I recommend three minor changes for the revised manuscript

[Reviewer]: In regards to FORLI: I still contend this FORLI does retrieve an HNO3 profile, so it is wrong to call it a "a stratosphere+troposphere integrated column product". Assuming the FORLI documentation is correct (see https://acsaf.org/docs/pum/Product_User_Manual_IASI_HNO3_Apr_2022.pdf) then the data product files contain *only HNO3 profiles*. Users have to compute the integrated quantity themselves.
It is perfectly fine to argue that this profile is best utilized as an integrated quantity (as the authors do, and in Ronsmans 2018) due to the limited sensitivity and level correlation. But the authors need to correctly describe it as such. Claiming that it is an "integrated column product" is misleading as this indicates a profile-scaling type retrieval which is not how FORLI works. The incorrect phrasing is used in the abstract and in line 88-89.

[Response]: We now refer to the FORLI product as "a stratosphere + troposphere correlated profile retrieval" in the abstract and on Line 94.

[Reviewer]: The new figure 2, showing the full averaging kernel matrices for each of the algorithm configurations, does help describe the change in algorithm behavior across configuration changes. What should be added (either as an additional figure, or new section of figure 2) is the prior and retrieved HNO3 profiles and the posterior uncertainty. This new plot would be very similar to the Figure 1 of the Ronsmans 2016 AMT paper. Thus, this manuscript would have plots similar to Figures 1 and 2 of Ronsmans 2016, and this is needed for readers to fully understand the difference between the two retrieval methodologies.

[Response]: We added the following text on Lines 308–314 to clarify our depiction of HNO3 as horizontal maps and not vertical profiles.

"*In Figure 2, the AKD profiles of the* ***R1***–***R4*** *configurations indicate that the CLIMCAPS retrieved HNO₃ profile (*$\hat{x}$*) will approximate the climatological* x_a *in the troposphere irrespective of* x_{true} *because the AKD (and therefore the SNR) approximates zeros below the tropopause. It is only in the lower stratosphere where the AKD profiles have peaks* $\gg 0$ *that* $\hat{x}$ *will significantly deviate from* x_a*. For this reason, we do not depict or promote the use of the vertica profiles because they are not representative of HNO₃ throughout the tropospheric atmosphere. Instead, we depict the retrieved HNO₃ as 2-D maps of lower stratospheric integrated values. To reiterate, the goal of our retrieval system is to retrieve stratospheric HNO₃ estimates that are independent of the tropospheric SNR in the measurements*."

[Reviewer]: The discussion explaining the choice of prior is confusing and misleading - this is following line 232 ("Future work could focus on re-evaluating this AFGL profile...".
There is an argument here that the retrieval will be biased high if the prior is larger than the true state, "because \hat{x} >= x_a by definition", which is not a true statement - term added to x_a in Equation 1 can be negative. I think what the authors mean here is that |\hat{x} - x_a| is always smaller than |x_true - x_a|, which means that \hat{x} is always biased toward the prior.

[Response]: We removed the statement “because \hat{x} >= x_a by definition” and clarified the discussion about the AFGL profile on Lines 236–249 as follows.

“*The AFGL HNO_3 profile represents a global average ranging between 0.01–1.0 ppb in the UTLS, which is orders of magnitude smaller than the stratospheric values retrieved for HNO_3 in the extratropics during wintertime. Future work could focus on re-evaluating this AFGL profile for use as HNO_3 x_a, but the solution is not a simple replacement with a different estimate. As depicted in Eq. 1, $\hat{x}$ depends on adding measurement SNR to x_a, which means that whenever x_a is high relative to the true state of x, $\hat{x}$ will be biased high. Given the large dynamic range of HNO_3 during the polar wintertime months, we argue that it is preferable for x_a to be very small so that the retrieved product depict elevated HNO_3 values only where the true state (x_{true}) is high. Stated differently: when x_{true} is very small then the corresponding IR measurement SNR for HNO_3 is very low and $\hat{x}$ will approximate x_a. Conversely, when x_{true} is large, the measurement SNR is large and $\hat{x}$ will have a large departure from x_a. . This is all the more important for a target variable like stratospheric HNO_3 that is very difficult to represent with an accurate x_a at each space and time retrieval footprint because of the lack of real-time observations. We can have confidence in CLIMCAPS HNO_3 retrievals because $\hat{x}$ will be low whenever x_{true} is low, and $\hat{x}$ will significantly depart from x_a only when the measurement SNR is large.*”

With those minor changes, the paper can be accepted.