
Global 3D Reconstruction of Clouds & Tropical Cyclones

Shirin Ermis

University of Oxford
shirin.ermis@physics.ox.ac.uk

Cesar Aybar

Universitat de València
cesar.aybar@uv.es

Lilli Freischem

University of Oxford
lilli.freischem@physics.ox.ac.uk

Stella Girtsou

National Observatory of Athens
National Technical University of Athens
girtsou.s@gmail.com

Kyriaki-Margarita Bintsi

Harvard Medical School and
Massachusetts General Hospital
kbintsi@mgh.harvard.edu

Emiliano Diaz Salas-Porras

Universitat de València
emdiazsal@gmail.com

Michael Eisinger

European Space Agency
michael.eisinger@esa.int

William Jones

University of Oxford
william.jones@physics.ox.ac.uk

Anna Jungbluth

European Space Agency
anna.jungbluth@esa.int

Benoit Tremblay

Environment & Climate Change Canada
benoit.tremblay@ec.gc.ca

Abstract

Accurate forecasting of tropical cyclones (TCs) remains challenging due to limited satellite observations probing TC structure and difficulties in resolving cloud properties involved in TC intensification. Recent research has demonstrated the capabilities of machine learning methods for 3D cloud reconstruction from satellite observations. However, existing approaches have been restricted to regions where TCs are uncommon, and are poorly validated for intense storms. We introduce a new framework, based on a pre-training–fine-tuning pipeline, that learns from multiple satellites with global coverage to translate 2D satellite imagery into 3D cloud maps of relevant cloud properties. We apply our model to a custom-built TC dataset to evaluate performance in the most challenging and relevant conditions. We show that we can – for the first time – create global instantaneous 3D cloud maps and accurately reconstruct the 3D structure of intense storms. Our model not only extends available satellite observations but also provides estimates when observations are missing entirely. This is crucial for advancing our understanding of TC intensification and improving forecasts.

1 Introduction

Tropical cyclones (TCs) are the most damaging and costly extreme events worldwide, with damages reaching billions of dollars per storm in the United States alone [1]. Although numerical and machine

learning (ML) weather forecasts for TCs have improved in recent decades, significant challenges remain to accurately forecast their paths and intensities [2]. Rapid intensification, a process in which TC winds increase by more than 30 knots in 24 hours, is particularly difficult to forecast [3, 4] and occurs in the majority of the most intense and damaging TCs [5]. Recent studies have shown that cloud microphysics in TCs can play an important role in their intensification [6], as the vertical structure of ice clouds causes radiative heating of the clouds and drives instability [7]. Observational studies of the influence of clouds on TC intensification [8] use the cloud profiling radar (CPR) aboard NASA’s CloudSat mission [9], which measures the vertical distribution and structure of clouds and their microphysical properties [10]. However, CloudSat is limited by its long revisit time (~ 16 days), narrow swath (1.4 km), and fixed time of day sampling. This also applies to lidar instruments (e.g., onboard NASA’s CALIPSO [11]) and novel satellites like ESA’s EarthCARE mission [12]. In contrast, geostationary imagery provides large coverage, typically every 10 minutes, but is restricted to measurements of cloud tops.

Deriving information in the vertical domain from observations of cloud tops is challenging, but can be achieved - for instance - through cloud tomography techniques like stereo photogrammetry [13, 14]. This involves collecting multiple perspectives of the same scene and using classical computer vision techniques to estimate depth and derive, for example, cloud top height, as is operationally done for NASA’s MISR mission [13]. However, simultaneous observations from multiple perspectives are rarely available for satellites, and classical cloud tomography provides limited information below cloud tops. Beyond photogrammetry, vertical information can be derived by combining aligned observations from imaging and profiling sensors. Motivated by the recent launch of ESA’s EarthCARE mission, Barker et al. developed a statistical pattern-matching algorithm to operationally extend observations from narrow vertical profiles and provide estimates of cloud volumes [15]. Following on this, recent research has demonstrated successes in the application of ML models to predict 3D volumes of CloudSat measurements from geostationary satellite imagery, including the use of U-Nets [16] for predicting radar reflectivity (Z) [17], ice water content (IWC) and ice crystal number concentration [18], and the use of vision transformers [19, 20]. ML is particularly suitable for this task, as it can effectively learn intricate spatial, temporal, and spectral patterns from large datasets. Building on these advances, we introduce a novel pre-training–fine-tuning framework that integrates data from multiple geostationary satellites to enable global, near real-time 3D prediction of key cloud and TC properties including Z , IWC, and droplet effective radius (r_e), facilitating detailed analysis of TC structure at high temporal cadence. Our model is based on a SWinMAE architecture [21], and encodes temporal and spatial context, including solar and satellite viewing geometry. We propose the first geospatially-aware ML model for global, near real-time 3D cloud reconstructions and 3D reconstruction of tropical cyclones.

2 Data

We compiled new multi-sensor datasets for 3D reconstruction of cloud/TC structure and microphysics, combining imagery from three geostationary satellites to achieve unprecedented diversity in viewing angles, cloud types, and geographic coverage, with co-located vertical profiles from CloudSat.

Geostationary satellite imagery. We use reflectance and brightness temperature (BT) data from three geostationary satellites: Meteosat Second Generation (MSG)/SEVIRI (centered at 0° longitude, 11 spectral channels, 3 km resolution at nadir, from 2004) [22], Himawari-8/AHI (centered at 140.7° E, 16 spectral channels, 2 km resolution at nadir, from 2015) [23], and NOAA’s GOES-16/ABI (centered at 75.2° W, 16 spectral channels, 2 km resolution at nadir, from 2018) [24]. Each geostationary imager has a field of view of $\pm 80^\circ$ which we limit to $\pm 45^\circ$ to reduce distortion effects. Full-disk scans are taken every 10 minutes by GOES-16 and Himawari, and every 15 minutes by MSG, enabling continuous monitoring of clouds and TCs.

Vertical profiles. We use vertical profiles of Z [25], IWC, and r_e [26] retrieved from CloudSat’s CPR [9] and CALIPSO’s lidar [11]. The full data record from 2006 to 2020 is used (daytime only from 2012 after CloudSat’s battery failure).

ML-ready datasets. We prepared three ML-ready datasets from the geostationary satellite imagery and vertical profiles: (1) a *pre-training dataset* consisting of 50,000 randomly sampled patches of 1024×1024 pixels per satellite; (2) a *clouds dataset* consisting of geostationary satellite imagery and spatially-temporally aligned CloudSat overpasses, and (3) a dedicated *TC dataset* consisting of

imagery and overpasses over tropical cyclones. For each image-profile pair, we align profiles of Z , IWC, r_e and cloud type classification through nearest neighbour averaging to the closest geostationary sensor pixel. More details on our datasets can be found in appendix tables 1, 2, 3.

3 Method

The objective of our model is to translate 2D multi-spectral imagery from geostationary satellites into 3D volumes of cloud properties. For each image-profile pair, metrics (including the model loss) are calculated only over the narrow ground-truth CloudSat measurement.

Data normalisation. To combine the diverse geostationary satellite sensors and create a unified input to our model, the 11 spectral channels with wavelengths closest to those of MSG/SEVIRI are selected from GOES/ABI and Himawari/AHI. Each spectral channel is normalised to a range of $[-1, 1]$ using min-max normalisation (reflectances: 0–100%; BT: 180–350 K). The target CloudSat vertical profiles are normalised to $[-1, 1]$ using min-max normalisation (Z : -30–20 dBz; IWC: 10^{-5} – 10 gm^{-3} ; r_e : 0–160 μm). IWC is log-normalised to account for the large skew in its distribution. From the original 125 vertical height levels in the CloudSat data, we remove the lowest 20 levels (below ground level) and upper 25 levels (above cloud level), leaving 80 height levels.

Baseline. We compare our model to a state-of-the-art approach for 3D cloud reconstruction [17]. This approach uses a 2D residual U-Net [27] of depth 4, with 32 channels in the initial convolution layer. The output layer produces 80 height levels.

Pre-training. A SWin transformer-based [28] masked autoencoder [SWinMAE: 21] is used for large-scale self-supervised pre-training on unlabelled geostationary imagery. Since CloudSat measures vertical profiles via a sun-synchronous orbit (i.e. measuring each location always at the same local time), our target data is inherently temporally biased. Pre-training on general cloud scenes sampled outside of CloudSat overpasses helps to overcome this constraint. During pre-training, we mask 50% of the input image, tasking the model to use spatial context to reconstruct missing information. Examples of image reconstructions are shown in appendix fig. 4. The SWin transformer backbone offers two main advantages over a traditional vision transformer: hierarchical feature extraction and computational efficiency. Combined, this enables the model to capture fine local structures and global mesoscale cloud organization. In each training step, we randomly crop 256×256 pixel patches from our larger 1024×1024 pre-training dataset. Each batch is satellite-consistent, but up to all three geostationary satellites are shown to the model during training. We compare two configurations: (i) spectral-only input, and (ii) spectral plus metadata embeddings (SWinSatMAE: time, coordinates, solar/satellite viewing angles) combining elements from SatMAE [29]. More training details can be found in appendix section 5.

Fine-tuning. For fine-tuning, we replace the MAE image reconstruction head with a task-specific 3D convolutional decoder that outputs a cloud property volume. We can perform either single-variable or multi-variable predictions using multiple output heads. In the latter case, we train the model to output Z , IWC, and r_e simultaneously, leveraging the shared structure and cross-correlation between these variables for improved predictions. An overview of our model pipeline is shown in fig. 1.

4 Results

Baseline comparison. In comparison to the U-Net, our model performs better at predicting Z , IWC, and r_e with lower root-mean-squared-error (RMSE) for all variables for both general cloud scenes and TCs (see appendix tables 6, 7). For TCs, the SWinSatMAE produces more accurate values, and better predicts cloud top and base height (fig. 2). Spatially, the SWinSatMAE produces more consistent predictions, with improvements over land and at higher satellite viewing angles (fig. 3).

Single vs multi-satellite: To evaluate the effects of including different sensors in our model training, we compare the baseline model to three U-Nets trained on each geostationary satellite separately. While the multi-satellite model has higher RMSE for MSG, it has lower RMSE for GOES and Himawari, indicating that the larger training dataset helps improve performance for these sensors, and the simple channel matching approach does not significantly affect model performance (see appendix table 4). The multi-satellite model improves predictions for TC regions.




Figure 1: Overview of our ML pipeline. We select the 11 closest-matched spectral channels from the 16 channels of GOES and Himawari to create a consistent model input. During pre-training, the image encoder and decoder learn cloud structures by reconstructing masked images. During fine-tuning a 3D decoder is trained using paired image-profile pairs. Multiple prediction heads are used to predict different cloud properties simultaneously.




Figure 2: 3D reconstructions of (a) Z , (b) IWC , and (c) r_e by the SWinSatMAE model for TC Dorian. The geostationary image from GOES channel 7 is shown under each 3D render, with the location of the CloudSat track marked in red. For validation purposes, the SWinSatMAE predictions along the CloudSat overpass are compared to the CloudSat retrievals and to the multi-variable U-Net baseline.




Figure 3: Spatial RMSE distribution of Z (top), IWC (middle), and r_e (bottom) predictions for the baseline model (left) and the SWinSatMAE model (middle), along with their difference (right).

Single vs multi-variable. U-Net models trained to predict a single variable each are compared to the baseline model. Overall, the multi-variable model produces better predictions across all variables for all metrics. For TCs, the multi-variable model has lower RMSE for Z and r_e (see appendix table 5).

Pre-training & encoding. We compare the pre-trained SWinSatMAE model to a SWinSatMAE trained from scratch in fine-tuning and a SwinMAE model without meta-data encoding. Overall, the pre-trained SWinSatMAE produces lower RMSE across all variables. For TCs, the SwinMAE produces better predictions, but the difference to the pre-trained SWinSatMAE is small (see appendix tables 6 and 7). When breaking down the metrics by cloud type, the SWinSatMAE performs better than the other models in all cloudy conditions for r_e , and in most cloud types for Z and IWC , but has worse metrics in clear skies (see appendix tables 6, 7, and 8).

5 Conclusion

Measurements of vertically-resolved cloud properties are essential for understanding the complex processes involved in TC intensification. To address this need, we curated an AI-ready dataset consisting of pairs of geostationary satellites images and CloudSat overpasses, including a TC-specific subset. Using this dataset, we developed and trained the SWinSatMAE model; an architecture inspired by SWin transformers and masked autoencoders, which advances the state-of-the-art in 3D cloud reconstruction by predicting multiple microphysical properties with higher accuracy than existing approaches. Beyond improvements in accuracy, the model leverages multiple geostationary satellites to enable near real-time, global 3D predictions of clouds (appendix fig. 10) and TCs (appendix fig. 11), even with limited paired observational data. In effect, the trained model enhances the observational capabilities of available geostationary satellites by enabling the inference of microphysical properties from the radiance measurements as if a virtual CloudSat were observing the same field of view.

While our model shows improvements already, there are certain areas where future research could lead to further advancement:

Error characterization. The model tends to smooth cloud edges, reflecting high uncertainty in these regions. Several loss functions—including Huber, total variation, and Gaussian mixture loss—were tested without yielding a definitive improvement. Future work should examine probabilistic or generative methods capable of sharpening cloud boundaries while maintaining overall accuracy. Current error characterizations are limited to region and cloud type; analyses by altitude, climatological regime, and local weather conditions can be carried out and inform future model improvements.

Validation. Additionally, a more comprehensive validation is essential before this model can be applied to downstream scientific or operational tasks. This requires comparison against independent sources of information. We identify three main categories: (i) airborne field campaigns providing in situ and remote-sensing measurements, (ii) ground-based remote-sensing networks, and (iii) satellite-based observations. Airborne campaigns such as ORCESTRA [30] provide in situ sampling of cloud microphysical properties, together with airborne radar, lidar, and polarimetric observations. These measurements can directly constrain estimates of reflectivity, bulk water content, and effective radius. Ground-based networks, such as the Atmospheric Radiation Measurement program [ARM: 31], deliver continuous datasets that include radar reflectivity profiles, ice and liquid water content retrievals from radar–lidar synergy, and effective radius estimates from combined radar–lidar–microwave approaches. Finally, satellite observations from EarthCARE [12] provide significant advantages over CloudSat, including improved lidar vertical resolution, the addition of Doppler velocities, and collocated broadband radiometry. Together, these enable more accurate and vertically resolved microphysical retrievals, as well as direct evaluation of the radiative consistency of model outputs.

Sensor dependence. A fully sensor-independent approach [e.g., 32] could improve multi-satellite generalization and exploit the full spectral range of modern sensors. Furthermore, expanded validation on TCs will be critical to improve reliability in the most challenging storm conditions. While numerical and ML forecasting models capture large-scale TC dynamics, they struggle to resolve the cloud processes most closely tied to rapid intensification. Satellites such as CloudSat offer valuable vertical information but lack the temporal resolution required for operational forecasting. By enabling near real-time 3D predictions of TC cloud properties from geostationary observations, our approach complements existing forecasting systems and offers new opportunities to better anticipate and mitigate the impacts of these devastating weather events.

Acknowledgments and Disclosure of Funding

This work has been enabled by Frontier Development Lab Earth Systems Lab (<https://eslab.ai/>)—a public / private partnership between the European Space Agency (ESA), Trillium Technologies, the University of Oxford and leaders in commercial AI supported by Google Cloud, Scan Computers, Nvidia Corporation and Pasteur Labs. The authors would also like to thank the reviewers and experts that provide advice throughout this project, including Julien Boussard, Gherardo Varando, Homer Durand, Milton Gomez, Johanna Mayer, Luis Gómez-Chova, Matteo Salvador, Alistair Francis, Mikolaj Czerkawski, Jacqueline Campbell, Emmanuel Johnson, Howard Barker, Sarah Brüning, Arthur Avenas, and Dominique Brunet.

References

- [1] NOAA. U.S. Billion-Dollar Weather & Climate Disasters 1980-2022, 2023.
- [2] Russell L. Elsberry, Hsiao-Chung Tsai, Wen-Hsin Huang, and Timothy P. Marchok. New Challenges for Tropical Cyclone Track and Intensity Forecasting in Unfavorable External Environment in Western North Pacific. Part I. Formations South of 20° N. *Atmosphere*, 16(2):226, February 2025.
- [3] Marie-Dominique Leroux, Kimberly Wood, Russell L. Elsberry, Esperanza O. Cayan, Eric Hendricks, Matthew Kucas, Peter Otto, Robert Rogers, Buck Sampson, and Zifeng Yu. Recent Advances in Research and Forecasting of Tropical Cyclone Track, Intensity, and Structure at Landfall. *Tropical Cyclone Research and Review*, 7(2):85–105, May 2018.
- [4] Eric A. Hendricks, Scott A. Braun, Jonathan L. Vigh, and Joseph B. Courtney. A summary of research advances on tropical cyclone intensity change from 2014-2018. *Tropical Cyclone Research and Review*, 8(4):219–225, December 2019.
- [5] J. A. Sippel. TROPICAL CYCLONES AND HURRICANES | Hurricane Predictability. In Gerald R. North, John Pyle, and Fuqing Zhang, editors, *Encyclopedia of Atmospheric Sciences (Second Edition)*, pages 30–34. Academic Press, Oxford, January 2015.
- [6] James H. Ruppert, Allison A. Wing, Xiaodong Tang, and Erika L. Duran. The critical role of cloud–infrared radiation feedback in tropical cyclone development. *Proceedings of the National Academy of Sciences*, 117(45):27884–27892, November 2020.
- [7] Allison A. Wing. Acceleration of Tropical Cyclone Development by Cloud-Radiative Feedbacks. *Journal of the Atmospheric Sciences*, 79(9):2285–2305, September 2022.
- [8] Tsung-Yung Lee and Allison A. Wing. Satellite-Based Estimation of the Role of Cloud-Radiative Interaction in Accelerating Tropical Cyclone Development. *Journal of the Atmospheric Sciences*, 81(6):959–982, June 2024.
- [9] Graeme L. Stephens, Deborah G. Vane, Ronald J. Boain, Gerald G. Mace, Kenneth Sassen, Zhien Wang, Anthony J. Illingworth, Ewan J. O’connor, William B. Rossow, Stephen L. Durden, Steven D. Miller, Richard T. Austin, Angela Benedetti, and Cristian Mitrescu. THE CLOUDSAT MISSION AND THE A-TRAIN: A New Dimension of Space-Based Observations of Clouds and Precipitation. *Bulletin of the American Meteorological Society*, 83(12):1771–1790, December 2002.
- [10] Graeme Stephens, David Winker, Jacques Pelon, Charles Trepte, Deborah Vane, Cheryl Yuhas, Tristan L’Ecuyer, and Matthew Lebsock. CloudSat and CALIPSO within the A-Train: Ten Years of Actively Observing the Earth System. *Bulletin of the American Meteorological Society*, 99(3):569–581, March 2018.
- [11] D. M. Winker, J. Pelon, J. A. Coakley, S. A. Ackerman, R. J. Charlson, P. R. Colarco, P. Flamant, Q. Fu, R. M. Hoff, C. Kittaka, T. L. Kubar, H. Le Treut, M. P. McCormick, G. Mégie, L. Poole, K. Powell, C. Trepte, M. A. Vaughan, and B. A. Wielicki. The CALIPSO Mission: A Global 3D View of Aerosols and Clouds. *Bulletin of the American Meteorological Society*, 91(9):1211–1230, September 2010.
- [12] T. Wehr, T. Kubota, G. Tzeremes, K. Wallace, H. Nakatsuka, Y. Ohno, R. Koopman, S. Rusli, M. Kikuchi, M. Eisinger, T. Tanaka, M. Taga, P. Deghaye, E. Tomita, and D. Bernaerts. The earthcare mission – science and system overview. *Atmospheric Measurement Techniques*, 16(15):3581–3608, 2023.
- [13] Ákos Horváth and Roger Davies. Simultaneous retrieval of cloud motion and height from polar-orbiter multiangle measurements. *Geophysical Research Letters*, 28(15):2915–2918, August 2001.
- [14] Lea Volkmer, Tobias Kölling, Tobias Zinner, and Bernhard Mayer. Consideration of the cloud motion for aircraft-based stereographically derived cloud geometry and cloud top heights. *Atmospheric Measurement Techniques*, 17(23):6807–6817, December 2024.

- [15] H. W. Barker, M. P. Jerg, T. Wehr, S. Kato, D. P. Donovan, and R. J. Hogan. A 3d cloud-construction algorithm for the earthcare satellite mission. *Quarterly Journal of the Royal Meteorological Society*, 137(657):1042–1058, April 2011.
- [16] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation, May 2015.
- [17] Sarah Brüning, Stefan Niebler, and Holger Tost. Artificial intelligence (AI)-derived 3D cloud tomography from geostationary 2D satellite data. *Atmospheric Measurement Techniques*, 17(3):961–978, February 2024.
- [18] Kai Jeggle, Mikolaj Czerkawski, Federico Serva, Bertrand Le Saux, David Neubauer, and Ulrike Lohmann. IceCloudNet: 3D reconstruction of cloud ice from Meteosat SEVIRI, October 2024.
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, June 2021.
- [20] Stella Girtsou, Emiliano Diaz Salas-Porras, Lilli Freischem, Joppe Massant, Kyriaki-Margarita Bintsi, Guiseppe Castiglione, William Jones, Michael Eisinger, Emmanuel Johnson, and Anna Jungbluth. 3D Cloud reconstruction through geospatially-aware Masked Autoencoders, January 2025. arXiv:2501.02035 [cs].
- [21] Zi'an Xu, Yin Dai, Fayu Liu, Weibing Chen, Yue Liu, Lifu Shi, Sheng Liu, and Yuhang Zhou. Swin mae: Masked autoencoders for small datasets. *Computers in biology and medicine*, 161:107037, 2023.
- [22] D. M. A. Aminou. MSG's SEVIRI Instrument. *ESA bulletin*, 111:15–17, August 2002.
- [23] Kotaro Bessho, Kenji Date, Masahiro Hayashi, Akio Ikeda, Takahito Imai, Hidekazu Inoue, Yukihiro Kumagai, Takuya Miyakawa, Hidehiko Murata, Tomoo Ohno, Arata Okuyama, Ryo Oyama, Yukio Sasaki, Yoshio Shimazu, Kazuki Shimoji, Yasuhiko Sumida, Masuo Suzuki, Hidetaka Taniguchi, Hiroaki Tsuchiyama, Daisaku Uesawa, Hironobu Yokota, and Ryo Yoshida. An Introduction to Himawari-8/9— Japan's New-Generation Geostationary Meteorological Satellites. *Journal of the Meteorological Society of Japan. Ser. II*, 94(2):151–183, 2016.
- [24] Timothy J. Schmit, Paul Griffith, Mathew M. Gunshor, Jaime M. Daniels, Steven J. Goodman, and William J. Leblair. A Closer Look at the ABI on the GOES-R Series. *Bulletin of the American Meteorological Society*, 98(4):681–698, July 2016.
- [25] Roger Marchand, Gerald G. Mace, Thomas Ackerman, and Graeme Stephens. Hydrometeor Detection Using Cloudsat—An Earth-Orbiting 94-GHz Cloud Radar. *Journal of Atmospheric and Oceanic Technology*, 25(4):519–533, April 2008.
- [26] Min Deng, Gerald. G. Mace, Zhien Wang, and Elizabeth Berry. CloudSat 2C-ICE product update with a new Ze parameterization in lidar-only region. *Journal of Geophysical Research: Atmospheres*, 120(23):12,198–12,208, 2015.
- [27] Foivos I. Diakogiannis, François Waldner, Peter Caccetta, and Chen Wu. ResUNet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 162:94–114, April 2020.
- [28] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows, August 2021.
- [29] Yezhen Cong, Samar Khanna, Chenlin Meng, Patrick Liu, Erik Rozi, Yutong He, Marshall Burke, David B. Lobell, and Stefano Ermon. SatMAE: Pre-training Transformers for Temporal and Multi-Spectral Satellite Imagery, January 2023.
- [30] Daniel Klocke, Marcus Dengler, Tim Carlsen, Robert Oscar David, Holger Baars, Annett Skupin, Sandrine Bony, Julien Delanoë, Silke Gross, Bjorn Stevens, Julia Windmiller, Allison Wing, Raphaela Vogel, and Geet George. Orcestra - organized convection and earthcare studies over the tropical atlantic. <https://orcestra-campaign.org/>.

- [31] G. M. Stokes and S. E. Schwartz. The atmospheric radiation measurement (arm) program: Programmatic background and design of the cloud and radiation test bed. *Bull. Amer. Meteor. Soc.*, 75:1201–1222, 1994.
- [32] Alistair Francis. Sensor Independent Cloud and Shadow Masking With Partial Labels and Multimodal Inputs. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–18, 2024.
- [33] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015.
- [34] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.

Appendix

Dataset Details

	Coverage	Product	Resolution	# Channels	# Files	Size
GOES	2018 - 2024	MCMIP	2 km @ SSP	16	50,000	793 GB
MSG	2004 - 2025	L1b	3 km @ SSP	11	50,000	646 GB
HIMAWARI	2015 - 2022	L1b	2 km @ SSP	16	50,000	930 GB

Table 1: Details of our pre-training dataset. We processed 50,000 patches of 1024×1024 that are randomly cropped to 256×256 pixels in each training step.

	Coverage	Resolution	# Channels	# Files	Size
GOES-CloudSat	2018 - 2024	2 km @ SSP	16	31,046	72 GB
MSG-CloudSat	2004 - 2025	3 km @ SSP	11	181,653	620 GB
HIMAWARI-CloudSat	2015 - 2022	2 km @ SSP	16	57,373	181 GB

Table 2: Details of our clouds dataset. We aligned the geostationary satellite imagery and corresponding CloudSat overpasses in space and time, and save crops of 256×256 pixels.

	Coverage	Resolution	# Channels	# Files	Size
GOES-CloudSat	2018 - 2024	2 km @ SSP	16	185	341 MB
MSG-CloudSat	N/A	3 km @ SSP	11	N/A	N/A
HIMAWARI-CloudSat	2015 - 2022	2 km @ SSP	16	518	1.4 GB

Table 3: We aligned the geostationary satellite imagery of tropical cyclones and corresponding CloudSat overpass in space and time, based on TC track information provided by the International Best Track Archive for Climate Stewardship (IBTrACS). We crop 256×256 pixels around each overpass. We consider all CloudSat overpasses within 256 km of the TC center. We focus on storms in the North Atlantic and Eastern Pacific for GOES, and the Western and Southern Pacific for HIMAWARI. Note that there are no cyclones in the MSG field-of-view. Our TC dataset is reserved exclusively for evaluation, providing a rigorous testbed for assessing model performance under the most intense storm conditions.

Training Details

We conducted our experiments on a single NVIDIA V100 GPU via Google Cloud, with batch size of 32 during pre-training and 8-16 during fine-tuning. We used the Adam optimizer [33] with a learning rate of 0.00015, using backpropagation [34]. Training was optimized via the Mean Squared Error (MSE) loss, while additional metrics like the Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index Measure (SSIM) were monitored. Checkpointing was used to save models with the lowest validation loss. Pre-training and fine-tuning ran for 50 and 100 epochs respectively. Regarding training times, the pre-training took around 6 hours, while fine-tuning ran for 5 hours - 3 days depending on the complexity of the model. Our U-Net baseline contains 1.9M trainable parameters, while our SWinSatMAE model contains 34.2M trainable parameters. Increasing the complexity of the U-Net to match or exceed the parameters of our model did not improve prediction results. Our U-Net was trained without dropout, 10^{-5} weight decay, and 4 up- and down-blocks of residual convolutions. Our SWinSatMAE model was pre-trained with a token size of 2×2 pixels, a masking window of 4×4 pixels, an attention window of 32×32 pixels, and 50% masking. During fine-tuning, we replaced the pre-training decoder with a custom SwinConv decoder that first reverses the operations of the Swin encoder, and then uses repeated blocks of ConvTranspose2D, ResidualConv, and ReLU activations to reach the final desired output of 256×256 pixels and 96 height levels. We then predict the n target variables and the final 80 height levels using n prediction heads, each made up of sequentially applied residual convolution, ReLU, convolution operations. We encode the latitude/longitude coordinates, time of measurement (fraction of the day and fraction of the year), satellite viewing angle (zenith and azimuth), and solar angle (zenith and azimuth) together with the positional encoding to make our model geospatially aware. We split our dataset by time, and allocate days 2-22 each month to

training, 24 - 26 to validation, and 28 - 31 to testing. We purposely leave a gap of 1 day to avoid data leakage. During fine-tuning on our clouds dataset, we filter our examples that contain less than 25% of cloudy columns.

Image Reconstructions During Pretraining




Figure 4: Comparison of masked, predicted, and original images (MSG example) during pre-training of our SWinSatMAE model. We pretrained our model for 50 epochs.

U-Net: Single Satellite vs. Multi-Satellite Input



(a) Model Input	Single Satellite		
Satellite	GOES	MSG	HIMAWARI
RMSE (dBZ)	5.44 ± 2.61	4.81 ± 2.21	6.14 ± 2.99
SSIM	0.75 ± 0.13	0.78 ± 0.11	0.71 ± 0.14
DICE	0.75 ± 0.13	0.88 ± 0.15	0.84 ± 0.16

(b) Model Input	Multi-Satellite			
Satellite	GOES	MSG	HIMAWARI	Combined
RMSE (dBZ)	5.22 ± 2.65	4.95 ± 2.17	6.03 ± 3.05	5.40 ± 2.69
SSIM	0.76 ± 0.13	0.77 ± 0.11	0.72 ± 0.15	0.75 ± 0.13
DICE	0.84 ± 0.18	0.87 ± 0.17	0.83 ± 0.17	0.85 ± 0.17



(c) Model Input	Single Satellite		
Satellite	GOES	MSG	HIMAWARI
RMSE (dBZ)	7.65 ± 4.42	N/A	10.82 ± 5.01
SSIM	0.65 ± 0.21	N/A	0.47 ± 0.22
DICE	0.82 ± 0.2	N/A	0.84 ± 0.15

(d) Model Input	Multi-Satellite			
Satellite	GOES	MSG	HIMAWARI	Combined
RMSE (dBZ)	8.24 ± 4.95	N/A	9.41 ± 4.33	8.89 ± 4.62
SSIM	0.62 ± 0.22	N/A	0.5 ± 0.21	0.55 ± 0.22
DICE	0.77 ± 0.28	N/A	0.84 ± 0.17	0.81 ± 0.23

Table 4: Comparison of a U-Net model, trained to predict Z using (a/c) a single satellite or (b/d) all three satellites as model input. Sub-tables (a) and (b) show the performance on our clouds test set, while (c) and (d) show the performance on our entire cyclone dataset, which was not seen during training. When training on each satellite individually, we use all spectral channels, i.e. 16 spectral channels for GOES and HIMAWARI, and 11 spectral channels for MSG. To combine the satellites into a unified dataset, we match the closest 11 spectral channels from each satellite and ignore resolution differences. We report the root-mean-squared error (RMSE; in dBZ), structural similarity index measure (SSIM; unitless) and Dice coefficient (unitless). All models were trained for 100 epochs, and the best validation loss checkpoint was chosen for inference. Despite differences in spectral characteristics and resolution of the different satellites, naive matching of satellites and spectral channels does not degrade model performance.




Figure 5: Comparison of prediction performance across the globe between U-Net models trained on our three satellites, MSG, GOES and Himawari individually (top), and a U-Net model trained on all three satellites together (Multi-Satellite U-Net, middle).

U-Net: Single Variable vs. Multi-Variable Target



(a) Model Target	Single Variable		
Variable	Z (dBZ)	Ice Water Content (g/m ³)	Effective Radius (μm)
RMSE (units)	5.40 ± 2.69	0.07 ± 0.08	10.99 ± 5.20
SSIM	0.75 ± 0.13	0.91 ± 0.07	0.52 ± 0.11
DICE	0.85 ± 0.17	0.30 ± 0.09	0.03 ± 0.05

(b) Model Target	Multi-Variable		
Variable	Z (dBZ)	Ice Water Content (g/m ³)	Effective Radius (μm)
RMSE (units)	5.10 ± 2.46	0.06 ± 0.08	10.30 ± 5.21
SSIM	0.77 ± 0.13	0.91 ± 0.07	0.55 ± 0.11
DICE	0.85 ± 0.17	0.30 ± 0.09	0.03 ± 0.04



(c) Model Target	Single Variable		
Variable	Z (dBZ)	Ice Water Content (g/m ³)	Effective Radius (μm)
RMSE (units)	9.03 ± 4.95	0.16 ± 0.16	14.80 ± 7.99
SSIM	0.57 ± 0.24	0.84 ± 0.12	0.48 ± 0.13
DICE	0.82 ± 0.20	0.31 ± 0.07	0.03 ± 0.03

(d) Model Target	Multi-Variable		
Variable	Z (dBZ)	Ice Water Content (g/m ³)	Effective Radius (μm)
RMSE (units)	8.55 ± 4.39	0.17 ± 0.16	14.66 ± 7.72
SSIM	0.57 ± 0.24	0.83 ± 0.12	0.49 ± 0.14
DICE	0.80 ± 0.19	0.31 ± 0.07	0.03 ± 0.04

Table 5: Comparison of a U-Net model, trained to predict either one variable (a/c) or all three variables jointly (b/d). Sub-tables (a) and (b) show the performance on our clouds test set, while (c) and (d) show the performance on our entire cyclone dataset, which was not seen during training. All models are trained on the closest matched spectral channels for GOES, MSG, and HIMAWARI. We report the root-mean-squared error (RMSE; in variable units), structural similarity index measure (SSIM; unitless), and Dice coefficient (unitless). All models were trained for 100 epochs, and the best validation loss checkpoint was chosen for inference.

Advanced Model Comparison



(a)	Z RMSE (dBZ)	Ice Water Content RMSE (g/m³)	Effective Radius RMSE (μm)
U-Net Baseline	5.10 ± 2.46	0.063 ± 0.081	10.3 ± 5.21
SWinMAE	4.21 ± 2.01	0.055 ± 0.071	7.93 ± 4.26
SWinSatMAE (no pre-training)	4.11 ± 1.98	0.055 ± 0.070	7.76 ± 4.21
SWinSatMAE (our model)	4.04 ± 1.93	0.053 ± 0.069	7.72 ± 3.99



(b)	Z RMSE (dBZ)	Ice Water Content RMSE (g/m³)	Effective Radius RMSE (μm)
U-Net Baseline	8.55 ± 4.39	0.166 ± 0.160	14.66 ± 7.72
SWinMAE	7.01 ± 2.61	0.144 ± 0.134	11.53 ± 5.12
SWinSatMAE (no pre-training)	7.08 ± 2.76	0.144 ± 0.136	11.74 ± 5.46
SWinSatMAE (our model)	7.08 ± 2.62	0.147 ± 0.148	11.55 ± 5.03

Table 6: **Clear & Cloudy**: Comparison of a U-Net model, a SWinMAE without encodings, our SWinSatMAE trained from scratch and including pre-training, across clear and cloudy pixels. We report the root-mean-squared error (RMSE; variable units) for (a) our clouds test set, and (b) for our entire cyclone dataset, which was not seen during training. All models are trained on the closest matched spectral channels for GOES, MSG, and HIMAWARI. All models were trained for 100 epochs, and the best validation loss checkpoint was chosen for inference.



(a)	Z RMSE (dBZ)	Ice Water Content RMSE (g/m³)	Effective Radius RMSE (μm)
U-Net Baseline	9.50 ± 5.30	0.082 ± 0.164	14.79 ± 13.08
SWinMAE	8.79 ± 5.11	0.075 ± 0.149	13.08 ± 11.47
SWinSatMAE (no pre-training)	8.57 ± 5.01	0.074 ± 0.149	12.91 ± 11.18
SWinSatMAE (our model)	8.46 ± 4.99	0.072 ± 0.145	12.40 ± 10.79



(c)	Z RMSE (dBZ)	Ice Water Content RMSE (g/m³)	Effective Radius RMSE (μm)
U-Net Baseline	12.75 ± 6.18	0.111 ± 0.179	16.52 ± 12.53
SWinMAE	11.40 ± 5.58	0.104 ± 0.160	13.70 ± 10.70
SWinSatMAE (no pre-training)	11.47 ± 5.61	0.103 ± 0.159	14.27 ± 11.10
SWinSatMAE (our model)	11.50 ± 5.65	0.101 ± 0.160	13.21 ± 10.49

Table 7: **Cloudy only**: Comparison of a U-Net model, a SWinMAE without encodings, our SWinSatMAE trained from scratch and including pre-training, across only cloudy pixels. We report the root-mean-squared error (RMSE; variable units) for (a) our clouds test set, and (b) for our entire cyclone dataset, which was not seen during training. All models are trained on the closest matched spectral channels for GOES, MSG, and HIMAWARI. All models were trained for 100 epochs, and the best validation loss checkpoint was chosen for inference.



	Z RMSE (dBZ)			
Cloud Type	U-Net Baseline	SWinMAE	SWinSatMAE (no pre-training)	SWinSatMAE (our model)
No Cloud	3.22 ± 1.52	2.44 ± 1.57	2.35 ± 1.44	2.37 ± 1.43
Cirrus	6.32 ± 2.91	5.65 ± 2.58	5.50 ± 2.47	5.25 ± 2.34
Altostratus	9.69 ± 4.56	8.66 ± 4.23	8.40 ± 4.04	8.23 ± 3.90
Alto cumulus	9.51 ± 4.71	9.04 ± 4.74	8.82 ± 4.73	8.83 ± 4.70
Stratus	7.91 ± 5.02	7.59 ± 5.13	7.29 ± 4.97	7.22 ± 4.84
Stratocumulus	10.88 ± 4.95	10.08 ± 4.92	10.03 ± 4.86	9.88 ± 4.87
Cumulus	9.94 ± 6.47	9.44 ± 6.29	9.21 ± 6.21	9.23 ± 6.27
Nimbostratus	14.12 ± 5.22	12.37 ± 4.93	11.93 ± 4.78	12.07 ± 4.85
Deep Convection	13.03 ± 4.82	12.11 ± 4.62	11.71 ± 4.51	11.64 ± 4.37

	Ice Water Content RMSE (g/m ³)			
Cloud Type	U-Net Baseline	SWinMAE	SWinSatMAE (no pre-training)	SWinSatMAE (our model)
No Cloud	0.005 ± 0.016	0.004 ± 0.016	0.004 ± 0.014	0.004 ± 0.018
Cirrus	0.028 ± 0.042	0.023 ± 0.030	0.023 ± 0.029	0.021 ± 0.026
Altostratus	0.095 ± 0.114	0.084 ± 0.097	0.082 ± 0.095	0.078 ± 0.092
Alto cumulus	0.026 ± 0.043	0.025 ± 0.039	0.024 ± 0.038	0.025 ± 0.039
Stratus	0.030 ± 0.078	0.029 ± 0.072	0.027 ± 0.069	0.026 ± 0.066
Stratocumulus	0.040 ± 0.115	0.037 ± 0.104	0.038 ± 0.106	0.034 ± 0.097
Cumulus	0.054 ± 0.107	0.052 ± 0.102	0.053 ± 0.103	0.055 ± 0.106
Nimbostratus	0.258 ± 0.188	0.232 ± 0.171	0.229 ± 0.172	0.224 ± 0.159
Deep Convection	0.593 ± 0.254	0.544 ± 0.242	0.539 ± 0.244	0.533 ± 0.235

	Effective Radius RMSE (μm)			
Cloud Type	U-Net Baseline	SWinMAE	SWinSatMAE (no pre-training)	SWinSatMAE (our model)
No Cloud	6.69 ± 3.04	4.70 ± 3.17	4.52 ± 2.88	5.03 ± 3.00
Cirrus	17.02 ± 6.35	15.85 ± 5.83	15.47 ± 5.82	14.34 ± 5.64
Altostratus	19.71 ± 12.36	17.48 ± 11.00	17.05 ± 10.60	16.19 ± 10.12
Alto cumulus	17.42 ± 13.19	15.80 ± 11.63	15.71 ± 11.50	15.68 ± 11.13
Stratus	9.45 ± 11.13	8.82 ± 10.75	8.86 ± 10.55	8.26 ± 10.12
Stratocumulus	8.75 ± 13.64	7.62 ± 12.32	7.55 ± 11.89	7.40 ± 11.68
Cumulus	9.38 ± 13.71	7.87 ± 11.48	7.88 ± 11.18	7.68 ± 10.94
Nimbostratus	23.90 ± 13.90	18.31 ± 10.10	17.77 ± 9.83	17.26 ± 9.33
Deep Convection	22.82 ± 9.33	19.88 ± 8.76	19.74 ± 8.74	18.81 ± 8.00

Table 8: Performance comparison across different cloud types. We report the root-mean-squared error (RMSE; variable units) for each cloud classification category in our clouds test set. All models were trained for 100 epochs, and the best validation loss checkpoint was chosen for inference.




Figure 6: Comparison of three of our models: SWinMAE (no geospatial encodings, left), SWinSatMAE (no pre-training, middle), and our chosen SWinSatMAE architecture that was pretrained on geostationary imagery before fine-tuning to predict CloudSat variables.




Figure 7: Radar reflectivity (first column), ice water content (second column), and droplet effective radius (third column) as retrieved by CloudSat (first row) along its swath through hurricane Dorian and as reconstructed by the single-variable U-Net baseline (second row) and the SWinSatMAE model without pre-training (third row). We refer the reader to fig. 2 for the context image of hurricane Dorian as well as the reconstructions by the multi-variable U-Net Baseline and the SWinSatMAE model with pre-training.




Figure 8: **Clear & Cloudy**: Comparison between the CloudSat radar reflectivity (mean per profile) and our model predictions. We consider all pixels, i.e. clear and cloudy for this plot.




Figure 9: Root-mean-square error (RMSE) and peak-signal-to-noise ratio (PSNR) during training.

Cloud and Tropical Cyclone Reconstructions




Figure 10: Using our model, we can –for the first time– generate global instantaneous 3D cloud maps.




Figure 11: Model prediction for tropical cyclone Dorian. The top shows a geostationary satellite image, middle our model reconstruction of radar reflectivity, and bottom the max column radar reflectivity of the 3D prediction.