

Interactive comment on “RadNet 1.0: Exploring deep learning architectures for longwave radiative transfer” by Ying Liu et al.

Ying Liu et al.

ying.liu.sweden@gmail.com

Received and published: 14 April 2020

article [utf8]inputenc color url

C1

Reply to the Reviewer 2 comments for: RadNet: Exploring deep learning architectures for longwave radiative transfer

We thank the reviewer for his/her comments.

1. The paper explores the use of neural networks architectures to model radioactive transfer in global circulation models. The paper is interesting and well written. Authors carried out thorough experiments and the conclusions about the use of NN in the field are interesting. Additionally, the source code for the implementation is publicly available, which is really useful. I have a few concerns about the experiments and how the performance is calculate.

Response: We thank the reviewer for their inputs and comments. We address the concerns below.

2. Speedups of table 2 are obtained by comparing the computational time of RRTMG and NN. I suppose that the baseline time for RRTMG is obtained in the Xeon CPU, and all the values are computed using this baseline. The

C2

objective of the paper is to evaluate the advantage of replace a RRTMG for a NN. To do that, a comparison in term of execution time is done. But, to be fair, the comparison of the NN should have been done with a RRTMG implementation for GPU. In the conclusions, authors say that "a complete rewrite of RRTMG for GPUs gives very similar performance gains as what we have achieved in our setup using NNs (Price et al., 2014; Mielikainen et al., 2016)". You should include evidences to support this affirmation. For example, in Mielikainen et al. author state that their implementation provides a speedup of 202x on a Tesla K40 GPU, compared a single-threaded Fortran counterpart running on Intel Xeon E5-2603 CPU. A higher speedup (370x) is indicated in Table 2, for NN model A and a batch size of 4096 running in GTX 1080. Both GPUs architectures are very different, so result are not directly comparable. In fact, single precision floating point (float32) performance is higher in the GTX 1080 than in the Tesla K40 (it's the opposite for double precision).

Response: We thank the reviewer for pointing this out. We agree that our comparison does not take into consideration differences in GPU architecture. We do not have access to a GPU implementation of RRTMG, which does not allow us to make a precise estimate of the kind pointed out by the reviewer.

The main point we were trying to make is that GPUs play a significant role in accelerating RadNet. We have qualified our conclusions to highlight this fact.

Lines Changed: 347–349

3. Related with the previous question, Intel Xeon CPU E3-1230 v5 is a quad-core processor (8 virtual cores). Is the

C3

execution of RRTMG single or multi-threaded? This should be indicated.

Response: The execution of RRTMG for the purposes of performance evaluation were single-threaded. We mentioned this in the text now.

Lines Changed: 244

4. More details should be included about how the time is measured. cProfile is used for calculate the execution time, but with times of only 0.31 ms, as indicated in table 2, the accuracy of the measure is not guaranteed, as a typical time granularity on Unix is 1 ms. Alternatives, as the time Python module, could provide more accurate measurements. Additionally, information about how many measurements have been done and how the final values have been computed should be included.

Response: We thank the reviewer for pointing out this omission. The execution time results are averaged from 10 measurements with execution of 100 000 predictions per measurement. The above text is added to the paper.

Lines Changed: 245–247

5. Introduction: To include a specific section with the related work, instead of including it in the introduction, could be beneficial

Response:

6. A recent work describing other GPU-based accelerating methods for the RRTMG_LW could be analyzed <https://www.mdpi.com/2076-3417/10/2/649>

C4

Response: We thank the reviewer for this reference. While they don't seem to achieve the same performance as previous GPU approaches, it is interesting nevertheless.

Lines Changed: 263, 346

7. Section 2.2. Is it realistic the perturbed dataset as has been created? Can be these values representative of a real situation?

Response: It is common to have random perturbations in the optical depth due to the presence of clouds, aerosols or horizontal transport of water vapour. In some sense, our perturbed dataset tries to capture this variability. However, we agree that random perturbations at all levels is an extreme case, and it can be taken as a very strong test on the ability of our models to generalize to new situations.

8. Section 4.4. In the single column model simulation, why the NN model F is used? Why not use model A or model C?

Response: The single column model uses a pressure level grid, which means that the pressure levels do not change during the course of the simulation. For this reason we use model F which interpolates the pressure levels onto a fixed pressure grid.

9. The developed code used TensorFlow, this should be indicated in the paper along with the version employed. It could be also interesting to describe the benefits of using this library in comparison with others, eg PyTorch.

Response: Yes, the implementation can be done in other mainstream machine learning library such as PyTorch, Keras, etc. They could also introduce slightly different speedups. We chose TensorFlow because it is well-supported by Google and it is a mainstream machine learning platforms, especially when we started

C5

this project. The purpose of this paper is to demonstrate that NN is capable of predicting radiative transfers with significant speedups. I believe this conclusion is still true when implementing the code in other libraries.

We have added the version employed in the manuscript.

Lines Changed: 172–173.

C6