We thank the reviewer for their comments. Our responses are below in red and quotations from the manuscript are in blue.

On the whole the authors have addressed the reviewers comments reasonably thoroughly and the revisions made are appropriate. Focusing on the approach (parameterization development) rather than the product (a specific parameterization) addresses many of the concerns, although this makes it less clear what the novel aspects of the study are.

My concern about the application of the approach has only partly been addressed. The use of dependent variables such as the oxidized species MHP and H2O2 as inputs prevents fully independent use of the parameterization, and means that it can only be applied reliably in chemical conditions close to those in which it was trained. This is fine for some applications (small changes, or source region tagging) but inappropriate for many others (e.g., different time periods or environments such as those illustrated in sections 4.1 and 4.2). A clearer statement of this is needed in the concluding section (see specific comments below). A key weakness of the approach is that it is not robust: there is no error checking or warning when conditions drift outside the training set, and the user can therefore never be quite sure how reliable their results are.

We have added the following text to emphasize that parameterizations of OH, such as the one described here, have been used successfully for decades. The novelty of our work is in the methodology to generate the parameterization, which is significantly less laborious than previous methods.

Line 71 - 72:

> For over thirty years, parameterizations of OH have provided a viable alternative to climatologies in helping to understand OH/CO/CH_4 feedbacks.

Line 107 - 109:

> Machine learning algorithms are one potential method to quickly and accurately generate a new parameterization of OH, offering an advance over the methods used in Duncan et al. (2000) and Spivakovsky et al. (1990).

Line 132 - 33:

> This represents a significant advance over previous, much more laborious, methodologies to generate a parameterization of OH.

We also note that the purpose of the parameterization, once integrated into a CCM framework, is to allow for multiple sensitivity runs, and not to capture all feedbacks on all species, such as H2O2 or MHP. We say this at Line 95 - 98:

> Through manipulation of the input parameters (i.e., chemical, meteorological, and solar irradiance variables) to the parameterization of OH, as well as emissions and dynamics, ECCOH can produce multiple, computationally cheap, sensitivity simulations that help deconvolve the causes of local to global trends and variations in OH, CO, and CH_4.

And in a new concluding paragraph Line 712 - 724:

> The methodology we present here allows for the quick generation of a parameterization of OH for use in a chemistry climate model to help disentangle the complicated relationship between CO, CH_4, and OH. The parameterization is designed for computationally inexpensive sensitivity runs and, as such, is not designed to capture feedbacks between OH and all of its chemical and dynamical drivers (e.g. H_2O_2 or MHP). Instead, if a user is interested in these feedbacks, they could use the results of the sensitivity tests to identify times, locations, or chemical regimes for a targeted full chemistry simulation. Likewise, the parameterization reflects the photochemical environments of the dataset on which it was trained. Therefore, the training dataset should be carefully chosen to reflect the goals of a given study. However, we have demonstrated that the sample parameterization outlined here accurately predicts the magnitude and spatial distribution of the large deviations in OH for the 2016 El Niño, an event that was not part of the training dataset. This result gives confidence in the fidelity of a parameterization developed with our methodology to simulate the spatial and temporal responses of OH to perturbations from large variations in the chemical, dynamical, and solar irradiance drivers of OH.

Finally, in regards to there not being error checks or warnings, this is something that can and should be applied once the parameterization is integrated into the larger modeling framework, as is done with the current ECCOH parameterization. We have added the following text (Line 687 - 689):

> Once integrated into a modeling framework, safeguards could be added to warn a user if parameterization input values fall outside of the bounds of the training dataset, as is done with the current ECCOH parameterization.

Fig 3: Would be better to have the same color scale on the two panels

We have updated the figure so that the limits of the colorbars are the same for both panels.

Fig 5 caption: "Colors have no specific meaning..." This is not a useful statement, and it would be better to state "Colors are assigned to the variables to permit easier comparison of the panels" (or something similar).

We have updated the figure caption with the suggested text.

Line 626-629: This sentence on application of the approach is speculative and potentially misleading. An experiment exploring the effect of a OH decrease would quickly take CH4 outside the range of the training dataset. The phrase "would require significant care to ensure valid results" should be replaced by a more honest assessment that the approach should not be used outside the training set. It would be more useful to include a statement that the training set should be chosen carefully to encompass all likely conditions under which the subsequent parameterization will be used.

We have removed the sentence and replaced it with the following (Line 630 - 632):

> As will be discussed in the following section, care should be taken in choosing the training dataset to ensure that it represents the full range of photochemical conditions on which the parameterization will be applied.

Line 683: As above, the bigger message here is to ensure that the training data spans all possible conditions for which the parameterization is likely to be used. The approach should come with this as a major health warning, and its value to future users relies on them fully understanding this.

In addition to emphasizing this point in the new final paragraph (reproduced above), we have added the following text (Line 685 - 687):

> This highlights the need to compare the distribution of any parameterization inputs to that of the training dataset to ensure that the training dataset fully encompasses the range of photochemical environments necessary for a given study.

Suggestions for avoiding the frailties of the approach that are highlighted in section 4.1 and 4.2 would be valuable.

One solution is using a training dataset created from multiple simulations covering a wide range of emissions and time periods, as we currently outline in the paragraph starting on Line 201 and reproduced below. We do not employ that technique in this paper, however, because it is not necessary for our planned applications of the parameterization. Another possible way to address these issues is with a warning when values go beyond the training bounds. We have added this to the text (Line 687 - 689) and have reproduced the text above. Finally, we note again that this is a tool for sensitivity simulations and not a chemical mechanism replacement, which would likely require different machine learning techniques.

Line 201 - 214:

> While we have used the publicly-available MERRA2 GMI dataset to train the sample parameterization described in this manuscript, the training data could come from any simulation or combination of self-consistent simulations that has output of the variables outlined in Table 1. These training datasets could come from existing simulations, which would greatly reduce computational expense, or from a training dataset tailored for the purposes of a given study. Even though we use daily-averaged training data for ECCOH, a user could train the parameterization with a dataset at any temporal resolution in order to make the parameterization compatible with a specific modeling platform or research goal. As discussed later, the parameterization performs best when applied to photochemical environments analogous to those on which it was trained. Therefore, users should carefully ensure that the training dataset reasonably encompasses the full range of photochemical environments necessary for a given sensitivity test or experiment. For example, as we will discuss further in Section 4, because the MERRA2 GMI training dataset only covers 1980 – 2018, it is inappropriate to use this for an application exploring changes in CH_4 from the pre-industrial period to 2100. Instead, a new training dataset covering that time period would be required.