<u>Review of Solid Earth - Discussions manuscript:</u>

Deep learning for fast simulation of seismic waves in complex media

by Ben Mosely, Tarje Nissen-Meyer and Andrew Markham

Review by Andrew Curtis, University of Edinburgh, UK.

Summary:

This manuscript presents interesting work which attempts to create useful methods for modelling and inverting acoustic waveforms for 1D (depth-varying) media, with or without the addition of an offset fault. It tests two architectures of neural network and shows that they are capable of modelling the waveforms to within a small error tolerance relative to the data magnitude, and contributes towards showing that similar networks might be used to invert recorded waveform data.

To my mind this manuscript presents really interesting results in the field of acoustic modelling of layered media, and a taster of potential future results for waveform inversion. It certainly makes a contribution to scientific research. However, in its current form it does not explain how it contributes to ***Solid Earth*** research. The manuscript also lacks an overview/introduction/review/discussion of much past research, and as such it is not clear whether the authors are aware of whether or how their work improves on past work in the area of Geophysics/solid earth research. The latter comment is particularly apparent with respect to waveform inversion: published work from the 1990's appears to achieve almost as much as this paper (this may not be true, but the authors need to explain why).

I also agree with most (but not all) of the comments from Reviewer 1 (Andrew Valentine), in particular his comment that a fair comparison for the layered-medium example would be between neural networks and other modelling methods that intrinsically assume a 1D Earth structure rather than full finite difference methods. I will therefore not repeat his comments here. However, I outline below a couple of places where I would perhaps add to his comments.

Best wishes to the Editor, Reviewer 1 and to the Authors,

Andrew Curtis.

Comments:

1. The main worry I have about this manuscript is a combination of (a) the method is acoustic and uses relatively simple acoustic models, (b) the authors themselves state that it will be difficult to extend the proposed methods to more complex models and particularly to solid (elastic) models, and (c) that these two points seem to imply that this method is not, as the authors propose, an advancement on solid Earth science but rather is a paper about toy problems in acoustics that can not be extended to elastic media. I am not saying that this is definitely what the authors believe, nor that it is true, but that is how their message came across to me as a reader. If true (and I would not actually be surprised if it is), this suggests that the paper might be more appropriate for an acoustics journal like JASA, rather than a solid Earth journal. Otherwise, the authors must better justify how this work advances solid Earth science.

It may be that the authors think that these methods could be extended to elastic media and to significantly more complex (heterogeneous) models, and they simply have not explained how. Or it may be that they think that this research closes off an avenue of research for solid Earth science which is useful to stop others from following – in effect they might decide to argue that they prove that this approach or network architecture will not be fruitful for the Earth sciences. Either is a positive step for science. One way or the other, the authors need to explain more clearly which (or which other message) we should take from their work, and why.

However, since this is a Discussion paper, in fact I think the best would be to rethink the paper slightly: I would begin by thinking through, and then presenting, a roadmap that *might* solve real modelling or inversion problems in the solid Earth sciences using models of more realistic complexity. Then show how this paper fits into that – either taking a first step in a possible direction to achieve it, or testing an avenue that while successful, turns out not to be practical for the real Earth. Either way, in the Discussion section they can explain how this work has advanced our state of knowledge about the overall strategy, and how we should move forward in future.

2. The introduction is interesting and reviews some of the appropriate material, but is very sparsely justified, and as also stated by Reviewer 1, it does not include many key references. In my view, every sentence of a scientific work must either be a logical deduction from previous text, must have been deduced/proven in another paper, or may be an argument based on the material in another paper; in that latter two cases that paper needs to be discussed and cited. In this paper, the Introduction cites very few references and therefore contains unjustified (in the sense of, 'not justified') statements. Examples include:

   One cannot write a paper on using neural networks to perform full waveform inversion (FWI) without citing Roeth and Tarantola (1994 – J. Geophys. Research). How does the FWI part of this paper improve on their work? That is not at all clear. There are many other papers using neural networks for imaging in Geophysics using waveforms or other types of data; you need to read and cite them, and describe how this work advances the field relative to those works. Currently the latter is not clear. The authors make the case for using neural networks for real-time applications – again first steps in this direction have already been taken (see Cao et al., 2019 – Geophysics, for example) and should be discussed.

3. The authors promote the fact that they use 'deep learning', and their application certainly fits into that category. However, they must at least discuss why this is a positive feature of the method, and cite previous Geophysical applications of deep learning to support that discussion. Deep learning is usually defined to be the use of 4 or more layers within a neural network. While I agree with Reviewer 1 that his previous work (Valentine and Trampert) was an example of deep learning, the first that I know of in Geophysics was in fact Devilee et al., (1999 – J. Geophys. Res) – which came from the same university.

   In my view there is therefore nothing new about the concept of deep learning: we were using it in Geophysics in the '90's. What *has* changed is the extent to which depth can be used to impose useful structure on networks (as the authors themselves have done in this manuscript – their Figure 9, and also in the paper cited by Reviewer 1); also the number of parameters that can now be used (the width of each layer) has increased hugely. In fact the number of parameters in the authors' application is relatively modest compared to some in machine learning literature, but is certainly comparable to other recent studies in Geophysics; the structure that the authors impose is both sensible and clearly useful in order to help to obtain stable results. These things should be discussed.

4. The paper appears to have committed the equivalent of an 'inverse crime'. If I understand correctly, the authors have trained networks in the forward and inverse directions using models with a certain parametrisation, and have tested the networks on models of exactly the same parametrization (Reviewer 1 touched on this too). While this would be reasonable if the models were themselves reasonably realistic, in a practical field like Geophysics I think it is necessary to test the networks on examples that lie outside of the range of the training set – not only using different models from those in the training set, but models that are not within the span of the algorithm used to generate the training set (as in the real Earth).

   For example, (a) Earp and Curtis (2019, arXiv – https://arxiv.org/abs/1907.00541) and (b) Earp et al. (2019, arXiv – https://arxiv.org/abs/1908.09588) perform (a) probabilistic travel time tomographic imaging, and (b) probabilistic surface wave inversion for averaged shear wave structures with depth, using deep neural networks. The test examples using in both cases are created using a finer parametrization than was used in the network training set – thus the actually structure of the (synthetically) 'true' Earth is not attainable by the networks; nevertheless, they can be used as a useful check of whether in such cases (as in the real Earth) the networks behave sensibly – giving results that are spatial averages in some sense of the 'true' structure. To be clear I do not think that the above references are perfect in this regard and could certainly be improved (e.g., could use even more complex models for tests); nevertheless the authors could usefully think about such tests for their work as it would strengthen the conclusions.

Although the above may be read as being rather critical, I would like to be clear that I do like this paper, and I think that it could create a usefully contribution to the field and be accepted for publication in Sold Earth. I just think that as it stands now, it is framed as a useful contribution to practical Solid Earth Science, yet it does not explain whether or how that is achieved. It may be that only a reframing of the paper is necessary, but I wonder whether a little more actual research is needed too in order to fulfill the papers own promise.