

1 **Probing the Skill of Random Forest Emulators for**
2 **Physical Parameterizations via a Hierarchy of Simple**
3 **CAM6 Configurations**

4 **Garrett C. Limon¹, Christiane Jablonowski¹**

5 ¹Department of Climate and Space Sciences and Engineering, University of Michigan, Ann Arbor, MI

6 **Key Points:**

- 7 • Random forests skillfully emulate simple physics schemes within NCAR's CAM6
8 atmosphere model
9 • Hierarchical approach probes the limitations of emulators as complexity increases

Corresponding author: Garrett C. Limon, glimon@umich.edu

10 **Abstract**

11 Machine learning approaches, such as random forests, have been used to effectively
 12 emulate various aspects of climate and weather models in recent years. The limitations
 13 to these approaches are not yet known, particularly with regards to varying complex-
 14 ity of the physical parameterization schemes being emulated within the climate model.
 15 Utilizing a hierarchy of model configurations, we explore the limitations of random for-
 16 est emulator skill using simplified model frameworks within NCAR’s Community Atmo-
 17 sphere Model, version 6 (CAM6). These include a dry model configuration, a moist ex-
 18 tension of the dry model, and an extension of the moist case that includes an additional
 19 convection scheme. With unique random forests being optimized for each tendency or
 20 precipitation rate across the hierarchy, we create a variety of emulators. Each model con-
 21 figuration is run with identical resolution and over the same time period. The models
 22 are then evaluated against the CAM6 output. All models show significant skill across
 23 each random forest emulator, often in-line with or exceeding similar approaches within
 24 the literature. In addition, as the CAM6 complexity is increased, the random forest skill
 25 noticeably decreases, regardless of the extensive tuning and training process each ran-
 26 dom forest goes through. This indicates a limit on the feasibility of random forests to
 27 act as physics emulators in climate models and encourages further exploration in order
 28 to identify that limit in the context of state-of-the-art climate model configurations.

29 **Plain Language Summary**

30 Machine learning has become an intriguing technique for replacing complicated as-
 31 pects of climate and weather models known as parameterizations, which account for pro-
 32 cesses like cloud interactions and rain. However, the limitations of machine learning tech-
 33 niques are not yet fully understood. We explore these limits using a specific machine learn-
 34 ing method and simplified climate model frameworks. The machine learning models are
 35 then carefully analyzed against the original climate model results. All of our machine
 36 learning models show impressive skill at recreating the original results. However, that
 37 skill is shown to noticeably decrease as the complexity of the climate model framework
 38 is increased. While this may be expected, it still indicates a limit on the feasibility of
 39 machine learning techniques to substitute for the complicated parameterizations within
 40 state-of-the-art climate models. Further investigation is needed to understand the via-
 41 bility of these methods being adopted into the simulation of the Earth system.

42 **1 Introduction**

43 In recent decades machine learning (ML) has become an intriguing tool for atmo-
 44 spheric scientists. It provides the unique ability to bridge data science with the phys-
 45 ical sciences in order to improve our understanding of the Earth system (Reichstein et
 46 al., 2019; Boukabara et al., 2021). While ML is still a relatively novel approach to ap-
 47 plications in climate science, there is already an abundance of research utilizing these
 48 techniques. Some examples include identifying mixed layer depths in the ocean via ob-
 49 servations (Foster et al., 2021), attributing model biases from physics-dynamics coupling
 50 in climate models (Yorgun & Rood, 2016), improving severe hail predictions over the US
 51 high plains (Gagne et al., 2017), post-processing bias corrections of weather forecasts (Chapman
 52 et al., 2019), and implementing corrective schemes like ‘nudging’ physics tendencies via
 53 coarse-graining or hindcasting (Bretherton et al., 2022; Watt-Meyer et al., 2021).

54 General Circulation Models (GCMs) are made up of a dynamical core, responsi-
 55 ble for the geophysical fluid flow calculations, and physical parameterization schemes,
 56 which estimate subgrid-scale processes that the dynamical core does not resolve. The
 57 latter are a source of significant bias and model uncertainty due to the heuristic nature
 58 of their development (Held, 2005; Stevens & Bony, 2013; Hourdin et al., 2017). Begin-

ning with the work of Krasnopolsky and Fox-Rabinovitz (2006) applying neural networks to climate and weather prediction model development, ML became an attractive candidate for augmenting the subgrid-scale physics schemes within weather and climate models. In recent years, ML techniques have already been shown to be capable of replicating parameterizations schemes to various degrees of effectiveness (Beucler et al., 2019; Yuval et al., 2020). Specifically, Ukkonen (2022) were able to develop ML emulators for radiative transfer processes, O’Gorman and Dwyer (2018) and Gentine et al. (2018) used ML to emulate moist convection processes, Gettelman et al. (2021) utilized neural networks to emulate a component in the micro-physics scheme within a GCM, Chantry et al. (2021) developed a nonorographic gravity wave drag emulator, and Rasp et al. (2018) and Brenowitz and Bretherton (2018) tackled a full physics emulator of cloud-resolving and near-global aquaplanet simulations, respectively, via neural networks. These are just a few examples showing both the promise of ML emulation and some limitations, particularly in regards to model stability and physical realism (Beucler et al., 2019; Yuval et al., 2021).

Our work is inspired by many of these recent studies into ML emulation for parameterization schemes, with a focus on multiple simplified physics configurations within version 6 of the Community Atmosphere Model (CAM6). CAM6 is the atmospheric GCM within the Community Earth System Model (CESM) (Danabasoglu et al., 2020) framework, developed by the National Center for Atmospheric Research (NCAR). In particular, we utilize a hierarchy of three physical forcing complexities, each with a well-defined increase in non-linearity associated with its mathematical expressions. The parameterization schemes begin with a dry model setup, referred to as HS hereon and described in Held and Suarez (1994). This is followed by a moist version of the HS scheme developed by Thatcher and Jablonowski (2016), referred to as TJ. Lastly, a modified version of the TJ scheme is used in which we couple a simple Betts-Miller (BM) convection scheme to the physics processes (Betts & Miller, 1986; Frierson, 2007). These three parameterization packages may be referred to throughout the papers as dry, moist, and convection, respectively. None of these physics schemes include topography or seasonal and diurnal cycles. We also keep our ML technique consistent, using random forests (Breiman, 1996). This allows for an investigation into the fundamental relationship between the degree of non-linearity within the parameterization scheme and the corresponding complexity of the random forest to effectively emulate the forcing.

In this work, we show that various physical forcing tendencies and the precipitation rate can be emulated by random forest models in an offline mode. We investigate how their skill depends on model complexity, as well as whether we can identify indications of limits on the feasibility of random forests for use in more complex model setups. In many cases, our ML models are shown to be highly skilled, both from a statistical perspective and from direct comparisons. We begin with an explanation of the three model configurations, our model run setup and data processing steps, and a background discussion on ML and the random forest techniques in section 2. This is followed by our results and discussion in section 3 before culminating with concluding thoughts in section 4.

2 Methods

2.1 CAM6 Configurations

2.1.1 Dry Scheme

The dry scheme is based on two forcing mechanisms as described in HS. The dissipation of the horizontal wind is represented by Rayleigh friction at the lower levels of the model (below 700 hPa),

$$\frac{\partial \vec{v}_h}{\partial t} = -\frac{1}{k_v(p)} \vec{v}_h \quad (1)$$

108 and radiation is mimicked by a Newtonian temperature relaxation described by

$$\left(\frac{\partial T}{\partial t}\right)_{\text{HS}} = -\frac{1}{k_T(\phi, p)}[T - T_{\text{eq}}(\phi, p)] \quad (2)$$

109 Here, $\partial/\partial t$ represents a sub-grid physics tendency (forcing) of a variable over a physics
 110 time step, \vec{v}_h is the horizontal velocity vector, T is the temperature, k_v and k_T are the
 111 dissipation coefficient and the relaxation coefficient, respectively, T_{eq} is a pre-defined equi-
 112 librium temperature profile, p is the pressure, and ϕ is the latitude. These forcings are
 113 coupled to the dry dynamical core and produce stable atmospheric fluid flow, trigger-
 114 ing quasi-realistic processes such as Rossby waves in the mid-latitudes. This model con-
 115 figuration comes implemented within CAM6's Simpler Models framework and is set with
 116 the 'FHS94' compset choice.

117 **2.1.2 Moist Scheme**

118 The moist physics scheme is similarly forced by Rayleigh friction and Newtonian
 119 temperature relaxation, this time with a modified equilibrium temperature profile and
 120 additional forcing mechanisms. These include heating and cooling due to large-scale con-
 121 densation, surface fluxes of latent and sensible heat, and a planetary boundary layer (PBL)
 122 mixing scheme via second order diffusion (Thatcher & Jablonowski, 2016), abbreviated
 123 as TJ16 later. The temperature forcing then takes the form

$$\left(\frac{\partial T}{\partial t}\right)_{\text{TJ}} = -\frac{1}{k_T(\phi, p)}[T - \tilde{T}_{\text{eq}}(\phi, p)] + \frac{L}{c_p}C + \frac{C_H|\vec{v}_a|(T_s - T_a)}{z_a} + \text{PBL Diffusion} \quad (3)$$

124 where \tilde{T}_{eq} is a modified equilibrium profile, L is the latent heat of vaporization, C is the
 125 large-scale condensation rate, c_p is the specific heat at constant pressure, C_H is the trans-
 126 fer coefficient for sensible heat, $|\vec{v}_a|$ is the horizontal wind speed at the lowest model level,
 127 T_s is the surface temperature, T_a is the temperature of the lowest model level, and z_a
 128 is the height of the lowest model level; the latter five are needed for the computation of
 129 the sensible heat flux at the surface. The mathematical details of the PBL diffusion of
 130 T are left out here and can be found in TJ16 and Reed and Jablonowski (2012). This
 131 model setup is similarly implemented within the Simpler Model options in CAM6 via
 132 the 'FTJ16' compset, which assumes an ocean-covered lower boundary with a prescribed
 133 sea surface temperature.

134 The inclusion of moisture brings an additional forcing tendency for specific humid-
 135 ity, which is similarly impacted by the large-scale condensation rate, the latent heat flux
 136 at the surface, and PBL diffusion

$$\left(\frac{\partial q}{\partial t}\right)_{\text{TJ}} = -C + \frac{C_E|\vec{v}_a|(q_{\text{sat},s} - q_a)}{z_a} + \text{PBL diffusion} \quad (4)$$

137 Here, q refers to the specific humidity in the atmosphere, C_E is the bulk transfer coef-
 138 ficient for water vapor, $q_{\text{sat},s}$ is the saturation specific humidity at the surface, and q_a
 139 is the specific humidity at the lowest model level. Again, mathematical details of the PBL
 140 diffusion of q are provided in TJ16. Additionally we chose to emulate the large-scale pre-
 141 cipitation rate given by

$$P_{\text{ls}} = \frac{1}{\rho_{\text{water}}g} \int_0^{p_s} C dp \quad (5)$$

142 where ρ_{water} is the density of water, g is gravity and p_s is the surface pressure.

143 **2.1.3 Convection Scheme**

144 The final step in our hierarchy couples the BM convection scheme to the TJ setup
 145 (Betts, 1986; Betts & Miller, 1986; Frierson, 2007). This configuration is not built into
 146 the CAM6 Simpler Model framework and required some minor modifications to the TJ

147 setup. The simplified BM technique follows the description by Frierson (2007) and we
 148 recommend their paper for a more complete description. To summarize, the resulting
 149 tendencies with the addition of the BM convection scheme can be written as

$$\left(\frac{\partial T}{\partial t}\right)_{\text{BM}} = -\frac{T - T_{\text{ref}}}{\tau} + \left(\frac{\partial T}{\partial t}\right)_{\text{TJ}} \quad (6)$$

150

$$\left(\frac{\partial q}{\partial t}\right)_{\text{BM}} = -\frac{q - q_{\text{ref}}}{\tau} + \left(\frac{\partial q}{\partial t}\right)_{\text{TJ}} \quad (7)$$

151

152

153

where τ is the convective relaxation time and T_{ref} and q_{ref} are reference temperature and specific humidity profiles. Within our implementation, the BM scheme is calculated first, before the rest of the TJ scheme.

154

155

156

157

158

159

160

161

162

The convection scheme utilizes regimes of precipitation due to warming, P_T , and precipitation due to drying, P_q . When we are in the regime of $P_T > 0$ and $P_q > 0$, ‘convection’ is triggered. Frierson (2007) described in detail how extra steps are taken with regards to the reference profiles in order for the convection scheme to ensure a conservation of enthalpy in the deep convection regime. The author also describes three techniques to handling shallow convection; in our work we use the “shallower” scheme, in which the reference temperature is further modified in order to lower the depth at which shallow convection occurs, for the shallow convection calculations. This is considered the simplest technique within the BM scheme that still allows for shallow convection to occur.

163

164

165

166

167

168

169

170

171

172

173

The BM convection scheme has a dependency on two coefficients: the relative humidity threshold for the reference temperature profile (RH_{BM}) and τ , the convective relaxation time. In order to choose these values, we examined various profiles of a variety of fields and compared them to fields from the CAM6 aquaplanet model (Williamson et al., 2012; Medeiros et al., 2016). Details on the aquaplanet and how it was used to identify our choices of RH_{BM} and τ can be found in Supporting Information Text S1. The aquaplanet configuration acts as a loose reference for these choices as it is a widely used model configuration in which the planet’s surface is covered by an ocean. This allows for surface-to-ocean interactions to become an integral component of the underlying physics. It is useful for exploring many aspects of geophysical fluid flow in a controlled model setting. The chosen values were $\tau = 4$ hr and $\text{RH}_{\text{BM}} = 0.7$.

174

2.2 Machine Learning

175

176

177

178

179

180

181

182

183

184

Broadly speaking, there are two categories of ML applications: supervised and unsupervised learning. Unsupervised learning encompasses tasks that attempt to identify general patterns in data, for example, clustering algorithms. Supervised learning strives to identify correlations or functional relationships between a labeled input and output. There are two primary tasks that can be done with supervised learning: classification and regression; the latter is applicable to emulating physical parameterizations. Regression is the process of estimating a functional relationship between a dependent variable, referred to as the label or what we are predicting, and one or more independent variables, referred to as features or input variables. With this framework in mind, we can think of regression as the process of identifying the function $\hat{g}(\vec{X})$ such that

$$\hat{g}(\vec{X}) \approx f(\vec{X}) \quad (8)$$

185

186

where $f(\vec{X})$ is the function we seek to identify and \vec{X} is the vector of input variables (features).

187

188

189

190

191

What separates modern machine learning techniques like neural networks, support vector machines, and random forests are their applications to nonlinear systems, providing methods for nonlinear regression tasks. In its simplest form, a physical parameterization is a nonlinear function that describes a tendency or precipitation rate (dependent variable) given the (independent) state variables. In the analogy to equation

192 8, the tendency would be f while the state variables makes up the vector \vec{X} and our trained
 193 ML model will be $\hat{g}(\vec{X})$.

194 We utilize random forests to emulate the parameterization schemes. A random for-
 195 est is an ensemble of decision trees, which can themselves be considered an ML technique.
 196 Decision trees identify thresholds among a branch network, forming a structure of con-
 197 ditional operations that produce a prediction (Breiman, 1996). Random forests are com-
 198 monly used in classification applications of ML, but have been shown to be effective for
 199 nonlinear regression tasks in atmospheric science as well (O’Gorman & Dwyer, 2018).
 200 Various trees in the forest are initialized at random and are then trained along side each
 201 other. The final result is an ensemble average of the results from all trees in the forest.

202 A random forest approach was chosen due to both its relative simplicity as an ap-
 203 plication of non-linear regression, along with its ability to inherently preserve underly-
 204 ing physical properties of the predicted fields. Since each individual tree produces an out-
 205 put that is within the scope of the training data, their average is also inherently within
 206 the scope of the data. This means that random forests cannot extrapolate a prediction
 207 outside of the range established by their training data. In the context of using ML tech-
 208 niques for physical science applications, this is a welcome restriction because it can avoid
 209 potential artifacts that could be inconsistent with the physics at play. For example, a
 210 random forest will inherently adhere to the non-negative property of precipitation, as
 211 it will have never encountered negative precipitation in its training data. This is in con-
 212 trast to techniques such as neural networks, which historically have trouble with extrap-
 213 olation and adhering to underlying physical constraints (Beucler et al., 2019). We de-
 214 veloped a streamlined workflow from data generation to training, testing, and analysis
 215 by utilizing CAM6’s built-in Simplified Models physics framework along with the Python
 216 libraries Xarray and scikitlearn (Pedregosa et al., 2011; Hoyer & Hamman, 2017). Xar-
 217 ray allows for easily manipulating data in the netCDF format, while scikitLearn is a well-
 218 maintained ML library that includes user-friendly random forest implementations for Python.

219 2.3 Model Setup and Data Preparation

220 The simple model configurations allow us to generate large quantities of model out-
 221 put to train our machine learning models. Working with CAM6, we utilize a finite vol-
 222 ume dynamical core with 30 pressure-based vertical levels and a model top at roughly
 223 2.2 hPa. The exact placement of the model levels is specified in Reed and Jablonowski
 224 (2012) (see their Appendix B). The model is run for 60-years with a latitude-longitude
 225 grid of resolution $1.9^\circ \times 2.5^\circ$ - simply referred to as 2-degree resolution and corresponds
 226 to roughly 200 km grid spacing. We output data for state variables, including temper-
 227 ature, surface pressure, specific humidity, and the diagnostic quantity relative humid-
 228 ity, once every week of the simulation just before the prognostic states are updated by
 229 the physics package. Additionally, we output the tendencies due to the physical param-
 230 eterization package after they are updated with the same output frequency. This is an
 231 important modification since by default both the state variables and physical tenden-
 232 cies are output after the physics update. We chose to output once per week in order to
 233 avoid the close correlations between the time snapshots. Strong correlations are present
 234 in data snapshots that are only separated by short time intervals, such as a day. This
 235 allows for our data to include a larger range of the functional space, while avoiding re-
 236 dundancies within the scope of the training data. It should be reiterated that our con-
 237 figurations do not include a diurnal or seasonal cycle, which allows us to be able to take
 238 weekly output without risking an incomplete representation of the functional space. For
 239 more complicated systems, care would need to be taken in choosing output intervals that
 240 effectively sample the functional space.

241 Here, we define the input fields for our ML models to be the state variables used
 242 by the underlying schemes, such as temperature and pressure. Similarly, the output fields

243 are the resulting tendency or precipitation rate being predicted. For preprocessing, we
 244 focus primarily on the shape of the data, input choices, and the distribution of the data
 245 between training and testing. The state variables and tendencies, using temperature (T)
 246 as an example, are generally output from the model in the shape

$$T(N_{\text{time}}, N_{\text{lev}}, N_{\text{lat}}, N_{\text{lon}})$$

247 where N_{time} , N_{lev} , N_{lat} , and N_{lon} correspond to the number of temporal snapshots,
 248 vertical levels, latitudes, and longitudes, respectively. Some variables are surface fields,
 249 such as the precipitation rates, and correspond to $N_{\text{lev}} = 1$. Due to the nature of the
 250 physical parameterizations being column-wise implementations in the atmospheric model,
 251 we carry this over as our feature/label dimension. This means our number of samples
 252 becomes

$$N_{\text{samples}} = N_{\text{time}} \times N_{\text{lat}} \times N_{\text{lon}}$$

253 The number of features becomes

$$N_{\text{features}} = N_{\text{lev}} \times N_{\text{input fields}}$$

254 again, where input fields include temperature, specific humidity, relative humidity, and
 255 pressure, among others. The number of labels becomes

$$N_{\text{labels}} = N_{\text{lev}} \times N_{\text{output fields}} = N_{\text{lev}}$$

256 where $N_{\text{output fields}} = 1$ for all cases in this work since we train a unique random for-
 257 est for each predicted tendency or precipitation rate. This was a conscious decision that
 258 allows for a robust investigation into the effectiveness of random forests for these em-
 259 ulation tasks as the functional form slowly increases in complexity within our hierarchy.
 260 This is in contrast to other similar efforts, such as Rasp et al. (2018) and Yuval et al.
 261 (2020), wherein a single ML model is trained to predict all fields of interest.

262 Finally, we partition the data into training and testing subsets. The training data
 263 comes from the first 50 years of the 60-year model run. We choose a selection of roughly
 264 15-20 million samples, which represents the majority of the available data from the 50
 265 years for training. This number depends primarily on the complexity of the random for-
 266 est parameters and the size and shape of the variable. For example, the moisture ten-
 267 dency is zero above roughly 250 hPa, which means we have six levels between 250 hPa
 268 and the model top that can be omitted from the process, resulting in significantly less
 269 data to be processed. Likewise, the precipitation rate is a surface field, which leads to
 270 significantly reduced computational cost for training since $N_{\text{labels}} = N_{\text{lev}} = 1$. This
 271 allows us to use closer to $N_{\text{samples}} \approx 20$ million for these emulators, which is just be-
 272 low the upper limit of our generated data. In contrast, the moist and convective tem-
 273 perature tendencies use 15 and 12 million samples, respectively. This was the largest we
 274 could use for these cases while remaining within our computational resource limits. The
 275 discrepancy between these two cases is a result of the size and complexity of each individu-
 276 ally-optimized random forest. The number of samples used in training for each case is included
 277 in tables S1 to S8 in the Supporting Information.

278 The testing data are used to quantify our model's ability to emulate the param-
 279 eterization. The testing data were not available during training and come from the fi-
 280 nal six years of the 60-year model run. It is important to evaluate model performance
 281 on data that the random forest has not seen while training in order to ensure our em-
 282 ulators do not show signs of overfitting. Overfitting in ML occurs when the ML model
 283 has been trained well on the subset of data that it has seen, but is unable to general-
 284 ize to a new set of data from the same source. Lastly, the ML algorithms need to have
 285 their hyperparameters tuned in order to obtain an optimized random forest architecture
 286 for the problem. This is an important part of the ML workflow and we currently utilize
 287 the SHERPA hyperparameterization library to accomplish it (Hertel et al., 2020). Fur-
 288 ther details about the process of hyperparameter tuning and the final choices of hyper-
 289 parameters can be found in tables S1 to S8 in the Supporting Information.

3 Results & Discussion

3.1 Snapshots & Mean Fields

Figures (1) and (2) show horizontal snapshots of the CAM6 (left column) output, our random forest ML predictions (middle column) and their differences (right column) for temperature and moisture tendencies, respectively. From top to bottom, the figures show each of the various physics schemes used: dry (Fig. (1) only), moist, and convection. We chose a snapshot from a randomly chosen time step at the model level closest to 850 hPa. The differences calculated in all plots are truth (CAM) subtracted from the random forest prediction. This means that positive and negative values correspond to over- and underestimations by the random forest. The snapshots in Figs. (1) and (2) immediately show how effective ML methods can be at emulating simple parameterization schemes in climate models for any given time step. These temporal snapshots allow us to appreciate the agreement between the CAM output and the random forest predictions, while still being able to identify areas of discrepancy. These snapshots also show how at a given time step, the ML prediction can reproduce the flow properties associated with baroclinic waves in the mid-latitudes. This is apparent in the heating tendency (1e, 1h) along the frontal zones, as well as decreasing moisture levels (2b, 2e) in these areas, corresponding to precipitation bands.

Figures (3) and (4) are zonally and temporally averaged temperature and specific humidity tendencies over the testing period of the final six years. The CAM6 physics and the random forest results are visually indistinguishable. The order of magnitude of the difference plots (right columns) is quite small in both figures. Additionally, Fig. (5) shows the same averaged field for the precipitation rates. The CAM6 output (blue) and the random forest ML predictions (dashed orange) overlay each other almost perfectly. The top row shows the large-scale precipitation rate and the bottom row the convective precipitation rate, while the left column corresponds to the moist case and the right to the convection case. We can also identify some of the same physical behaviors that were obvious in the snapshots from Figs. (1) and (2), such as the heating bands in the mid-latitudes corresponding to the peaks around 40N & 40S in Figs. (3e) and (3h). In addition, the intense precipitation regions in the tropics are emulated well by the random forests as displayed in Figs. (5a, c). This precipitation is correlated with the intense tropical heating peaks in Figs. (3e, h).

The minor differences between ML predictions and the CAM6 output in Figs. (1 - 5) somewhat mirror minor artifacts that could arise through other common numerical changes to a GCM, such as dynamical core grid choices. Further, when we incorporate the zonal-mean time-means in Figs. (3), (4), and (5) these subtle discrepancies are averaged out and we obtain plots that become virtually indistinguishable for most cases. We also begin to see a hint that as we increase the complexity of the schemes, the random forest's skill begins to decrease. Figure (3) shows the difference plots between the predicted result and the original CAM6 output in the right column for all three cases of the temperature tendencies. The regions of larger magnitude differences for the convection case in plots (3f) and (3i) appear throughout a majority of the domain, while the dominant regions in the dry case in plot (3c) are more-so bound to the poles and equatorial regions. There is also an increase in magnitude throughout the difference plots between Figs. (3f) and (3i), as well as (4c) and (4f). This is another indication of a noticeable decrease in skill as complexity increases.

In Fig. (5), the emulated precipitation rates are even less distinguishable in the mean fields. The various peaks in the zonal-mean time-mean plots in Fig. (5) align closely with the areas of 'drying' in Fig. (4). This is in particular true for the equatorial region in both cases, dominant in the moist case, as well as in the mid-latitudes in the convection case. We also notice that there is a negligible difference in performance between the moist and convection cases' large-scale precipitation emulator. This is due to the fact that by adding

the BM convection scheme to the moist physics, we do not impact the calculation of the large-scale precipitation. Instead, the resulting large-scale precipitation rate in the convection case is impacted only by the fact that the convection scheme, which is called first, has already removed a significant amount of moisture from the atmosphere. Therefore the overall amount of precipitation that accumulates from the large-scale scheme is less and more concentrated in the regions that did not meet the criteria for convection as described in the BM scheme. Mathematically, the large-scale precipitation scheme has not changed and we can see that the random forest maintains virtually identical skill in this case across the two schemes.

3.2 Point-wise Comparison

Next, we show one-to-one scatter plots of the results from CAM and the random forests in Figs. (6) and (7) for the temperature and specific humidity tendencies at a level near 850 hPa and precipitation rates, respectively. This is a metric that allows for an effective visualization of the spread of our predictions. If the emulator were to produce the exact results as the original scheme, the points on these plots would follow the one-to-one line $y = x$, shown in orange. One-to-one scatter plots have been shown in related papers, such as O’Gorman and Dwyer (2018), Rasp et al. (2018), and Han et al. (2020) for various metrics and fields. Figure (6) contains the temperature tendencies in the top row and the moisture in the bottom row for both the moist case (left column) and convection case (right column). Figure (7) shows the scatter plots for each precipitation rate, oriented in the same configuration as Fig. (5) above. Each plot in Figs. (6) and (7) contains the $y = x$ (one-to-one) line in orange along with a least squares linear fit in blue. The least squares is calculated via the Python library NumPy and is used here to illustrate how closely the predictions align with, or deviate from, the $y = x$ line.

We also include a panel of histograms in Figs. (8) and (9) corresponding to Figs. (6) and (7), respectively’ where N denotes the total number of test data points of the model level closest to 850 hPa or the surface (precipitation rates). These are plotted on a log scale in order to better visualize the histograms, since the data is saturated around the central bin, corresponding to the $y = x$ lines in the scatter plots. The histograms were inspired by the findings in Han et al. (2020) and help to illustrate how our scatter plots are dominated by points that fall along the $y = x$ line. Taking into account the difference between shown metrics and model configurations, our results with the one-to-one scatter plots show exceedingly skillful ML emulators, in line with, if not superior to, what we find in the literature for similar work. This is something to keep in mind while we discuss potential limitations arising from increasing complexity.

For both of the large-scale precipitation rate emulators in Figs. (7a, b), the $y = x$ and least-squares fit lines overlap almost completely, and correspondingly, we have minimal spread in the underlying scatter plot. The convective precipitation (7c) plot shows the most visual spread among the precipitation rate scatter plots. Along these same lines, both tendencies in Figs. (6) and (8) show significantly more spread in the convection case over the moist case. However, we note that the all of the histograms in Figs. (8) and (9) show the overwhelming majority of point-wise differences fall within the first few bins and that while outliers occur, they are exceedingly rare. The black dashed lines convey the percentage of instances contained within them. Each case indicates at least 95% of the data within the black dashed lines, and in some cases over 97%, as indicated in the legends. All plots appear to have a slight bias to underestimate the extreme precipitation, as indicated by the direction of the deviation of the regression fit from the one-to-one line along with the consistent under estimated tails in the upper right region of all panels in Fig. 7). This is likely due to the inability for a random forest to predict a value that is not within the range of its training data set, as discussed in Section 2.2.

3.3 R^2 Investigation

Another performance metric is the coefficient of determination, or, R^2 . We calculate R^2 contours over the time and zonal dimensions, given by the formula

$$R^2(:, :) = 1 - \frac{\sum_t \sum_\lambda [\text{CAM}(t, :, :, \lambda) - \text{ML}(t, :, :, \lambda)]^2}{\sum_t \sum_\lambda [\text{CAM}(t, :, :, \lambda) - \overline{\text{CAM}}(:, :)]^2} \quad (9)$$

where λ is the longitudinal dimension, the numerator is referred to as the residual sum of squares and the denominator is the variance of the CAM6 output. The average in the calculation, indicated by $\overline{\text{CAM}}$, is a zonal-mean time-mean over the testing data set. R^2 can simply be understood as a measurement of how well a regression model has learned the functional relationship between the input and the predicted output based on the true output. The closer to one, the better the R^2 . It should be noted here that the R^2 can take negative values whenever the errors in the predictions are larger than the variance in the original data. In general, this may be interpreted as a model that cannot identify, or has not ‘learned’, the functional relationships at play. This approach was inspired by Figs. (1) and (7) in O’Gorman and Dwyer (2018), where in the author shows a panel of R^2 contours for temperature tendencies for various training scenarios also using random forests to emulate the tendencies.

We display a panel of R^2 plots for all of our tendencies in Fig. (10) and precipitation rates in Fig. (11). The plots show areas of significant skill, ranging from 0.8 to 0.99, for all cases. All fields also show regions of low R^2 values, around 0.4 and under, however, the overwhelming majority of the plots show at least $R^2 > 0.7$. The work in this paper is not meant to be a direct comparison to the work in O’Gorman and Dwyer (2018), as we use different model configurations and an entirely different GCM. It is still worth noting how this works as an effective reference due to the many similarities in both ML approach and aspects being emulated. We note again that all of our trained emulators show skill in line with various other examples of similar work within the literature, such as O’Gorman and Dwyer (2018), as well as Yuval et al. (2020).

The R^2 panels in Figs. (10) and (11) reveal a wide variety of aspects. For example, as we increase the complexity of our system, the random forest’s global effectiveness decreases with regards to the R^2 skill. Excluding (10a), from left-to-right we increase in complexity from the moist case to the convection case, and in doing so we notice the impact on the R^2 skill globally. In (10c) there are more regions of $R^2 \leq 0$ in the upper atmosphere, which was not seen in (10b). Along the same line, two pockets of $R^2 \approx 0.3$ form around the tropics in (10e), while (10d) shows $R^2 > 0.7$. These all support the expected conclusion that the effectiveness of random forests for emulating complex functions in climate models are impacted by the increasing complexity of those schemes.

We also note that the R^2 calculation can be an unreliable metric in regimes where there is minimal activity, which does occur in the white regime of Figs. (10a), (10d), and (10e). This is because when the values in the variance and the sum of squares are both functionally zero, they are still seen as floating point numbers of extremely small order and the equation above can lead to the result

$$R^2(:, :) \approx 1 - \frac{10^{-6}}{10^{-13}} \approx 1 - 10^7 \ll 0 \quad (10)$$

For the dry case in plot (10a), this occurs around the equator in the mid-atmosphere. Similarly, this occurs in the upper atmosphere for the moisture tendencies in (10d) and (10e). For the dry case, there is very little heating or cooling occurring on average in this area of the atmosphere and for the moist case there is in fact no forcing occurring at upper levels. However, due to the nature of floating point numbers the R^2 calculation identifies these regimes as areas of poor skill. This is an example of a weakness in R^2 as a metric of regression skill, rather than a reflection of a weakness in the ML model for these particular cases.

439 There is virtually no noticeable drop-off in skill between the moist and the convec-
 440 tion cases for the large-scale precipitations rates (Fig. 11a,b). This is again because their
 441 functional relationships are identical in each case. However, we see that the convective
 442 precipitation scheme suffers a minor loss in skill when compared to the large-scale pre-
 443 cipitation emulators. The emulator is still among the highest skill model, with $R^2 > 0.78$
 444 globally. However, this is in-line with what we observe with the tendencies in Fig. (10)
 445 and the expected result: that as we increase complexity, the effectiveness of the random
 446 forest decreases, even if only marginally for some cases inspected in this work.

447 3.4 Skill Variation

448 Various aspects of the ML training process impact the skill of our emulators. A com-
 449 mon example of this is the idea of feature importance. Feature importance is the inves-
 450 tigation into the relative importance of various input parameters for the skillfulness of
 451 an ML model. It is important that we do not use every possible quantity as inputs as
 452 this increases the computational demand of training these emulators. We know what in-
 453 put fields are used to calculate the functions that we emulate, as discussed in section 2.1.
 454 These tend to include the temperature, pressure, specific heat, and heat fluxes for ex-
 455 ample. One input field that we investigated more closely was relative humidity (RH).
 456 Since RH is not an explicit variable used in calculating the physics tendencies and pre-
 457 cipitation rates, would including it improve performance? Figure (12) shows the R^2 com-
 458 parison of explicitly including the RH (left) and not including it (right), using identical
 459 random forest setups, trained independently, for the moist specific humidity tendency.
 460 The random forest shows skill without the inclusion of the RH field, however it is sig-
 461 nificantly improved upon with the inclusion of the RH.

462 From a pure data science perspective, it may not be apparent that the RH field will
 463 improve the performance since it is not an explicit variable used in the functional form
 464 of the parameterization. From the atmospheric science perspective, this is to be expected
 465 since relative humidity is an important indicator of changing moisture levels in the at-
 466 mosphere. It is also an indicator of supersaturation ($RH > 100\%$) in the large-scale pre-
 467 cipitation algorithm. In turn, this impacts the make up of moisture in the atmosphere
 468 which indirectly influences all predicted fields. This is an example of how important it
 469 is to approach ML problems in the physical sciences with our physical intuition in mind.

470 We also assessed how the model is dependent on the amount of training data points
 471 provided. Figure (13) shows the impact that training data can have on the moisture ten-
 472 dency, as we decrease the amount of training data used for otherwise identical random
 473 forest emulators. Our current models all used around 12 to 20 million samples in which
 474 to train on, as outlined in Supporting Information tables S1 to S8. In this specific test,
 475 20 million samples led to the most skillful R^2 field as well as the highest overall R^2 at
 476 0.876. When decreasing the number of samples we see a decrease in skill in Fig. (13),
 477 as expected. It does appear with a minimum of 500,000 samples, we are able to effec-
 478 tively emulate the moisture tendency with $R^2 > 0.8$ globally. However, once we decrease
 479 that number by another order of magnitude we lose much of our skill, including a genu-
 480 inely poor emulator for only 5,000 training samples. We also find that as we increase
 481 above five million samples, the R^2 begins to level out close to 0.9, indicating that roughly
 482 five or 10 million samples is a sufficient minimum for our work. Investigations like these
 483 can help to guide us as we attempt to find the balance between skillful ML models and
 484 computationally feasible approaches.

485 4 Concluding Thoughts & Applications to Future Work

486 Individual random forests are developed and trained for emulating temperature ten-
 487 dencies, specific humidity tendencies, large-scale precipitation rates, and convective pre-
 488 cipitation rates from physical parameterization schemes within three simple physics model

489 configurations within the CAM6 model framework. The configurations built upon one
 490 another in a hierarchy of complexity, beginning with the dry case, before moving to the
 491 moist and culminating with the moist case coupled with a convection scheme. We uti-
 492 lized a broad scope of training data over a 60-year model run for each configuration and
 493 allowed the SHERPA hyperparameter optimization tool to freely optimize each individ-
 494 ual random forest. This allowed us to create robust emulators in order to probe the lim-
 495 its of skill achievable via random forest emulators of sub-grid physics schemes in an of-
 496 fline mode.

497 All of our emulators showed significant skill when tested on the test data over the
 498 final six years of model output. With the given metrics, our emulators showed results
 499 at least as skillful as other similar examples within the literature, while in many cases
 500 outperforming similar work. This statement bares reminding that in a majority of cases
 501 our model configurations were less-complex than we find in the current literature, so di-
 502 rect comparisons are not possible. We also note the benefits of using random forests rather
 503 than other methods, such as deep learning, since random forests inherently preserve many
 504 physical laws. An example being the non-negativity of precipitation due to the inabil-
 505 ity for random forests to make a prediction outside of the range of the data it was trained
 506 on.

507 Further, we identify indications of limitations in using random forests for emulat-
 508 ing physical parameterization schemes, even within our highly simplified hierarchy of con-
 509 figurations. In many cases, there were noticeable decreases in skill as the complexity of
 510 the physics scheme was increased. This raises interesting questions about where random
 511 forests fit within the overall approach to the community’s interest in using ML to em-
 512 ulate these sub-grid processes and otherwise augment aspects of climate and weather mod-
 513 els. Balancing the trade-offs between physical realism, computational efficiency, and model
 514 complexity should inform the choice of ML technique, especially when looking forward
 515 towards state-of-the-art model complexity. Random forests are unlikely to remain as skill-
 516 ful, particularly in full-tendency emulators like ours, in cases with such a substantial in-
 517 crease in overall complexity. In regards to this work, our next intriguing step will be to
 518 couple the random forests to the CAM6 implementation and analyze how they perform
 519 in an online mode. Investigating whether the rare, yet present, outliers impact the sta-
 520 bility of the coupled model will provide further insight into where random forests may
 521 fit into the future of data science-augmented climate and weather models. Based solely
 522 on this work, we expect it to likely fall within the emulation of individual components
 523 that can take advantage of the inherent conservation of physical laws, such as convec-
 524 tion schemes or the radiation parameterization.

525 Acknowledgments

526 This work was made possible by the National Science Foundation’s Graduate Research
 527 Fellowship Program and the NOAA grant NA17OAR4320152(127). We would like to ac-
 528 knowledge high-performance computing support from Cheyenne (doi:10.5065/D6RX99HX)
 529 and the Casper data analysis server. These resources were provided by the Computa-
 530 tional and Information Systems Laboratory of the National Center for Atmospheric Re-
 531 search (NCAR), sponsored by the National Science Foundation. The CAM6 output data
 532 used for all three cases of machine learning in the study are available at Limon (2022)
 533 and all machine learning related scripts at https://github.com/gclimon/ml_climate.

534 References

- 535 Betts, A. K. (1986). A new convective adjustment scheme. Part I: Observational and
 536 theoretical basis. *Quart. J. Roy. Meteor. Soc.*, *112*, 677–692.
- 537 Betts, A. K., & Miller, M. J. (1986). A new convective adjustment scheme. Part II:
 538 Single column tests using GATE wave, BOMEX, and arctic air-mass data sets.

- 539 *Quart. J. Roy. Meteor. Soc.*, *112*, 693–709.
- 540 Beucler, T., Rasp, S., Pritchard, M., & Gentine, P. (2019). Achieving Conserva-
541 tion of Energy in Neural Network Emulators for Climate Modeling. *arXiv*. Re-
542 trieved from <http://arxiv.org/abs/1906.06622> (arXiv:1906.06622v1)
- 543 Boukabara, S. A., Krasnopolsky, V., Penny, S. G., Stewart, J. Q., McGovern, A.,
544 Hall, D., . . . Hoffman, R. N. (2021). Outlook for exploiting artificial intelli-
545 gence in the Earth and environmental sciences. *Bulletin of the American Mete-
546 orological Society*, *102*(5), E1016–E1023. doi: 10.1175/BAMS-D-20-0031.1
- 547 Breiman, L. (1996). Bagging Predictors. *Machine Learning*, *24*(2), 123–
548 140. Retrieved from <https://doi.org/10.1007/BF00058655> doi:
549 10.1007/BF00058655
- 550 Brenowitz, N. D., & Bretherton, C. S. (2018). Prognostic Validation of a Neural
551 Network Unified Physics Parameterization. *Geophys. Res. Lett.*, *45*(12), 6289–
552 6298. doi: 10.1029/2018GL078510
- 553 Bretherton, C. S., Henn, B., Kwa, A., Brenowitz, N. D., Watt-Meyer, O., McGib-
554 bon, J., . . . Harris, L. (2022). Correcting Coarse-Grid Weather and Climate
555 Models by Machine Learning From Global Storm-Resolving Simulations. *J.
556 Adv. Model. Earth Syst.*, *14*(2). doi: 10.1029/2021MS002794
- 557 Chantry, M., Christensen, H., Düben, P., & Palmer, T. (2021). Opportunities
558 and challenges for machine learning in weather and climate modelling: hard,
559 medium and soft AI. *Phil. Trans. R. Soc. A*, *379*(2194), 20200083. doi:
560 10.1098/rsta.2020.0083
- 561 Chapman, W. E., Subramanian, A. C., Delle Monache, L., Xie, S. P., & Ralph,
562 F. M. (2019). Improving atmospheric river forecasts with machine learning.
563 *Geophys. Res. Lett.*, *46*(17–18), 10627–10635. doi: 10.1029/2019GL083662
- 564 Danabasoglu, G., Lamarque, J.-F., Bacmeister, J., Bailey, D. A., DuVivier, A. K.,
565 Edwards, J., . . . Strand, W. G. (2020). The Community Earth System Model
566 Version 2 (CESM2). *J. Adv. Model. Earth Syst.*, *12*(2), e2019MS001916. doi:
567 10.1029/2019MS001916
- 568 Foster, D., Gagne, D. J., & Whitt, D. B. (2021). Probabilistic Machine Learn-
569 ing Estimation of Ocean Mixed Layer Depth From Dense Satellite and
570 Sparse In Situ Observations. *J. Adv. Model. Earth Syst.*, *13*(12), 1–33. doi:
571 10.1029/2021MS002474
- 572 Frierson, D. M. W. (2007). The Dynamics of Idealized Convection Schemes and
573 Their Effect on the Zonally Averaged Tropical Circulation. *J. Atmos. Sci.*, *64*,
574 1959–1976.
- 575 Gagne, D. J., McGovern, A., Haupt, S. E., Sobash, R. A., Williams, J. K., & Xue,
576 M. (2017). Storm-Based Probabilistic Hail Forecasting with Machine Learning
577 Applied to Convection-Allowing Ensembles. *Weather and Forecasting*, *32*(5),
578 1819–1840. doi: 10.1175/WAF-D-17-0010.1
- 579 Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G., & Yacalis, G. (2018). Could ma-
580 chine learning break the convection parameterization deadlock? *Geophys. Res.
581 Lett.*, *45*(11), 5742–5751. doi: 10.1029/2018GL078202
- 582 Gettelman, A., Gagne, D. J., Chen, C.-C., Christensen, M. W., Lebo, Z. J., Morri-
583 son, H., & Gantos, G. (2021). Machine Learning the Warm Rain Process. *J.
584 Adv. Model. Earth Syst.*, *13*(2), e2020MS002268. doi: [https://doi.org/10.1029/
585 2020MS002268](https://doi.org/10.1029/2020MS002268)
- 586 Han, Y., Zhang, G. J., Huang, X., & Wang, Y. (2020). A moist physics pa-
587 rameterization based on deep learning. *J. Adv. Model. Earth Syst.*, *12*(9),
588 e2020MS002076. doi: 10.1029/2020MS002076
- 589 Held, I. M. (2005, November). The gap between simulation and understanding in cli-
590 mate modeling. *Bull. Amer. Meteor. Soc.*, *86*, 1609–1614.
- 591 Held, I. M., & Suarez, M. J. (1994, October). A proposal for the intercomparison
592 of the dynamical cores of atmospheric general circulation models. *Bull. Amer.
593 Meteor. Soc.*, *75*(10), 1825–1830.

- 594 Hertel, L., Collado, J., Sadowski, P., Ott, J., & Baldi, P. (2020). Sherpa: Robust
595 Hyperparameter Optimization for Machine Learning. *SoftwareX*, 24. Retrieved
596 from [10.1016/j.softx.2020.100591](https://doi.org/10.1016/j.softx.2020.100591)
- 597 Hourdin, F., Mauritsen, T., Gettelman, A., Golaz, J.-C., Balaji, V., Duan, Q., ...
598 Williamson, D. (2017). The art and science of climate model tuning. *Bull.*
599 *Ameri. Meteor. Soc.*, 98(3), 589–602. doi: 10.1175/BAMS-D-15-00135.1
- 600 Hoyer, S., & Hamman, J. (2017). xarray: N-D labeled Arrays and Datasets in
601 Python. *Journal of Open Research Software*, 5(1), 10. doi: 10.5334/jors.148
- 602 Krasnopolsky, V. M., & Fox-Rabinovitz, M. S. (2006). Complex hybrid models
603 combining deterministic and machine learning components for numerical cli-
604 mate modeling and weather prediction. *Neural Networks*, 19(2), 122–134. doi:
605 10.1016/j.neunet.2006.01.002
- 606 Limon, G. C. (2022). *Simple Physics CAM6 Dataset for Training Machine Learning*
607 *Algorithms* (Tech. Rep.). University of Michigan - Deep Blue Data. Retrieved
608 from <https://doi.org/10.7302/r6v3-s064> doi: 10.7302/r6v3-s064
- 609 Medeiros, B., Williamson, D. L., & Olson, J. G. (2016). Reference aquaplanet cli-
610 mate in the Community Atmosphere Model, version 5. *J. Adv. Model. Earth*
611 *Syst.*, 8(1), 406–424.
- 612 O’Gorman, P. A., & Dwyer, J. G. (2018). Using Machine Learning to Parame-
613 terize Moist Convection: Potential for Modeling of Climate, Climate Change,
614 and Extreme Events. *J. Adv. Model. Earth Syst.*, 10(10), 2548–2563. doi:
615 10.1029/2018MS001351
- 616 Pedregosa, F., Varoquaux, G., Gramfort, A., & Michel, V. (2011). Scikit-learn:
617 Machine Learning in Python. *Journal of Machine Learning Research*, 12(Oct),
618 2825–2830.
- 619 Rasp, S., Pritchard, M. S., & Gentine, P. (2018). Deep learning to represent sub-grid
620 processes in climate models. *Proceedings of the National Academy of Sciences*,
621 115(39), 9684–9689. doi: 10.1073/pnas.1810286115
- 622 Reed, K. A., & Jablonowski, C. (2012). Idealized tropical cyclone simulations of in-
623 termediate complexity: A test case for AGCMs. *J. Adv. Model. Earth Syst.*, 4,
624 M04001.
- 625 Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N.,
626 & Prabhat. (2019). Deep learning and process understanding for data-driven
627 earth system science. *Nature*, 566, 196 – 204. doi: 10.1038/s41586-019-0912-1
- 628 Stevens, B., & Bony, S. (2013). What Are Climate Models Missing? *Science*,
629 340(6136), 1053–1054. doi: 10.1126/science.1237554
- 630 Thatcher, D. R., & Jablonowski, C. (2016). A moist aquaplanet variant of the Held-
631 Suarez test for atmospheric model dynamical cores. *Geoscientific Model Devel-*
632 *opment*, 9(4), 1263–1292. doi: 10.5194/gmd-9-1263-2016
- 633 Ukkonen, P. (2022). Exploring Pathways to More Accurate Machine Learning Em-
634 ulation of Atmospheric Radiative Transfer. *J. Adv. Model. Earth Syst.*, 14(4),
635 1–19. doi: 10.1029/2021ms002875
- 636 Watt-Meyer, O., Brenowitz, N. D., Clark, S. K., Henn, B., Kwa, A., McGibbon, J.,
637 ... Bretherton, C. S. (2021). Correcting Weather and Climate Models by
638 Machine Learning Nudged Historical Simulations. *Geophys. Res. Lett.*, 48(15),
639 1–10. doi: 10.1029/2021GL092555
- 640 Williamson, D. L., Blackburn, M., Hoskins, B. J., Nakajima, K., Ohfuchi, W., Taka-
641 hashi, Y. O., ... Stratton, R. (2012). *The APE Atlas* (NCAR Technical
642 Note Nos. NCAR/TN-484+ STR). Boulder, Colorado: National Center for
643 Atmospheric Research. doi: 10.5065/D6FF3QBR
- 644 Yorgun, M. S., & Rood, R. B. (2016). A Decision Tree Algorithm for Investigation
645 of Model Biases Related to Dynamical Cores and Physical Parameterizations.
646 *J. Adv. Model. Earth Syst.*, 8, 1769–1785. doi: 10.1002/2016MS000657
- 647 Yuval, J., O’Gorman, P. A., & Hill, C. N. (2020). Stable machine-learning param-
648 eterization of subgrid processes for climate modeling at a range of resolutions.

649
650
651
652
653

Nat. Commun., 11, 3295. doi: 10.1038/s41467-020-17142-3
 Yuval, J., O’Gorman, P. A., & Hill, C. N. (2021). Use of Neural Networks for
 Stable, Accurate and Physically Consistent Parameterization of Subgrid At-
 mospheric Processes With Good Performance at Reduced Precision. *Geophys.*
Res. Lett., 48(6), 1–11. doi: 10.1029/2020gl091363

654 **Figures**

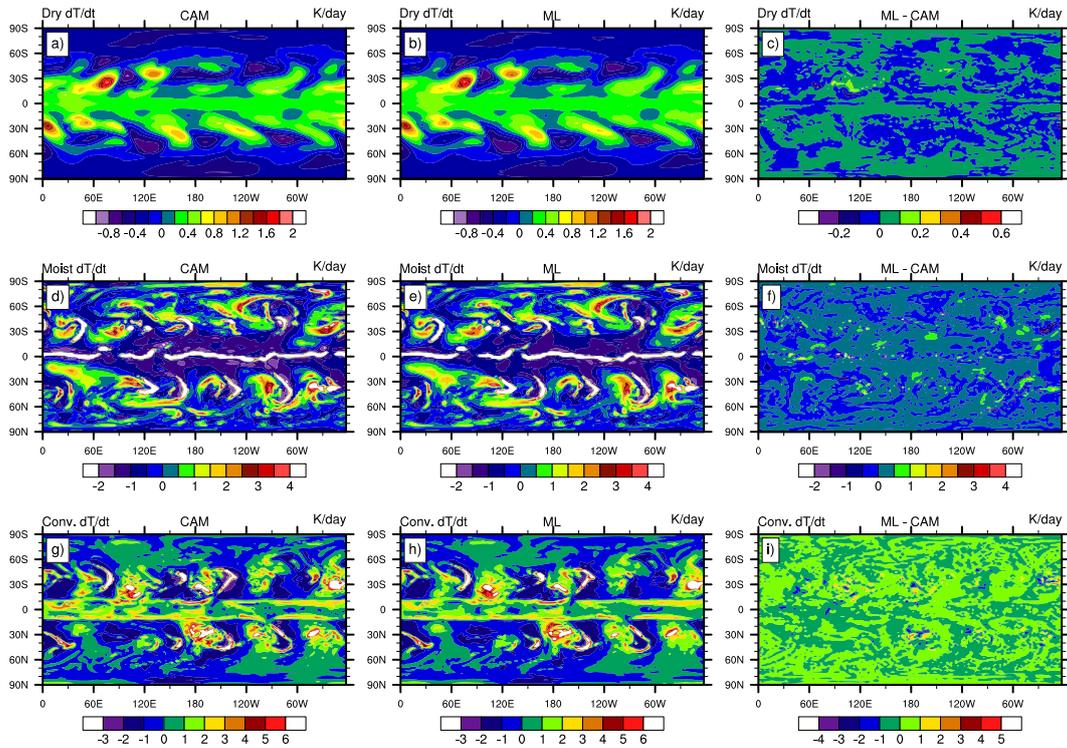


Figure 1: Time step snapshot plots of CAM6 and ML predicted temperature tendencies and their differences near 850 hPa. Top row corresponds to the dry case, middle row to the moist case, and bottom to the convective case. Left column corresponds to CAM6 output, middle the random forest ML predictions, and right depicts their respective differences.

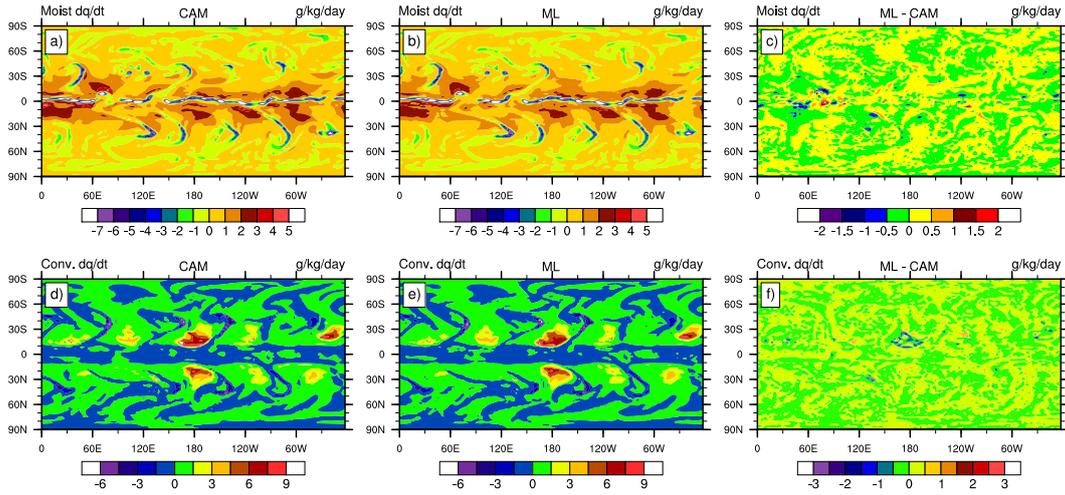


Figure 2: Time step snapshot plots of CAM6 and ML predicted specific humidity tendencies and their differences near 850 hPa. Top row corresponds to the moist case and bottom to the convective case. Left column corresponds to CAM6 output, middle the random forest ML predictions, and right depicts their respective differences.

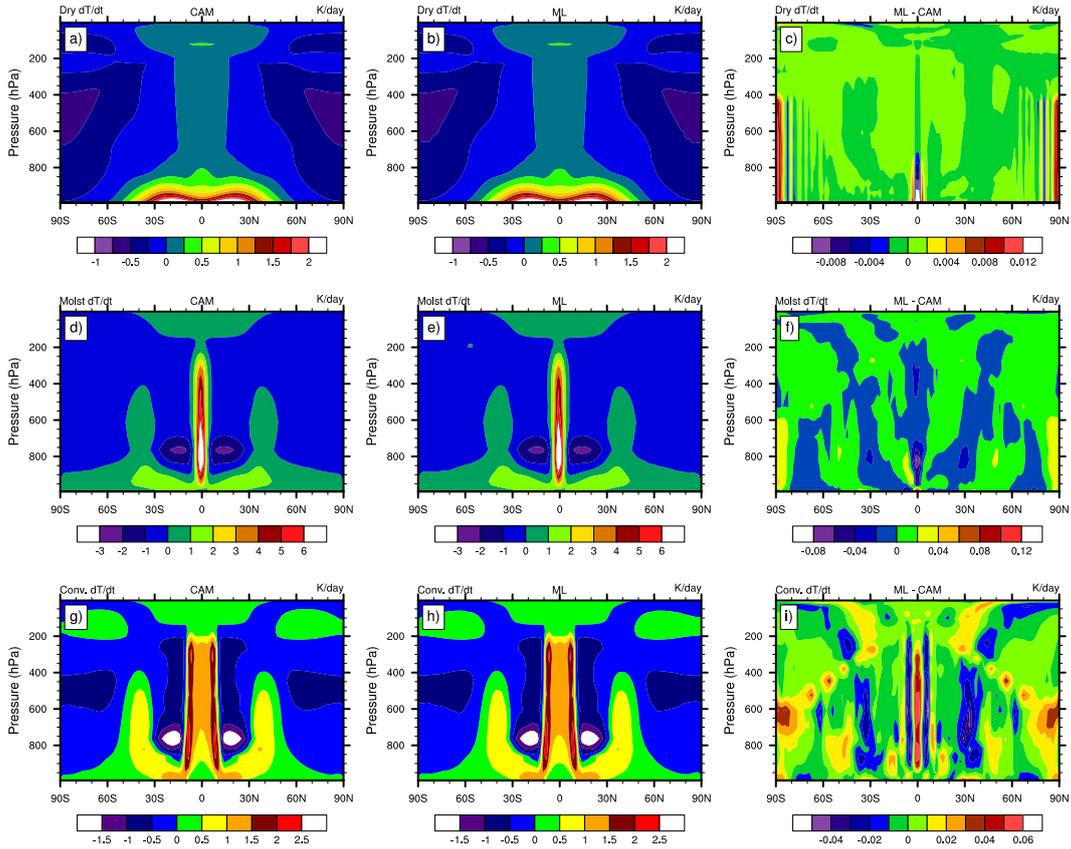


Figure 3: Zonal-mean time-mean plots of original and predicted temperature tendencies and their differences over the full testing data set. Top row corresponds to the dry case, middle row to the moist case, and bottom to the convective case. Left column corresponds to CAM6 output, middle the random forest ML predictions, and right depicts their respective differences.

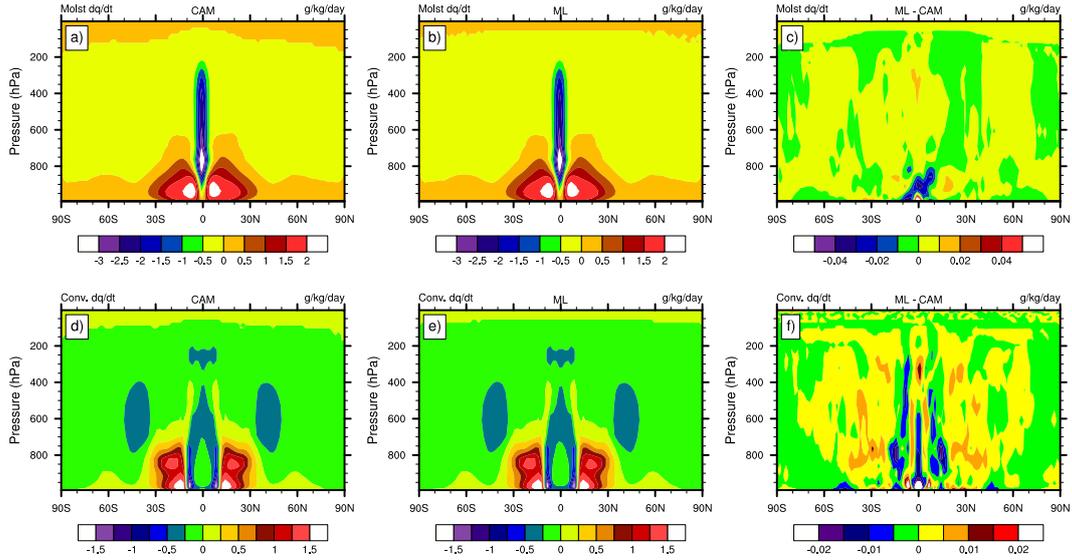


Figure 4: Zonal-mean time-mean plots of original and predicted moisture tendencies and their differences over the full testing data set. Top row corresponds to the moist case, and bottom to the convective case. Left column corresponds to CAM6 output, middle the random forest ML predictions, and right depicts their respective differences.

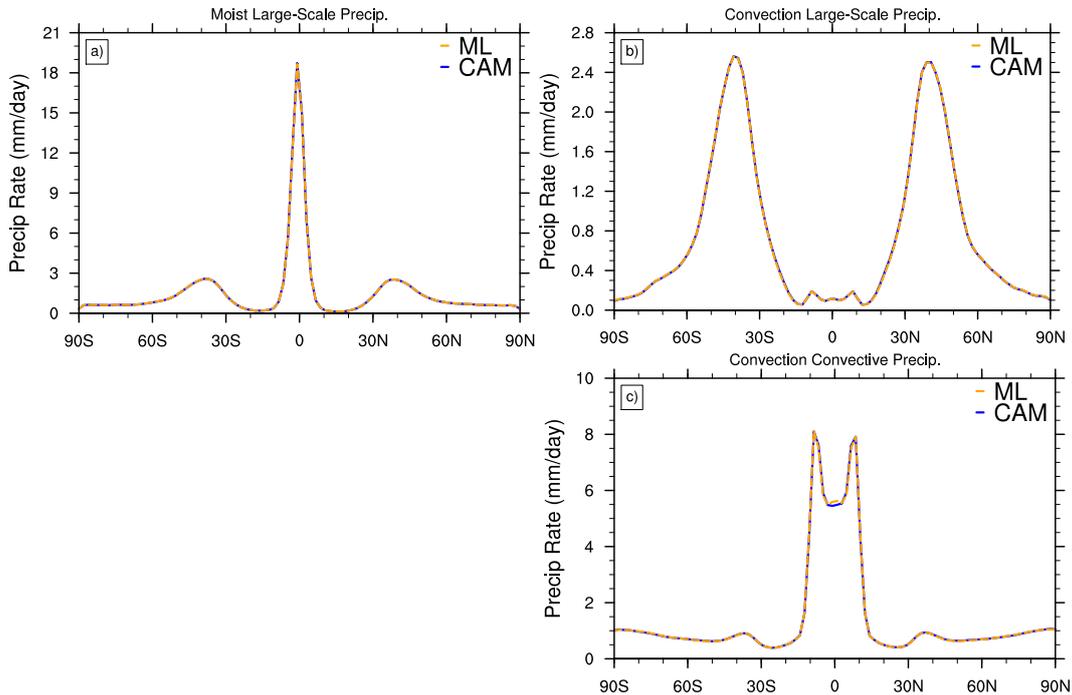


Figure 5: Zonal-Mean Time-Mean plots of CAM6 (blue) and ML predicted (orange) precipitation totals over the full testing data set. Top row corresponds to the large-scale precipitation (eq. 5) and bottom to the convective precipitation. Left column corresponds to moist case, while the right is the convective case.

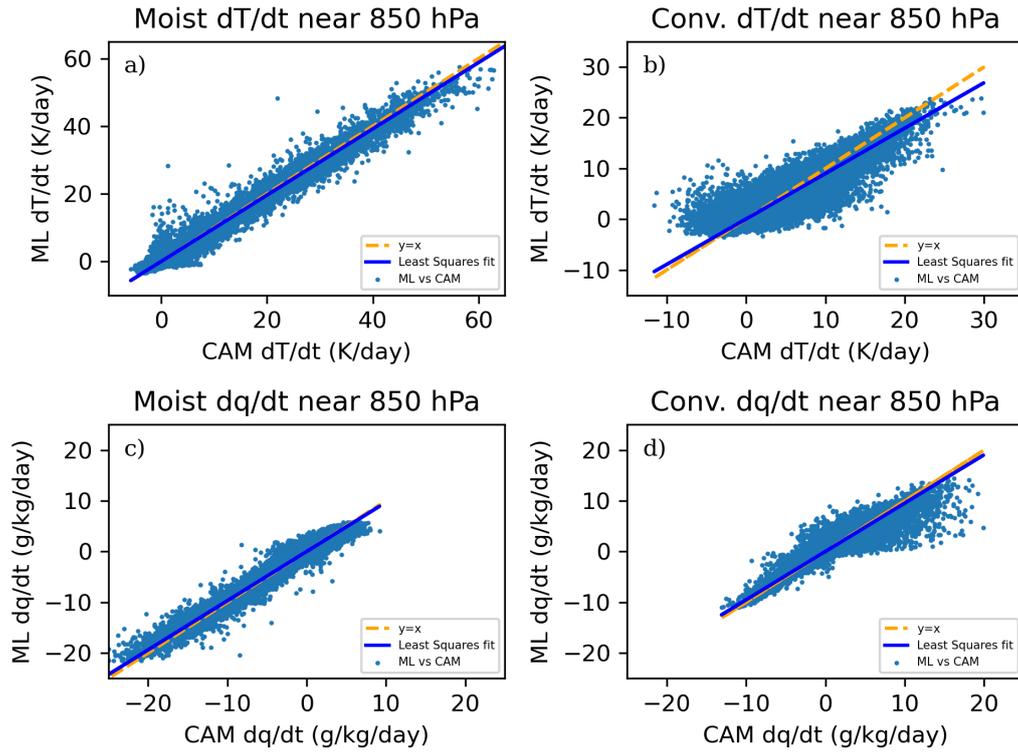


Figure 6: Direct comparison scatter plots for ML predicted values (y-axis) against original CAM6 output (x-axis) for all horizontal grid points near 850 hPa over the testing data for (a) moist-case temperature tendency, (b) convection-case temperature tendency, (c) moist-case moisture tendency, and (d) convection-case moisture tendency.

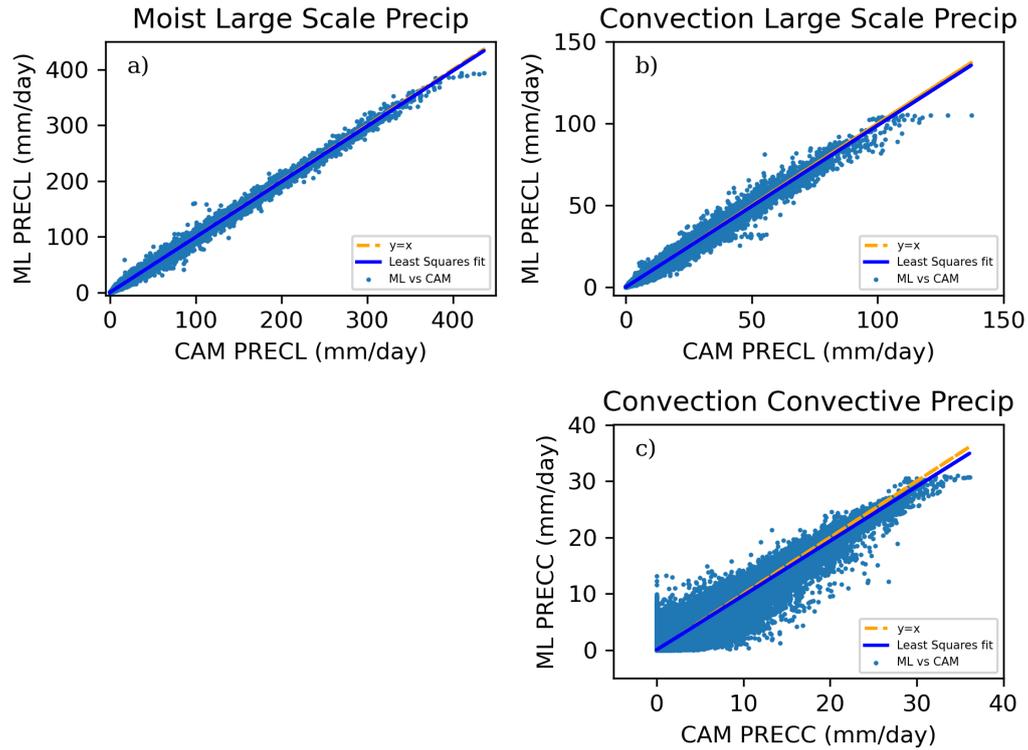


Figure 7: Direct comparison scatter plots for ML predicted values (y-axis) against original CAM6 output (x-axis) for all horizontal grid points near 850 hPa over the testing data for (a) moist-case large-scale precipitation, (b) convection-case large-scale precipitation, and (c) convection-case convective precipitation.

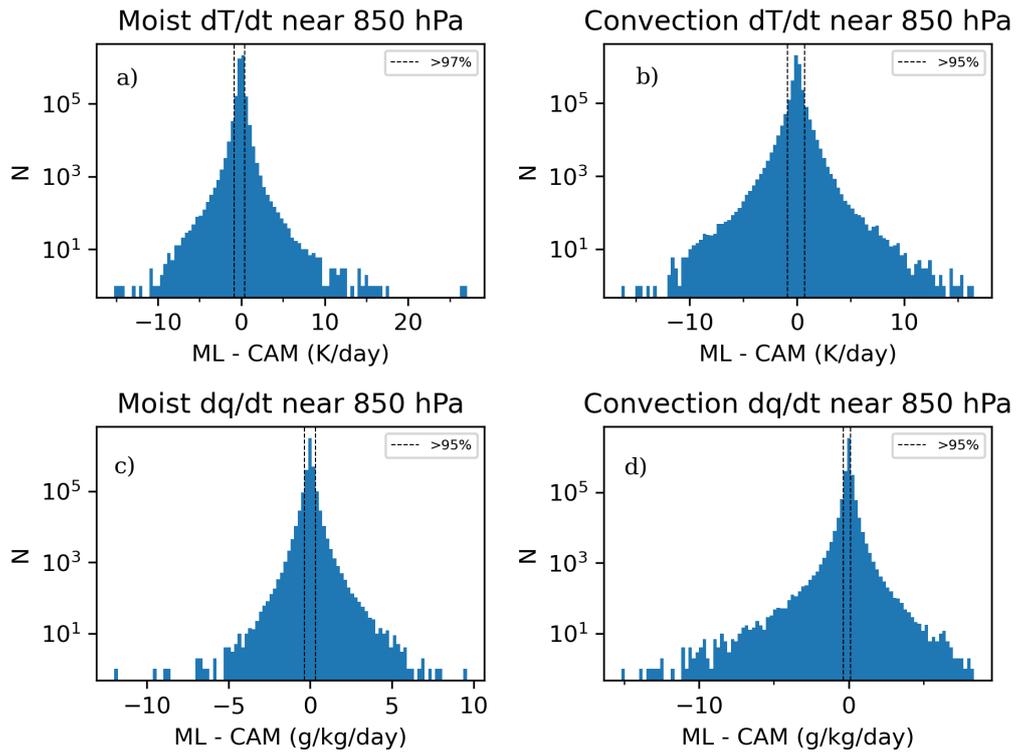


Figure 8: Histograms of the point-wise difference (ML - CAM6) for the temperature (top) and specific humidity (bottom) tendencies, corresponding to the scatter plots in Fig. (6) on a log scale using 100 bins. Percentage of data contained within the black dashed lines are indicated in individual legends and conveys the significance of the log-scale and the overwhelming majority of instances occurring near the peak of the distribution.

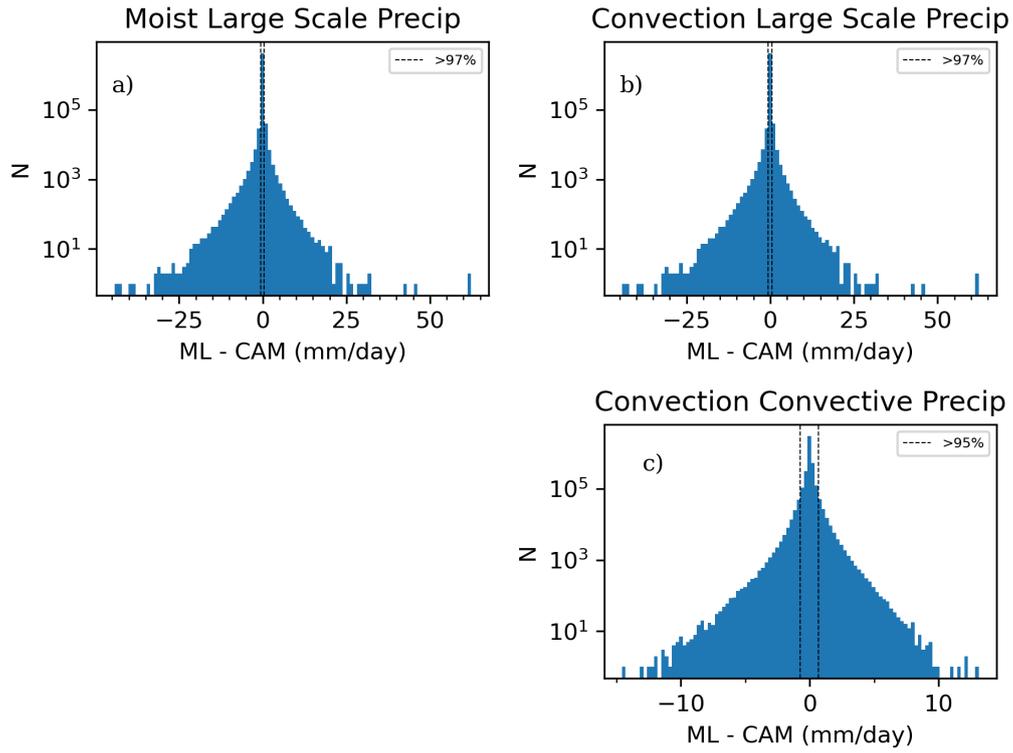


Figure 9: Histograms of the point-wise difference (ML - CAM6) for the precipitation rates corresponding to the scatter plots in Fig. (7) on a log scale using 100 bins. Percentage of data contained within the black dashed lines are indicated in individual legends and conveys the significance of the log-scale and the overwhelming majority of instances occurring near the peak of the distribution.

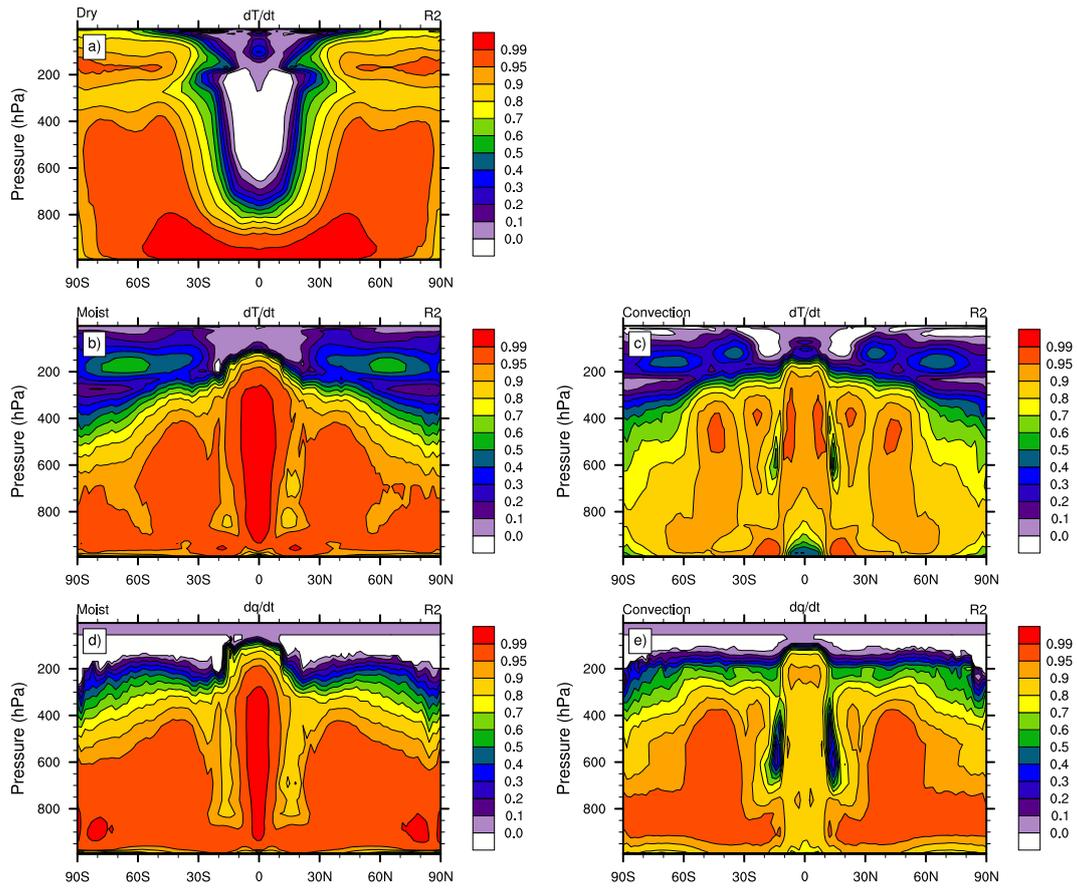


Figure 10: R^2 calculations over the zonal & temporal dimensions via eq. 10. (a) Dry temperature tendency, (b) moist temperature tendency, (c) convection temperature tendency, (d) moist moisture tendency, and (e) convection moisture tendency.

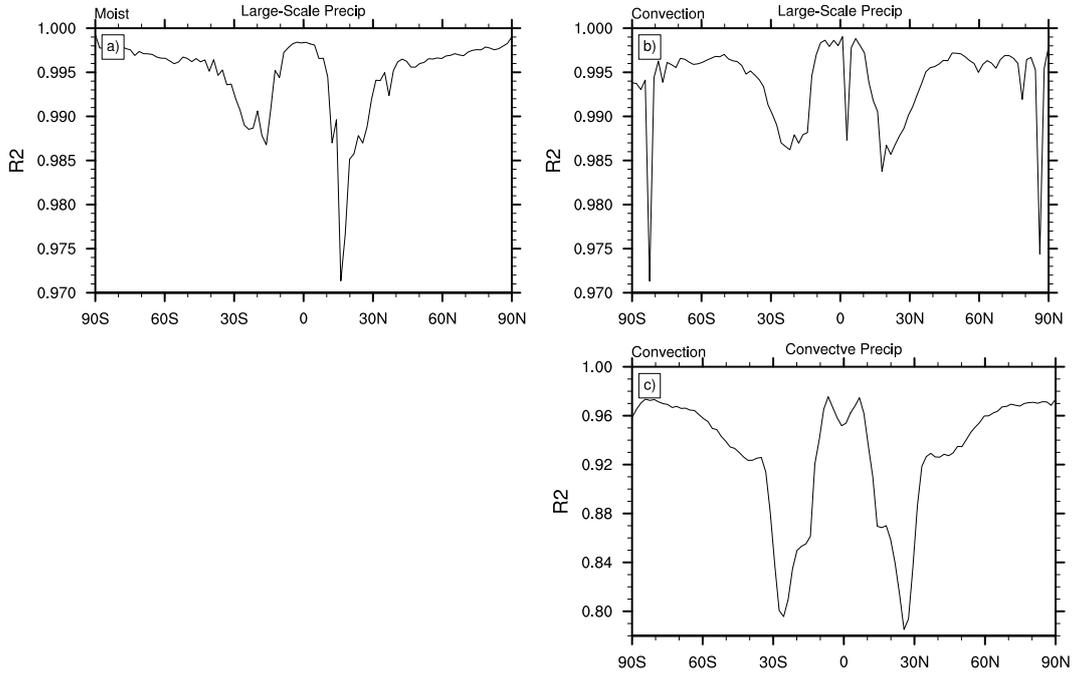


Figure 11: R^2 calculations over the zonal & temporal dimensions via eq. 10. (a) moist large-scale precipitation, (b) convection large-scale precipitation, and (c) convection convective precipitation.

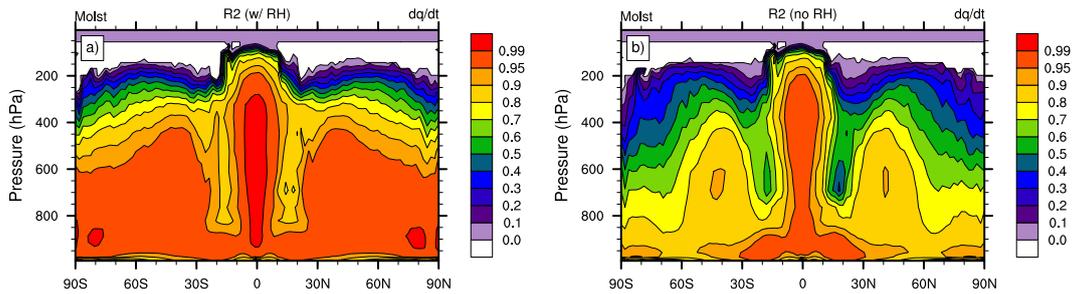


Figure 12: Comparison of R^2 plot - as defined in Fig. (10) - with (a) and without (b) relative humidity as a feature for ML prediction of the moisture tendency for the moist case. Figure (12a) reproduces Fig. (10d).

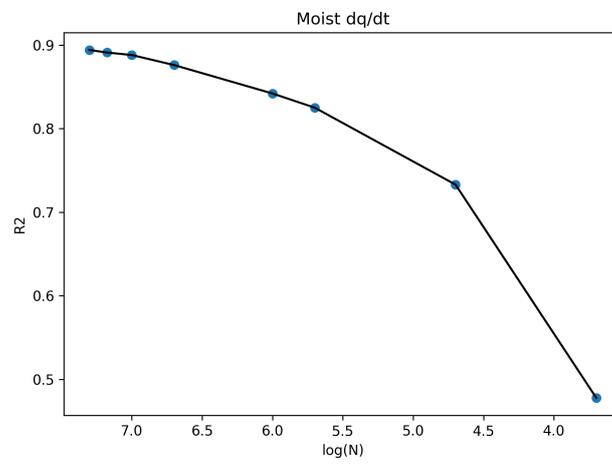


Figure 13: Globally-averaged R^2 value (y-axis) for ML prediction of the moisture tendency for the moist case using varying number (N) of training samples (x-axis) on a log scale.