

1 Structural Forecasting for Short-term Tropical Cyclone Intensity Guidance

2 Trey McNeely,^{a,d} Pavel Khokhlov, Niccolò Dalmasso,^b Kimberly M. Wood,^c and Ann B. Lee^a

3 *^a Carnegie Mellon University Department of Statistics and Data Science*

4 *^b J.P. Morgan AI Research*

5 *^c Mississippi State University Department of Geosciences*

6 *^d Microsoft AI Development Acceleration Program*

7 *Corresponding author:* Ann B. Lee, annlee@stat.cmu.edu

8 ABSTRACT: Because geostationary satellite (Geo) imagery provides a high temporal resolution
9 window into tropical cyclone (TC) behavior, we investigate the viability of its application to short-
10 term probabilistic forecasts of TC convective structure to subsequently predict TC intensity. Here,
11 we present a prototype model which is trained solely on two inputs: Geo infrared imagery leading
12 up to the synoptic time of interest and intensity estimates up to 6 hours prior to that time. To estimate
13 future TC structure, we compute cloud-top temperature radial profiles from infrared imagery and
14 then simulate the evolution of an ensemble of those profiles over the subsequent 12 hours by
15 applying a Deep Autoregressive Generative Model (PixelSNAIL). To forecast TC intensities at
16 hours 6 and 12, we input *operational* intensity estimates up to the current time (0 h) and simulated
17 future radial profiles up to +12 h into a “nowcasting” convolutional neural network. We limit our
18 inputs to demonstrate the viability of our approach and to enable quantification of value added by
19 the observed and simulated future radial profiles beyond operational intensity estimates alone. Our
20 prototype model achieves a marginally higher error than the National Hurricane Center’s official
21 forecasts despite excluding environmental factors, such as vertical wind shear and sea surface
22 temperature. We also demonstrate that it is possible to reasonably predict short-term evolution of
23 TC convective structure via radial profiles from Geo infrared imagery, resulting in interpretable
24 structural forecasts that may be valuable for TC operational guidance.

25 **SIGNIFICANCE STATEMENT:** This work presents a new method of short-term probabilistic
26 forecasting for tropical cyclone (TC) convective structure and intensity using infrared geostationary
27 satellite observations. Our prototype model's performance indicates that there is some value
28 in observed and simulated future cloud-top temperature radial profiles for short-term intensity
29 forecasting. The non-linear nature of machine learning tools can pose an interpretation challenge,
30 but structural forecasts produced by our model can be directly evaluated and thus may offer helpful
31 guidance to forecasters regarding short-term TC evolution. Since forecasters are time-limited in
32 producing each advisory package despite a growing wealth of satellite observations, a tool that
33 captures recent TC convective evolution and potential future changes may support their assessment
34 of TC behavior in crafting their forecasts.

35 **1. Introduction**

36 Tropical cyclones (TCs) are powerful, organized systems that pose a major risk to coastal
37 populations. Though many statistical models provide forecast guidance on future TC intensity
38 change (e.g., the Statistical Hurricane Intensity Prediction Scheme [SHIPS]; DeMaria and Kaplan
39 1999), direct measurement of most predictors such as relative humidity or vertical wind shear used
40 in such models is impossible due to the development of TCs over open ocean far from land-based
41 observing networks (Gray 1979). Many predictors must be inferred through a combination of
42 remote observation and dynamic models of ocean and atmospheric behavior.

43 Infrared (IR; 10.3-10.7 μ m) imagery from geostationary (Geo) satellites such as the Geostationary
44 Operational Environmental Satellites (GOES) provides one of the few regular high-resolution
45 observations of TC behavior over the open ocean with a historical record spanning decades (Knapp
46 and Wilkins 2018; Janowiak et al. 2020). Furthermore, modern Geo IR platforms such as GOES-16
47 provide observations at even greater spatial and temporal resolution (Schmit et al. 2017). Since
48 cloud-top temperature is related to cloud-top height, low IR temperatures tend to indicate higher
49 cloud tops and thus stronger convection, and convective structures are known to be related to TC
50 intensity (Dvorak 1975; Olander and Velden 2007).

51 In light of this growing record of satellite observations, a broad array of recent works have
52 explored the wealth of information contained in the *spatio-temporal* structure of Geo IR imagery.
53 The Dvorak technique and more recent Advanced Dvorak Technique (ADT) have long related Geo

54 IR imagery to TC intensity (Dvorak 1975; Olander and Velden 2007), and more recent work has
55 leveraged neural networks to improve the nowcasting accuracy of the ADT (AI enhanced Dvorak
56 Technique; Olander et al. 2021). Here, we define “nowcasting” as estimating the current TC
57 intensity based on intensity estimates up to 6 h prior and IR features up to the current time (0
58 h). Spatial analyses of IR imagery have been leveraged to improve forecasts of TC eye formation,
59 a process related to intensification (DeMaria 2015; Knaff and DeMaria 2017). The deviation
60 angle variance (DAV) technique, a measure of convective organization in IR imagery, contains
61 valuable information for short-term (≤ 24 h) TC intensity guidance (Hu et al. 2020). The shape and
62 evolution of Geo IR radial profiles is known to relate to intensity and intensity change respectively
63 (Sanabia et al. 2014; McNeely et al. 2020). In this work, we utilize the *evolution over time* of
64 radial profiles (see Figure 1) to jointly forecast short-term TC intensity and structure changes.
65 We leverage deep auto-regressive (AR) generative models to construct *interpretable and high-*
66 *resolution* structural probabilistic forecasts, which display entire functions rather than time series
67 of thresholded quantities, such as pixel counts beneath a given temperature threshold.

68 Concurrent with the rise of high-resolution Geo IR imagery is the growing application of
69 convolutional neural networks (CNNs), powerful tools for performing prediction tasks with images
70 as input. Predicting TC intensity from Geo IR data is an obvious candidate application; indeed,
71 there are dozens of such works in the machine learning literature applying CNNs to this problem,
72 including Pradhan et al. (2017); Combinido et al. (2018); Lee et al. (2019); Tian et al. (2020); Wang
73 et al. (2020); and Zhang et al. (2021). These models achieve reasonable forecast accuracy via the
74 traditional machine learning framework with a CNN taking IR imagery as input to directly predict
75 intensity by, e.g., minimizing the average squared-error loss on independent test data. Explainable
76 AI approaches may then use methods such as layer-wise relevance propagation, saliency maps, and
77 activation maps to better understand *how* the model produced its point estimate (McGovern et al.
78 2019; Ebert-Uphoff and Hilburn 2020). For an example of explainable CNN-based TC intensity
79 forecasting in the meteorological literature, see Griffin et al. (2022).

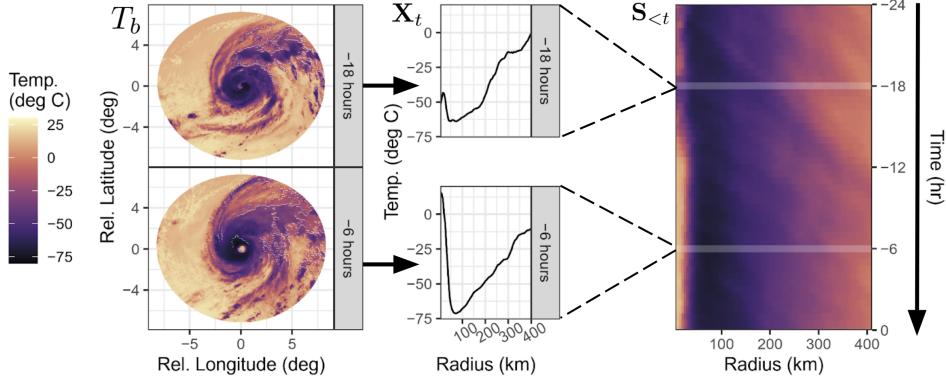
80 Our proposed pipeline takes a different approach to explainability—one which remains compat-
81 ible with the above tools for insight into the relationships leveraged by CNNs. Our approach (i)
82 utilizes a dimensionality-reducing functional transformation of IR imagery prior to analysis, and
83 (ii) provides 12-hour ensemble forecasts of TC convective structure in addition to TC intensity.

84 First, we extract *scientifically-motivated functional features*, reducing the dimension of the
85 problem (from 2-D images over time to 1-D functions over time) in a directly interpretable summary,
86 rather than directly relying on the CNN to extract salient features from (high-dimensional and low-
87 sample size) raw Geo IR imagery. These rich summary functions are derived from the ORB suite:
88 Organization (e.g., DAV as a function of radius), Radial structure (e.g., the radial profiles examined
89 in this work), and Bulk morphology (e.g., pixel counts as a function of a temperature threshold).
90 Temporal sequences of radial profiles are highly relevant to both intensity and intensity change
91 (Sanabia et al. 2014; McNeely et al. 2020, 2022). Temporal changes in these sequences of profiles
92 can be visualized via Hovmöller diagrams, which are more readily digestible by users than inferring
93 temporal patterns from animations of satellite imagery.

94 Second, we provide a probabilistic *structural forecast*, a prediction of an ensemble of possible
95 TC convective evolution, rather than directly predicting future intensity from past IR structure and
96 TC intensity. Our novel approach to intensity guidance via Geo IR imagery results in interpretable
97 intensity forecasts such as “our model predicts short-term intensification due to the potential
98 emergence of an eye-eyewall structure in the next 12 hours”. Though methods such as layer-wise
99 relevance propagation can provide further insight into the CNN’s use of structural forecasts, the IR
100 structural forecasts themselves are the core of our proposed intensity guidance pipeline.

101 Figure 2 outlines Section 3 via a schematic diagram of the structural forecasting to intensity
102 forecasting pipeline. There are three main subsections:

- 103 (a) *Structural trajectories via ORB*. First, we apply the ORB framework (McNeely et al. 2019,
104 2020) to observed IR imagery to create a “structural summary” (Figure 1) of the spatio-
105 temporal evolution of the present and recent past TC structure.
- 106 (b) *Structural forecasting with a deep autoregressive generative model*. Next, we propagate the
107 observed IR structure up to 12 hours forward in time via a deep pixel-autoregressive model,
108 which stochastically simulates an ensemble of possible trajectories of IR radial profiles.
- 109 (c) *Forecasting TC intensity via convolutional neural networks*. Finally, we input the observed
110 structure, the forecasted structure, and TC intensity up to 6 hours prior to the current time into
111 a nowcasting model to estimate the current intensity; we choose CNNs because they are easy
112 to train and commonly used for image data. By filling in the missing $t+6$ hour and $t+12$ hour



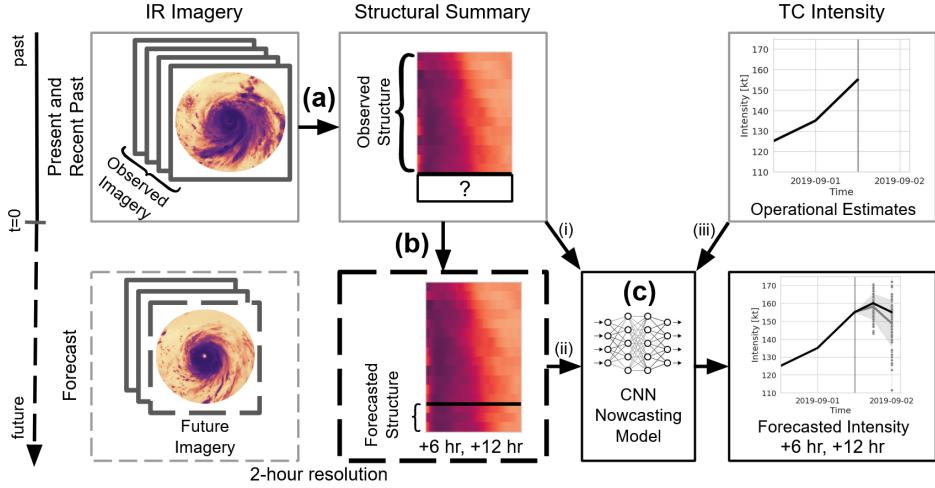
130 **FIG. 1. Evolution of TC IR Imagery as Structural Trajectories.** The raw data at each time t is a sequence of
 131 TC-centered cloud-top temperature images from GOES. We quantify the image at time t by its radial profile (X_t).
 132 The sequence of consecutive radial profiles, sampled every 2 hours, defines a *structural trajectory* or Hovmöller
 133 diagram ($S_{<t}$).

113 structure, we can then extend the nowcasting model from a nowcast for time t (i.e., hour 0) to
 114 a forecast at time $t + 6$ hours and then to time $t + 12$ hours.

115 Section 4 details the results of our prototype forecasting pipeline. The final Geo IR-based TC
 116 intensity guidance provides inherent measures of uncertainty and insight into the potential TC
 117 structural changes that influence a given forecast. The results in this work use proof-of-concept
 118 structural forecasting and a pipeline that relies solely on persistence predictors (i.e., prior intensity
 119 estimates) together with observed past and simulated future radial profiles; no environmental factors
 120 such as vertical wind shear or ocean heat content are included at this time. We demonstrate that a
 121 purely autoregressive prototype achieves a useful degree of forecasting accuracy.

122 **2. Data**

123 Our model relies on two data sources: sequences of Geo IR imagery captured by GOES satellites
 124 and past TC intensity. For training and verification (that is, model selection), we use NOAA's
 125 Hurricane Database 2 (HURDAT2; Landsea and Franklin 2013) because that database provides
 126 the post-season best estimates of TC intensities. For forecasting, we rely on *operational* TC
 127 intensity estimates, the CARQ entries from the Naval Research Laboratory's Automated Tropical
 128 Cyclone Forecast (ATCF) operational "A-deck" files (Sampson and Schrader 2000) to assess model
 129 performance under real-time conditions.



145 **FIG. 2. TC Intensity Guidance via Structural Forecasting.** Outline of the structural forecasting to intensity
 146 guidance pipeline. (a) ORB functions are used to quantify the evolution of spatio-temporal convective structure,
 147 linking IR imagery to Structural Summaries. (b) We generate structural forecasts by projecting the ORB functions
 148 into the future via a deep autoregressive model, thereby filling in the missing +6 hour and +12 hour structure. (c)
 149 A CNN nowcasting model then forecasts intensities at +6 to +12 hours from three sources of inputs: (i) observed
 150 structure, (ii) forecasted structure, and (iii) operational intensity estimates up through the current time.

134 GOES IR imagery is available through NOAA's Merged IR (MERGIR) database (Janowiak et al.
 135 2020) at 30-minute \times 4-km resolution over the North Atlantic (NAL) basin from 2000–2020. For
 136 each TC, we download \sim 2,000 km \times 2,000 km “stamps” of IR imagery centered on the TC location
 137 at a 30-minute temporal resolution. Figure 1 (left) shows two such stamps after an 800-km radius
 138 mask is applied. For this work, we sample the 30-minute data at 2-hour resolution because of
 139 periodic corruption of the imagery in the MERGIR database (Liu 2021).

140 During training, we linearly interpolate TC location and intensity from HURDAT2 to obtain
 141 locations and intensities for non-synoptic times; however, model assessment is restricted to synoptic
 142 times. We include TC lifetimes between the first synoptic time at which intensity reaches at least
 143 35 kt and the last synoptic time at which intensity is at least 35 kt; note that this can result in the
 144 inclusion of TCs $<$ 35 kt if the TC decays and then re-intensifies.

151 Finally, we rely on NHC's official forecast verification to assess our model's performance.
 152 We also draw on the SHIPS developmental database's 200-850-hPa vertical wind shear values
 153 calculated within a 200-800-km annulus from the TC center as reference during model validation
 154 due to the known impact of shear on TC convective structure (DeMaria 2018).

155 **3. Methods**

156 As outlined in Figure 2, we first construct a summary of IR structural evolution (**a**; Section
157 3a). We then train a stochastic autoregressive model, which is an explicit likelihood model (of
158 structural trajectories) that we can use to simulate probable IR structural evolution (**b**; Section 3b).
159 Finally, we combine observed and forecasted structure with operational intensity estimates up to
160 and including the current time to provide interpretable short-term intensity guidance, based solely
161 on Geo IR imagery and operational intensity estimates (**c**; Section 3c).

162 *a. Structural Trajectories via ORB*

163 Operational forecasting of TC intensity is a human-in-the-loop process and thus places a premium
164 on guidance interpretability. In this spirit, the ORB framework (Organization, Radial structure,
165 Bulk morphology) summarizes 2-D imagery via continuous 1-D functions to enable static vi-
166 sualization of spatio-temporal patterns in TC development via Hovmöller diagrams (Hovmöller
167 1949). Our past work focused on the rich quantification of spatial information in Geo IR imagery
168 (McNeely et al. 2019, 2020). More recently, we demonstrated the value of *temporal* patterns in
169 ORB functions (McNeely et al. 2022), specifically the radial profile.

170 The radial profile of brightness temperature $\overline{BT}(r) = \frac{1}{2\pi} \int_0^{2\pi} T_b(r, \theta) d\theta$ captures the structure of
171 cloud-top brightness temperature (T_b) as a function of radius r from the TC center and serves as an
172 easily interpretable description of the depth and location of convection near the TC core (Sanabia
173 et al. 2014; McNeely et al. 2020). The radial profiles are computed at 5-km resolution from 0-400
174 km ($d = 80$) (Figure 1, center); we denote the summary of convective structure at each time t by
175 \mathbf{BT}_t . The *structural trajectory* is then defined as the 24-hour sequence of present and 12 preceding
176 radial profiles at a 2-hour resolution:

$$\mathbf{S}_{<t} = (\mathbf{BT}_{t-24h}, \mathbf{BT}_{t-22h}, \dots, \mathbf{BT}_t) \quad (1)$$

177 We visualize such a trajectory with a Hovmöller diagram (see Figure 1, right).

178 McNeely et al. (2022) demonstrated a relationship between TC intensity change and Hovmöller
179 diagrams of radial profiles. However, the radial profile, if averaged over all angles, will disregard
180 asymmetry within the original 2-D images, which can degrade performance for cases affected

181 by strong vertical wind shear. In this work, we instead compute a separate radial profile for
182 each geographic quadrant (NE, NW, SE, SW) to capture asymmetries via the differences between
183 quadrants. We use geographic quadrants instead of motion-relative or shear-relative quadrants
184 because the directions of motion and shear are unstable when the magnitudes of those vectors are
185 small.

186 *b. Structural Forecasting via Deep Autoregressive Generative Model*

187 The crucial step in our guidance framework is the propagation of radial profiles into the near
188 future. The Hovmöller diagram captures the spatio-temporal evolution of the TC over an extended
189 period of time; that is, we can summarize TC development by an easily-interpretable image. By
190 treating the structural trajectory as an image, where the y-axis corresponds to the passage of time,
191 forecasting radial profiles becomes equivalent to an image completion problem. That is, we predict
192 the missing pixels at the bottom of an image (forecasted structure) given those at the top (observed
193 structure). Image completion is an active research area in machine learning; here we focus on a
194 state-of-the-art model in the class of pixel-autoregressive models (Van Oord et al. 2016).

195 Pixel-autoregressive models impose an ordering on the pixels of an image, such as a raster-scan
196 ordering (left-to-right, top-to-bottom). Let each pixel in a four-quadrant radial profile trajectory be
197 represented by $\mathbf{x}_i := (x_{i1}, x_{i2}, x_{i3}, x_{i4})$, where i is the index in the raster scan. The pixel-AR approach
198 factors the joint distribution of pixel values in the image as a product of conditionals,

$$p(\mathbf{x}_1, \dots, \mathbf{x}_n) = \prod_{i=1}^n p(\mathbf{x}_i | \mathbf{x}_{i-1}, \dots, \mathbf{x}_1), \quad (2)$$

199 where the probability of each pixel value is conditioned on all previous pixels in the raster scan.
200 Then, to generate a new radial profile, one simulates repeatedly from $p(\mathbf{x}_i | \mathbf{x}_{i-1}, \dots, \mathbf{x}_1)$; due to the
201 raster-scan ordering, the distribution of a given pixel is not influenced by elements further down
202 the sequence, hence enforcing causality.

203 The challenge of how to estimate the conditional likelihoods $p(\mathbf{x}_i | \mathbf{x}_{i-1}, \dots, \mathbf{x}_1)$ has given rise
204 to many flavors of pixel-autoregressive models, including PixelRNN (Van Oord et al. 2016),
205 PixelCNN (Van den Oord et al. 2016), PixelCNN++ (Salimans et al. 2017), and PixelSNAIL
206 (Chen et al. 2018). This work utilizes the last model, PixelSNAIL. There are two main ingredients

207 in the model: (i) causal convolution and (ii) self-attention. Causal convolution utilizes the same
 208 convolutional feature extraction outlined in Section c but masks each convolution so that each
 209 element in the raster sequence only receives information from previously generated sequences
 210 (e.g., Figure 3). Purely convolutional models, however, are restricted to small neighborhoods
 211 of pixels, leading to only a finite receptive field (area of the source image involved in a given
 212 convolution), and thus struggle with long-range dependencies in the conditional $p(\mathbf{x}_i|\mathbf{x}_{i-1}, \dots, \mathbf{x}_1)$.
 213 PixelSNAIL, on the other hand, features a self-attention mechanism that leads to unbounded
 214 receptive fields with pinpointed access to information far away in the sequence; see Chen et al.
 215 (2018) for details on the PixelSNAIL architecture.

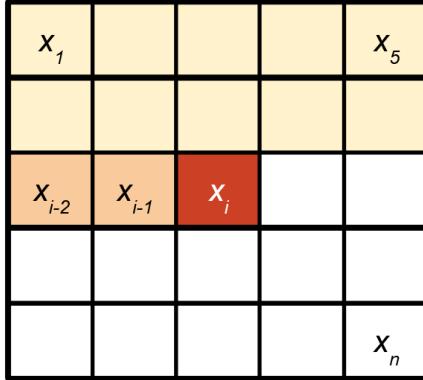
216 To ensure that cloud-top temperatures remain bounded, we rescale the data to the range $\mathbf{X} \in (0, 1)^4$
 217 and work in the logit-transformed space, $\mathbf{Z} = \log(\mathbf{X}/(1 - \mathbf{X}))$; while values of \mathbf{Z} are unbounded,
 218 \mathbf{X} remains bounded in $(0, 1)^4$ and are then transformed back to the temperature range observed
 219 in the training data. We model the density $p(\mathbf{z}_i|\mathbf{z}_{i-1}, \dots, \mathbf{z}_1)$ as 4 independent mixtures of logistic
 220 distributions, one for each quadrant. That is, for pixel i in quadrant q ,

$$p(z_{iq}|\mathbf{z}_{i-1}, \dots, \mathbf{z}_1) = \sum_{k=1}^K \pi_{qk} f(z_{iq}; \mu_{qk}, s_{qk}) \quad (\text{mixture distribution}),$$

where $f(z_{iq}; \mu_{qk}, s_{qk}) = \frac{g_{qk}(z)}{s_{qk}(1 - g_{qk}(z))^2}$, $g_{qk}(z) = \exp(-(z - \mu_{qk})/s_{qk})$,
 and $\sum_{k=1}^K \pi_{qk} = 1$.

221 Thus, with four quadrants and K mixture components, the distribution $p(\mathbf{z}_i|\mathbf{z}_{i-1}, \dots, \mathbf{z}_1)$ has $4(3K -$
 222 1) parameters: a mean μ , a scale s , and a mixture coefficient π for each quadrant, with the constraint
 223 that the mixture coefficients in each quadrant sum to one. With $K = 3$ mixture components, this
 224 results in 32 parameters total. Each draw from the distribution is transformed back to the bounded
 225 space via the relationship $\mathbf{X} = \frac{1}{1 + \exp(-\mathbf{Z})}$, then rescaled to the range of values observed in the input
 226 radial profiles.

227 This autoregressive model enables stochastic simulation of structural trajectories based on Geo
 228 IR persistence. For a given synoptic time, we can simulate many trajectories from the observed



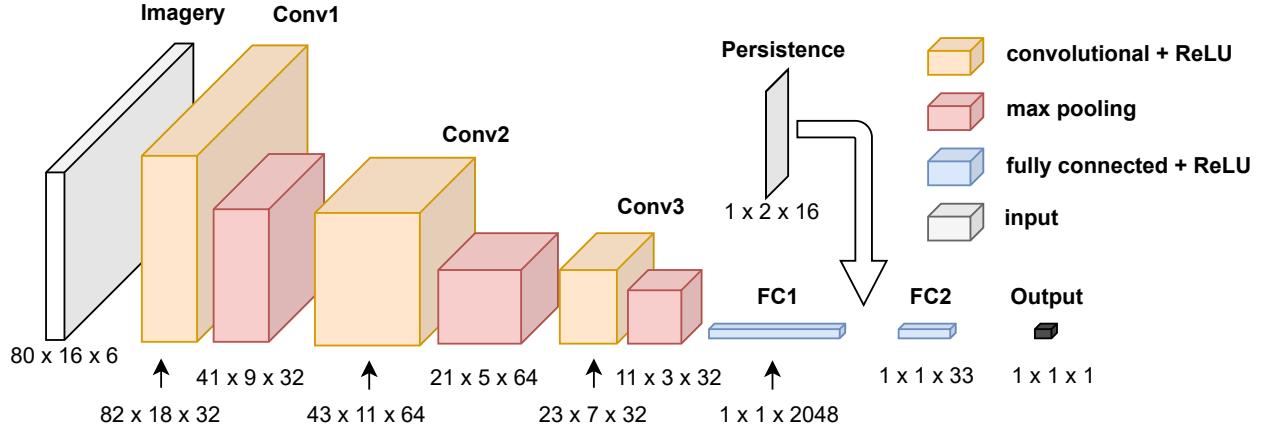
234 **FIG. 3. Masking in Pixel Autoregression.** Illustration of raster-scan ordering and the causal masking.
 235 Convolutions at index i only have access to pixel values in previous rows (earlier time points, color coded by
 236 yellow), and pixel values in the same row but to the left of pixel x_i (same time point, color coded by orange).

229 history and then feed each potential trajectory through the nowcasting model to obtain the associated
 230 intensity guidance. Via multiple simulations per forecast time, an ensemble forecast provides a
 231 measure of uncertainty in both structural trajectories and intensities while also offering insight
 232 into cases where the model over- or under-estimates intensity. For example, overestimates may be
 233 caused by too-low profile temperatures or overestimated symmetry between quadrants.

237 The structural forecasting model is trained on TCs from 2000-2012, with 2013-2020 withheld
 238 for testing. We train the model using input radial profiles calculated every 2 hours but test on
 239 synoptic times. Because AR models are likelihood-based, we can directly calculate and minimize
 240 the negative log-likelihood (NLL), a measure of the model’s ability to generalize well on withheld
 241 data.

242 *c. Nowcasting TC Intensity via Convolutional Neural Networks*

243 Traditional linear models are attractive for reasons of interpretability and good performance in
 244 low sample size settings. However, linear models often struggle to capture the complex, time-
 245 varying processes which drive TCs. It is also unclear how to include the radial profile Hovmöller
 246 diagrams as inputs to a linear model without sacrificing interpretability. In this work, we instead
 247 consider a simple convolutional neural network to map observed IR trajectories (\mathbf{S}_t) to current
 248 intensities (Y_t). Because we treat time as a spatial dimension in these diagrams and a structural
 249 trajectory is represented as an image, a CNN will leverage both spatial and temporal patterns in
 250 the data.



258 **FIG. 4. Nowcasting Model Architecture.** The convolutional neural network used for nowcasting consists
 259 of three convolution-max pool layer pairs (Conv#), fully connected layers (FC#), and a concatenation with
 260 persistence features between regressions.

251 CNNs operate by two main elements: convolutional layers and fully-connected layers. The
 252 convolutional layers first convolve each layer (here, each quadrant) with a library of filters (i.e.,
 253 matrices whose entries are learned parameters); some of these filters may resemble familiar
 254 matrices, such as gradient approximators (e.g., Sobel matrices). After each convolutional layer,
 255 the image is pooled to reduce the image size and increase the receptive field of the next set of
 256 convolutions. In the final step, the results of all convolutions are passed into a fully-connected layer
 257 which approximates the relationship between the convolutional feature map and the response.

261 We augment the traditional CNN in two ways. First, we add two layers which encode the
 262 radial and temporal location of pixels within the image; regular CNNs are translation invariant,
 263 whereas patterns in TCs have different meaning depending on their location and time of occurrence
 264 (corresponding to column versus row index, respectively, in the Hovmöller diagram). Second, we
 265 augment the model output with relevant TC persistence features (intensities and intensity changes
 266 up to 6 hours prior to time t) before passing them to a second fully-connected layer. The final
 267 model is then

$$Y_t = f_{\text{nwct}}(\mathbf{S}_{<t}, Y_{t-30h}, Y_{t-24h}, \dots, Y_{t-6h}, \Delta Y_{t-30h}, \Delta Y_{t-24h}, \dots, Y_{t-6h}) + \epsilon_t, \quad (3)$$

268 where $\mathbf{S}_{<t}$ is the 24-h structural trajectory up to current time (see Equation 1); the Y 's are intensity
269 data at 6-h resolution; the ΔY 's are intensity changes interpolated (ip) from 6-h data to a 2-h
270 resolution (e.g., $\Delta Y_{t-30h} = Y_{t-30h} - Y_{t-32h}^{ip}$), and ϵ is the prediction error.

271 Like the structural forecasting model, the nowcasting model is trained on TCs from 2000-2012,
272 here by minimizing the mean squared error. The model is trained on data with a 2-hour resolution
273 (rather than synoptic times alone) with intensities linearly interpolated to those times; we do not
274 include non-synoptic times in the test TCs (2013-2020). The details of the CNN architecture are
275 given in Figure 4.

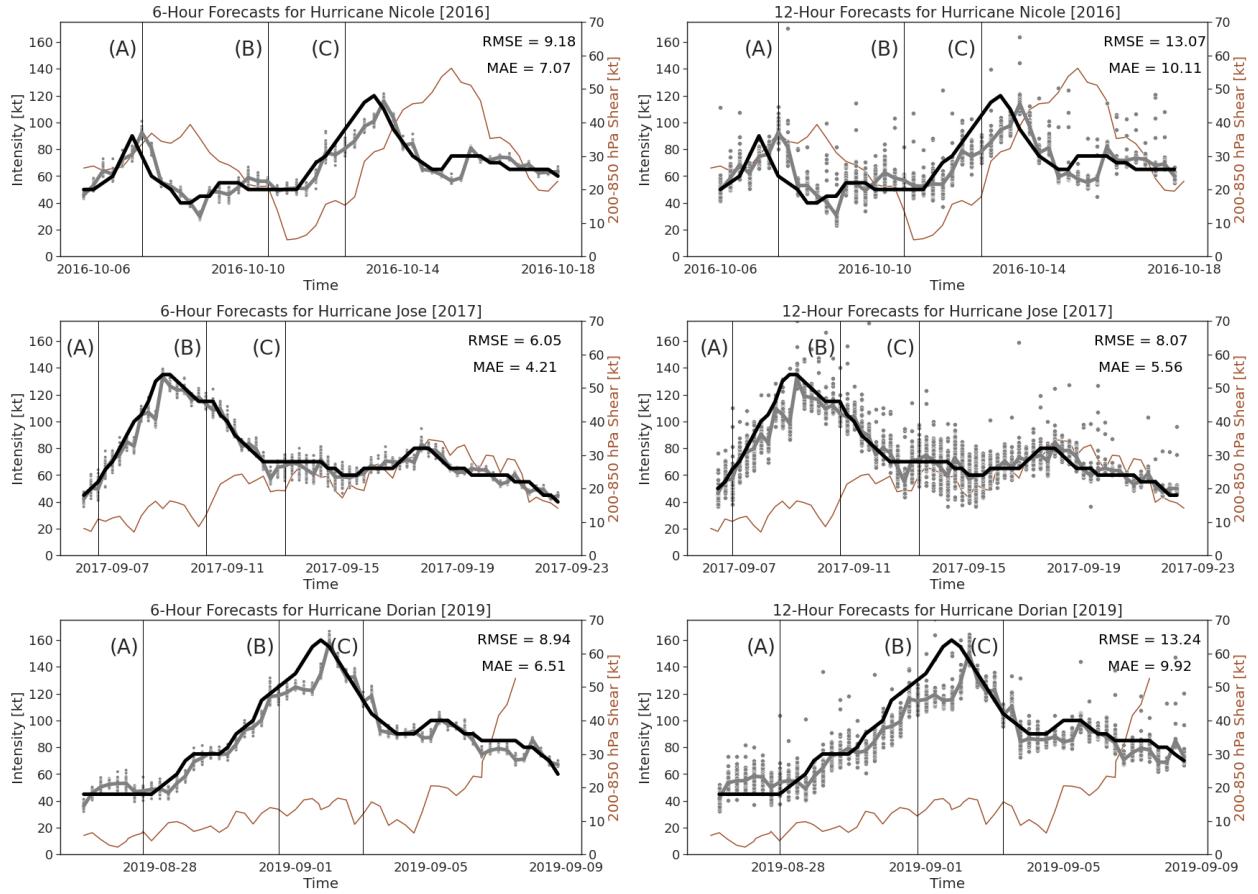
276 *d. From Nowcasting to Forecasting*

277 Section 3c defines a nowcasting model for estimating intensity at time t (i.e., hour 0) by training
278 on post-season (best-track or HURDAT2) intensities from -30 h to -6 h and imagery from -30 h
279 to 0 h. After we have trained and validated the nowcasting model to estimate 0-h intensities, we
280 apply the CNN nowcasting model to TC intensity forecasting. To forecast intensity at time $t+6$
281 h, we need the intensities at times $\leq t$ (in this work, we use operational intensities drawn from
282 CARQ in the A-deck files when generating TC intensity *forecasts*) and structural trajectories at
283 times $\leq t+6$ h (observed at times $\leq t$ and simulated at times from $t+2$ h to $t+6$ h). Using the
284 structural forecasting model in Section 3b, we simulate many possible trajectories from times $t+2$
285 hr to $t+6$ hr. Each of these possible future trajectories is then passed to the nowcasting model to
286 obtain a separate intensity forecast, giving an ensemble of possible intensities.

287 Our proposed framework for intensity forecasts at +6 and +12 h has two primary benefits: (i) by
288 providing an additional *structural* forecast, we provide insight into potential TC evolution predicted
289 by the model, such as deepening convection or the emergence of an eye; (ii) because the structural
290 forecast is stochastic, we can straightforwardly assess the uncertainty in structural evolution over
291 time and the associated uncertainty in the intensity forecasts.

292 **4. Model Results**

293 We first demonstrate the performance of our proposed model on specific cases (Hurricanes Jose
294 [2017], Nicole [2016], and Dorian [2019]) in Section a, discussing both accuracy and the insight

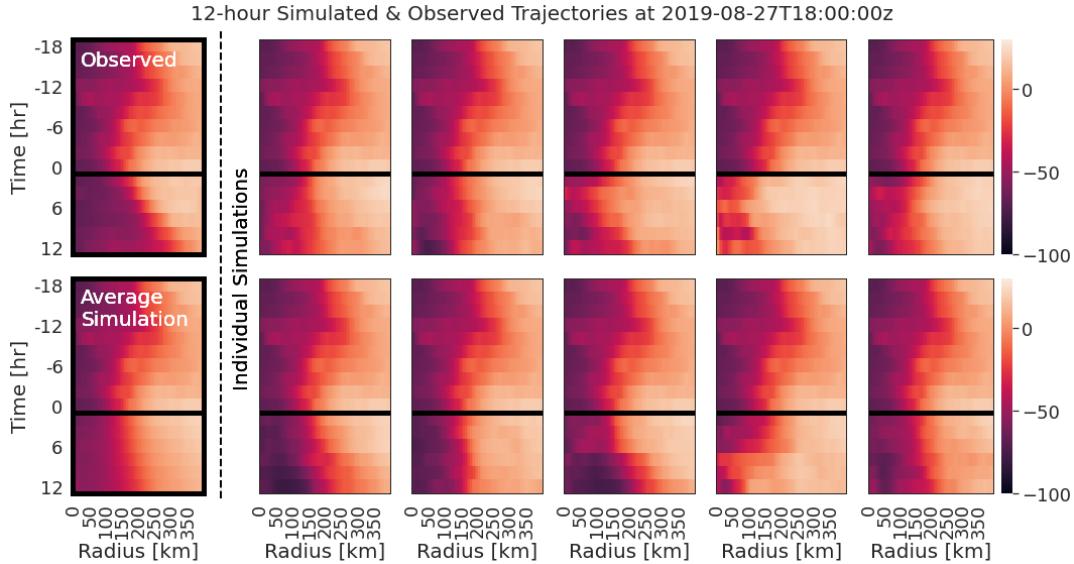


297 **FIG. 5. Case Studies:** Comparison of the 6-hour (left) and 12-hour (right) intensity forecasts (gray) from the
 298 structural forecasting model versus observed intensities (black). Vertical wind shear also shown (gold). The
 299 solid gray lines indicate the mean of 64 simulations at each time point. Reported errors indicate the error of
 300 the average model prediction, not individual simulations. The model captures the behavior of Hurricanes Jose
 301 [2017; center] and Dorian [2019; bottom] fairly well, but struggles with both rapid weakening events exhibited
 302 by Hurricane Nicole [2016; top] as well as low-intensity maintenance periods in Hurricanes Nicole and Dorian.

295 provided by structural forecasts at 6- and 12-hour lead times. We then assess the performance of
 296 the model during 2013-2020 in the North Atlantic basin at 6- and 12-hour lead times in Section b.

303 *a. Case Studies*

304 We examine Hurricane Jose (2017) due to the presence of high vertical wind shear which
 305 produces convective asymmetries not captured by the azimuthally-averaged radial profiles (i.e., not
 306 quadrant-based) of McNeely et al. (2022). Hurricane Nicole (2016) was selected due to undergoing

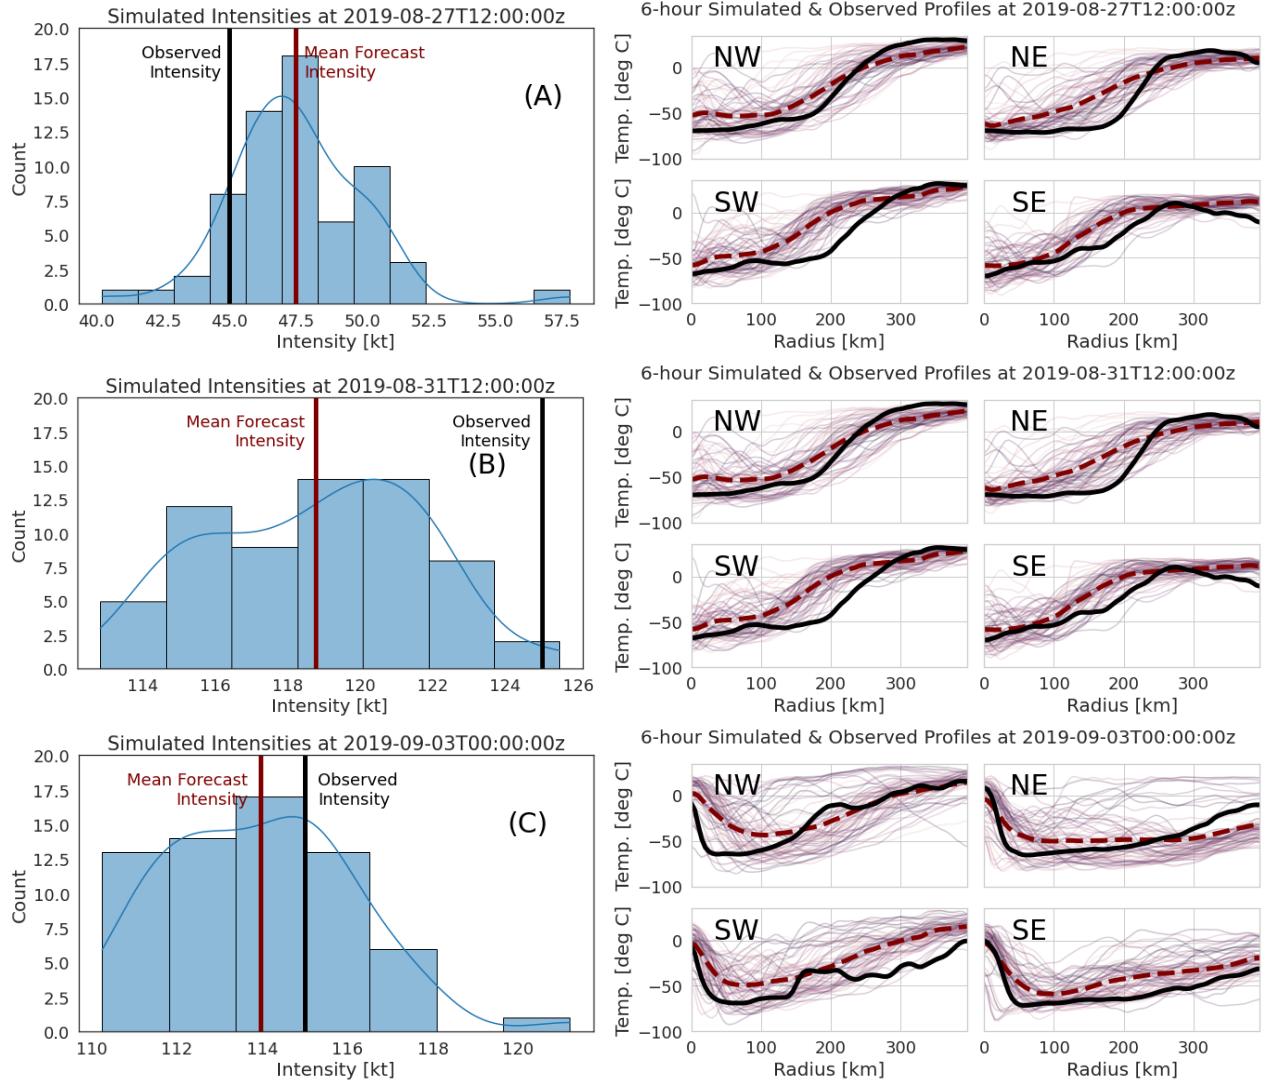


316 **FIG. 6. Observed and Simulated Trajectories, Hurricane Dorian [2019]:** The *observed* structural trajectory
 317 is shown in the top left corner. To the right, we see 10 *individual simulations* of radial profiles (averaged over
 318 all quadrants) at 12-hour lead times. All radial profiles above the black horizontal line are observed, while
 319 profiles below the black line are simulated. The bottom left corner shows the *arithmetic mean* over 64 simulated
 320 trajectories.

307 two rapid intensification and two rapid weakening events. Finally, Hurricane Dorian (2019) was a
 308 powerful TC with many *in situ* observations.

309 (i) *Intensities* Figure 5 shows the 6-hour forecasts based on 64 independently simulated struc-
 310 tural trajectories per synoptic time. Because the structural forecasts are currently based entirely
 311 on persistence—no environmental fields, such as 200-850-hPa vertical wind shear, have been
 312 included—we expect the guidance to be most useful in the short term (6- to 12-hour time frame).
 313 The steadier development of Hurricanes Jose and Dorian are well-modeled, but the swift intensity
 314 changes exhibited by Hurricane Nicole as well as the rapid intensification period of Hurricane
 315 Dorian both prove challenging to capture.

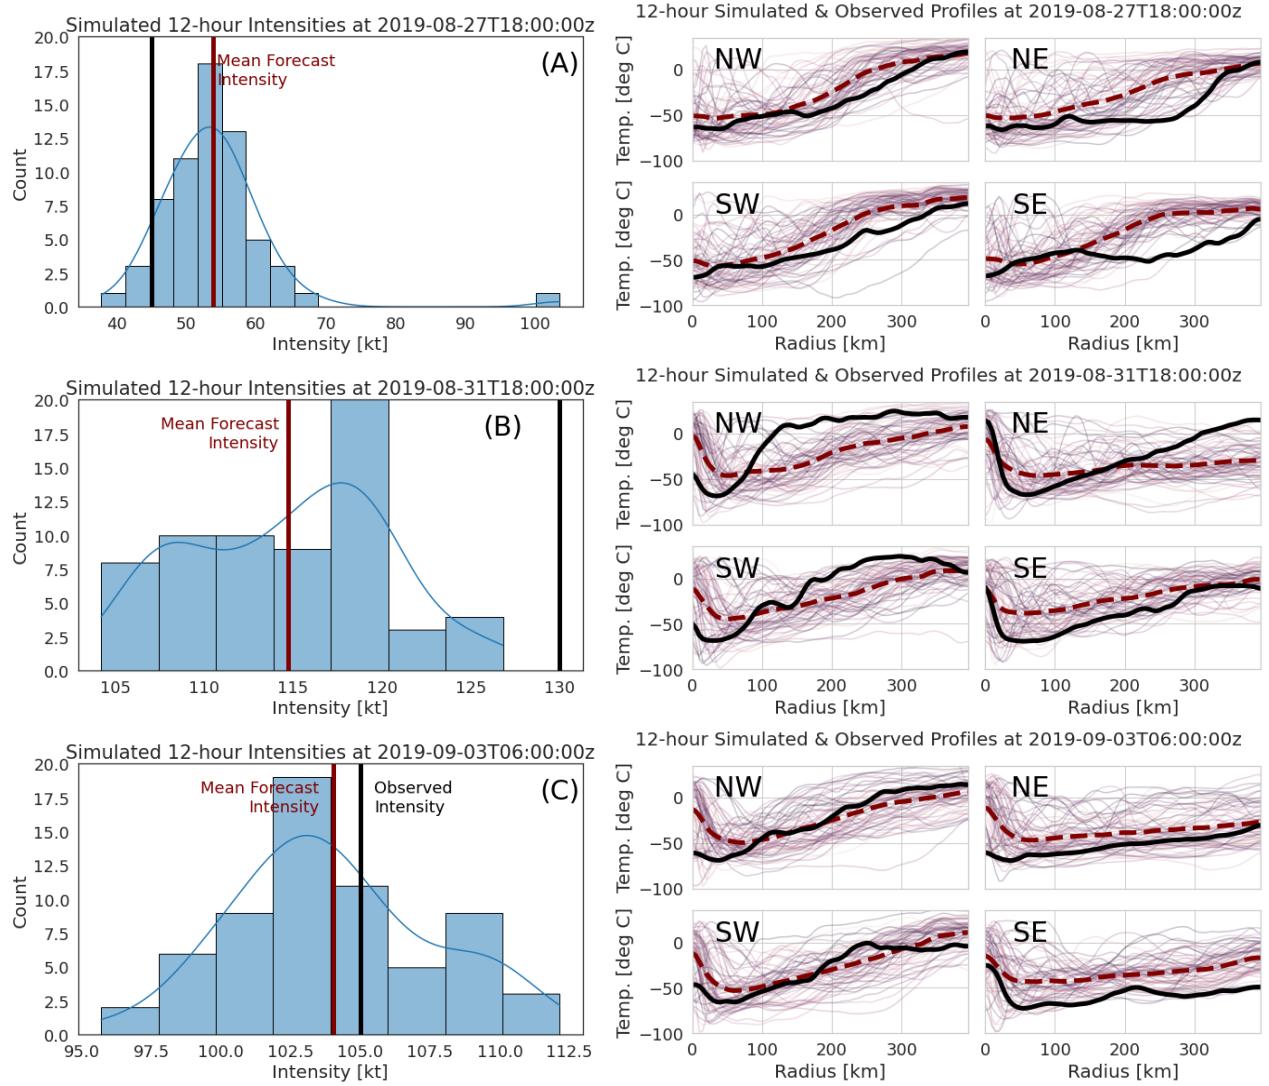
321 Extending lead time to 12 hours increases the variation among individual simulations, but the
 322 average simulated intensity continues to roughly track the observed intensities. The rapid intensity
 323 change events exhibited by Hurricane Nicole are challenging to forecast with only IR persistence.
 324 However, the model follows Hurricane Jose's evolution relatively well, indicating that the model
 325 has value as-is at 12-hour lead times.



329 **FIG. 7. Hurricane Dorian [2019] 6-h Guidance:** (Left) Distribution of forecasted intensities with observed
 330 (black) and average forecast (red) intensities marked. The distribution of the 64 intensity forecasts in the ensemble
 331 is approximated by a histogram (bar plot) and by a kernel density estimate (blue curve). (Right) Simulated profiles
 332 by quadrant with observed profiles represented by solid black curves, and averaged simulated profiles represented
 333 by dashed red curves.

326 (ii) *Diagnostics* While the end goal of intensity guidance models is ultimately prediction of TC
 327 intensity, our structural forecasting pipeline adds valuable diagnostic insight into structural factors
 328 contributing to its predictions.

340 Figure 6 demonstrates the three-step (12-hour lead time) structural forecast for Hurricane Dorian
 341 valid for 18 UTC 27 August during a period in which it maintained 45-kt intensity. The final



334 **FIG. 8. Hurricane Dorian [2019] 12-h Guidance:** Model bias during intensification (center row) is more
 335 exaggerated at 12-hour lead times. (Left) Distribution of forecasted intensities with observed (black) and average
 336 forecast (red) intensities marked. The distribution of the 64 intensity forecasts in the ensemble is approximated
 337 by a histogram (bar plot) and by a kernel density estimate (blue curve). (Right) Simulated profiles by quadrant
 338 with observed profiles represented by solid black curves, and averaged simulated profiles represented by dashed
 339 red curves.

342 6 rows of each Hovmöller diagram are simulated from the structural forecast model; in this
 343 figure, we average the four quadrants for ease of visualization (see Appendix A, Figure B10, in
 344 the supplementary materials for Hurricane Dorian structural forecasts broken down by quadrant).

345 Cloud-top temperature magnitude tends to be underestimated, but the expansion of cloud coverage
346 during this 12-h period is captured across most simulations.

347 Figure 7 demonstrates the 6-hour forecasting guidance available at individual synoptic times.
348 The average simulated profiles in each quadrant tend to track observed profiles reasonably well,
349 although they tend to predict too flat a curve and too symmetric an eye. Figure 8 shows the same
350 information but for the 12-hour lead time. Here, model biases tend to be amplified by longer lead
351 times. We note that the emergence of an eye *is* captured in trajectory (B), even 12 hours out.

352 Similar figures for Hurricanes Jose and Nicole are provided in the supplemental material. In
353 general, the structural forecast follows the observed profile, even at 12-hour lead times. We did
354 not perform any data augmentation during training (e.g., rotation) in order to preserve dominant
355 geographic patterns (e.g., the prevalence of TC convection sheared eastward and northeastward in
356 the North Atlantic), but it is possible that augmentation by rotating TCs would improve simulation
357 fidelity, as it has been shown to improve accuracy in other TC intensity forecasting applications
358 such as Griffin et al. (2022).

359 *b. Model Verification*

360 The same models are used to produce 16 simulated trajectories with associated intensity guidance
361 for each synoptic time from 2013-2020 at the 6- and 12-hour lead times. (We use 16 rather than
362 64 simulations when validating over the entire 8-year period for computational reasons.) Intensity
363 predictions provided via averaging the 16 simulations are validated against HURDAT2 best-track
364 intensities, and past TC intensity values provided as input to the model come from operational
365 estimates (CARQ) to emulate real-time performance.

366 Table 1 (left) reports the performance of the simulated trajectories versus forecast lead time in
367 terms of root mean variance (RMV), mean absolute deviation (MAD), and bias averaged over all
368 quadrants and radii. Let $\overline{BT}_{qi}(r)$ denote the i^{th} simulated profile in quadrant q and $\tau_q(r)$ denote

369 the true profile in quadrant q . Then,

$$\text{RMV} = \left(\frac{1}{400 \text{ km}} \int_0^{400} \frac{1}{4n} \sum_{q=1}^4 \sum_{i=1}^n (\overline{BT}_{qi}(r) - \tau_q(r))^2 dr \right)^{\frac{1}{2}}, \quad (4)$$

$$\text{MAD} = \frac{1}{400 \text{ km}} \int_0^{400} \frac{1}{4n} \sum_{q=1}^4 \sum_{i=1}^n |\overline{BT}_{qi}(r) - \tau_q(r)| dr, \quad (5)$$

$$\text{Bias} = \frac{1}{400 \text{ km}} \int_0^{400} \frac{1}{4n} \sum_{q=1}^4 \sum_{i=1}^n \overline{BT}_{qi}(r) - \tau_q(r) dr. \quad (6)$$

370 The above are defined for simulations at a single simulation time; to combine over multiple
 371 simulation times, we average MAD and bias, and average RMV in quadrature. The measures
 372 of noise (RMV and MAD) are large even for the shortest forecast lead times; increasing the
 373 simulation size beyond 16 will reduce the impact of this noise on the average forecast. Bias,
 374 meanwhile, becomes steadily more negative with time for our structural forecasts (top left). This
 375 will lead to overestimates of intensity, as low IR temperatures are generally associated with stronger
 376 TCs. Persistence IR forecasts (bottom left) offer a less biased IR forecast on average but higher
 377 overall errors in structure at all lead times.

378 Table 1 (right) reports verification statistics for intensity guidance using the traditional definitions
 379 for root mean squared error (RMSE), mean absolute error (MAE), and bias. As expected, the
 380 negative bias in structural forecasts manifests as a positive bias in intensity guidance.

399 Tables 2 and 3 assess the performance of our intensity guidance via structural forecasting at
 400 12-hour lead times and compare it to the NHC's official forecast verification from 2013-2019 due
 401 to availability of verification data at time of writing; note that this is a subset of the times reported
 402 in Table 1, consisting of cases where both our structural forecasts and NHC official verification
 403 are available. Overall, the RMSE of the structural forecast is about 1.1 kt larger than the NHC
 404 official forecast error as computed by RMSE, and structural forecasts produce roughly twice the
 405 bias (1.1 vs -0.6 kt). The structural forecast sees unchanged MSE with increasing 200-850-hPa
 406 vertical wind shear; the bias, however, increases with increasing wind shear (Table 2). This trend is
 407 expected, as the model does not include wind shear as a predictor but instead relies on the positive
 408 correlation between shear and asymmetry in IR imagery (as captured by radial profiles computed
 409 by quadrant). The NHC official forecast error exhibits a similar, if less pronounced, trend in bias

410 with increasing shear. The direction of shear seems more important, with both our model and the
 411 NHC official forecast performing most poorly for NW shear (6% of cases) and best for SW shear
 412 (9% of cases). The NE and SE cases dominate the overall model performance since they comprise
 413 the remaining 85% of the data set. The disparity between different shear magnitudes and directions
 414 could be alleviated in a model which utilizes environmental predictors.

415 Table 3 demonstrates similar error trends for both official forecasts and our structural forecasts.
 416 Errors tend to increase with TC intensity and with rate of intensification or weakening. The
 417 structural model produces higher bias for weaker TCs and lower bias for stronger TCs. Similarly,
 418 the structural model tends to overestimate intensities during weakening and underestimate them
 419 during intensification. The model errors are comparable to NHC official forecast errors during
 420 periods of maintenance and intensification (although bias is higher); it is periods of weakening
 421 which tend to be poorly modeled by the structural forecast. We suspect that the inclusion of
 422 environmental information could improve fidelity in weakening cases; see Section 6 on Future
 423 Work Directions for a discussion of such avenues for model improvement.

432 *c. Variable Importance in Intensity Forecasts*

433 Our model results show that structural forecasts result in 6-hour and 12-hour intensity predictions
 434 of comparable accuracy to NHC official forecasts. For insight into how much our model relies

Trajectory Verification							
Model	Lead time	2 h	4 h	6 h	8 h	10 h	12 h
Structural	RMV	10.5C	14.7C	17.4C	22.2C	25.1C	26.9C
	MAD	7.1C	10.5C	12.8C	15.6C	17.6C	19.0C
	Bias	-0.6C	-2.2C	-4.0C	-6.3C	-7.9C	-9.3C
Persistence	RMV	12.0C	16.2C	19.7C	22.7C	28.2C	31.7C
	MAD	8.2C	11.1C	13.6C	15.8C	18.5C	20.5C
	Bias	-0.4C	0.6C	2.0C	3.5C	5.5C	6.8C

Intensity Verification		
Lead time	6 h	12 h
RMSE	4.9kt	9.5kt
MAE	3.5kt	7.1kt
Bias	0.6kt	2.3kt

381 **TABLE 1. Overall Model Verification at $n = 16$:** (Left) Trajectory verification of structural forecasts, compared
 382 to IR persistence forecasts where the radial profiles are fixed at their 0 h values. Simulation noise (root mean
 383 variance and mean absolute deviation) grows rapidly in the first 6 hours; bias increases in magnitude steadily.
 384 We note that persistence offers a less biased IR forecast on average, but higher overall errors in structure at all
 385 lead times. (Right) Intensity verification vs HURDAT2 best-track intensities from 2013-2020 at each lead time.

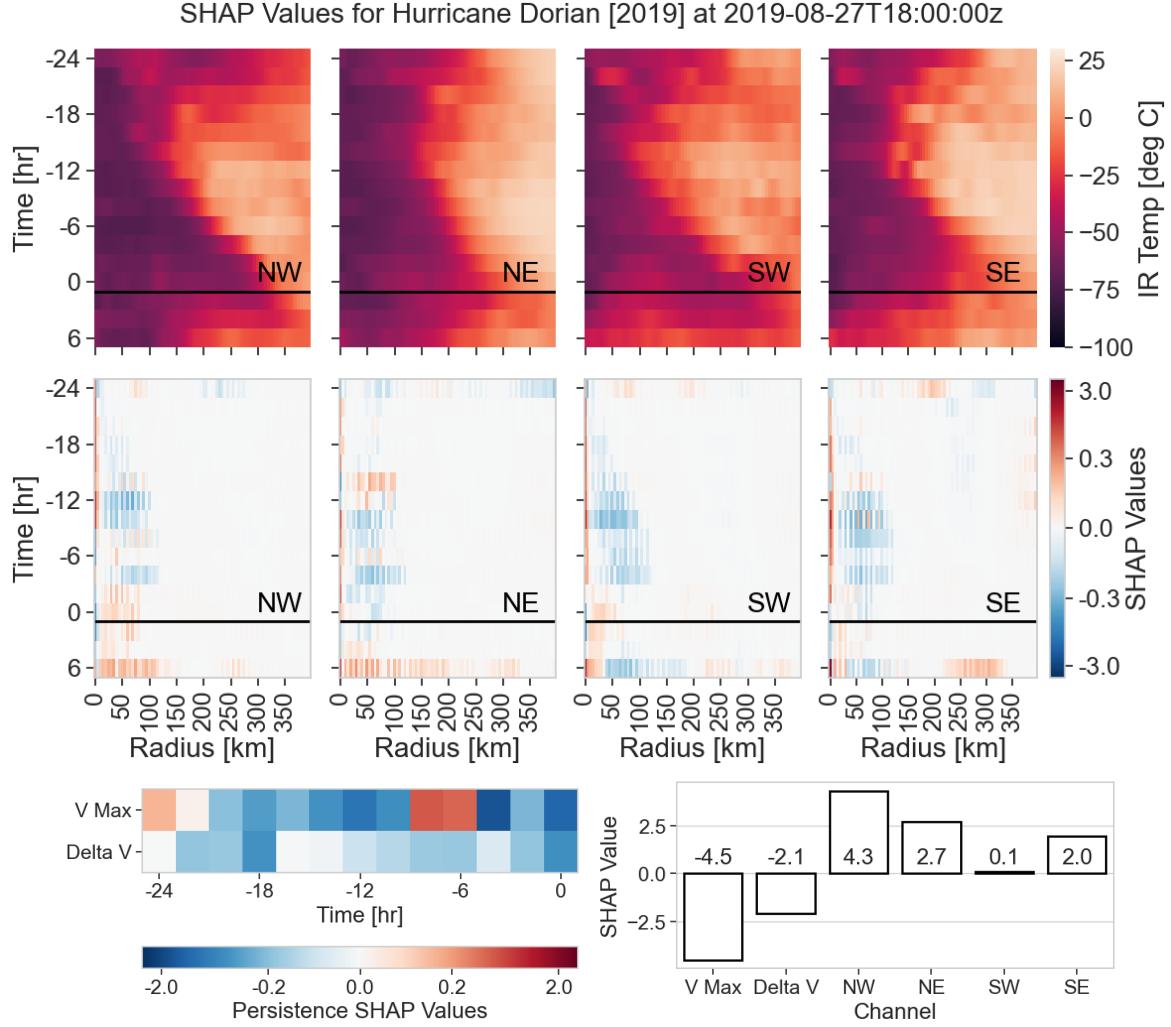
Shear Magnitude				Shear Direction			
Shear	Structural	NHC OFCL	N	Shear Direction	Structural	NHC OFCL	N
	12-h RMSE/MAE/Bias				12-h RMSE/MAE/Bias		
0-10kt	9.3/7.3/-0.7 kt	9.3/6.4/-1.1 kt	243	SW	8.6/6.1/-0.1 kt	7.5/5.4/-1.1 kt	106
10-20kt	9.7/6.9/1.0 kt	8.1/5.7/-0.5 kt	509	SE	9.1/6.9/0.3 kt	8.5/5.9/-0.5 kt	440
20+kt	8.6/6.3/2.1 kt	7.5/5.1/-0.6 kt	439	NE	9.3/6.6/2.2 kt	7.9/5.4/-0.6 kt	575
Total	9.2/6.8/1.1 kt	8.1/5.6/-0.6 kt	1,191	NW	10.6/8.2/-2.0 kt	8.9/6.1/-0.9 kt	70
				Total	9.2/6.8/1.1 kt	8.1/5.6/-0.6 kt	1,191

386 TABLE 2. **Intensity Guidance Verification Relative to Shear:** Model verification binned by 200-850-hPa
 387 vertical wind shear, reported as RMSE/MAE/Bias. (Left) The performance of the structural forecasting model
 388 does not change meaningfully relative to wind shear magnitude, while the NHC official forecast performs better
 389 in higher shear environments. The structural forecast has comparable performance to the NHC official forecasts
 390 in low-shear environments. (Right) The performance of the structural model does vary with shear direction.
 391 Both the NHC forecasts and the structural model produce higher errors for NW shear (6% of cases).

TC Intensity				Intensity Change			
Category	Structural	NHC OFCL	N	Evolution	Structural	NHC OFCL	N
	12-h RMSE/MAE/Bias				12-h RMSE/MAE/Bias		
Tropical Depression	5.5/4.4/3.0 kt	5.9/4.2/-3.1 kt	112	Weakening	12.1/8.9/7.7 kt	7.8/5.4/2.9 kt	263
Tropical Storm	7.2/5.6/2.8 kt	6.4/4.4/-0.6 kt	567	Maintenance	6.7/5.3/3.3 kt	6.1/4.3/-0.0 kt	497
Hurricane	10.3/7.7/0.6 kt	9.3/6.7/-0.6 kt	355	Intensifying	9.8/7.3/-5.5 kt	10.2/7.2/-3.5 kt	431
Major Hurricane	14.0/10.7/-5.5 kt	11.8/8.4/0.9 kt	157				
Total	9.2/6.8/1.1 kt	8.1/5.6/-0.6 kt	1,208	Total	9.2/6.8/1.1 kt	8.1/5.6/-0.6 kt	1,208

392 TABLE 3. **Intensity Guidance Verification by TC Intensity:** Model verification split out by intensity and
 393 intensity change, reported as RMSE/MAE/Bias. (Left) Both the structural and NHC official forecasts struggle
 394 more with intense storms, which are rarer. The structural forecast has much stronger bias, which is expected due
 395 to the heavy influence of persistence features in the absence of environmental predictors. (Right) Similarly, both
 396 forecasts perform best during maintenance periods (6-hour change ≤ 5 kt in magnitude), overestimate during
 397 weakening, and underestimate during intensification. The bias is more pronounced in the structural forecast due
 398 to the absence of environmental predictors.

435 on IR inputs and prior intensities when making predictions, we compute a saliency map (also
 436 known as pixel attribution) for each input. There are varied definitions for saliency, including
 437 occlusion-based approaches such as SHAP explainability values (Lundberg and Lee 2017), LIME
 438 values (Ribeiro et al. 2016), and gradient-based approaches.



424 **FIG. 9. SHAP Variable Importance Map for 6-Hour Intensity Forecast:** (Top) IR channels of CNN
 425 nowcasting model for Hurricane Dorian [2019] with observed IR structure above the horizontal black lines
 426 and +6 h forecasted structure from deep auto-regressive generative model below the horizontal lines. (Center)
 427 Pixel-wise SHAP variable importance of IR inputs on the 6-h intensity forecast. (Bottom Left) SHAP variable
 428 importance of VMax (linearly interpolated operational intensity estimates) and Delta V (2-h rate of change of
 429 operational intensity estimates) on the 6-h intensity forecast. (Bottom Right) Aggregated SHAP values over each
 430 channel, indicating IR features contributing to intensity forecasts to a degree comparable to persistence features.

431

439 Figure 9 (center) shows a map of the SHAP importance or contribution of each pixel of the IR
 440 observed and forecasted imagery on the 6-hour intensity forecast for Hurricane Dorian [2019].
 441 The bottom-left panel shows the SHAP values for prior intensity and prior intensity change. The

442 bottom-right panel shows aggregated SHAP values for each input channel. From this result and
443 a similar analysis with SHAP variable importance maps for Hurricane Jose [2017] and Hurricane
444 Nicole [2016] in Appendix A and gradient-based saliency maps in Appendix B of Supplementary
445 Materials, we conclude that (i) IR imagery contributes to the intensity forecasts to a degree
446 comparable to persistence features, (ii) forecasted infrared imagery from our deep autoregressive
447 generative model plays a more important role than observed past imagery in the TC intensity
448 forecasts, (iii) the current and past presence/absence of an eye is generally the key feature of
449 a storm, and (iv) the core temperatures outside of the eye play a significant role for intensity
450 forecasting.

451 **5. Discussion and Conclusion**

452 This paper demonstrates a novel interpretable approach to short-term TC intensity guidance
453 trained solely on intensity estimates up to 6 hours prior to the current time and IR observations
454 up to 0 h. We specifically leverage spatial characteristics of TC convection as captured by radial
455 IR profiles. By forecasting an ensemble of +6 h and +12 h trajectories of TC IR structure with
456 radial profiles computed over four geographic quadrants, we obtain reasonable estimates of future
457 +6 h and +12 h TC intensity while simultaneously capturing and enabling visualization of signals
458 in convective structure relevant to those future intensities. We focus on interpretable, physically-
459 based factors to facilitate understanding of the model’s performance (e.g., upcoming intensification
460 corresponds with decreasing cloud-top temperatures in the structural forecast). The approach
461 outlined here has the potential for further improvement by adopting other network architectures
462 for structural forecasts and by including environmental predictors provided in real time by SHIPS
463 guidance. Though testing on years of cases takes time, an individual forecast for a single TC can
464 be obtained in minutes on a single GPU, indicating the potential for the eventual use of this model
465 as part of the available TC guidance suite in an operational setting.

466 **6. Future Work Directions**

467 *a. Improving the Network Architecture for Structural Forecasts*

468 The PixelSNAIL approach provides reasonable simulations of TC IR structural evolution up to
469 12-hour lead times. However, there exists a wealth of alternate deep autoregressive generative

470 models, each of which can be designed and trained in innumerable ways. Likewise, deep autore-
471 gressive models are not the only generative models available. Simulation could be carried out
472 via vector autoregression on a low-dimensional projection of profiles (e.g., principal component
473 analysis, Fourier bases, etc.), generative adversarial networks (GANs; Creswell et al. 2018), or
474 transformers (e.g., temporal fusion transformers for multihorizon forecasting (Lim et al. 2021) and
475 spatiotemporal transformers (Grigsby et al. 2021)). The PixelSNAIL architecture was chosen to
476 demonstrate the value and feasibility of structural forecasting for intensity guidance.

477 *b. Calibrating the Probability Distribution of Structural Forecasts*

478 Our structural forecasts are probabilistic in nature, taking the form of *probability distributions*
479 over future structural trajectories $\mathbf{S}_{>t}$. In the current work, we apply a standard machine learning
480 approach of fitting a model by minimizing a loss function (in this case the negative log likelihood).
481 A good probabilistic forecast, however, should be *conditionally calibrated*. That is, the probability
482 of a particular event (in our case, specific radial profiles 6-12 hours into the future), given or
483 “conditional on” a particular history of evolution and other predictors, should match the predicted
484 probability of the same event. This is essentially saying that draws from the forecasting model
485 should be indistinguishable from actual observations, if all relevant conditions are the same. Dey
486 et al. (2022) recently proposed a new method for adjusting or “recalibrating” probabilistic forecasts,
487 so that they will have this property. Indeed, one can potentially apply their procedure sequentially to
488 each autoregressive component $p(Z_{iq}|\mathbf{Z}_{i-1}, \dots, \mathbf{Z}_1)$, for pixel $i = 1, \dots, n$, and quadrant $q = 1, 2, 3, 4$,
489 so as to obtain a conditionally calibrated density over structural trajectories $\mathbf{S}_{>t}$ given present and
490 past observations; see Discussion in Dey et al. (2022).

491 *c. Inclusion of Environmental Variables*

492 The PixelSNAIL model presented here is a purely autoregressive process; that is, it simulates
493 future structural features using only past IR imagery as an input. The inclusion of environmental
494 variables known to impact TCs such as vertical wind shear, atmospheric moisture, or sea surface
495 temperature may improve the accuracy of the forward simulation of radial profiles, particularly of
496 structural evolution beyond 12 hours. Such factors can be added to the PixelSNAIL architecture as
497 additional input layers via values provided by SHIPS which are not forecasted by the model. These

498 inputs would then serve as “guiderails” for simulated structural evolution with potential to better
499 capture the effects of such factors on profile asymmetry. Despite these limitations, our prototype
500 model (which is derived solely from prior and present TC intensity estimates and Geo IR imagery
501 alongside forecasted TC structure using a very simple network architecture) provides reasonable
502 short-term structural and intensity forecasts comparable to NHC forecasts at 6- and 12-h lead times.
503 The inclusion of environmental variables in the nowcasting model is likely to improve its intensity
504 forecasts, which would then be compared to SHIPS forecasts as well as NHC official forecasts, the
505 latter of which are crafted using SHIPS and other guidance.

506 *Acknowledgments.* Part of this research was done as an Independent Study in Spring 2021 while
507 Pavel Khokhlov was a Master in Machine Learning student at Carnegie Mellon University. We are
508 grateful to Microsoft for providing Azure computing resources for this work. The authors would
509 like to thank Katerina Fragkiadaki for a discussion on deep generative networks, and Galen Vincent
510 for many helpful comments on the research. This work is supported in part by NSF DMS-2053804,
511 NSF PHY-2020295, and the C3.ai Digital Transformation Institute.

512 *Data availability statement.* Code to generate Geo IR radial profiles from openly available data
513 can be found at <https://github.com/ihmcneely/ORB2sample>, which draws from the MER-
514 GIR database openly available from NASA at https://disc.gsfc.nasa.gov/datasets/GPM_
515 MERGIR_1/summary. The HURDAT2 best track database is openly available from the NHC at
516 <https://www.nhc.noaa.gov/data>, while the ATCF operational best tracks (B-deck) are openly
517 available from the National Center for Atmospheric Research at <http://hurricanes.ral.ucar.edu/repository/>. Official forecast verification files are openly available from the NHC at
518 <https://www.nhc.noaa.gov/verification/>. Finally, SHIPS developmental data is openly
519 available from the Cooperative Institute for Research in the Atmosphere at https://rammb.cira.colostate.edu/research/tropical_cyclones/ships/developmental_data.asp.

522 **References**

523 Chen, X., N. Mishra, M. Rohaninejad, and P. Abbeel, 2018: Pixelsnail: An improved autoregressive
524 generative model. *International Conference on Machine Learning*, PMLR, 864–872.

- 525 Combinido, J. S., J. R. Mendoza, and J. Aborot, 2018: A convolutional neural network approach
526 for estimating tropical cyclone intensity using satellite-based infrared images. *2018 24th Inter-*
527 *national Conference on Pattern Recognition (ICPR)*, IEEE, 1474–1480.
- 528 Creswell, A., T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, 2018:
529 Generative adversarial networks: An overview. *IEEE Signal Processing Magazine*, **35** (1),
530 53–65.
- 531 DeMaria, M., 2018: SHIPS developmental database file format and predictor descriptions (2018).
532 Tech. rep., Colorado State University. URL http://rammb.cira.colostate.edu/research/tropical_cyclones/ships/docs/ships_predictor_file_2018.doc, Developmental Data.
- 533
- 534 DeMaria, M., and J. Kaplan, 1999: An updated statistical hurricane intensity prediction scheme
535 (SHIPS) for the Atlantic and Eastern North Pacific basins. *Weather and Forecasting*, **14** (3),
536 326–337.
- 537 DeMaria, R., 2015: Automated tropical cyclone eye detection using discriminant analysis. Ph.D.
538 thesis, Colorado State University.
- 539 Dey, B., D. Zhao, J. A. Newman, B. H. Andrews, R. Izbicki, and A. B. Lee, 2022: Calibrated pre-
540 dictive distributions via diagnostics for conditional coverage. *arXiv preprint arXiv:2205.14568*.
- 541 Dvorak, V. F., 1975: Tropical cyclone intensity analysis and forecasting from satellite im-
542 agery. *Monthly Weather Review*, **103** (5), 420–430, [https://doi.org/10.1175/1520-0493\(1975\)103<0420:TCIAAF>2.0.CO;2](https://doi.org/10.1175/1520-0493(1975)103<0420:TCIAAF>2.0.CO;2).
- 543
- 544 Ebert-Uphoff, I., and K. Hilburn, 2020: Evaluation, tuning, and interpretation of neural networks
545 for working with images in meteorological applications. *Bulletin of the American Meteorological
546 Society*, **101** (12), E2149–E2170.
- 547 Gray, W. M., 1979: Hurricanes: Their formation, structure and likely role in the tropical circulation.
548 *Meteorology over the tropical oceans*, **155**, 218.
- 549 Griffin, S. M., A. Wimmers, and C. S. Velden, 2022: Predicting rapid intensification in north atlantic
550 and eastern north pacific tropical cyclones using a convolutional neural network. *Weather and
551 Forecasting*.

- 552 Grigsby, J., Z. Wang, and Y. Qi, 2021: Long-range transformers for dynamic spatiotemporal
553 forecasting. *arXiv preprint arXiv:2109.12218*.
- 554 Hovmöller, E., 1949: The trough-and-ridge diagram. *Tellus*, **1** (2), 62–66.
- 555 Hu, L., E. A. Ritchie, and J. S. Tyo, 2020: Short-term tropical cyclone intensity forecasting from
556 satellite imagery based on the deviation angle variance technique. *Weather and Forecasting*,
557 **35** (1), 285–298.
- 558 Janowiak, J., B. Joyce, and P. Xie, 2020: NCEP/CPC L3 half hourly 4km global (60S
559 - 60N) merged IR v1. NASA Goddard Earth Sciences Data and Information Services
560 Center, URL https://disc.gsfc.nasa.gov/datacollection/GPM_MERGIR_1.html, <https://doi.org/10.5067/P4HZB9N27EKU>.
- 561
- 562 Knaff, J. A., and R. T. DeMaria, 2017: Forecasting tropical cyclone eye formation and dissipation
563 in infrared imagery. *Weather and Forecasting*, **32** (6), 2103–2116.
- 564 Knapp, K. R., and S. L. Wilkins, 2018: Gridded satellite (GridSat) GOES and CONUS data. *Earth*
565 *System Science Data*, **10** (3), 1417–1425, <https://doi.org/10.5194/essd-10-1417-2018>.
- 566 Lee, J., J. Im, D.-H. Cha, H. Park, and S. Sim, 2019: Tropical cyclone intensity estimation
567 using multi-dimensional convolutional neural networks from geostationary satellite data. *Remote*
568 *Sensing*, **12** (1), 108.
- 569 Lim, B., S. Ö. Arik, N. Loeff, and T. Pfister, 2021: Temporal fusion transformers for interpretable
570 multi-horizon time series forecasting. *International Journal of Forecasting*, **37** (4), 1748–1764.
- 571 Liu, Z., 2021: private communication.
- 572 Lundberg, S. M., and S.-I. Lee, 2017: A unified approach to interpreting model predictions.
573 *Advances in neural information processing systems*, **30**.
- 574 McGovern, A., R. Lagerquist, D. J. Gagne, G. E. Jergensen, K. L. Elmore, C. R. Homeyer, and
575 T. Smith, 2019: Making the black box more transparent: Understanding the physical implications
576 of machine learning. *Bulletin of the American Meteorological Society*, **100** (11), 2175–2199.
- 577 McNeely, T., A. B. Lee, D. Hammerling, and K. Wood, 2019: Quantifying the spatial structure of
578 tropical cyclone imagery.

- 579 McNeely, T., A. B. Lee, K. M. Wood, and D. Hammerling, 2020: Unlocking GOES: A statistical
580 framework for quantifying the evolution of convective structure in tropical cyclones. *Journal of*
581 *Applied Meteorology and Climatology*, **59** (10), 1671–1689.
- 582 McNeely, T., G. Vincent, A. B. Lee, R. Izicki, and K. M. Wood, 2022: Detecting distributional
583 differences in labeled sequence data with application to tropical cyclone satellite imagery. *arXiv*
584 *preprint arXiv:2202.02253*.
- 585 Olander, T., A. Wimmers, C. Velden, and J. P. Kossin, 2021: Investigation of machine learning
586 using satellite-based advanced dvorak technique analysis parameters to estimate tropical cyclone
587 intensity. *Weather and Forecasting*, **36** (6), 2161–2186.
- 588 Olander, T. L., and C. S. Velden, 2007: The advanced dvorak technique: Continued development of
589 an objective scheme to estimate tropical cyclone intensity using geostationary infrared satellite
590 imagery. *Weather and Forecasting*, **22** (2), 287–298.
- 591 Pradhan, R., R. S. Aygun, M. Maskey, R. Ramachandran, and D. J. Cecil, 2017: Tropical cyclone
592 intensity estimation using a deep convolutional neural network. *IEEE Transactions on Image*
593 *Processing*, **27** (2), 692–702.
- 594 Ribeiro, M. T., S. Singh, and C. Guestrin, 2016: ” why should i trust you?” explaining the
595 predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD international conference*
596 *on knowledge discovery and data mining*, 1135–1144.
- 597 Salimans, T., A. Karpathy, X. Chen, and D. P. Kingma, 2017: Pixelcnn++: Improving the
598 pixelcnn with discretized logistic mixture likelihood and other modifications. *arXiv preprint*
599 *arXiv:1701.05517*.
- 600 Sanabia, E. R., B. S. Barrett, and C. M. Fine, 2014: Relationships between tropical cy-
601 clone intensity and eyewall structure as determined by radial profiles of inner-core in-
602 frared brightness temperature. *Monthly Weather Review*, **142** (12), 4581–4599, <https://doi.org/10.1175/MWR-D-13-00336.1>.
- 604 Schmit, T. J., P. Griffith, M. M. Gunshor, J. M. Daniels, S. J. Goodman, and W. J. Lebair, 2017: A
605 closer look at the ABI on the GOES-R series. *Bulletin of the American Meteorological Society*,
606 **98** (4), 681–698.

- 607 Tian, W., W. Huang, L. Yi, L. Wu, and C. Wang, 2020: A cnn-based hybrid model for tropical
608 cyclone intensity estimation in meteorological industry. *IEEE Access*, **8**, 59 158–59 168.
- 609 Van den Oord, A., N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, and Coauthors, 2016: Con-
610 ditional image generation with pixelcnn decoders. *Advances in neural information processing*
611 *systems*, **29**.
- 612 Van Oord, A., N. Kalchbrenner, and K. Kavukcuoglu, 2016: Pixel recurrent neural networks.
613 *International conference on machine learning*, PMLR, 1747–1756.
- 614 Wang, C., Q. Xu, X. Li, and Y. Cheng, 2020: Cnn-based tropical cyclone track forecasting
615 from satellite infrared images. *IGARSS 2020-2020 IEEE International Geoscience and Remote*
616 *Sensing Symposium*, IEEE, 5811–5814.
- 617 Zhang, C.-J., X.-J. Wang, L.-M. Ma, and X.-Q. Lu, 2021: Tropical cyclone intensity classification
618 and estimation using infrared satellite images with deep learning. *IEEE Journal of Selected*
619 *Topics in Applied Earth Observations and Remote Sensing*, **14**, 2070–2086.

1 Supplemental Material for “Structural Forecasting for Short-term Tropical
2 Cyclone Intensity Guidance”

3 Trey McNeely,^{a,d} Pavel Khokhlov, Niccolò Dalmasso,^b Kimberly M. Wood,^c and Ann B. Lee^a

4 ^a *Carnegie Mellon University Department of Statistics and Data Science*

5 ^b *J.P. Morgan AI Research*

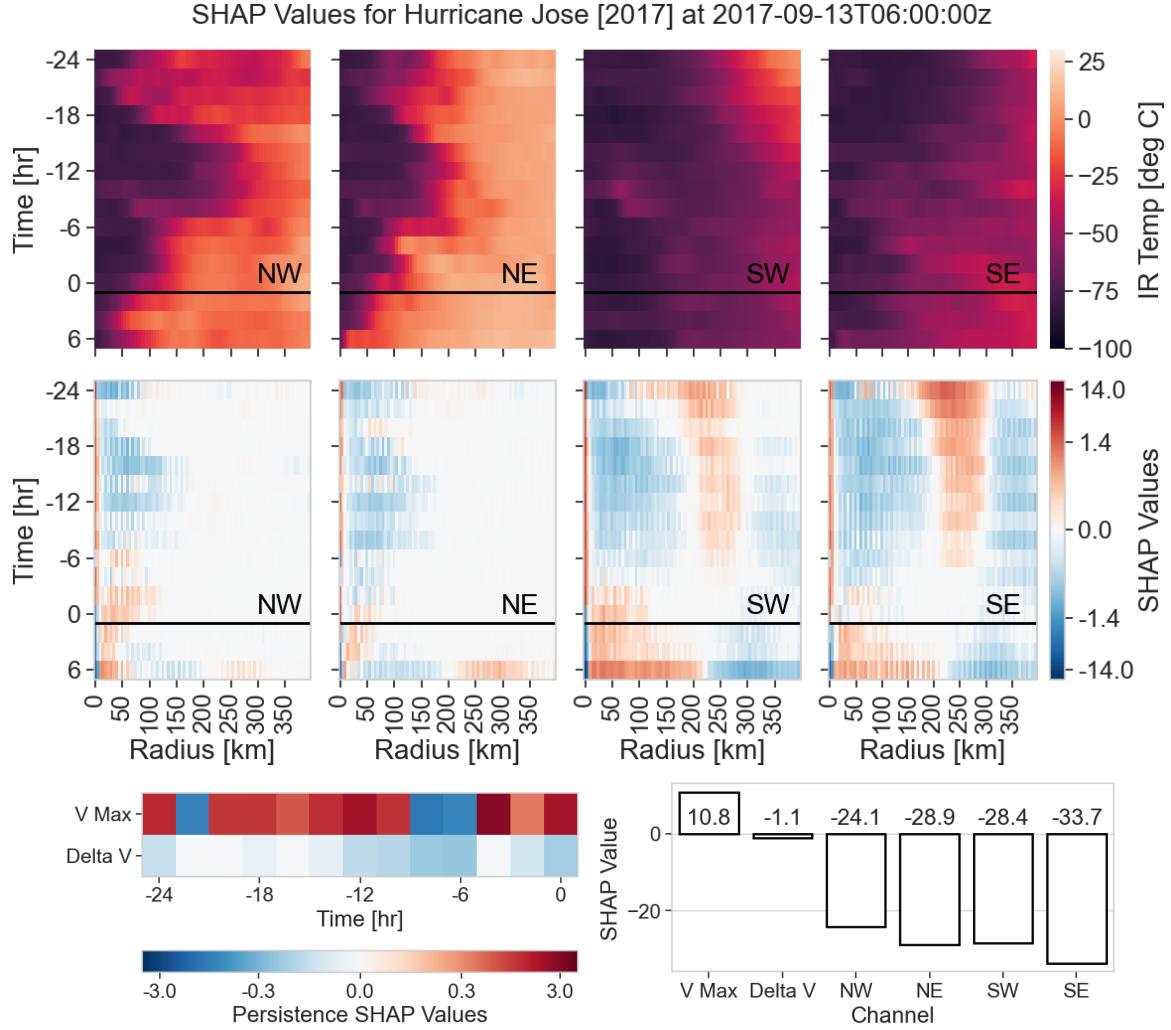
6 ^c *Mississippi State University Department of Geosciences*

7 ^d *Microsoft AI Development Acceleration Program*

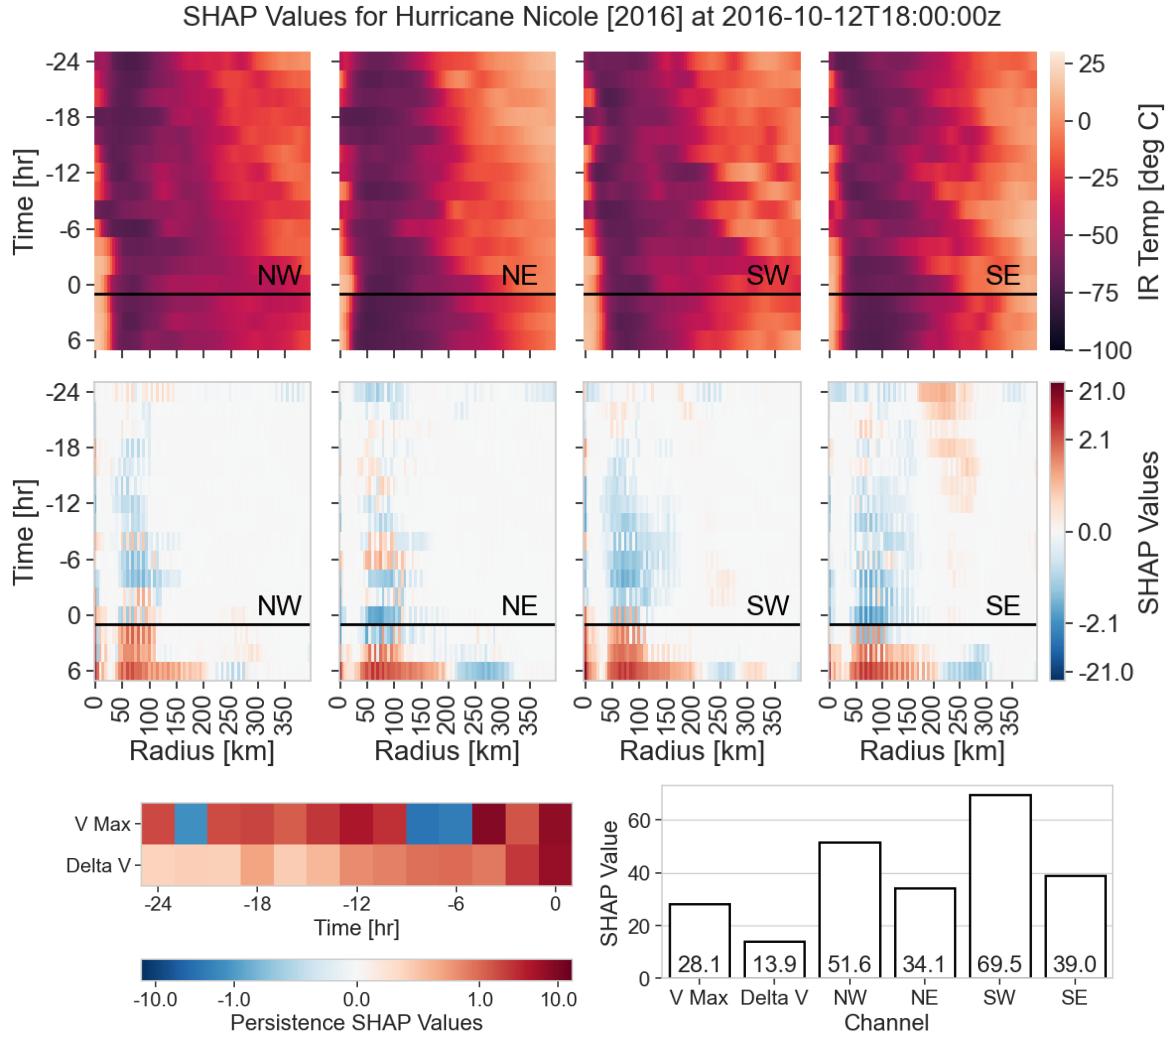
8 Corresponding author: Ann B. Lee, annlee@stat.cmu.edu

APPENDIX A

Additional SHAP Variable Importance Maps for TC Intensity Forecasts



11 **FIG. A1. SHAP Variable Importance Map for 6-Hour Intensity Forecast:** (Top) IR channels of CNN
 12 “nowcasting” model for Hurricane Jose [2017] with observed IR structure above the horizontal black lines, and
 13 +6 hour forecasted structure from deep auto-regressive generative model below the horizontal lines. (Center)
 14 Pixel-wise SHAP variable importance of IR inputs on 6-hour intensity forecast. (Bottom Left) SHAP variable
 15 importance of VMax (linearly interpolated operational intensity estimates) and Delta V (2-h rate of change of
 16 operational intensity estimates) on the 6-h intensity forecast. (Bottom Right) Aggregated SHAP values over each
 17 channel showing IR features contributing significantly to the intensity forecasts.



18 FIG. A2. **IR Channels and SHAP Variable Importance Map for 6-Hour Intensity Forecast:** As figure A1,
 19 but for Hurricane Nicole [2016].

APPENDIX B

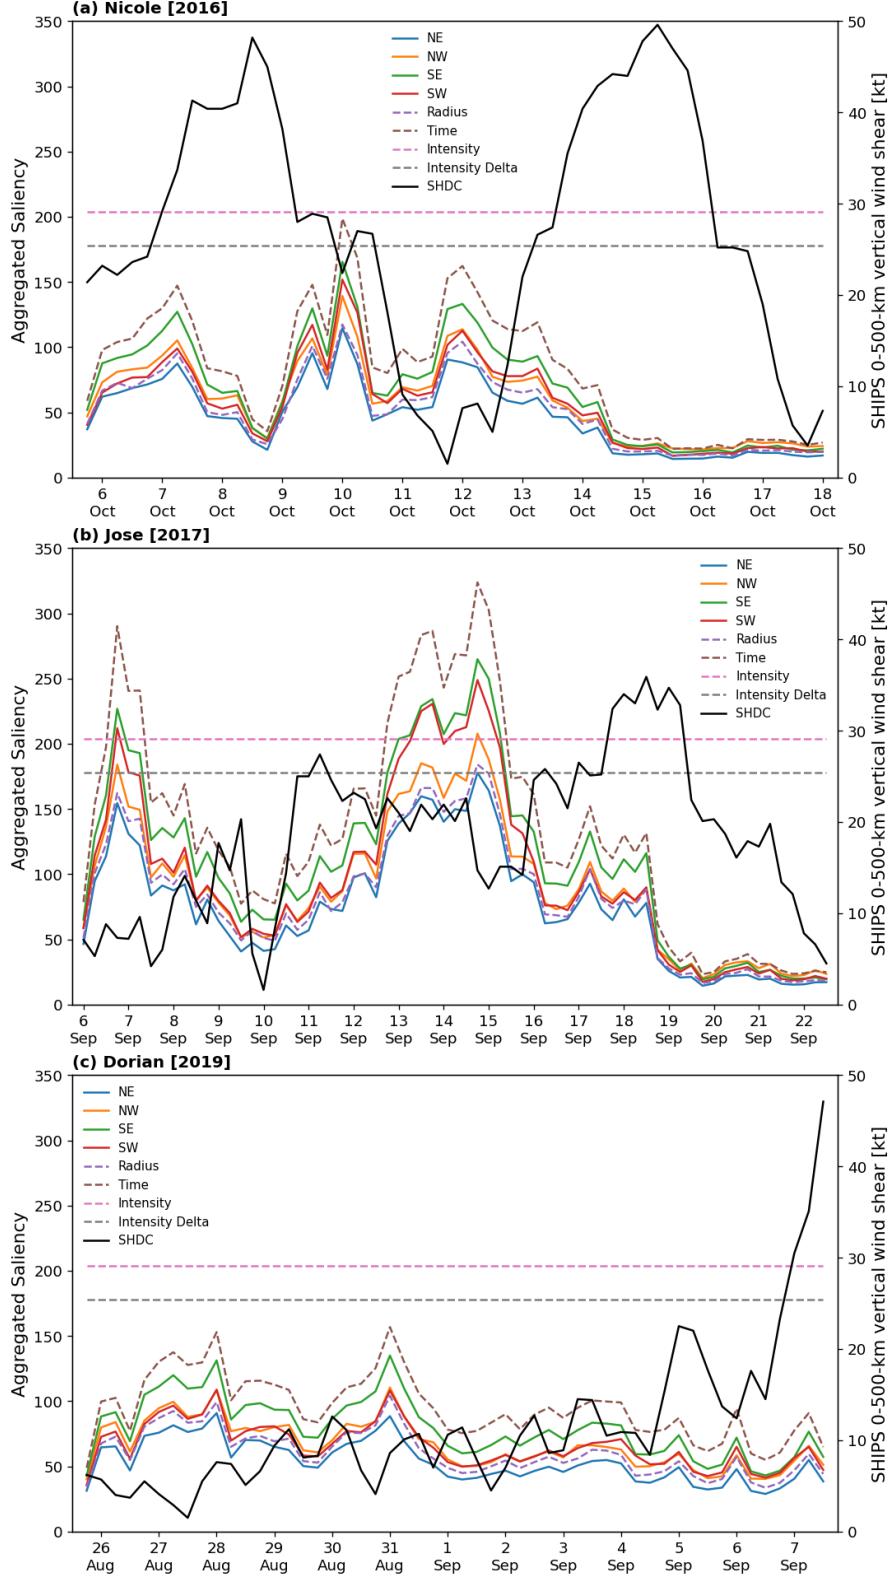
Input Saliency for Forecasting Model

22 Our model results in Section 4 showed that structural forecasts result in 12-hour intensity
 23 predictions of comparable accuracy to NHC official forecasts. Here we apply a simple gradient-
 24 based approach to provide some insight as to how much the model relies on different IR inputs
 25 when making predictions. The gradient describes how much a feature contributes to the model
 26 response Y . More specifically, we define the saliency $S_i(\mathbf{x})$ of the i^{th} pixel or feature by

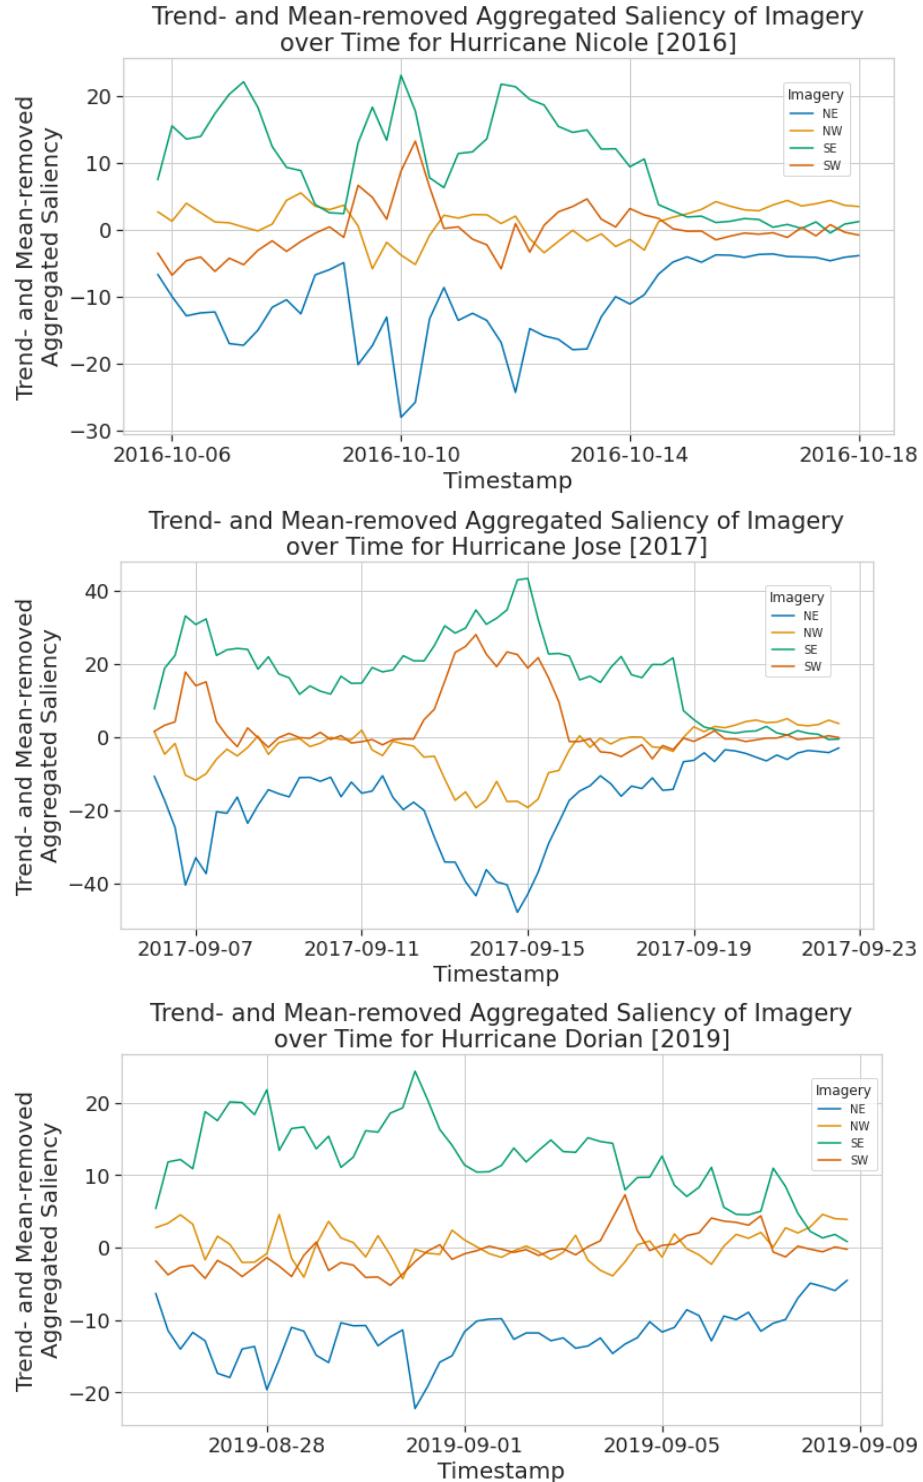
$$S_i(\mathbf{x}) = \left| \frac{\partial Y}{\partial x_i} \Big|_{\mathbf{x}} \right|, \quad (\text{B1})$$

27 where \mathbf{x} denotes the total input.

28 In order to visualize the overall impact of each input channel (four IR quadrants, the radius
 29 channel, the time channel, the observed prior intensity, and the observed prior intensity change) on
 30 the forecasted future intensity, we “aggregate” the saliency, summing over all pixels in each channel.
 31 Figure B1 shows the saliency aggregated by channel over time for each of our three example TCs.
 32 Note that because prior intensity/intensity change are not included in the convolutional layers, they
 33 are linear, and thus have a fixed saliency; because the model is nonlinear in the other channels, the
 34 saliency varies over time with the IR inputs. Of particular note is that *the aggregated saliency of*
 35 *the IR input channels is comparable to the persistence features*, indicating that the model does not
 36 simply rely on persistence to make its predictions but instead makes use of the structural forecasts.
 37 Figure B2 shows the same values for IR channels with the mean and trend removed, demonstrating
 38 that the model tends to rely more heavily on convective structure in the southern quadrants, and
 39 particularly the southeast quadrant.



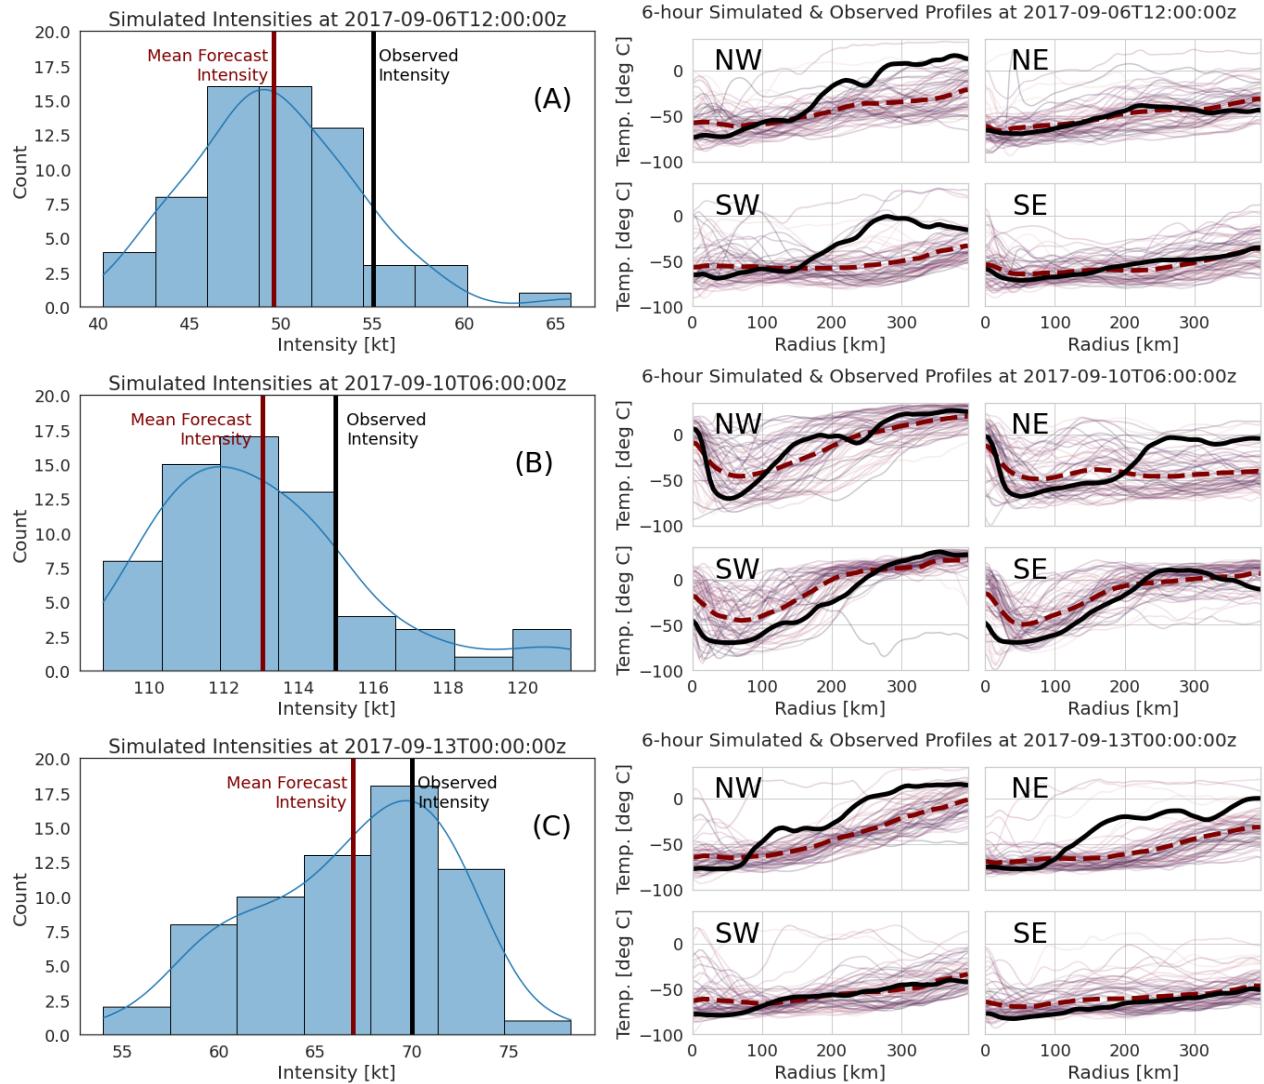
40 FIG. B1. Saliency per input layer in the model. Saliency values over time indicate that the model for TC
 41 intensity forecasts utilizes *image* inputs to a degree comparable to prior intensity values.



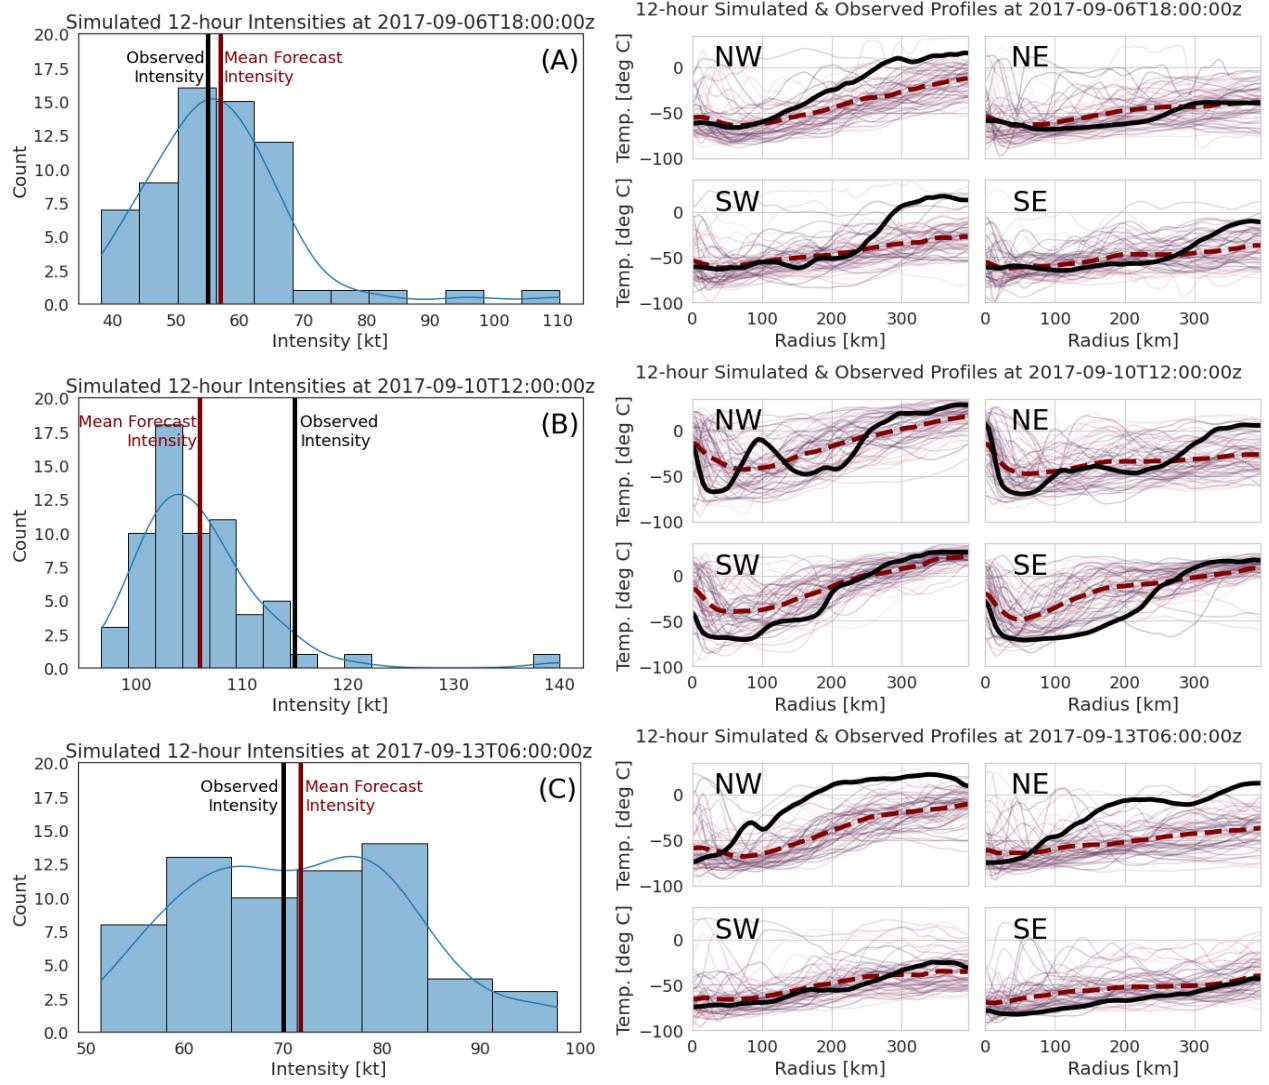
42 FIG. B2. Trend- and mean-removed saliency for image inputs. Comparing the saliency values over time
 43 indicates 1) the southern quadrants are more heavily utilized by the forecasting model, and 2) the degree to which
 44 the model relies on one quadrant over another is not constant (e.g. Hurricane Jose [2017] around Sept. 14, 2017).

APPENDIX C

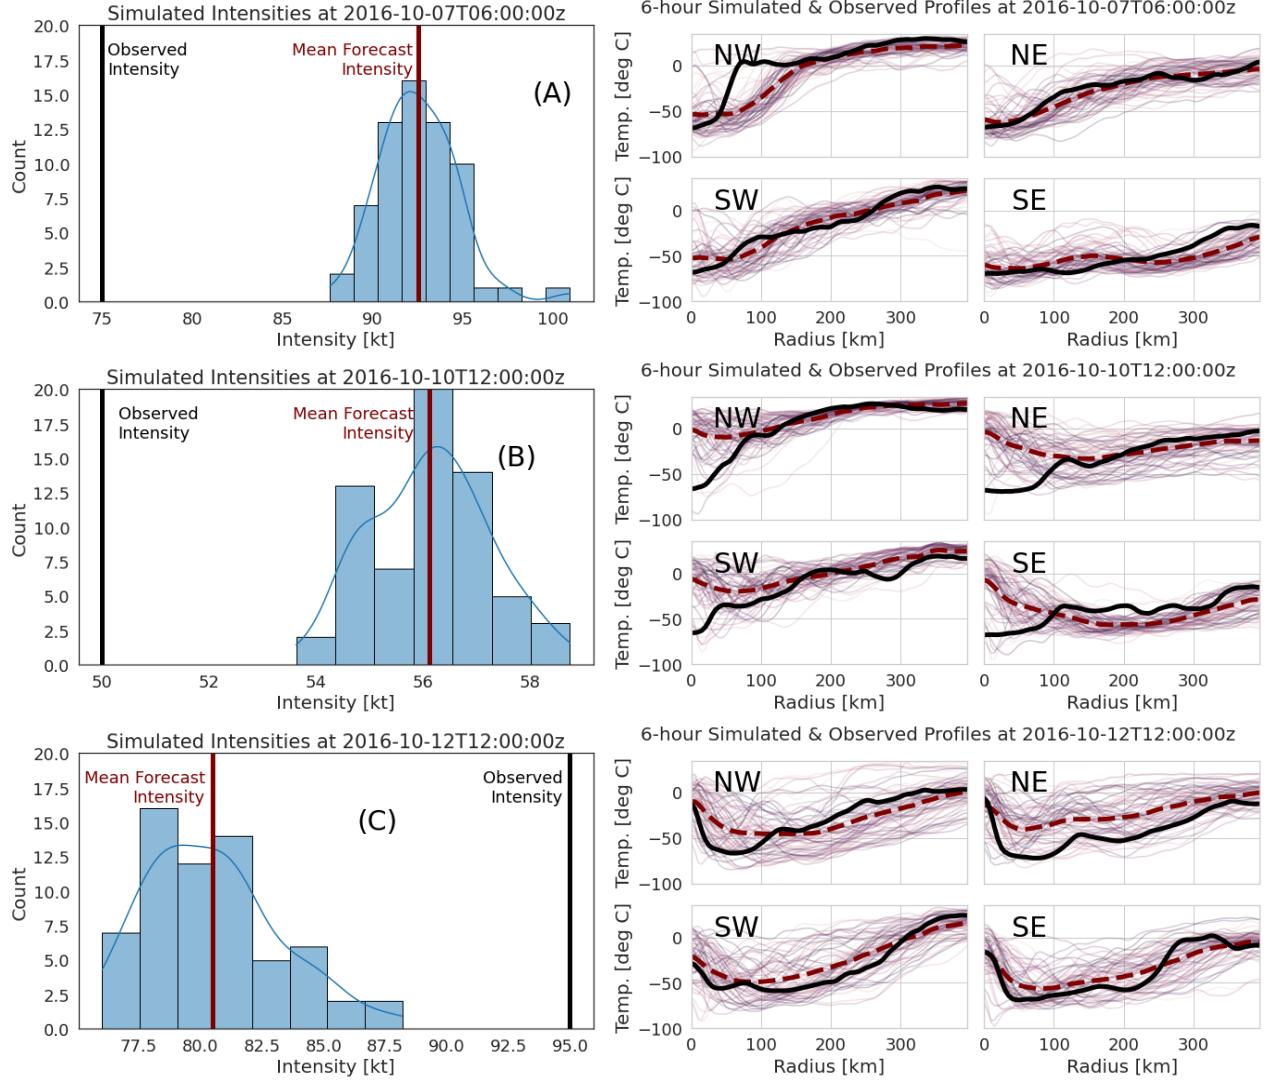
Additional Forecast Materials for Case Studies



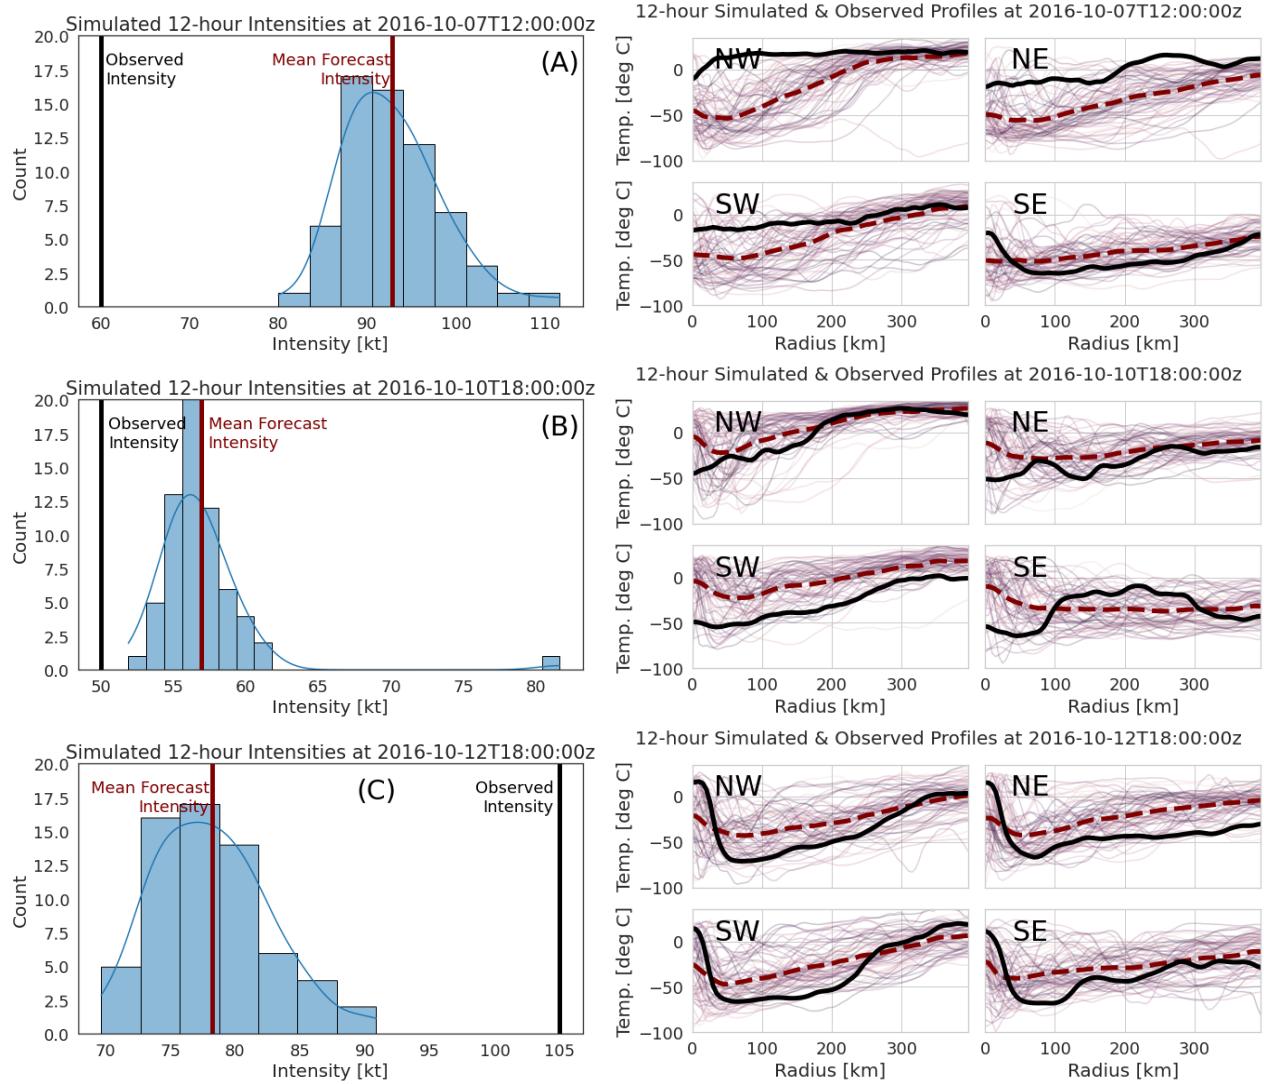
47 **FIG. C1. Hurricane Jose [2017] 6-hr Guidance:** Hurricane Jose was subjected to vertical wind shear out of
 48 the west/northwest due to Hurricane Irma; the structural forecasts tend to underestimate temperatures in the NW
 49 quadrant and thus overestimate TC intensity. (Left) Distribution of forecasted intensities with observed (black)
 50 and average forecast (red) intensities marked. (Right) Simulated profiles by quadrant with observed profiles
 51 represented by solid black curves, and averaged simulated profiles represented by dashed red curves.



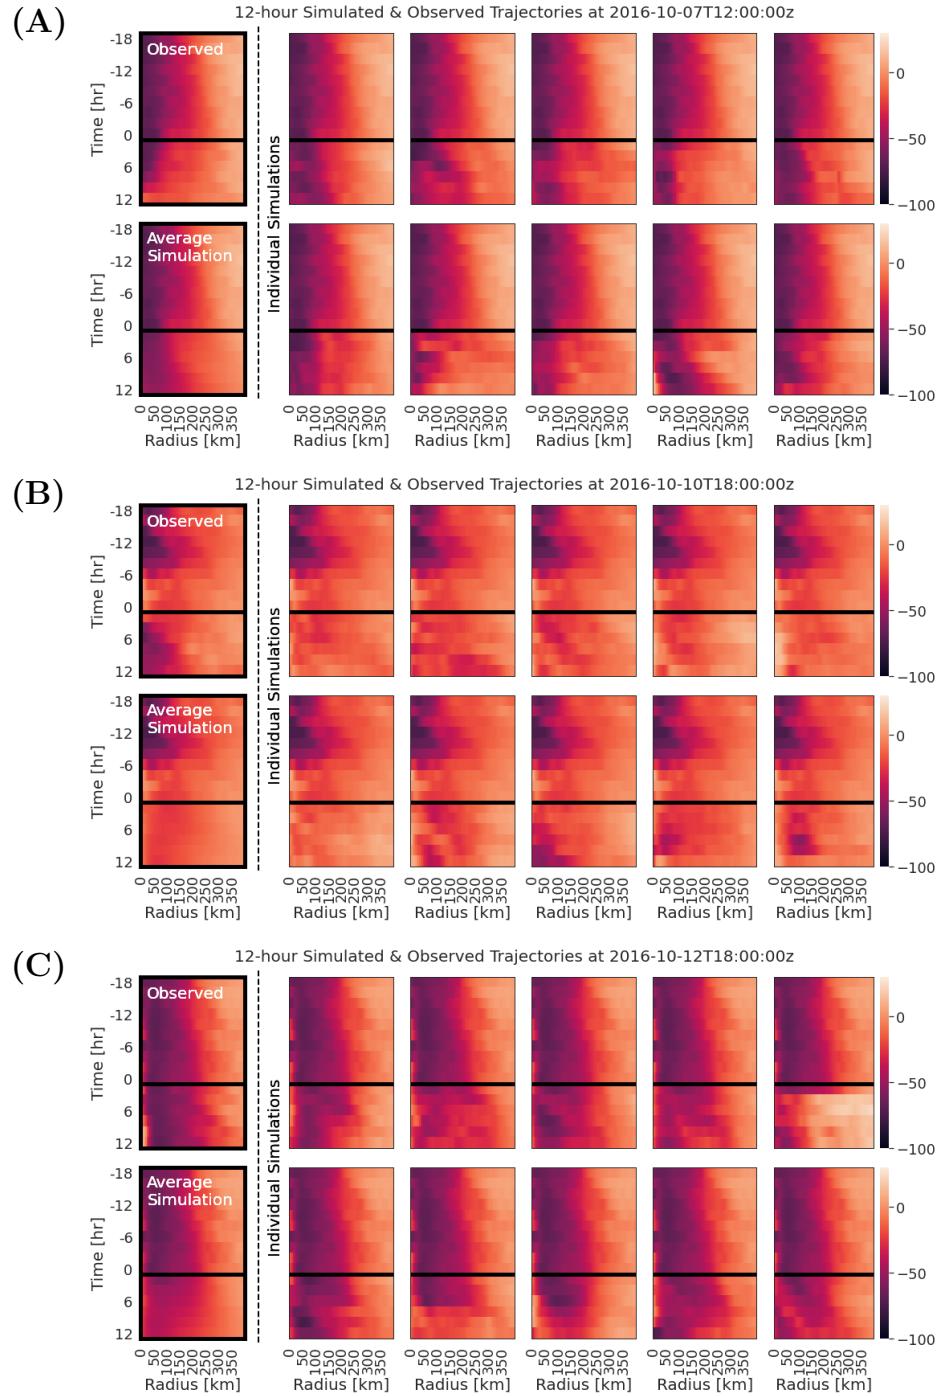
52 **FIG. C2. Hurricane Jose [2017] 12-hr Guidance:** At 12-hour lead times, this TC's evolution is still
 53 well modeled save for a handful of individual overestimates. (Left) Distribution of forecasted intensities with
 54 observed (black) and average forecast (red) intensities marked. (Right) Simulated profiles by quadrant with
 55 observed profiles represented by solid black curves, and averaged simulated profiles represented by dashed red
 56 curves.



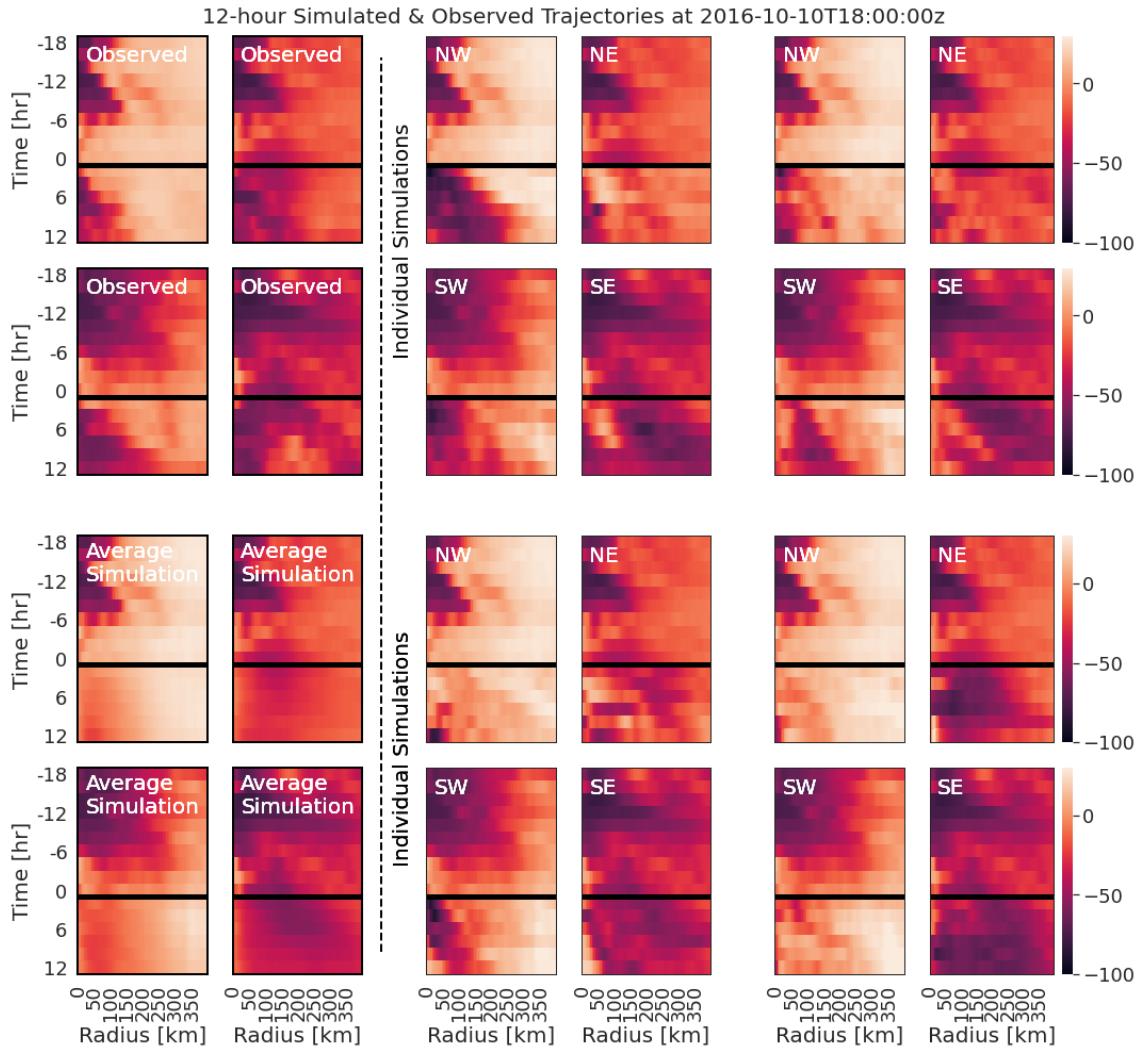
57 **FIG. C3. Hurricane Nicole [2016] 6-hr Guidance:** Hurricane Nicole underwent several rapid intensity change
 58 events. At 6 UTC 7 October, the structural forecast models the cloud top temperatures well, but this is insufficient
 59 to predict the extreme change from intensification to weakening even at 6-hour lead times. (Left) Distribution
 60 of forecasted intensities with observed (black) and average forecast (red) intensities marked. (Right) Simulated
 61 profiles by quadrant with observed profiles represented by solid black curves, and averaged simulated profiles
 62 represented by dashed red curves.



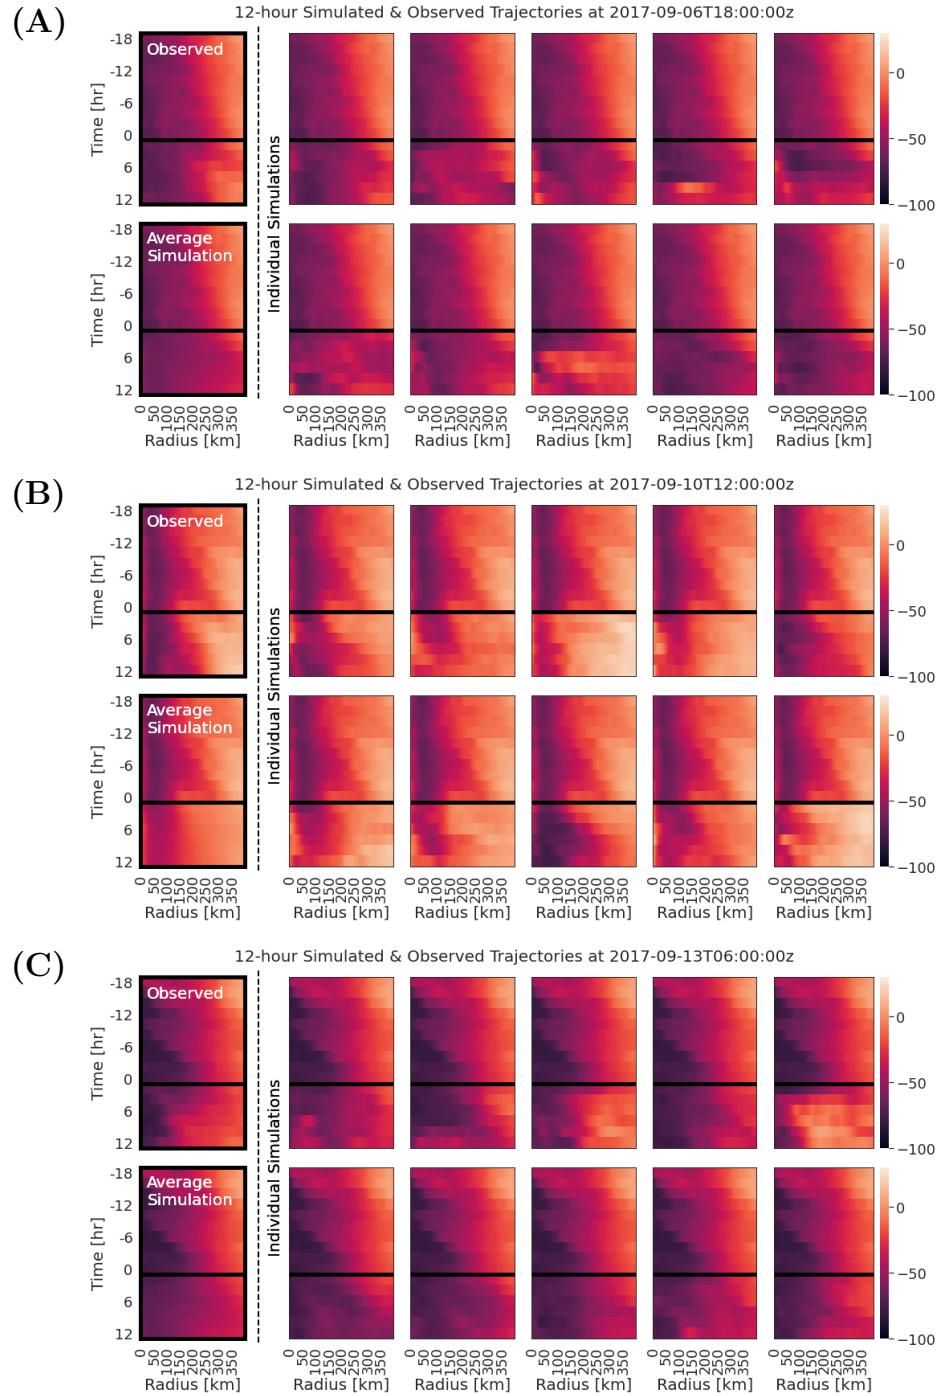
63 **FIG. C4. Hurricane Nicole [2016] 12-hr Guidance:** Model biases during the quick shifts from intensification
 64 to weakening and during steady intensification are exacerbated at 12-hour lead times. (Left) Distribution of
 65 forecasted intensities with observed (black) and average forecast (red) intensities marked. (Right) Simulated
 66 profiles by quadrant with observed profiles represented by solid black curves, and averaged simulated profiles
 67 represented by dashed red curves.



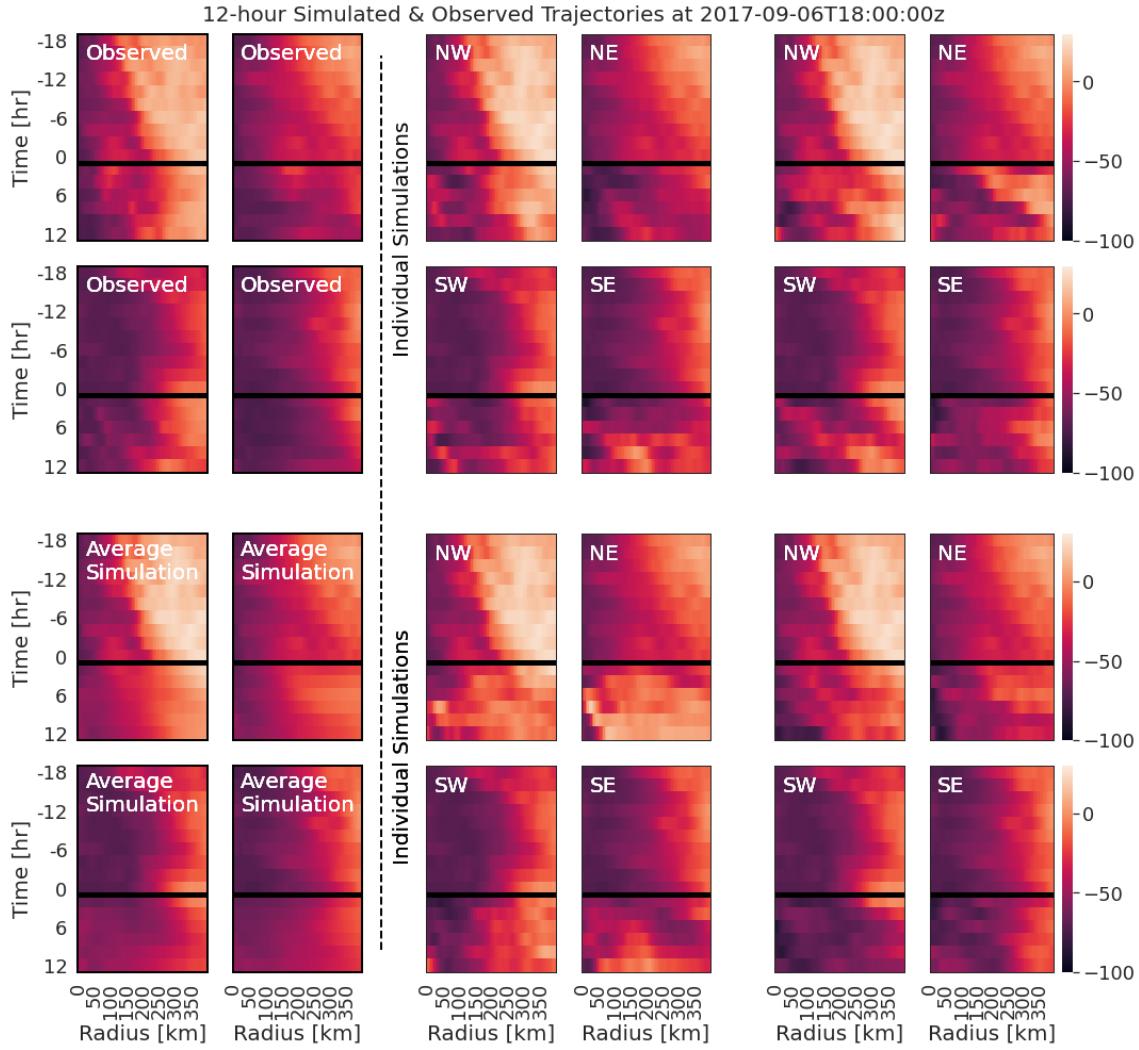
68 **FIG. C5. Hurricane Nicole [2016] 12-hr Structural Forecasts:** The *observed* structural trajectory is shown in
 69 the top left corner of each row. To the right, we see 10 *individual simulations* of radial profiles (averaged over all
 70 quadrants) at 12-hour lead times. All radial profiles above the black horizontal line are observed, while profiles
 71 below the black line are simulated. The bottom left corner shows the *average simulation* over 64 simulated
 72 trajectories.



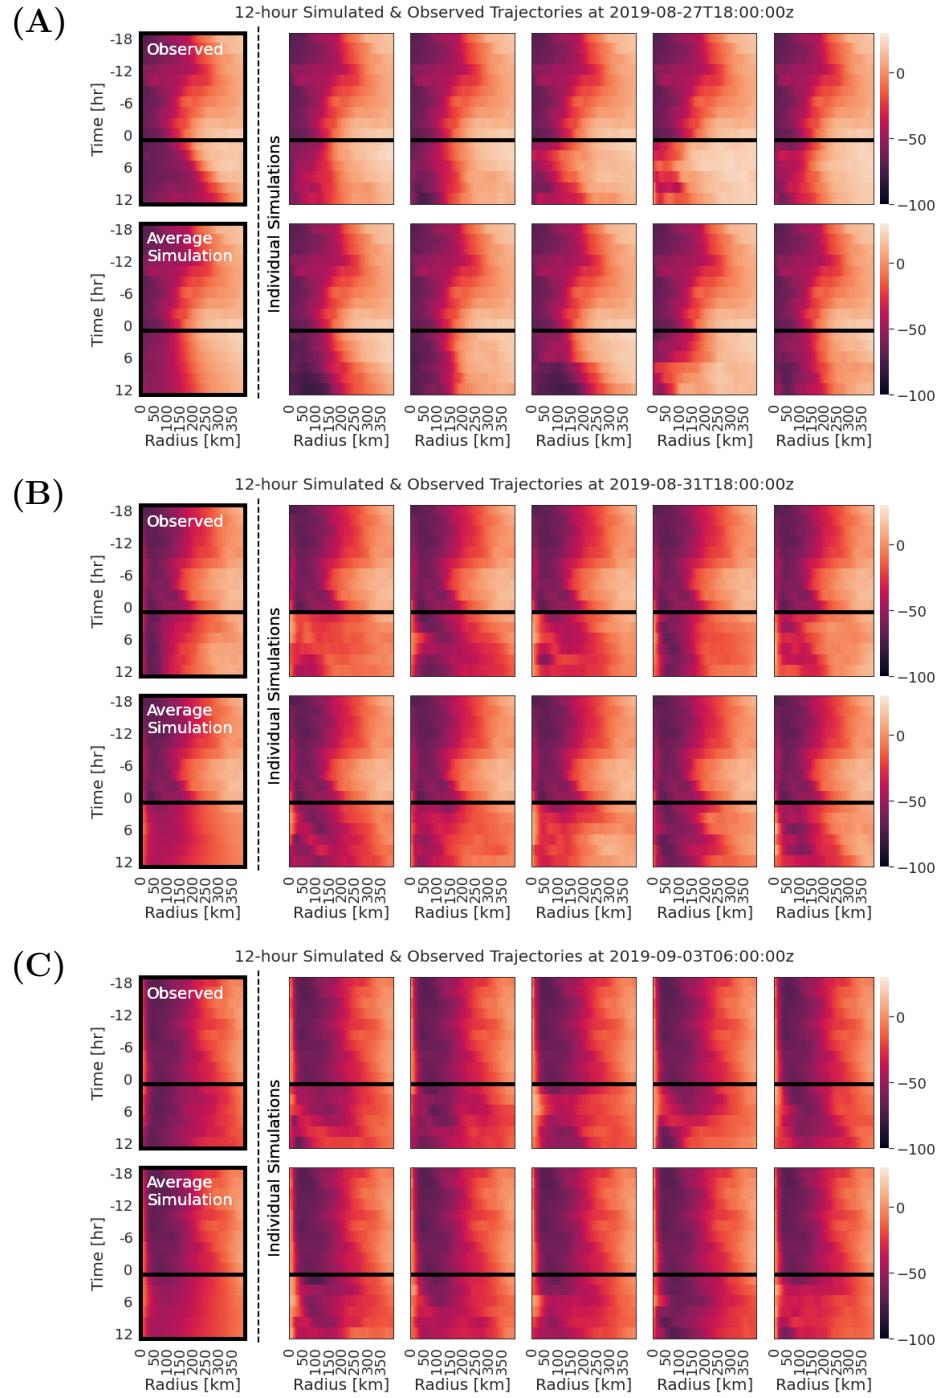
73 **FIG. C6. Hurricane Nicole [2016] Forecasts By Quadrant (B):** As Figure C5, but broken down by quadrant
 74 for example (B) only. The *observed* structural trajectory is shown in the top left corner. To the right, we see four
 75 *individual simulations* of radial profiles by quadrant at 12-hour lead times. All radial profiles above the black
 76 horizontal line are observed, while profiles below the black line are simulated. The bottom left corner shows the
 77 *average simulation* over 64 simulated trajectories.



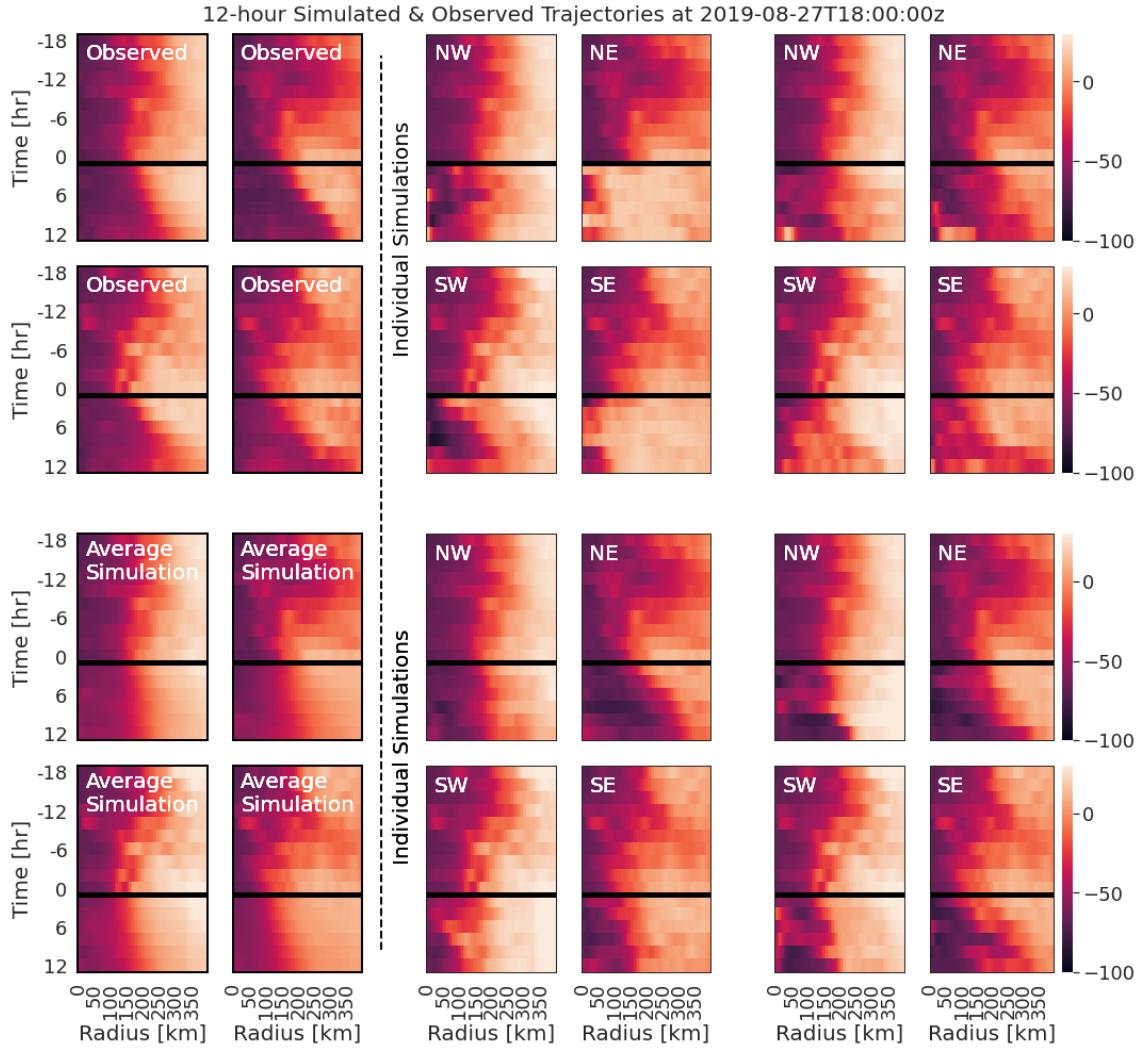
78 **FIG. C7. Hurricane Jose [2017] 12-hr Structural Forecasts:** The *observed* structural trajectory is shown in
 79 the top left corner of each row. To the right, we see 10 *individual simulations* of radial profiles (averaged over all
 80 quadrants) at 12-hour lead times. All radial profiles above the black horizontal line are observed, while profiles
 81 below the black line are simulated. The bottom left corner, shows the *average simulation* over 64 simulated
 82 trajectories.



83 **FIG. C8. Hurricane Jose [2017] Forecasts By Quadrant (A):** As Figure C7, but broken down by quadrant
 84 for example (A) only. The *observed* structural trajectory is shown in the top left corner. To the right, we see four
 85 *individual simulations* of radial profiles by quadrant at 12-hour lead times. All radial profiles above the black
 86 horizontal line are observed, while profiles below the black line are simulated. The bottom left corner shows the
 87 *average simulation* over 64 simulated trajectories.



88 **FIG. C9. Hurricane Dorian [2019] 12-hr Structural Forecasts:** The *observed* structural trajectory is shown
 89 in the top left corner of each row. To the right, we see 10 *individual simulations* of radial profiles (averaged over
 90 all quadrants) at 12-hour lead times. All radial profiles above the black horizontal line are observed, while profiles
 91 below the black line are simulated. The bottom left corner, shows the *average simulation* over 64 simulated
 92 trajectories.



93 **FIG. C10. Hurricane Dorian [2019] Forecasts By Quadrant (A):** As Figure C9, but broken down by quadrant
 94 for example (A) only. The *observed* structural trajectory is shown in the top left corner. To the right, we see four
 95 *individual simulations* of radial profiles by quadrant at 12-hour lead times. All radial profiles above the black
 96 horizontal line are observed, while profiles below the black line are simulated. The bottom left corner shows the
 97 *average simulation* over 64 simulated trajectories.