Response to reviewers for the manuscript:
"Automated shear-wave splitting analysis for single- and multi- layer anisotropic media"

Key:
Reviewer comments – black
Author responses – red

Summary:

We thank the reviewers for such thorough reviews of the manuscript. Their comments have definitely improved the work, especially regarding clarification in a number of areas. Aside from clarification points, the reviewers highlighted two main points that were particularly important to address: (1) that the new multi-layer splitting method presented will not be applicable for teleseismic data; and (2) that they would appreciate some benchmarking of our results against other popular software packages. Regarding this first point, we very much appreciate reviewer A bringing this to our attention. We have now looked again in detail at the teleseismic case and indeed are in agreement with the reviewer that our method would indeed not be applicable for any teleseismic scenario that we can currently envisage. As such, we have changed the text to reflect this, adding a section on the limitations of both our multi-layer method and multi-layer methods generally when applied to the deep earth. This doesn't stop us presenting the new multi-layer method, since it is likely applicable for a number of settings that exhibit strong anisotropy for high frequency seismic wavefields (e.g. glaciers, volcanoes, hydrocarbon reservoirs etc). Regarding the second point, we now provide a direct comparison to the majority of major shear-wave splitting packages already used by the community for two end-member examples: an icequake and an SKKS phase.

Further details on all our responses and associated changes to the manuscript can be found below, as well as in a file containing all the tracked changes. Note that line numbers correspond to the tracked changes document. We hope that the paper is now in good shape for publication.

Reviews:

Reviewer A:

I have read and reviewed "Automated shear-wave splitting analysis for single- and multi-layer anisotropic media" by Hudson et al. They present a new software package coded in python to facilitate shear-wave splitting analyses from teleseismic to icequake data, therefore aiming to be the only tool one would need to study the anisotropy of a medium one is interested in. They also present a new method to facilitate multi-layer splitting which likely often exists in the Earth's interior and offer a full automatization. In my opinion, a software package for shear-wave splitting analysis for all types of earthquake data in python as long been missing. Accordingly, I am very happy to see this manuscript describing one.

However, there is a severe misconception for their new multi-layer splitting method: the method they propose claims to be able to remove i.e. the effect of crustal anisotropy from teleseismic data. Unfortunately, as far as I understand the method, it is unable to do that since it requires the anisotropy to be strong compared to the dominant period of the phase (which the authors explicitly state in the conditions), so that a complete split of wave trains exist. This is not the case for teleseismic data and will rarely be the case for i.e. slab earthquakes and possibly even crustal seismicity depending on their frequency. Accordingly, their method is likely only relevant for the case they discuss here, which are icequakes (another example I could think of would a very high frequency earthquake in the presence of very strong anisotropy). As someone who routinely uses teleseismic and local earthquakes to apply shear-wave splitting analysis, I assume the average user looking for a shear-wave splitting software will not be working on icequakes. Accordingly, I would ask the authors to discuss their method with a very clear disclaimer on the applicability of their method (which I guess won't be applicable in 95% of datasets). To remedy this situation, the authors could implement another 2-layer method as done in other software packages. This would allow the user to chose either one of the methods as needed.

This is an excellent point and one that we hadn't sufficiently comprehended previously. We are grateful to the reviewer for prompting us to think about it further. We are now in agreement with the reviewer that the various conditions/assumptions of our method make it likely not applicable for isolating crustal anisotropy from teleseismic signals for the vast majority of, if not all, situations. We have therefore removed any suggestion that this is possible using our method. We have added a very clear disclaimer at various points in the manuscript, as well as an entire subsection (Section 3.4.3) discussing the challenges of multi-layer splitting measurements and the limitations of our method. We debated including the more traditional 2-layer method included in other packages (the apparent or simultaneous multi-layer splitting method), but do not implement it in the software itself. The reason for this is that we want the package to focus on providing measurements for single ray-paths and the other multi-layer method in the literature requires many ray-paths to perform the inversion. Hopefully the reviewers are happy with the philosophy (now stated clearly in the introduction as "Here, we describe SWSPy, a new, open-source software package for shear-wave splitting analysis, specifically to accurately and efficiently measure shear-wave velocity anisotropy for individual earthquake-receiver ray-paths." (L62-63)). Further details of specific changes made to reflect the reviewer's comments can be found below.

As far as I can judge without having looked at the source code, the software itself seems to work well. From the manuscript, it is sometimes hard to understand e.g. how time windows are chosen, what the input needs to be, whether this code has a graphical user interface, whether it requires user-interaction etc. I found the online documentation for this code which answers some of these questions, but it would be beneficial if the manuscript was more communicative about it.

To strengthen the manuscript further, I would like to see a better comparison between existing packages and this one. This is particularly useful when methods / selection criteria etc. are different (or the same for that matter). Furthermore, if authors want to appeal to a general audience interested in shear-wave splitting, I suggest to keep the icequake examples to a minimum and emphasize teleseismic or crustal applications where possible.

Overall, I would certainly like to try out this new software package myself in the near future and suggest that after revisions, the manuscript will make a good contribution to Seismica.

These points are all addressed in detail below. We very much welcome anyone to try out the new package, indeed that is why we developed it and made it open-source. If/when the reviewer does try out the package, we very much welcome any feedback on usage etc.

Line by Line comments

Line 29: Please add here that fast shear-wave then travels parallel to the anisotropic fast axes of the medium and that the slow wave is orthogonal to that and that this leads to significant energy on the transverse component of a seismogram.

Added a sentence to communicate this.

Line 33: I would argue the predominant usage of shear-wave splitting is to infer mantle flow regimes. Please add some examples from the lower & upper mantle.

Good point. We've now added the following text: "In the mantle, one can infer mantle flow in both the upper mantle (Hein et al., 2021; Fontaine et al., 2007; Long et al.…), and the lower mantle (Reiss et al., 2019; Creasy et al., 2021; Wolf et al., 2022; Asplet et al., 2023), as well as image shear and mineral transitions…" (L34-37). Hopefully that gives readers a flavour that mantle flow can be imaged using shear-wave splitting.

Line 49: For better reasoning, please dedicate a short paragraph on previous packages before you describe your own. This gives you the opportunity to describe the (long) history of shear-wave splitting packages e.g. for local seismicity Teanby et al., 2004; MFast - Savage et al., 2010; & Spingos et al., 2020; for core-refracted phases SplitLab by Wüstefeld et al., 2007 & extension for multi-event analysis by Grund et al., 2017; SplitRacer (Reiss et al., 2017) and extension to full automatization by Link et al., 2022 (I hope this list is complete, but these are the most important ones that come to mind). Afterwards, you can describe what you aim to accomplish with your package compared to the others (i.e., advantages of coding language, applicability, multiple layers etc.).

Good idea. We initially didn't include this as we didn't want to weigh down the paper with a comprehensive review, but can see the value of a short summary. We have now included this as suggested (L51-60).

Line 66: Since you also refer to a different method later, maybe you could quickly summarize here the main methods such as eigen values, energy minimization, correlation rotation and splitting intensity. This would provide a clearer distinction of what is method and what is your implementation in the code. Also, Silver and Chan are the only paper I would cite when talking about the eigenvalue method, as the other ones are simply implementations of the same concept (maybe simply indicate this?) except for Walsh et al., 2013 who actually found an error in the original formulation of Silver and Chan.

We have now added a brief description of each method and cited the eigenvalue method more appropriately.

Line 72: A flow chart of processing steps would add to the readability here.

Added a flow diagram (Figure 3).

Line 85: I would be very careful here. Any upsampling may give the false impression of 'more' accuracy, but you can never provide results more accurate than the original sample rate. I would ask you to rethink if that option should be available at all.

Yes, good point. We now emphasise this in the text: "upsampling… …will still be fundamentally limited by the sampling-rate of the native data, so should be used with caution" (L114-115). We still include it as an option for the user since it can allow for more precise measurements, even if they are not more accurate. Hopefully that addition is sufficient to address this point.

Line 96: Maybe replace 'situations' with 'settings' or 'cases'

Replaced with settings.

Line 100: While you define LQT and ZNE, you never define BPA (you use it in the caption of Figure 2); here the reader will question what the A- and B-planes are. I see you are referring to Walsh et al., but you cannot expect the reader of this paper to look up their naming scheme.

Thanks, that is a good point. We now define these planes/components where they are first introduced in the text and in the caption of Figure 2.

Line 110: How do you define a null-measurement? Or what's the codes criterion to detect such?

Good point. We have now added a line defining a null-measurement ("where one can be confident that no splitting is observed") in Section 2.1.5. Here, we also detail the code's three preferred metrics for quantifying the quality of a measurement. These metrics are by no means perfect, and so we hope that others in the community will help contribute to the development of further metrics going forwards for quantifying the quality of shear-wave splitting results.

Line 107: It seems odd that the first time you refer to Figure 3 it is to subpanel f; most journals will expect you to discuss figures in order of their appearance. If subpanel f is important here, then please move it to its own figure if you don't want to discuss the rest of figure 3 here.
Fair point. We've just removed the sentence that refers to it, as it is unnecessary and obvious what the phi-dt space is when one comes to look at the examples. Apologies for that. The original reason for this is that we wanted to keep the figures as true to the actual software outputs as possible.

Line 112: How do you choose any window you mention in this paragraph? User interaction by clicking on seismograms, a fixed time from an S-wave pick, a statistical analysis? E.g. Link et al. (2022) has a statistical analysis tool choosing the 'best' window. I completely agree that multiple windows are necessary but I don't understand how the code choses them.

This is a very good point and has been the source of most confusion for users to date. Thanks for raising it. The user specifies the window positions relative to the S-phase arrival time. We now explicitly describe this in Section 2.1.4, ensuring complete clarity by also providing a figure defining the various required parameters (Figure 3 in the updated text). Hopefully it is now clear how the windows are defined.

We have also cited Link et al. (2022) in Section 2.1.4 and pointed to their example of another, perhaps better but more specific method for automatically defining the windows.

Line 113: SplitRacer (Reiss et al., 2017) and its extension (Link et al., 2022) also use a multi-window analysis, the first only assuming the statistical distribution, the latter using a boot-strapping approach Line 131: How is this different or similar to the automatization by Link et al., 2022?

Their approach (if I have understood correctly) is to use spectral analysis to automatically correct for predicted XKS phase arrivals. Our method assumes that shear-wave arrival times and their associated uncertainties are adequately defined. We have chosen not to focus on improving S-wave phase arrival time picks, as we deem that to beyond the current scope of the package. However, future users may wish to add this functionality via an SWSPy submodule in the future. We now highlight this in Section 2.1.4.

Line 177ff: What is missing in this section is a clear distinction between the existing methods which are based on inverting multiple waveforms from differing azimuths at the same station (Özalaybey and Savage (1994; also see the code implementation of Reiss et al., 2017 or Link et al., 2022) and the suggested method here which works on only one waveform at a time (which brings its own set of limitations – please discuss here or later).

Apologies for not communicating the distinction between our method and that of others. There is a fundamental difference in the philosophical approach of our method to that of others. Our method assumes independence of each source-receiver ray-path from every other path, and seeks to measure how multiple layers of anisotropy affect that individual ray path. Özalaybey and Savage (1994) invert for two layers simultaneously using a theoretical relationship between layer parameters and apparent splitting parameters, This is for each independent source-receiver ray-path, but is non-unique and so they get around this issue by inverting for many sources at many azimuths arriving at the same receiver, aka stacking (hopefully the text already clarifies their method). They do also investigate stacking multiple results for a single station. Reiss et al. (2017) do something similar, using multiple waveforms from differing azimuths, in an inversion I would liken to a simplified implementation of anisotropy tomography. There is a fundamental difference in approach between using individual measurements vs. stacking. While there is much value in these methods, one requires suitable data (many phase arrivals with similar incidence and differing azimuths) such as teleseismic phases, whereas our method can be applied to any data. The other reason why we think our method is novel and worth sharing is that results from our measurements can be straight-forwardly used for anisotropy tomography in a similar context to, say, travel-time tomography. It would be more challenging/limited to do this with stacked data since much of the directional information has already been used in the initial inversion and "ray-paths" no longer exist in the same way for a tomography.

We now hopefully communicate this point in the text, explicitly stating/bounding our aims for the software package and highlighting the differences from other methods. (See Section 2 for an updated clarification). We also now include an entire subsection (Section 3.4.3) of the results dedicated to emphasising the challenges associated with multi-layer splitting and how our method breaks down for teleseismic type data.

Line 180: I agree with the fact that an inversion with more degrees of freedom is less constrained but whether it is 'poor' or not has to be judged on the merit of its individual outcome.

Fair point. We've now changed "more poorly constrained result" to "less well constrained result".

Line 188: Please insert a paragraph here to also convey stylistically that this is a new element.

Thanks, good point. We now start this as a new paragraph.

Line 202: Unfortunately, this rules out all applications for multiple splitting analysis of core-refracted shear-waves, because an average dominant period of a core-refracted phase would be somewhere between 8-12 s (with longer periods possibly still contributing) while most dominant mantle layers cause a delay time of about 1 s.

Yes. That is unfortunate. We now include an entire subsection (Section 3.4.3) of the results dedicated to emphasising the challenges associated with multi-layer splitting and how our method breaks down for teleseismic type data. We attempted to try and overcome these limitations but came to the conclusion that it is just too unconstrained a problem with the assumptions we make for our method needed to break the non-uniqueness. While we do not include another multi-layer method (for the philosophical reasons described in a later point), it would be straight forward for users to use SWSPy as part of a multi-layer multi-source apparent splitting inversion as described by others in the literature. We now briefly make this point in Section 4.1.

Following your method, you propose to 'cut' the time series into two parts and perform the analysis successively - this will only work for very small periods / high frequencies if the anisotropy is much stronger. E.g. Rümpker & Ryberg (2000) Figure 1 shows a nice summary of how the particle motion of multiple split shear-waves look depending on frequency. Individual wave trains of multiple split shear-waves are only clearly separated if the frequency of the wave is high (= small period). It is only then the particle motion shows a ‘cross’ (wave trains clearly separated) and not an ellipse. Accordingly, I don’t see how you could apply your multi-layer method to teleseismic data; or even local seismicity if we assume the crust hast a contribution of a few tenths of seconds the most (e.g. 0.2 s for a lower layer, the earthquake S-wave dominant period has to be higher than 5 Hz).

Thanks for emphasising this point about our multi-layer method not being applicable to teleseismic signals. Indeed, we had not thought carefully enough about the applicability for such a case and overemphasised the multi-layer method applicability in the text. We now more strongly emphasise the limitations of our multi-layer method in the text, with a specific subsection now dedicated to this (Section 3.4.3), as well as summarising this in the limitations section of the discussion (Section 4.1). Thanks for pointing us to the literature. I think the reviewer is actually referring to Rumpker and Silver (1998), which we now cite to emphasise how important it is for the multi-layer method that the fast and slow shear-waves from the deepest layer must have delay-times greater than the dominant period. However, we still hold to the statement that there are a number of physical scenarios for which the multi-layer splitting method is relevant, since it is proven to at least work for a synthetic case and a glacier example.

Line 204: This is such a specific criterion, it rules out many cases. For comparison, the ‘traditional’ 2-layer approach does not have that problem.

True. However, the ‘traditional’ 2-layer approach requires using multiple source-receiver ray-path pairs to perform an inversion to derive the multiple layer properties from the many apparent splitting measurements. Our intention in this study is to focus on single source-receiver measurements, in order to provide maximum flexibility for subsequent analysis (e.g. tomography). We now address this in the text, explicitly describing our approach and the reasons for it (L62-63, L73-74, L246-250, L590-593). This does not preclude users using the package to calculate apparent splitting parameters and setting up their own inversion, which we’d describe as effectively a 1D tomographic inversion. We just decided not to focus on any tomography/multiple ray-path stacking requirement for any of the algorithm. Hopefully now that we explicitly explain the limits of our scope, that is acceptable for the reviewer. We very much thank them for their comments and did seriously consider implementing such algorithms directly in the package.

Line 213: Your criteria are not only stringent, they assume a complete split of the waveform which happens so rarely that your application on icequakes may be one of the only ones it will work.

Maybe this is the case. However, given that the multi-layer splitting algorithm is only one component of the overall paper, we think it is ok to present it with its limitations and let the community explore its applicability. We think that it will likely be applicable in some other strongly anisotropic environments, e.g. shales and volcanoes, although we have yet to test it in such settings. Hopefully the single real example is sufficient to at least evidence its merit in the real-world. Glacier seismic anisotropy is in itself a growing field and so even if it is only applicable in that environment, it will still be useful for a number of studies going forward.

Line 220: How do you chose the two windows?

We've made the choice of windows more explicit by adding the text: "controlled by the apparent delay-time, δt_{app}".

Line 244: In conclusion to your multi-layer analysis, it would be more useful if you also implemented a computationally 'heavy' two-layer method based on an apparent inverse splitting operator from two individual layers as given by Özalaybey and Savage (1994) or Rümpker & Silver (1998) and implemented in e.g. SplitRacer (Reiss et al., 2017; Link et al., 2022). This would allow the user the choice between both methods and opt for the computationally heavy one in all the cases where your analysis will not be possible (which I assume will be the vast majority).

Thanks for this point. The code actually has the main components of such a method already implemented. However, although we really value these methods, we still choose not to include such a method here as it involves using multiple source-receiver pair measurements, which fundamentally changes the nature/premise of the package. We've previously mentioned the logic/philosophy behind this. Hopefully that is acceptable for the reviewer.

Line 259: I suggest screenshots of python code belong in the user manual.

We think that Seismica encourage these code snippets for papers that detail software. Hopefully it is ok to keep them in, although we are happy to take them out if needed.

Line 314: What are P and A components? (This goes back to missing explanation earlier)

They are now defined earlier in the text (where first introduced and in the coordinate definitions figure).

Line 325: If your code relies on the clustering, or if it is substantial for the interpretation, then please state shortly here (or above) how it works.

It isn't particularly substantial for the interpretation, just one metric to provide an idea of whether a result is "good" or not. However, we have added a little text here to provide the reader with some additional tips on what they should look for. Hopefully this clarifies things. There is also additional information on the clustering method in the methods section (albeit heavily referencing Teanby et al. (2004), since we don't see much point in repeating their detailed description of the method).

Line 346: 'core-refracted'

Thanks. Change made.

Line 369: This is a very interesting section and I agree with the authors that this analysis should be done more often.

Thanks!

Line 409: In your model, do you assume that between anisotropic layers is an isotropic layer? The time delay between both split phases seems rather large. In any case, would the proposed method still work if anisotropic layers were immediately stacked ontop of each other?

No, we assume an immediate transition. However, as the reviewer says, the method is approximately ambivalent to the presence of an isotropic layer. The time does look quite large, but this is an effect of the source-time function used. If one looks at the peak-to-peak of the various arrivals, from the first peak they are 0.2 s, 0.5 s, 0.7 s, respectively for each subsequent peak. These correspond to fast/slow, slow/fast, slow/slow for layer-1/layer-2 (with the first peak being fast/fast). Hopefully that makes sense. These times are correct.

Line 431: I appreciate this example where the proposed two-layer method clearly works well; however, I would suggest to use an example that is not that highly specialized, because very few users of this package will be using it on icequakes. An example from a volcano or subduction zones would convey the applicability much better.

We'd have to explicitly find another example first. While we'd love to present a number of examples, we have yet to explore for other situations where we observe these phenomena. Hopefully we can leave that as an area of future work. Our aim with this paper is simply to show that the method works for at least one real-world example.

Line 463: Please explain why a shallow velocity gradient is important.

Thanks for pointing this out, as we meant "steep" rather than shallow. We have now changed the text to: "allowing SWSPy to likely be useful for measuring anisotropy using borehole data or settings without a significant steep velocity gradient that refracts waves towards vertical incidence".

Line 469: Maybe reference to the existing anisotropy tomography codes from i.e. the group of Maureen Long or Sebastien Chevrot?

Good point. Added more references here.

Line 472: As outlined above, this will not be possible with the method you propose. A core-refracted phase split by mantle anisotropy and again in the crust will never completely separate so that you could apply your method. Please rephrase or only keep it if you add another two-layer method.

Have rephrased to: "A further advance provided by SWSPy is the ability to measure multi-layer anisotropy *under certain conditions*". We have also modified the entire paragraph to remove any implication that the method is applicable for teleseismic phases.

Line 483: Link et al.,2022 outline how one can tune these parameters in an automatization

Indeed. Their automation approach is similar to ours. We now point to this paper at this point in the text as it nicely backs up the use of such metrics for automation for a range of SWS applications.

Line 488: One's limitation is another ones' gain! I think a python package for general SWS splitting in python was long overdue and nicely fits within the seismological community's usage of python as the go-to language. There are very few Julia packages for seismology and I don't see the majority of seismologists wanting to use C.

We very much agree. We feel we had to put some such statement in to appease the developers of packages in such languages, which do have their own advantages.

Figure 1: Maybe I misunderstand the sketch, but to me it looks like the fast shear-wave (red colors) has propagated less (so is slower) than the slow shear-wave as indicated by blue colors. Shouldn't the red colored fast shear-wave always be on-top of the blue, slow shear-wave? Is this a problem of the perspective? Please make sure this is correct.

Thanks for drawing our attention to this! You are indeed correct. I think we must have become so absorbed in making the sketch that we got the colours wrong. We have now rectified the issue. Thanks for pointing it out.

Figure 3: I know the main author has significant experience with ice quakes, but the audience of this paper will likely not. Accordingly, I would suggest to use a more general example, i.e. from a local earthquake.

The reviewer makes a good point here. However, we'd prefer to keep the icequake example in as it is such a clear example of anisotropy and flows nicely into the multi-layer method results. In other figures we show results from a range of other settings. Hopefully that is ok.

Figure 4: Is there a reason you don't show the results of the cluster analysis here?

The reason is that we hope that the reader has picked up the concept of what the cluster analysis looks like from the previous figure. We don't typically use the clustering result plots regularly for interpretation of good splitting results and we didn't want to over clutter the remaining figures in the paper, instead choosing to focus on the results that best demonstrate a particular point. Also, all the figures are reproducible via notebooks within the package examples, where the clustering analysis for any event can be seen. Hopefully that is ok.

Recommendation: Revisions Required

Reviewer C:

Review of Hudson et al., 2023

The authors of Hudson et al. (2023) present a new Python software package to conduct shear-wave splitting on any source-receiver combination and any S wave. They add new features not currently used in other shear-wave splitting codes, such as parallelization, multi-layer measurements, S-wave polarization, and clustering analysis. I think this code would be beneficial to many seismologists who conduct shear-wave splitting. Installation was easy and simple. I was also able to run one of the examples from the code and there were no errors. However, I do have some minor concerns that the authors should address to appeal to potential users of this code.

We thank the reviewer for both their comments on the manuscript and on the code itself. We address their comments in detail below.

Recommendations for description of the code:

Onboarding to HPC is suggested in the paper but not expanded on. The authors should supply information or some examples of how to run this code on HPC.

This is a fair point. The reason for not doing this is that HPCs are all different, and although we could provide an example for perhaps the most popular HPC job management system (SLURM), it would likely not be exactly applicable for any individual HPC. However, we do now include a section in the documentation about how one would go about running the software on a HPC (via an anaconda environment managed by SLURM). We have also reworded the manuscript text to be more general about how the software could be used on a HPC ("and is parallelised so can maximise the potential of modern computers and High Performance Computing (HPC) architecture" (L65-66)).

Line 356 – a more thorough comparison with SHEBA would be preferable rather than one SKKS measurement. The authors should benchmark their code with another research group's results. For example, the authors could choose a station from Wuestefeld et al. (2008) using Splitlab or Reiss and Ruempker (2017) - SplitRacer.

We now include a section comparing to other packages (see Section 3.5).

Also, the online documentation is quite limited. This is not required for publication, but end users would appreciate more details in the documentation of how to implement their own datasets. For example, more description is needed for preferred file formats, earthquake meta data, and station meta data. The online documentation I am referring to is in the README.md file: "Full documentation for the package is available [here](https://swspy.readthedocs.io/en/latest/)."

Thanks for the feedback. We have attempted to improve the documentation and continue to do so as more users provide feedback on the software and its usability.

Questions about methods:

Why do the authors only use the eigenvalue method for splitting? I think the authors could expand in the future on other shear wave splitting methods.

We use the eigenvalue method since it is typically a relatively stable method and is probably the most widely adopted method. However, we have also implemented a cross-correlation method for calculating certain quality metrics. At present, we don't plan on actively supporting other methods, but given the modular nature of the code, it would be relatively simple for other users to contribute other/new methods to the package. We have now added a few lines of text communicating this to the reader from the outset of the paper (L93-96).

Since cross-correlation method is used for the quality factors, why is it not included in the splitting analysis?

Good question. It would be trivial to include in the outputs. We deemed it would be confusing for inexperienced users to have to deal with so many outputs. Also, in practice we don't always calculate the quality factors as it involves passing around large arrays of data, which can be memory intensive. I guess this is a subjective point, and we can see the advantages of passing such data. Hopefully the reviewer is happy if we don't pass this data for now, but that we could actively add an option to pass such data if we receive multiple requests for such data outputs in the future.

Other questions/comments:

Line 157 – the direction of fault slip is not necessarily aligned with the fault slip. This can occur over certain azimuths. I think this sentence should be corrected or be more explicit on the S-wave polarization patterns. S-wave polarizations can be calculated based on forumlations in Aki and Richards textbook. I recommend the authors comparing their polarizations to equations in Aki and Richards (2002: equations 4.89-4.91). The authors could include predicted S-wave polarizations and compare to their measurements. I think this would make their statement more robust and grounded.

Our understanding of Eq. 4.89 to 4.91 is that they describe the directional variation in displacement. For a pure double-couple source, the S-wave source polarisation (fault slip) doesn't vary with direction of radiation, with the radiation pattern only describing the variation in amplitude of the S-wave. In other words, the 3D S-wave source polarisation remains constant regardless of azimuth and inclination angles, with only the scalar amplitude across all three components varying by the same amount. This holds theoretically, but also is evident in the data we present. If it varied with azimuth, then one would expect a lot of scatter on the source polarisation rose plots. Hopefully that clarifies things.

Equation 1 – the authors should specify what phi is. I believe it is fast-axis polarization, but it could be confusing to a reader.

Thanks for pointing this out. Although the parameter is introduced earlier in the text, we agree that it would be helpful to define it explicitly here too. We have now done this in the text.

Line 136 – what are the alpha values? I couldn't find where they were specifically defined in the paper

They are defined at this point in the text as the uncertainty in the relevant labelled/subscript parameter. Also, we follow typical convention that alpha generally signifies uncertainty, unless otherwise indicated. However, to aid clarity in the text, we now use brackets () rather than commas in the text here. Hopefully that makes it clearer.

Figure 3 – what is the P and A amplitudes? Is that P-wave? I didn't see a clear definition in the text. I don't these are common terms that most readers would know what A and P amplitudes are.

Good point. P and A amplitudes are the amplitudes in the polarisation and null directions, respectively. We now clarify this in the caption of Figure 3.

Figure 7 – the text on the side is small and of low image quality. I suggest making that text larger. I don't think it would look well resolved if the paper was printed. Most of the figures have this small text with some important information. I would suggest most of the texts in the figures need to be bigger.

Good point. We've made the text in this figure bigger, and where appropriate throughout the paper. Some of the other figure text is still small. This is because we wanted to plot the direct outputs of the code at some points. The amount of information in these summary figures inevitably results in small text. Hopefully the figures are all now clear though.

Recommendation: Revisions Required
