Review of gmd-2024-207: "JuWavelet – Continuous Wavelet Transform and Stockwell-transform for gravity wave analysis" by Jörn Ungermann and Robert Reichert

Overview

In this study, the authors introduce and describe a new Python package called JuWavelet, which implements both the 1-D, 2-D and 3-D continuous wavelet transform (CWT) using the Morlet wavelet and also the widely-used *S*-transform. These tools are widely used in the geoscience for, among many other applications, the analysis of atmospheric gravity wave perturbations in a variety of geophysical datasets, such as those from satellite, radar and other remote sensing techniques.

By creating this simple and efficient software package, the authors have provided a valuable tool for the community, doubly so because it is coded in the open-source Python language. Their description of the analysis processes and its mathematical foundations are well articulated and their test applications are thorough. The figures are of very high quality representing 1-D, 2-D and 3-D test cases on synthetic, modelled and observed geophysical datasets.

Overall I recommend this study for publication after the authors have considered my comments below, which are minor but I think they would improve the paper by more accurately clarifying technical aspects, introducing some of the geophysical concepts, better referencing existing methods and literature, and helping the interpretation of results by future readers.

I have "General Comments" and "Specific Comments" which it would be nice if the authors could respond to, and then "Minor Typos and Suggestions" for which I do not require a response.

General Comments

1. Journal Scope: I notice that this article is submitted to Geophysical Model Development, which at first was a little surprising to me as it is perhaps better suited to Atmospheric Measurement Techniques? However, I will leave the editor to decide whether this is the right EGU journal and I certainly do not want to trigger another review cycle with AMT for the authors, as the paper is, in my view almost ready for acceptance.

2. The paper describes an analysis method specifically focused on the analysis of gravity waves in geophysical datasets, but doesn't actually mention what gravity waves are, how they are typically detected, why we would want to measure them, or any of the atmospheric science aspects of gravity wave remote sensing, detection or analysis. This is fine if the paper is *only* describing a generalised wavelet software package (which I guess JuWavelet is), but because it goes further to apply it to gravity wave measurements (as is mentioned even in the title) the manuscript should at least give some kind of overview of atmospheric gravity waves in modelling and observations and what parameters we are interested in. It's obvious to me as a gravity wave scientist but this information is essential for non-specialists who might not be familiar with gravity waves, or who might want to apply JuWavelet in other fields, and I think would strengthen the narrative of the paper. The authors could also briefly mention some other potential applications for JuWavelet beyond gravity waves, if they wish (I'm sure there are many possible uses in the geosciences).

3. Amplitude underestimation in the *S*-transform: As the authors know, for an infinitely long time series containing an infinitely long sinusoidal wave, the *S*-transform is able to measure the instantaneous amplitude of this wave perfectly. However, in the atmosphere gravity waves almost always occur in small packets containing maybe only a few wavecycles. In this case, when analysed with the ST, the convolution of the gravity wave packet's "envelope" with the Gaussian window in the ST results in an underestimation of the wave amplitude in the ST coefficients at that frequency due to spectral leakage, even if the wave is monochromatic. This underestimation is derived analytically for a hypothetical monochromatic gravity wave "packet" analysed by the 1-D, 2-D and 3-D *S*-transform in the Appendix A1 of Hindley et al. (2019). They found that for 1-D the effect is negligible, but it can be quite significant for higher dimensions. It can also be mitigated by adjusting the scaling parameter, similar to the authors' "free parameter k" (see below). Although the authors briefly mention this effect in Sect. 4.3, it would be really interesting for the authors to discuss (a) to what extent this amplitude underestimation happens in their application of the *S*-transform and (b) if they have any ideas or methods that they could use to counteract or avoid this underestimation in their approach. It's no problem that it occurs, it should just be mentioned in the paper a bit more clearly and I would be very interested to hear the authors comment on the issue and if they have any ideas or solutions.

4. How does JuWavelet deal with the ambiguity of wave directions in 2-D/3-D data? Specifically, when a wave packet has k_x, k_y and is ambiguous with a wave with $-k_x, -k_y$, what does JuWavelet measure? Does it limit only to analysing for angles of 0 to π or similar?

Specific Comments

1. **l.15 and elsewhere:** "Stockwell transform" - Minor point, but the *S*-transform is technically just called the *S*-transform (Stockwell et al., 1996), the S does not stand for Stockwell, strictly. But I appreciate the community often refers to it as such. Some journals even require it written *S*-transform using italic, based on their journal style.

2. **l.15:** "flavour" - Some authors have debated over the years whether the *S*-transform (ST) is actually a modified type of CWT (Gibson et al., 2006), or whether is is a localised version of the Fourier transform, particularly in its discrete form (Stockwell, 2007). I have no feeling either way but the authors should mention that it has some concepts (like absolute referenced phase) which are quite different from a CWT, and that the ST "wavelet" does not have zero-mean so it is not strictly admissible in the CWT, but this is minor semantics. From a practical application point of view in the geosciences, one of the major differences between the ST and the CWT is that the coefficients of the ST can be directly interpreted as wave amplitude, whereas the CWT coefficients more closely resemble wavelet power (although I acknowledge this is mentioned later in the manuscript regarding Fig. 2). These are all subtleties and semantics, so we do not require a full exploration of these points, but the authors should ensure that they are consistent and mention some of the key differences that would be useful for the readers who might apply JuWavelet to their data.

3. **l.19:** The authors mention the 2-D *S*-transform of Hindley et al. (2016) but neglect the 3-D *S*-transform of Hindley et al. (2019), the code for which is actually N-dimensional and applies the 1-D, 2-D, 3-D and 4-D *S*-transform. I understand the narrative justification for JuWavelet though and there is definitely need for a multi-dimensional CWT/ST package in Python, but the authors could be more complete with their references to the literature. The authors should also probably mention the S3D method described by Lehmann et al. (2012) and Ern et al. (2017), if nothing else perhaps to simply avoid

forgetting their colleagues at Jülich! Actually, mentioning S3D works well in the narrative because JuWavelet overcomes the "cubes" limitations of S3-D and (possibly?!) the discretisation limitations of the *S*-transform of Hindley et al. (2019), but more on this below.

4. **l.40:** "atoms" - not clear what atoms are in this context? Do they authors mean axioms? Even so that would not be quite accurate. Maybe just use "functions".

5. **Sect. 3:** "free parameter k" - If I understand correctly, this is a very nice description by the authors of how (in their formulation) the *S*-transform basis functions are similar to the Morlet wavelet but with the free parameter k set to 2π. It would be useful to mention, for consistency with Stockwell's original formulation (and to help readers who are less familiar!), that this is simply the same as scaling the Gaussian window in the ST with a standard deviation equal to $1/f$, where f is the analysing frequency. For example, I'm not sure I understand correctly whether k is the value that scales the Gaussian window with frequency, or whether it is some multiple of that frequency? k also does not appear in the Eqn on l.90 so perhaps it would be clearer for the reader if this was written explicitly, please rephrase.

6. **Sect. 3:** Further to the above point, it would be useful if the authors also discussed the effect of adjusting the ST's free parameter k by some multiple of 2π to achieve improved spatial (spectral) localisation at the expense of spectral (spatial) localisation. Other studies have experimented with this depending on their geophysical dataset to achieve the desired results, such as Fritts et al. (1998); Pinnegar and Mansinha (2003); Hindley et al. (2016, 2019). Typically these studies describe this adjustment as a scaling parameter c such that the Gaussian window scales as c/f, referring to Stockwell's original formulation as mentioned in the previous comment. Does JuWavelet have this capability (I think it does) and did the authors experiment with it? In any case, it would be useful to mention this capability and when it might be applicable, to help users who may want to apply JuWavelet to their data. For example, Hindley et al. (2019) found that setting $c = 1/4$ (or $c = 4$, depending on whether the ST is written in the spatial or spectral or spectral domain) achieved improved spatial localisation in 3-D analysis of AIRS satellite observations of gravity waves. Because Hindley et al. (2019) only considered the dominant wave at each spatial location x, y, z, it did not matter very much if the spectral peak was broadened because the peak was still in the same location in the spectral domain f_x, f_y, f_z. Do the authors find the same with JuWavelet? It would be useful if the authors commented on this, and if they varied this free parameter when they analysed their 1-D, 2-D and 3-D datasets, and if they found improvements over previous methods, which I expect they do.

7. **Sect. 4:** Figure 3 is discussed in the text before Figure 2, consider rearranging.

8. **Sect. 4.3:** It would be nice to include a nod to Hindley et al. (2016) for Fig. 5, given the very close resemblance to their 2-D test case. It would also be nice to include a line plot of wavelength in versus wavelength out, and amplitude in versus amplitude out in order to assess the capability of JuWavelet for different wave scales in a dataset such as this. For example, Hindley et al. (2016) and Hindley et al. (2019) both showed that there is not a perfect 1 to 1 measurement for all wave scales for the ST in a test like this, but I would interested to see if the CWT mode of JuWavelet recovers the wavelengths perfectly, or if it has discretisation limitations like the Hindley et al. (2019) *S*-transform. If not, that would be a major strength.

9. **l.160-161:** This method of collapsing the 4-D spectrum by selecting the dominant frequency at each amplitude follows the approach of Hindley et al. (2016) and Hindley et al. (2019) for 6-D (also l.214), so it would be nice to include a reference because I don't recall seeing this method as common practise in the geosciences before these papers. Feel free to disagree, it's not essential.

10. **l.165-166:** Somewhat related to the point 6 above, the authors mention here that adjusting the Morlet parameter (is this k?) "can deliver more accurate amplitudes at the expense of spectral resolution". I'm not sure I fully understand the Morlet parameter then. I had naively thought that the setting $k = 2\pi$ was equivalent to running the analysis in the *S*-transform mode (see the Eqn. on line 90), where the standard deviation of the analysing Gaussian window is $\sigma = 1/f$. I also hadn't considered that there could also be a scaling parameter in CWT mode that could adjust the spatial-spectral resolution, as is done for the ST. It would be very useful for the reader if the authors could explain this better, especially if a user is going to be using JuWavelet, they need to be aware of exactly how these options work.

11. **l.167-169:** I'm not it's enough to say that "the wavelengths/directions of the synthetic packets are well produced", the authors should be more quantitative. The easiest way is including a simple table or extra panel in Fig. 5 that shows the amplitude, wavelengths, phases, directions before and after the JuWavelet analysis.

12. **l.167-169:** Further on this point, as mentioned above does the JuWavelet formulation suffer from the discretisation limitations that one encounters when using e.g. the *S*-transform as derived from the discrete Fourier transform (DFT)? An *S*-transform based on the DFT (such as that first written in code by Stockwell) is very fast but is limited to discrete frequency voices, and might struggle to accurately measure low frequency waves like synthetic wave #8, which has a relatively large wavelength compared to the physical size of the image. Does JuWavelet also have this limitation, or is this solved by the derivation based on the CWT? If not, then this should be mentioned in the manuscript as it is a major advantage of JuWavelet (depending on the associated runtime).

13. **Sect 4.4:** Figure 7 is mentioned before Figure 6, please consider rearranging.

14. **l.177:** I'm intrigued by the *aspect* parameter. Is it not possible to achieve this result by setting the range of horizontal and vertical scales of the analysing wavelet (or scales and angles) to a given range, or does JuWavelet not have this functionality and it's better to stretch the data to get a more 1:1 aspect ratio of the waves? Also, does the stretching process actually resample the input data to make the JuWavelet input data larger and more "square", and does this have an effect on the runtime or memory requirements? As I said, I'm intrigued.

15. **l.180:** I'm not sure what is meant by "please see panel titles in App. B" and how this is related to the authors point. Do the authors expect the reader to run the code in order to read these - I am guessing the panel titles they are referring to are $f\lambda_x = \text{period[idx[0]]/np.cos(theta[idx[1]]):3.0fkm}$ and $f\lambda_z = \text{period[idx[0]]}/(\text{cwt['aspect']*np.sin(theta[idx[1]])}):3.1f\text{km}$.
    I would say that having these written in Python in the Appendix is not sufficient for the reader in terms of explaining how to recover the horizontal and vertical wavelengths when setting *aspect* $\neq 1$. The authors should provide a clear equation in the text of how to recover these wavelengths when using different values of *aspect*, unless the JuWavelet package automatically calculates this? The reader should not have to either read or run the Python code to understand how this is done, it feels like the authors are cutting corners a little bit in the description. Maybe my Python is rusty but I also couldn't work out what the 3.0 and 3.1 were referring to.

16. **l.183:** "filtered" - can the range of scales or angles in the CWT be filtered *before* the transform is computed to save runtime?

17. **l.187:** The authors should provide a little more information on this clustering algorithm, I think it's a very powerful addition to the software package to be able to detect e.g. the most important N waves

in a given image using this clustering algorithm, and could have widespread use in the geosciences. Just a little more information about which algorithm is used, how it is applied and what exactly it clusters - scales/angles, amplitudes or both?

18. **Sect. 4.5:** Figure 7 is mentioned before Figure 6, consider rearranging.

19. **Sect. 4.5:** The use of a segmentation algorithm with JuWavelet to decompose, segment and then reconstruct the overlapping mountain waves is powerful and impressive! It would be interesting to apply this to the synthetic wave field in Fig. 5, or for cases where waves have more similar wavelengths than the mountain wave example shown, which seems a little easy. No need to include such an example, but a quick discussion of its strengths and limitations would be appreciated. On that point, what would happen if the the input data contained two gravity wave packets with very similar wavelengths and angles that were separated in distance and in opposite corners of the image? How would the segmentation algorithm cope with that case in the spectral domain? Would the two waves be recorded as one wave?

20. **l.209-211:** The azimuth and zenith angles seem quite coarse, is this to save runtime? Also, regarding the aspect ratio - if one can specify the input scales and angles, why is the aspect ratio required to stretch the data in a given direction? I apologise if I have not understood correctly. Perhaps this could be clearer in the manuscript.

21. **l.223:** "can be simply realized" - somehow I feel like parallelising the analysis of individual scales on the same input data like this is possible, but it might not be quite as "simple" as the authors suggest! It would interesting to include some kind of information about runtime for JuWavelet, although I appreciate this is relative to hardware and input. Even so, some ballpark numbers for runtime, or typical numbers of scales and angles, or some general advice on best practise would be appreciate for readers deciding if JuWavelet could work for their data. For example, should users always be prepared to have enough memory to generate a 6-D CWT/ST object if their input data is 3-D? This is important to know when analysing, for example, high resolution modelling output in 3-D. It would be useful there is an option to avoid the creating the 6-D object in memory and select only the dominant wave at each spatial location to output 3-D objects, for example.

22. **l.227:** "using the Morlet wavelet." - I'm interested to know if the use of different wavelets, such as higher order complex Gaussian wavelets, could improve the amplitude estimation for wave packets in the JuWavelet CWT/ST analysis? No need to include this information in the manuscript, unless the authors are interested or they have some useful thoughts on the matter that might be worth including.

23. **Appendix A:** This may be more of a comment for the typesetting stage, but EGU journals are still two-column format as far I know. Therefore, the authors may struggle to show their 2-D and 3-D equations clearly, such as those on l.245-253. One solution could be to generalise the formulae to multiple dimensions by writing spatial and spectral vectors such as x, y, z and $\omega_x, \omega_y, \omega_z$ as $\boldsymbol{x} = (x_1,\ x_2,\ \cdots\ x_N)$ and $\boldsymbol{\omega} = (\omega_1,\ \omega_2,\ \cdots\ ,\ \omega_n)$, as was done in Hindley et al. (2019, their Eqns. 1 and 2) for an EGU journal. No problem if the authors are not concerned about the formatting, running the equation over multiple lines is also a solution.

24. **l.230:** This paper introduces the software package JuWavelet, but they don't actually mention where and how interested readers can download and use the package, or how it might be made available and under what conditions. In a method paper such as this, this information should definitely be included.

Minor Typos and Suggestions

1. **l. 25** "brevity's sake" -> "brevity"
2. **l.33:** "precise" -> "(or more precisely, a function)"

References

M. Ern, L. Hoffmann, and P. Preusse. Directional gravity wave momentum fluxes in the stratosphere derived from high-resolution airs temperature data. Geophy. Res. Lett., 44(1):475–485, 2017. doi: 10.1002/2016GL072007.

D. C. Fritts, D. M. Riggin, B. B. Balsley, and R. G. Stockwell. Recent results with an mf radar at mcmurdo, antarctica: Characteristics and variability of motions near 12-hour period in the mesosphere. Geophy. Res. Lett., 25(3):297–300, 1998. ISSN 1944-8007. doi: 10.1029/97GL03702.

P. C. Gibson, M. P. Lamoureux, and G. F. Margrave. Letter to the editor: Stockwell and wavelet transforms. Journal of Fourier Analysis and Applications, 12(6):713–721, 2006. ISSN 1069-5869. doi: {10.1007/s00041-006-6087-9}.

N. P. Hindley, N. D. Smith, C. J. Wright, D. A. S. Rees, and N. J. Mitchell. A two-dimensional stockwell transform for gravity wave analysis of AIRS measurements. Atmospheric Measurement Techniques, 9(6): 2545–2565, June 2016. doi: 10.5194/amt-9-2545-2016.

N. P. Hindley, C. J. Wright, N. D. Smith, L. Hoffmann, L. A. Holt, M. J. Alexander, T. Moffat-Griffin, and N. J. Mitchell. Gravity waves in the winter stratosphere over the southern ocean: high-resolution satellite observations and 3-d spectral analysis. Atmospheric Chemistry and Physics, 19(24):15377–15414, 2019. doi: 10.5194/acp-19-15377-2019.

C. I. Lehmann, Y.-H. Kim, P. Preusse, H.-Y. Chun, M. Ern, and S.-Y. Kim. Consistency between fourier transform and small-volume few-wave decomposition for spectral and spatial variability of gravity waves above a typhoon. Atmos. Meas. Tech., 5(7):1637–1651, 2012. doi: 10.5194/amt-5-1637-2012.

C. R. Pinnegar and L. Mansinha. The S-transform with windows of arbitrary and varying shape. Geophysics, 68(1):381–385, JAN-FEB 2003. doi: {10.1190/1.1543223}.

R. G. Stockwell. A basis for efficient representation of the s-transform. Digital Signal Processing, 17(1):371 – 393, 2007. ISSN 1051-2004. doi: http://dx.doi.org/10.1016/j.dsp.2006.04.006.

R. G. Stockwell, L. Mansinha, and R. P. Lowe. Localization of the complex spectrum: the s transform. IEEE Transactions on Signal Processing, 44(4):998–1001, Apr 1996. ISSN 1941-0476. doi: 10.1109/78.492555.