
Generative Modeling of High-resolution Global Precipitation Forecasts

James Duncan

UC Berkeley

Berkeley, CA 94720

jpduncan@berkeley.edu

Shashank Subramanian

Lawrence Berkeley National Laboratory

Berkeley, CA 94720

shashanksubramanian@lbl.gov

Peter Harrington

Lawrence Berkeley National Laboratory

Berkeley, CA 94720

pharrington@lbl.gov

Abstract

Forecasting global precipitation patterns and, in particular, extreme precipitation events is of critical importance to preparing for and adapting to climate change. Making accurate high-resolution precipitation forecasts using traditional physical models remains a major challenge in operational weather forecasting as they incur substantial computational costs and struggle to achieve sufficient forecast skill. Recently, deep-learning-based models have shown great promise in closing the gap with numerical weather prediction (NWP) models in terms of precipitation forecast skill, opening up exciting new avenues for precipitation modeling. However, it is challenging for these deep learning models to fully resolve the fine-scale structures of precipitation phenomena and adequately characterize the extremes of the long-tailed precipitation distribution. In this work, we present several improvements to the architecture and training process of a current state-of-the art deep learning precipitation model (FourCastNet) using a novel generative adversarial network (GAN) to better capture fine scales and extremes. Our improvements achieve superior performance in capturing the extreme percentiles of global precipitation, while comparable to state-of-the-art NWP models in terms of forecast skill at 1–2 day lead times. Together, these improvements set a new state-of-the-art in global precipitation forecasting.

1 Introduction

Precipitation is a fundamental climate phenomenon with major impact on crucial infrastructure such as food, water, energy, transportation, and health systems [1]. On timescales relevant to weather prediction, making accurate forecasts of precipitation and extreme events is critical to the planning and management of these systems, along with disaster preparedness for extreme precipitation events, which are greatly amplified by climate change [2]. Unfortunately, extreme precipitation remains one of the most challenging atmospheric phenomena to forecast accurately, due to high spatiotemporal variability and the myriad of complex multi-scale, multi-phase processes that govern its behavior [3–5].

State-of-the-art NWP models, such as the Integrated Forecast System¹ (IFS), produce operational forecasts by combining physics-based PDE solvers with assimilated observations from a variety

¹<https://www.ecmwf.int/en/forecasts/documentation-and-support>

of sources. The complexity of moisture physics requires many parameterizations in these models for processes like turbulent mixing, convection, subgrid clouds and microphysics [6], and such parameterizations can lead to large biases in NWP precipitation forecasts [7]. As a result, global precipitation forecasts generally achieve inadequate forecast skill [5] and, hence, there has been an increasing interest in fully data-driven solutions, primarily using deep learning, in recent years.

Data-driven models can be orders of magnitude faster with the potential to learn complex parameterizations between input and output function spaces directly from data, reducing model bias. With such models, major advances have been made in the area of precipitation “nowcasting”, where forecasts are made over limited spatial regions with lead times on the order of minutes to hours. Deep learning models trained directly on radar and/or satellite observations now outperform traditional methods for nowcasting [8–10]. However, until recently, there has been limited progress for models predicting precipitation at larger spatiotemporal scales (e.g., over the full globe up to days in advance), mainly due to computational limitations on resolution [11]. FourCastNet [12] is the first deep-learning-based global weather model running at $\sim 30\text{km}$ scale, which outperforms IFS in terms of precipitation forecast skill up to ~ 2 day lead times and is the current state-of-the-art. However, despite using a dedicated network just for precipitation due to its unique challenges, FourCastNet predictions still lack fine-scale details (see Figure 1) and thus underestimate the extreme percentiles of the precipitation distribution.

In this work, we aim to overcome some of these limitations using generative models to advance the state-of-the-art in deep learning-based precipitation forecasts. In particular, our contributions are as follows: (i) we apply a state-of-the-art generative adversarial network [13] that integrates multi-scale semantic structure and style information, allowing us to synthesize physically realistic fine-scale precipitation features; (ii) we show that capturing fine-scale phenomena leads to improved predictions of extreme precipitation while also preserving forecast skill, attaining comparable skill at 1–2 lead day times with respect to IFS.

2 Methods

2.1 Dataset

We replicate the data preparation pipeline of the original FourCastNet precipitation model [12], relying on the European Center for Medium-Range Weather Forecasting (ECMWF) global reanalysis dataset ERA5 [14], which combines archived atmospheric observations with physical model outputs. Following FourCastNet’s time steps of length 6 hours, we model the 6hr accumulated total precipitation TP (this also makes for easier comparisons against IFS, which archives TP forecasts in 6 hourly accumulations as well). A sample snapshot of TP , in log-normalized units for easy visualization of fine-scale details, is shown in the bottom-right inset of Figure 1. We use the years 1979-2015 and 2016-2017 as training and validation sets, respectively, and set aside 2018 as a test set. We refer the reader to the original FourCastNet paper [12] for further details on the dataset.

2.2 Model

Given the success of adversarial learning for high-resolution precipitation models in localized regions [8, 15, 16], we explore the utility of conditional generative adversarial networks (cGANs) for modeling TP over the entire globe using prognostic atmospheric variables as conditional input for operational diagnosis. In particular, we adopt the TSIT architecture [13] for our task, due to its success in varied image tasks and ability to flexibly condition intermediate features at a variety of scales in the generator. TSIT’s generator G employs symmetric downsampling and upsampling paths, fusing features from the downsampling path via feature-wise adaptive denormalization, and uses multi-scale patch-based discriminators [17] for adversarial training. We take a constant latitude embedding channel and the 20 atmospheric variables output by FourCastNet (corresponding to various different prognostic variables such as wind velocities, total column water vapor and temperature at different pressure levels) as input to the network, as shown in Figure 2. Randomness is injected at intermediate scales via elementwise additive noise, and we can thus generate “zero-shot” ensembles for probabilistic forecasting given a single input, which we explore in Appendix A. We refer the reader to [13] for additional details on the TSIT architecture, and list hyperparameters in Appendix A, together with the particulars of our setup. In this work, we focus on three model versions for demonstrating the impact of the adversarial training: (i) *FourCastNet*: the precipitation baseline model from [12], (ii) L_1 -only:

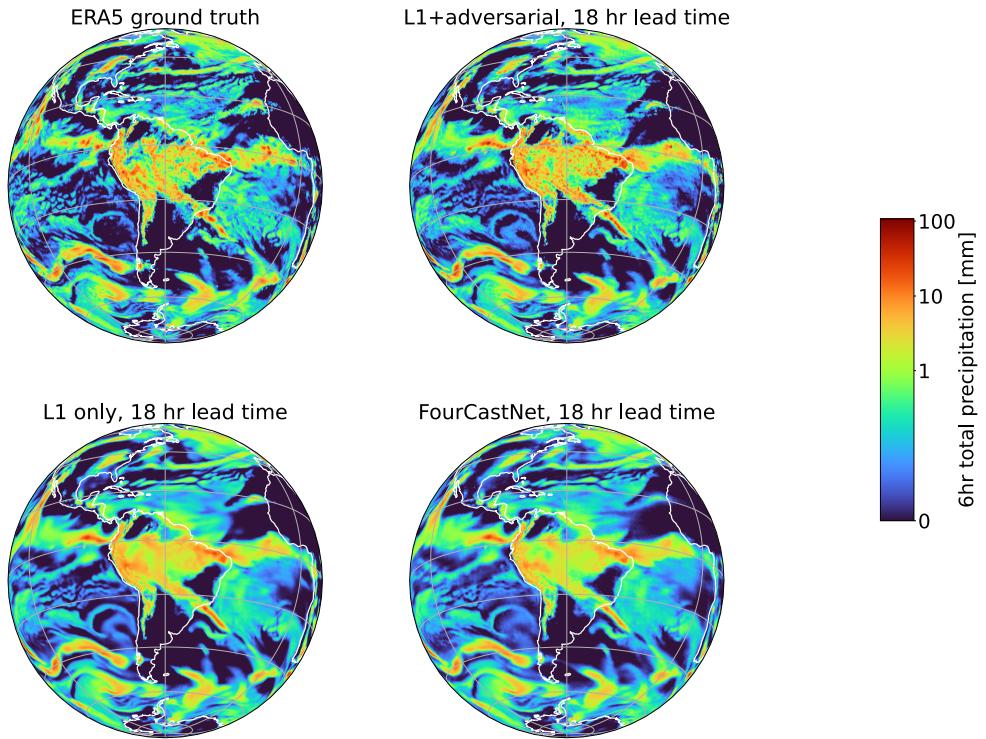


Figure 1: Visualization of precipitation forecasts at 18hr forecast lead time, comparing fine-scale features between FourCastNet, our \mathcal{L}_1 -only model, our adversarial+ \mathcal{L}_1 model, and the ground truth ERA5 over South America. We observe that the adversarially trained model shows finer-scale structures and matches the ground truth more accurately.

TSIT model with a simple (non-adversarial) \mathcal{L}_1 loss, and (iii) *adversarial+ \mathcal{L}_1* : TSIT model with adversarial training.

3 Results

We observe that, qualitatively, the adversarial training procedure clearly improves the perceptual realism and fine-scale detail represented in precipitation forecasts, as presented in the visualizations in Figure 1 which compares TP from the 2018 test set over South America at 18hr lead time. We confirm these results quantitatively by computing a 1D power spectrum of the TP forecasts along the East-West direction, again at 18hr lead time, averaged over several initial conditions in the test set². The results are plotted in Figure 3a, where we find that the adversarial learning framework effectively matches the the ground truth spectrum, specifically at higher wavenumbers (and hence finer spatial scales), outperforming all other models. Beyond assessing fine scales, the general task of identifying metrics for evaluating precipitation forecast quality is itself a challenge [4, 18]. For example, Anomaly Correlation Coefficient (ACC) is widely used to assess forecast skill, but it is biased towards larger scales, so smoothing functions can artificially inflate this metric (we illustrate this with simple gaussian blur in Appendix B). Furthermore, ACC is generally insensitive to extremes which is a key quantity to capture in fields such as precipitation. While our model is comparable to the IFS in terms of ACC on 1-2 day lead times (see Appendix B), we emphasize here the model’s ability to more faithfully capture the extreme TP values that are of interest to weather and climate stakeholders. We plot the extreme percentiles of our forecasts, again at 18hr lead time, in Figure 3b, binning logarithmically to emphasize the tail of the distribution [19] and averaging over the 2018 test set. Clearly, the model better replicates the long tail of the TP distribution, nearly matching the

²The initial conditions are spaced apart by 2 days as a rough estimate of the temporal decorrelation

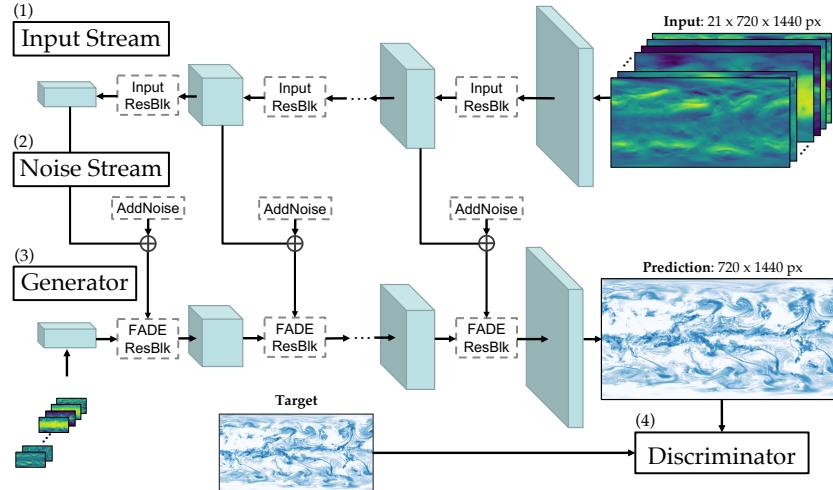


Figure 2: Training architecture of the *adversarial+ \mathcal{L}_1* model, a modified version of TSIT [13]. The model learns multi-scale representations of the input, perturbed by injected noise, and generates stochastic images with fine-scale detail promoted by a multi-scale discriminator.

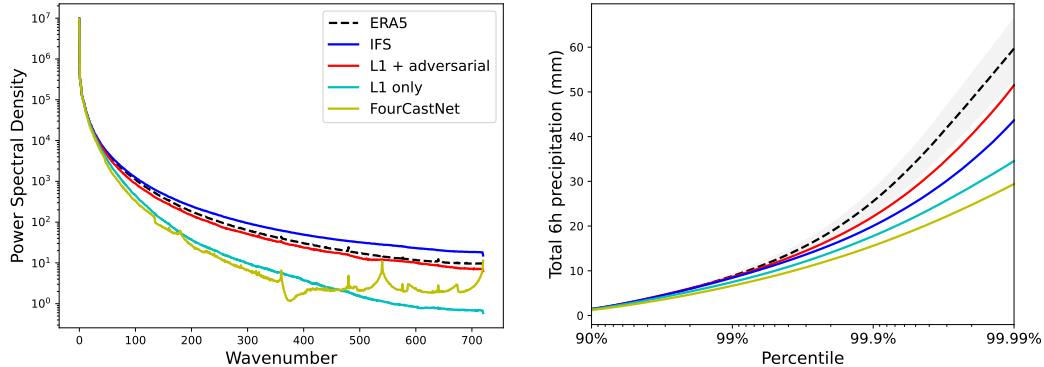


Figure 3: Analysis of precipitation power spectrum (left) and extreme percentiles (right) at 18hr lead time. Results are averaged over 178 initial conditions spread uniformly across 2018, and the gray shaded region shows the ERA5 variability in percentiles over initial conditions.

ground truth percentiles within the variability over initial conditions in the test set, and outperforms IFS in this regard as well. We show additional results on the TP distribution in Appendix B.

4 Conclusions

Deep-learning-based models like FourCastNet show great promise in data-driven forecasting of precipitation (TP), but exhibit smooth features and underestimate extreme events. In this work, we demonstrate the ability of GAN-based models to resolve small-scale details and synthesize physically realistic extreme values in the tail of the TP distribution, outperforming IFS (a leading NWP model) on both counts while retaining competitive forecast skill in terms of ACC. Our model sets a new state-of-the-art in global TP forecasting, and can be complementary to existing localized, short-timescale, high-resolution precipitation models which require initial or boundary conditions from NWP (or similar) models as input.

Acknowledgments and Disclosure of Funding

The authors would like to thank Jaideep Pathak and Karthik Kashinath for helpful discussions. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231

References

- [1] IPCC. *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK, 2022.
- [2] Lans P. Rothfusz, Russell Schneider, David Novak, Kimberly Klockow-McClain, Alan E. Gerard, Chris Karstens, Gregory J. Stumpf, and Travis M. Smith. FACETs: A Proposed Next-Generation Paradigm for High-Impact Weather Forecasting. *Bulletin of the American Meteorological Society*, 99(10):2025 – 2043, 2018.
- [3] Di Chen, Aiguo Dai, and Alex Hall. The Convective-To-Total Precipitation Ratio and the “Drizzling” Bias in Climate Models. *Journal of Geophysical Research: Atmospheres*, 126(16):e2020JD034198, 2021. e2020JD034198 2020JD034198.
- [4] Francisco J. Tapiador, Rémy Roca, Anthony Del Genio, Boris Dewitte, Walt Petersen, and Fuqing Zhang. Is Precipitation a Good Metric for Model Performance? *Bulletin of the American Meteorological Society*, 100(2):223 – 233, 2019.
- [5] Jun-Ichi Yano, Michał Z. Ziemiański, Mike Cullen, Piet Termonia, Jeanette Onvlee, Lisa Bengtsson, Alberto Carrassi, Richard Davy, Anna Deluca, Suzanne L. Gray, Víctor Homar, Martin Köhler, Simon Krichak, Silas Michaelides, Vaughan T. J. Phillips, Pedro M. M. Soares, and Andrzej A. Wyszogrodzki. Scientific Challenges of Convective-Scale Numerical Weather Prediction. *Bulletin of the American Meteorological Society*, 99(4):699 – 710, 2018.
- [6] Peter Bechtold, Richard Forbes, Irina Sandu, Simon Lang, and Maike Ahlgrimm. A major moist physics upgrade for the IFS. pages 24–32, 07 2020.
- [7] David A. Lavers, Shaun Harrigan, and Christel Prudhomme. Precipitation Biases in the ECMWF Integrated Forecasting System. *Journal of Hydrometeorology*, 22(5):1187 – 1198, 2021.
- [8] Suman Ravuri, Karel Lenc, Matthew Willson, Dmitry Kangin, Remi Lam, Piotr Mirowski, Megan Fitzsimons, Maria Athanassiadou, Sheleem Kashem, Sam Madge, et al. Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597(7878):672–677, 2021.
- [9] Sylwester Klocek, Haiyu Dong, Matthew Dixon, Panashe Kanengoni, Najeeb Kazmi, Pete Luferenko, Zhongjian Lv, Shikhar Sharma, Jonathan Weyn, and Siqi Xiang. MS-nowcasting: Operational Precipitation Nowcasting with Convolutional LSTMs at Microsoft Weather. In *NeurIPS 2021 Workshop on Tackling Climate Change with Machine Learning*, December 2021.
- [10] Lasse Espeholt, Shreya Agrawal, Casper Sønderby, Manoj Kumar, Jonathan Heek, Carla Bromberg, Cenk Gazen, Rob Carver, Marcin Andrychowicz, Jason Hickey, et al. Deep learning for twelve hour precipitation forecasts. *Nature Communications*, 13(1):1–10, 2022.
- [11] Stephan Rasp and Nils Thuerey. Data-Driven Medium-Range Weather Prediction With a Resnet Pretrained on Climate Simulations: A New Model for WeatherBench. *Journal of Advances in Modeling Earth Systems*, 13(2):e2020MS002405, 2021. e2020MS002405 2020MS002405.
- [12] Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators. *arXiv preprint arXiv:2202.11214*, 2022.
- [13] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. TSIT: A Simple and Versatile Framework for Image-to-Image Translation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 206–222, Cham, 2020. Springer International Publishing.

[14] Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla, Xavier Abellan, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot, Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger, Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux, Philippe Lopez, Cristina Lupu, Gabor Radnóti, Patricia de Rosnay, Iryna Rozum, Freja Vamborg, Sébastien Villaume, and Jean-Noël Thépaut. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.

[15] Jussi Leinonen, Daniele Nerini, and Alexis Berne. Stochastic super-resolution for downscaling time-evolving atmospheric fields with a generative adversarial network. *IEEE Transactions on Geoscience and Remote Sensing*, 59(9):7211–7223, 2020.

[16] Ilan Price and Stephan Rasp. Increasing the accuracy and resolution of precipitation forecasts using deep generative models. In *International Conference on Artificial Intelligence and Statistics*, pages 10555–10571. PMLR, 2022.

[17] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.

[18] L. Ruby Leung, William R. Boos, Jennifer L. Catto, Charlotte A. DeMott, Gill M. Martin, J. David Neelin, Travis A. O’Brien, Shaocheng Xie, Zhe Feng, Nicholas P. Klingaman, Yi-Hung Kuo, Robert W. Lee, Cristian Martinez-Villalobos, S. Vishnu, Matthew D. K. Priestley, Cheng Tao, and Yang Zhou. Exploratory Precipitation Metrics: Spatiotemporal Characteristics, Process-Oriented, and Phenomena-Based. *Journal of Climate*, 35(12):3659 – 3686, 2022.

[19] Benjamin Fildier, William D. Collins, and Caroline Muller. Distortions of the Rain Distribution With Warming, With and Without Self-Aggregation. *Journal of Advances in Modeling Earth Systems*, 13(2):e2020MS002256, 2021. e2020MS002256 2020MS002256.

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

[21] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved Techniques for Training GANs. *Advances in neural information processing systems*, 29, 2016.

[22] N. P. Klingaman, G. M. Martin, and A. Moise. ASoP (v1.0): a set of methods for analyzing scales of precipitation in general circulation models. *Geoscientific Model Development*, 10(1):57–83, 2017.

Appendices

A Additional training details and probabilistic inference

A.1 Hyperparameters & architecture details

The \mathcal{L}_1 -only and *adversarial*+ \mathcal{L}_1 models both use a modified version of the architecture described in [13]. We list hyperparameters in Table 1, and refer the reader to [13] for further details on the hyperparameters.

During training of the *adversarial*+ \mathcal{L}_1 model, a fixed latitude-embedding channel is appended to the 20 prognostic input variables used in [12] and passed to the input stream at full-resolution (720px \times 1440px). The entry point (layer $k = 0$) to the input stream is a standard convolutional residual layer [20] with 64 output channels (pictured in the upper right of Figure 2). Next, these intermediate features are downsampled to 512px \times 1024px using nearest neighbor interpolation and passed to the next residual layer of similar design, doubling the number of output channels to 128. Starting here, the output features are stored for multi-scale synthesis in the generator at a later stage, with increasing depth of field until the final layer ($k = 8$). The input stream uses downsampling at the beginning of each residual block, and the channel dimension is doubled after each residual block up to $d = 1024$ at layer ($k = 4$), resulting in a final representation with dimensions 4px \times 8px \times 1024 channels.

Table 1: Hyperparameters for *adversarial*+ \mathcal{L}_1 model, with any differences in the case of the \mathcal{L}_1 -only model listed between square brackets.

Hyperparameter	Value for adversarial training [\mathcal{L}_1 -only, if different]
Global batch size	64
Optimizer	Adam
Initial learning rate	2.5×10^{-4} [6.5×10^{-4}]
LR schedule num. epochs	15 constant + 6 linear decay [14 + 0]
β_1, β_2	0.85, 0.95
λ_{feat}	0.5 [N/A]
Number of upsampling blocks	8
Input stream normalization	Spectral instance
Generator normalization	FADE + spectral batch
Discriminator normalization	Spectral instance [N/A]
Multi-scale discriminators	4
Discriminator layers	6
Padding	Zeros

The generator mirrors the input stream, but travels in the opposite direction, from global to local representations. First, we downsample the 21 input fields to the same spatial dimensions as the final representation from the input stream. Then, a fully-connected layer takes in the downsampled input and returns the starting set of 1024 features (lower left of Figure 2), matching the full dimensions of the corresponding input stream features. These features then pass through the first of eight feature-adaptive denormalization (FADE) layers, integrating multi-scale semantic information from the input stream at the corresponding points in the generator.

To further promote diversity in the generator’s outputs, the noise stream draws random gaussian noise which it adaptively scales in a feature-specific, multi-scale manner. This noise is added to the input stream’s features prior to synthesis with the current output representation in each FADE layer. The generator’s intermediate features are then upsampled before passing on to the next layer. Once the generator reaches the penultimate resolution of 512px \times 1024px by 64 channels, it upsamples to the full resolution, applies Leaky ReLU with negative slope of 0.2, passes the features through a final convolutional layer for single-channel output, and applies the ReLU activation function to enforce precipitation non-negativity.

The multi-scale discriminator consists of four separate discriminators, each processing the same input features at a different spatial scale. One discriminator operates at full-resolution, while the others process inputs which are downsampled by progressive factors of 2 along the spatial dimensions. Each discriminator takes as input the real or predicted *TP* field, along with the input atmospheric variables fed to the generator (which get concatenated along the channel dimension) for improved training stability. Each discriminator’s output is a 2D grid of predictions classifying patches of the input as real or fake. This “patchGAN” approach helps the model focus more on local texture and fine-scale details [17].

In addition to the adversarial loss term using the discriminators’ outputs, a feature-matching loss [21] term is used to improve training stability for the generator. We refer the reader to [13, 17] for further specifics on the FADE and multi-scale patch-based discriminator architectures. The architecture of the \mathcal{L}_1 -only model resembles that of *adversarial*+ \mathcal{L}_1 as pictured in Figure 2, but with the noise stream and multi-scale discriminator components removed.

A.2 Probabilistic ensemble forecasts

Through the injection of adaptive randomness by the noise stream, outputs from the *adversarial*+ \mathcal{L}_1 model are non-deterministic. Hence, we can output a number of repeated predictions and form a probabilistic ensemble forecast for a single timepoint, as depicted in Figure 4, even if we use a single control forecast for the other atmospheric variables input to the generator.

An alternate method for attaining multi-modal predictions, first introduced in the TSIT paper [13], is a variational auto-encoder (VAE) approach in which the bottom level of the generator is a latent

space given by an auxiliary encoder and shaped by an additional KL-divergence loss. In preliminary tests we found that this was unable to attain sufficient variability in predictions, and instead opted for directly perturbing the generator synthesis pipeline with our additive noise stream. However, further investigation is needed to determine the optimal method for multi-modal predictions in this application.

We present an example ensemble forecast in Figure 4, plotting mean and standard deviation per gridbox over $n = 100$ ensemble members. Qualitatively, we observe highest variance over challenging locations like coastlines and mountainous regions, which aligns with our expectations as such areas exhibit high spatiotemporal variability and sharp, fine-scale features for TP . In future work we hope to quantitatively assess the quality of such ensemble forecasts, as rapid probabilistic forecasting from such a model would be a major step forward.

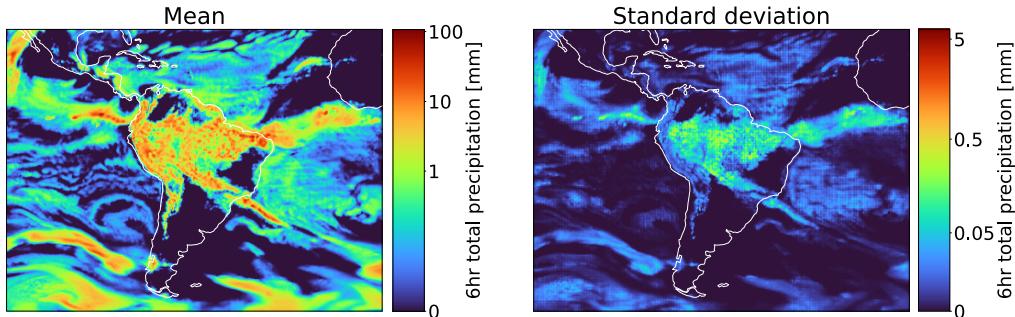


Figure 4: Through the introduction of randomness in the form of added noise between the input stream and generator, probabilistic forecasts are created by repeated inference runs with the same inputs, leading to an ensemble of unique outputs. We can then use the mean of this set of stochastic outputs as a final ensemble prediction for a given timepoint (left panel), and analyze the ensemble’s standard error field to quantify forecast uncertainty at particular geographic locations (right panel).

B Additional metrics and analysis

In Figure 5, we plot the fractional contribution [22, 18] of binned precipitation rates at 18hr lead time to the total precipitation in each gridbox over all initial conditions in 2018. This assesses how much of the total precipitation comes from light versus heavy precipitation, and we find that our GAN model achieves the best agreement with the ground truth in general. Notably, the IFS model exhibits more drizzle bias [3] for rates between 0.1-10 mm/6hrs, though we find that both IFS and our model overestimate the central mode of the fractional distribution at ~ 5 mm/6hrs.

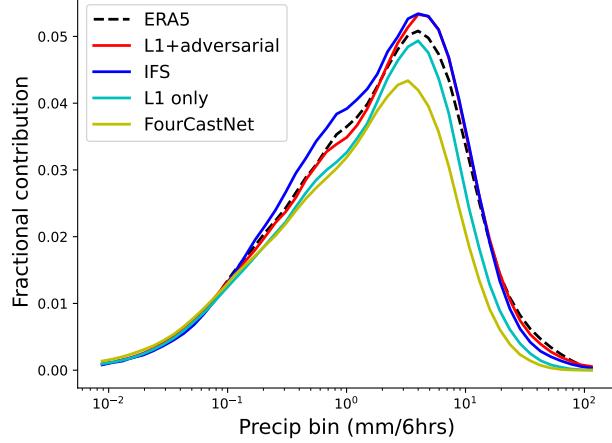


Figure 5: Fractional contributions of precipitation rates at 18hr lead time compared over the models in question and the ERA5 ground truth. Our model more faithfully captures the ground truth in general, but overestimates the central mode, similar to the IFS (which also shows more drizzle bias)

Finally, in Figure 6, we demonstrate tradeoffs between global metrics like ACC and quantifying extreme events. To do so, we conduct a simple experiment—we add gaussian blurring to the total precipitation outputs after inference and re-compute two metrics: ACC and extreme percentiles. While smoothing predictably leads to worse performance in capturing extremes, we observe that it does lead to higher ACC values. This could possibly be because ACC is a global metric and smoothing can help correct highly localized errors—this can inflate the quality of the forecast if we only focus on global behaviour. Hence, identifying the right metrics to characterize precipitations is very important. We leave evaluating our forecasts with a comprehensive set of diagnostic tools that focus on different aspects of modeling precipitation to future work.

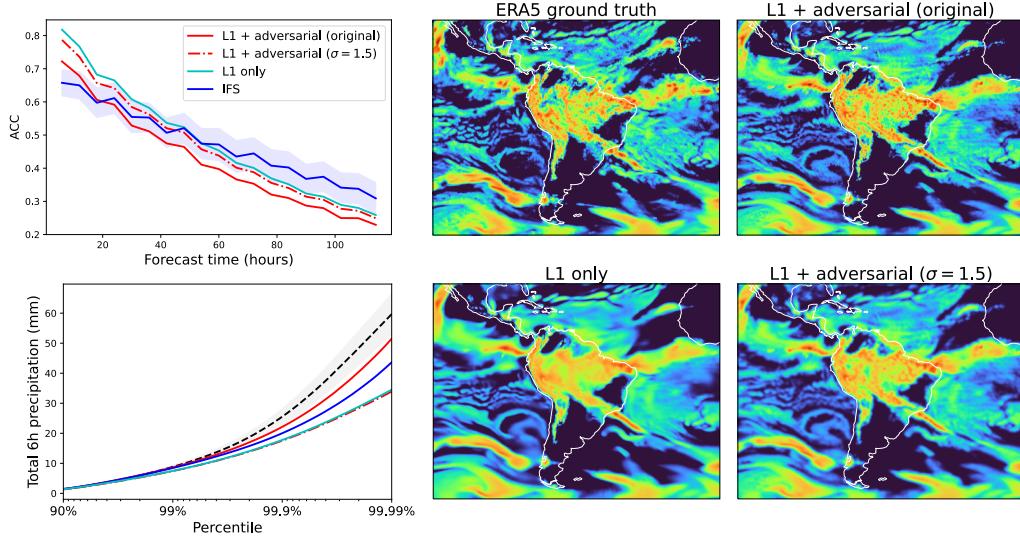


Figure 6: Applying a small amount of gaussian blurring ($\sigma = 1.5$) to the outputs from the adversarial+ \mathcal{L}_1 model leads to increased ACC over the entire forecast sequence as compared to the unaltered outputs (upper left panel). At the same time, the ability to capture extremes is degraded, as shown in the lower left panel. ACC and percentiles are averaged over 178 initial conditions, as in Figure 3, with variability over ACC additionally included for IFS.