PIANO: Physics-informed Dual Neural Operator for Precipitation Nowcasting

Seokhyun Chin
California Institute of Technology
Pasadena, CA, 91125
schin@caltech.edu

Junghwan Park
TelePIX
Seoul, South Korea, 07331
brian897743@gmail.com

Woojin Cho
TelePIX
Seoul, South Korea, 07331
woojin.py@gmail.com

Abstract

Precipitation nowcasting, key for early warning of disasters, currently relies on computationally expensive and restrictive methods that limit access to many countries. To overcome this challenge, we propose precipitation nowcasting using satellite imagery with physics constraints for improved accuracy and physical consistency. We use a novel physics-informed dual neural operator (PIANO) structure to enforce the fundamental equation of advection-diffusion during training to predict satellite imagery using a PINN loss. Then, we use a generative model to convert satellite images to radar images, which are used for precipitation nowcasting. Compared to baseline models, our proposed model shows a notable improvement in moderate (4mm/h) precipitation event prediction alongside short-term heavy (8mm/h) precipitation event prediction. It also demonstrates low seasonal variability in predictions, indicating robustness for generalization. This study suggests the potential of the PIANO and serves as a good baseline for physics-informed precipitation nowcasting.

1 Introduction

As global warming intensifies, extreme precipitation events are becoming more common, leading to significant issues such as severe flooding, soil erosion, landslides, reduction of agricultural productivity, and increased health risks [Ombadi et al., 2023, Tabari, 2020]. To mitigate the impact of such problems, accurate precipitation nowcasting is essential.

Traditional precipitation prediction relies heavily on Numerical Weather Prediction (NWP) models, but these suffer from heavy computational requirements and low spatial resolution, which hinders the prediction of heavily localized weather events. To combat this issue, many data-driven weather prediction methods have emerged. Models like Pangu-Weather [Bi et al., 2023], FourCastNet [Kurth et al., 2023], and GraphCast [Lam et al., 2023] have shown comparable performance to traditional NWP methods in data-driven approaches. However, these weather forecasting models rely heavily on computationally expensive ERA5 [Hersbach et al., 2020] reanalysis data for predictions, thus still require supercomputers for training. Many teams have explored radar-image-based precipitation nowcasting methods [Andrychowicz et al., 2023, Ravuri et al., 2021, Zhang et al., 2023] as a remedy, but these models apply only to regions with radar coverage, limiting their applicability to developing countries and other areas that lack coverage.

Precipitation nowcasting using satellite imagery emerges as a viable alternative, as satellite images have near-global coverage and are easily accessible. Indeed, a variety of studies (e.g., [Lebedev et al., 2019, Park et al., 2025, Gruca et al., 2022]) have already explored this possibility. However, these studies fail to take into account the foundational physics equations of atmospheric dynamics: an idea that has proven to improve predictive capabilities of AI-based weather prediction [Kochkov et al.,

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Machine Learning and the Physical Sciences.

2024, Verma et al., 2024]. To this end, we propose a proof-of-concept PIANO (Physics-Informed duAl Neural Operator) for precipitation nowcasting based on satellite imagery that aims to produce physically consistent precipitation forecasts while maintaining overall accuracy.

2 Proposed Methods: Physics-informed Dual Neural Operator

This section details the proposed PIANO, a novel architecture designed for satellite-based precipitation nowcasting. The core of PIANO lies in its unique dual structure, which decouples the task of data-driven temporal forecasting from the inference of latent physical dynamics. This separation allows for the explicit integration of physical principles, specifically the advection-diffusion equation, as a powerful inductive bias to regularize the learning process.

2.1 Model Architecture

As shown in the overall pipeline (Figure 1), the two operators run in series for the application of physics-informed neural networks (PINN) loss [Raissi et al., 2019, Dabrowski et al., 2023, Cho et al., 2024] and data driven loss. The Time-stepping Neural Operator (T-NO) is a data-driven operator that learns to advance the state forward in time, while the Velocity-extraction Neural Operator (V-NO) is a physics-informed operator that learns the underlying flow field responsible for evolving the satellite imagery. By decoupling these tasks into separate modules, we can enforce physical consistency via V-NO without constraining the expressiveness of T-NO.

T-NO takes input satellite images $u_{t-s+1:t}$: $\{u_{t-s+1}, \cdots, u_{t-1}, u_t\}$ and ψ where each $u \in \mathbb{R}^{C\times H\times W}$, represents a satellite image and $\psi \in \mathbb{R}^{H\times W}$ represents the digital elevation map. It outputs frames $\hat{u}_{t+1:t+s}$, which the V-NO takes as input to output $\hat{v}_{t:t+s}$. Both $\hat{u}_{t+1:t+s}$ and $\hat{v}_{t+1:t+s}$ are used to enforce a data-driven and PINN loss in a fashion similar to the physics-informed neural operator and the physics-informed deep neural operator [Goswami et al., 2023, Li et al., 2024]. The loss function is described further in the section below.

Figure 1: **Proposed method: PIANO.** The framework consists of two distinct neural operators: V-NO, which learns the mapping for velocity fields, and T-NO, which learns the mapping for time-stepping dynamics.

We then use a Pix2Pix [Isola et al., 2017] ResNet [He et al., 2016] trained using a Generative Adversarial Network [Goodfellow et al., 2014] structure with a ResNet generator and a ResNet-based discriminator network to transform predicted satellite images to radar images that can be used for precipitation nowcasting. The generative structure is motivated by the choice of [Park et al., 2025], whose work we aim to improve with this study.

2.2 Model Training

We use a two-step training strategy where the V-NO and T-NO are pre-trained parallelly and combined for further fine tuning. In our study, T-NO takes eight time steps as input and output predictions of the next eight time steps, and V-NO takes eight time steps and outputs eight velocity fields.

Governing Equation: Advection-Diffusion Equation (Eq. (1)**)** Because we aim to model precipitation, we assume that the satellite images will evolve according to the governing advection-diffusion equation, which can be modeled as [Wesseling, 2001]:

$$\frac{\partial u}{\partial t} = \nabla \cdot (D\nabla u - \vec{v}u) + R, \tag{1}$$

where $\partial u/\partial t$ represents the local rate of change of the quantity u (in our case, the input of each layer in the satellite image), $\vec{v}$ is the velocity vector, D is the diffusion coefficient and R is the source term. This equation can be discretized in time using finite differences, which allows us to achieve the following equation:

$$u_{t+1} = u_t + D\nabla^2 u_t - \vec{v}_t \cdot (\nabla u_t) - u_t(\nabla \cdot \vec{v_t}) + R. \tag{2}$$

We utilize finite differencing due to the algorithmic efficiency and simplicity [Zhu et al., 2019]. In implementation, the values $u_{0:8}, D$, and R are all scalar matrices of size $H \times W$, and the gradient and Laplacian of each value can be calculated using a convolution matrix [Woods, 2012]. Under this implementation, the value R represents the input of the quantity within each region and the diffusivity constant D represents the diffusion between two regions. The beginning constants for the two values are 0 for R and 1 for D, respectively.

We train V-NO to learn the velocity $\vec{\hat{v}}_{0:8}, D, R$ from the satellite images $u_{0:8}$ using the PINN loss only, which is written as:

$$\mathcal{L}_{PDE} = \sum_t ||u_{t+1} - u_t - D\nabla^2 u_t + \vec{\hat{v}}_t \cdot (\nabla u_t) + u_t(\nabla \cdot \vec{\hat{v}}_t) - R||^2 \tag{3}$$

The T-NO does not explicitly encode physics; it learns the spatio-temporal patterns of satellite imagery evolution purely from data. To train T-NO, we use a standard supervised reconstruction loss comparing the predicted frames to the ground-truth future frames. Let $\hat{u}_t$ denote the T-NO prediction for time t and u_t the corresponding ground truth. We minimize the mean squared error (MSE) over all predicted time steps in the training sequences, $\mathcal{L}_{data} = \sum_t ||\hat{u}_t - u_t||^2$.

Finally, the V-NO and T-NO are combined and are trained with both the MSE and physics-informed loss functions, which can be written as:

$$\mathcal{L}_{total} = \sum_t ||\hat{u}_t - u_t||^2 + \alpha||\hat{u}_{t+1} - \hat{u}_t - D\nabla^2 \hat{u}_t + \vec{\hat{v}}_t \cdot (\nabla \hat{u}_t) + \hat{u}_t(\nabla \cdot \vec{\hat{v}}_t) - R||^2. \tag{4}$$

2.3 Dataset

We use the Sat2RDR [Park et al., 2025] dataset, which provides GK2A satellite infrared (10.5 µm), upper- (6.3 µm) and lower-level (7.3 µm) water vapor, plus DEM data of South Korea. We subsample January–August 2020–2024, excluding September–December due to missing 0600-hour values, which makes training unbalanced temporally if incorporated correctly (i.e., times 0000-0500 would not be trained to a similar extent). For PIANO training, patches of size 768 × 768 are randomly cropped, with season-aware sampling [Park et al., 2025] yielding twenty frames per month. We reserve 2023 for validation and 2024 for testing. For the Pix2Pix ResNet, radar and satellite images are cropped to 224 × 224 and augmented with flips. Combined precipitation prediction performance is evaluated on January-August 2022 and 2023.

2.4 Baseline models

We select the two SOTA models relevant to satellite-image-based precipitation nowcasting, which are PhyDNet [Guen and Thome, 2020, Pihrt et al., 2022] and the Neural Prediction Model (NPM) [Park et al., 2025]. We do not select radar-to-radar (e.g., the model of Gao et al. [2023]) or fusion-to-radar (e.g., the model of Andrychowicz et al. [2023]) models due to it not performing our specific task. Models are trained on a singular RTX 5090 with a batch size of three.

2.5 Evaluation Metrics

In evaluating the precipitation prediction accuracy, we use the critical success index (CSI), as done by Andrychowicz et al. [2023] and Park et al. [2025]. The CSI is calculated as follows:

$$CSI = \frac{TP}{TP + FP + FN}, \tag{5}$$

Table 1: Comparison of NPM, PhyDNet and our proposed method with CSI(↑).

Method	CSI 4 mm								CSI 8 mm							
	1h	2h	3h	4h	5h	6h	7h	8h	1h	2h	3h	4h	5h	6h	7h	8h
NPM	0.750	0.752	0.753	0.754	0.755	0.756	0.756	0.755	0.599	0.599	0.599	0.600	0.600	0.600	0.600	0.599
PhyDNet	0.737	0.743	0.744	0.747	0.752	0.756	0.760	0.763	0.601	0.611	0.617	0.626	0.638	0.646	0.652	0.654
Ours	0.757	0.758	0.760	0.757	0.760	0.763	0.763	0.763	0.611	0.612	0.613	0.611	0.612	0.616	0.616	0.614

where True Positives (TP) denote the number of correctly predicted precipitation events, False Positives (FP) denote incorrectly predicted precipitation events, and False Negatives (FN) denote the number of missed precipitation events. We specifically use the CSI of moderate (more than 4mm/h; CSI 4mm) and heavy (more than 8mm/h; CSI 8mm) precipitation as metrics as they are known to be more challenging.

3 Results and Discussions

In this section, we evaluate the validity of our proposed PIANO and compare the performance of it by comparing it against several baseline methods.

V-NO outputs Figure 2 demonstrates the vector field outputted due to the V-NO. It can be noticed that all images have a generally rightward (note the bottom right parts of images) and downward trend (note the middle regions of images) in velocity, which is in line with the previously known trend of the atmosphere over South Korea [Baek et al., 2015].

Figure 2: Visualization of the output of the V-NO on January 1st, 2024.

Precipitation Nowcasting Performance Table 1 shows the CSI of each model when paired with the GAN. Our model is able to predict heavier precipitation events of 4mm with more skill than both baseline models, demonstrating our approach's validity and potential. Additionally, it performs well on predicting the first two hours of 8mm precipitation, although it becomes less accurate than PhyDNet in further prediction windows. This is a trend also observed in the work of Park et al. [2025], where PhyDNet outperforms the NPM in later time frames. It is to note that the CSI values calculated from our models are incredibly larger than those reported by Andrychowicz et al. [2023] and Park et al. [2025]. This may be due to the differently sampled data during training and evaluation of Sat2RDR performance that may have led to relatively simpler predictions.

Figure 3: Performance comparison of CSI 4mm of PIANO by month.

Comparison of month Figure 3 shows consistent CSI 4mm performance across all months, indicating strong precipitation nowcasting across the time domain. While there are some seasonal variations, the degree of change is low compared to the NPM, which showed incredibly high variation up to 0.1 [Park et al., 2025].

Sensitivity Analysis Because the value of α in Eq. (4) is important in determining training, we modify the value of α and train varying models to obtain results shown in Table 2.

Table 2: Satellite Prediction MSE of different iterations of PIANO

Model type	1h	2h	3h	4h	5h	6h	7h	8h
PIANO ($\alpha = 5.0$)	0.247	0.269	0.316	0.369	0.424	0.480	0.535	0.592
PIANO ($\alpha = 1.0$)	0.205	**0.252**	**0.309**	**0.366**	**0.419**	**0.471**	**0.520**	**0.569**
PIANO ($\alpha = 0.2$)	0.205	0.261	0.321	0.376	0.431	0.482	0.532	0.580
PIANO ($\alpha = 0.0$)	**0.202**	0.260	0.320	0.375	0.430	0.480	0.529	0.576

It is clear that the value of α must be selected carefully; both over ($\alpha = 5$) and under ($\alpha = 0.2$) constraining the T-NO produces poor results. However, if α is selected well, using the V-NO and adding PINN loss improves performance of the model in predicting further lead times. However, it is to note that the T-NO produces highest performance at a lead time of one hour. This is likely because the PINN loss does not constrain the prediction value of the first hour at all, therefore the PINN loss does not help improve performance.

4 Conclusion

We presented PIANO, a physics-informed dual neural operator framework for satellite-based precipitation nowcasting that cleanly separates two roles: V-NO estimates the velocity field over satellite imagery, while T-NO advances the state forward in time. Empirically, PIANO demonstrates competitive accuracy and seasonal consistency while remaining simple to train and deploy on satellite imagery, demonstrating its potential for future development. Focus on further optimization of model hyperparameters and comprehensive evaluations against strong baselines will be key to fully realizing PIANO's capabilities.

References

Mohammed Ombadi, Mark D. Risser, Alan M. Rhoades, and Charuleka Varadharajan. A warming-induced reduction in snow fraction amplifies rainfall extremes. *Nature*, 619(7969):305–310, June 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06092-7.

Hossein Tabari. Climate change impact on flood and extreme precipitation increases with water availability. *Scientific Reports*, 10(1), August 2020. ISSN 2045-2322. doi: 10.1038/s41598-020-70816-2.

Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, July 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06185-3.

Thorsten Kurth, Shashank Subramanian, Peter Harrington, Jaideep Pathak, Morteza Mardani, David Hall, Andrea Miele, Karthik Kashinath, and Anima Anandkumar. Fourcastnet: Accelerating global high-resolution weather forecasting using adaptive fourier neural operators. In *Proceedings of the Platform for Advanced Scientific Computing Conference*, PASC '23, page 1–11. ACM, June 2023. doi: 10.1145/3592979.3593412.

Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Oriol Vinyals, Jacklynn Stott, Alexander Pritzel, Shakir Mohamed, and Peter Battaglia. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, December 2023. ISSN 1095-9203. doi: 10.1126/science.adi2336.

Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. The era5 global reanalysis. *Quarterly journal of the royal meteorological society*, 146(730):1999–2049, 2020.

Marcin Andrychowicz, Lasse Espeholt, Di Li, Samier Merchant, Alexander Merose, Fred Zyda, Shreya Agrawal, and Nal Kalchbrenner. Deep learning for day forecasts from sparse observations. 2023. doi: 10.48550/ARXIV.2306.06079.

Suman Ravuri, Karel Lenc, Matthew Willson, Dmitry Kangin, Remi Lam, Piotr Mirowski, Megan Fitzsimons, Maria Athanassiadou, Sheleem Kashem, Sam Madge, Rachel Prudden, Amol Mandhane, Aidan Clark, Andrew Brock, Karen Simonyan, Raia Hadsell, Niall Robinson, Ellen Clancy, Alberto Arribas, and Shakir Mohamed. Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597(7878): 672–677, September 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03854-z.

Yuchen Zhang, Mingsheng Long, Kaiyuan Chen, Lanxiang Xing, Ronghua Jin, Michael I. Jordan, and Jianmin Wang. Skilful nowcasting of extreme precipitation with nowcastnet. *Nature*, 619(7970):526–532, July 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06184-4.

Vadim Lebedev, Vladimir Ivashkin, Irina Rudenko, Alexander Ganshin, Alexander Molchanov, Sergey Ovcharenko, Ruslan Grokhovetskiy, Ivan Bushmarinov, and Dmitry Solomentsev. Precipitation nowcasting with satellite imagery. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2680–2688, 2019.

Young-Jae Park, Doyi Kim, Minseok Seo, Hae-Gon Jeon, and Yeji Choi. Data-driven precipitation nowcasting using satellite imagery. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(27):28284–28292, April 2025. ISSN 2159-5399. doi: 10.1609/aaai.v39i27.35049.

Aleksandra Gruca, Federico Serva, Llorenç Lliso, Pilar Rípodas, Xavier Calbet, Pedro Herruzo, Jiří Pihrt, Rudolf Raevskyi, Petr Šimánek, Matej Choma, Yang Li, Haiyu Dong, Yury Belousov, Sergey Polezhaev, Brian Pulfer, Minseok Seo, Doyi Kim, Seungheon Shin, Eunbin Kim, Sewoong Ahn, Yeji Choi, Jinyoung Park, Minseok Son, Seungju Cho, Inyoung Lee, Changick Kim, Taehyeon Kim, Shinhwan Kang, Hyeonjeong Shin, Deukryeol Yoon, Seongha Eom, Kijung Shin, Se-Young Yun, Bertrand Le Saux, Michael K Kopp, Sepp Hochreiter, and David P Kreil. Weather4cast at neurips 2022: Super-resolution rain movie prediction under spatio-temporal shifts. In Marco Ciccone, Gustavo Stolovitzky, and Jacob Albrecht, editors, *Proceedings of the NeurIPS 2022 Competitions Track*, volume 220 of *Proceedings of Machine Learning Research*, pages 292–313. PMLR, 28 Nov–09 Dec 2022.

Dmitrii Kochkov, Janni Yuval, Ian Langmore, Peter Norgaard, Jamie Smith, Griffin Mooers, Milan Klöwer, James Lottes, Stephan Rasp, Peter Düben, Sam Hatfield, Peter Battaglia, Alvaro Sanchez-Gonzalez, Matthew Willson, Michael P. Brenner, and Stephan Hoyer. Neural general circulation models for weather and climate. *Nature*, 632(8027):1060–1066, July 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07744-y.

Yogesh Verma, Markus Heinonen, and Vikas Garg. ClimODE: Climate and weather forecasting with physics-informed neural ODEs. In *The Twelfth International Conference on Learning Representations*, 2024.

Maziar Raissi, Paris Perdikaris, and George 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*, 378:686–707, 2019.

Joel Janek Dabrowski, Daniel Edward Pagendam, James Hilton, Conrad Sanderson, Daniel MacKinlay, Carolyn Huston, Andrew Bolt, and Petra Kuhnert. Bayesian physics informed neural networks for data assimilation and spatio-temporal modelling of wildfires. *Spatial Statistics*, 55:100746, 2023.

Woojin Cho, Minju Jo, Haksoo Lim, Kookjin Lee, Dongeun Lee, Sanghyun Hong, and Noseong Park. Parameterized physics-informed neural networks for parameterized pdes. *arXiv preprint arXiv:2408.09446*, 2024.

Somdatta Goswami, Aniruddha Bora, Yue Yu, and George Em Karniadakis. *Physics-Informed Deep Neural Operator Networks*, page 219–254. Springer International Publishing, 2023. ISBN 9783031366444. doi: 10.1007/978-3-031-36644-4_6.

Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM / IMS J. Data Sci.*, 1(3), May 2024. doi: 10.1145/3648506.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.

Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Pieter Wesseling. *Principles of Computational Fluid Dynamics*. Springer Berlin Heidelberg, 2001. ISBN 9783642051463. doi: 10.1007/978-3-642-05146-3.

Yinhao Zhu, Nicholas Zabaras, Phaedon-Stelios Koutsourelakis, and Paris Perdikaris. Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, 394:56–81, October 2019. ISSN 0021-9991. doi: 10.1016/j.jcp.2019.05.024.

John W. Woods. *Image Enhancement and Analysis*, page 223–256. Elsevier, 2012. ISBN 9780123814203. doi: 10.1016/b978-0-12-381420-3.00007-2.

Vincent Le Guen and Nicolas Thome. Disentangling physical dynamics from unknown factors for unsupervised video prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11474–11484, 2020.

Jiří Pihrt, Rudolf Raevskiy, Petr Šimánek, and Matej Choma. Weatherfusionnet: Predicting precipitation from satellite data. *arXiv preprint arXiv:2211.16824*, 2022.

Zhihan Gao, Xingjian Shi, Boran Han, Hao Wang, Xiaoyong Jin, Danielle C. Maddix, Yi Zhu, Mu Li, and Bernie Wang. Prediff: Precipitation nowcasting with latent diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

Seoin Baek, Heetae Kim, and Hyun Chang. Optimal hybrid renewable power system for an emerging island of south korea: The case of yeongjong island. *Sustainability*, 7(10):13985–14001, October 2015. ISSN 2071-1050. doi: 10.3390/su71013985.

A Results in satellite to satellite image prediction

Figure A.1: Visualization of the ground truth satellite IR band (a) and the prediction by the PIANO (b), NPM (c) and PhyDNet (d) for January 1st, 2024, 0900-1600 hrs

Figure A.2: Visualization of satellite WV at 6.3 μm band (a) and the prediction by the PIANO (b), NPM (c) and PhyDNet (d) for January 1st, 2024, 0900-1600 hrs

Figure A.3: Visualization of satellite WV at 7.3 μm band (a) and the prediction by the PIANO (b), NPM (c) and PhyDNet (d) for January 1st, 2024, 0900-1600 hrs

B Results in satellite to radar transformation

Figure B.1: Visualization of difference between ground truth radar and PIANO predictions(a), NPM predictions (b) and PhyDNet predictions(c) for April 23rd, 2024, 1400-2100 hrs. The date and time was selected due to the recorded high precipitation within that window