

PrecipDiff: Leveraging Image Diffusion Models to Enhance Satellite-Based Precipitation Observations

Ting-Yu Dai^{1,2}, Hayato Ushijima-Mwesigwa²

¹ The University of Texas at Austin

² Fujitsu Research of America, Inc.

funnyengineer@utexas.edu, hayato@fujitsu.com

Abstract

A recent report from the World Meteorological Organization (WMO) highlights that water-related disasters have caused the highest human losses among natural disasters over the past 50 years, with over 91% of deaths occurring in low-income countries. This disparity is largely due to the lack of adequate ground monitoring stations, such as weather surveillance radars (WSR), which are expensive to install. For example, while the US and Europe combined possess over 600 WSRs, Africa, despite having almost one and half times their landmass, has fewer than 40. To address this issue, satellite-based observations offer a global, near-real-time monitoring solution. However, they face several challenges like accuracy, bias, and low spatial resolution. This study leverages the power of diffusion models and residual learning to address these limitations in a unified framework. We introduce the first diffusion model for correcting the inconsistency between different precipitation products. Our method demonstrates the effectiveness in downscaling satellite precipitation estimates from 10 km to 1 km resolution. Extensive experiments conducted in the Seattle region demonstrate significant improvements in accuracy, bias reduction, and spatial detail. Importantly, our approach achieves these results using only precipitation data, showcasing the potential of a purely computer vision-based approach for enhancing satellite precipitation products and paving the way for further advancements in this domain.

Introduction

Water-related disasters, such as landslides, floods, and droughts, constitute a significant majority of natural disasters. The past five decades have witnessed over 11,000 reported weather-related disasters globally, tragically claiming over two million lives and inflicting \$3.64 trillion in economic losses (WMO 2021). Developing countries disproportionately bear this burden, accounting for over 91% of these fatalities. However, advancements in early warning systems and disaster management have yielded a nearly threefold reduction in fatalities between 1970 and 2019, underscoring the life-saving potential of technological progress.

Precipitation monitoring forms a cornerstone of effective early warning systems, proving crucial for forecasting and mitigating weather-related hazards. Weather surveil-

lance radar (WSR) systems offer high-resolution data on precipitation intensity and distribution, invaluable for short-term forecasting and nowcasting. However, the prohibitive cost of WSRs, often reaching millions of dollars per unit, leads to uneven global distribution. For example, the U.S. and Europe, with a combined population of roughly 1 billion, possess nearly 700 WSRs. In stark contrast, Africa, with a larger population and a landmass more than 1.5 times, operates fewer than 40 WSRs (Cambridge 2023).

Satellite-based precipitation products (SPPs) offer a valuable alternative, providing continuous spatiotemporal estimation with global coverage. Products such as the Tropical Rainfall Measuring Mission (TRMM), the Climate Prediction Center Morphing Technique (CMORPH) (Joyce et al. 2004), the Global Precipitation Measurement (GPM) mission (Hou et al. 2014), and Global Satellite Mapping of Precipitation (GSMaP) (Kubota et al. 2020), along with advancements like the Integrated Multi-satellite Retrievals for GPM (IMERG), have significantly improved precipitation monitoring capabilities.

However, SPPs still face challenges. These products often exhibit biases compared to ground-based observations due to various factors, including limitations in sensor technology, retrieval algorithms, cloud cover, and the inherent complexities of precipitation processes. Their indirect estimation nature, coupled with limitations in spatiotemporal sampling, instrument capabilities, and retrieval algorithms, can lead to biases and relatively low accuracy. Comprehensive evaluation and bias correction are therefore crucial, especially for early warning systems (Yin et al. 2008; Ji and Chen 2012; Yang and Luo 2014; Lu et al. 2018a,b). Furthermore, the coarse spatial resolution of current SPPs (typically 0.1° or lower) limits their applicability in urban settings where finer resolutions are required (Berne et al. 2004).

This work addresses these two key challenges of SPPs: bias correction and spatial resolution enhancement through downscaling. The main contributions of this paper are:

1. We formulate both bias correction and downscaling as a residual learning problem and propose using diffusion models to solve them in a unified framework.
2. We introduce the first computer vision-based bias correction algorithm, which uses the diffusion model to learn the biases of SPPs.

3. We demonstrate the practical application of diffusion models for operational datasets. Moving a step beyond previous related work, that primarily focused on simulated data, we showcase their effectiveness in real-time downscaling tasks using state-of-the-art deep learning technique
4. Extensive experiments demonstrate the effectiveness of our proposed method in both bias correction and downscaling tasks, even without incorporating other related variables such as elevation, temperature, humidity, etc, which is the common practice for tackling these problems.

Background

Bias Correction in SPPs

Accurate bias correction is essential for enhancing the reliability of SPPs and ensure their effective application in diverse domains, including hydrological modeling, climate monitoring, and disaster preparedness.

Numerous methods have been proposed for correcting biases in SPPs, broadly categorized into statistical, machine learning, and hybrid approaches. Traditional statistical techniques encompass linear regression (Wilks 2011), quantile mapping (Scheuerer and Hense 2014), multiplicative bias correction (Wilks 2011), and power transformation (Wilks 2011). Machine learning methods, on the other hand, offer the flexibility to capture nonlinear relationships between SPP estimates and rain gauge data, potentially leading to more accurate corrections. Examples of machine learning approaches include Support Vector Machines, random forests, and deep learning techniques (Chen et al. 2019). Despite significant progress in bias correction methodologies, it remains an active area of research within the field of SPP development and application.

Challenges in Precipitation Downscaling

Precipitation downscaling presents unique challenges compared to downscaling other meteorological variables like temperature or wind. Its highly variable nature in both space and time, influenced by factors such as orography, land-sea contrasts, and convective processes, makes capturing fine-scale precipitation patterns difficult, particularly in regions with complex terrain or limited observational networks.

Furthermore, extreme precipitation events, crucial for flood risk assessments and water resource management, are often poorly represented in downscaled datasets. This is primarily due to their rarity and the inherent limitations of both dynamical and statistical downscaling methods in capturing these infrequent high-intensity events (Hewitson et al. 2014). While temperature and other variables typically exhibit smoother spatial patterns amenable to downscaling, precipitation often involves abrupt changes over short distances, further complicating the modeling process. These challenges underscore the need for more sophisticated downscaling techniques that can effectively account for the complexities of precipitation processes and provide more reliable, high-resolution estimates for climate impact assessments.

Diffusion Models

Diffusion models (Ho, Jain, and Abbeel 2020; Sohl-Dickstein et al. 2015) have emerged as a powerful class of generative models, demonstrating remarkable success in image synthesis, surpassing even the quality of GANs in certain applications. Their applicability extends to both unconditional (Ho, Jain, and Abbeel 2020; Nichol and Dhariwal 2021; Song et al. 2020) and conditional (Ho et al. 2022; Nichol et al. 2021; Nichol and Dhariwal 2021; Preechakul et al. 2022; Ramesh et al. 2022; Rombach et al. 2022) generation settings. Beyond image synthesis, diffusion models have rapidly found utility in diverse domains, including audio and video generation, image segmentation, language translation, and recently, climate downscaling.

A denoising diffusion probabilistic model (DDPM) () employs two Markov chains: a forward process that gradually injects noise into the data and a reverse process that learns to progressively remove noise to recover the original data distribution. Formally, given a data distribution $x_0 \sim q(x_0)$, the forward Markov process generates a sequence of random variables x_1, \dots, x_T with transition kernel $q(x_t|x_{t-1})$. Leveraging the Markov property, the joint distribution is given by:

$$q(x_1, \dots, x_T) = \prod_{t=1}^T q(x_t|x_{t-1}). \quad (1)$$

In DDPMs, the transition kernel is commonly defined as:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I), \quad (2)$$

where $\beta_t \in (0, 1)$ is a hyperparameter controlling the variance schedule. This forward process gradually adds noise to the original data until it becomes essentially pure noise.

To generate new data samples, a noise vector is first sampled from the prior distribution. The reverse process, parameterized by a prior distribution $p(x_T) = \mathcal{N}(x_T; \mathbf{0}, \mathbf{I})$ and a learnable transition kernel $p_\theta(x_{t-1}|x_t)$, then gradually removes the noise. This kernel is defined as:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (3)$$

where θ represents the model parameters, and the mean $\mu_\theta(x_t, t)$ and variance $\Sigma_\theta(x_t, t)$ are parameterized by deep neural networks.

Methodology

Residual-Based Diffusion Model Training

Inspired by Corrdiff (Mardani et al. 2024), we shift our focus from directly generating high-resolution (HR) precipitation data to learning the residual difference between low-resolution (LR) and HR data. Direct HR generation using diffusion models often results in a loss of intensity and detail in cloud formations and weather fronts. Instead, we train our model to generate the residual, conditioned on the LR data, using the EDM scheduler and sampler. This residual is then added to the interpolated LR image to obtain the downscaled precipitation data. A classic U-Net architecture is employed for denoising the noisy residual based on the EDM

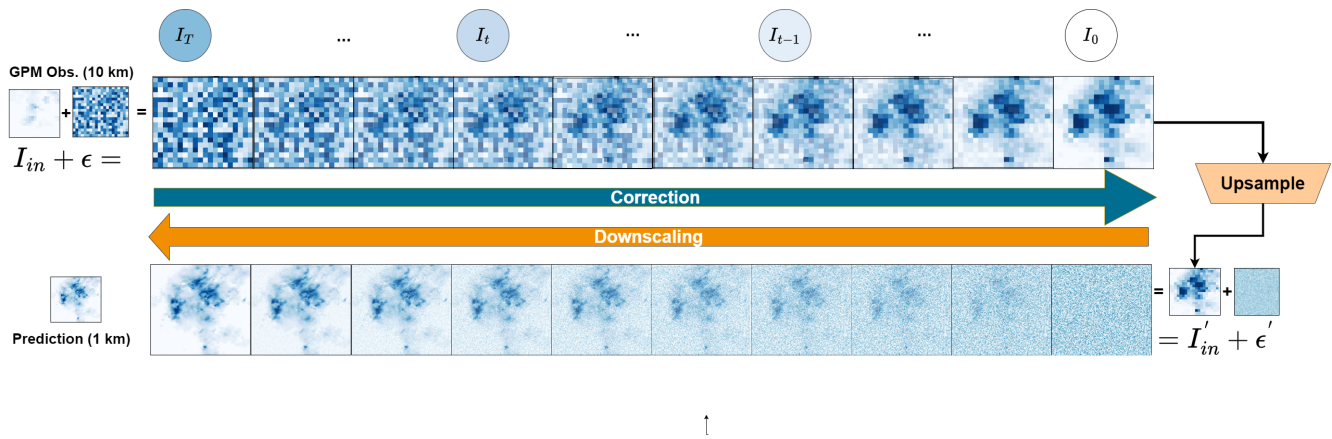


Figure 1: Overview flowchart for inference process. Satellite precipitation data are corrected by the first diffusion process and then upsampled using a linear method. Second diffusion process are applied to synthesized downscaled predictions to create high-resolution data.

noise scheduling algorithm. A key assumption in this residual training approach is the availability of high-quality LR data as a foundation.

Training and Inference Pipeline

Figures 1 and 2 illustrate the separate processes of training and inference, respectively. Our framework addresses two primary tasks: bias correction and downscaling.

Bias Correction This task aims to mitigate discrepancies between satellite-based GPM and radar observations, which arise from differences in measurement techniques and processing algorithms. As shown in Figure 2, the correction model is trained on residuals calculated by subtracting IMERG data from LR MRMS data. IMERG data is used as conditioning information to guide the denoising of noisy residuals.

Downscaling For downscaling, MRMS data is first downsampled to the same resolution as satellite observations, creating an LR representation. Residuals are then calculated by comparing the original MRMS data with the downsampled LR MRMS data, representing the information lost during downsampling. The same residual diffusion learning process is applied to this task, with the coarsened radar observations serving as conditioning information.

Figure 1 depicts the operational workflow. First, satellite GPM observations are calibrated using the trained bias correction diffusion model. The corrected data is then upsampled to 1 km resolution using linear interpolation and fed as conditioning information to the downscaling diffusion model. The model generates the residual, which is added to the corrected and upsampled data to produce the final high-resolution precipitation prediction. This pipeline enables the generation of corrected and downsampled satellite precipitation data in real-time, with accuracy comparable to radar products and at the same spatial resolution.

Bias Correction Model

The diffusion model for bias correction learns to synthesize the distribution difference between LR satellite estimates and coarsened HR radar data. This addresses the inherent bias between satellite and radar observations due to their different measurement principles and locations. The model effectively corrects this systematic bias, aligning satellite observations more closely with radar observations. To achieve this, radar observations are first coarsened to the same resolution as satellite data using maximum value interpolation to preserve precipitation intensity. Residual training is then applied using the coarsened radar data and the corresponding satellite precipitation data.

Downscaling Model

After bias correction, the same residual training strategy is employed for downscaling. The diffusion model is trained on the residual between coarsened radar precipitation and the original high-resolution radar data, aiming to capture fine-scale details present in the radar data. Coarsened radar observations serve as conditioning information, guiding the model to generate the denoised residual. Importantly, the downscaling model operates independently of the bias correction stage, as the corrected satellite observations may still exhibit differences compared to LR radar images. By training solely on radar data, we minimize uncertainties stemming from input data discrepancies, enabling the model to effectively enhance the spatial resolution of corrected satellite precipitation estimates.

Experiments

Data

To closely mimic real-world scenarios, we utilize two operational precipitation datasets: the IMERG dataset as our LR satellite observations and the MRMS dataset as our HR radar observations.

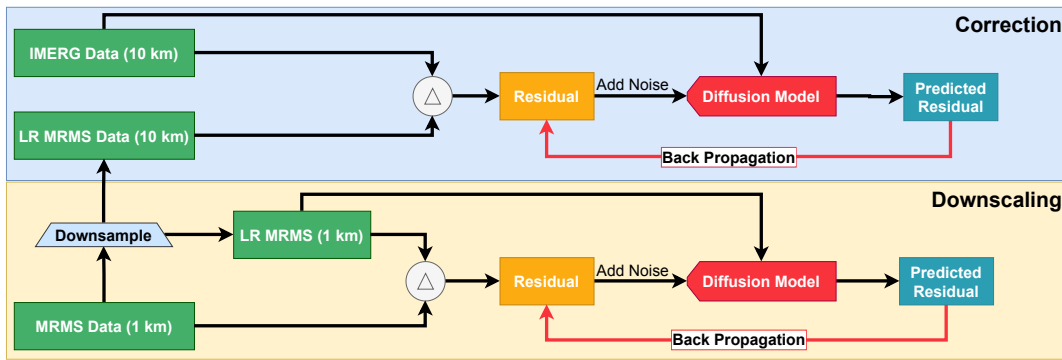


Figure 2: Overview flowchart for training correction and downscaling diffusion models. The correction model is trained at 10 km resolution using noisy residuals as input and IMERG data as conditioning information. Residuals at this stage are calculated by subtracting IMERG data from LR MRMS data. The downscaling model is trained at 1 km resolution, with residuals computed by subtracting interpolated LR MRMS data from original MRMS data.

IMERG The Integrated Multi-satellitE Retrievals for the Global Precipitation Measurement (GPM) mission (IMERG) is a product of the U.S. GPM Science Team, providing globally covered precipitation estimates with approximately 10 km resolution based on the GPM satellite constellation. IMERG offers three product qualities: early, late, and final run, reflecting data latency and quality. This study utilizes the final run product from version 07.

MRMS The Multi-Radar Multi-Sensor (MRMS) system is a sophisticated operational hydrometeorological framework. It integrates observations and information from various radar networks, satellites, and surface observational systems to generate precipitation estimates at 1 km resolution. MRMS offers various temporal frequencies (daily to 2-minute) and data quality levels, including products fused with ground rain gauge data. To align with the temporal scale of IMERG and leverage the most comprehensive ground information, we utilize the hourly Q3 Multi-sensor Pass 2 precipitation product.

Experiment Setup

Our study focuses on the Seattle region, characterized by frequent rain events. We use data from both IMERG and MRMS for Year 2022 and 2023. To focus on rain events, we filter out data samples with a low proportion of rain, specifically those where the proportion of zero precipitation values within the selected region is below 0.5. This results in approximately 1600 data samples. We then randomly split the dataset into training (90%) and testing (10%) sets. To evaluate the effectiveness of our diffusion model approach for both correction and downscaling tasks independently, we design three experiments: correction task, downscaling task, and the unified correction and downscaling task.

Unified correction and downscaling task This experiment simulates a real-world application scenario. First, we correct the IMERG data using the trained correction model. Then, we use this corrected IMERG data as input to the trained downscaling model. Finally, we compare the resulting downscaled IMERG product with the original HR

MRMS data to evaluate the overall effectiveness of our framework.

Correction task The target is the residual between the original IMERG data and the coarsened (LR) MRMS data. The input (conditioning image) is the original IMERG data.

Downscaling task The target is the residual between the LR MRMS data and the original (HR) MRMS data. The input (conditioning image) is the LR MRMS data.

Model Configurations and Training The correction model is trained on 20x20 pixel image patches using a 3-block U-Net with latent channels of (32, 64, 128) while the downscaling model is trained on 184x200 pixel images (corresponding to the 1 km resolution target) using a 4-block U-Net with latent channels of (128, 256, 256, 512). Both models are trained for 2000 epochs with a learning rate of $2e-4$ and a batch size of 16, distributed across two Quadro RTX 8000 GPUs. Total training time is approximately 10 hours.

Baseline Model We compare our correction model with a supervised Swin2SR (Conde et al. 2022) model to assess the effectiveness of our approach against a state-of-the-art supervised method. The Swin2SR model is trained with the same configuration and a comparable number of trainable parameters as the correction diffusion model, using original IMERG data as input and LR MRMS data as the target.

Inference During inference, we employ the second-order EDM stochastic sampler proposed by (Karras et al. 2022) to sample the residual. The correction model first generates a residual, which is added to the original IMERG data to produce the corrected IMERG product. This corrected product then serves as input to the downscaling model, which generates a second residual. Adding this residual to the corrected IMERG data yields the final downscaled IMERG product. Both models utilize 25 steps to solve the reverse SDE process during inference.




Figure 3: A comparison of the error distribution between original IMERG and Corrected IMERG. The errors are calculated based on the LR MRMS observations, with all pixels flattened for the comparison.

Method	RMSE	CRPS	CC	SSIM
Bicubic LR	3.7087	1.9914	0.4593	0.5222
Swin2SR	2.2842	1.8493	0.5685	0.3418
PrecipDiff (Ours)	1.7969	1.2464	0.4961	0.3772

Table 1: Evaluation for correction task: **IMERG to LR MRMS**. RMSE, CRPS, Correlation coefficient, and SSIM assess the overall performance and image synthesis quality between LR MRMS data. The denoted bold values indicate the best performance.

Results & Discussions

We evaluate our models using four metrics:

- Root Mean Squared Error (RMSE): Measures overall performance and the magnitude of errors.
- Continuous Ranked Probability Score (CRPS): Equivalent to Mean Absolute Error (MAE) for deterministic predictions, assessing overall accuracy. Comparing RMSE and CRPS helps determine if the model primarily improves extreme values (lower RMSE) or general pixel values.
- Pearson Correlation Coefficient (CC): Evaluates the linear relationship between predictions and ground truth, assessing the model’s ability to capture overall precipitation patterns.
- Structural Similarity Index Measure (SSIM): Measures the similarity of image structures, assessing the model’s ability to reproduce spatial details.

Bias Correction

The goal of bias correction is to calibrate satellite observations to a scale comparable with ground-based radar data. Directly applying a diffusion model from IMERG to MRMS proved challenging due to consistent biases between these operational datasets. Therefore, we focused on predicting LR MRMS data from IMERG data. Initially, we experimented with a Swin2SR model (Mardani et al. 2024) for bias correction. While this supervised approach yielded promising metrics, the predictions were overly smooth, resulting in the loss of distinct storm shapes. To address this, we




Figure 4: A comparison of the error distributions between LR MRMS and downsampled MRMS. The errors are calculated based on the original MRMS observations, with sampling from the pixels flattened for the comparison.

Method	RMSE	CRPS	CC	SSIM
LinearRegression	0.8642	0.5999	0.9347	0.7378
PrecipDiff (Ours)	0.4543	0.2685	0.9386	0.8249

Table 2: Evaluation for downscaling task: **LR MRMS to MRMS**. RMSE, CRPS, Correlation coefficient, and SSIM are used to evaluate the overall performance and image synthesis quality between original MRMS measurements. The denoted bold values indicate the best performance.

employed our proposed residual diffusion model for bias correction. As shown in Table 1, our model outperforms both the Swin2SR model and bicubic linear interpolation in RMSE and CRPS. One possible explanation for the weaker performance relative to Swin2SR in CC and Bicubic LR in SSIM is that our approach incorporates the sampled residual into the initial Low-Res images. This integration might disturb the overall image structure, causing to lower scores in these two metrics that assess overall similarity.

Figure 3 compares the error distributions of the original and corrected IMERG data relative to LR MRMS data. IMERG tends to overestimate lower precipitation values and underestimate higher values. Our corrected predictions effectively mitigate the overestimation issue and reduce variation in higher values, although improvements in underestimation are limited. We also experimented with correcting IMERG data at high resolution by interpolating IMERG and using the residuals between the interpolated data and HR MRMS for training. However, this approach yielded suboptimal results due to the abundance of zero-value pixels in HR MRMS data, which led the model to produce unrealistically low precipitation values.

Downscaling

To evaluate the downscaling performance independently of distribution shifts, we conducted an experiment to downscale LR MRMS to HR MRMS. This experiment aimed to assess the model’s ability to add fine-scale details to LR data. Table 2 demonstrates significant improvements across all metrics for our diffusion model compared to linear regression.

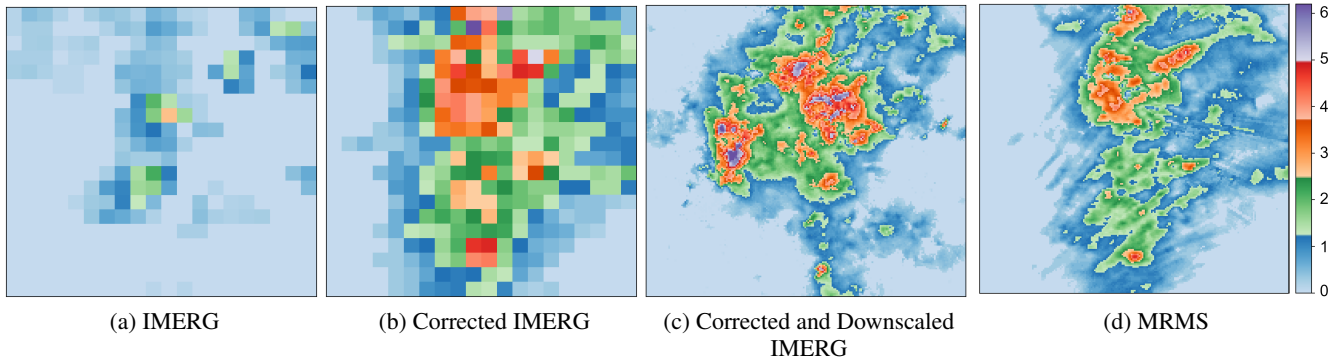


Figure 5: Illustration of a unified test for enhancing the satellite precipitations. The IMERG is calibrated by the corrector resulting in (b) corrected IMERG, and (c) is sampled by conditioning on the corrected IMERG. A comparison is made between (c) corrected and downscaled IMERG and (d) MRMS data.

Data Comapred with MRMS	RMSE	CRPS	CC	SSIM
IMERG	3.3410	1.7264	0.5417	0.6344
Corrected IMERG	1.3597	0.9701	0.6615	0.5736
Corrected and Dowscaled IMERG	1.0846	0.7432	0.6426	0.5766

Table 3: Evaluation for Unified task: **IMERG to MRMS**. RMSE, CRPS, CC, and SSIM are used to evaluate the overall performance and synthesized image quality between original MRMS data across all stages.




Figure 6: Unified correction and downscaling: The error distribution comparison between IMERG, corrected IMERG and downscaled IMERG observations.

Figure 4 compares the error distributions of LR MRMS and downscaled MRMS relative to HR MRMS. To simulate real-world scenarios where precise ground truth values are unavailable at every location, we randomly sampled nearby pixels within the LR pixel area. The results show that our downscaled predictions exhibit errors tightly centered around zero, indicating high accuracy, while LR MRMS shows a wider error distribution.

Unified Correction and Downscaling

To evaluate the combined performance of our framework, we applied both the correction and downscaling models sequentially to downscale IMERG data to the resolution of MRMS observations. Figure 5 illustrates the results, showcasing the detailed cloud-like formations and precipitation

patterns captured by our method. Notably, the correction model plays a crucial role in providing a more accurate foundation for downscaling.

Figure 6 compares the error distributions of IMERG, corrected IMERG, and downscaled IMERG relative to MRMS data. Corrected IMERG significantly reduces bias and variation across all precipitation intensities. Downscaled IMERG performs particularly well at lower precipitation values, effectively refining details and capturing subtle variations. However, the increased number of zero-value pixels in the downscaled product leads to some underestimation at higher precipitation intensities.

Table 3 highlights that the downscaled IMERG product exhibits the best performance to MRMS data in terms of RMSE and CRPS, indicating its suitability for both visual interpretation and numerical weather modeling applications. For correlation-focused metrics i.e. CC and SSIM, since the our method is to tweak the LR image with the sampled residual, the overall correlation is lower after the correction task, which is a potential area for further exploration.

Out-of-Distribution Evaluation

To assess the generalizability of our models, we applied them to data from three additional cities in the United States: New York, Portland, and San Jose (Figures 7 and 8). The downscaling model effectively reproduced fine-scale details in these new regions, indicating good generalization capabilities. However, the correction model’s performance was less consistent, likely due to regional variations in precipitation characteristics and the influence of local rain gauge data assimilation in the IMERG and MRMS products.

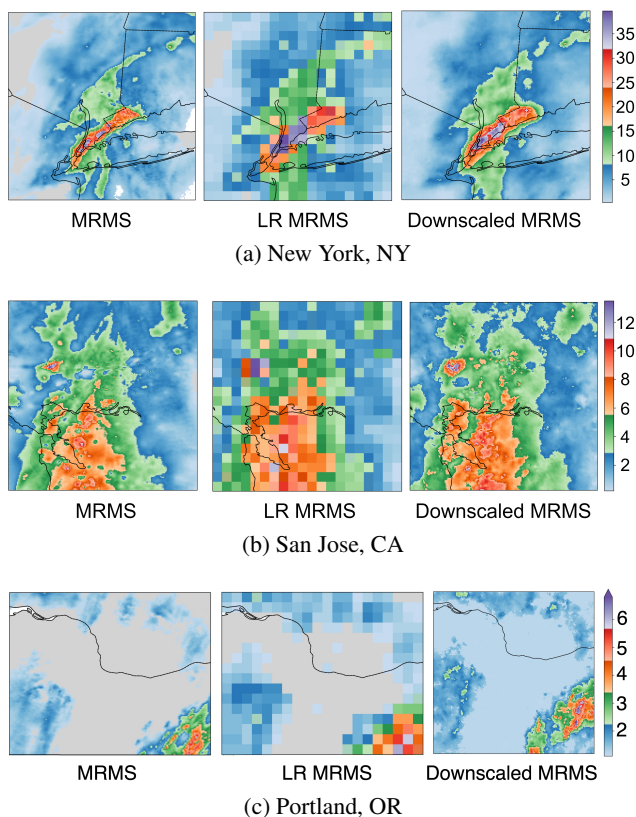


Figure 7: Downscaling with Seattle-data trained model on other regions

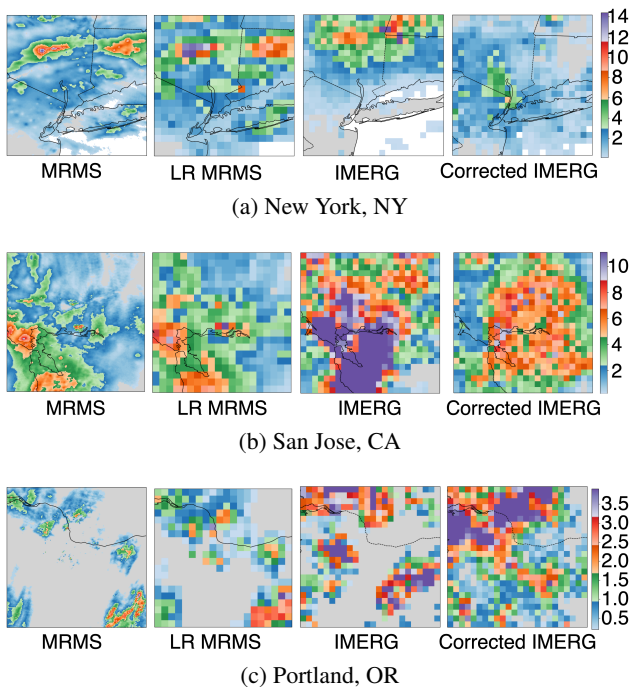


Figure 8: Correction using a model trained with Seattle data for other regions

Modeling Precipitation with ONLY Precipitation

One of the primary objectives of this work is to explore the strengths of purely computer vision methods for enhancing precipitation data. Precipitation is a complex variable to work with. This is partly due to the fact that it is sparsely spatially distributed compared to other meteorological variables, like temperature and wind direction. In particular, precipitation is a result of meteorological interactions. Due to such features, it is compelling to see if neural networks can learn the underlying mechanisms. We realize that directly predicting precipitation can be ineffective due to its sparsity. Models tend to underestimate actual values and struggle to preserve the original data due to its stochastic nature. This limitation is particularly evident in our experiments with the correction model which works in low resolution and demonstrates success due to the reduced sparsity.

It is also worth noting that the inherent uncertainty in precipitation distributions poses a challenge for traditional synthesis methods that rely on Gaussian distribution to capture the skewness. To overcome this, we propose learning the residuals between the target and the input, which is based on our finding that the residuals exhibit less skewed distributions compared to the original precipitation values. Consequently, the diffusion model can concentrate on generating fine-scale precipitation details rather than reconstructing the entire sample. However, it is important to note that this approach assumes the availability of a LR data that accurately represents the large-scale statistical properties corresponding to the HR data. In particular, this highlights the critical role of the correction model, which proved to be more influential than originally anticipated, especially when working with operational datasets like IMERG and MRMS. With these datasets, we encountered challenges derived from observational biases to different instruments and various data processing algorithms. These factors introduce regional variations in data distributions, even though both datasets aim to estimate the same variable. Future work will prioritize addressing these distribution shifts more comprehensively.

Conclusion

This study introduces PrecipDiff, a novel framework leveraging diffusion models and residual learning to address the critical need for high-resolution precipitation data, particularly in regions with limited ground-based monitoring. PrecipDiff effectively tackles two key challenges: correcting biases in satellite precipitation products (SPPs) and downscaling SPP estimates to finer resolutions. Our approach demonstrates substantial improvements in bias reduction and spatial detail, as independent tasks and as a unified framework.

A key innovation of PrecipDiff is its reliance solely on precipitation data, demonstrating the potential of purely data-driven computer vision techniques for enhancing SPPs. This has significant implications for data-scarce regions heavily reliant on satellite observations for applications like disaster preparedness and water resource management. To our knowledge, this is the first study to apply diffusion models exclusively to operational precipitation datasets and for bias correction of SPPs.

References

- Berne, A.; Delrieu, G.; Creutin, J.-D.; and Obled, C. 2004. Temporal and spatial resolution of rainfall measurements required for urban hydrology. *Journal of Hydrology*, 299(3-4): 166–179.
- Cambridge. 2023. Death tolls from climate disasters will ‘balloon’ without investment in Africa’s weather stations. <https://www.cam.ac.uk/research/news/death-tolls-from-climate-disasters-will-balloon-without-investment-in-africas-weather-stations>. Accessed: 2024-08-15.
- Chen, Y.; Li, J.; Wang, X.; and Zhang, J. 2019. Deep learning for precipitation nowcasting: A review. *Atmospheric Research*, 225: 134–153.
- Conde, M. V.; Choi, U.-J.; Burchi, M.; and Timofte, R. 2022. Swin2sr: Swinv2 transformer for compressed image super-resolution and restoration. In *European Conference on Computer Vision*, 669–687. Springer.
- Hewitson, B.; Daron, J.; Crane, R.; Zermoglio, M.; and Jack, C. 2014. Interrogating empirical-statistical downscaling. *Climatic change*, 122: 539–554.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; Saharia, C.; Chan, W.; Fleet, D. J.; Norouzi, M.; and Salimans, T. 2022. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47): 1–33.
- Hou, A. Y.; Kakar, R. K.; Neeck, S.; Azarbarzin, A. A.; Kummerow, C. D.; Kojima, M.; Oki, R.; Nakamura, K.; and Iguchi, T. 2014. The global precipitation measurement mission. *Bulletin of the American meteorological Society*, 95(5): 701–722.
- Ji, X.; and Chen, Y. 2012. Characterizing spatial patterns of precipitation based on corrected TRMM 3 B 43 data over the mid Tianshan Mountains of China. *Journal of Mountain Science*, 9: 628–645.
- Joyce, R. J.; Janowiak, J. E.; Arkin, P. A.; and Xie, P. 2004. CMORPH: A method that produces global precipitation estimates from passive microwave and infrared data at high spatial and temporal resolution. *Journal of hydrometeorology*, 5(3): 487–503.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35: 26565–26577.
- Kubota, T.; Aonashi, K.; Ushio, T.; Shige, S.; Takayabu, Y. N.; Kachi, M.; Arai, Y.; Tashima, T.; Masaki, T.; Kawamoto, N.; et al. 2020. Global Satellite Mapping of Precipitation (GSMaP) products in the GPM era. *Satellite Precipitation Measurement: Volume 1*, 355–373.
- Lu, X.; Tang, G.; Wei, M.; Yang, L.; and Zhang, Y. 2018a. Evaluation of multi-satellite precipitation products in Xinjiang, China. *International journal of remote sensing*, 39(21): 7437–7462.
- Lu, X.; Wei, M.; Tang, G.; and Zhang, Y. 2018b. Evaluation and correction of the TRMM 3B43V7 and GPM 3IMERGM satellite precipitation products by use of ground-based data over Xinjiang, China. *Environmental earth sciences*, 77: 1–18.
- Mardani, M.; Brenowitz, N.; Cohen, Y.; Pathak, J.; Chen, C.-Y.; Liu, C.-C.; Vahdat, A.; Kashinath, K.; Kautz, J.; and Pritchard, M. 2024. Residual Diffusion Modeling for Km-scale Atmospheric Downscaling.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.
- Preechakul, K.; Chatthee, N.; Wizadwongsa, S.; and Suwajanakorn, S. 2022. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10619–10629.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Scheuerer, M.; and Hense, A. 2014. Multivariate quantile mapping for downscaling regional climate model simulations. *Journal of Climate*, 27(1): 153–169.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Wilks, D. S. 2011. *Statistical methods in the atmospheric sciences*. Academic Press.
- WMO, W. M. O. 2021. Weather-related disasters increase over past 50 years, causing more damage but fewer deaths. Accessed: 2024-08-15.
- Yang, Y.; and Luo, Y. 2014. Using the back propagation neural network approach to bias correct TMPA data in the arid region of Northwest China. *Journal of Hydrometeorology*, 15(1): 459–473.
- Yin, Z.-Y.; Zhang, X.; Liu, X.; Colella, M.; and Chen, X. 2008. An assessment of the biases of satellite rainfall estimates over the Tibetan Plateau and correction methods based on topographic analysis. *Journal of Hydrometeorology*, 9(3): 301–326.