This paper develops a geographically weighted gaussian process emulator for anthropogenic PM2.5 using global mean atmospheric CO2 concentration as a proxy for greenhouse gas driven climate change. Combined with a selection of anthropogenic emissions and surface CH4 concentrations, the emulator is able to reproduce a subset of GCHP simulated PM2.5 considered to be anthropogenic in nature. The gaussian process predictions are produced at a fraction of the computational cost, reducing inference time from thousands of CPU hours to seconds.

The choice of Gaussian Process regression is well justified, particularly given its ability to represent non linear chemistry and provide predictive uncertainty estimates, which enhances the practical utility of the framework.

Overall, reported performance metrics, including high R2 values and low mean absolute error, suggest that the emulator performs impressively within the tested domain. I believe the manuscript makes a valuable contribution and merits publication. However, I outline below what I consider to be a key limitation, and I suggest that additional clarification or running of some additional simulations would help with the scope of applicability demonstrated by the experiments presented.

Major comments:

Climate effects are parameterised exclusively through atmospheric CO_2 concentration, while methane is included solely as a chemical driver. This separates methane's chemical and radiative roles and implicitly assumes that CO2 adequately represents total radiative forcing across scenarios. Although the authors note that CO2 may misrepresent total radiative forcing when greenhouse gas trends decouple, no targeted experiments are presented to quantify the resulting systematic error. It would strengthen the manuscript to evaluate emulator performance under scenarios in which forcing composition diverges from the training data. Particularly in high CH4 pathways.

While the framework appears generalisable, the present emulator is trained on meteorology drawn from a single IGSM CAM trajectory. It would be helpful for the authors to clarify whether they view this emulator instance as portable across alternative climate trajectories, or whether retraining would be required under different forcing compositions or climate sensitivities. This is particularly important so as to give an idea as to how often the training simulations would need to be run.

Minor comments:

The author mentions that CO2 is a good proxy as it directly impacts SOA production, but this seems a bit out of place to the story the author is telling as the authors specifically omit SOA from their PM2.5. Also TRF effects on other emission areas such as wetlands and wildfires are ignored.

A sentence on how GCHPs wildfire emissions are being handled would be useful as the authors consider all BC and OC to be anthropogenic.

I would like to see where the 4 scenarios lie relative to the LHS samples. Are they meaningfully different from all the training simulations?