

Downscaling Precipitation with Bias-informed Conditional Diffusion Model

Ran Lyu
Virginia Tech
ran22@vt.edu

Linhan Wang
Virginia Tech
linhan@vt.edu

Yanshen Sun
Virginia Tech
yansh93@vt.edu

Hedanqiu Bai
Texas A&M University
baisy@tamu.edu

Chang-Tien Lu
Virginia Tech
clu@vt.edu

Abstract—Climate change is intensifying rainfall extremes, making high-resolution precipitation projections crucial for society to better prepare for impacts such as flooding. However, current Global Climate Models (GCMs) operate at spatial resolutions too coarse for localized analyses. To address this limitation, deep learning-based statistical downscaling methods offer promising solutions, providing high-resolution precipitation projections with a moderate computational cost. In this work, we introduce a bias-informed conditional diffusion model for statistical downscaling of precipitation. Specifically, our model leverages a conditional diffusion approach to learn distribution priors from large-scale, high-resolution precipitation datasets. The long-tail distribution of precipitation poses a unique challenge for training diffusion models; to address this, we apply gamma correction during preprocessing. Additionally, to correct biases in the downscaled results, we employ a guided-sampling strategy to enhance bias correction. Our experiments demonstrate that the proposed model achieves highly accurate results in an $8\times$ downscaling setting, outperforming previous deterministic methods. The code and dataset are available at Github.

Index Terms—Deep learning, Denoising Diffusion Probability Models, Statistical Downscaling, Climate Modeling

I. INTRODUCTION

Due to climate change, there is a growing demand for reliable weather and climate simulation at local scales [10]. The current de facto technique Global Climate Models (GCMs) [8] can simulate the Earth’s response to varying atmospheric greenhouse gas (GHG) emissions scenarios. However, the outputs from GCMs are often coarse and lack the granularity [7] needed to understand local weather patterns and their impacts on specific sectors such as agriculture, water resources, and food security.

One of the foundational methods is statistical downscaling, which enhances coarse climate model outputs by predicting fine-resolution data based on statistical relationships between low-resolution and high-resolution observations. This approach shares similarities with image super-resolution techniques in computer vision. A historically significant image super-resolution method is the Super-Resolution Convolutional Neural Network (SRCNN). Although SRCNN has been widely adopted and dominated the field for nearly a decade, it suffers from limitations such as a restricted receptive field and insufficient network depth, which impede its ability to capture global context and effectively extract features [13]. Additionally, SRCNN struggles with poor adaptation to

long-tail data distributions [5], such as precipitation, making it unsuitable for tasks like $8\times$ downscaling.

Recently, denoised diffusion models [9], [16] have emerged as the dominant deep generative models for images due to their comprehensive coverage of data distributions and high-quality outputs. Improved Denoising Diffusion Probabilistic Models (DDPMs) [14] are generative models that produce high-quality samples, achieve competitive log-likelihoods with simple modifications, enable faster sampling through learned variances, and scale effectively with model capacity and compute. Furthermore, diffusion models outperform state-of-the-art generative models like GANs [6] in image synthesis tasks by leveraging improved architectures and classifier guidance for enhanced fidelity and diversity, achieving superior FID scores across multiple resolutions while maintaining efficient sampling and better distribution coverage [3].

In this work, we aim to adapt conditional diffusion models for the task of precipitation downscaling. However, precipitation downscaling has its unique challenges. **1)** The distribution of precipitation is inherently non-normal, posing difficulties for traditional deep learning models to capture its long-tail characteristics. **2)** Bias is a critical metric for precipitation downscaling, as small inaccuracies can significantly impact downstream applications. To tackle these issues, we introduce a Bias-aware Guided Sampling (BGS) approach to systematically reduce bias during the downscaling process, improving the overall accuracy and reliability of the generated high-resolution precipitation data.

II. METHOD

In this section, we first introduce the definition and statistical downscaling. Then we explain how to train a conditioned denoised diffusion model for statistical downscaling. Finally, we introduce the motivation and formulation of our two core innovations: Gamma Correction and Bias-aware Guided Sampling.

A. Problem Setting and Conditional Diffusion Models

Statistical downscaling involves generating high-resolution climate variables from low-resolution counterparts using statistical methods. This process closely resembles the image super-resolution task in computer vision. [15].

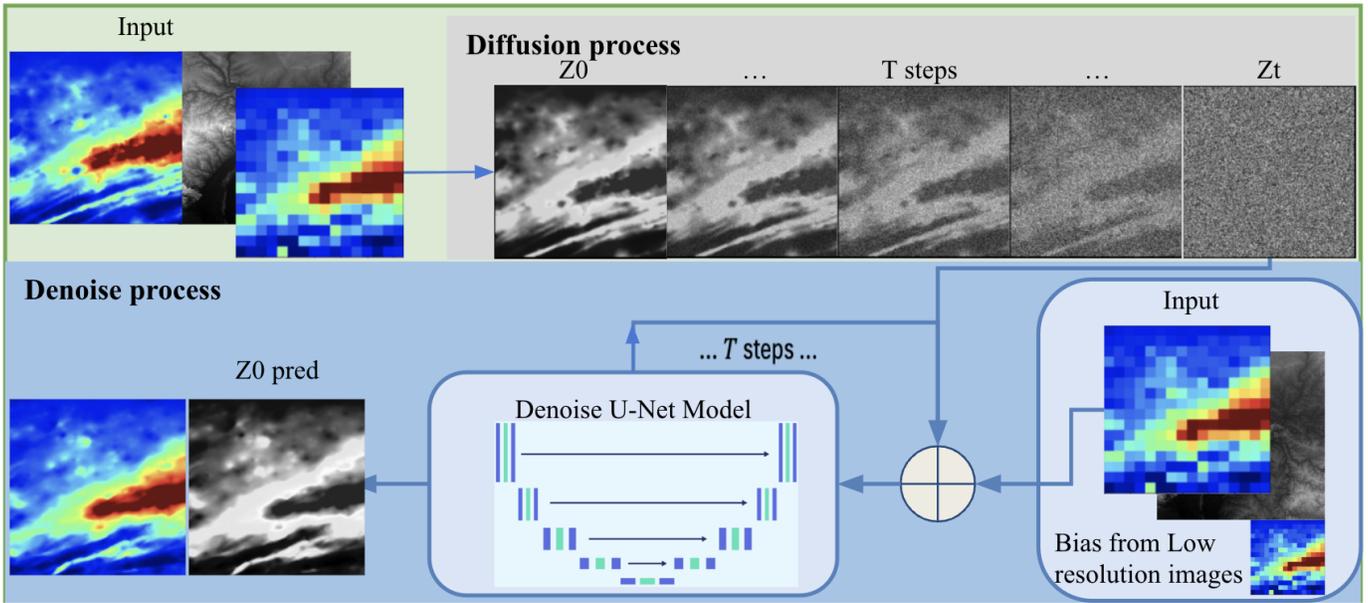


Fig. 1. Overall pipeline of the proposed framework: In the training phase, a noise-corrupted high-resolution image, a low-resolution image, and a topography image are concatenated and input into a U-Net model. The U-Net is trained to predict the noise in the corrupted high-resolution precipitation image. In the prediction phase, the trained U-Net iteratively denoises precipitation images, transforming pure noise into high-resolution precipitation. Our bias-informed sampling strategy quantifies and reduces the bias between the corrupted high-resolution precipitation and the low-resolution input at each denoising step.

We are given a dataset of low-high resolution precipitation pairs, denoted $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$, which represent samples drawn from an unknown conditional distribution $p(\mathbf{y}|\mathbf{x})$. Here \mathbf{y} represents high-resolution precipitation and \mathbf{x} represents low-resolution precipitation. Due to the high correlation between topography information and precipitation [2], we concatenate it with \mathbf{x} as the input for our model.

The conditional diffusion model generates a high resolution precipitation \mathbf{y}_0 in T refinement steps. Starting with an initial noise precipitation $\mathbf{y}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the model progressively refines the precipitation through successive iterations $(\mathbf{y}_{T-1}, \mathbf{y}_{T-2}, \dots, \mathbf{y}_0)$ using learned conditional transition distributions $p_\theta(\mathbf{y}_{t-1}|\mathbf{y}_t, \mathbf{x})$ such that $\mathbf{y}_0 \sim p(\mathbf{y}|\mathbf{x})$.

B. Gamma correction

Distribution of precipitation naturally deviates from Gaussian distribution, which is favored by deep learning models. Inspired by the similarity between distributions of precipitation and low-light images [11], we apply Gamma Correction on both low and high resolution precipitation data in preprocessing stage. The equation of Gamma Correction is

$$\hat{a} = a^\gamma, \gamma = 0.15 \quad (1)$$

This approach addresses the issue of large precipitation areas dominating the optimization process, thereby effectively enhancing the model's performance.

C. Bias-aware Guided Sampling(BGS)

Classifier-guided sampling is a technique used in [3]. This approach can improve the performance of diffusion model

without slowing down the training process. We propose to incorporate this technique in our sampling process.

While the diffusion process conditions on low-resolution image \mathbf{x} , we assume that the information of \mathbf{x} is not fully leveraged due to indirect connection between \mathbf{x} and \mathbf{y} . Moreover, the \mathbf{x} comes from GCMs, which are expert models with minimum bias from real climate variables. To fully leverage this prior knowledge, we propose a Bias-aware Guided Sampling approach [3]. Specifically, we use a L2 norm $f = \|\mathbf{y}_t - \mathbf{x}\|_2$ to quantify the bias between \mathbf{y}_t and \mathbf{x} and use its gradients to steer the generation process at each diffusion step. Formally,

$$\mathbf{y}_{t-1} = \frac{1}{\sqrt{\alpha_t}}(\mathbf{y}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{y}_t, \mathbf{x}, t)) - w \nabla f \quad (2)$$

where $w = 100$ is a hyperparameter we searched by experiments.

III. EXPERIMENT AND RESULTS

A. Dataset

The high-resolution precipitation data(4km) is from PRISM Climate data [1]. The low-resolution precipitation data(32km) is bilinearly downsampled 8x from the high-resolution data. We use the data between 2000-2018 for training and 2019-2022 for evaluation.

B. Evaluation metrics

Following DeepSD [17], we evaluate our method and baselines using three metrics: RMSE, correlation, and bias. RMSE measures the average error, correlation assesses the linear relationship and bias quantifies systematic errors. Both

TABLE I
EVALUATION OF DOWNSCALING METHODS ON PRISM DATASET

Methods	RMSE(mm/day)	Corr	Bias(mm/day)
Interpolation	3.134	0.939	-0.0053
SRCNN	5.921	0.828	0.0660
Ours w/o BGS, Topo	3.019	3.943	0.0544
Ours w/o BGS	<u>2.991</u>	0.944	-0.0548
Ours	2.972	0.945	-0.0389

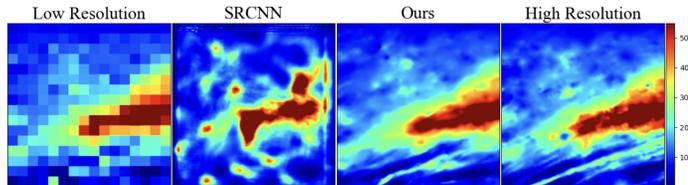


Fig. 2. One example of $16 \times 16 \rightarrow 128 \times 128$ downscaling of precipitation.

RMSE and bias are in mm/day. The interpolation method uses bilinear interpolation to upscale the original low-resolution image to high resolution, serving as our baseline. SRCNN performs poorly in the evaluation metrics, with a significantly increased RMSE and noticeably low correlation, highlighting its limitations in effectively downscaling precipitation data.

C. Implementation details

Our code is based on improved diffusion [14]. We treat both the high and low resolution precipitations as 128×128 grayscale images. All the precipitations are cropped from the same region for convenience. We use AdamW [12] and learning rate $3e-4$. We train our model using one A100 for 1000 epochs, which takes around one day.

D. Quantitative and qualitative results

The quantitative results are presented in Tab I. As shown, our method outperforms the deterministic baseline SRCNN [4], [17], and both the incorporation of topography information (Topo) and Bias-aware Guided Sampling (BGS) prove effective in this setting.

As shown in Fig 2, while SRCNN struggles to generate high-resolution precipitation under an $8 \times$ super-resolution setting. The distribution and visual performance of the SRCNN image significantly deviate from the high-resolution ground truth, while our method generates physically and visually plausible results with fine details.

IV. CONCLUSION

In this work, we proposed a bias-informed conditional diffusion model for precipitation downscaling, addressing the challenges posed by the long-tail distribution of precipitation and biases in low-resolution input data. By incorporating gamma correction during preprocess and introducing Bias-aware Guided Sampling (BGS), our method achieved significant improvements in accuracy and bias correction compared to baseline methods. The experimental results demonstrated the effectiveness of our approach in producing high-resolution precipitation maps with fine details, making it a promising solution for localized climate impact

studies. Future work will explore extending this framework to other climate variables and integrating physical constraints for enhanced generalizability.

REFERENCES

- [1] Christopher Daly, Michael Halbleib, Joseph I Smith, Wayne P Gibson, Matthew K Doggett, George H Taylor, Jan Curtis, and Phillip P Pasteris. Physiographically sensitive mapping of climatological temperature and precipitation across the conterminous united states. *International Journal of Climatology: a Journal of the Royal Meteorological Society*, 28(15):2031–2064, 2008.
- [2] Christopher Daly, Ronald P Neilson, and Donald L Phillips. A statistical-topographic model for mapping climatological precipitation over mountainous terrain. *Journal of Applied Meteorology and Climatology*, 33(2):140–158, 1994.
- [3] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [4] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pages 184–199. Springer, 2014.
- [5] Vitaly Feldman and Chiyuan Zhang. What neural networks memorize and why: Discovering the long tail via influence estimation. *Advances in Neural Information Processing Systems*, 33:2881–2891, 2020.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [7] William J Gutowski, Paul Aaron Ullrich, Alex Hall, L Ruby Leung, Travis Allen O’Brien, Christina M Patricola, RW Arritt, MS Bukovsky, Katherine V Calvin, Zhe Feng, et al. The ongoing need for high-resolution regional climate models: Process understanding and stakeholder information. *Bulletin of the American Meteorological Society*, 101(5):E664–E683, 2020.
- [8] Nicholas G Heavens, Daniel S Ward, and MM Natalie. Studying and projecting climate change with earth system models. *Nature Education Knowledge*, 4(5):4, 2013.
- [9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [10] Bipin Kumar, Kaustubh Atey, Bhupendra Bahadur Singh, Rajib Chattopadhyay, Nachiketa Acharya, Manmeet Singh, Ravi S Nanjundiah, and Suryachandra A Rao. On the modern deep learning approaches for precipitation downscaling. *Earth Science Informatics*, 16(2):1459–1472, 2023.
- [11] Xiangsheng Li, Manlu Liu, and Qiang Ling. Pixel-wise gamma correction mapping for low-light image enhancement. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(2):681–694, 2023.
- [12] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [13] Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard Zemel. Understanding the effective receptive field in deep convolutional neural networks. *Advances in neural information processing systems*, 29, 2016.
- [14] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021.
- [15] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4713–4726, 2022.
- [16] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [17] Thomas Vandal, Evan Kodra, Sangram Ganguly, Andrew Michaelis, Ramakrishna Nemani, and Auroop R Ganguly. DeepSD: Generating high resolution climate change projections through single image super-resolution. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*, pages 1663–1672, 2017.