

Interactive comment on “The HadGEM3-GA7.1 radiative kernel: the importance of a well-resolved stratosphere” by Christopher J. Smith et al.

Christopher J. Smith et al.

c.j.smith1@leeds.ac.uk

Received and published: 15 June 2020

This manuscript documents a new set of “radiative kernels,” calculations of the sensitivity of radiative flux to atmospheric state, developed from a current-generation climate model with a domain reaching well into the stratosphere. The construction of the kernels was motivated by a desire to understand the fast response of stratospheric temperature to changes in carbon dioxide concentration and the authors demonstrate the value added by the new kernels. The construction of the kernels is described and their accuracy and generality assessed.

The data are well worth publishing. They require substantial computational resources to produce, extend the vertical domain in an almost-unique way, and use an accu-

C1

rate radiative transfer code. The free availability of the data has been verified. The manuscript is effective at documenting how the kernels are produced, providing enough details for readers to understand and potentially replicate the steps. It is also effective at motivating why this implementation is useful, noting that the diagnosis of the fast climate response (the adjustment) to increased concentrations of carbon dioxide depends importantly on having a deep vertical domain. Beyond a few small points of expression noted below the manuscript could be most improved by more context for the uninitiated and a more general treatment of some ideas.

»» Thank you for your positive overall comments and we are pleased that you agree that the data and description paper is worth publishing in a form close to present.

The introduction, which introduces the concept of and motivation for a radiative kernel, may be more general than is needed for the present manuscript. The generality makes it open to objections as to how the ideas are expressed. The general idea of a kernel is the ability to compute flux perturbations from state perturbations. As originally implemented by Soden, Shell, and others, these were restricted to specific characteristics of atmospheric state (air and surface temperature, water vapor, surface albedo, and excluding clouds) based partly on prioritization and partly on based on the availability of data. The does not establish a “standard” (line 47) nor does it exclude in principle other variables from being relevant (line 28). Readers may also wonder how the general material on the use of kernels (lines 39-54) is directly relevant to the construction of the present kernels.

»» Thank you for these suggestions. In our opinion the introduction is not overly long at present so we are inclined to keep lines 39-54 in the paper as general background. We note that others have taken a different approach and have assumed more background knowledge (e.g. Prendergrass et al 2018). Readers familiar with radiative kernels could easily skip over the introduction. Those not familiar may welcome it, and as the kernel method is used increasingly outside the climate feedback community in which it was developed, it may make the paper more self-contained.

C2

»» Line 47 has been changed to “Cloud adjustments and feedbacks cannot be determined directly using atmospheric state kernels”. The previous wording could be taken to imply that cloud adjustments/feedbacks could not be calculated at all using only atmospheric kernels, which was not the intention.

»» Line 28: after this sentence, included “Although other (non-cloud) variables may also be relevant, the majority of adjustments are expected to be captured under this framework (Vial et al., 2013).”

The authors might revisit the introduction and focus it more tightly on the subject of the manuscript. This might include a not too profound explanation of how kernels can be used to diagnose both feedbacks and adjustments, and an explanation as to why yet another set of kernels might be desirable (i.e. the material that begins section 4). Care should be taken not to confuse routine practice with standardization.

»» To improve the motivation, the first paragraph from section 4 has been assimilated into the introduction. Following the previous response, we believe the present level of background is appropriate.

Section 2:

It would be worth noting explicitly that these kernels rely on two almost distinct aspects of HadGEM: the radiation code SOCRATES run at low spectral resolution, and the climatology of atmospheric state including clouds, even if experience shows relatively weak dependence on the background state.

This section has been updated to include a sentence:

»» "The kernel is therefore dependent on two aspects of the HadGEM3-GA7.1 model: the pre-industrial background climatology (including clouds), and the broadband version of the radiation code."

The authors might also explain some of their choices and any expected impacts. These might include the choice to develop kernels for pre-industrial conditions, the relatively

C3

highly-resolved vertical structure and coarse horizontal resolution of the simulations, and the high time resolution.

»» In the updated section 2 we have added a few sentences at opportune points, justifying some of the methodological choices. More explanation for each is given below.

»» The pre-industrial conditions were chosen for the kernel base state as these kernels were designed to be used first and foremost with the RFMIP piClim-X forcing experiments ($X = 4\times\text{CO}_2$, present-day GHG, present-day aerosols, present-day land-use and present-day total anthropogenic), covering both negative and positive forcing relative to the pre-industrial. The base climate of the target models is pre-industrial except for the perturbed component(s) and differences are taken with respect to piClim-control (an atmosphere-only pre-industrial control run). Therefore, a pre-industrial climatology for the kernel is appropriate. Other choices were of course possible such as linearising around present-day conditions.

»» The relatively high (for a GCM) vertical resolution is the default configuration of the HadGEM3-GA7.1 model, and its importance for stratospheric adjustment is already stated.

»» The horizontal resolution ($1.875^\circ \times 1.25^\circ$) is the lowest resolution used for HadGEM3 and UKESM1 in CMIP6. Many other CMIP6 models that are ultimately the target of the radiative kernels (excluding HighResMIP) are on similar resolutions to this. We are not aware of an exhaustive list, but Table 1 in Smith et al. (2020) lists those used in RFMIP, so it is not apparent that an increase in resolution in the base climate would improve accuracy after the kernels are re-gridded to the resolution of target climate models. But the main reason why a finer resolution wasn't considered is computing time. The “MM” resolution of HadGEM3 has an atmospheric grid of $0.83^\circ \times 0.56^\circ$ with 5 times as many grid points, so running the offline radiation model would have taken 5 times longer. To do this in a reasonable amount of time would require

C4

a reduction in the vertical or time resolution (or use of HPC). The high vertical resolution of these kernels are one of their strengths, and it would be difficult to reduce the number of temporal radiation calls as explained below.

»» A two-hour time step was used as instantaneous rather than time-mean climate output is needed to run the offline radiative transfer code (a discussion of why time-mean output does not work well is in Bellouin et al, 2020; section 4). Two hours was considered fine enough to avoid biases by undersampling the diurnal cycle of temperature, humidity and clouds, and shortwave solar geometry. On the last point, using longer timesteps like 6 hours between shortwave calls leads to some longitudes receiving substantially more incoming solar radiation than others over the course of the year. This effect could be achieved alternatively by using a timestep that does not divide 24 (e.g. 22, 23, 25 or 26 hours) and running several years of the climate model to sample diurnal, seasonal and interannual variability, such as is sometimes done in PRP calculations.

Section 4 illustrates the added value of the new kernels quite nicely. The use case is important but a little narrow. Is the value also added for other greenhouse gas forcings?

»» Thank you for your positive comments here and agree it is beneficial to demonstrate why the kernels are useful. In response to an earlier comment we moved the first paragraph to the introduction.

»» It is indeed useful for other GHG forcings. We generalise from CO₂ to GHGs in a few places in Section 4, although the focus of the results will still be on CO₂ as we have the double-call results from IPSL-CM6A-LR for this experiment.

»» We didn't re-run every kernel calculation using the piClim-ghg experiment in this paper. However, in supplementary figure 1 in Smith et al. (2020) we compare four of the kernels for all RFMIP experiments, showing that HadGEM3, ECMWF-Oslo and GFDL kernels show larger stratospheric adjustments for piClim-ghg than the CCSM kernel which is in accordance with the 4xCO₂ results in this paper. In attachment to

C5

this response we show the temperature profile for the GHG experiment, which moves in the same direction as for 4xCO₂ but with a lower magnitude.

[Figure B: stratospheric temperature differences from piClim-ghg minus piClim-control for 13 models in RFMIP].

»» Also note that in figure 4 we have removed the missing top level from GFDL-CM4, which was previously displayed as zero rather than missing.

Section 5 is the least organized and clear and the section seems to assume a lot of background knowledge. The point of the section is to demonstrate the accuracy and applicability of the kernels. The narrative should be constructed to as to make this goal clear, explain how accuracy and applicability can be assessed, and finally to demonstrate the results.

»» Thank you for this suggestion. It has been re-written to introduce the aim of the exercise (the discussion on IRF appears too soon and is given too much weight) before showing the results. Following suggestions from reviewer #1 we have also included a comparison of the results here to a previous study from CMIP5. To address the assumption of background knowledge on interpretation of section 5, we note that the introduction section may come in useful here to the uninitiated reader, which is a reason why we feel it is appropriate to keep most of it.

Lines 2-3: "the utility of radiative kernels. . . is most appropriate" The last word isn't quite right. Utility can be greater or less but not appropriate.

»» Thanks for pointing out the confusing wording. Changed "most appropriate" to "greatest".

Line 23: Kernels represent derivatives of flux with respect to state, not differential equations

»» A slight terminological liberty taken on our part. We'll keep the notation of eq. (1) as it is used by others (e.g. Shell et al., (2008) where they have used (F-Q) in place

C6

of R, and Huang et al. (2017)) and provides a nice concise representation but explain that it isn't strictly a differential equation.

Line 32: climate model (not mode)

»» Typo corrected - thank you

Line 84: the equation should have units

»» Updated to confirm that 10000 and p_{thick} are both in units of Pa here.

Line 113: The sudden appearance of PDRMIP may confuse the uninitiated

»» PDRMIP (Precipitation Driver and Response Model Intercomparison Project) acronym is now introduced, which was an oversight in the first submission.

Line 125: the limitations of low-topped kernels are presumably independent of whether the state comes from a "climate model" or any other source

»» Agreed: revised this sentence to be more general: "For kernels built from underlying atmospheric profiles where the top of the profile is not sufficiently high or with too coarse a resolution in the stratosphere, this additional upper stratospheric cooling is missed."

Line 161: cars break down - what is meant here is "decomposition" or similar

»» As part of the re-write of section 5, this sentence will be revised.

Line 163: "ways of calculating the residual can be obtained" is a confusing phrasing

»» Agree this is not meaningful: changed to "Two different ways of calculating the residual exist."

Interactive comment on Earth Syst. Sci. Data Discuss., <https://doi.org/10.5194/essd-2019-254>, 2020.

C7




Fig. 1. Figure B: stratospheric temperature differences from piClim-ghg minus piClim-control for 13 models in RFMIP

C8