

Emulation of cloud microphysics in a climate model

W. Andre Perkins¹, Noah D. Brenowitz²,
Christopher S. Bretherton¹, Jacqueline M. Nugent³

¹Allen Institute for Artificial Intelligence

²NVIDIA

³University of Washington

Key Points:

- We build an emulator to replace the Zhao-Carr Fortran microphysics scheme in FV3GFS
- The integrated emulator sustains high skill throughout a 1-year simulation
- Tailoring the ML architecture to the structure of the underlying scheme greatly improves the online behavior of the emulator

Corresponding author: W. Andre Perkins, andrep@allenai.org

Abstract

We present a machine learning based emulator of a microphysics scheme for condensation and precipitation processes (Zhao-Carr) used operationally in a global atmospheric forecast model (FV3GFS). Our tailored emulator architecture achieves high skill ($\geq 94\%$) in predicting condensate and precipitation amounts and maintains low global-average bias ($\leq 4\%$) for 1 year of continuous simulation when replacing the Fortran scheme. The stability and success of this emulator stems from key design decisions. By separating the emulation of condensation and precipitation processes, we can better enforce physical priors such as mass conservation and locality of condensation, and the vertical dependence of precipitation falling downward, using specific network architectures. An activity classifier for condensation imitates the discrete-continuous nature of the Fortran microphysics outputs (i.e., tendencies are identically zero where the scheme is inactive, and condensate is zero where clouds are fully evaporated). A temperature-scaled conditional loss function ensures accurate condensate adjustments for a high dynamic range of cloud types (e.g., cold, low-condensate cirrus clouds or warm, condensate-rich clouds). Despite excellent overall performance, the emulator exhibits some deficiencies in the uppermost model levels, leading to biases in the stratosphere. The emulator also has short episodic skill dropouts in isolated grid columns and is computationally slower than the original Fortran scheme. Nonetheless, our challenges and strategies should be applicable to the emulation of other microphysical schemes. More broadly, our work demonstrates that with suitable physically motivated architectural choices, ML techniques can accurately emulate complex human-designed parameterizations of fast physical processes central to weather and climate models.

Plain Language Summary

In this study, we create computer code that uses machine learning to mimic a weather model's algorithm for handling how clouds form and rain falls. When used in the weather model to replace this algorithm, our machine learning code is highly accurate in simulations for a whole year. We achieve this by making smart code design choices. We split the code into two parts: one for cloud formation and one for rain and snow. This allows us to better build important aspects of these processes into the machine learning approach. For instance, clouds form where it is moist and evaporate when it gets dry, and rain and snow fall downward. Our code learns cloud behavior based on temperature to ensure it works both for cold, thin clouds high up in the sky and warm, thick clouds closer to the ground. Our work shows a path for suitably-designed machine learning code to eventually replace important parts of weather and climate models, but also that this path still requires careful human design respecting known physical principles.

1 Introduction

Atmospheric models combine fluid dynamics integrated on a discrete global grid with parameterizations of unresolved physical processes for weather and climate prediction. These parameterizations, encompassing phenomena such as cloud formation, precipitation, and radiative transfer, are crafted by experts and typically blend theoretical foundations with empirical relationships to capture interactions between various atmospheric processes. The ongoing development and refinement of these components require a careful balance between accuracy and efficiency to achieve high-fidelity simulations using limited computational resources.

Over the past few decades, advances in machine learning have led to substantial investments in computing facilities that combine more traditional CPU-based computing resources with accelerators such as GPUs. This shift in computational infrastructure has motivated the atmospheric modeling community to explore ways to capitalize on these newer resources to speed up simulations. The fluid dynamics algorithms imple-

63 mented in atmospheric models can often be recoded for more efficient GPU computa-
 64 tion using compiler directives or domain-specific language extensions (Dahm et al., 2023).
 65 However, the column-based physics parameterizations often involve more complex logic
 66 and data dependences that do not naturally fit into this paradigm.

67 An alternative approach to accelerating the physical components of atmospheric
 68 models is the creation of machine-learned emulators. Emulators are machine learning
 69 (ML) models trained directly on the inputs and outputs of a specific component, aim-
 70 ing to provide a seamless replacement of the original scheme. This strategy offers a nat-
 71 ural path to speed up model operation on accelerator-based compute resources, which
 72 are optimized to run ML workloads. Consequently, most emulation studies have focused
 73 on radiative transfer (Chevallier et al., 1998; Krasnopolsky et al., 2005, 2010; Veerman
 74 et al., 2021; Ukkonen et al., 2020), the most expensive subcomponent in the typical at-
 75 mospheric physics suite. However, recent studies have also emulated deep convection (O’Gorman
 76 & Dwyer, 2018), gravity wave drag (Chantry et al., 2021), atmospheric chemistry (Keller
 77 & Evans, 2019; Kelp et al., 2022; Schreck et al., 2022), and details of the warm rain pro-
 78 cess (Gettelman et al., 2021).

79 Emulation also serves as an excellent test bed for ML approaches that aim to im-
 80 prove on existing physical parameterizations, such as those using fine-resolution data to
 81 train corrective ML models (e.g., Brenowitz & Bretherton, 2019; Rasp et al., 2018; Yu-
 82 val & O’Gorman, 2020; Bretherton et al., 2022). Typically, these learn improvements to
 83 the combined suite of physical parameterizations, e.g. radiation, microphysics, turbu-
 84 lence and surface exchange, cumulus convection and orographic drag. Emulation of in-
 85 dividual component physical processes is clearly posed as a supervised learning task, so
 86 it can be used to explore skill bounds, quirks, and optimal architectural choices for em-
 87 ulating an entire parameterization suite.

88 The cloud microphysics scheme plays a central role in atmospheric modeling, man-
 89 aging rapid phase changes such as condensation, evaporation, and precipitation. It is tightly
 90 coupled to the model dynamics through latent heat release. We are not aware of past
 91 studies using ML to emulate an entire microphysics scheme, perhaps due to its lower com-
 92 putational cost compared to radiation. Nevertheless, it is a key part of emulating the
 93 combined physical parameterization suite and exposes a variety of ML challenges that
 94 are relevant to that broader problem. It is also a fast-acting process, producing local-
 95 ized atmosphere heating and drying tendencies that are much larger than for radiation.
 96 Thus, emulation of a representative microphysics scheme is a worthy complement to em-
 97 ulation of radiation parameterizations. It can provide valuable insights into the poten-
 98 tial and challenges of ML emulators of atmospheric physical processes.

99 In this work, we train an ML model to emulate the Zhao and Carr (1997, ZC) mi-
 100 crophysics scheme. This scheme was used for many years in the Global Forecast System
 101 (GFS) model by the U. S. National Centers for Environmental Prediction (NCEP). Here,
 102 it is included in a recent version of GFS that uses the FV3 dynamical core (Harris & Lin,
 103 2013), which we call the FV3GFS global atmospheric model. The ZC scheme, with only
 104 one prognostic condensate variable, seemed to be a simple machine learning target. How-
 105 ever, for a variety of reasons, developing a successful emulator of this scheme proved more
 106 challenging than anticipated, and required several architectural choices relevant to em-
 107 ulating other more complex microphysical parameterizations with many more prognos-
 108 tic hydrometeor types.

109 In Section 2, we describe the emulator architecture, training data, and integration
 110 into the FV3GFS model. In Section 3, we demonstrate that the emulator serves as a sta-
 111 ble, skillful replacement to the original Fortran Zhao-Carr microphysics scheme, with low
 112 global average bias for at least 1 year of simulation. Despite impressive overall perfor-
 113 mance, the emulator induces regional biases in the uppermost model levels— in our ex-
 114 perience, a relatively common online issue with ML integrated as one component in con-

115 ventional atmospheric models (e.g., Brenowitz & Bretherton, 2019; Clark et al., 2022).
 116 In Section 4, we discuss the major decisions that influenced the emulator’s performance
 117 and address some remaining challenges and limitations of our approach.

118 In accordance with AGU’s AI tool policy, the authors acknowledge the use of Ope-
 119 nAI’s ChatGPT-4 tool to help edit the manuscript draft for clarity, conciseness, and gram-
 120 matical correctness. All suggestions provided by the AI tool were reviewed and edited
 121 by the authors for correctness and consistency. The plain language summary was gener-
 122 ated by prompting the tool for a generally accessible version of our written abstract
 123 and then edited by the authors.

124 **2 Methods**

125 In this work, we utilize the FV3GFS global atmospheric model (Harris & Lin, 2013),
 126 which is currently used by NOAA for operational weather forecasting. FV3GFS com-
 127 bines the FV3 nonhydrostatic finite-volume dynamical core with a suite of physical pa-
 128 rameterizations developed for the Global Forecast System (GFS). For the simulations
 129 presented here, the FV3GFS model is run on a C48 cubed-sphere grid (approximately
 130 200 km horizontal grid spacing) with 79 vertical levels.

131 Within FV3GFS, we target the emulation of the Zhao-Carr (ZC) microphysics (Zhao
 132 & Carr, 1997), which was used in the operational version of GFS until 2018. The ZC mi-
 133 crophysics scheme predicts changes in cloud condensate, precipitation, and the associ-
 134 ated heating and moistening rates at each grid point in a vertical column, based on state
 135 inputs. The scheme divides the prediction into two subroutines: one calculating the lo-
 136 cal condensate change via grid-scale condensation (gscond) and the other calculating col-
 137 umn precipitation and associated condensate adjustments (precpd). Figure 1 shows a
 138 graphical depiction of the information flow through the ZC microphysics subroutines.
 139 The scheme diagnoses the phase partitioning of cloud condensate into liquid and ice at
 140 each step based on temperature and the presence of overlying ice cloud. Furthermore,
 141 it diagnoses the downward precipitation flux and its phase partitioning into rain and snow
 142 at each grid level during each time step. Appendix A gives further details.

143 The ZC scheme initially seemed appealing for ML emulation due to its simplicity,
 144 featuring only a single prognostic hydrometeor type: the cloud water mixing ratio. De-
 145 spite the initial appearance of simplicity, the schematic (Fig. 1) illustrates that the ZC
 146 scheme is architecturally more complex than we anticipated due to the implicit depen-
 147 dence on the column thermodynamic state sampled within the previous time step. Fur-
 148 thermore, vertically and temporally nonlocal phase partitioning of condensate does not
 149 appear as an explicit output of the scheme, despite its use by other parameterizations.
 150 These subtleties add considerable time-consuming challenge to the accurate ML emu-
 151 lation of the ZC scheme.

152 To emulate the ZC scheme, we employ hooks to interact with the Fortran model
 153 via the package `call_py_fort` (https://github.com/nbren12/call_py_fort). This pack-
 154 age enables users to call functions and interact with selected Fortran state fields within
 155 an initialized Python environment, giving access to the comprehensive suite of ML and
 156 data tools available in Python and accelerating ML prototyping and testing.

157 **2.1 Training Data**

158 We generate the training dataset by initializing 30-day simulations from GFS anal-
 159 ysis on the first day of each month in 2016, saving fields every 5 hours to sample the di-
 160 urnal cycle. A list of all stored fields is shown in Table S1. We reserve three months of
 161 data for validation during training (February, June, and September). The training dataset

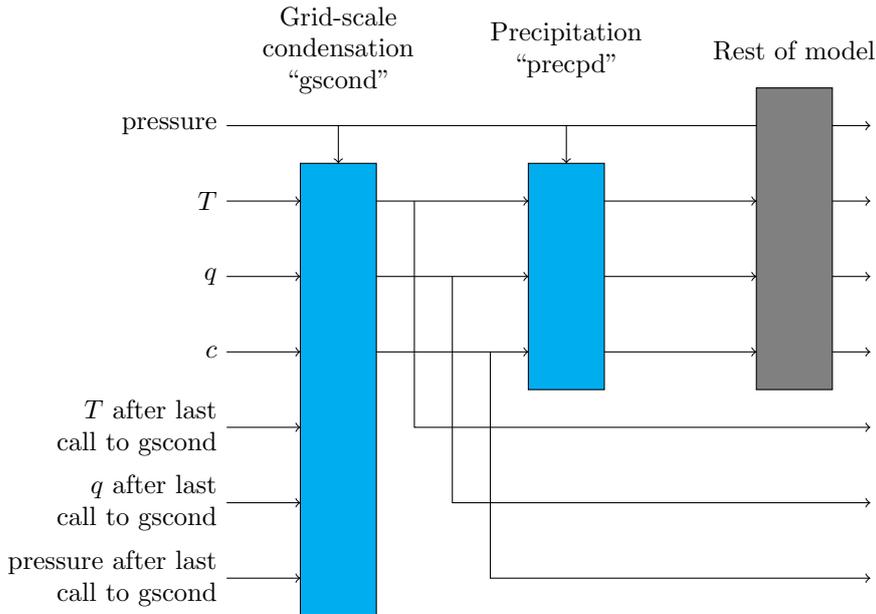


Figure 1. Information flow of the Zhao-Carr microphysics within FV3GFS for a single time step. Inputs of a given scheme are represented as inward arrows. The “after last call to gscond” inputs are used to compute a relative humidity tendency that encompasses the rest of the model and precpd. This approach to computing the tendency effectively adds three new state variables to the model.

162 includes 1080 global snapshots consisting of $48^2 \times 6 = 13824$ atmospheric columns, to-
 163 taling nearly 15 million samples.

164 From the saved training data, we derive the target increments for the ZC micro-
 165 physics that we seek to emulate. The total change, denoted as $\Delta = \Delta_g + \Delta_p$, is the
 166 sum of the two subroutine updates from gscond (Δ_g) and precpd (Δ_p). Both gscond and
 167 precpd calculate updates for temperature (T), specific humidity (q), and the cloud wa-
 168 ter mixing ratio (c); precpd also diagnoses the amount of surface precipitation (P) dur-
 169 ing the time step. We note that the use of tendencies in this manuscript refers to the sub-
 170 routine increment divided by the model time step (15 minutes).

171 Figure 2 displays an example transect of tendencies of the target data for clouds
 172 and humidity along the 100°W meridian. The gscond cloud water tendency (Fig. 2a; $\Delta_g c$)
 173 can be positive (condensation) or negative (evaporation), depending on local thermo-
 174 dynamic state. Active regions in this snapshot include the boundary layer of the sub-
 175 tropical Pacific and free-tropospheric weather features (e.g., convection or frontal zones)
 176 over land. Because cloud water tendency involves a phase change between water vapor
 177 and condensate, the corresponding tendencies of temperature ($\Delta_g T$) and specific humid-
 178 ity ($\Delta_g q$) exhibit similar patterns to the cloud water tendency. The gscond tendencies
 179 for these three fields are fully determined by grid-local thermodynamic state, with the
 180 exception of one vertically non-local flag, which influences the diagnostic decomposition
 181 between liquid and ice clouds and the resulting latent heating tendency. That flag in-
 182 dicates whether mixed-phase clouds with temperatures between 0° and -15°C are over-
 183 laid by contiguous ice cloud colder than -15°C .

184 The corresponding precpd condensate tendency transect (Fig. 2b; $\Delta_p c$) shows losses
 185 due to autoconversion of thicker clouds to precipitation. Regions of positive precpd va-

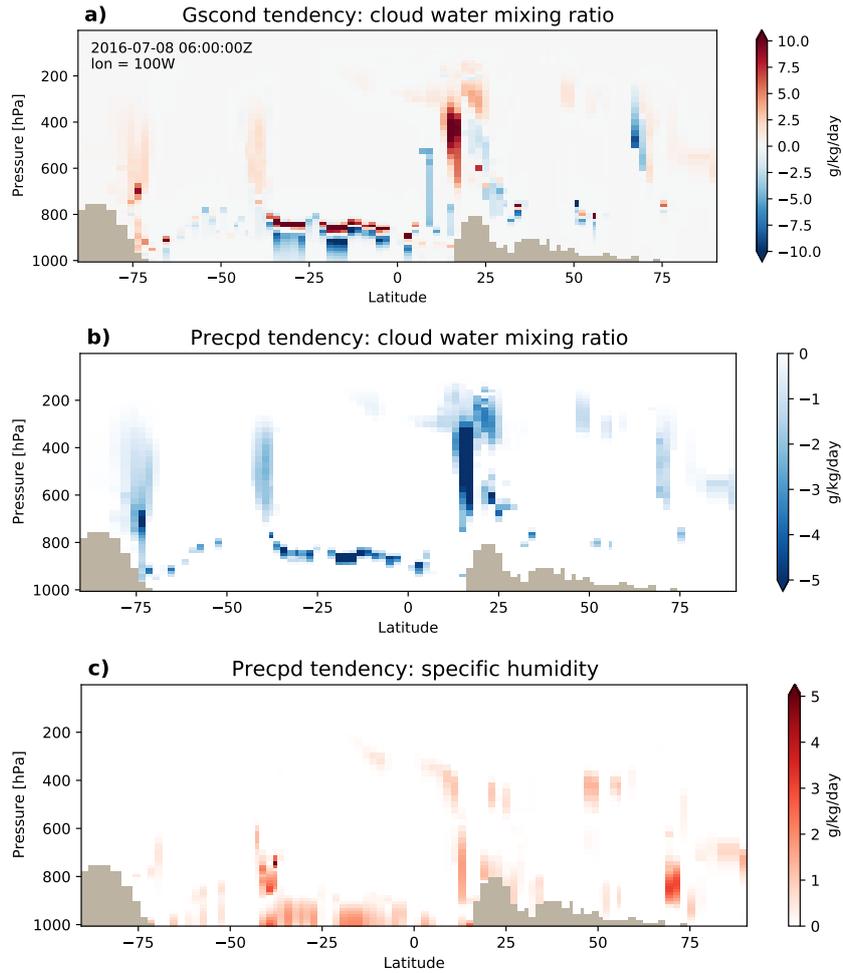


Figure 2. Latitude–pressure transects along longitude 100°W for a sample Zhao-Carr microphysics step on July 8th, 2016 at 06Z showing: (a) the condensation rate from gscond, (b) the conversion rate of cloud to precipitation in precpd, and (c) the precipitation evaporation rate in precpd. Transect data has been interpolated to pressure levels from model levels for presentation.

186 por tendency (Fig. 2c; $\Delta_p q$) are due to the evaporation of precipitation falling from over-
 187 lying grid layers.

188 These transects highlight two general challenges for emulating microphysics. First,
 189 the microphysics scheme is not active at the majority of grid points. It produces a range
 190 of adjustments to the state fields where clouds or precipitation are present, but elsewhere,
 191 the tendencies should be exactly zero. Second, the condensation scheme can generate
 192 large condensate increments throughout the troposphere despite the humidity being orders
 193 of magnitude smaller in the upper troposphere than near the surface.

194 Some other general considerations are also important for ML microphysics emu-
 195 lation. For instance, clouds are very sensitive to relative humidity. A small error in pre-
 196 dicted water vapor or temperature can significantly impact clouds and precipitation. Sec-
 197 ond, cirrus clouds with small condensate mixing ratios can be as radiatively important
 198 as liquid water clouds with hundred-fold higher condensate mixing ratios. Thus, an ML
 199 emulator must accurately predict a large range of condensate tendencies to skillfully re-

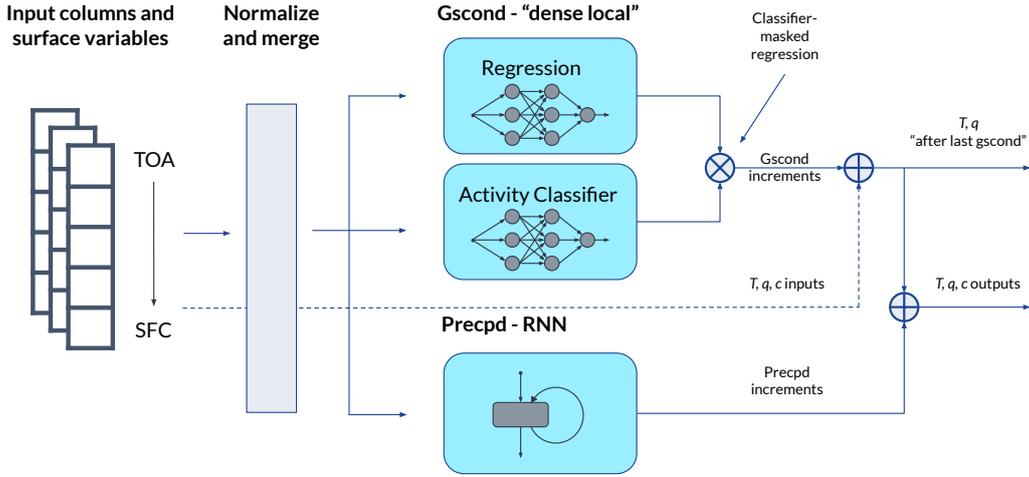


Figure 3. A schematic of the ZC microphysics emulation architecture.

200 produce the original model’s climate. Third, complete cloud evaporation/sublimation
 201 is common; to obtain this outcome in a model time step requires the condensate tendency
 202 to exactly remove all cloud condensate in a grid box. Lastly, microphysical tendencies
 203 are a combination of local (e.g., condensation) and non-local (e.g., precipitation) pro-
 204 cesses and constraints. An emulation scheme must replicate these dependencies to yield
 205 accurate and physically consistent results.

206 These factors heavily influenced the final design of our emulation methodology, which
 207 we detail in the following section. We elaborate on the sensitivity of results to these choices
 208 and discuss remaining challenges in Section 4.

209 2.2 Emulator Architecture

The emulation model architecture is shown in Figure 3. Separate emulators for *gscond* and *precpd* take a total of 13 input variables, including the same set of inputs as the Fortran ZC scheme: T , q , c , and surface pressure as well as the “after last *gscond*” values of T , q , and surface pressure. We provide additional inputs of air pressure and pressure thickness of the atmospheric layer, as well as derived inputs of relative humidity (RH), and log-scaled q , c , and q after last *gscond*. Each input is normalized:

$$x'_j = (x_j - \mu_j)/\sigma \quad (1)$$

210 and combined to form input channels for the emulation models. The mean, μ_j , is a sam-
 211 ple mean at level j using 150,000 random columns from the training data. The scaling
 212 factor, σ , is calculated using the standard deviation over all per-level centered $(x_j - \mu_j)$
 213 values in the same sample. This scaling enhances training stability and conveniently down-
 214 weights inputs from the upper levels, where the microphysics scheme is less active. Sur-
 215 face variables are normalized as a single level and then broadcast to 79 levels when merged
 216 into model inputs to simplify general usage. The same input data are passed to all three
 217 of the emulator subcomponents.

218

2.2.1 Condensation emulator

219
220
221
222
223
224
225
226
227

In the condensation subroutine (gscond), net condensation $\Delta_g c$ at a given point in an atmospheric column is physically determined by the thermodynamic inputs at that same level, a property we refer to as grid-point locality. The gscond emulator takes advantage of this property by applying a single MLP to each grid point, which we refer to as a dense-local model. The MLP is 2 layers of 256 channels, each with ReLU activation. It takes in 79-level \times 13-channel inputs, applies the model to each level, and outputs a single column (79 \times 1) through a linear readout layer. We train the gscond dense-local regressor for 50 epochs using the Adam optimizer with a learning rate of 0.0001. We use a mean squared error (MSE; Eq. 2) based loss (Eq. 3).

$$\text{MSE}(a, b) = \frac{1}{N} \sum_{i=1}^N (a_i - b_i)^2 \quad (2)$$

$$L = \text{MSE}(\tilde{y}, \hat{y}) + \lambda \cdot \text{MSE}(c'_g, \hat{c}'_g) \quad (3)$$

$$\tilde{y} = \frac{\Delta_g c - \tilde{\mu}(T_{in})}{\tilde{\sigma}(T_{in})} \quad (4)$$

$$\hat{y} = f(x) \quad (5)$$

$$c_g = \Delta_g c + c_{in} \quad (6)$$

$$\hat{c}_g = \hat{y} \sigma(T_{in}) + \mu(T_{in}) + c_{in} \quad (7)$$

228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246

The target increment in the loss (\tilde{y} , Eq. 4) is conditionally scaled due to a physical expectation that cloud properties depend strongly on temperature (Fig. S1). To accurately emulate cold cirrus clouds, which typically have little condensate and correspondingly small condensate increments, and also emulate warm liquid clouds, which can have hundred-fold larger condensate increments, the loss function normalizes to be sensitive in both cases. The scaling terms for the mean $\tilde{\mu}(T_{in})$ and standard deviation $\tilde{\sigma}(T_{in})$ represent a piecewise interpolation based on the input temperature T_{in} . We compute the underlying interpolation function by calculating binned mean and standard deviation values after grouping samples of $\Delta_g c$ into 50 linearly-spaced bins between the minimum and maximum input temperature. We optimize the gscond emulator $\hat{y} = f(x)$ to predict temperature-scaled increments (\tilde{y}) as functions of the grid point features x . These increments are descaled into a predicted post-gscond condensate amount (\hat{c}_g , Eq. 7) by adding the de-scaled increment to the input condensate amount. We include a post-gscond condensate MSE in the loss (Eq. 3) using the normalized condensate amounts (c'_g, \hat{c}'_g) scaled by $\lambda = 50000$ to make the loss contribution $O(1)$. The addition of the final condensate value to the loss function improves validation MSE for the unscaled condensate increment by over 80%. This likely happens because the final condensate term gives additional weight to warm-cloud condensation. The remaining state increments for T and q are determined at runtime from the predicted $\Delta_g c$ value (see Section 2.3).

247
248
249
250
251
252

We train an activity classifier to handle the mixed discrete-continuous nature of the condensation scheme, i.e., the need to force the emulator prediction to either (i) zero tendency when there should be no cloud change during the time step, or (ii) the exact tendency to fully evaporate cloud condensate present at the beginning of the time step. The classifier model employs the same dense-local architecture as the regressor, but predicts four target variables to identify the following classes:

253
254
255
256

- $\Delta_g c = 0$,
- $c_g = 0$ and $\Delta_g c \neq 0$,
- $c_g \neq 0$ and $\Delta_g c > 0$, and
- $c_g \neq 0$ and $\Delta_g c < 0$.

257 The first two cases, corresponding to situations (i) and (ii) above, together usually ac-
 258 count for 80% or more of the outcomes depending on the level (Fig. S2). During infer-
 259 ence, the model constrains $\Delta_g c$ when the classifier identifies either of the first two cases.
 260 Otherwise, the regressor makes the condensate prediction. We train the classifier using
 261 categorical cross-entropy loss with the same hyperparameters as the regressor, except
 262 for an increased learning rate of 0.001. After training, the classifier is approximately 98%
 263 accurate over all classes and levels (Table S2). Please refer to Section 4.1 for a more in-
 264 depth discussion on the impacts of the conditional loss function and activity classifier.

265 **2.2.2 Precipitation emulator**

266 The diagnostic precipitation scheme (precpd) generates precipitation through au-
 267 toconversion of cloud condensate in upper levels. The precipitation falls and can either
 268 evaporate in lower layers or reach the surface. To enforce this downward dependence in
 269 the precpd emulator by construction, we use a recurrent neural network (RNN) that re-
 270 curses over vertical layers starting at the top of atmosphere (see schematic in Fig. S3).
 271 A single RNN layer,

$$h_{j+1} = (W_h h_j + W_x x_j + b)^+, \quad (8)$$

272 uses the same normalized inputs, x'_j , as the gscond emulator where $j \in [0, 79]$ and $j =$
 273 0 is the top of the atmosphere. In this form, h_j is the RNN hidden state at level j , W_h
 274 represents trainable weights for the recursion on hidden state, W_x are the trainable weights
 275 for inputs, b is the bias, and $(\cdot)^+$ represents a ReLU activation function. We stack two
 276 hidden 256-channel layers followed by a level-independent linear readout layer ($\hat{y}_j = Ah_j +$
 277 b) to predict the increments $\Delta_p T$, $\Delta_p q$, and $\Delta_p c$. This construction ensures that only
 278 inputs x_i from levels at and above level j ($i \leq j$) can affect RNN predictions at level
 279 j . We embed additional constraints within the precpd emulator such that it converts clouds
 280 to precipitation ($\Delta_p c \leq 0$), that it evaporates precipitation ($\Delta_p q \geq 0$ and $\Delta_p T \leq 0$),
 281 and that the final cloud is non-negative ($c_p \geq 0$). The RNN loss includes the MSE for
 282 the normalized increments (using Eq. 1 instead of conditional normalization) and the
 283 MSE of the normalized post-precpd output for each variable scaled such that the indi-
 284 vidual contributions are $O(1)$. The surface precipitation rate (P) is diagnosed from the
 285 net loss in total column water at runtime using:

$$P = - \sum_{j=0}^{78} (\Delta_p c_j + \Delta_p q_j) \cdot \Delta p_j / g, \quad (9)$$

286 where for each level j , $(\Delta_p c_j + \Delta_p q_j)$ is the local water change due to autoconversion
 287 and evaporation, Δp_j is the input pressure thickness of the atmospheric layer, and g is
 288 gravity.

289 **2.3 Prognostic runs**

290 The utility of a microphysics emulator ultimately depends on its performance when
 291 used within the atmospheric model as a substitute for the human-designed parameter-
 292 ization it is trained to replace. Specifically, the emulator should not cause catastrophic
 293 model failures, it should consistently provide a skillful representation of the original mi-
 294 crophysics behavior, and it should have a minimal impact on the integrated statistics (i.e.,
 295 the climate) of the underlying model. To test this, we embed the ZC microphysics em-
 296 ulator in FV3GFS and run a series of prognostic tests using two model configurations:
 297 one with the emulator as the active microphysics scheme (online) and a baseline with
 298 the Fortran microphysics active (offline). In each case, we run the inactive component

299 in a diagnostic mode (“piggybacked”; Grabowski, 2019) and save the resulting tenden-
300 cies for comparison.

301 To evaluate the skill and climate impact of the emulated microphysics, we initial-
302 ize 30-day simulations in each calendar month from February 2016 to January 2017. The
303 initializations are taken from the end of the training data simulations, testing both model
304 configurations on atmospheric states independent of the training data. We compute skill
305 scores for all microphysics tendencies (ΔT , Δq , Δc ; converted to a tendency by divid-
306 ing increments by $\Delta t = 900$) and P using a modified R^2 score $1 - \sum(\hat{y} - y)^2 / \sum y^2$.
307 A score of 1 indicates a perfect emulation, while a value of 0 or lower indicates an em-
308 ulator worse than a no-information prediction. We also compute the bias of the micro-
309 physics outputs and the atmospheric state over all levels and times of the 12 simulations.
310 To assess long-term stability, we simulate a full year using the emulator in place of the
311 ZC microphysics and check the global averages and bias for evidence of any climate drifts.

312 The last step in applying the emulator as part of an online simulation is to apply
313 final physical limiters and constraints and generate the full set of outputs for the em-
314 ulated ZC microphysics. For the gscond emulator, we compute the increments $\Delta_g T$ and
315 $\Delta_g q$ through local conservation of the net condensation. First, we limit the net conden-
316 sation based on moisture availability using:

$$\Delta_g c = \begin{cases} \max(-c_{in}, \Delta_g c), & \text{if } \Delta_g c < 0 \\ \min(q_{in}, \Delta_g c), & \text{if } \Delta_g c > 0 \end{cases}. \quad (10)$$

317 Then, the change in water vapor mirrors the change in condensate ($\Delta_g q = -\Delta_g c$) and
318 the temperature change is determined via latent heating ($\Delta_g T = (L_v/c_p)\Delta_g c$), where
319 L_v is the latent heat of vaporization and c_p is the specific heat of air at constant pres-
320 sure. This is an approximation, as some phase changes in ZC occur between ice and vapor,
321 releasing additional latent heat; however, these phase changes are not fully locally
322 determined and our efforts to use a posthoc determination of ice cloud latent heating ef-
323 fects slightly degraded online emulator skill. For online application, we set the top 5 lev-
324 els of gscond increments to zero since the ZC microphysics scheme is never active in those
325 stratospheric levels and noise issues in ML-predicted condensate increments arise in these
326 levels (see Section 4.2 for further discussion). Finally, we add the increments to the cor-
327 responding input state variable to obtain fields after gscond ($T_g = T_{in} + \Delta_g T$, $q_g =$
328 $q_{in} + \Delta_g q$, and $c_g = c_{in} + \Delta_g c$).

329 The precipd increment constraints are directly integrated into the ML model as de-
330 scribed earlier. We derive the surface precipitation (Eq. 9), and then add the precipd in-
331 crements to the post-gscond values to generate the final scheme outputs ($T_p = T_g +$
332 $\Delta_p T$, $q_p = q_g + \Delta_p q$, and $c_p = c_g + \Delta_p c$).

333 3 Results

334 We begin with the top-level results of our ZC emulation 30-day runs in Table 1.
335 The offline skill scores for all emulated quantities are nearly perfect at $\sim 99\%$, with low
336 root mean-square error (RMSE) values and biases that are 1–2 orders of magnitude smaller
337 than the RMSE (i.e., a small component of the error).

338 Online skill is a strict test where deviations from a realistic physical state can cause
339 the diagnostic Fortran microphysics to output large state adjustments or even crash. Nev-
340 ertheless, when the emulator is used online, it maintains high skill scores with only a ~ 1 –
341 5% reduction compared to the offline case. Predicted cloud water tendencies show the
342 lowest average performance at 94%, which is still quite high for a sparse and highly sen-
343 sitive tendency field. The corresponding tendency RMSEs of emulator tendencies vs. pig-
344 ggybacked Fortran tendencies are roughly double those of the offline configuration, ex-

ZC Output	Offline			Online		
	Skill score	RMSE	Bias	Skill Score	RMSE	Bias
ΔT [K/day]	0.99	0.42	-0.03	0.98	0.58	-0.02
Δq [mg/kg/day]	0.995	110	3.0	0.99	200	-1.1
Δc [mg/kg/day]	0.99	140	-1.0	0.94	330	-0.7
P [mm/day]	0.998	0.21	-0.02	0.97	0.77	0.02

Table 1. Skill metrics for the ZC microphysics emulator outputs compared to the Fortran microphysics outputs for the offline (Fortran driving) and online (emulator driving) configurations. All table metrics are calculated for twelve 30-day runs initialized at the start of each calendar month and then averaged together.

cept for P , where the tendency RMSE is nearly four times larger. The larger online error result is an expected outcome due to detrimental feedbacks between the model and the ML emulator that cannot be accounted for when using offline training. The biases remain small in the online case, suggesting no systematic breakdown of the emulator behavior from the diagnostic Fortran microphysics.

We compare the time-averaged atmospheric state averaged across the twelve 30-day online simulations with identically initialized baseline simulations to show that the emulator produces little mean-state drift when used in FV3GFS in place of the original ZC microphysics. Figure 4 depicts zonal averages of the online bias of the emulator-based simulation compared to the baseline simulation, which have been interpolated from model level to pressure coordinates to display biases at a true relative height. Table 2 gives global average area- and mass-weighted bias for selected output fields.

Cloud water is a key output of the microphysics scheme. Its zonal average mixing ratio (Fig. 4a, b) has the largest absolute bias near the surface in Antarctica, ~ 6 mg/kg. This bias is relatively large for the characteristically cold, dry air there. Outside of the Antarctic, the cloud water biases are ~ 3 mg/kg or less— a much smaller relative change from the baseline— and are generally positive, except for a negative bias in the tropical upper troposphere. The global-mean cloud water bias is small— 0.2 mg/kg, an approximately 2% increase compared to the baseline state (Table 2). These cloud changes result in $O(1\%)$ changes to the outgoing top-of-atmosphere longwave (-1.4 W/m²) and shortwave radiation ($+1.3$ W/m²), but in total the changes largely cancel out.

Figure 4d depicts the online bias in RH, which displays a small shift towards saturation in the middle-to-lower troposphere. The largest biases in RH ($>10\%$) occur in the Antarctic upper atmosphere near the large gradient in drying near the tropopause. There are also similar albeit smaller positive RH biases in the tropics and Arctic tropopause regions. Overall, the global-mean RH shows a small positive bias of 0.8% (Table 2), congruent with the small positive cloud water bias.

The zonal average temperature has a small cold bias of up to -1.5 K in the high latitudes. Between 50°S – 50°N , this bias is weakened or even slightly reversed at some pressures, but there is a thin layer of warm bias up to 1 K near the tropopause. The zonal temperature biases largely cancel out when averaged globally over the 30-day runs (Table 2).

Lastly, the total surface precipitation (emulated ZC microphysics + convective sources) has a slight positive bias of 0.03 mm/day, a 1% increase from the baseline simulation (Table 2). Fig. 5a depicts the online zonal average surface precipitation just from the ZC microphysics component. The emulated ZC precipitation production is nearly identical to the baseline simulation owing to the high emulation skill of Δq and Δc , but produces

Field	Bias	Baseline mean
Air temperature [K]	-0.1	251
Specific humidity [mg/kg]	-0.7	2590
Relative humidity [%]	0.8	45.5
Cloud water [mg/kg]	0.2	9.6
Total surface precipitation [mm/day]	0.03	3.04
Upward shortwave at TOA [W/m ²]	1.3	91.9
Upward longwave at TOA [W/m ²]	-1.4	237
Total outgoing radiation at TOA [W/m ²]	-0.06	329

Table 2. Global average online biases and baseline means for selected state fields averaged over all 30-day simulations.

ZC Output	Online skill	
	1-year run	30-day runs avg.
ΔT	0.98	0.98
Δq	0.98	0.99
Δc	0.94	0.94
P	0.97	0.97

Table 3. Online skill score for 1-year online simulation compared against the skill scores averaged across the twelve 30-day runs initialized across the calendar year.

382 0.02 mm/day less global precipitation than the baseline ZC scheme. This bias must mostly
383 be associated with state drift rather than offline emulator errors, because the piggybacked
384 Fortran ZC scheme, which is applied to the online emulator state, diagnoses slightly less
385 precipitation than the online emulator, especially in the Northern Hemisphere storm track.
386 The Fortran convection parameterization also responds to the slight emulator-induced
387 state changes by producing a global mean convective precipitation increase of 0.05 mm/day.

388 The instantaneous precipitation-rate distribution based on all grid columns and sam-
389 pling times (Fig. 5b) corroborates this analysis. It shows that the emulator overproduces
390 light precipitation (< 0.1 mm/day) compared to the piggybacked Fortran scheme, but
391 these two schemes agree well at most higher precipitation rates, and their small discrep-
392 ancies don't explain the online emulator differences from the baseline simulation. Instead,
393 the largest precipitation rate bins (~100 mm/day) suggest that the online emulator-driven
394 simulation shifts to fewer states that support heavy precipitation events compared to the
395 baseline simulation.

396 **3.1 1-year continuous simulation**

397 The monthly-initialized runs show the embedded ZC emulator is stable for at least
398 30 days during all calendar months of the year, with low biases. To further explore the
399 long-term fidelity of emulator-based simulations, we present results from a continuous
400 1-year integration starting in July 2016. We ran two simulations, one masking only the
401 top 5 levels of the gscond increments (i.e., setting the increments to 0) and the other mask-
402 ing the top 5 levels of both gscond and precip increments. We found adding the mask
403 to the top 5 levels of the precip scheme reduced the number and severity of transient
404 tendency skill dropouts (Fig. B1) for the 1-year simulation. Both online simulations ran
405 with online emulation for the full year. We present results for the top 5 gscond and precip
406 increment configuration due to better performance. We discuss the unresolved sensitiv-
407 ity of the emulator to the upper levels in Section 4.2.

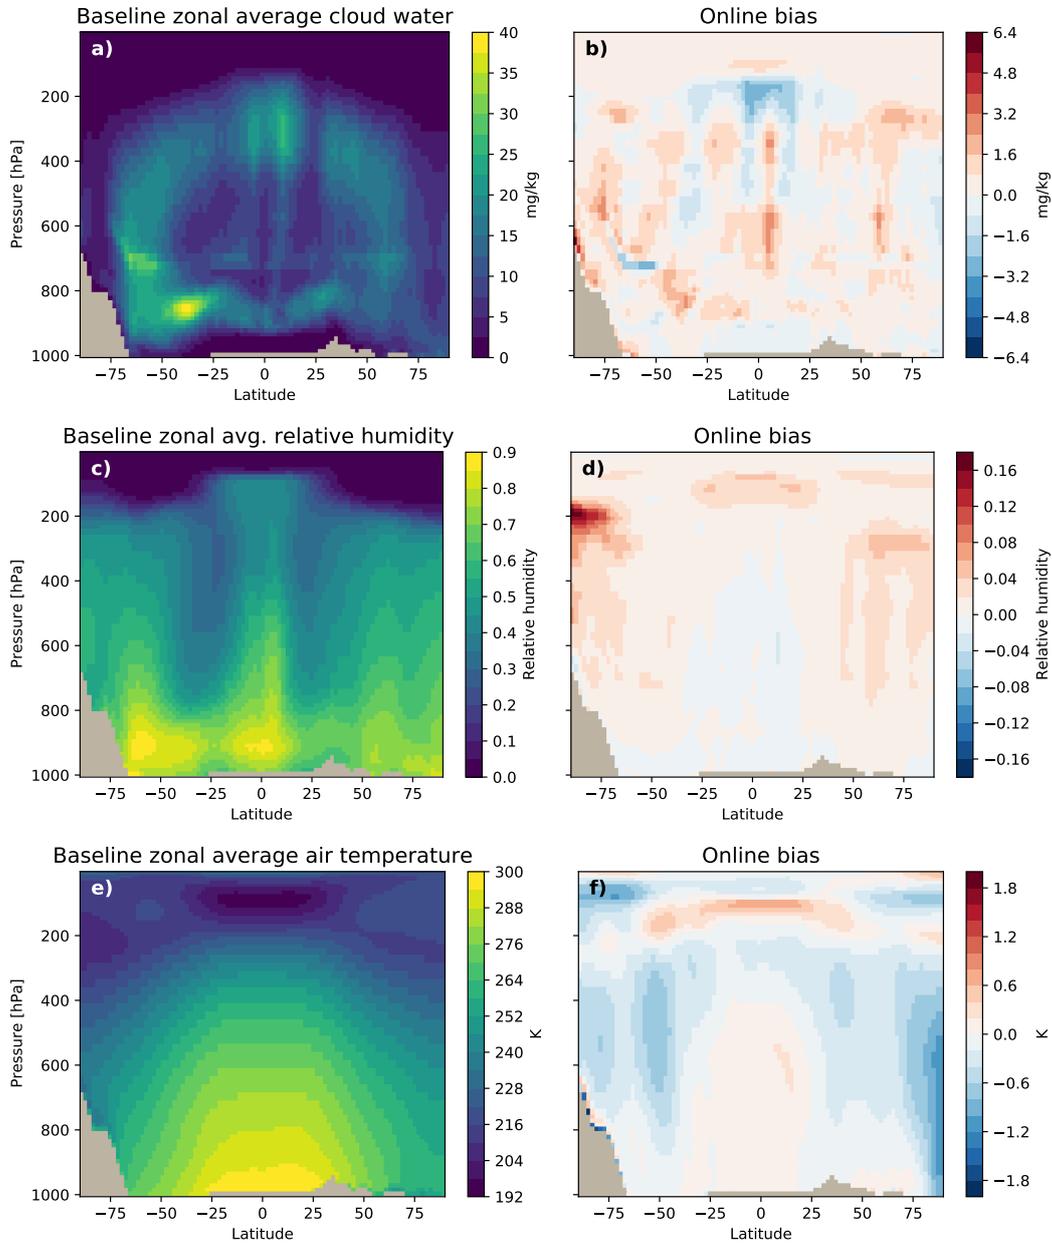


Figure 4. Latitude–pressure sections of zonal and time average state from baseline Fortran simulations (left) and online bias of simulation using the emulator (right) for cloud water mixing ratio (a, b), relative humidity (c, d), and air temperature (e, f). Averages are over twelve 30-day simulations initialized in each month of the calendar year, using values vertically interpolated from model levels.

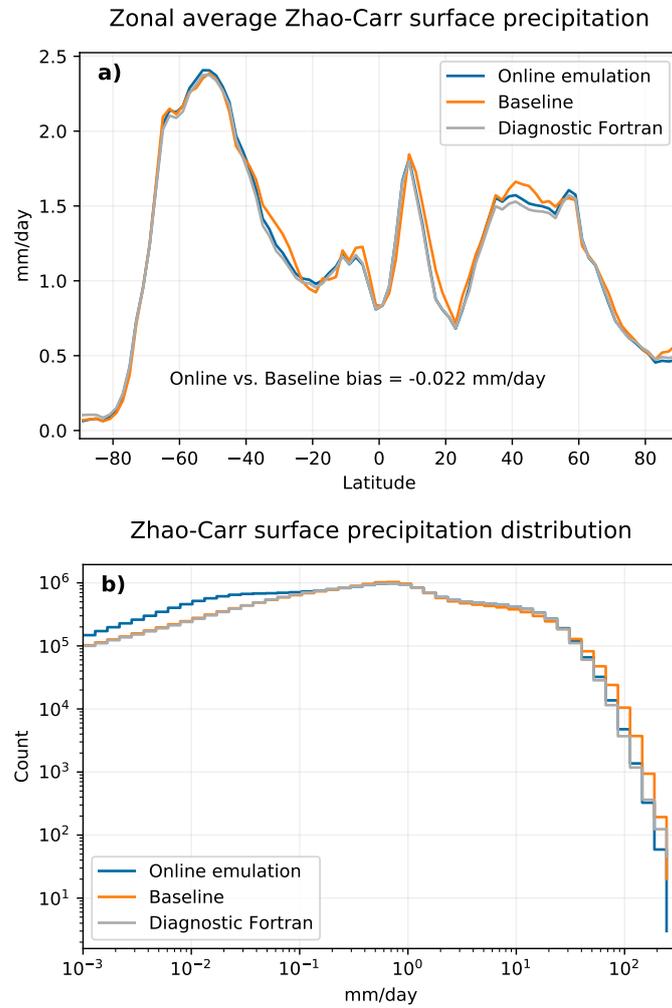


Figure 5. (a) Zonal average surface precipitation rate from ZC microphysics compared between the online emulator (blue) baseline Fortran (orange) and diagnostic Fortran microphysics (grey), which is generated diagnostically using inputs from the online emulation state. (b) Surface precipitation rate distribution compared between the same schemes. Shown quantities are calculated from twelve 30-day simulations initialized at the beginning of each calendar month.

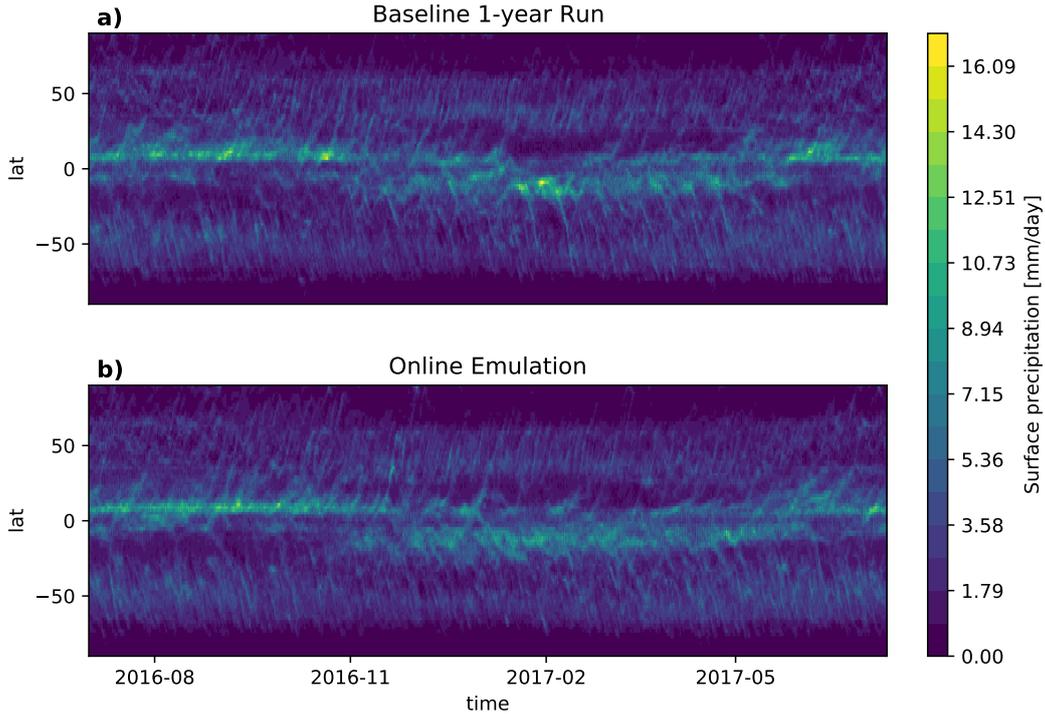


Figure 6. Time–latitude plots of the instantaneous surface precipitation rate saved every 3 hours from the 1-year (a) baseline and (b) online emulation simulations.

408 The online skill metrics for the 1-year continuous run are, reassuringly, almost identical
 409 to the average of the 30-day runs (Table 3). A time–latitude plot of total surface
 410 precipitation (Fig. 6) compares the baseline and online emulation runs, demonstrating
 411 the emulation retains the spatiotemporal character of the baseline precipitation (and pre-
 412 cipitating clouds by proxy) throughout the seasonal cycle. A slight reduction in the largest
 413 precipitation events for the online emulation is apparent in the tropics; we already noted
 414 this issue for the month-long simulations in Fig. 5b. Some global-annual-average biases
 415 (Table 4) are somewhat larger than in the 30-day runs: T (-0.3 K), RH (1.9%), and net
 416 TOA outgoing radiation (-0.4 W/m²; the difference of a -2.1 W/m² outgoing longwave
 417 bias and a 1.6 W/m² reflected shortwave bias). Absolute cloud water and surface pre-
 418 cipitation biases remain similar to those of the 30-day runs. Cloud water and RH have
 419 the largest relative bias from the baseline simulation at $\sim 4\%$, respectively.

420 The zonal average biases of T and RH from the 1-year emulator-based simulation
 421 are very small in the troposphere but become more significant in the polar stratosphere
 422 (Fig. 7). In this region, large negative cold biases (as low as -8 K) are co-located with
 423 positive RH biases up to 30%. The temperature bias appears within the first few months
 424 of the simulation and stabilizes for the rest of the simulation. We further investigated
 425 these biases and found that both the gscond and precpd emulators have deficiencies in
 426 the dry, cold polar stratosphere. Within a few hours after the start of the simulation,
 427 the gscond emulator produces too much condensate because the emulator predicts con-
 428 densation for what the Fortran piggybacked microphysics diagnoses should mostly be
 429 evaporation at marginal relative humidities (40–50%; Fig. S4). We have confirmed that
 430 the gscond bias drift is unrelated to precpd or the classifier. We hypothesize that the
 431 tendency drift is likely related to a subtle online shift in some characteristics of the in-
 432 put distribution specific to this region.

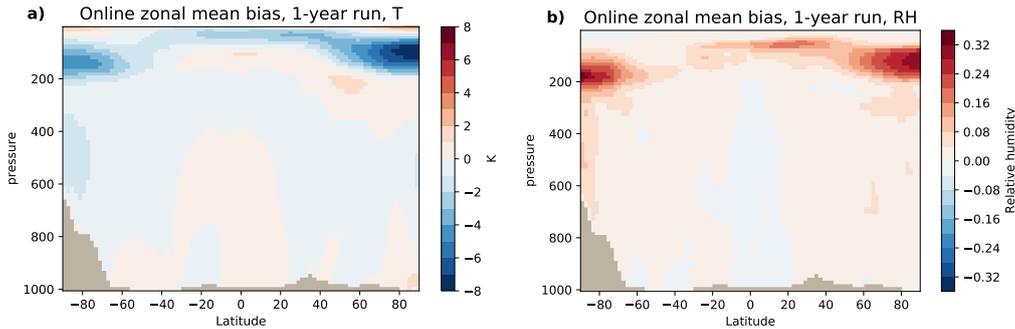


Figure 7. Zonal mean bias of the 1-year online emulation simulation for (a) temperature and (b) relative humidity.

Field	Bias	Baseline mean
Air temperature [K]	-0.3	247
Specific humidity [mg/kg]	17.2	2680
Relative humidity [%]	1.9	45.6
Cloud water [mg/kg]	0.2	7.6
Surface precipitation [mm/day]	0.03	3.03
Upward shortwave at TOA [W/m^2]	1.6	92.1
Upward longwave at TOA [W/m^2]	-2.1	237
Total outgoing radiation at TOA [W/m^2]	-0.44	329

Table 4. As in Table 2 but for the 1-year simulation.

433 The precipd emulator’s shortcomings in the polar stratosphere are evident from of-
 434 fine diagnosis. Specifically, errors from the emulator’s noise floor produce evaporation
 435 despite no falling precipitation (Fig. S5) in this region. This is a particular failing of the
 436 the single-scaling loss normalization (Eq. 1), where optimization fails to minimize the
 437 large relative errors in the polar stratosphere. The errors produce a directional bias due
 438 to constraints imposed in the model architecture ($\Delta_p q > 0$ and $\Delta_p T < 0$) and a lack
 439 of enforced conservation. As they grow, these biases in the high-latitude stratosphere likely
 440 feed back with radiation and the atmospheric circulation before ultimately equilibrat-
 441 ing.

422 4 Challenges and choices

443 In this section, we highlight key decisions that led to a skillful, stable, and low-bias
 444 emulation, as well as some remaining challenges. From the outset, our goal was to use
 445 simpler ML models with the potential for general applicability in emulating atmospheric
 446 physics parameterizations. However, the path to the final emulator necessitated several
 447 problem-specific choices to successfully emulate the ZC microphysics scheme.

448 4.1 Key decisions

449 One of the most influential decisions was to target subcomponents of the micro-
 450 physics scheme, specifically grid-scale condensation (gscond) and precipitation (precipd).
 451 Initial attempts to encapsulate the total ZC scheme tendency increments in a single model
 452 yielded high offline skill, but the online integration often resulted in difficult-to-interpret
 453 failures that crashed the simulation. This is a common failure mode when training mod-
 454 els outside of the environment in which they are deployed (e.g., Brenowitz & Brether-

run type	gscond arch.	precpd arch.	ΔT	Δq	Δc	P
offline	dense-local	RNN	0.99	0.995	0.99	0.998
	dense-local	dense-column	0.99	0.99	0.97	0.99
	dense-column	dense-column	0.97	0.98	0.95	0.99
online	dense-local	RNN	0.98	0.98	0.95	0.98
	dense-local	dense-column	0.74	0.76	0.01	0.01
	dense-column	dense-column	-0.39	-0.46	-0.07	0.17

Table 5. Sensitivity of emulation skill to the use of general vs. prior-informed model architectures. “Dense-column” refers to a fully connected MLP with 2 hidden layers of 256 width and a linear readout layer. “Dense-local” and “RNN” refer to the architectures described in the methods section.

455 ton, 2019). Separating the subcomponents simplifies the enforcement of physical priors
456 through model architecture design or output postprocessing.

457 Following component separation, we observed substantial improvements in online
458 emulation skill by incorporating physically informed architectures. For the gscond em-
459 ulator, we enforce grid point locality (i.e., dependence only on the grid point-local ther-
460 modynamic state) by using a dense-local MLP that does not mix any vertical informa-
461 tion. For the precpd emulator, we enforce the downward dependence (i.e., rain falls down-
462 ward) using an RNN that recurses downward over a vertical column. Table 5 displays
463 the offline and online skill for a single 30-day run initialized in July, comparing the per-
464 formance of the informed architectures to a reasonable uninformed default for atmospheric
465 model process parameterization— a dense MLP combining features over the entire grid
466 column to predict the full column increments. While these dense-column models exhibit
467 high skill offline (always >95%), they fail online when continuously integrating on the
468 atmospheric state. Replacing the RNN used for precpd emulation with a dense-column
469 architecture that does not enforce downward dependence reduces cloud and precipita-
470 tion skill to nearly 0%, even when using the physically informed gscond architecture. Us-
471 ing dense-column models for both subroutines results in negative skill (i.e., worse than
472 zero-increment predictions) for all variables except surface precipitation.

473 The discrete-continuous nature of outputs from some atmospheric physics param-
474 eterizations (e.g., for microphysics) poses a unique challenge for emulation. Neural net-
475 work regressors have difficulty producing exact zeros, since they are trained to a certain
476 degree of precision and will produce noise below that threshold. This can complicate on-
477 line integration, particularly for a microphysical scheme, where the local thermodynamic
478 state may be quite sensitive to small changes in condensate or humidity, especially in very
479 cold regions (e.g., Antarctica or the upper troposphere). For this reason, we introduced
480 the activity classifier described in Sect. 2.2.1 into the gscond emulator. Figure 8 illus-
481 trates the need for such a classifier by comparing cloud distributions from simulations
482 with and without a classifier to a baseline run. By day 15 after initialization, the con-
483 densate histogram shows that the emulation scheme without an activity classifier accu-
484 mulates small values of cloud water (≤ 2 mg/kg) at many grid points. Including a clas-
485 sifier within the gscond emulator to constrain the microphysical activity resolves this is-
486 sue. Based on the good performance of the 30-day online simulations and non-locality
487 of the precipitation scheme, we decided not to pursue an activity mask for the precpd
488 emulator. However, the erroneous T and q precpd increments in the polar stratosphere
489 contributing to biases in the 1-year run suggest a classifier might be helpful overall.

490 The final choice important to the success of the ZC emulator involved optimizing
491 the model to predict condensate increments that span many orders of magnitude. As de-

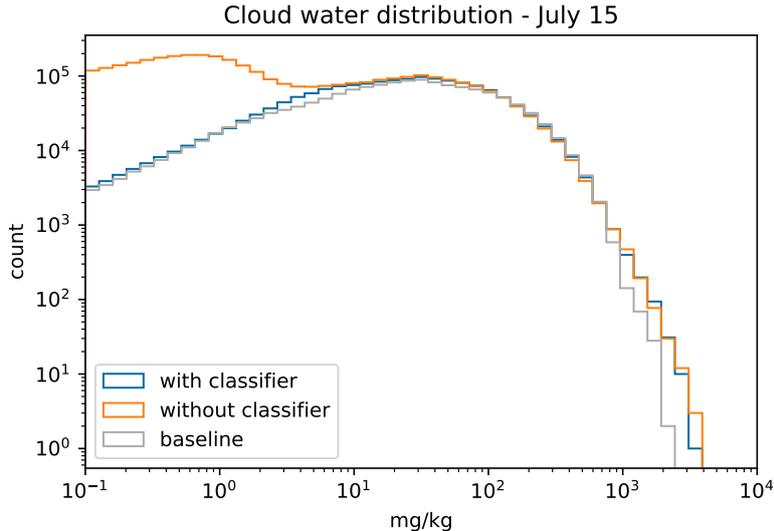


Figure 8. Cloud water mixing ratio distributions compared between three configurations: online emulation with a `gscond` activity classifier (blue), online emulation without an activity classifier (orange), and a baseline simulation (grey). Samples are taken from 8 3-hourly snapshots across day 15 of a 30-day simulation initialized on July 1.

492 scribed in Sect. 2.2.1, we used a temperature-dependent scaling in the `gscond` loss function,
 493 ensuring proportionate errors across a large range of local microphysical states. Model-
 494 level scaling is insufficient to handle this because a given model level may span a broad
 495 range of temperatures (e.g., the tropical boundary layer vs. the Antarctic plateau).

496 In addition to the conditional scaling, we added select rescaled input values (RH,
 497 log-scaled q and c) into the emulator inputs. Removing log-scaled inputs negatively im-
 498 pacts offline skill in polar and upper-level model regions (not shown). Including RH as
 499 an input increased skill and reduced condensate biases, particularly in the Antarctic re-
 500 gion. For example, by day 5 of a July 1 initialized simulation, the emulator using RH
 501 as an input has an Antarctic average column-integrated condensate of 87 g/m^2 compared
 502 to a baseline value of 79 g/m^2 . When not including RH, the average Antarctic column-
 503 integrated condensate value is 154 g/m^2 by July 5, roughly double the baseline value.
 504 Despite the overlap of the additional inputs, we believe they help reduce errors in cold-
 505 cloud regions by allowing the emulator discern vertical position, which is removed by per-
 506 level demeaning in the input normalization (Eq. 1). We conducted an experiment to rein-
 507 troduce the vertical information by adjusting the input normalization for air pressure
 508 to remove the column mean instead of the per-level mean from each level. This config-
 509 uration also increased offline skill and largely removed the Antarctic condensate bias with-
 510 out the need for RH, but was generally more sensitive to skill dropouts when used on-
 511 line.

512 4.2 Remaining challenges

513 In developing our emulation scheme, online simulations commonly presented un-
 514 expected challenges that needed to be addressed. Certain months, primarily October and
 515 November, tended to have lower online skill ($\sim 85\text{--}90\%$ compared to $\sim 93\text{--}96\%$) for clouds
 516 and precipitation compared to other months. The lower aggregate skill in these months
 517 was mainly due to significant `precipd` autoconversion misses (“skill dropouts”) during con-
 518 vective events for a few low-latitude columns (see Appendix B for an example). These

519 skill dropouts start in the mid-troposphere near the freezing level and quickly affect the
 520 entire upper troposphere. The emulator recovers in the affected grid columns within a
 521 few hours or, at worst, a few days.

522 To minimize such dropouts, we employed a strategy of training an ensemble of emu-
 523 lators initialized with varying random seeds (e.g., as in Clark et al., 2022) and then se-
 524 lect combinations of `gscond` and `precpd` emulators with the best online skill during the
 525 most problematic months of October and November. While this approach does not guar-
 526 antee prevention of severe skill dropouts during other months or in a year-long simula-
 527 tion, it consistently produces stable, low-bias emulators with high skill.

528 We still do not have a foolproof approach for designing emulators without occa-
 529 sional skill dropouts. For instance, the emulator configuration that gave the most skill-
 530 ful 1-year online simulation (masking the top 5 levels of increments from both `gscond`
 531 and `precpd`) produces a substantial skill dropout in a 30-day simulation initialized at the
 532 start of December, leading to a December Δc skill = 54%, while the original `gscond`-only
 533 top 5 mask configuration has no issues (December Δc skill = 94%).

534 Altogether, this suggests the need for further refinement of the architectural de-
 535 sign and training choices, such as whether recursion from the top model level is neces-
 536 sary, whether additional measures should be adopted to reduce sensitivity to the upper
 537 levels, or whether more training data are needed to handle the few convective events on
 538 the edges of the data distribution.

539 To handle the large dynamic range of condensate increments, we use temperature
 540 scaling in the `gscond` loss function. While this is generally very beneficial, especially in
 541 tandem with the `gscond` classifier, it does not prevent the emulator from occasionally cre-
 542 ating spurious cloud in the uppermost model levels. These levels lie in the stratosphere,
 543 where temperature increases with height. Warmer temperatures lead to larger-amplitude
 544 condensate “noise”, which the emulator later struggles to remove. Because there should
 545 never be any cloud in the top-most levels, we pragmatically resolved this by masking `gscond`
 546 increments in the top 5 model levels. However, as seen in the 1-year simulation polar strato-
 547 spheric biases, a few issues remain related to emulator deficiencies in the upper levels.

548 While the current manuscript focuses on the development and evaluation of a ro-
 549 bust, accurate ZC emulator, we recognize that speed of execution is a paramount con-
 550 sideration for emulator adoption, especially in operational settings. The current code in-
 551 frastructure was designed for flexibility and ease of testing new ideas, rather than for op-
 552 timal speed. In its current unoptimized state, the model with online emulation runs ap-
 553 proximately 30% slower (~ 5.8 s/time step) than to the original C48 simulation (~ 4.8
 554 s/time step) even when using available GPUs (4x Nvidia T4). Variable transfer between
 555 Python and Fortran adds around 7% to the run time. The remaining slowdown is likely
 556 related to choices in model architecture, such as shallow depth and sequential RNN steps,
 557 which lead to low GPU utilization ($< 10\%$). We believe that it will be possible to design
 558 ML emulators of more complex microphysical schemes that are more speed-competitive
 559 with the Fortran code which they aim to replace.

560 5 Conclusions

561 We have successfully developed an emulator to replace a simple Fortran microphysics
 562 scheme (Zhao-Carr) in FV3GFS, which controls grid-scale condensation (`gscond`) and
 563 precipitation (`precpd`) processes. Our findings demonstrate that when used online as a
 564 replacement for the Fortran scheme, the emulator maintains high skill ($\geq 94\%$) with low
 565 global-average bias (on the order of 1% or less) and remains stable for at least one year
 566 of continuous simulation. To our knowledge, this is the first successful emulation of a bulk
 567 microphysics scheme, and the first successful online emulation of a fast-timescale atmo-
 568 spheric parameterization central to global atmospheric forecasting.

569 A key contributor to the success of our emulator was tailoring its architecture to
 570 the underlying physical processes. By creating separate emulators for `gscond` and `precpd`,
 571 we enforce grid point locality and conservation for the condensation scheme, and we use
 572 an RNN to impose downward dependence in the atmosphere associated with falling pre-
 573 cipitation. This greatly improves the emulator’s skill, especially when used online. Adding
 574 an activity classifier to the condensation emulator alleviated issues of excess condensate
 575 related to the discrete-continuous nature of the tendencies and field outputs. Using a temperature-
 576 scaled conditional loss function for the `gscond` emulator and providing re-scaled inputs
 577 to all emulators helped maintain skill across the high dynamic range of condensate and
 578 humidity tendencies that must be accurately predicted to simulate cloud processes through-
 579 out the global atmosphere.

580 As with any ML-based emulation problem, achieving perfection is difficult, and the
 581 current scheme is no exception. In 1-year online integrations, biases develop in the po-
 582 lar stratospheric temperature and humidity fields. These regions challenge the ML train-
 583 ing because they have distinctly different local environments than the rest of the atmo-
 584 sphere and comprise a small fraction of the emulator’s training data. Further improve-
 585 ments could clearly be made, but are beyond the scope of this paper, which was to demon-
 586 strate the feasibility of a skillful ML microphysics emulator for online use. For instance,
 587 a natural possibility that we did not have time to implement would be to explicitly pre-
 588 dict precipitation flux at every model interface, which carries all the nonlocality in the
 589 microphysics. The hidden state of the `precpd` RNN is a skillful but imprecise proxy for
 590 this design, causing potential biases and drifts because physical constraints are imper-
 591 fectly respected (e.g., that the evaporation of precipitation in any model level cannot ex-
 592 ceed the downward flux of precipitation into that model level).

593 A compounding difficulty in the present work and generally for physics emulation
 594 is the inability to train emulation schemes directly in the context of their deployment
 595 within an atmospheric model. Fortran tooling for ML applications is challenging com-
 596 pared to the Python, but is still required for current atmospheric models. We utilize a
 597 Python package (`call_py_fort`) that provides an exceptional solution for interactive pro-
 598 totyping, but is not optimized for computational efficiency. Modeling frameworks on the
 599 horizon may simplify this process of ML integration and speed the development path to
 600 emulators that perform well online (Schneider et al., 2017; Dahm et al., 2023).

601 Our results stress the importance of evaluating the online performance for any pro-
 602 posed emulator, as it is straightforward to produce skillful offline models that may not
 603 perform well when integrated back into the model. It is also important to recognize that
 604 the development of emulators that perform well online is a challenging and time-consuming
 605 endeavor. If efficiency is the only goal, it may sometimes be more practical to invest in
 606 porting existing codes to run on GPUs, for example, as emulation requires significant
 607 human effort and problem-specific tuning.

608 Despite the challenges, our method and results are a proof-of-concept that machine
 609 learning techniques can effectively emulate fast physical processes central to the dynam-
 610 ics in weather and climate models. While our focus has been on a specific microphysics
 611 parameterization, we hope that the illustration of our problem-specific decisions will in-
 612 form the application to similar or more complex physical schemes. With further research
 613 and development, emulation techniques can continue to contribute to improved skill and
 614 efficiency of weather and climate models.

615 Appendix A Zhao-Carr Microphysics

616 This scheme handles both phase changes—condensation and evaporation—and pre-
 617 cipitation processes. Tendencies due to the former are typically 10x larger in magnitude.

618 The prognostic variables used by the scheme are the temperature T , specific humidity
619 q , and a combined cloud water/ice mixing ratio c .

620 The gscond scheme handles evaporation of cloud and condensation. Evaporation
621 of cloud is given by $E_c = \frac{1}{\Delta t} \max[\min[q_s(f_0 - f), c], 0]$. f is relative humidity. f_0 is a
622 critical relative humidity threshold which Zhao and Carr (1997) describe as “empirically
623 set to 0.75 over land and 0.90 over ocean.” q_s is the saturation specific humidity.

624 Condensation C_g on the other hand is given by a more complex formula involving
625 a relative humidity tendency. See Eq. (8) of Zhao and Carr (1997). Both formulas de-
626 pend only on the thermodynamic state of a single (x, y, z) location, but there is some
627 non-local dependence on the assumed phase of the cloud and the corresponding latent
628 heating rate.

The precpd scheme handles the conversion of cloud into rain/snow and the evap-
oration of the latter as it falls through the atmosphere. Broadly speaking, it can be writ-
ten as the following:

$$\begin{aligned} E_{rr} &= E_r(T, f, P_r) \\ E_{rs} &= E_r(T, f, P_s) \\ P &= P(T, f, c, P_r, P_s) \\ P_{sm} &= P_{sm}(T, f, c, P_r, P_s) \\ P_r &= \int_{p_t}^P (P - E_{rr}) dp/g \\ P_s &= \int_{p_t}^P (P_{sm} - E_{rs}) dp/g. \end{aligned}$$

629 Most of the formulas are proportional to rainfall P_r and snowfall P_s rates at a given level,
630 though are some rate constants that depend exponentially on temperature. p_t is the pres-
631 sure at the top of the atmosphere.

632 Appendix B Precpd emulator skill dropouts

633 Over the course of refining the emulation methodology, we observed larger variabil-
634 ity in the online skill scores of cloud and precipitation predictions, despite minimal-or-
635 no changes in emulator training or runtime configuration. In this section, we discuss the
636 primary source of that variance, which we refer to as skill dropouts. As an example, Fig-
637 ure B1 displays the surface precipitation skill over time for two 1-year simulations. When
638 the top 5 layer increment mask is adjusted from application to only gscond to both gscond
639 and precpd, the severity of skill dropouts decreases markedly.

640 Upon closer examination of the skill dropouts, the precpd emulator appears to be
641 the source of the issue. We focus on the dropout about 6 months into the gscond-only
642 masking experiment to illustrate this point. In this case, a cluster of columns near the
643 Maritime Continent is responsible for most of skill reduction. By removing the five grid
644 columns with the largest tendency errors, the overall snapshot skill goes from approx-
645 imately 0% to over 70%. When examining the tendency profiles from the column with
646 the largest errors (Fig. B2), the gscond emulator largely matches the diagnostic Fortran,
647 while the precpd emulator completely misses the autoconversion of condensate to pre-
648 cipitation in middle-and-upper levels. Leading up to this time step, we have confirmed
649 that gscond remains skillfull, while precpd skill degrades (not shown). The gscond em-
650 ulator retaining skill throughout this event suggests that a process outside of the ZC scheme,
651 such as deep convection, adds condensate throughout the column. The precpd emula-
652 tor then fails to precipitate the added condensate.

653 Overall, we hypothesize that the skill dropouts are associated with training data
654 insufficiency related to intense convection and/or unconstrained sensitivities of the RNN

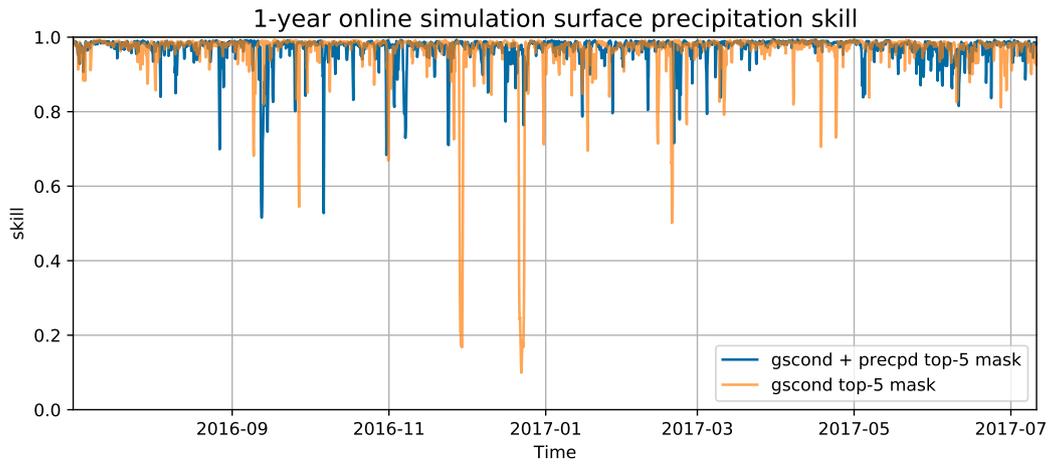


Figure B1. Surface precipitation skill over the 1-year online simulation for two increment masking configurations: (orange) gscond-only top 5 layer masking and (blue) gscond and precpd top 5 layer masking.

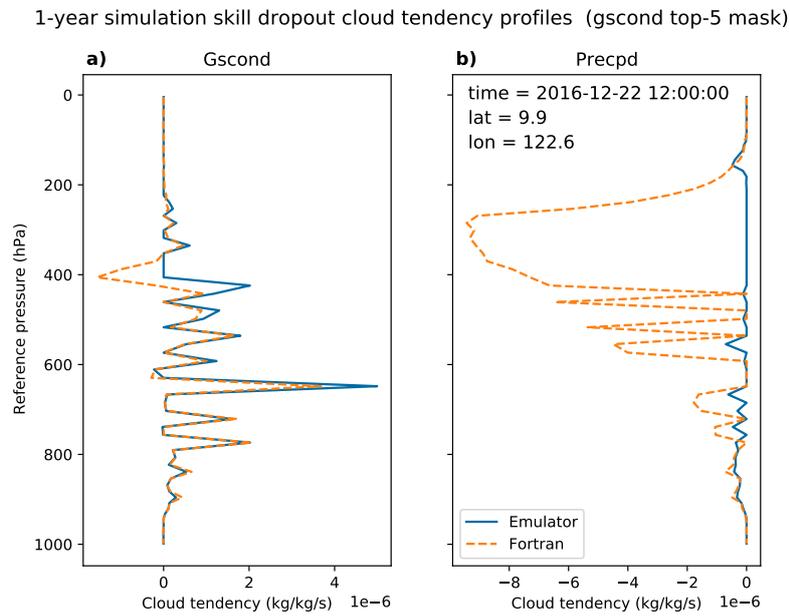


Figure B2. Vertical tendency profiles from the (a) gscond and (b) precpd schemes during the December 22nd 12 UTC skill dropout event in the gscond top 5 layer increment mask 1-year simulation. Each subcomponent panel shows the condensate tendencies predicted from the emulator (blue) and the diagnostic Fortran (orange dashed) for the selected column with the largest errors.

655 to upper-level inputs. It is encouraging that despite the magnitude of the misses, the ZC
 656 emulators resolve the issue in a few days or less for all cases observed (e.g., see Fig. S6).
 657 We also note that the dropouts tend to be confined to only a few grid columns, typically
 658 occurring in the tropics or subtropics. The isolated spatial extent of the skill dropout
 659 sources highlights the challenge in achieving consistently high skill in our chosen met-
 660 rics throughout the simulations. It also demonstrates how quickly the skill can deteri-
 661 orate if even a few predictions degrade.

662 Glossary

663 **dense-local** An MLP that takes in a single vertical level of inputs from a column and
 664 produces outputs for that same level. The vertical independence makes it "local".
 665 **gscond** The gridscale condensation component of Zhao-Carr microphysics
 666 **precpd** The precipitation component of Zhao-Carr microphysics
 667 **skill dropout** A temporary reduction in the online skill metric calculated between the
 668 emulator tendencies and the diagnostic Fortran tendencies

669 Acronyms

670 **ML** machine learning
 671 **MLP** multi-layer perceptron (feed-forward neural net)
 672 **RNN** recurrent neural net
 673 **ZC** Zhao-Carr

674 Appendix C Open Research

675 The code and configurations used to produce training data, train ML models, and
 676 run FV3GFS simulations are available on Github (<https://github.com/ai2cm/zc-emulation-manuscript>)
 677 and archived on Zenodo (<https://doi.org/10.5281/zenodo.7976184>). The data and
 678 docker images to reproduce results with the code are available on Zenodo (<https://doi.org/10.5281/zenodo.7976184>).

679 Acknowledgments

680 We thank the Allen Institute for Artificial Intelligence for their support of this work and
 681 for hosting Jacqueline Nugent as a summer intern during a portion of this project. We
 682 would also like to acknowledge NOAA-GFDL, NOAA-EMC, and the UFS community
 683 for providing publicly available code and data to initialize and run the FV3GFS atmo-
 684 spheric model.

685 References

- 686 Brenowitz, N. D., & Bretherton, C. S. (2019). Spatially Extended Tests of a Neu-
 687 ral Network Parametrization Trained by Coarse-Graining. *Journal of Ad-
 688 vances in Modeling Earth Systems*, 11(8), 2728–2744. Retrieved 2023-05-25,
 689 from <https://onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001711>
 690 (_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2019MS001711>) doi:
 691 10.1029/2019MS001711
- 692 Bretherton, C. S., Henn, B., Kwa, A., Brenowitz, N. D., Watt-Meyer, O.,
 693 McGibbon, J., ... Harris, L. (2022). Correcting Coarse-Grid
 694 Weather and Climate Models by Machine Learning From Global Storm-
 695 Resolving Simulations. *Journal of Advances in Modeling Earth Sys-
 696 tems*, 14(2), e2021MS002794. Retrieved 2023-05-25, from [https://
 697 onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002794](https://onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002794) (_eprint:

- 698 <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2021MS002794> doi:
699 10.1029/2021MS002794
- 700 Chantry, M., Hatfield, S., Dueben, P., Polichtchouk, I., & Palmer, T. (2021).
701 Machine Learning Emulation of Gravity Wave Drag in Numerical
702 Weather Forecasting. *Journal of Advances in Modeling Earth Sys-*
703 *tems*, 13(7), e2021MS002477. Retrieved 2023-05-25, from [https://](https://onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002477)
704 onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002477 (_eprint:
705 <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2021MS002477>) doi:
706 10.1029/2021MS002477
- 707 Chevallier, F., Chérury, F., Scott, N. A., & Chédin, A. (1998, November). A Neu-
708 ral Network Approach for a Fast and Accurate Computation of a Longwave
709 Radiative Budget. *Journal of Applied Meteorology and Climatology*, 37(11),
710 1385–1397. Retrieved 2023-05-25, from [https://journals.ametsoc.org/](https://journals.ametsoc.org/view/journals/apme/37/11/1520-0450_1998_037_1385_anna.2.0.co.2.xml)
711 [view/journals/apme/37/11/1520-0450_1998_037_1385_anna.2.0.co.2.xml](https://journals/apme/37/11/1520-0450_1998_037_1385_anna.2.0.co.2.xml)
712 (Publisher: American Meteorological Society Section: Journal of Applied
713 Meteorology and Climatology) doi: 10.1175/1520-0450(1998)037<1385:
714 ANNAFA>2.0.CO;2
- 715 Clark, S. K., Brenowitz, N. D., Henn, B., Kwa, A., McGibbon, J., Perkins,
716 W. A., ... Harris, L. M. (2022). Correcting a 200 km Resolution Cli-
717 mate Model in Multiple Climates by Machine Learning From 25 km
718 Resolution Simulations. *Journal of Advances in Modeling Earth Sys-*
719 *tems*, 14(9), e2022MS003219. Retrieved 2023-05-25, from [https://](https://onlinelibrary.wiley.com/doi/abs/10.1029/2022MS003219)
720 onlinelibrary.wiley.com/doi/abs/10.1029/2022MS003219 (_eprint:
721 <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2022MS003219>) doi:
722 10.1029/2022MS003219
- 723 Dahm, J., Davis, E., Deconinck, F., Elbert, O., George, R., McGibbon, J., ...
724 Fuhrer, O. (2023, May). Pace v0.2: a Python-based performance-portable
725 atmospheric model. *Geoscientific Model Development*, 16(9), 2719–2736. Re-
726 trieved 2023-05-25, from [https://gmd.copernicus.org/articles/16/2719/](https://gmd.copernicus.org/articles/16/2719/2023/)
727 [2023/](https://gmd.copernicus.org/articles/16/2719/2023/) (Publisher: Copernicus GmbH) doi: 10.5194/gmd-16-2719-2023
- 728 Gettelman, A., Gagne, D. J., Chen, C.-C., Christensen, M. W., Lebo,
729 Z. J., Morrison, H., & Gantos, G. (2021). Machine Learning the
730 Warm Rain Process. *Journal of Advances in Modeling Earth Sys-*
731 *tems*, 13(2), e2020MS002268. Retrieved 2023-05-25, from [https://](https://onlinelibrary.wiley.com/doi/abs/10.1029/2020MS002268)
732 onlinelibrary.wiley.com/doi/abs/10.1029/2020MS002268 (_eprint:
733 <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2020MS002268>) doi:
734 10.1029/2020MS002268
- 735 Grabowski, W. W. (2019, September). Separating physical impacts from natural
736 variability using piggybacking technique. *Advances in Geosciences*, 49, 105–
737 111. Retrieved 2023-05-25, from [https://adgeo.copernicus.org/articles/](https://adgeo.copernicus.org/articles/49/105/2019/)
738 [49/105/2019/](https://adgeo.copernicus.org/articles/49/105/2019/) doi: 10.5194/adgeo-49-105-2019
- 739 Harris, L. M., & Lin, S.-J. (2013, January). A Two-Way Nested Global-Regional
740 Dynamical Core on the Cubed-Sphere Grid. *Monthly Weather Review*, 141(1),
741 283–306. Retrieved 2023-05-25, from [https://journals.ametsoc.org/view/](https://journals.ametsoc.org/view/journals/mwre/141/1/mwr-d-11-00201.1.xml)
742 journals/mwre/141/1/mwr-d-11-00201.1.xml (Publisher: American Meteoro-
743 logical Society Section: Monthly Weather Review) doi: 10.1175/MWR-D-11
744 -00201.1
- 745 Keller, C. A., & Evans, M. J. (2019, March). Application of random forest regression
746 to the calculation of gas-phase chemistry within the GEOS-Chem chemistry
747 model v10. *Geoscientific Model Development*, 12(3), 1209–1225. Retrieved
748 2023-05-25, from <https://gmd.copernicus.org/articles/12/1209/2019/>
749 (Publisher: Copernicus GmbH) doi: 10.5194/gmd-12-1209-2019
- 750 Kelp, M. M., Jacob, D. J., Lin, H., & Sulprizio, M. P. (2022). An Online-
751 Learned Neural Network Chemical Solver for Stable Long-Term Global
752 Simulations of Atmospheric Chemistry. *Journal of Advances in Model-*

- 753 *ing Earth Systems*, 14(6), e2021MS002926. Retrieved 2023-05-25, from
 754 <https://onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002926>
 755 (_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2021MS002926>)
 756 doi: 10.1029/2021MS002926
- 757 Krasnopolsky, V. M., Fox-Rabinovitz, M. S., & Chalikov, D. V. (2005, May).
 758 New Approach to Calculation of Atmospheric Model Physics: Accurate
 759 and Fast Neural Network Emulation of Longwave Radiation in a Climate
 760 Model. *Monthly Weather Review*, 133(5), 1370–1383. Retrieved 2023-
 761 05-25, from [https://journals.ametsoc.org/view/journals/mwre/133/
 762 5/mwr2923.1.xml](https://journals.ametsoc.org/view/journals/mwre/133/5/mwr2923.1.xml) (Publisher: American Meteorological Society Section:
 763 Monthly Weather Review) doi: 10.1175/MWR2923.1
- 764 Krasnopolsky, V. M., Fox-Rabinovitz, M. S., Hou, Y. T., Lord, S. J., & Belochitski,
 765 A. A. (2010, May). Accurate and Fast Neural Network Emulations of Model
 766 Radiation for the NCEP Coupled Climate Forecast System: Climate Simula-
 767 tions and Seasonal Predictions. *Monthly Weather Review*, 138(5), 1822–1842.
 768 Retrieved 2023-05-25, from [https://journals.ametsoc.org/view/journals/
 769 mwre/138/5/2009mwr3149.1.xml](https://journals.ametsoc.org/view/journals/mwre/138/5/2009mwr3149.1.xml) (Publisher: American Meteorological Soci-
 770 ety Section: Monthly Weather Review) doi: 10.1175/2009MWR3149.1
- 771 O’Gorman, P. A., & Dwyer, J. G. (2018). Using Machine Learning to Pa-
 772 rameterize Moist Convection: Potential for Modeling of Climate, Cli-
 773 mate Change, and Extreme Events. *Journal of Advances in Modeling
 774 Earth Systems*, 10(10), 2548–2563. Retrieved 2023-05-25, from [https://
 775 onlinelibrary.wiley.com/doi/abs/10.1029/2018MS001351](https://onlinelibrary.wiley.com/doi/abs/10.1029/2018MS001351) (_eprint:
 776 <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2018MS001351>) doi:
 777 10.1029/2018MS001351
- 778 Rasp, S., Pritchard, M. S., & Gentine, P. (2018, September). Deep learning to repre-
 779 sent subgrid processes in climate models. *Proceedings of the National Academy
 780 of Sciences*, 115(39), 9684–9689. Retrieved 2023-05-25, from [https://
 781 www.pnas.org/doi/full/10.1073/pnas.1810286115](https://www.pnas.org/doi/full/10.1073/pnas.1810286115) (Publisher: Proceedings
 782 of the National Academy of Sciences) doi: 10.1073/pnas.1810286115
- 783 Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth System
 784 Modeling 2.0: A Blueprint for Models That Learn From Observations
 785 and Targeted High-Resolution Simulations. *Geophysical Research Let-
 786 ters*, 44(24), 12,396–12,417. Retrieved 2023-05-25, from [https://
 787 onlinelibrary.wiley.com/doi/abs/10.1002/2017GL076101](https://onlinelibrary.wiley.com/doi/abs/10.1002/2017GL076101) (_eprint:
 788 <https://onlinelibrary.wiley.com/doi/pdf/10.1002/2017GL076101>) doi:
 789 10.1002/2017GL076101
- 790 Schreck, J. S., Becker, C., Gagne, D. J., Lawrence, K., Wang, S., Mouchel-
 791 Vallon, C., ... Hodzic, A. (2022). Neural Network Emulation of
 792 the Formation of Organic Aerosols Based on the Explicit GECKO-
 793 A Chemistry Model. *Journal of Advances in Modeling Earth Sys-
 794 tems*, 14(10), e2021MS002974. Retrieved 2023-05-25, from [https://
 795 onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002974](https://onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002974) (_eprint:
 796 <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2021MS002974>) doi:
 797 10.1029/2021MS002974
- 798 Ukkonen, P., Pincus, R., Hogan, R. J., Pagh Nielsen, K., & Kaas, E. (2020). Ac-
 799 celerating Radiation Computations for Dynamical Models With Targeted
 800 Machine Learning and Code Optimization. *Journal of Advances in Mod-
 801 eling Earth Systems*, 12(12), e2020MS002226. Retrieved 2023-05-25, from
 802 <https://onlinelibrary.wiley.com/doi/abs/10.1029/2020MS002226>
 803 (_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2020MS002226>)
 804 doi: 10.1029/2020MS002226
- 805 Veerman, M. A., Pincus, R., Stoffer, R., van Leeuwen, C. M., Podareanu, D., &
 806 van Heerwaarden, C. C. (2021, February). Predicting atmospheric optical
 807 properties for radiative transfer computations using neural networks. *Philo-*

- 808 *sophical Transactions of the Royal Society A: Mathematical, Physical and*
809 *Engineering Sciences*, 379(2194), 20200095. Retrieved 2023-05-25, from
810 <https://royalsocietypublishing.org/doi/10.1098/rsta.2020.0095>
811 (Publisher: Royal Society) doi: 10.1098/rsta.2020.0095
- 812 Yuval, J., & O’Gorman, P. A. (2020, July). Stable machine-learning parameteri-
813 zation of subgrid processes for climate modeling at a range of resolutions. *Na-*
814 *ture Communications*, 11(1), 3295. Retrieved 2023-05-25, from [https://www](https://www.nature.com/articles/s41467-020-17142-3)
815 [.nature.com/articles/s41467-020-17142-3](https://www.nature.com/articles/s41467-020-17142-3) (Number: 1 Publisher: Nature
816 Publishing Group) doi: 10.1038/s41467-020-17142-3
- 817 Zhao, Q., & Carr, F. H. (1997, August). A Prognostic Cloud Scheme for Oper-
818 ational NWP Models. *Monthly Weather Review*, 125(8), 1931–1953. Re-
819 trieved 2023-01-30, from [https://journals.ametsoc.org/view/journals/](https://journals.ametsoc.org/view/journals/mwre/125/8/1520-0493_1997_125_1931_apcsfo_2.0.co_2.xml)
820 [mwre/125/8/1520-0493_1997_125_1931_apcsfo_2.0.co_2.xml](https://journals.ametsoc.org/view/journals/mwre/125/8/1520-0493_1997_125_1931_apcsfo_2.0.co_2.xml) (Publisher:
821 American Meteorological Society Section: Monthly Weather Review) doi:
822 10.1175/1520-0493(1997)125<1931:APCSFO>2.0.CO;2