

We thank both reviewers for their careful reading of the manuscript and their insightful comments, which have encouraged us to make revisions that have significantly strengthened the analysis. We provide a detailed response to their comments below, with the original comments from the reviewers in **black**, our responses in **blue**, and changes to the manuscript in **green**. Line numbers correspond to the revised manuscript. Based on the reviewers' comments, we have made the following major changes to our analysis and the manuscript:

- 1) We updated the discussion and analysis of the latent space representations to provide more connection between the learned latent variables and typical signal analysis methods.
- 2) We have added a comparison between the VAE analysis and a PCA analysis to demonstrate the potential value added by the non-linear dimensionality reduction method
- 3) We have replaced the section on outlier detection with a section on using cosine similarity as a metric to compare the within-class and between-class similarity of L-II signals.

We have also updated the github repository so that our analysis can now be easily reproduced using Google Colab.

\*\*\*\*\*

#### **Reviewer 1:**

The paper by Doshi and Lamb introduces an unsupervised machine learning approach to better understand the structure of absorbing aerosols using L-II signals from the SP2. Using a variational autoencoder (VAE) the authors are able to extract a compressed latent feature vector of the L-II signals, and use this for outlier detection and enhanced identification of distinct aerosol populations (even outperforming previous tests using significant feature engineering). The paper is generally well written and I appreciate the conciseness of everything. Before fully recommending the paper for publication, I have a handful of questions/comments I'd like to see addressed surrounding latent feature physical interpretations, the dimensionality reduction methodology, outlier detection approach, and generalizability. Further, several of the figures should be updated to match the specifications set forth by EGU (i.e., enhanced text/label size throughout and improved color choices for visibility) to improve general readability.

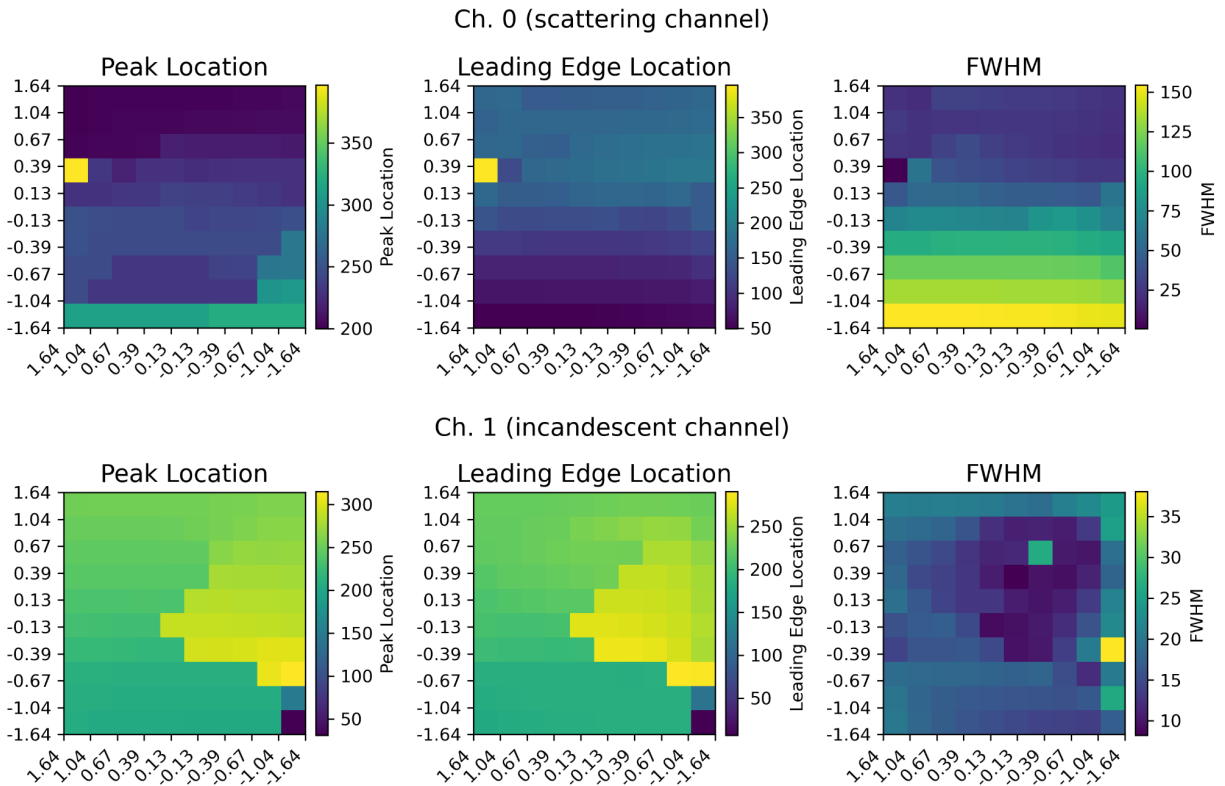
We thank the reviewer for their careful reading of the manuscript and for providing very useful critical feedback that has helped us to improve our analysis and the robustness of our results.

#### **General Comments:**

1. **Physical Interpretations:** At the core of this research project is the latent feature vector [Z1, Z2] produced by the VAE, however it wasn't clear to me throughout the analysis what these values represent physically? I.e., what does this compressed representation mean wrt. aerosols? This is always a challenge in DR projects, but I think it is important to consider since this vector is used throughout the remainder of the results for visualization/interpreting different classes/outlier etc., and we would like to have some confidence that it is learning a real physical feature in the L-II signals, and not some

spurious measurement artifact from the SP2 for instance. I understand Section 3 talks a bit towards this point, but I'd like to see more detail from the authors examining this in more detail, and if possible, provide some validation of their interpretation to ancillary data (if available).

We have now expanded Section 3 to further discuss the physical interpretability of the latent space dimension. In particular, we added an additional figure (Figure 7) to the manuscript analysing the latent manifold signals to indicate what types of variance are encoded in the latent space variables.



**Figure 7: Signal Analysis of Latent Manifolds for the L-II Signals.** The main peak location, peak leading edge location (defined as the leading edge of the largest peak at half maximum), and full-width half maximum for the decoded L-II signals for Ch. 0 (top row) and Ch. 1 (bottom row). The peak location, leading edge location, and FWHM width are all given in terms of digitalized time points (out of 400 for the NOAA SP2).

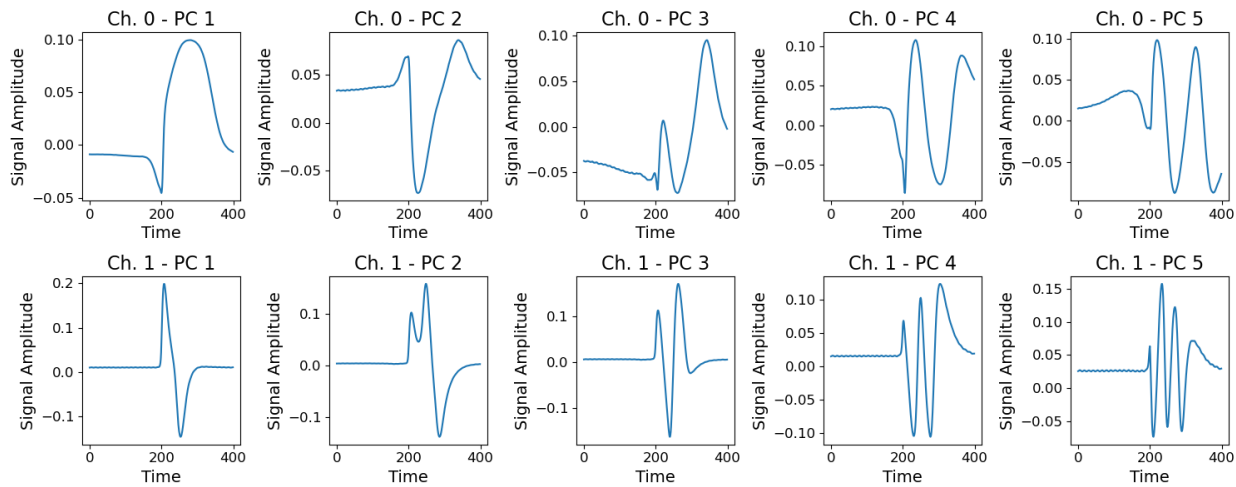
We have added the following discussion to the manuscript in lines 293 - 301:

To more quantitatively assess the types of variability learned by the VAEs for Ch. 0 and Ch. 1, we analyze the decoded signals associated with each decile, using the scipy package (Virtanen et al. 2022). We use the find\_peaks function in scipy to find the largest peak in each time series, and

then determine the peak location (in terms of digitized time points). We also use the `peak_widths` function to determine the width of the full width half maximum (FWHM) of this peak, as well as the digitized time point of the leading edge of the peak at half maximum (which we refer to as leading edge location). The values associated with the decoded signals of the latent manifolds for both Ch. 0 and Ch. 1 are shown in Figure 7. For Ch. 0 in particular it is clear that both the leading edge location and the FWHM are strongly correlated with  $z_2$ . For Ch. 1, the peak location and leading edge location have their highest values (associated with later times in the L-II signals) towards the center of the manifold. The signals at the center of the manifold are also associated with more narrow peaks in terms of their FWHM.

2. **Dimensionality Reduction Methods:** Building on my previous point, VAEs are great at capturing nonlinear relationships and for providing a smooth manifold, but the latent embeddings are typically less interpretable than the EOFs presented by a PCA for instance. While I understand this problem likely benefits from nonlinear DR like VAEs, UMAP etc., did you attempt to fit a PCA to the same data? What EOFs were produced and how did they differ from those using the VAE? This should be easy to fit, and I feel would be a useful/interesting comparison for this project to include as it provides a linear baseline for comparison, motivating the need for a more computationally expensive technique like the VAE.

We have now included the discussion of the PCA analysis of the SP2 signals, and included a Jupyter notebook of this analysis in our github repository. The characteristic EOFs for the first 5 principal components for Ch.0 and Ch.1 are shown below (Figure 3):



**Figure 3. First 5 Principal Components of the L-II signals.** Top row: Ch. 0 (Scattering channel). Bottom row: Ch. 1 (Incandescent channel)

We have expanded the discussion throughout the paper to include the comparison between the VAE and a PCA analysis. In particular, we have now added additional figures of the aerosols in terms of the first two PCs of Ch. 0 and Ch. 1 (Figure 4). We also compare PCA analysis with the

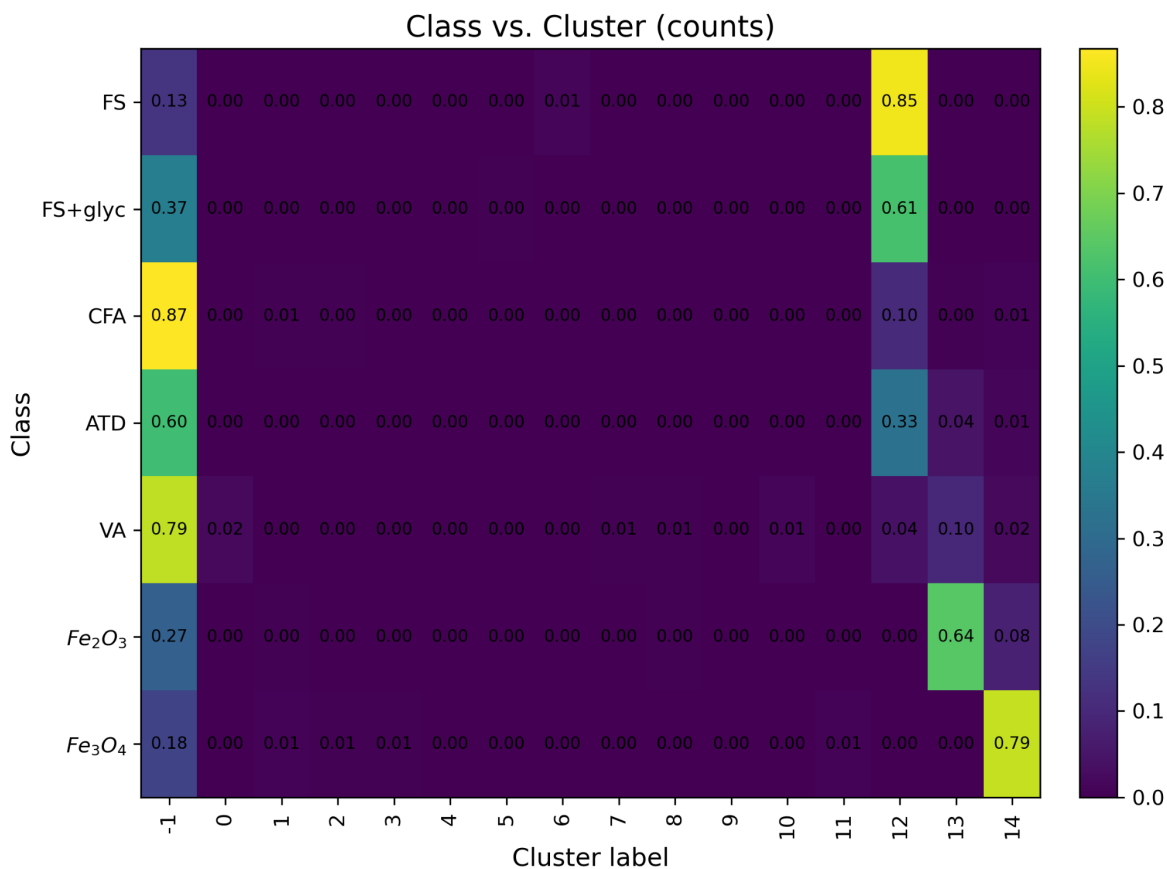
2D and 3D VAE in Figures 10 (the cosine similarity metric between classes) and Figures 11 (the random forest prediction of aerosol class).

3. **Outlier Detection:** I like this idea for outlier detection (Section 4) using the latent manifolds from the VAE. However, I am not totally convinced that the centroid + euclidean distance approach is the most robust approach for this. As the authors illustrate in Figs. 5/6, the manifolds from the latent vector don't produce circular/spherical distributions across the embedding plane. When dealing with continuous real-world observations, we instead typically see these much more abstract and diverse shapes, where a simple euclidean distance from the centroid might not capture a true outlier, but instead, another relevant mode in the data. I would recommend the authors consider including a discussion/comparison to other approaches for outlier detection (e.g., HDBSCAN) which looks at the density of the cases and how they are distributed across the scene.

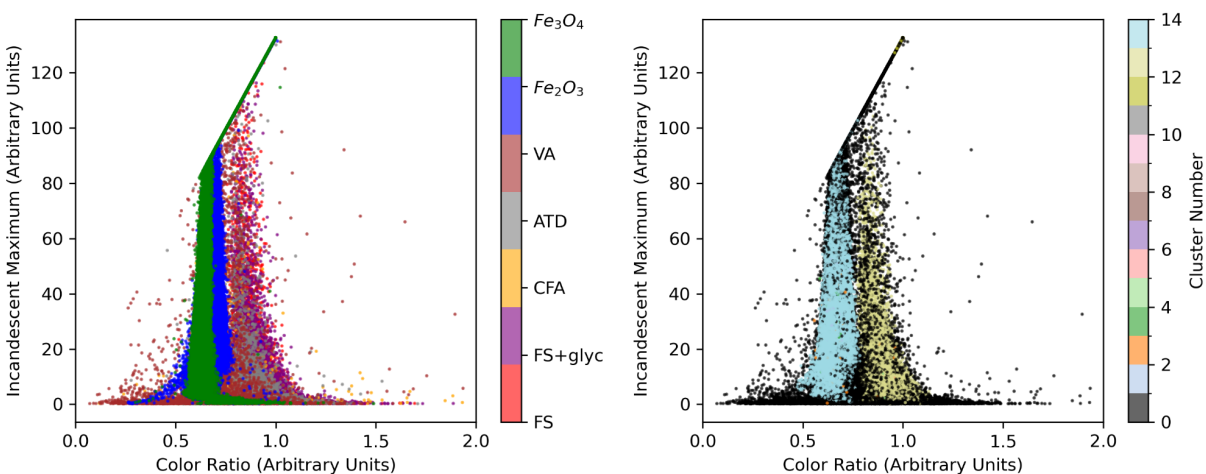
We explored using HBSCAN, and found it promising but not conclusive in terms of unsupervised identification of various aerosol classes.

When applying HDBSCAN to the 2D latent features from both the scattering and incandescent channels as well as the color temperature ratio, the incandescent maximum, and the scattering maximum, with a minimum cluster size of 50, we did find promise in unsupervised identification of iron oxide aerosols from fullerene soot. However, we did not find total separability in terms of identifying the individual classes, as shown in Figure R1. However, HDBSCAN did do a good job of identifying aerosols that seemed to be outliers in color temperature vs. incandescent maximum space (Figure R2), suggesting that this is a promising approach that should be explored in more detail in future research.

We agree that the Euclidean distance approach has its limitations in terms of identifying outliers from the latent space manifolds, given their complex structures. We have therefore chosen to deemphasize the outlier detection analysis (and instead discuss this as a promising future research direction). We instead replaced this section with a new analysis of the cosine similarity of the embeddings from the various dimensionality reduction techniques (PCA, 2D and 3D VAE) within and between aerosol classes, which provides a quantitative metric for how similar (or different) aerosols within and between each class are in terms of their response in the SP2.



**Figure R1:** Class vs. cluster identification by HDBSCAN.



**Figure R2:** Incandescent maximum vs. color ratio space for the L-II signals, colored by aerosol class (left) and HDBSCAN cluster label (right). Black points in the HDBSCAN cluster labeled plot indicate L-II signals that were identified as outliers (i.e. not belonging to any of the HDBSCAN clusters).

4. **Generalizability:** Since this work is based on laboratory data, do you expect the results of the VAE manifolds to change if applied to actual atmospheric observations? How robust is this to different regional climates/periods, and how stable would you expect the VAE manifold to remain? It would be nice to see some additional details or some text in the discussion addressing the applicability of this approach and these results to new, unseen data.

The VAE will depend on the data that it is trained on, as well as the training process itself, since the initial weights of the neural networks are randomly selected. However, once trained, the VAE can be applied in its inference mode to any new data set. Thus the VAE could be trained on laboratory samples and then could be applied to ambient measurements, allowing a comparison between the latent representations of ambient aerosols and laboratory proxies.

We have now added this paragraph to Section 2.3 in lines 225-231:

Since the specific latent representations are learned from the distribution of data that the VAE is trained on, the VAE in general will not learn the same latent representation from a new data set. Training on a representative sample of laboratory measurements (as we do here) can provide a wide range of possible samples for the generative model. However, once trained, the VAE can be applied in its inference mode to any new data set, providing the same mapping between input signals and lower dimensional representations. For example, the VAE could be trained on laboratory samples and then can be applied to ambient measurements, allowing a comparison between the latent representations of ambient aerosols and laboratory proxies.

#### **Specific Comments:**

1. **Title:** I would remove “(VAE)” from the title

We have updated the title to “Unsupervised Classification of Absorbing Aerosols Detected by the Single Particle Soot Photometer”, as we have now expanded the analysis to include PCA.

2. **Figure 1:** Can you add some additional details to the figure description explaining the SP2 (outside of just the citation).

We have expanded the description in the caption to read:

Aerosols sampled by the SP2 are introduced into the cavity of an ND:YAG laser (the beam is indicated as red curved lines), where they scatter light and potentially incandesce in the laser beam. Four photomultiplier tubes (indicated in the figure as the red detector, blue detector, and scattering light detectors) acquire time series of the light scattered or emitted by the aerosol particles as they interact with the laser beam.

3. **Line 36:** Should be “incandesce”, right?

Yes, corrected.

4. **Lines 48-49:** This felt like a really abrupt change in the introduction, moving from the SP2 to unsupervised machine learning. I would suggest reworking these paragraphs to flow more clearly/logically.

Great point. We have added an additional paragraph to discuss the typical usage of the SP2 in measuring aerosols during field campaigns, and the large amount of data that the SP2 acquires, which requires automated data processing methods in lines 50-56:

The SP2 is typically used in both ground-based and airborne field campaigns to monitor ambient aerosol particles (e.g., Schwarz et al. 2006, Schwarz et al. 2010, Lamb et al. 2018, Lamb et al. 2021, Katich et al. 2023). During a typical research flight in a source region, the SP2 might acquire signals for ~4 million individual aerosol particles, necessitating automated detection techniques to process L-II signals from the instrument. Typical automated methods to process L-II signals use curve-fitting techniques and calibrations from laboratory studies to determine aerosol particle mass and coating thickness assuming a Mie-theory core shell framework (e.g. Gao et al. 2007, Schwarz et al. 2006, Schwarz et al 2010). As an alternative approach to traditional SP2 data processing techniques, Lamb (2019) alternatively explored a supervised machine learning approach.

5. **Lines 49-63:** I recognize this is a follow-on from Lamb (2019), but are there other references you could include throughout this section looking at supervised/unsupervised ML in aerosols to better motivate this work? The paper would likely benefit from broadening the current set of references to better situate itself in current literature.

Yes, this is a good point. We have now included a paragraph in the introduction on the recent use of machine learning for atmospheric observations from other types of in situ aerosol and cloud particle instrumentation besides the SP2 in lines 57-66:

Machine learning is increasingly being applied to other types of specialized instrumentation used to detect in situ aerosol and cloud particles, providing new insights into the properties of atmospheric aerosol and cloud particle populations in large data sets. While initial efforts have focused on supervised machine learning (for example, for the Wideband Integrated Bioaerosol Sensor (WIBS) (Rusket et al. 2017; 2018), the Particle Analysis by Laser Mass Spectrometry (PALMS) instrument (Christopoulos et al. 2018; Zawadowicz et al. 2017), and for cloud particle imager probe images (e.g. Przybylo et al. 2022), more recent efforts have also explored unsupervised and self-supervised machine learning approaches, sometimes in combination with co-sampled environmental observations (e.g. Allwayin et al. 2024; Ko et al. 2025). As large datasets of atmospheric observations from in situ observations are increasingly being compiled across field campaigns and sampling conditions, these methods promise to lead to new physical insights into the atmospheric processes that underlie these observations (Lamb et al. 2025).

Allwayin, N., Larsen, M.L., Glienke, S. and Shaw, R.A., 2024. Locally narrow droplet size distributions are ubiquitous in stratocumulus clouds. *Science*, 384(6695), pp.528-532.



Ko, J., Govindarajan, H., Lindsten, F., Przybylo, V., Sulia, K., van Lier-Walqui, M. and Lamb, K., 2025. Understanding Ice Crystal Habit Diversity with Self-Supervised Learning. *Tackling Climate Change with AI Workshop, Conference on Neural Information Processing Systems, 2025*. *arXiv preprint arXiv:2509.07688*.

Lamb, K.D. and Singer, C. and Loftus, K. and Morrison, H. and Powell, M. and Ko, J. and Buch, J. and Hu, A. and van Lier Walqui, M. and Gentine, P. Perspectives on Systematic Cloud Microphysics Scheme Development with Machine Learning. *Journal of Advances in Modeling Earth Systems* (under review), 2025.

Przybylo, V.M., Sulia, K.J., Schmitt, C.G. and Lebo, Z.J., 2022. Classification of Cloud Particle Imagery from Aircraft Platforms Using.

6. **Lines 57-59:** You write “unsupervised machine learning” three times in three lines and it reads as quite repetitive, I would also recommend reworking this portion of the introduction to flow more cleanly.

We have reworked this paragraph (lines 75-84) to connect more clearly to the previous paragraph on limitations of supervised methods:

On the other hand, unsupervised machine learning algorithms discern inherent patterns and correlations in data sets, without requiring predefined categories. These methods have not previously been applied to the problem of classifying L-II signals, and they may address the limited generalization performance of supervised methods to ambient aerosol populations where laboratory proxies are not available or only partially replicate the characteristics of ambient aerosols. Here we explore how unsupervised machine learning can be applied to L-II signals from the SP2, with the goal of identifying different populations of aerosols based on the information in their L-II signals alone. We also investigate whether unsupervised methods can provide insights into the variability of aerosols of different types based on their L-II response in the SP2. Our analysis provides insight into the amount of independent information that can be gained from L-II signals in terms of identifying the composition of refractory aerosols that reach detectable incandescence in the SP2, and also provides a starting point for future studies on how data-driven methods might provide new insights into the sources and variability of aerosol populations.

7. **Line 67:** I know you define VAE in the abstract, but I typically recommend redefining all abbreviations within the text itself, so it can stand on its own.

Updated.

8. **Lines 153-155:** I am less familiar with the experimental setup for aerosol-specific model training, but are there any concerns about overfitting from temporal autocorrection using a random 50/25/25 split over a segmented approach?

The individual aerosol detection events (what we are referring to here as the L-II signals) in the SP2 are determined by when the voltage on the scattering detector passes a specific threshold



value due to an aerosol particle passing through the center of the laser beam. Based on this trigger, a predetermined time series length of 80 microseconds (400 time point signal) is saved as the signal associated with an individual aerosol particle. This pre-processing is done before the machine learning algorithms are applied, and thus we expect each L-II signal to represent an independent triggering event. Therefore we do not expect any temporal auto-correlation will play a role in the analysis of these data sets, as each L-II signal represents an individual aerosol detection event, and we treat each signal as an independent sample.

To clarify this in the text, we have added this sentence to the Dataset description in lines 97-99:

Each L-II signal represents an individual aerosol detection event in the SP2, and we treat each signal as an independent sample in applying the machine learning algorithms.

9. **Figure 2:** I like this figure, but the text is really small/hard to read in its current state

We have updated the text size so that it is easier to read.

10. **Line 232:** You mention hyperparameter selection here, but don't go into details about the process. What approach was used/what search space was evaluated? I summary table with final tuned values would also be a useful addition to consider.

We were referring to varying the size of the latent space from 2 - 5 variables, which we have discussed in detail in the manuscript. We expanded the discussion here to now include an estimation of the intrinsic dimension of the time series for Ch. 0 and Ch. 1, which further solidifies our conclusion that 3 variables is likely the maximum number needed for the VAE to capture the variance in our data set in lines 342-346:

We additionally estimated the intrinsic dimension of the latent variables associated with Ch. 0 and Ch. 1 for  $n=2,3,4$ , and 5 latent variables. The intrinsic dimension is the minimum number of variables required to represent the variability in the dataset. We used the intrinsic dimension algorithms in sci-kit dimension (Bac et al. 2021), and found that independent of the intrinsic dimension algorithm used, the estimated intrinsic dimension was approximately 3 for the time series associated with both channels.

We did not extensively explore the hyperparameters in the VAE or the random forest because our focus here was on demonstrating the potential value of an unsupervised learning approach for identifying useful features in the SP2 signals that might help to differentiate between different types of aerosols, rather than optimizing performance on this specific data set.

11. **Figures 5/6:** Similar to my previous comment about label size. Also the yellow cases are really hard to see on the white background. It might also be useful to apply a set of coloblind-friendly colors here so the "All classes" figure is more easily digested.

We have adjusted the font size and updated the colors to be more visible.

12. **Figure 7:** Font size comment again

We have updated this section to focus on the cosine similarity calculation of the embeddings in each class, so we have removed this figure from the updated version of the paper.

13. Line 361-362: I appreciate the inclusion of the GitHub link with code for reproducibility, but when examining it, there are no included datasets to test with? Unless I am missing this step, I would recommend taking a subset of your \*.npy data depending on data size and making it available on Colab through gdown (or some similar alternative). Then users can simply run your notebook online to reproduce the primary results.

As mentioned in the Code and Data availability section, the data sets used in this study are available in a Zenodo repository as numpy arrays: <https://doi.org/10.5281/zenodo.15800436>.

We have now refactored the code and set up the github so that the Jupyter notebooks can be opened directly in Colab and the datasets can be immediately downloaded to Colab from Zenodo in order to facilitate reproducibility.