

Interactive comment on “Surface and top-of-atmosphere radiative feedback kernels for CESM-CAM5” by Angeline G. Pendergrass et al.

Angeline G. Pendergrass et al.

apgrass@ucar.edu

Received and published: 9 January 2018

We thank both Reviewers for their attention to our manuscript and for their feedback. We have revised the manuscript, and we think it is substantially improved, thanks in part to the Reviewers' comments. Please find our responses below, as well as a version of the manuscript with tracked changes as a supplementary attachment.

1 Response to Anonymous Reviewer 1

We thank Reviewer 1 for their attention to our manuscript and for their feedback. We have incorporated your feedback and expanded the validation, and hopefully improved

C1

the manuscript. In this document, your comments are italicized and our responses are in plain type.

The authors present a concise manuscript documenting and describing radiative feedback kernels for the top-of-atmosphere and surface produced with the CESM-CAM5. These kernels are likely to contribute to scientific advancement since they are one of few sets of documented kernels that also provide surface radiative responses and because they are produced from a different model and radiation code (although some clarification may be needed, as seen in detailed comments later). The authors also provide an assortment of example scripts that can be used for application of the kernels, which will be very helpful to the community. Nonetheless, I think the paper/data can be improved with regard to a few general things:

1) Some of the details in producing the kernels and applying them are lost due to the very concise nature of the text. While some of these details may be obvious to those with experience in producing/using kernels, there is not quite enough basic information and detail for first-time kernel users. The authors provide several references pointing toward more information, but I don't think it would hurt or greatly increase the length of the text to give some of that information here. See specific comments below.

Thank you for your constructive feedback about how to improve the text, in particular places where more information for those unfamiliar would be useful.

2) The provided data and scripts are a bit disorganized. If possible, I recommend putting all of the relevant data and example code in one place, rather than spanning two webpages. It may also help to organize the files into folders or categories based on their type (e.g., kernel files, forcing, relevant code, demos). On a more technical note, the readme file currently on the zenodo webpage should point to the actual filenames of all scripts that are described under “Additional scripts. . .”

We would of course also prefer a consolidated location for the data and the software,

C2

but unfortunately we do not have access to a single repository that will provide long-term persistence for both data and software that meets the needs of our project. Specifically, we feel it is important to be able to include updates to the software as it is disseminated and the users find bugs. Github is an appropriate place to do this. However, the Github repository does not allow uploads of greater than 2 GB per project, and our data are 2.6 GB. The repository where the data are provided, the Earth System Grid (ESG), is widely used in the climate modeling community, but it is not possible to provide updates to the software there, so it is not an appropriate place to keep the software – though we do include a minimal set of scripts there, along with the data.

The README file included with the Zenodo data link does not contain the exact names of scripts that are found in the Github repository because the software repository can be updated over time, while the data repository cannot. For example, in response to offline comments associated we have included an additional script. We will continue to update the scripts in the event that users find additional code bugs.

Below are a series of specific comments, related to both scientific content and presentation.

Section 2, page 2, L6: It appears this radiation scheme is based on the CAM4, which is quite different and somewhat outdated compared to that in CAM5. Was the choice of this scheme because no other currently available kernels make use of it (e.g., I believe kernels based on ECHAM6 use a similar radiation scheme as CAM5), thereby increasing structural diversity among available kernels? It would be helpful to say a bit more here with regard to this.

The radiation scheme used for these kernels is RRTMG. It is used in CAM5, but not CAM4. Thank you for pointing out this ambiguity in our text, which probably arose because the published article about PORT that we cite was based on CAM4. PORT was originally developed for CAM4, and here we use a version that was ported to CAM5. The original manuscript stated that the CAM5 microphysics were ported (this was a

C3

substantial technical challenge) and though we did not state it in the first manuscript version, the radiation scheme was also ported (but this was straightforward and not as memorable). We now state this explicitly and also cite the RRTMG radiation scheme.

Section 2, page 2, L8: Please explicitly state here what these “other necessary fields” are.

To calculate the kernels, the fields that are varied are temperature, mixing ratio, and four surface albedo parameters. For the radiative forcing calculations, carbon dioxide, methane, and a set of aerosol fields are also varied. Running the radiative calculation also requires the full instantaneous model 3-d atmospheric state to run, including clouds. This is now stated explicitly in the text.

Section 2.1, first paragraph: State here that the simulations are only conducted for one year.

We now include this statement.

Section 2.1, L16: It is unclear where the number 63 comes from.

The control, surface temperature, and albedo are each one kernel-month of calculation, and air temperature and moisture are each 30 (one per level). This is stated in the revised text.

Section 2.1, L21-22: It should be mentioned that these perturbations are done at each grid cell.

We now include this.

Section 2.1, L22: Please clarify if these perturbations are computed with respect to each corresponding control 3-hour timestep.

C4

They are computed as monthly-mean differences. However, it should be noted that because these are absolute differences and averaging and differencing are linear operations (addition and subtraction, with division only by a constant), the order in which the averaging is done has no effect on the result. The order is chosen for computational efficiency (holding active storage as few 3-hourly instantaneous output fields as possible). The monthly mean difference of three-hourly perturbed and control fluxes is exactly equal to the difference of monthly mean three-hourly perturbed and control fluxes. Nonetheless, this is now included in the text.

General computation comment: When going from the full kernel field (3-hourly, grid cell, level resolved) to the global averages presented here in tables, what is the procedure? For example, flux responses are first averaged in time (to monthly then annual) and space, and then vertically integrated with pressure weighting? I feel this should be mentioned.

As discussed above, averaging absolute quantities in space and time is a linear operation, and this is also relevant here - the order of space and time averaging does not affect the result. Operations relevant to application of the kernels that are not interchangeable are, indeed, the pressure weighting (when pressure-coordinate kernels are used), and also normalization by global-mean surface temperature change. Pressure weighting is implicit to the kernels on the native CESM hybrid-sigma coordinate and is only employed when the kernels are interpolated to pressure levels. Pressure weighting is relevant for Figures 1 and 2, where the zonal-mean vertical structure of the kernels is presented, and no other figures or tables. Soden et al (2008) and Block and Mauritsen (2013) may have used pressure weighting in their calculations (which we reproduce in Table 2). As for normalizing by global mean surface temperature change, this is done in the final feedback calculations presented in Table 2 and Figure 6.

In the revised manuscript, we now state, "The global, annual mean change in radiative flux for each feedback is calculated and then normalized by the change in global-mean

C5

surface temperature for each ensemble member; then, the ensemble average is calculated."

To describe the procedure for calculating Figures 1 and 2, the revised manuscript states, "To create these plots, the kernels are regridded to standard CMIP5 pressure levels in the troposphere including, pressure weighting (a vertical regridding script is available at <https://github.com/apendergrass/cam5-kernels>), and then the zonal and annual means are calculated."

Fig. 1 caption: This caption has some repetition and missing information. The Fig. 2 caption is fine and should be replicated for Fig. 1 but with obvious modifications.

This seems to have been a proofreading error. Thank you for pointing it out; we have now added the additional information to the Figure 1 caption.

Section 2, general: Can you expand a bit on how the patterns shown in the figures come about physically (e.g., the larger temperature kernel magnitudes in the tropics and the multi-peak structure in Fig. 1e), or at least provide references to previous work that has done so?

We avoided interpreting our data because this is beyond the scope of ESSD. We now provide two references to previous work that discusses the patterns.

Fig. 3 caption: Should note that all panels are showing LW+SW, assuming that is actually the case.

"Net" radiative forcing is LW+SW. We have now included LW+SW parenthetically to be more clear.

Section 3, paragraph 2: It's a bit fuzzy what is done here. Are the kernels multiplied by the responses of the changing fields between the two years (with a procedure anal-

C6

ogous to section 5)? Please expand. Also, explicitly state that these calculations are compared with the model-calculated flux responses between the same years.

We have overhauled the validation in the revised manuscript, both in response to your comments as well as those from Reviewer 2. We have revised this paragraph so that the validation procedure is more clear, as suggested.

Figures 4 and 5 captions: State that what is being shown is the global mean.

These figures no longer appear in the manuscript; analogous information is now available in the updated Table 1. The row labels in the updated table state that the errors represent the global mean.

Section 3, paragraph 3 and related figures (4 and 5): Wouldn't the model-produced all-sky flux responses between 2006 and 2096 include cloud feedbacks as well? If so, what is the value of comparing those responses with corresponding all-sky kernel estimates that do not include the cloud changes (as stated on L11)?

Thank you for prompting us to elaborate on the focus of the validation. As mentioned above, we have overhauled the validation in the revised version of the manuscript. Both clear-sky and all-sky fluxes are calculated by the model. All-sky fluxes do ultimately determine climate feedbacks, but the method that is typically used to calculate the cloud feedback, the adjusted cloud radiative effect technique, results in errors of all-sky and clear-sky fluxes (in Wm^{-2}) which are exactly equal.

Section 4, L22-23: change "based on changes in TOA..." to "based on model-computed changes in TOA. . ."

This paragraph no longer appears in the revised manuscript. Instead, the error due to internal variability is diagnosed explicitly in the previous section.

C7

Section 4, L24: Insert "approximately" or symbols, e.g., "~", in front of these values.

As with the previous comment, this paragraph no longer appears in the revised manuscript.

Section 5, first paragraph: It would be useful to 1) give some brief qualitative discussion about how these calculations are conducted (e.g., that the kernels are first monthly averaged then multiplied by monthly-resolved climatological changes in the fields) and 2) Include here a reference to section 6 where one can find more information about relevant example code.

Thank you for pointing out how we can make better guide our readers. We have made both of these changes.

Section 5, L6: Could you clarify why one needs long-term mean water vapor mixing ratios?

The moisture kernels are normalized by climatological moisture so that the difference follows the logarithm of moisture. The logarithm of moisture changes more linearly with warming than moisture itself, as noted by Soden et al (2008), so this increases the accuracy of the fluxes estimated with kernels over using the unnormalized change in mixing ratio.

Table 2 caption: Suggest changing the "Here" heading to "CAM5"

We removed the word "here" (CAM5 already appears earlier in the sentence).

Section 5, L25: Unsure why long timescales are more similar to your calculations since your kernels are only computed for one year. Please explain.

Block and Mauritsen (2013)'s "long" timescale is simply the long-term climate response. Their alternative short timescale focuses on the first 20 years of a rapid transient re-

C8

sponse. The climate change from 2006 to 2096 is expected to be more similar to the developed long-term response from their “long” timescale, and thus it is the appropriate comparison here.

The kernels are only calculated for one year (in our kernels as well as those from Block and Mauritsen 2013), rather than, say, 30, or even individually for each experiment, because the computational expense and the human expense of managing the calculation is large. The purpose of radiative kernels is to provide a computationally efficient shortcut to decompose radiative fluxes into their contributions from temperature, moisture, and albedo.

We now clarify in the text that the long timescale refers to the application of the kernels, rather than the kernels themselves.

Section 5, L30, with regard to “0.02,” perhaps also specify which feedback this is and give this value as a percentage of the mean as well. I also recommend reiterating here that this estimate does not account for potential variability among kernels themselves if computed from different years/ensemble members.

This paragraph no longer appears in the revised manuscript. Instead, we diagnose the error due to internal variability explicitly using the ensemble in the revised Section 3.

Section 6, L4: Please again specify what the four radiative kernels are. Also note that they are provided as monthly averages in the files.

We now include this and also state what the two forcings are.

Typos/grammar:

Thank you for your careful reading of the manuscript. We have made all of the recommended changes.

C9

Section 2, page 2, L25: change “has units” to “have units”

Section 3, page 4, L5: insert “of” between “decomposition” and “the”

Section 3, page 4, L6: Suggest changing “linear in. . .” to “linear with respect to. . .”

Section 5, page 5, L9-10: Should “the changes in top-of-atmosphere radiative feedbacks” just be “the top-of-atmosphere radiative feedbacks”?

Section 5, page 5, L16: “m(Pendergrass, 2017a)odel”

C10

2 Response to Anonymous Reviewer 2

We thank Reviewer 2 for their thorough and critical feedback on our manuscript. We have incorporated your feedback and expanded the validation, and hopefully improved the manuscript. In this document, your comments are italicized and our responses are in plain type.

This paper presents a new set of radiation kernels that can be used for feedback analysis. This affords a useful means for GCM inter-comparison and for understanding their differences in climate sensitivity. I recommend publication after the following issues are properly addressed.

Two major suggestions are: 1. Provide comparisons between this new set of kernels and other kernels and note their differences. This would serve the community greatly to understand whether (and how) feedback determined using the kernel method is sensitive to the kernel dataset used.

Comparison between datasets is certainly important and useful. We touch on this briefly in Table 2. Comparison of spatial and vertical structure across different kernels is explored in Shell et al (2008). Furthermore, this is the topic of a separate, forthcoming manuscript, which includes extensive comparison across many kernel datasets, which the first author is involved in. We intend the present manuscript to focus on documenting the CAM5 kernels, consistent with the scope of ESSD, leaving comparison among kernels to the forthcoming manuscript. That said, in the process of preparing that manuscript, we have made comparisons with other existing kernels and are confident that our kernels do compare reasonably with others, despite that this analysis is not included in the present manuscript. Also of potential relevance to your comment, see below for how the zonal mean kernels as shown in Figs. 1 and 2 now are much more consistent with previously published kernels, after resolving a plotting error.

C11

2. Document more extensive validation tests. Currently only global mean values from one case (2096-2006) are reported (Fig. 4, 5); it is not clear whether the percentage errors reported are representative. Error statistics based on global maps of radiation changes and time series of global means would make a more rigorous assessment.

In the revised version of the manuscript, we have overhauled and extended the validation. Because the radiative forcing varies over time in the RCP8.5 scenario that was used to generate the kernels, validating the kernels themselves (rather than the forcing fields, which are secondary) is only tractable for the years 2006 and 2096. But we now make use of the other 39 members of the CESM1 large ensemble to validate against the changes from 2006 to 2096 in the members that were not used to generate the kernels. In Fig. 3 of the revised manuscript, we document the global-mean absolute error of TOA and surface LW and SW changes in radiative flux for each of the ensemble members. We have updated Table 1 to document the error of global mean kernel-estimated flux change and global mean absolute error of kernel-estimated flux changes from ensemble member 1.

Additional comments

Page 1, Line 21. It is not a correct statement. New TOA and surface kernels computed from EARI[sic] atmosphere have been made available by Huang et al. (2017). This should be referenced here.

Thank you for pointing out this omission. In the revision, we qualify that Previdi (2010) are the only model-based surface kernels, but Huang et al (2017) provide reanalysis-based kernels.

Page 2, Line 7. Note there is strong inter-annual variations of atmospheric states, e.g., El Nino vs. La Nina, which affects kernel values quantitatively. In relation to Suggestion 2 above, it is worth discussing, and if possible demonstrating, (in)accuracy in feedback, e.g., in the central tropical Pacific region, related to this issue.

C12

In Fig. 4 of the revised manuscript, we show maps of the mean error across these 39 ensemble members to document the spatial pattern of error.

Fig. 1. I am surprised to see the temperature sensitivity maximize in the middle tropical troposphere (~600hPa) here, which is noticeably different from other kernel datasets (e.g., Soden et al. 2008, Huang et al. 2017) - that may be due to misplaced clouds in CAM?

Please see response below.

Fig. 2. I am also surprised by that the surface radiation sensitivity to temperature and humidity doesn't maximize in the lowermost atmospheric layer. These results do not agree with radiative transfer-based expectations.

Thank you for pointing out these inconsistencies with existing literature. Upon revisiting our code (and at the prompting of kernel user Paulo Ceppi), we realized that the way we incorporated surface pressure into the grid before taking the zonal average was the cause of this difference. We have updated Figures 1 and 2 accordingly, and they are now much more in line with other published zonal-mean kernels. The bug had a small effect in one script in the provided software, which we have now updated as well. To be clear, this was only a plotting error and has no effect on the kernels themselves, nor on any results in hybrid-sigma coordinates (Figs. 3-6 and Tables 1 and 2).

Please also note the supplement to this comment:

<https://www.earth-syst-sci-data-discuss.net/essd-2017-108/essd-2017-108-AC1-supplement.pdf>

Interactive comment on Earth Syst. Sci. Data Discuss., <https://doi.org/10.5194/essd-2017-108>, 2017.