In this manuscript, the authors introduce a new set of all-sky and clear-sky, top-of-atmosphere (TOA) and surface radiative kernels, generated with the RRTMG radiative transfer model for 5 years of input fields from ERA5 reanalysis data. The authors incorporate these kernels into a more general inter-comparison of the magnitude and structure of existing sets of radiative kernels, highlighting the sensitivity of radiative kernels to the input climate state fiends used to compute them. Along these lines, they also highlight the sensitivity of radiative kernels to inter-annual variability in the climate state, taking advantage of the fact that the ERA5 kernels have been computed for 5 years of data, over a period with notable changes in ENSO and sea ice coverage.

The manuscript is comprehensive and well written. More multi-kernel comparison analysis is certainly needed, so this will be a welcomed addition to the literature. However, I feel the paper suffers in a few ways by trying to balance all three tasks (introduce a new kernel, multi-kernel comparison, sensitivity of kernels to variability) in one paper. For instance, I think the analysis of inter-annual variability is the most valuable part of the work, but the analysis is not well connected to the multi-kernel comparison, and the analysis is not as in depth as it could be. Additionally, it's not clear why we need the ERA5 kernels when they don't seem all that different from the others, and the kernels based on ERA-Interim reanalysis data (ERAi) previously developed by the second author of this paper were also calculated for 5 years of data and could have been used for the inter-annual variability analysis instead. Given the journal, I think justifying the new data product is important. My comments below reflect these concerns and hopefully add some additional, useful explanation to my points. Given these, I think the paper deserves consideration for publication, pending major revisions.

Ryan Kramer

General

1) As noted above, some additional justification for producing the ERA5 kernels is necessary, particularly given that the ERAi kernels exist, which use similar input data and a similar radiative transfer code, also for 5 years of data. For example, is the improvement of RRTMG (ERA5 kernels) over RRTM (the ERAi kernels) large enough to warrant new kernels? Even if so, the sensitivity of kernels to RT is not really a focus of the analysis here. Does ERA5 have more realistic climate state fields than ERAi? The two kernels were computed over different periods, so maybe that is reason to make new kernels? But the period for the ERAi kernels (2008-2012) also has notable swings in ENSO, so I don't quite see the advantage of the later period used for ERA5 kernels.

1a) Related, what is the justification for developing radiative kernels from reanalysis when the fields are available from models and observations? Arguably reanalysis offers a happy middle between the two. They may not be pure observations, but they do have the full diurnal cycle that most satellite observations do not have. This may be important for diagnosing feedbacks in models, where the model fluxes are also a response to the full diurnal cycle. But there is also the argument that, in order to diagnose the true feedback, model feedbacks should be diagnosed with a kernel developed from models and observed feedbacks should be diagnosed with appropriate observations. What is the value of reanalysis-based kernels in that context? Some discussion along these lines would be really valuable to a community often confused about what kernels they should be using.

2) This may be the only example where an "older" (ERAi) and "updated" (ERA5) radiative kernel were developed by the same research group using similar RT codes. This could be a really powerful tool for the multi-kernel comparison analysis and should be exploited here, but has not been yet. If the second author still has access to the ERAi kernel input data, I would like this team to include a more in-depth comparison of the ERAi and ERA5 kernels in the context of the multi-kernel intercomparison. Throughout the current manuscript, the authors highlight examples of large multi-kernel differences, tying them to potential differences in the underlying climate input data. For a given example, is the spread also evident in differences between the ERAi and ERA5 kernels? If so, the authors should analyze the climate input fields directly to reveal specifics about why the kernels differ. A few specific comments in the section below try to prompt this type of analysis. Among other groups, I think this could be extremely useful for the ECMWF developers of ERAi and ERA5, who are always trying to understand the biases and limitations of their product, giving your work exposure to an additional, large community.

3) Section 3.1: The authors should rethink the presentation and the focus of discussion in this section. First, general descriptions of the sign, basic explanation of the causes of the sign, and the zonal-mean vertical structure of kernels are discussed here for the new ERA5 kernels, but these topics have been covered extensively for other kernels and the new kernels don't seem to deviate from that. Therefore, it seems redundant to repeat that information here. This is true even for surface radiative kernels, where the structure and sign were covered by Kramer et al. 2019 a and b (see refs below). The description of the horizontal spatial structure of the kernels is newer however, and worthy of focus in this section. Second, the number of figures and figure panels in this section is also overwhelming for the reader and should be consolidated. Given these points, I would instead:

-For Figures 1-8, make the first figure or two just the spatial maps of each kernel, with a title for each subplot that describes which kernel we are looking at (e.g. all sky, surface temp, clear-sky SW WV, etc.).
-Given their prevalence elsewhere (e.g. Soden et al. 2008; Block and Mauritsen 2013; Kramer et al. 2019a,b, Smith et al. 2021), the latitude-pressure subplots of the ERA5 kernels can be combined and put in supplemental material.
-Any subplot currently referring to the intercomparison across existing kernels should be saved for new, separate figures placed in Section 3.2, where that material is discussed in the text.

This above list is just a recommendation. I'm sure there are other ways of reorganizing the plots to improve manuscript readability, but some reorganization is necessary.

Kramer, R. J., A. V. Matus, B. J. Soden, and T. S. L'Ecuyer, 2019: Observation-Based Radiative Kernels From CloudSat/CALIPSO. JGR Atmospheres, 124, 5431–5444, https://doi.org/10.1029/2018jd029021.

Kramer, R. J., B. J. Soden, and A. G. Pendergrass, 2019: Evaluating Climate Model Simulations of the Radiative Forcing and Radiative Response at Earth's Surface. Journal of Climate, 32, 4089–4102, https://doi.org/10.1175/jcli-d-18-0137.1.

4) It is evident that clouds play an important role in determining the spatial pattern of the all-sky radiative kernels. More description of the type of clouds impacting the kernels would be very useful and novel. For example, cloud vertical extent? Cloud base height (for sfc kernels)? Optical properties? A deeper analysis of the ERA5 cloud fields would be helpful here. And if the fields are still available for the ERAi kernels, even better.

5) I think the analysis of inter-annual variability in the kernel is the most interesting contribution of this paper to the literature. It deserves a more prominent place in the title, abstract, and conclusion section.

6) The inter-annual variability analysis feels disjointed from the introduction of the new kernels in Section 3.1, the multi-kernel comparison, and the feedback estimate section. It is evident that the ERA5 kernels are sensitive to inter-annual variability, but does this really matter for overall kernel spread? For instance, the ERA5 all-sky TOA Ts kernel is clearly sensitive to interannual variability in the Eq. Pacific at ~Longtiude 180, but this doesn't appear to be a particularly noteworthy area of inter-kernel differences in fig 1e. It may very well be important, but the analysis at present doesn't offer enough of a connection to make that point. Comparing the inter-annual variability in the ERA5 vs ERAi kernels in more detail may be a useful starting point to help make the connection between this section and the others.

7) After the interesting analysis showing the sensitivity of kernels to inter-annual variability spatially, the feedback analysis in Section 4.2 and 4.3 mostly just focuses on global-mean values. While this type of analysis is valuable in a general sense for kernel users, it doesn't quite fit well with this paper, particularly because the ERA5 kernels don't stand out as being more accurate or unique. Instead, I'd like to authors to focus more on multi-kernel differences in the feedback spatial patterns. Given the large focus on the pattern effect and the influence of regional feedbacks on the global-mean in recent years, I think that could be particularly citeable. We know kernels are generally in agreement in the global-mean (especially for the TOA). But what about for kernel spread in estimates of the regional feedbacks?

8) It is tough to pick out valuable information from your tables of TOA and Surface feedbacks. The authors should turn those into summary figures (e.g. dot plots or scatter plots used by e.g. Smith et al. 2018 supplemental, or Zelinka et al. 2020) where possible.

9) There are now many observation-based kernels in the literature from CloudSat/CALIPSO observations, CERES CCCM products, AIRS, and a bunch of others specific to surface albedo kernels that are not included in the multi-kernel analysis here. I think it's an open question whether these observational kernels should be included in a comparison with model-based kernels or not. The authors should provide a brief reason for not incorporating them (or should include them if they feel its appropriate), especially since the Kramer et al. and Thorsen et al. reference papers are cited in the text.

Specific or Minor Comments

Line 311-312 and Appendix: We discussed similar RT issues regarding the surface temperature kernel for surface fluxes in Kramer et al. 2019 (JClim). The authors can cite and/or refer to it for some additional support. We also argued that, as noted in the author's current Appendix and elsewhere, similar RT issues can bias the lowest level of the surface flux Ta kernel in an equal and opposite manner as the Ts kernel, thereby allowing a kernel like CAM5 to be correct in its estimate of the vertically integrated Temperature feedback, but for the wrong reason. This likely explains why CAM5 and ERA5 kernels agree in Figure A2c and A2d. Your text somewhat gets to this point, but I'd call it out directly as a warning to kernel users: Some kernels may give you the correct temp. feedback for the wrong reason. Presumably this is true for the TOA temperature feedback too, but the contribution from the surface and near surface layers to that vertically integrated feedback are small, so maybe it doesn't matter much?

Kramer, R. J., B. J. Soden, and A. G. Pendergrass, 2019: Evaluating Climate Model Simulations of the Radiative Forcing and Radiative Response at Earth's Surface. Journal of Climate, 32, 4089–4102, https://doi.org/10.1175/jcli-d-18-0137.1.

Line 316-319 and more generally: Bright and O'Holleran (2019) and Donohoe et al. (2020) performed nice, detailed comparisons of surface albedo kernels. These papers should be cited and their work should be put into context of the author's own results within the present manuscript. This would also be useful for the section on diagnosing radiative feedback spread, since the authors show that the kernels give quite different results in the poles. Riihela et al. (2021) also did a comparison of surface albedo kernels in the context of sea ice states, and should be cited somewhere in the present manuscript.

Bright, R. M., and T. L. O'Halloran, 2019: Developing a monthly radiative kernel for surface albedo change from satellite climatologies of Earth's shortwave radiation budget: CACK v1.0. Geosci. Model Dev., 12, 3975–3990, https://doi.org/10.5194/gmd-12-3975-2019.

Donohoe, A., E. Blanchard-Wrigglesworth, A. Schweiger, and P. J. Rasch, 2020: The Effect of Atmospheric Transmissivity on Model and Observational Estimates of the Sea Ice Albedo Feedback. Journal of Climate, 33, 5743–5765, https://doi.org/10.1175/jcli-d-19-0674.1.

Riihelä, A., R. M. Bright, and K. Anttila, 2021: Recent strengthening of snow and ice albedo feedback driven by Antarctic sea-ice loss. Nat. Geosci., 14, 832–836, https://doi.org/10.1038/s41561-021-00841-x.

Line 316-227 and more generally: Aligning with general comment 4 above, the author's assumption that clouds are the cause of the discrepancies discussed here is likely right, but this has been alluded to before in past work. The author's have a unique opportunity to prove it directly by using the kernel's cloud input data directly in the analysis. Using the ERA5 and ERAi cloud fields can the authors confirm their assumption that cloud fields mattter? Or provide a more detailed analysis? Among all the potential sources of kernel differences, clouds seem to warrant deeper investigation.

Figure 2f: The ERAi and ERA5 kernels are noticeably different below ~800mb. Why?

Figure 3e and 3k: There are large standard deviations relative to the magnitude of the ERA5 kernels at certain locations within the vertical kernel structure, but for anything above ~950mb, the kernel magnitude is small relative to the lowermost atmospheric levels, and likely does not contribute much to the vertically-integrated quantity. A zoomed in version of these plots, highlighting the important standard deviation across kernels in the lowermost levels, would be informative.

Figure 3f and 3l. Why is the ERAi kernel so much larger at the lowest levels near the surface than the ERA5 kernel (and the other kernels)? This is true for both all-sky and clear-sky. Could vertical resolution of the kernels be playing a role? We discuss the potential influence of resolution in the Appendix of Kramer et al. (2019, JClim).

Interestingly, there is also a fairly large difference between ERA5 and ERAi in the all-sky LW WV kernel for surface fluxes (figure 5L), but it only shows up in the all-sky kernel, not the clear-sky. Does that suggest clouds matter more for the LW WV kernel in explaining differences than they do for the LW Ta kernel, relative to other potential sources of kernel spread?

Line 368: Is this the sensitivity of the surface to the vertically integrated kernel change? Only the sensitivity to a certain level of WV change? It is not obvious from the figure 9 caption either. Some rethinking of how the kernels are described in the text here and elsewhere would be helpful. Maybe shorthand acronyms or some other naming convention would be helpful.

Plot 9: I really like this figure. I'm not sure anyone else has shown the sensitivity of these temperature and water vapor kernels to variability spatially yet. But I think it would be helpful to go one step further and show what level of WV and cloud variability is impacting the temporal variability of the kernels most. And does that particular level help explain the multi-kernel spread in those kernels, evident in Figure 1-8? Or is inter-annual variability not enough to explain the kernel spread? This comments connects with my general comment #6 above.

Plot 9: The TOA SW WV kernel is a potentially interesting case where the largest multi-kernel spread (tropics around 500mb in Figure 6k) sits above the level at which the ERA5 kernel is the strongest (closer to 800mb in Figure 6j). Building on my comment above, how does the importance of inter-annual variability play into the large kernel spread at this ~500mb level? And does the spread at 500mb actually matter much for the vertically-integrated quantity? This level of detail could be useful for e.g. modeling centers trying to connect TOA biases to particular biases in their climate states.

Figure 10: I struggle to see spatially where the variability in water vapor is having an effect on the kernels shown in the other subplots. Maybe only for the NE coast of Greenland? Should cloud changes be shown in this plot instead? I'd give more detailed evidence about why water vapor is important here.

Line 412-419: Some explanation of why you need both abrupt4xCO2 and piCLim-4xCO2 (e.g. to remove rapid adjustments) is needed, since most people just use abrupt4xCO2 with Gregory regression to get feedbacks.

Line 436 and Equation 5: What does the clear-sky residual term mean physically and why does it matter for computing cloud feedback? Although the author's math in Equation 6 works out to be the same as what everyone else uses, I think the way they've introduced this method, and terminology used, is less common. Some extra detail would be helpful.

Line 463-464: Block and Mauritsen (2013) can be cited here for their nice discussion and analysis of the non-linearity of the surface albedo kernel in 4xCO2 runs.

Block, K., and T. Mauritsen, 2013: Forcing and feedback in the MPI-ESM-LR coupled model under abruptly quadrupled CO2. J. Adv. Model. Earth Syst., 5, 676–691, https://doi.org/10.1002/jame.20041.

Last Column Table 3: The multi-model values differ somewhat across the different kernels, but not the associated standard deviations, which look to be essentially the same for all rows of the column. Does this suggest the kernels can estimate feedbacks differently in a systematic manner across all models, but they do not necessarily estimate the magnitude of the feedback spread differently? In other words, the kernel may get the model-mean feedback value wrong, but the spread in feedbacks correct? This is worth noting if so.

Line 487-491: Since you are not using cloud radiative kernels, the inter-kernel differences in cloud feedback must come from the difference between all-sky and clear-sky kernels (cloud masking). Can you identify which of the kernel terms is the culprit? From a related discussion see text around figs 9-11 in Kramer et al. (2019, JGR), for example.

Kramer, R. J., A. V. Matus, B. J. Soden, and T. S. L'Ecuyer, 2019: Observation-Based Radiative Kernels From CloudSat/CALIPSO. JGR Atmospheres, 124, 5431–5444, https://doi.org/10.1029/2018jd029021.

Line 769-793: This is a really nice description of how you develop the water vapor kernel. I'd highlight in the main text that you have included this section in appendix. I think many kernel users and developers are looking for a description like this and will turn to it in the future. There are often question about this calculation.