

## ***Interactive comment on “Physically Regularized Machine Learning Emulators of Aerosol Activation” by Sam J. Silva et al.***

**Anonymous Referee #2**

Received and published: 15 January 2021

Silva et al. discuss the use of machine learning emulators to predict the activation of aerosol into cloud droplets. For this they test three machine learning algorithms – ridge regression, gradient boosted trees, and deep neural networks – and use them with and without an a priori estimate of aerosol activation based on two physical parameterizations of different complexity (denoted the Twomey scheme and the ARG scheme, respectively). They find that two of the tested machine learning emulators (gradient boosted trees and neural network) outperform the parameterizations even without imposing physical constraints, and all methods outperform the regular parameterization when combined with the physical parameterizations.

The paper is well written and the study is within the scope of the journal. The machine learning approach seems generally sound (see minor points below) and the authors

C1

provide a detailed description of the methodology. I thus recommend the manuscript for publication after addressing the following minor points:

- Table 1: Please provide a long name for each parameter, e.g., temperature, pressure, etc.
- Figure 2: Is there a reason why the performance of the Twomey scheme is not shown? Also, please state in the figure (or at least in the figure label) that these are the results for the physically naïve emulators.
- Section 4: it would be useful to show the machine learning statistics for both the training data and the test data to demonstrate that the models are not suffering from overfitting.
- On line 260, I think it should say: “tend to perform. . .”
- Section 4.4: The weak performance of the naïve and Twomey regularized emulator under low hygroscopicity regimes seems somewhat surprising. Doesn't this imply that the training data did not include enough training data capturing a low hygroscopicity environment? Based on Table 1, the hygroscopicity range of the training data spans 0 to 1.2, so the training data should capture this range at least to some extent. Or does this result suggest that the hygroscopicity value should be log-transformed to give higher weights to the lower bound? Given that low hygroscopicity values are not uncommon in the real atmosphere, this should be addressed a bit more convincingly in the revised version of the manuscript.
- Figure 7: is it possible to also show the performance of the regular parameterizations (without any machine learning)? This would help demonstrate the value of adding the machine learning correction to these parameterizations.

---

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2020-393>, 2020.

C2