0.361</b>	0.399	0.502	0.513	0.529	0.499	0.435
FSS(1km) (↑)	0.375	0.163	0.394	0.432	0.414	<b>0.451</b>	0.430	0.428
FSS(10km) (↑)	0.467	0.211	0.456	0.493	0.463	<b>0.534</b>	0.510	0.481
FSS(20km) (↑)	0.522	0.248	0.498	0.534	0.508	<b>0.591</b>	0.565	0.521
AUC (↑)	0.454	0.378	0.516	0.538	0.510	<b>0.567</b>	0.532	0.520

curacy. Future research will aim to further enhance the model's capabilities by incorporating additional physical constraints, such as the impact of air temperature on extreme precipitation events, and by evaluating the model's performance across diverse geographical settings.

## REFERENCES

- [1] H. Tabari, "Climate change impact on flood and extreme precipitation increases with water availability," *Scientific Reports*, vol. 10, no. 1, p. 13768, 2020.
- [2] R. Prudden, S. Adams, D. Kangin, N. Robinson, S. Ravuri, S. Mohamed, and A. Arribas, "A review of radar-based nowcasting of precipitation and applicable machine learning techniques," *arXiv preprint arXiv:2005.04988*, 2020.
- [3] R. Imhoff, C. Brauer, A. Overeem, A. Weerts, and R. Uijlenhoet, "Spatial and temporal evaluation of radar rainfall nowcasting techniques on 1,533 events," *Water Resources Research*, vol. 56, no. 8, p. e2019WR026723, 2020.
- [4] S. Ravuri, K. Lenc, M. Willson, D. Kangin, R. Lam, P. Mirowski, M. Fitzsimons, M. Athanassiadou, S. Kashem, S. Madge *et al.*, "Skilful precipitation nowcasting using deep generative models of radar," *Nature*, vol. 597, no. 7878, pp. 672–677, 2021.
- [5] Y. Zhang, M. Long, K. Chen, L. Xing, R. Jin, M. I. Jordan, and J. Wang, "Skilful nowcasting of extreme precipitation with nowcastnet," *Nature*, 2023.
- [6] K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian, "Accurate medium-range global weather forecasting with 3d neural networks," *Nature*, 2023.
- [7] S. Pulkkinen, D. Nerini, A. A. Pérez Hortal, C. Velasco-Forero, A. Seed, U. Germann, and L. Foresti, "Pysteps: An open-source python library for probabilistic precipitation nowcasting (v1. 0)," *Geoscientific Model Development*, vol. 12, no. 10, pp. 4185–4219, 2019.
- [8] H. Bi, M. Kyrliuk, Z. Wang, C. Meo, Y. Wang, R. Imhoff, R. Uijlenhoet, and J. Dauwels, "Nowcasting of extreme precipitation using deep generative models," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [10] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [11] J. Jing, Q. Li, X. Ding, N. Sun, R. Tang, and Y. Cai, "Aenn: A generative adversarial neural network for weather radar echo extrapolation," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 42, pp. 89–94, 2019.
- [12] J. Willard, X. Jia, S. Xu, M. Steinbach, and V. Kumar, "Integrating physics-based modeling with machine learning: A survey," *arXiv preprint arXiv:2003.04919*, vol. 1, no. 1, pp. 1–34, 2020.
- [13] K. Kashinath, M. Mustafa, A. Albert, J. Wu, C. Jiang, S. Esmailzadeh, K. Azizzadenesheli, R. Wang, A. Chattopadhyay, A. Singh *et al.*, "Physics-informed machine learning: case studies for weather and climate modelling," *Philosophical Transactions of the Royal Society A*, vol. 379, no. 2194, p. 20200093, 2021.
- [14] R. Sluiter, *Interpolation methods for the climate atlas*. KNMI De Bilt, The Netherlands, 2012. [Online]. Available: <https://cdn.knmi.nl/knmi/pdf/bibliotheek/knmi/TR335.pdf>
- [15] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers *et al.*, "The era5 global reanalysis," *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020.
- [16] B. S. Murphy, "Pykrige: development of a kriging toolkit for python," in *AGU fall meeting abstracts*, vol. 2014, 2014, pp. H51K–0753.
- [17] M. A. Oliver and R. Webster, "Kriging: a method of interpolation for geographical information systems," *International Journal of Geographical Information System*, vol. 4, no. 3, pp. 313–332, 1990.
- [18] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [19] A. Daw, M. Maruf, and A. Karpatne, "Pid-gan: A gan framework based on a physics-informed discriminator for uncertainty quantification with physics," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 237–247.
- [20] Z. Pu and E. Kalnay, "Numerical weather prediction basics: Models, numerical methods, and data assimilation," *Handbook of hydrometeorological ensemble forecasting*, pp. 67–97, 2019.
- [21] P. Esser, R. Rombach, and B. Ommer, "Taming transformers for high-resolution image synthesis," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 12 873–12 883.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [23] A. Van Den Oord, O. Vinyals *et al.*, "Neural discrete representation learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [24] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [25] R. V. Rohli and C. Li, *Meteorology for Coastal Scientists*. Springer Nature, 2021.
- [26] K. P. Georgakakos and R. L. Bras, "A hydrologically useful station precipitation model: 1. formulation," *Water Resources Research*, vol. 20, no. 11, pp. 1585–1596, 1984.
- [27] H. De Bruin, "From Penman to Makkink," in *Evaporation and Weather: Technical Meeting 44, Ede, The Netherlands 25 March 1987. The Hague, Netherlands. 1987. p 5-31. 1 fig, 4 tab, 34 ref.*, 1987.
- [28] W. Yan, D. Hafner, S. James, and P. Abbeel, "Temporally consistent transformers for video generation," *arXiv preprint arXiv:2210.02396*, 2022.
- [29] C. Meo, A. Roy, M. Lică, J. Yin, Z. B. Che, Y. Wang, R. Imhoff, R. Uijlenhoet, and J. Dauwels, "Extreme precipitation nowcasting using transformer-based generative models," *arXiv preprint arXiv:2403.03929*, 2024.