An aerosol activation metamodel of v1.2.0 of the pyrcel cloud parcel model: Development and offline assessment for use in an aerosol-climate model

Daniel Rothenberg[1] and Chien Wang[1]

[1]Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA
Correspondence to: Daniel Rothenberg (darothen@mit.edu)

Abstract. We describe an emulator of a detailed cloud parcel model which has been trained to assess droplet nucleation from a complex, multi-modal aerosol size distribution simulated by a global aerosol-climate model. The emulator is constructed using a sensitivity analysis approach (polynomial chaos expansion) which reproduces the behavior of the targeted parcel model across the full range of aerosol properties and meteorology simulated by the parent climate model. An iterative technique using aerosol fields sampled from a global model is used to identify the critical aerosol size distribution parameters necessary for accurately predicting activation. Across the large parameter space used to train them, the emulators estimate cloud droplet number concentration (CDNC) with a mean relative error of 9.2% for aerosol populations without giant cloud condensation nuclei (CCN), and 6.9% when including them. Versus a parcel model driven by those same aerosol fields, the best-performing emulator has a mean relative error of 4.6%, which is comparable with two commonly-used activation schemes also evaluated here (which have mean relative errors of 2.9% and 6.7%, respectively). We identify the potential for regional biases in modeled CDNC, particularly in oceanic regimes, where our best-performing emulator tends to over-predict by 7%, whereas the reference activation schemes range in mean relative error from -3% to 7%. The emulators which include the effects of giant CCN are more accurate in continental regimes (mean relative error of 0.3%), but strongly over-estimate CDNC in oceanic regimes by up to 22%, particularly in the Southern Ocean. The biases in CDNC resulting from the subjective choice of activation scheme could potentially influence the magnitude of the indirect effect diagnosed from the model incorporating it.

1 Introduction

Aerosols play a critical role in the climate system by interacting with radiation through several different mechanisms. Depending on their composition, aerosol particles can directly scatter or absorb incoming solar radiation, leading to a direct radiative effect and rapid changes in the energy budgets of the surface and atmosphere. Additionally, aerosol particles mediate the production of clouds by providing surface area on which water vapor may condense to form droplets. Through this second pathway, changes in the aerosol population perturb the radiative properties of clouds by altering their microstructure and lifecycle, thereby impacting the planetary radiative budget. Despite decades of focused research by the scientific community, the radiative forcing produced through this second pathway, known as aerosol-cloud interactions, remains one of the largest

uncertainties in understanding contemporary and future climate change on both regional and global scales (Boucher et al., 2013).

To include this second pathway, contemporary Earth System Models predict cloud droplet number concentration (CDNC) by evaluating the nucleation of droplets (aerosol activation) from their simulated aerosol fields. As a result, these models can resolve aerosol-climate indirect effects which arise when anthropogenic aerosol emissions influence cloud microphysical and optical properties through impacting the CDNC burden. The interactions between aerosol particles, water vapor, and cloud droplets are often described using the conceptual model of a possibly-entraining, adiabatic cloud parcel (e.g. Feingold and Heymsfield, 1992; Nenes et al., 2001; Ervens et al., 2005; Topping et al., 2013). However, it is not practical to directly include these calculations in global models, because of their computational complexity and because parcel theory describes a process occurring on spatial and temporal scales much finer than those resolved by climate model grids. Some efforts have sought to incorporate fine-scale information about aerosol-cloud interactions by embedding higher-resolution models within each grid cell of a global model (for example, the "multi-scale modeling framework" of Wang et al., 2011), but these are also too coarse to resolve the fine scale motions associated with parcel theory. To resolve these issues, global models employ aerosol activation parameterizations.

A comprehensive review of the development of these parameterizations is provided by Ghan et al. (2011). In broad terms, there are two families of such schemes: look-up tables based on detailed parcel model calculations (e.g. Saleeby and Cotton, 2004; Ward et al., 2010), and physically-based parameterizations (e.g. Twomey, 1959; Ghan et al., 1993; Abdul-Razzak and Ghan, 2000; Cohard and Pinty, 2000; Nenes and Seinfeld, 2003; Fountoukis and Nenes, 2005; Ming et al., 2006; Shipway and Abel, 2010). Physically-based approaches often rely on empirical fits or tuning derived from parcel model calculations, but are preferable for inclusion in regional or global climate models, because they are usually applicable to an arbitrary description (modal or sectional) of the aerosol size distribution. New look-up tables must be computed whenever additional aerosol parameters are introduced; this is a problem because a look-up table's size grows exponentially as the number of parameters included increases. But look-up tables offer some advantages, including computational efficiency and the ability to account for biases in estimating CDNC, such as kinetic limitations on droplet growth (Nenes et al., 2001), the co-condensation of organic vapors (Topping et al., 2013), and competition for water vapor uptake in the presence of multiple aerosol modes (Simpson et al., 2014).

Rothenberg and Wang (2016) previously used statistical emulation to conduct a sensitivity analysis of a detailed parcel model, focusing on a single, lognormal aerosol mode. Other studies have also used emulation approaches to study response surfaces of both cloud parcel models (Partridge et al., 2011) and global models (Carslaw et al., 2013), but Rothenberg and Wang further developed a framework for using their emulators as an activation parameterization, which they compared to a high-dimensional extension of traditional look-up tables. Compared to two other physically-based schemes, the emulator-activation parameterizations tended to predict CDNC more accurately versus a reference parcel model, particularly in regimes with weak updraft speeds and high aerosol number concentration, where the traditional parameterizations performed the most poorly.

In this present work, we extend this approach to develop a set of metamodels trained for the aerosol and meteorology parameter space simulated by a global aerosol-climate model. The resulting metamodels can thus be directly used as activation parameterizations inside a global climate to predict online CDNC, given information about the aerosol size distributions and sub-grid scale meteorology in each model grid-box. To refine the parameter space used for emulation, we present an analysis

5 of how the size distribution parameters of each aerosol mode in our model contribute to activation dynamics and droplet nucleation. We then evaluate the performance of our emulator-parameterizations versus two physically-based schemes which are used in the vast majority of contemporary global models (see Table 3 of Ghan et al., 2011), and assess the impact of including each of these schemes on the computational expense of our global model.

2 Emulation of aerosol activation

10 ### 2.1 Parent aerosol-climate model

We seek to derive aerosol activation emulators for the MultiMode, two-Moment, Mixing-state resolving Model of Aerosols for Research of Climate (MARC; version 1.0.1 here) (Kim et al., 2008, 2014). MARC builds on the NCAR Community Earth System Model (CESM; version 1.2.2 here), which is a fully-coupled global climate model with sub-components for simulating climate processes in the land, ocean, atmosphere, and sea ice domains. In particular, MARC replaces the default modal aerosol

15 treatment (Liu et al., 2012) in the atmosphere component of the CESM, the Community Atmosphere Model (CAM5; version 5.3 here Neale et al., 2012), with a scheme which simultaneously resolves both an external mixture of different aerosol species and internal mixtures between others (Wilson et al., 2001).

Table 1 summarizes the aerosol modes predicted by MARC, which includes tri-modal sulfate (nucleation ["NUC"], Aitken ["AIT"], and accumulation ["ACC"]), discrete modes for pure black carbon ("BC") and a generic organic carbon species

20 ("OC"), and two "mixed" modes, one of each for mixed sulfate-black-carbon ("MBS") and sulfate-organic-carbon particles ("MOS"). For each of these seven modes, MARC predicts total particle mass (M) and number concentrations (N) for a corresponding lognormal size distribution with a prescribed width (geometric standard deviation; σ_g), as well as the total mass of carbon in each of the MOS and MBS modes. Additionally, sea salt ("SSLTn") and dust ("DSTn") particle size distributions are each described within MARC using a 4-bin, single-moment scheme with fixed particle sizes and an assumed σ_g. Each mode

25 has a prescribed particle hygroscopicity follow κ-Köhler theory (Petters and Kreidenweis, 2007) except for the MOS mode, which has a composition-dependent κ computed assuming that the carbon and sulfate in the particle form a simple internal mixture. Unlike the MOS mode, the MBS mode assumes a core-shell structure with sulfate coating a black carbon nucleus, and has a fixed hygroscopicity corresponding to that of sulfate, $\kappa = 0.507$. The sea salt modes are assumed to be comprised of NaCl with $\kappa = 1.16$, and the dust modes are assumed to be a mixture of minerals with a hygroscopicity of $\kappa = 0.14$ (Scanza

30 et al., 2014). The OC and BC modes are assumed to be non-hygroscopic and not significant players in aerosol activation.

The aerosol size distributions predicted by MARC interact with both radiation and cloud microphysics. MARC uses a two-moment, stratiform cloud microphysics scheme (Morrison et al., 2008) which includes an explicit source of cloud droplet formation via aerosol activation. This is facilitated by means of a physical parameterization which takes as input both the phys-

ical and chemical properties of the ambient aerosol as well as limited information about meteorology such as, the distribution of subgrid-scale vertical velocities. In CAM5 (as used by MARC), a single characteristic updraft velocity (V) diagnosed from the grid cell turbulent kinetic energy (TKE) provided by the moist turbulence scheme (Park and Bretherton, 2009) and assumed to be isotropic is used to estimate droplet nucleation following Ghan et al. (1997) and Lohmann et al. (1999), such that

$$5 \quad V = \overline{V} + \sqrt{\frac{2}{3}\text{TKE}}$$

where $\overline{V}$ is the large-scale resolved updraft velocity. Furthermore, we limit $0.2\,\text{m}\,\text{s}^{-1} < V < 10\,\text{m}\,\text{s}^{-1}$ because the processes driving turbulence are not represented well in MARC, particularly those driven by cloud-top radiative cooling above the planetary boundary layer (Ghan et al., 1997). Morales and Nenes (2010) and West et al. (2014) have explored how using different characteristic updraft velocities to represent subgrid-scale variability can influence simulated aerosol-cloud interactions; in

10   particular, West et al. (2014) showed that using a similar TKE-based parameterization produced more realistic spatial and temporal variability in V, but tends to produce an unrealistically high frequency of its minimum-permissible value. We further assume that activation occurs in non-entraining, adiabatic updrafts which carry air up and into the base of stratiform clouds.

2.2   Aerosol activation parameter space

MARC requires 24 parameters to fully describe its 15-mode aerosol size distribution, although not all of these modes are

15   important for aerosol activation. Activation calculations additionally require three meteorological parameters (temperature, pressure, and V). To assess this parameter space, we sample 70 snapshots of instantaneous aerosol and meteorology fields produced by a simulation of MARC forced with present-day aerosol and precursor gas emissions (Kim et al., 2014), and with interactive dust (Albani et al., 2014) and sea salt (Mårtensson et al., 2003) emissions. Only grid-cells located below $500\,\text{mbar}$ and between 70S and 70N were included in this assessment, since this is where liquid clouds and droplet nucleation are most

20   prevalent in the model.

As shown in Fig. 1, V tends to take the lower bound of $0.2\,\text{m}\,\text{s}^{-1}$ nearly 50% of the time in this sampling dataset and rarely exceed $1\,\text{m}\,\text{s}^{-1}$. On average, in continental regimes V is slightly larger ($0.41\,\text{m}\,\text{s}^{-1}$ vs $0.32\,\text{m}\,\text{s}^{-1}$ in ocean regimes), but has a long, positive tail maxing out between $3\,\text{m}\,\text{s}^{-1}$ to $4\,\text{m}\,\text{s}^{-1}$. The upper bound on V is never reached in our model output sampling.

25   The distributions of key aerosol size distribution parameters are shown in Fig. 2. Because particle size is a critical parameter in assessing aerosol activation, we show distributions of μ_g, the geometric mean radius, instead of M. Many size distribution parameters can vary over several orders of magnitude across the globe in even a single timestep. Values for N and μ_g in each mode tend to co-vary, though there is significant heterogeneity in the overall aerosol mixing state (not shown here). For most modes, the parameter distributions have long tails where N and μ_g are both small. In these cases, few particles are likely to

30   activate due to their small size, and would yield few cloud droplets even if they did. The number concentrations of DST and SSLT mode particles are typically very small compared to those in the ACC, MOS, and MBS modes. Aitken-mode particles in

MARC are generally small but numerous, based on how the mode is defined in the model (Table 1). We note that on average, MOS-mode particles have an intermediate κ of 0.27 ± 0.04 (standard deviation) which varies with local sulfate abundance.

The emulation method used by Rothenberg and Wang (2016) is designed to work with an arbitrary input parameter space. However, we restrict the aerosol activation emulation parameter space by eliminating parameters which exert little or no influence on the activation process. This reduces the number of parcel model simulations necessary to train emulators as well as their complexity. The OC and BC modes are assumed to be hydrophobic and hence cannot serve as cloud condensation nuclei (CCN); NUC particles are typically too small to serve as CCN, although they could be present in large enough number to influence supersaturation and droplet nucleation. However, where nucleation mode particles are usually most abundant when larger sulfate particles (especially ACC) are also present in large concentrations and have a greater influence on activation dynamics. We neglect size distribution parameters from these modes to build our emulators.

To further reduce the emulation parameter space, we assess the relative importance of each individual aerosol mode and its influence on activation dynamics using an ensemble of iterative, single-mode activation calculations using a detailed reference parcel model (Rothenberg and Wang, 2016). For each of our n aerosol modes, we pick a test mode and run a parcel model simulation to compute the supersaturation maximum (S_{max}) achieved in an updraft in which that mode is embedded. The mode which produces the minimum S_{max} is said to be the "dominant" mode, and we record its size distribution. We then remove that aerosol from the original set of n modes, and re-visit each of the $n-1$ remaining modes, including now the first dominant mode along with each new test mode. The end result of $n-1$ iterations is a sorting of the modes based on their contribution to reducing S_{max} in the parcel model simulations.

We apply this algorithm to a set of 50,000 aerosol size distribution and meteorology parameters taken from our reference MARC simulation. Overall, ACC is the dominant mode in 96.5% of the sample cases. Infrequently, MOS and small dust (DST1) are the dominant mode, accounting for all of the remaining cases. When ACC dominates the activation dynamics, MBS/MOS or smallest sea salt mode (SSLT1) is the second-most dominant (in 10.3%, 36.2%, and 52.8% of cases, respectively). In 85% of all the sample cases, three of ACC, MOS, MBS, or SSLT01 comprise the top three dominant modes.

Figure 3 illustrates the potential error in calculating both S_{max} and N_{act} (number concentration of activated particles or CDNC) for each iteration of the activation calculations relative to using the complete MARC aerosol population, aggregated by the first dominant mode. Using a subset of the modes leads to an over-prediction of N_{act} due to predicting a high value for S_{max}. This is consistent with the physics of the activation problem; the presence of more aerosol surface area on which condensation can occur tends to deplete water vapor in the ascending parcel more quickly, suppressing the development of a higher S_{max}. But as Fig. 3 shows, this over-prediction decreases rapidly as additional modes are included in the calculation. Part of this decrease is related to the fact that adding modes in each iteration captures a higher fraction of the total aerosol number; on average, the first dominant mode contains $70\% \pm 27\%$ of the total aerosol number, which increases to $80\% \pm 17\%$ and $89\% \pm 13\%$ after adding the second and third modes. Following this increase in fraction of the total aerosol number included by the dominant mode set in each iteration is a decrease in the absolute error in N_{act} relative to the full aerosol population, with an average of less than $1\,\mathrm{cm}^{-3}$ and max of $57\,\mathrm{cm}^{-3}$ by the third iteration. Although giant CCN particles are well-represented

in MARC by dust and sea salt modes, they are typically too few in number to largely influence droplet nucleation except in rare cases where other modes are not abundant.

Based on the sampling and iterative calculations presented in this section, we define the activation emulation parameter space as in Table 2. We consider the full range of potential values for V, but restrict further the emulation to include the most frequent "dominant" aerosol modes. Table 2 further defines ranges considered for each of the aerosol and meteorological parameters, and indicates the percentile of the sampling data at which the extreme values of each range occur. In all cases we include upper bounds above the 98th percentile of the sampling data. For the giant CCN modes (DST and SSLT), we restrict the lower ranges of N but at values where small changes in N could only have very small contributions to changes in N_{act}.

2.3 Emulator construction

The following sections briefly describe the chosen cloud parcel model and emulation technique, polynomial chaos expansion. For more details on both techniques and their application, we refer the reader to Rothenberg and Wang (2016), which derives activation emulators for a simplified, single lognormal aerosol mode.

2.3.1 Parcel Model

Adiabatic cloud parcel models are a standard modeling tool for detailed assessments of aerosol activation and other studies focused on the composition of atmospheric particulates (Seinfeld and Pandis, 2006). In such a model, a constant-speed updraft drives adiabatic cooling in a closed, zero-dimensional air parcel within which are any number or configuration of aerosol particles. Initially prescribed a temperature, pressure, and water vapor content, the cooling parcel eventually develops a supersaturation with respect to water vapor. In a sufficiently supersaturated environment, water vapor condenses on particulate surfaces. This condensation releases latent heat and slows down the cooling of the parcel, but more importantly acts to sink water vapor mass. In terms of producing supersaturation, these counter-acting processes can be expressed

$$\frac{dS}{dt} = \alpha(T,P)V - \gamma(T,P)\frac{dw_c}{dt} \tag{1}$$

where V is the updraft speed, $\alpha(T,P) = (gM_wL/c_pRT^2) - (gM_a/RT)$ and $\gamma(T,P) = (PM_a/e_sM_w) + (M_wL^2/c_pRT^2)$ are functions weakly dependent on temperature and pressure (Leaitch et al., 1986), M_w and M_a are the molecular weights of water and air, L is the latent heat of vaporization of water, c_p is the specific heat of dry air at constant pressure, R is the universal gas constant, g is the acceleration due to gravity, e_s is the saturation vapor pressure, and w_c is the liquid cloud water mass mixing ratio (please refer to Appendix A of Rothenberg and Wang, 2016, for more details).

At some time t, the balance between heating due to latent heat release and cooling due to the parcel's adiabatic ascent will approximately balance such that $\frac{dS}{dt} = 0$ and a supersaturation maximum, S_{max}, will occur. Thereafter, S generally decreases, relaxing to some value close to unity as condensation drives droplet growth, quenching the ambient water vapor surplus. Beyond this point, the aerosol bifurcates into two populations: new cloud droplets which will continue to grow due to condensation

and eventually collision and coalescence, and interstitial haze particles which may have taken up some water according to their hygroscopic mass, but upon which further condensation is not thermodynamically favorable.

2.3.2 Polynomial Chaos Expansion

We emulate the behavior of the detailed parcel model by applying the probabilistic collocation method (PCM; Tatang et al., 1997). PCM is a method of polynomial chaos expansion which seeks to construct a model response surface by mapping input parameters related to the initial conditions and behavior of a model to some response measured from the model. This process yields a computationally efficient yet accurate reproduction of the model.

The PCM is a non-intrusive technique which does not require modifications to an existing model in order to be applied. Instead, the PCM treats the original, full-complexity model as a black box and the chosen set of M' input parameters as independent, random variables, $\mathbf{X} = X_1, \ldots, X_{M'}$, each with an associated probability density function. This PDF is used as a weighting function to derive a family of orthogonal polynomials which are used as the bases for the polynomial chaos expansion to be constructed, ϕ. Using a finite number of these bases and choosing some model response, R, we write the polynomial chaos expansion as

$$R \approx \sum_{j=0}^{P} \alpha_j \phi_j(\mathbf{X}) \tag{2}$$

Such an expression has $N_t = P + 1 = (M' + p)!/(M'!p!)$ total terms, since a given chaos expansion of order p will contain $p+1$ basis terms for each input parameter and combinations thereof. The coefficients α_j are computed by evaluating the original model at a set of particular set of sample points, recording the response of the model, and solving a regression problem. Those sample points are generated by taking the roots of the orthogonal polynomials associated with each of the input parameters and their random variables.

In order to compute the polynomial chaos expansions, we use the Design Analysis Kit for Optimization and Terascale Applications (DAKOTA; Adams et al., 2014), version 6.1. This software automates the process of generating input parameter sets, sampling the full-complexity model to be emulated, and constructing the polynomial chaos expansion. Furthermore, it provides many useful statistical properties of the sample dataset and the chaos expansions themselves.

2.3.3 Application to MARC aerosol

We extend the idealized calculations in Rothenberg and Wang (2016) by applying the parcel model and chaos expansion technique previously described to construct emulators of aerosol activation suitable for use in MARC. Using the parameter set and value ranges summarized in Table 2, we consider two different cases for emulation: a "main" case which includes just the ACC, MBS, and MOS modes, and a "gCCN" (giant CCN) scheme which adds in dust and sea salt modes. We treat these two cases separately because of the non-linear influence of gCCN on activation dynamics and cloud microphysics, which critically depends on the ambient CCN burden (Feingold et al., 1999).

For emulation, all the aerosol size distribution parameters are transformed using a logarithm, since they can take on values that span several orders of magnitude. We then construct uniform distributions with the associated ranges of values for each transformed parameter (Table 1). Using these uniform distributions, we cast all of the input parameters as independent random variables with a uniform probability density function covering the range of each corresponding distribution.

5 This methodology represents a compromise between using high-fidelity representation of the aerosol size distribution parameters for our emulation, and the desire to build an emulator that can later be used in a GCM. However, we note that the distributions of aerosol parameters simulated by MARC are neither normal nor independent from one another. For instance, over remote maritime regions, total aerosol number concentration tends to be small but dominated by sea salt and small sulfate particles. In contrast, continental regions with anthropogenic emissions may feature much higher burdens of carbonaceous

10 aerosol. Using both numerically-generated orthogonal polynomials and statistical transformations, both of these complications can be handled directly (Isukapalli, 1999; Feinberg and Langtangen, 2015), but not without a trade-off. First, regions of the multi-dimensional input parameter space with low but non-zero probability of occurrence are unlikely to be sampled when building the emulator. For a sufficiently linear response function, this may not be an issue. More importantly, it becomes difficult to serialize the resulting chaos expansion in a form that can be re-produced for use later. This requires storing both the

15 weights (coefficients) of the numerically-generated orthogonal polynomials from the chaos expansion and some representation of how to re-construct the transformed joint probability distribution formed from the input parameter probability distributions. Employing uniform probability distributions for each input parameter and assuming independence solves both of these problems, since canonical orthogonal polynomials with well-known techniques to efficiently generate them can be used, and few additional, costly transformations of the parameter samples are necessary to run the emulator. We study the impact of

20 this choice by evaluating samples of input parameter sets drawn from their true joint probability distribution (as simulated by MARC) in Section 3.

These parameters are used to drive parcel model simulations where we record $\log_{10} S_{\max}$ as the response variable. This value can then be used to diagnose the number concentration of droplets nucleated by assuming that any particles which experience their Köhler theory-predicted critical supersaturation activate. We note that although this does not resolve the issue of kinetic

25 limitations on droplet growth and its potential to cause an under-prediction in droplet number (Nenes et al., 2001), unlike existing activation schemes, our emulator accounts for the feedback of these effects on $S_{\max}$, so one avenue of its impact is lessened.

The emulators constructed through this process are functions which map $\log_{10}(S_{\max})$ to a set of values from our input parameter space,

$$\log_{10}(S_{\max}) = f(\log_{10} N_{\mathrm{ACC}}, \log_{10} N_{\mathrm{MOS}}, \log_{10} N_{\mathrm{MBS}},$$
$$\log_{10} \mu_{\mathrm{ACC}}, \log_{10} \mu_{\mathrm{MOS}}, \log_{10} \mu_{\mathrm{MBS}}, \kappa_{\mathrm{MOS}}, \log_{10} V, P, T[, m$$

30 $$\log_{10} N_{\mathrm{DST01}}, \log_{10} N_{\mathrm{DST02}}, \log_{10} N_{\mathrm{SSLT01}}]) \qquad (3)$$

From a prediction of the $S_{\max}$ achieved in an ascending parcel with the conditions passed to the emulator, we can then diagnose aerosol activation by re-writing the lognormal size distribution for each mode as a function of critical supersaturation (Ghan et al., 2011) to yield an expression

$$N_{\text{act}} = \sum_{i=1}^{n} \frac{N_{t,i}}{2} \left(1 - \text{erf} \left[2 \ln \left(\frac{S_{m,i}}{S_{\max}} \right) / (3\sqrt{2} \ln \sigma_{g,i}) \right] \right) \tag{4}$$

5   where $S_{m,i}$ is the critical supersaturation for a particle of radius $\mu_{g,i}$ from mode i.

3 Evaluation of Emulators

We evaluate our emulators by applying them to both a synthetic sample of input parameters as well as real samples drawn from a MARC simulation. In all of our comparisons, we study third and fourth order chaos expansions both excluding ("main") and including ("gCCN") the coarse dust and sea salt modes.

10   As a reference, we compute activation statistics for each sample using both a detailed parcel model and two widely-used activation schemes. The first parameterization, by Abdul-Razzak and Ghan (2000) (ARG), uses a pseudo-analytical solution to an integro-differential equation derived from the original adiabatic parcel model system; the second, by Morales Betancourt and Nenes (2014b) (MBN), applies an iterative scheme to partition the aerosol population into two subsets, and uses different limits on the underlying analytical formulas to derive a maximum supersaturation. Because it requires a sequence of iterations to run, the MBN scheme is more computationally expensive than the ARG scheme, but has the potential to include more

15   detailed links between particle composition and condensation (Kumar et al., 2009) or entrainment into the parcel (Barahona and Nenes, 2007). Both the ARG and MBN schemes rely on some parameters fit to parcel model simulations conceptually similar (but different in implementation) to the one emulated here.

3.1 Input Parameter Space Sampling

20   Using the parameter space defined in Table 2, $n = 10,000$ sample parameter sets were drawn using maximin Latin Hypercube Sampling (LHS). This randomized sampling method helps to ensure that the full aerosol and meteorology parameter space is studied while assessing its performance.

Figure 4 compares the performance of each emulator and the two reference activation scheme against parcel model simulations using all of the LHS samples for the "main" aerosol parameter sets. In the simulations, higher updraft speeds are nearly

25   always associated with a much higher supersaturation maximum. For the emulators, accuracy tends to increase on average going from the 3rd order (Fig. 4-a) to the 4th order (Fig. 4-b) scheme, although there is higher variance in the relative error compared to the parcel model at higher updraft speeds. With respect to the driving updraft speed, though, there isn't a consistent mode of bias—on average, the relative error is very low. The same does not hold true for the two reference schemes. The ARG scheme (Fig. 4-c) tends to mis-predict supersaturation maxima at higher updraft speeds but is relatively well-calibrated

30   at lower updraft speeds, yielding a lower supersaturation maximum. On the other hand, the MBN scheme (Fig. 4-d) is more

accurate and better-calibrated than either of the emulators or the ARG scheme, especially at higher updraft speeds, but tends to spuriously over-estimate S_{max} for weak updraft speeds.

Figure 5 extends this evaluated to diagnosed CDNC (N_{act}) nucleated for each predicted S_{max}. For all the schemes, there can be substantial differences between the parcel model and each parameterization. This is particularly the case in regimes which produce smaller total CDNC, either due to a lower driving updraft velocity, or a lower total aerosol number available to activate. The MBN scheme tends to consistently nucleate a higher number of droplets with respect to the parcel model, especially in situations which should have very few droplets, such as those where the total aerosol concentration is below $10\,cm^{-3}$. Although it does not have as consistent of a bias, the ARG scheme can both egregiously over-predict and under-predict CDNC, with these biases exaggerated at lower updraft speeds. By comparison, the emulators show much less overall bias. The mean error for the emulators follows that of S_{max} and is small, but there is variance which tends to impart a small low or high bias on its estimates.

Both of these sets of plots are repeated in Figs. 6 and 7, but for the "gCCN" experiment. Qualitatively, the results for all parameterizations are very similar, with the same overall biases - especially for the ARG and MBN parameterizations. The emulators tend not to perform as well overall in the "gCCN" cases, although they are still the most highly-calibrated scheme and do not have the velocity-regime errors that the MBN scheme has. In both the "main" and "gCCN" parameter sets, the MBN scheme tends to more regularly predict too many cloud droplets, save for polluted regimes giving rise to $100\,cm^{-3}$ droplets where that bias reverses and the scheme has a tendency to under-predict droplet number. Neither the emulators nor the ARG show this same tendency in bias.

These differences in bias are most likely related to the choice of parcel model used in testing and building the ARG and MBN schemes; because each scheme relies on some empirical tuning to parcel model calculations, details in the implementation of each parcel model which influence its sensitivity should show up in ensemble evaluations of each activation scheme. The "gCCN" case is more taxing to simulate with parcel models using a Lagrangian description of the particle size distribution, because condensational growth is computed for each particle bin simultaneously. The stiffness ratio in this case will be extremely large, as the liquid water uptake by small particles in the main aerosol modes is much slower than those in the giant CCN modes. Although modern ODE solvers can automatically handle these scenarios, the subjective choice of which particular solver and how to discretize the giant CCN population (how many bins per mode) could influence the sensitivity of S_{max} to changes in the model inputs and account for the differences observed here.

To better summarize the results in Figs. 4 to 7, summary statistics on the error of each scheme versus their corresponding parcel model calculations are shown in Table 3. In both sampling cases, all of the parameterizations show a strong linear correlation (r^2) between their predictions and the result of the parcel model. The emulators (PCM Order p) predict S_{max} with lower overall absolute and relative error, but with a much higher variance (not shown here). However, that lower error does not always translate into the emulators being the most accurate absolute predictors of N_{act}. For the "gCCN" parameters, the ARG scheme predicts N_{act} with a lower average mean relative error. In both parameter sets, the MBN scheme is the least accurate compared to the parcel model used in these sampling calculations.

3.2 MARC Aerosol Sampling

Although the sampling in the previous section fully exhausts the input parameter space over which aerosol activation may need to be assessed, it undoubtedly samples from aerosol and meteorological conditions which may not be likely to occur in the real world. To better understand the performance and potential bias of the emulators developed here and the existing activation schemes, then, we also studied a sample of $n = 10,000$ aerosol and meteorology parameter sets drawn directly from the MARC simulation previously used to study the distributions of simulated aerosol parameters, in contrast with the previous LHS sample. Error statistics from this sample, since they use model-produced parameter sets, are more representative of the real-world performance of the emulators and physically-based activation schemes. All of the schemes were evaluated again using these parameter sets and the same detailed parcel model. This includes the "main" and "gCCN" emulators, which allows us to identify the importance of including the dust and sea salt modes as predictors in the chaos expansions. The parameters in these sets occasionally include values outside the ranges defined in Table 2 and studied in the previous section. These cases are more frequently associated with very low total aerosol number concentration, especially over the ocean where anthropogenic aerosols are limited and natural aerosols—which have a lower overall number burden—dominate. Because the distributions of the parameters of aerosol samples from oceanic and continental grid cells differ in this fundamental way, we break down the following analysis to reflect those differences. As in the previous sampling experiment, summary statistics on the performance of each emulator, alongside the ARG and MBN schemes, are detailed in Table 4.

Qualitatively, all of the activation schemes perform similarly when evaluated against the MARC parameters as compared to the more generic sampling in the previous section. Figure 8 summarizes distributions of relative error in N_{act} over land and ocean for each scheme. Neither the ARG nor the MBN scheme show much difference in error for the two regimes, although on average, the MBN tends to under-predict N_{act}. This under-prediction usually occurs in regimes with higher updraft speeds and thus higher overall droplet number concentrations. In conditions with weaker updraft speeds, the MBN scheme instead tends to slightly over-predict N_{act}. The ARG scheme is well-calibrated in both regimes.

The emulators derived here do not fare as well as the physically-based parameterizations when using the MARC samples. Both 3rd order schemes tend to over-predict droplet number over oceans, and under-predict it over land, but with an extremely large variance extending to $\pm 100\%$. However, including the effects of giant CCN measurably improves the performance of the 3rd order emulators in oceanic regimes. Increasing the order of the emulator also has a significant impact on their accuracy; the 4th order scheme which neglects giant CCN produces smaller absolute error than the ARG and MBN schemes on average, and shows little bias between land and ocean regimes, indicating good convergence with its parent parcel model. On the other hand, the gCCN scheme has not yet converged by including 4th order terms, even while its mean error statistics improve. Particularly troublesome is a secondary mode of extreme under-prediction of droplet number of oceanic regimes, but this metric is deceptive. For very low total aerosol number concentrations—with particle number in the single-digits per cubic centimeter—the 4th order "gCCN" scheme tends to predict half as many droplets as parcel model calculations indicate should form. This typically occurs when one or more of the input size distribution parameters (in particular, the number concentration) for the natural aerosol falls below the minimum threshold where the emulator was trained. When the emulators encounter inputs

greater (lower) than these thresholds, they hold them to the maximum (minimum) value in its training range. This follows the assumption that the bounds for each parameter cover the entire range over which activation is sensitive to changes in that input. As a result, activation should be relatively insensitive to changes near the maximum or minimum values in the range for each parameter. With respect to number concentration, this must be the case; populations with fewer than 10^{-3} particles cm^{-3} offer very little surface area for condensation and water vapor sink, and simply cannot exert a strong influence over the developing supersaturation in the parcel. That the 4th order "gCCN" emulator produces too high of sensitivity in this regime suggests that statistical over-fitting is occurring near the extremes of the input parameter space.

To contextualize these differences in N_{act} bias over different geographical regimes, Fig. 9 re-maps the testing samples back to the original MARC grid. Here, the difference in regional biases becomes much clearer. Virtually everywhere, the MBN scheme is biased a little low, but there is no systematic difference in this bias between land or ocean, or by geographical areas. The ARG scheme and the 4th order "main" scheme show a different pattern; the ocean-land contrast in bias is clearly visible in the northern hemisphere. Furthermore, the bias is typically positive over maritime regions, but negative over regions with anthropogenic aerosol influence. In particular, these regions include Europe and southeastern Asia - where aerosol distributions are dominated by anthropogenic sulfate and black carbon - and over north central Africa - where the aerosol is a mixture of both dust and organic carbon emissions from biomass burning. In the zonal average, the main_4 scheme is virtually identical to the ARG scheme. However, both cases as well as with MBN, there are larger biases over the southern parts of the oceans, where the aerosol is predominantly comprised of sea salt and smaller sulfate particles produced indirectly through the emission of DMS.

Figure 9 also illustrates the poor performance of the gCCN_4 scheme, which under-predicts N_{act} nearly everywhere, but especially so in the southern portions of the ocean basins. The consistent under-prediction in this region explains the bimodal distribution over the ocean hinted at in Fig. 8. The gCCN_4 scheme does not perform too dissimilarly than the other schemes over regions with anthropogenic pollution or with mostly dust aerosol.

3.3 Implementation in MARC and Computational Expense

The ARG scheme is the original activation parameterization used within CAM5.3 to assess cloud droplet nucleation. We implemented the MBN scheme as an alternative in MARC, as well as an interface for chaos expansion-based schemes. To use the emulators derived in this work, one must provide a NetCDF file which contains at least three pieces of information:

- an m-length vector of coefficients corresponding to each of the m terms in the expansion

- a $2{\times}\mathrm{p}$-shape matrix which defines the lower and upper bounds for each of the input parameters

- a $m \times p$ matrix of integers correspond to the p variable expansion orders for each of the m terms in the expansion

MARC caches these vectors and matrices in memory at start-up, just as it caches several time-invariant terms used in both the ARG and the MBN schemes for each of the CCN-providing aerosol modes.

To estimate the impact of each scheme on MARC's performance, we performed a set of three-month simulations initialized with fully spun-up aerosol and meteorology fields from a previous experiment. The simulations were conducted using 480 MPI tasks with two threads allocated to each task. Using the default configuration of MARC with the ARG scheme, the atmosphere component of the model averaged $6.1\,\mathrm{s}$ per model day. The MBN scheme averaged 7% longer per model day, while the emulators tended to be comparable to the ARG scheme. Per model day, both the "main" schemes were comparable to within 0.4% of the ARG scheme's performance, with the higher-order scheme costing an additional 0.16%. Similarly, the "gCCN" schemes also compared similarly with the ARG scheme; the 3rd-order scheme 0.15% faster than the ARG scheme, but the 4th-order scheme was 3% slower.

Adding additional parameters to the chaos expansion underpinning the emulators would continue to add overhead to each evaluation by increasing the number of terms in the expansion. However, a larger penalty is incurred by increasing the expansion order for a given set of parameters, because this produces a much larger increase in the number of terms added to the expansion than adding a single parameter for the same order expansion. An assessment of the offline implementations of each scheme used in the analysis in the previous section yielded similar results.

4   Discussion and Conclusions

In this work, we extend the metamodeling technique of Rothenberg and Wang (2016) in order to assess activation of a complex, multi-modal aerosol mixture simulated by a modern aerosol-climate model. Simultaneously, we characterize the performance of both our new emulators for aerosol activation and two widely-used schemes from the literature, focusing on that same high-dimensional, complex aerosol parameter space. To identify the most important factors impacting activation in that complex parameter space, we apply a physically-based approach to assess the sensitivity of activation statistics to the composition of the aerosol size distribution. Finally, we explore contrasts between aerosol and meteorology regimes over land and ocean, noting the potential for different biases in assessed cloud droplet number depending on the choice of activation scheme used in a particular global modeling application.

In ensembles of iterative calculations using a large sample of aerosol size distributions from a coupled aerosol-climate model, we note that typically, a single mode tends to dominate activation or otherwise strongly predict the total number of droplets nucleated. This approach to understanding the sensitivities of activation dynamics on the underlying aerosol population is distinct from previously-published approaches in the literature. For instance, Karydis et al. (2012) and Morales Betancourt and Nenes (2014a) apply an adjoint approach to derive the sensitivity of aerosol activation to perturbations in input parameters supplied to activation schemes. Detailed calculations using this approach yield a map of local sensitivities or gradients in the relationship between, for example, N_{act} and one input parameter while holding all others constant, and are thus difficult to interpret. The iterative calculations performed here aim instead to address the global sensitivity of activation to configurations of an aerosol population.

In terms of predicting CDNC, the accumulation mode sulfate (ACC) alone serves as a good proxy for the activity of a full aerosol population in many cases, including in the presence of giant CCN and a wide swath myriad updraft regimes. However,

it is known that giant CCN exert a larger influence on precipitation formation in cleaner regimes (Feingold et al., 1999; Yin et al., 2000), and thus where ACC is abundant (especially over continents) the impacts of giant CCN on N_{act} can be muted. This result is also model-dependent in some sense; the ACC mode in MARC is not only ubiquitous, but may inadvertently (and subjectively) take a range of mean particle sizes for which aerosol activation is especially sensitive. At the same time, the coarse dust and sea salt modes in MARC, on average, hold too small a number concentration to dramatically impact activation calculations save for remote maritime regions far removed from anthropogenic sources. However, the presence of sea salt as one of the modes most frequently ranked in the top three influencers of activation agrees with previous results which indicate the presence of giant CCN can influence activation dynamics (e.g. Barahona et al., 2010).

The fact that a single mode can place such a strong constraint on aerosol activation is useful for attempts seeking to extend look-up table methods for building parameterizations. If two modes—an accumulation-size and a coarse-size—accurately predict aerosol activation, then one can constrain the look-up table to just a few key aerosol size distribution parameters. The inclusion of variable aerosol composition would still likely make employing a look-up table in a global model unwieldy, though, necessitating more sophisticated approaches such as the metamodeling technique adopted here.

When sampling against the full training parameter space, our emulators perform capably. Neglecting the influence of the giant CCN modes, the mean relative error in predicting $\log_{10} S_{\max}$ for the emulators is less than 1% , which translates to mean relative errors for N_{act} of 9.2% and 8.9%. Including the giant CCN mode appears, at first, to dramatically increase the performance of the emulator, bringing those same metrics down to 0.3% and 6.9% for the 4th-order scheme. Relative to the ARG and MBN schemes, the emulators are much more accurate on average when compared to our reference parcel model. Instead, we emphasize that the comparison of our emulators with the ARG and MBN scheme is motivated as part of a broader attempt to understand how the fundamental activation process dictates the simulated relationship between aerosol burden and CDNC in a global climate model. This relationship is critical for understanding uncertainty in the indirect effect produced by any given model, since the simulated background CDNC strongly correlates with its strength (Hoose et al., 2009).

Assessing the relative performance of activation schemes which, for all intents and purposes, perform extremely well at reproducing their own reference parcel models, is a critical step in establishing the parametric uncertainty in translating aerosol to droplet numbers and which underlies uncertainty in global model estimates of the indirect effect.

For this reason, we supplemented the evaluation of our emulators by using a second set of input parameter samples drawn from aerosol fields simulated by an aerosol-climate model. In contrast with previous studies, we use instantaneous fields in lieu of monthly or annual averages for our samples. Activation is inherently a fast process; because the microphysics schemes in aerosol-cloud models directly account for a tendency of new droplets formed via nucleation, the activation parameterization in any model will be called every time-step and in every grid-cell where clouds are occurring. Assessing activation schemes using temporally-averaged aerosol fields risks missing some combinations of input parameters and limiting the range of values for which the scheme will need to accurately perform.

Most of the emulators and schemes tested here perform differently in oceanic and continental regimes, owing to the relative abundance of natural and anthropogenic aerosols in each. When focusing on the narrower range of aerosol parameters present in MARC (in comparison with the larger parameter space on which the emulators were trained), the emulators which explicitly

account for giant CCN perform poorly, especially in maritime regimes dominated by sea salt. However, their counterpart performs nearly identical to the ARG scheme, showing a slight over-prediction of N_{act} in maritime regimes and a slight under-prediction over continents. In the global average, the emulator agrees better with the detailed parcel model than the ARG scheme. By comparison, the MBN scheme, while prone to under-predicting N_{act} in both regimes in these calculations, shows

5   far less variance in its mis-prediction. This would suggest the MBN scheme actually performs extremely well—it is simply calibrated against a different baseline. Both the ARG and MBN schemes were developed using parcel models conceptually-similar to the one employed here, but which differ in the details of their implementation. Furthermore, both schemes use some parameters computed from empirical fits to their associated parcel models. As a result, we do not expect perfect agreement between all the schemes evaluated her, and instead note the importance of the relative differences between the CDNC simulated

10   for each one. In particular, the MBN scheme does not show a large difference in relative error versus our parcel model between ocean and land regions, suggesting it is appropriately sensitive to a large range of different aerosol populations.

The results presented here have important implications for global modeling studies seeking to quantify uncertainty in the aerosol indirect effect on climate. While different activation schemes generally perform equally well when faced with idealized sets of input parameters (Ghan et al., 2011), their application in coupled aerosol-climate models may not be straightforward.

15   Relative to parcel model calculations, activation schemes can likely show biases in predicting cloud droplet number in different regions of the world owing to spatial heterogeneity in the underlying aerosol and meteorology parameter distributions. This, in turn, could lead to biases in cloud radiative forcing and diagnosis of the indirect effect. For instance, Ghan et al. (2011) performed a pair of GCM experiments using two different schemes and observed a 10% difference in the global average CDNC, which produced a $0.2\,\mathrm{W\,m^{-2}}$ difference in the indirect effect. Gantt et al. (2014) similarly showed large regional differences

20   in CDNC when using a different set of activation schemes, leading to a spread $0.9\,\mathrm{W\,m^{-2}}$ in global average shortwave cloud forcing. These impacts on CDNC and indirect effect resulting from using different activation schemes will necessarily be model-dependent, since the formulation of the basic activation diagnostic in each model is intertwined with regional and global variability in their simulated aerosol size distributions.

Future work should seek to systematically assess the differences in cloud microphysical processes and aerosol-cloud inter-

25   actions arising from choice of activation schemes in aerosol-climate models. As this work illustrates, employing emulators of detailed parcel model calculations which include complex chemical and physical effects on activation will aid with this task, since additional effects (e.g. changes in droplet surface tension due to organic surfactants (Abdul-Razzak and Ghan, 2004)) can more easily be incorporated into parameterizations built using this method. However, this work must proceed in tandem with efforts to place strong constraints on the climatology and variability of cloud droplet number concentration across regions and

30   meteorological regimes. The synthesis of these two lines of work may provide the necessary constraints to diagnose systematic biases in the representation of fundamental aerosol-cloud interactions in global aerosol-climate models and thus reconcile the disagreement between model- and satellite-derived estimates of the indirect effect.

Appendix A: Code and Data Availability

A git repository archiving the scripts used to generate the chaos expansions can be found at https://github.mit.edu/darothen/marc_pcm_activation. For the convenience of the reader, an up-to-date commit (c71f8ca9bd4) has been included in the Supplementary Materials. Dependencies of these scripts are recorded in the README file therein. The sampling datasets generated for analysis in this work are archived with DOI 10.5281/zenodo.60937. The source code and documentation for the pyrcel cloud parcel model are archived with DOI 10.5281/zenodo.46127, and can be accessed from http://github.com/darothen/pyrcel.

Acknowledgements. The work in this study was supported by the National Science Foundation Graduate Research Fellowship Program under both NSF Grant 1122374 and NSF Grant AGS-1339264; the National Research Foundation of Singapore through the Singapore–MIT Alliance for Research and Technology and the interdisciplinary research group of the Center for Environmental Sensing and Modeling; and the U.S. Department of Energy, Office of Science (DE-FG02-94ER61937). We thank Steve Ghan (PNNL) and Athanasios Nenes (Georgia Tech) for reference implementations of their activation parameterizations. We would also like to thank Graham Mann and three anonymous reviews for comments that helped improve this manuscript.

References

Abdul-Razzak, H. and Ghan, S. J.: A parameterization of aerosol activation 2. Multiple aerosol types, Journal of Geophysical Research, 105, 6837, doi:10.1029/1999JD901161, http://www.agu.org/pubs/crossref/2000/1999JD901161.shtmlhttp://dx.doi.org/10.1029/1999jd901161, 2000.

5   Abdul-Razzak, H. and Ghan, S. J.: Parameterization of the influence of organic surfactants on aerosol activation, Journal of Geophysical Research: Atmospheres, 109, n/a–n/a, doi:10.1029/2003JD004043, http://doi.wiley.com/10.1029/2003JD004043, 2004.

Adams, B. M., Ebeida, M. S., Eldred, M. S., Jakeman, J. D., Swiler, L. P., Stephens, J. A., Vigil, D. M., Wildey, T. M., Bohnhoff, W. J., Dalbey, K. R., Eddy, J. P., Hu, K. T., Bauman, L. E., and Hough, P. D.: DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.0 User's Manual, Tech. rep.,

10   Sandia National Laboratories, Albuquerque, New Mexico, 2014.

Albani, S., Mahowald, N. M., Perry, A. T., Scanza, R. A., Zender, C. S., Heavens, N. G., Maggi, V., Kok, J. F., and Otto-Bliesner, B. L.: Improved dust representation in the Community Atmosphere Model, Journal of Advances in Modeling Earth Systems, pp. n/a–n/a, doi:10.1002/2013MS000279, http://doi.wiley.com/10.1002/2013MS000279, 2014.

Barahona, D. and Nenes, A.: Parameterization of cloud droplet formation in large-scale models: Including effects of entrainment, Journal of

15   Geophysical Research, 112, D16 206, doi:10.1029/2007JD008473, http://doi.wiley.com/10.1029/2007JD008473, 2007.

Barahona, D., West, R. E. L., Stier, P., Romakkaniemi, S., Kokkola, H., and Nenes, A.: Comprehensively accounting for the effect of giant CCN in cloud activation parameterizations, Atmospheric Chemistry and Physics, 10, 2467–2473, doi:10.5194/acp-10-2467-2010, http://www.atmos-chem-phys.net/10/2467/2010/, 2010.

Boucher, O., Randall, D., Artaxo, P., Bretherton, C. S., Feingold, G., Forster, P., Kerminen, V.-M., Kondo, Y., Liao, H., Lohmann, U., Rasch,

20   P., Satheesh, S. K., Sherwood, S., Stevens, B., and Zhang, X. Y.: Clouds and Aerosols, in: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, edited by Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S. K., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P. M., chap. 7, pp. 571–657, Cambridge University Press, Cambridge, United Kingdom, and New York, NY, USA, 2013.

Carslaw, K. S., Lee, L. a., Reddington, C. L., Pringle, K. J., Rap, a., Forster, P. M., Mann, G. W., Spracklen, D. V., Woodhouse,

25   M. T., Regayre, L. a., and Pierce, J. R.: Large contribution of natural aerosols to uncertainty in indirect forcing., Nature, 503, 67–71, doi:10.1038/nature12674, http://www.ncbi.nlm.nih.gov/pubmed/24201280, 2013.

Cohard, J.-M. and Pinty, J.-P.: A comprehensive two-moment warm microphysical bulk scheme. I: Description and tests, Quart. J. Roy. Meteor. Soc., 126, 1815–1842, doi:10.1002/qj.49712656613, http://doi.wiley.com/10.1002/qj.49712656614, 2000.

Ervens, B., Feingold, G., and Kreidenweis, S. M.: Influence of water-soluble organic carbon on cloud drop number concentration, Journal of

30   Geophysical Research D: Atmospheres, 110, 1–14, doi:10.1029/2004JD005634, 2005.

Feinberg, J. and Langtangen, H. P.: Chaospy: An open source tool for designing methods of uncertainty quantification, Journal of Computational Science, 11, 46–57, doi:10.1016/j.jocs.2015.08.008, http://linkinghub.elsevier.com/retrieve/pii/S1877750315300119, 2015.

Feingold, G. and Heymsfield, A. J.: Parameterizations of Condensational Growth of Droplets for Use in General Circulation Models, Journal of the Atmospheric Sciences, 49, 2325–2342, http://journals.ametsoc.org/doi/abs/10.1175/1520-0469(1992)049{%}3C2325:

35   POCGOD{%}3E2.0.CO;2, 1992.

Feingold, G., Cotton, W. R., Kreidenweis, S. M., and Davis, J. T.: The Impact of Giant Cloud Condensation Nuclei on Drizzle Formation in Stratocumulus: Implications for Cloud Radiative Properties, Journal of the Atmospheric

Sciences, 56, 4100–4117, doi:10.1175/1520-0469(1999)056<4100:TIOGCC>2.0.CO;2, http://journals.ametsoc.org/doi/abs/10.1175/1520-0469{%}281999{%}29056{%}3C4100{%}3ATIOGCC{%}3E2.0.CO{%}3B2, 1999.

Fountoukis, C. and Nenes, A.: Continued development of a cloud droplet formation parameterization for global climate models, Journal of Geophysical Research, 110, D11 212, doi:10.1029/2004JD005591, http://www.agu.org/pubs/crossref/2005/2004JD005591.shtml, 2005.

5 Gantt, B., He, J., Zhang, X., Zhang, Y., and Nenes, A.: Incorporation of advanced aerosol activation treatments into CESM/CAM5: Model evaluation and impacts on aerosol indirect effects, Atmospheric Chemistry and Physics, 14, 7485–7497, doi:10.5194/acp-14-7485-2014, http://www.atmos-chem-phys.net/14/7485/2014/, 2014.

Ghan, S. J., Leung, L. R., Easter, R. C., and Abdul-Razzak, H.: Prediction of cloud droplet number in a general circulation model, Journal of Geophysical Research: Atmospheres, 102, 21 777–21 794, doi:10.1029/97JD01810, http://doi.wiley.com/10.1029/97JD01810, 1997.

10 Ghan, S. J., Abdul-Razzak, H., Nenes, A., Ming, Y., Liu, X., Ovchinnikov, M., Shipway, B., Meskhidze, N., Xu, J., and Shi, X.: Droplet nucleation: Physically-based parameterizations and comparative evaluation, Journal of Advances in Modeling Earth Systems, 3, M10 001, doi:10.1029/2011MS000074, http://dx.doi.org/10.1029/2011MS000074, 2011.

Ghan, S. J. S., Chung, C. C. C., and Penner, J. E. J.: A parameterization of cloud droplet nucleation part I: single aerosol type, Atmospheric Research, 30, 198–221, doi:10.1016/0169-8095(93)90024-I, http://www.sciencedirect.com/science/article/pii/016980959390024Ihttp://linkinghub.elsevier.com/retrieve/pii/016980959390024I, 1993.

15 Hoose, C., Kristjánsson, J. E., Iversen, T., Kirkevåg, a., Seland, Ø., and Gettelman, a.: Constraining cloud droplet number concentration in GCMs suppresses the aerosol indirect effect, Geophysical Research Letters, 36, L12 807, doi:10.1029/2009GL038568, http://doi.wiley.com/10.1029/2009GL038568, 2009.

Isukapalli, S. S.: Uncertainty analysis of transport-transformation models, Ph.D. thesis, Rutgers, The State University of New Jersey, 1999.

20 Karydis, V. A., Capps, S. L., Russell, A. G., and Nenes, A.: Adjoint sensitivity of global cloud droplet number to aerosol and dynamical parameters, Atmospheric Chemistry and Physics, 12, 9041–9055, doi:10.5194/acp-12-9041-2012, http://www.atmos-chem-phys.net/12/9041/2012/, 2012.

Kim, D., Wang, C., Ekman, A. M. L., Barth, M. C., and Rasch, P. J.: Distribution and direct radiative forcing of carbonaceous and sulfate aerosols in an interactive size-resolving aerosol–climate model, Journal of Geophysical Research, 113, D16 309, doi:10.1029/2007JD009756, http://www.agu.org/pubs/crossref/2008/2007JD009756.shtml, 2008.

25 Kim, D., Wang, C., Ekman, A. M. L., Barth, M. C., and Lee, D.-I.: The responses of cloudiness to the direct radiative effect of sulfate and carbonaceous aerosols, Journal of Geophysical Research: Atmospheres, pp. n/a–n/a, doi:10.1002/2013JD020529, http://doi.wiley.com/10.1002/2013JD020529, 2014.

Kumar, P., Sokolik, I. N., and Nenes, a.: Parameterization of cloud droplet formation for global and regional models: including adsorption activation from insoluble CCN, Atmospheric Chemistry and Physics, 9, 2517–2532, doi:10.5194/acp-9-2517-2009, http://www.atmos-chem-phys.net/9/2517/2009/, 2009.

30 Leaitch, W. R., Strapp, J. W., Isaac, G. A., and Hudson, J. G.: Cloud droplet nucleation and cloud scavenging of aerosol sulphate in polluted atmospheres, Tellus B, 38B, 328–344, doi:10.1111/j.1600-0889.1986.tb00258.x, http://www.tellusb.net/index.php/tellusb/article/view/15141, 1986.

35 Liu, X., Easter, R. C., Ghan, S. J., Zaveri, R., Rasch, P., Shi, X., Lamarque, J.-F., Gettelman, A., Morrison, H., Vitt, F., Conley, A., Park, S., Neale, R., Hannay, C., Ekman, a. M. L., Hess, P., Mahowald, N., Collins, W., Iacono, M. J., Bretherton, C. S., Flanner, M. G., and Mitchell, D.: Toward a minimal representation of aerosols in climate models: description and evaluation in the Community Atmosphere

Model CAM5, Geoscientific Model Development, 5, 709–739, doi:10.5194/gmd-5-709-2012, http://www.geosci-model-dev.net/5/709/ 2012/, 2012.

Lohmann, U., Feichter, J., Chuang, C. C., and Penner, J. E.: Prediction of the number of cloud droplets in the ECHAM GCM, Journal of Geophysical Research, 104, 9169–9198, http://onlinelibrary.wiley.com/doi/10.1029/1999JD900046/full, 1999.

5   Mårtensson, E. M., Nilsson, E. D., de Leeuw, G., Cohen, L. H., and Hansson, H.-C.: Laboratory simulations and parameterization of the primary marine aerosol production, Journal of Geophysical Research: Atmospheres, 108, n/a–n/a, doi:10.1029/2002JD002263, http://doi.wiley.com/10.1029/2002JD002263, 2003.

Ming, Y., Ramaswamy, V., Donner, L. J., and Phillips, V. T. J.: A New Parameterization of Cloud Droplet Activation Applicable to General Circulation Models, Journal of the Atmospheric Sciences, 63, 1348–1356, doi:10.1175/JAS3686.1, http://journals.ametsoc.org/doi/abs/

10   10.1175/JAS3686.1, 2006.

Morales, R. and Nenes, a.: Characteristic updrafts for computing distribution-averaged cloud droplet number and stratocumulus cloud properties, Journal of Geophysical Research, 115, D18 220, doi:10.1029/2009JD013233, http://doi.wiley.com/10.1029/2009JD013233, 2010.

Morales Betancourt, R. and Nenes, A.: Understanding the contributions of aerosol properties and parameterization discrepancies to droplet number variability in a global climate model, Atmospheric Chemistry and Physics, 14, 4809–4826, doi:10.5194/acp-14-4809-2014, http:

15   //www.atmos-chem-phys.net/14/4809/2014/, 2014a.

Morales Betancourt, R. and Nenes, A.: Droplet activation parameterization: the population-splitting concept revisited, Geoscientific Model Development, 7, 2345–2357, doi:10.5194/gmd-7-2345-2014, http://www.geosci-model-dev.net/7/2345/2014/, 2014b.

Morrison, H., Gettelman, A., Morrison, H., and Ghan, S. J.: A New Two-Moment Bulk Stratiform Cloud Microphysics Scheme in the Community Atmosphere Model, Version 3 (CAM3). Part I: Description and Numerical Tests, Journal of Climate, 21, 3642–

20   3659, doi:10.1175/2008JCLI2105.1, http://journals.ametsoc.org/doi/abs/10.1175/2008JCLI2116.1http://journals.ametsoc.org/doi/abs/10.1175/2008JCLI2105.1, 2008.

Neale, R. B., Gettelman, A., Park, S., Chen, C.-c., Lauritzen, P. H., Williamson, D. L., Conley, A. J., Kinnison, D., Marsh, D., Smith, A. K., Vitt, F., Garcia, R., Lamarque, J.-f., Mills, M., Tilmes, S., Morrison, H., Cameron-smith, P., Collins, W. D., Iacono, M. J., Easter, R. C., Liu, X., Ghan, S. J., Rasch, P. J., and Taylor, M. a.: Description of the NCAR Community Atmosphere Model (CAM 5.0). NCAR Technical

25   Notes., Ncar/Tn-464+Str, p. 214, 2012.

Nenes, A. and Seinfeld, J. H.: Parameterization of cloud droplet formation in global climate models, Journal of Geophysical Research, 108, 4415, doi:10.1029/2002JD002911, http://www.agu.org/pubs/crossref/2003/2002JD002911.shtml, 2003.

Nenes, A., Ghan, S., Abdul-Razzak, H., Chuang, P. Y., and Seinfeld, J. H.: Kinetic limitations on cloud droplet formation and impact on cloud albedo, Tellus Series B-Chemical and Physical Meteorology, 53, 133–149, doi:10.1034/j.1600-

30   0889.2001.d01-12.x, http://www.tellusb.net/index.php/tellusb/article/view/16569http://www.tellusb.net/index.php/tellusb/article/view/16569{%}5Cn{%}3CGotoISI{%}3E://WOS:000169810700003, 2001.

Park, S. and Bretherton, C. S.: The University of Washington shallow convection and moist turbulence schemes and their impact on climate simulations with the community atmosphere model, Journal of Climate, 22, 3449–3469, doi:10.1175/2008JCLI2557.1, 2009.

Partridge, D. G., Vrugt, J. a., Tunved, P., Ekman, a. M. L., Gorea, D., and Sorooshian, a.: Inverse modeling of cloud-aerosol interactions

35   – Part 1: Detailed response surface analysis, Atmospheric Chemistry and Physics, 11, 7269–7287, doi:10.5194/acp-11-7269-2011, http://www.atmos-chem-phys.net/11/7269/2011/, 2011.

Petters, M. D. and Kreidenweis, S. M.: A single parameter representation of hygroscopic growth and cloud condensation nucleus activity, Atmospheric Chemistry and Physics, 7, 1961–1971, doi:10.5194/acp-7-1961-2007, http://www.atmos-chem-phys.net/7/1961/2007/http://www.atmos-chem-phys.net/13/1081/2013/acp-13-1081-2013.htmlhttp://www.atmos-chem-phys.net/13/1081/2013/, 2007.

Rothenberg, D. and Wang, C.: Metamodeling of Droplet Activation for Global Climate Models, Journal of the Atmospheric Sciences, 73, 1255–1272, doi:10.1175/JAS-D-15-0223.1, http://journals.ametsoc.org/doi/abs/10.1175/JAS-D-15-0223.1http://journals.ametsoc.org/doi/10.1175/JAS-D-15-0223.1, 2016.

Saleeby, S. M. and Cotton, W. R.: A Large-Droplet Mode and Prognostic Number Concentration of Cloud Droplets in the Colorado State University Regional Atmospheric Modeling System (RAMS). Part I: Module Descriptions and Supercell Test Simulations, Journal of Applied Meteorology, 43, 182–195, doi:10.1175/1520-0450(2004)043<0182:ALMAPN>2.0.CO;2, http://journals.ametsoc.org/doi/abs/10.1175/1520-0450(2004)043{%}3C0182:ALMAPN{%}3E2.0.CO;2, 2004.

Scanza, R. A., Mahowald, N., Ghan, S., Zender, C. S., Kok, J. F., Liu, X., and Zhang, Y.: Modeling dust as component minerals in the Community Atmosphere Model: development of framework and impact on radiative forcing, Atmospheric Chemistry and Physics Discussions, 14, 17749–17816, doi:10.5194/acpd-14-17749-2014, http://www.atmos-chem-phys-discuss.net/14/17749/2014/, 2014.

Seinfeld, J. H. and Pandis, S. N.: Atmospheric Chemistry and Physics: From Air Pollution to Climate Change, vol. 2nd, Wiley, http://www.knovel.com/knovel2/Toc.jsp?BookID=2126, 2006.

Shipway, B. J. and Abel, S. J.: Analytical estimation of cloud droplet nucleation based on an underlying aerosol population, Atmospheric Research, 96, 344–355, doi:10.1016/j.atmosres.2009.10.005, http://linkinghub.elsevier.com/retrieve/pii/S0169809509002798, 2010.

Simpson, E., Connolly, P., and McFiggans, G.: An investigation into the performance of four cloud droplet activation parameterisations, Geoscientific Model Development, 7, 1535–1542, doi:10.5194/gmd-7-1535-2014, http://www.geosci-model-dev.net/7/1535/2014/, 2014.

Tatang, M., Pan, W., Prinn, R., and McRae, G.: An efficient method for parametric uncertainty analysis of numerical geophysical models, Journal of Geophysical Research, 102, 925–932, http://www.agu.org/pubs/crossref/1997/97JD01654.shtml, 1997.

Topping, D., Connolly, P., and McFiggans, G.: Cloud droplet number enhanced by co-condensation of organic vapours, Nature Geoscience, 6, 443–446, doi:10.1038/ngeo1809, http://www.nature.com/doifinder/10.1038/ngeo1809, 2013.

Twomey, S.: The nuclei of natural cloud formation part II: the supersaturation in natural clouds and the variation of droplet concentration, Pure and Applied Geophysics, http://www.springerlink.com/index/Q0N7L477832V7170.pdf, 1959.

Wang, M., Ghan, S., Easter, R., Ovchinnikov, M., Liu, X., Kassianov, E., Qian, Y., Gustafson Jr., W. I., Larson, V. E., Schanen, D. P., Khairoutdinov, M., and Morrison, H.: The multi-scale aerosol-climate model PNNL-MMF: model description and evaluation, Geoscientific Model Development, 4, 137–168, doi:10.5194/gmd-4-137-2011, http://www.geosci-model-dev.net/4/137/2011/, 2011.

Ward, D. S., Eidhammer, T., Cotton, W. R., and Kreidenweis, S. M.: The role of the particle size distribution in assessing aerosol composition effects on simulated droplet activation, Atmospheric Chemistry and Physics, 10, 5435–5447, doi:10.5194/acp-10-5435-2010, http://atmos-chem-phys-discuss.net/10/4189/2010/acpd-10-4189-2010.pdfhttp://www.atmos-chem-phys.net/10/5435/2010/, 2010.

West, R. E. L., Stier, P., Jones, A., Johnson, C. E., Mann, G. W., Bellouin, N., Partridge, D. G., and Kipling, Z.: The importance of vertical velocity variability for estimates of the indirect aerosol effects, Atmospheric Chemistry and Physics, 14, 6369–6393, doi:10.5194/acp-14-6369-2014, http://www.atmos-chem-phys.net/14/6369/2014/, 2014.

Wilson, J., Cuvelier, C., and Raes, F.: A modeling study of global mixed aerosol fields, Journal of Geophysical Research, 106, 34081–34108, doi:10.1029/2000JD000198, http://www.agu.org/pubs/crossref/2001/2000JD000198.shtml, 2001.
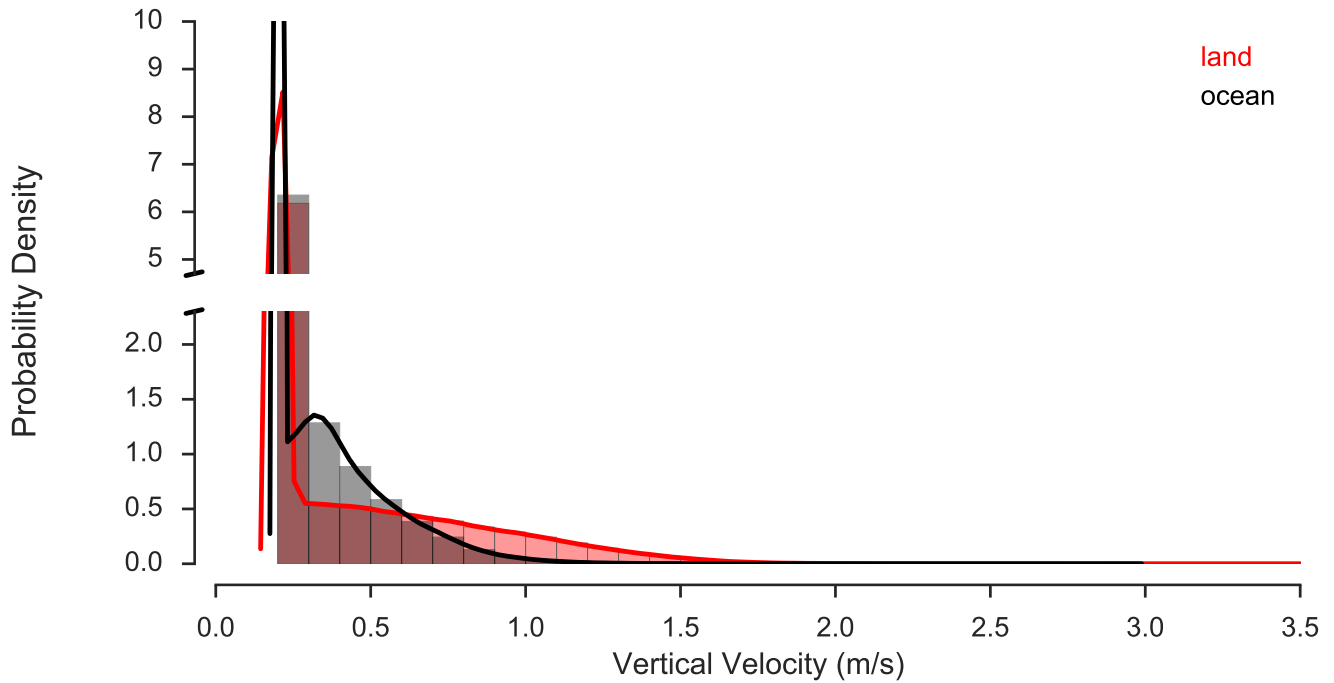
Figure 1. Distributions of model-predicted instantaneous sub-grid scale vertical velocity for near-surface (below 700 mb) grid-cells broken down by land (red) and ocean (black) regimes.

Yin, Y., Levin, Z., Reisin, T. G., and Tzivion, S.: The effects of giant cloud condensation nuclei on the development of precipitation in convective clouds — a numerical study, Atmospheric Research, 53, 91–116, doi:10.1016/S0169-8095(99)00046-0, http://linkinghub. elsevier.com/retrieve/pii/S0169809599000460, 2000.
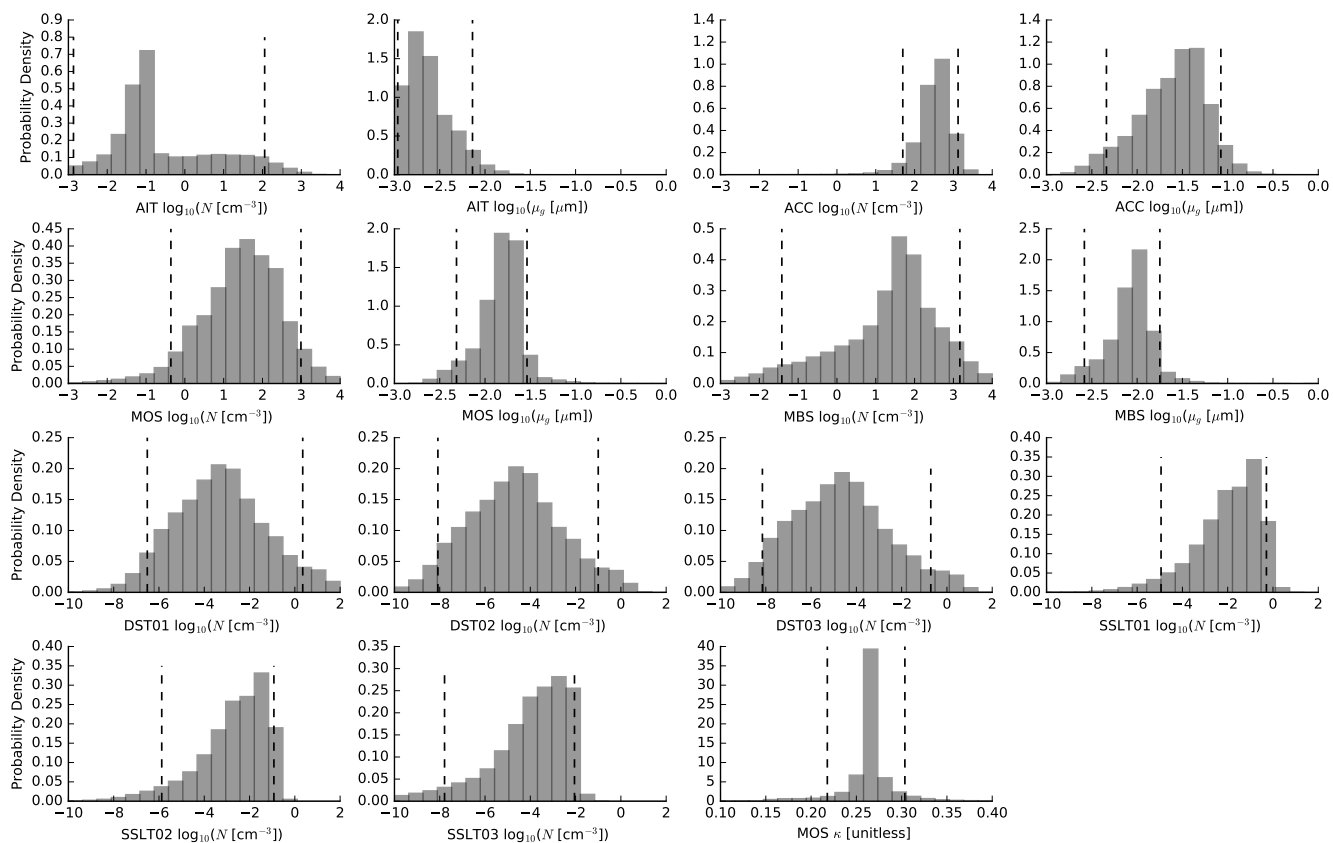
Figure 2. Distributions of aerosol size distribution parameters for key modes simulated by MARC. Vertical dashed lines indicate the 5th and 95th percentiles of the sampling distribution for each parameter. Note that each mode name corresponds to those in Tables 1 and 2.
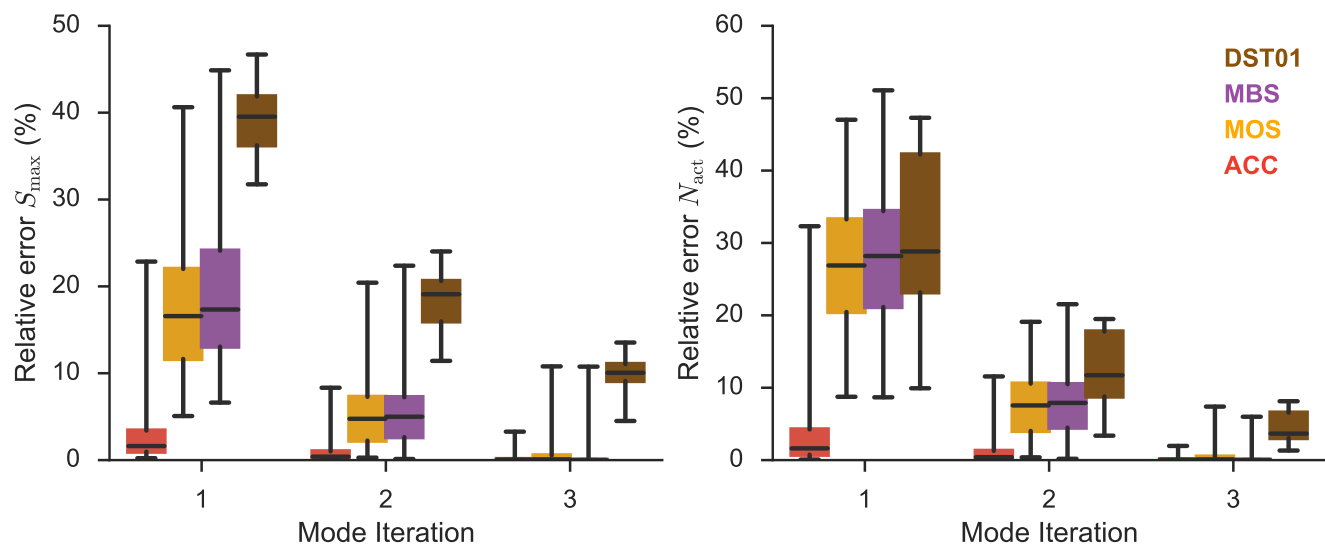
Figure 3. Relative errors in S_{max} (left) and N_{act} (right) in subsequent iterations of the iterative activation calculations. The coloring of each box indicates which mode was the first or dominant one. In each boxplot, box encompasses the interquartile range and the whiskers extend to the 1st and 99th percentiles in the corresponding sub-sample. Outliers beyond these percentiles are not plotted. Labels in legend correspond to accumulation mode (ACC), mixed organic-sulfate mode (MOS), mixed black-carbon sulfate mode (MBS) and smallest dust bin (DST01).
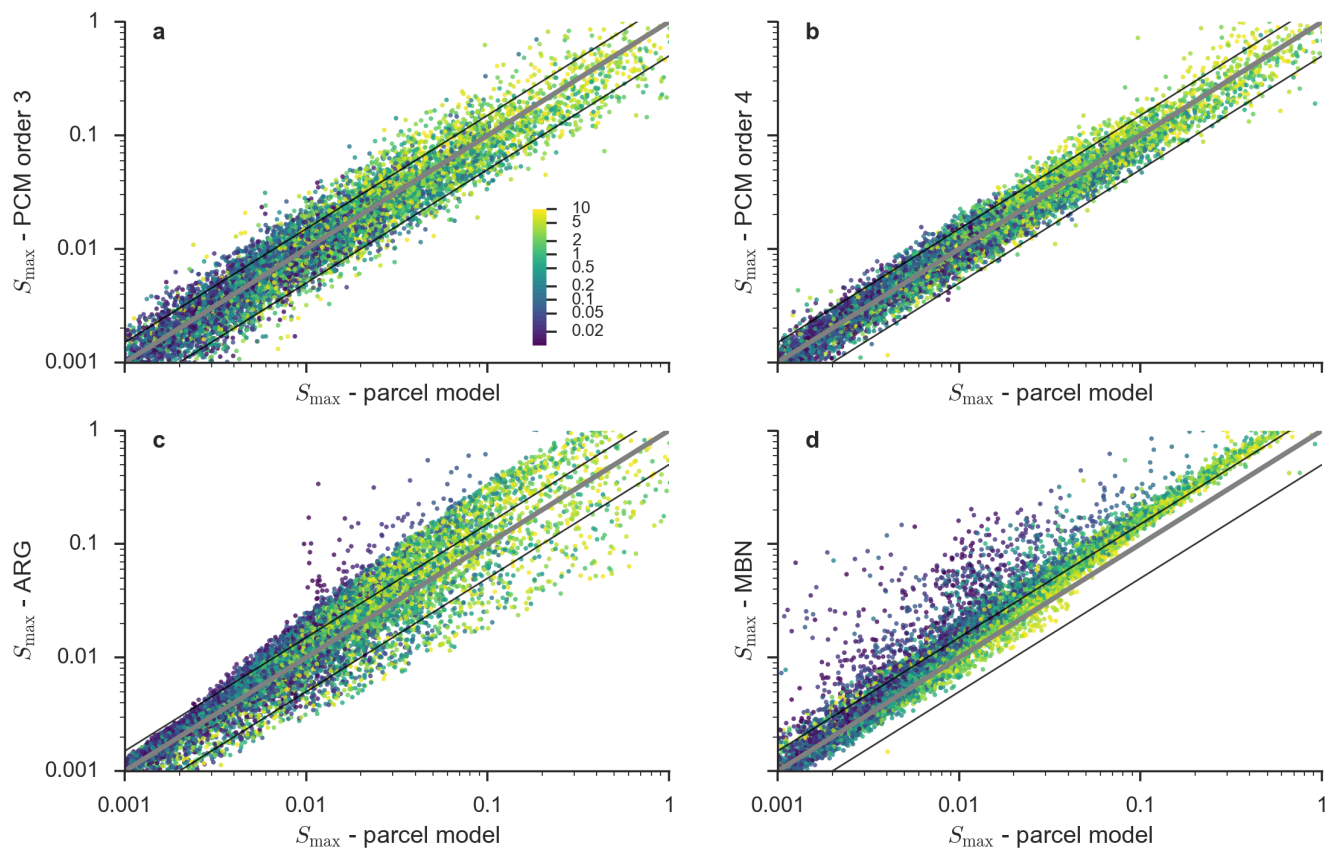
Figure 4. Predicted supersaturation maxima (%)) from parcel model and activation parameterizations - 3rd-order emulator (a), 4th-order emulator (b), ARG (c) and MBN (d). The "main" aerosol parameter set (excluding the dust and sea salt as predictor modes) were utilized here. Glyphs are shaded to denote updraft velocity corresponding to each sample draw (in m/s), and are consistent for each panel. Solid black lines denote a factor-of-2 difference between predicted values from parcel model and corresponding parameterization evaluations
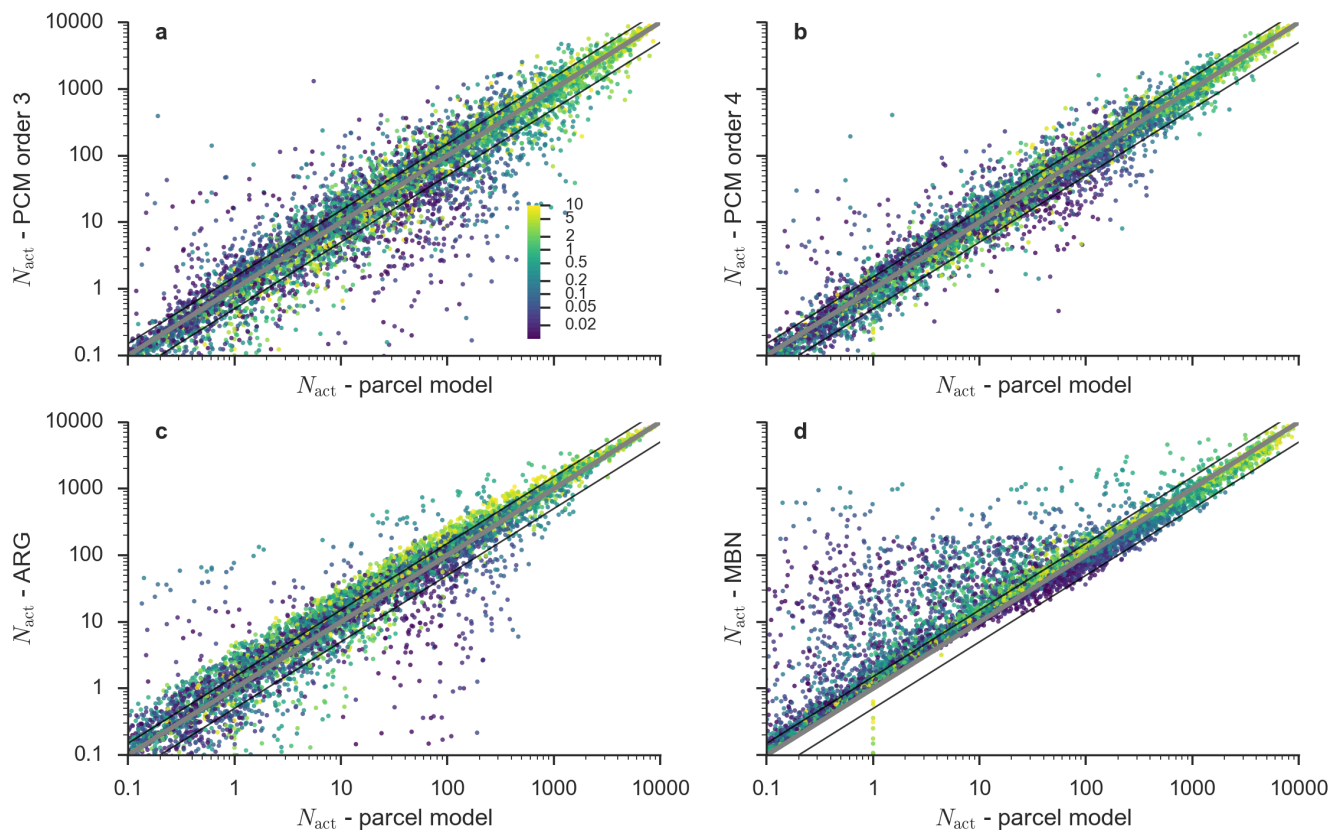
Figure 5. Like Figure 4, but plotting the predicted droplet number concentration (cm^{-3}) nucleated for the aerosol "main" parameter set
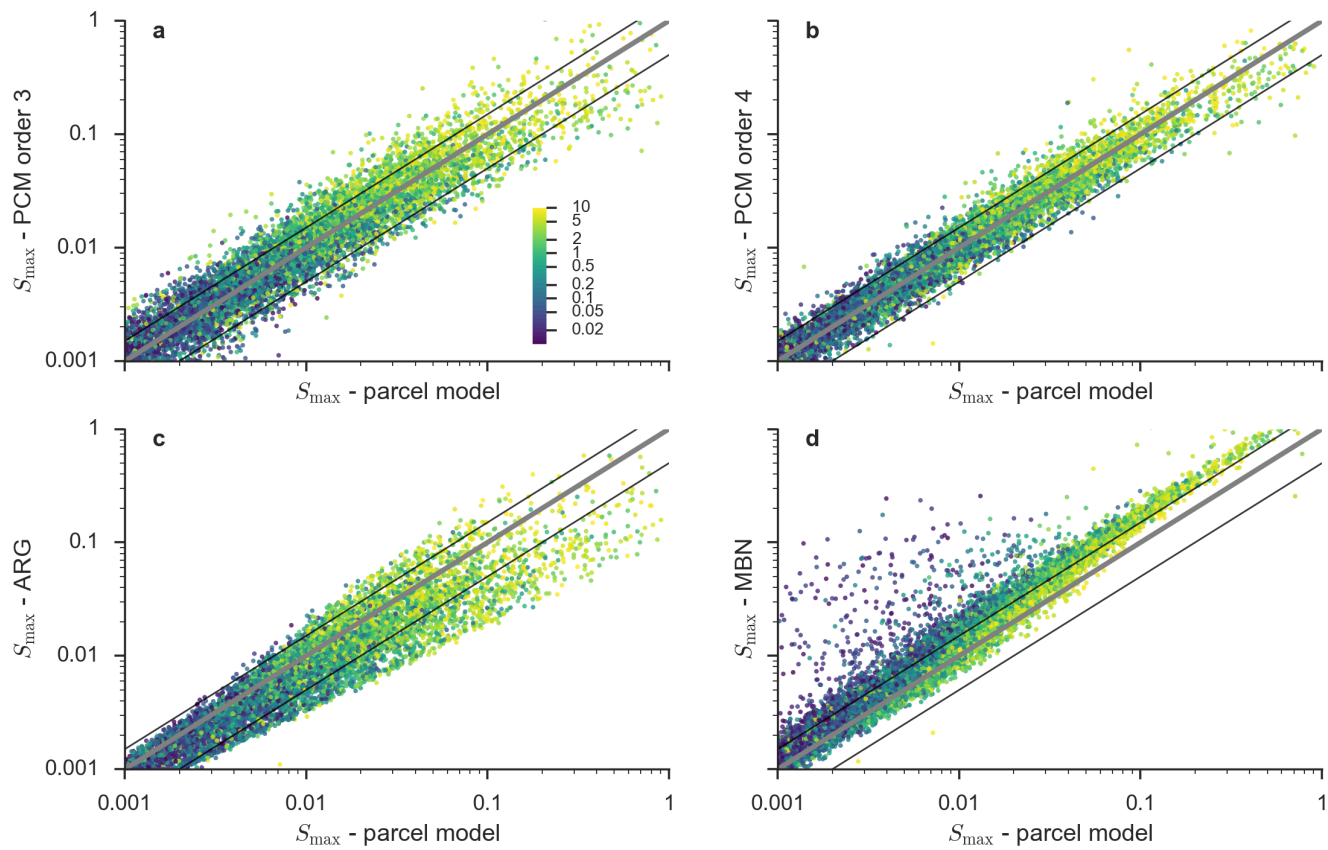
Figure 6. Like Figure 4, but for the "gCCN" parameter set
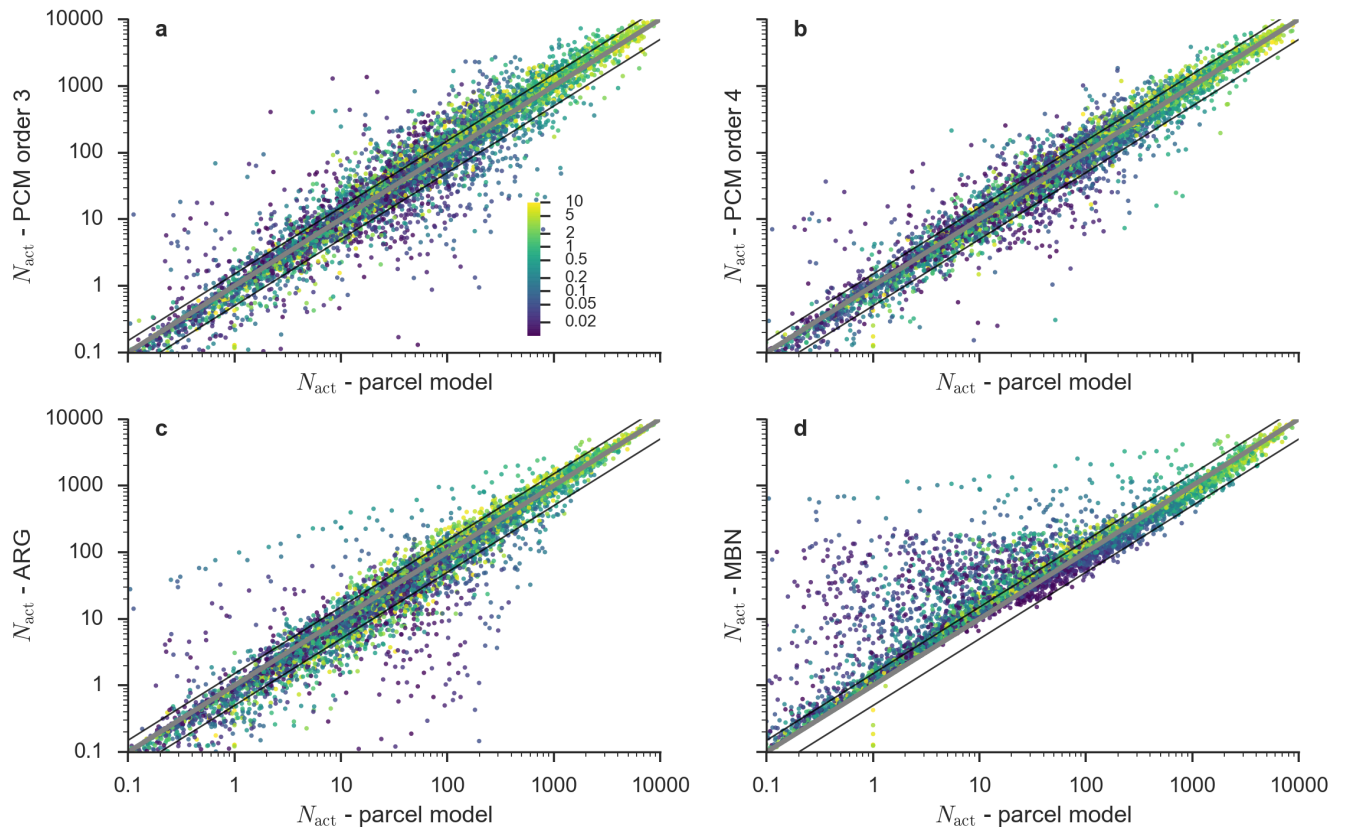
Figure 7. Like Figure 5, but for the "gCCN" parameter set
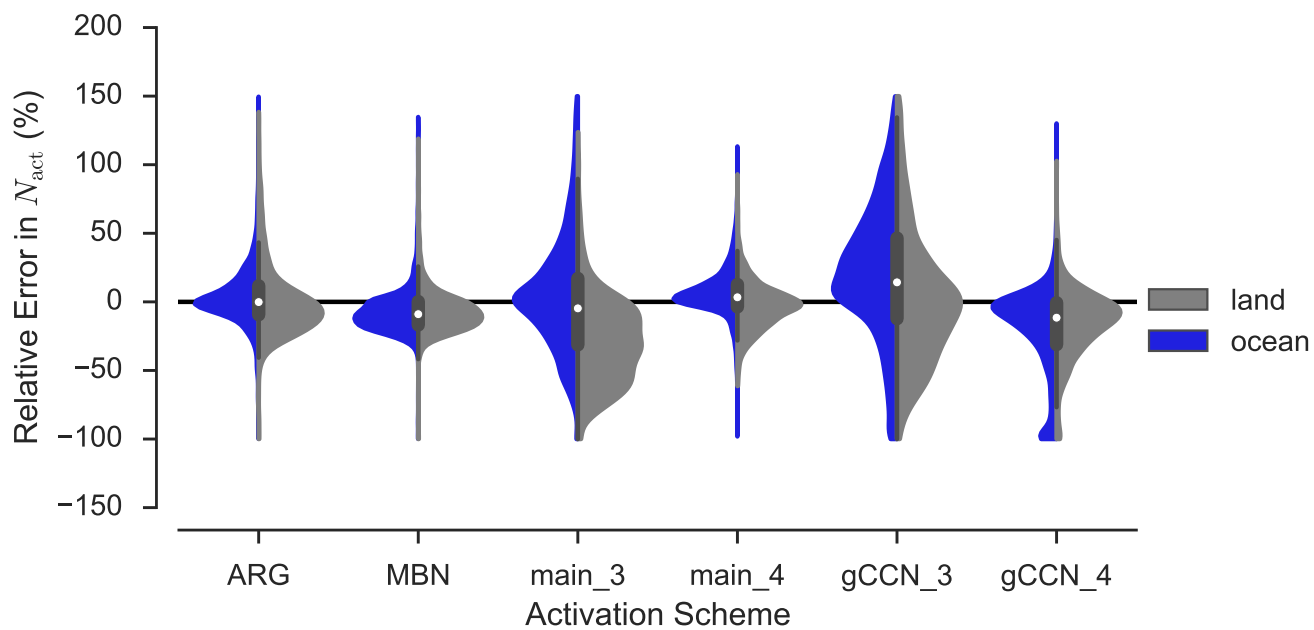
Figure 8. Distributions of relative error in scheme prediction of N_{act} versus detailed parcel, evaluated using samples taken from instantaneous MARC aerosol size distribution and meteorology and colored by geographical regime. The long tail of each distribution is clipped at the extrema for each scheme. The box plot in the center of each distribution shows the median and inter-quartile range of the total distribution of both land and ocean samples for each scheme.
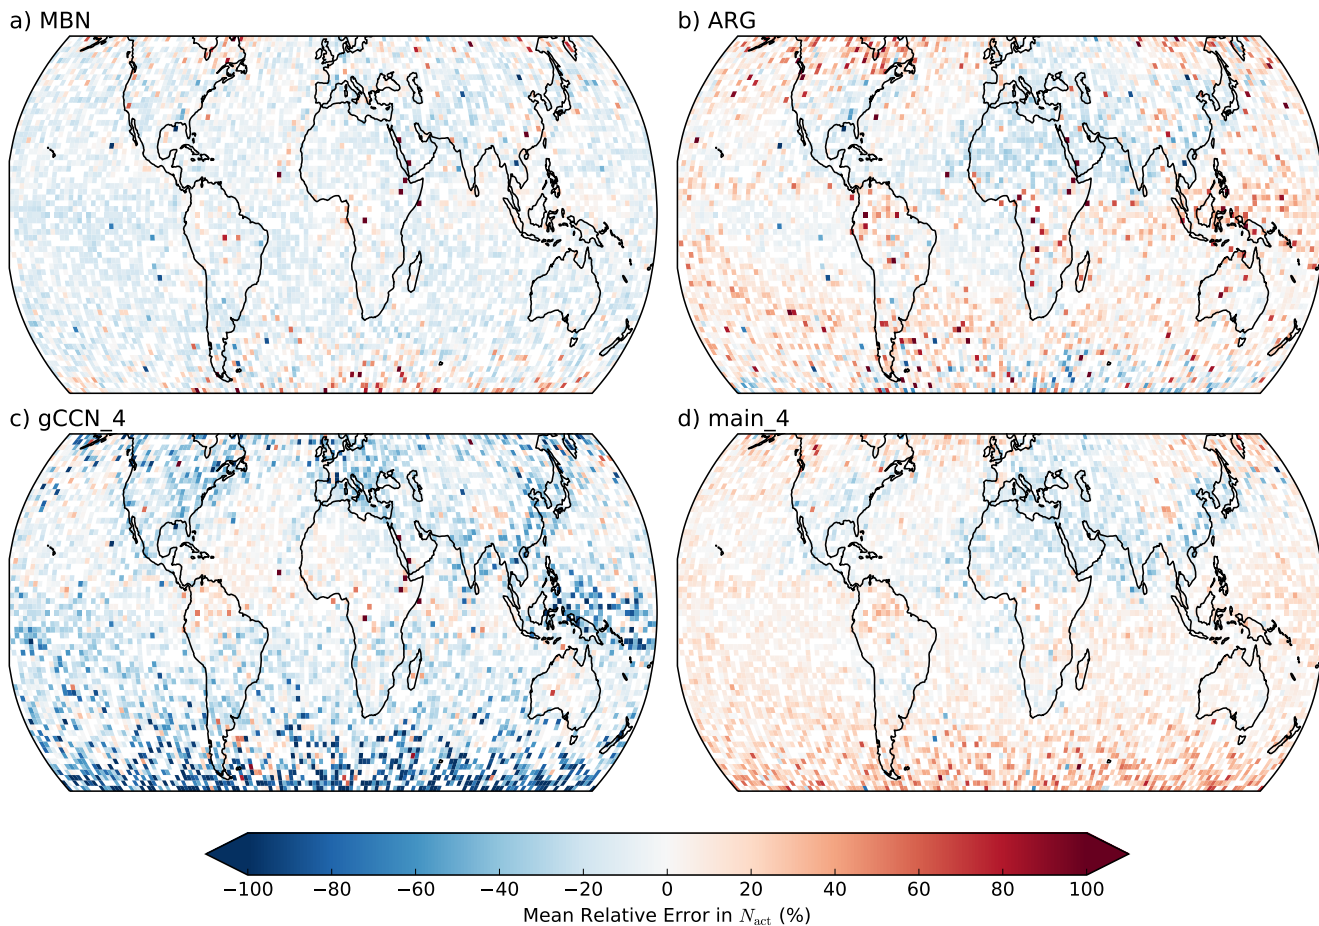
Figure 9. Mean relative error in scheme prediction of N_{act} versus detailed parcel model plotted against location on globe where those samples originated. At each grid location, all samples across timesteps and vertical levels (below 700 mb) are averaged together to compute the mean.

Aerosol Mode	Geometric Mean Particle Diameter (µm)	Geometric Std Deviation (σ_g)	Density ($\mathrm{g\,cm^{-3}}$)	Hygroscopicity (κ)
NUC	0 to 0.005 84	1.59	1.8	0.507
AIT	0.005 84 to 0.031	1.59	1.8	0.507
ACC	>0.031	1.59	1.8	0.507
OC	-	2.0	2.0	10^{-10}
MOS	-	2.0	†	†
BC	-	2.0	2.0	10^{-10}
MBS	-	2.0	2.0	0.507
DST01	0.16	1.4	-	0.14
DST02	0.406	1.4	-	0.14
DST03	0.867	1.4	-	0.14
DST04	1.656	1.4	-	0.14
SSLT01	0.5	1.59	-	1.16
SSLT02	2.0	1.37	-	1.16
SSLT03	5.0	1.41	-	1.16
SSLT04	15.0	1.22	-	1.16

Table 1. MARC aerosol mode size distribution and composition parameters. The MOS mode (†) has a composition-dependent density and hygroscopicity which is computed using the internal mixing state of organic carbon and sulfate present at a given grid-cell and timestep.

Symbol	Description	Lower Bound	Upper Bound
logN_ACC	Log of accumulation mode sulfate number concentration (cm^{-3})	-3 (1.2)	4 (100)
logN_MOS	Log of mixed sulfate-organic carbon number concentration (cm^{-3})	-5 (1.5)	4 (99.9)
logN_MBS	Log of mixed sulfate-black carbon number concentration (cm^{-3})	-5 (1.6)	4 (99.8)
logN_DST01*	Log of 0.16 micron dust particle number concentration (cm^{-3})	-5 (18.2)	2 (99.8)
logN_DST02*	Log of 0.4 micron dust particle number concentration (cm^{-3})	-5 (38.9)	1 (99.9)
logN_SSLT01*	Log of 0.5 micron sea salt particle number concentration (cm^{-3})	-5 (3.6)	1 (100)
logmu_ACC	Geometric mean size of accumulation mode (micron)	-3 (0.1)	0 (98.9)
logmu_MOS	Geometric mean size of mixed sulfate-organic carbon mode (micron)	-3 (0.06)	-1 (98.3)
logmu_MBS	Geometric mean size of mixed sulfate-black carbon mode (micron)	-3 (0.1)	-1 (98.5)
kappa_MOS	Hygroscopicity of mixed sulfate-organic carbon mode	0.1	0.6
log_V	Log of updraft velocity (m/s)	-2	1
T	Temperature (K)	240	310
P	Pressure (Pa)	50000	105000

Table 2. Input parameter space and bounds on associated uniform probability density functions used to derive polynomial chaos expansions for MARC activation. For the lower and upper bounds on the aerosol size distribution parameters, the parenthetical values denote the percentile of the distribution for that parameter at which the bound occurs. All terms are present for the main expansion; terms affixed with an (*) are added for the gCCN expansion.

		$\log_{10} S_{\max}$				N_{act}			
exp	scheme	MAE	MRE	NRMSE	r^2	MAE	MRE	NRMSE	r^2
main	ARG	0.18	-3.26	0.10	0.94	40.14	25.39	0.15	0.98
	MBN	0.20	-11.79	0.18	0.81	59.05	44.95	0.30	0.90
	PCM Order 3	0.16	0.59	0.09	0.95	72.54	9.20	0.31	0.90
	PCM Order 4	0.10	-0.60	0.06	0.98	45.47	8.89	0.19	0.96
gCCN	ARG	0.17	8.54	0.09	0.93	37.41	-3.92	0.15	0.98
	MBN	0.20	-9.58	0.17	0.78	56.03	33.30	0.31	0.89
	PCM Order 3	0.16	0.59	0.08	0.95	81.19	15.14	0.34	0.87
	PCM Order 4	0.10	0.36	0.06	0.98	50.99	6.90	0.23	0.94

Table 3. Summary statistics for error in supersaturation maxima and droplet number nucleated predicted by emulators and activation parameterization relative to corresponding simulations with a detailed parcel model. From left-to-right, each column represents the coefficient of determination (r^2), mean absolute error (MAE), mean relative error (MRE), and the normalized root-mean-square error (NRMSE)

	$\log_{10} S_{\max}$				N_{act}			
scheme	MAE	MRE	NRMSE	r^2	MAE	MRE	NRMSE	r^2
ARG	0.05	-0.16	0.03	0.92	25.5	2.87	0.16	0.94
MBN	0.06	0.05	0.05	0.71	26.7	-6.68	0.19	0.93
main Order 3	0.12	0.42	0.08	0.33	64.7	-1.81	0.44	0.59
main Order 4	0.04	-0.31	0.02	0.96	24	4.59	0.19	0.93
gCCN Order 3	0.14	-1.84	0.09	0.18	75.6	20.9	0.47	0.52
gCCN Order 4	0.12	4.61	0.10	-0.19	44.3	-19.9	0.33	0.76

Table 4. Same as Table 3, but for the sampling study using MARC aerosol and meteorology parameter sets.