
Structural Forecasting for Tropical Cyclone Intensity Prediction: Providing Insight with Deep Learning

Trey McNeely¹ Niccolò Dalmasso¹ Kimberly M. Wood² Ann B. Lee¹

¹Department of Statistics & Data Science, Carnegie Mellon University

²Department of Geosciences, Mississippi State University
imcneely@stat.cmu.edu

Abstract

Tropical cyclone (TC) intensity forecasts are ultimately issued by human forecasters. The human in-the-loop pipeline requires that any forecasting guidance must be easily digestible by TC experts if it is to be adopted at operational centers like the National Hurricane Center. Our proposed framework leverages deep learning to provide forecasters with something neither end-to-end prediction models nor traditional intensity guidance does: a powerful tool for monitoring high-dimensional time series of key physically relevant predictors and the means to understand how the predictors relate to one another and to short-term intensity changes.

1 Introduction

Tropical cyclones (TCs) are powerful, highly organized storm systems that help transfer energy from the upper levels of the world’s oceans to the atmosphere. Increasing coastal populations, concurrent with rising ocean temperatures due to climate change, increase the dangers (such as storm surge and damaging winds) posed by TCs. Thus, accurate prediction of TC trajectories and intensities becomes an ever more critical component of disaster preparation and response. Track forecasting has made great strides since the 1990s, but intensity forecasting, especially in the short-term 12-h and 24-h windows, have seen less pronounced improvements [5, 3]. Cases of rapid intensity change (30-knot intensity changes in 24 hours), in particular, have proved difficult to predict [13, 14, 28].

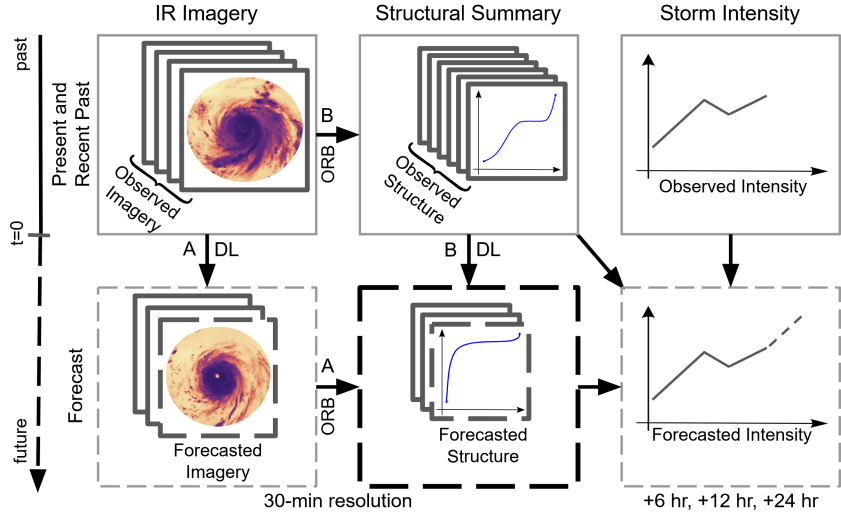
TC prediction and science rely in part on the relationships between the external environment, internal structure, and behavior of these storms. For example, vertical wind shear (an external factor) and TC convective patterns (an internal factor) are both known to relate to TC intensities [6, 14, 12]. Forecasters can access a wide array of real-time TC observations of both types, but new public forecasts must be issued every six hours by expert human forecasters. This human-in-the-loop pipeline requires that any forecasting guidance must be easily understood by TC experts. If the output of a method cannot be digested by scientists in a handful of minutes, it will not be adopted by stakeholders at operational centers such as the National Hurricane Center (NHC).

The NHC and other operational forecast centers use both physics-based models (dynamical models) and data-driven models (statistical models) [3]. The Statistical Hurricane Intensity Prediction Scheme (SHIPS [4]) is a particularly successful statistical-dynamical model which largely relates area averages of environmental fields (such as vertical wind shear) and cloud-top temperatures from infrared (IR) imagery to future TC intensity change. This IR imagery is available at high spatial and temporal resolutions from satellites like the Geostationary Operational Environmental Satellites (GOES¹ [26, 16]). Structure in IR imagery has long been known to relate to TC intensity (see the Dvorak technique [6, 20]), but the area-averaged values commonly used in statistical models discard critical spatial information. The evolution of spatial structure in such fields correlates with intensity change, but unlocking this rich source of information in IR imagery without sacrificing interpretability remains an open problem.

We propose a framework to incorporate evolving spatial information into intensity forecasts by utilizing deep learning (DL) in the high-dimensional time series setting while remaining cognizant of the needs of the end users: forecasters and scientists. While our end goal is the prediction of short-term intensity changes, our framework

¹<https://www.star.nesdis.noaa.gov/goes/>

Figure 1: Framework for combining deep learning (DL) with structural functions (ORB) to provide a structural forecast via two pathways. Final forecasts of intensity (bottom-right) are based on both the observed intensity (top-right) and the observed and forecasted structure (center). The intermediate structural forecast is the key contrast with traditional end-to-end DL methods.



will offer something neither end-to-end models (black-box prediction of intensity directly from IR imagery) nor existing operational forecasts can: **structural summaries** or interpretable quantification of the convective structure of TCs in the form of one-dimensional functions, together with **structural forecasts** of the dynamic evolution of these summaries. This new setup will provide forecasters with a powerful tool for monitoring high-dimensional time series of key physically relevant predictors (like the eye-eyewall structure and symmetry of deep convection relative to the TC center) with the means to understanding how the predictors relate to one another and to TC intensity change.

Figure 1 describes our approach. In this figure, the top row indicates observed information, while the bottom row indicates forecasts; the final intensity forecast (bottom-right) depends on the past intensity (top-right) and a “structural forecast” (bottom-center) — the predicted short-term evolution of TC structure itself. We obtain structural forecasts via two paths (A and B). Each approach extracts a set of interpretable functions and uses DL to propagate the structural state of the TC into the near future. These parallel pathways allow forecasters to check the structural forecasts for agreement via quantitative error metrics in addition to checking them for physical plausibility directly examining both forecasted imagery and structure. Path A evolves the original IR imagery via DL and then computes structural functions on the forecasted imagery. Path B instead computes structural functions first then evolves them directly via DL. Traditional end-to-end DL approaches step directly from the observed IR imagery and intensity history to an intensity forecast; our framework’s inclusion of interpretable structural forecasts as an intermediary step sets it apart from these traditional machine learning applications.

2 Data

Geostationary satellite observations, including the Geostationary Operational Environmental Satellites (GOES), provide high spatial (≤ 4 km) and temporal (≤ 30 min) resolution imagery of the Atlantic and Pacific TC basins. Such observations are consistently available, unlike aircraft reconnaissance or land-based instruments such as radar. We will focus on longwave infrared ($\sim 10.7 \mu\text{m}$), which provides estimates of cloud-top temperature. Because temperature generally decreases with height in the troposphere, low cloud-top temperatures typically indicate regions of stronger thunderstorms and thus deeper convection. The data collected consist of $\sim 200,000$ IR images ($\sim 400 \times 400$ pixels) from 656 unique TCs in the North Atlantic and eastern North Pacific between 2000 and 2019. Storm location and intensity are drawn from the NHC’s HURDAT2 database [18].

3 Providing Insight to TCs via Deep Learning and Forecasted Structure

The ultimate goal of this work is to leverage evolving spatial structures in TCs to better understand storm behavior, specifically short-term storm intensity change (6- to 24-hour time frame). We would like to answer the question “Can we predict short-term intensity change using interpretable functions from GOES-IR?”

Preliminary Work. In prior work [8], we have shown that a set of interpretable structural functions drawn from GOES-IR imagery can nowcast (that is, predict ongoing) rapid intensity changes as well as SHIPS environmental predictors. Our approach extracted a set of ORB functions summarizing the Organization, Radial structure, and Bulk morphology of a TC image as continuous functions (of, e.g., the radius r from the TC center or the threshold

c of level sets). The ORB functions were then compressed via a principal component analysis (PCA), and the PC coefficients served as inputs to logistic lasso models which classified the storms as rapidly changing in intensity or not. The (generalized) linear model allowed the user to directly relate the probability of a rapid TC intensity change to changes in key physical predictors like the eye-eyewall structure and symmetry of deep convection in the storm core. Our initial ORB suite performed as well as a subset of area-averaged SHIPS environmental predictors, while the combination of the two sets outperformed SHIPS alone. For reliable intensity forecasting, however, we need a richer suite of ORB functions, as well as a tool for projecting high-dimensional ORB functions $X_t \in \mathbb{R}^d$ (where d is very large) into the future without a prior dimension reduction.

Structural Summaries. In the proposed work, we will develop a richer suite of structural summary functions that also includes center-independent structural functions, measures of the spatial structure of vector fields, and additional satellite observations such as water vapor ($\sim 6.5\mu\text{m}$) imagery. Crucially, we will compare our final intensity forecasts based on these structural functions with intensity forecasts resulting from end-to-end DL models; in this way, we can assess the richness of the feature suite and quantify information lost in compression w.r.t. the root mean squared (RMS) intensity prediction errors. Instead of focusing on exhaustive feature design, we will predominantly explore new classes of functions to improve the range of physical structures quantified by ORB.

Structural Forecasting. In lieu of an end-to-end deep learning model, we will project our structural summaries of the TC 6, 12, and 24 hours into the future via DL; instead of only answering “How strong will the TC be in 6/12/24 hours?” we will also model “What will the TC look like in 6/12/24 hours?” This approach will provide the critical next step in the prediction pipeline, which enables forecasters to examine the structural forecast before relying on the intensity model. Since they can compare the forecast to the original satellite imagery, this improves stakeholder trust in the model, providing additional clues to the emergence of unrealistic TC structures or trajectories. Furthermore, this framework mimics the utilization of numerical weather models by statistical-dynamical intensity guidance; such models use physical laws rather than statistical learning to forecast the state of the atmosphere, but the final intensity guidance (e.g., SHIPS) still draws on summaries of this atmospheric forecast in the same way we will draw on a structural forecast.

We will generate structural forecasts via two pathways. The first of these (in Figure 1, pathway A) steps the TC imagery forward in time, then computes ORB functions at each time step. This could be achieved using convolutional neural network architectures used in frame-to-frame prediction (as in [1, 9]; see [21] for an exhaustive review). As video prediction tasks can be challenging, a modified version of this pathway could have deep convolutional architectures forecast intensity directly [23, 15, 10]; i.e., an end-to-end model going straight from the top left to the bottom right of Figure 1. The second route (in Figure 1, pathway B) will directly step the ORB functions forward in time, thus optimizing for the outcome of interest (TC structure). While linear auto-regressive models such as ARIMA [2] are appropriate, we plan on leveraging recurrent neural networks [7], which have been shown to be successful in multivariate time series prediction tasks [29, 22, 17, 24] (see [19] for a recent review). Comparing the output of pathways A and B will allow us to compare the impact of using interpretable structural functions on the accuracy of structural forecasts against those automatically extracted by deep learning models.

From Structure to Intensity. Traditional linear models have historically been attractive to forecasters and scientists for reasons of interpretability and good performance in low sample size settings. However, linear models often struggle to capture the complex, time-varying processes which drive these storms. We here propose an approach that leverages DL to handle the high-dimensional structural forecasting problem, while the relationship between structure and intensity is handled by additive models (generalized additive models [11], sparse additive models [25], etc.). Additive models are attractive due to their combination of high-capacity, ease of interpretation, and straightforward visualizations. Importantly, these additive models have access to a structural forecast; the heavy lifting of high-dimensional time series prediction is handled by the deep learning models. In addition to this regression problem, comparison of trajectories via spectral clustering will be used to identify modes of evolution [27]. Such identification of similar trajectories in historical TCs can provide forecasters and scientists alike with analogous TCs against which to compare the evolution of new storms.

Evaluation Metrics. We will have two quantitative metrics for the performance of this framework. First, we will compare the resultant RMS intensity errors and bias at 6, 12, and 24 hours to both end-to-end DL and to the official NHC forecasts for those times. Second, we will compare the L_2 -distance between the resultant structural forecasts from the two pathways (to assess “sufficiency” of tracking only ORB functions) as well as distance from the ground truth structural functions. We acknowledge that a model based solely on satellite imagery with limited atmospheric and oceanographic data will not outperform state-of-the-art intensity prediction schemes on its own. However, a successful model which relates structural evolution to intensity change will be an invaluable addition to the toolbox utilized by forecasters and scientists analyzing TCs.

Acknowledgement This work is supported in part by the NSF AI Planning Institute for Data-Driven Discovery in Physics, NSF PHY-2020295.

References

- [1] Mohammad Babaeizadeh, Chelsea Finn, Dumitru Erhan, Roy H. Campbell, and Sergey Levine. Stochastic variational video prediction. *arXiv e-prints*, page arXiv:1710.11252, October 2017.
- [2] George.E.P. Box and Gwilym M. Jenkins. *Time Series Analysis: Forecasting and Control*. Holden-Day, 1976.
- [3] John P. Cangialosi, Eric Blake, Mark DeMaria, Andrew Penny, Andrew Latta, Edward Rappaport, and Vijay Tallapragada. Recent Progress in Tropical Cyclone Intensity Forecasting at the National Hurricane Center. *Weather and Forecasting*, 35(5):1913–1922, 08 2020.
- [4] Mark DeMaria and John Kaplan. An updated statistical hurricane intensity prediction scheme (SHIPS) for the Atlantic and Eastern North Pacific basins. *Weather and Forecasting*, 14(3):326–337, 1999.
- [5] Mark DeMaria, Charles R. Sampson, John A. Knaff, and Kate D. Musgrave. Is tropical cyclone intensity guidance improving? *Bulletin of the American Meteorological Society*, 95(3):387–398, 2014.
- [6] Vernon F. Dvorak. Tropical cyclone intensity analysis and forecasting from satellite imagery. *Monthly Weather Review*, 103(5):420–430, 1975.
- [7] Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.
- [8] Author et. al. Unlocking GOES: A statistical framework for quantifying the evolution of convective structure in tropical cyclones. *Journal of Applied Meteorology and Climatology*, 2020.
- [9] Jean-Yves Franceschi, Edouard Delasalles, Mickaël Chen, Sylvain Lamprier, and Patrick Gallinari. Stochastic latent residual video prediction. Proceedings of the International Conference on Machine Learning, 2020.
- [10] Sophie Giffard-Roisin, Mo Yang, Guillaume Charpiat, Christina Kumler Bonfanti, Balázs Kágl, and Claire Monteleoni. Tropical cyclone track forecasting using fused deep learning from aligned reanalysis data. *Frontiers in Big Data*, 3:1, 2020.
- [11] Trevor J Hastie and Robert J Tibshirani. *Generalized additive models*, volume 43. CRC press, 1990.
- [12] Liang Hu, Elizabeth A Ritchie, and J Scott Tyo. Short-term tropical cyclone intensity forecasting from satellite imagery based on the deviation angle variance technique. *Weather and Forecasting*, 35(1):285–298, 2020.
- [13] John Kaplan, Mark DeMaria, and John A. Knaff. A revised tropical cyclone rapid intensification index for the Atlantic and Eastern North Pacific basins. *Weather and Forecasting*, 25(1):220–241, 2010.
- [14] John Kaplan, Christopher M Rozoff, Mark DeMaria, Charles R Sampson, James P Kossin, Christopher S Velden, Joseph J Cione, Jason P Dunion, John A Knaff, Jun A Zhang, et al. Evaluating environmental impacts on tropical cyclone rapid intensification predictability utilizing statistical models. *Weather and Forecasting*, 30(5):1374–1396, 2015.
- [15] S. Kim, H. Kim, J. Lee, S. Yoon, S. E. Kahou, K. Kashinath, and M. Prabhat. Deep-hurricane-tracker: Tracking and forecasting extreme climate events. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1761–1769, 2019.
- [16] K. R. Knapp and S. L. Wilkins. Gridded satellite (GridSat) GOES and CONUS data. *Earth System Science Data*, 10(3):1417–1425, 2018.
- [17] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long- and short-term temporal patterns with deep neural networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18*, page 95–104, New York, NY, USA, 2018. Association for Computing Machinery.
- [18] Christopher W. Landsea and James L. Franklin. Atlantic hurricane database uncertainty and presentation of a new database format. *Monthly Weather Review*, 141(10):3576–3592, 2013.
- [19] Bryan Lim and Stefan Zohren. Time series forecasting with deep learning: A survey. *arXiv preprint arXiv:2004.13408*, 2020.
- [20] Timothy L Olander and Christopher S Velden. The advanced dvorak technique (adt) for estimating tropical cyclone intensity: Update and new capabilities. *Weather and Forecasting*, 34(4):905–922, 2019.
- [21] Sergiu Oprea, Pablo Martinez-Gonzalez, Alberto Garcia-Garcia, John Alejandro Castro-Vargas, Sergio Orts-Escolano, Jose Garcia-Rodriguez, and Antonis Argyros. A review on deep learning techniques for video prediction. *arXiv e-prints*, page arXiv:2004.05214, April 2020.
- [22] Yao Qin, Dongjin Song, Haifeng Cheng, Wei Cheng, Guofei Jiang, and Garrison W. Cottrell. A dual-stage attention-based recurrent neural network for time series prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI'17*, page 2627–2633. AAAI Press, 2017.

- [23] Evan Racah, Christopher Beckham, Tegan Maharaj, Samira Ebrahimi Kahou, Mr. Prabhat, and Chris Pal. Extremeweather: A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3402–3413. Curran Associates, Inc., 2017.
- [24] Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7785–7794. Curran Associates, Inc., 2018.
- [25] Pradeep Ravikumar, John Lafferty, Han Liu, and Larry Wasserman. Sparse additive models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(5):1009–1030, 2009.
- [26] Timothy J Schmit, Paul Griffith, Mathew M Gunshor, Jaime M Daniels, Steven J Goodman, and William J Lebar. A closer look at the ABI on the GOES-R series. *Bulletin of the American Meteorological Society*, 98(4):681–698, 2017.
- [27] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- [28] Kimberly M Wood and Elizabeth A Ritchie. A definition for rapid weakening of North Atlantic and Eastern North Pacific tropical cyclones. *Geophysical Research Letters*, 42(22):10–091, 2015.
- [29] Rose Yu, Stephan Zheng, Anima Anandkumar, and Yisong Yue. Long-term forecasting using higher order tensor RNNs. *arXiv e-prints*, page arXiv:1711.00073, October 2017.