5th October 2021.

Author Response to RC2 for manuscript "Proper orthogonal decomposition of ice velocity identifies drivers of flow variability at Sermeq Kujalleq (Jakobshavn Isbrae)" submitted to The Cryosphere.

I, and on the behalf of my co-authors, would like to thank Dr Bryan Riel for their insightful and constructive review of our manuscript. These comments will undoubtably improve the accuracy and clarity of the paper, and hopefully make this first application of POD in this setting more convincing to the community.

We would like to specifically draw attention to two themes of the Dr Riel's comments which we see as strongly interlinked: that our results are influenced by the "travelling waves" recently described by Dr Riel in this journal; and that we should expand our study area to a wider area. By specifically looking at a small area we act minimise the first effect, in addition to minimising the potential effect of spatially variable velocity image coverage. The travelling wave has a reported speed of 0.4 km/day, yet study area is 4 km long, and the image pairs used are typically separated by 11 days. Hence, we expect any wave to travel entirely through our study area and not be resolved by the velocity product on this scale. The strong dominance of Mode 1 in explaining the dataset variance is, as Dr Riel points out, evidence that the area does at this scale broadly move "in phase". Our study area is ~11 km flowline-distance from the terminus most-retreated position where the effect, according to Riel (2021, TC) Figure 2, is becoming less pronounced. We do accept, however, that as the velocity product also uses overlapping pairs we might expect some contamination of the modes, but this is unlikely to have an annual periodicity and is therefore more likely to be accounted for in the higher (i.e., low variance) modes. In the revised manuscript we would endeavour to make this justification explicitly.

Please see below our detailed response to comments. The Reviewer comments are italicised and indented, and immediately followed by our response in normally formatted text.

Best regards,
David W. Ashmore

RC2.

Overview:

> *The manuscript "Proper orthogonal decomposition of ice velocity identifies drivers of flow variability at Sermeq Kujalleq (Jakobshavn Isbræ)" explores the application of POD to ice velocity time series in order to efficiently decompose time series into orthogonal spatial and temporal modes. This decomposition could potentially isolate velocity signals originating from distinct forcing mechanisms, such as changes in terminus position or enhanced basal sliding due to subglacial hydrology. These signals may or may not be readily apparent from the raw velocity time series, so POD could potentially reveal key signals to enhance our understanding of ice dynamics.*
>
> *Overall, the manuscript was well-written, the methods were clearly described, and the results and supporting data are quite intriguing. POD and its variants have always been powerful analysis tools for spatiotemporal data, and this paper demonstrates the utility of POD to the increasing availability of velocity data over the ice sheets. However, POD also has well known limitations, primarily due to the assumption of orthogonality of the modes. While orthogonality doesn't necessarily affect lower-order reconstruction of the data, it does leave users vulnerable to over-interpretation of the separate modes, and I believe the authors need to explore these vulnerabilities a bit*

further. I describe in detail the implications of orthogonality below. My other main comment is concerned with the error analysis, which I also describe below.

Thank-you for your generally positive assessment of our manuscript, and we are pleased that you consider our findings intriguing and that we demonstrate the utility of the POD technique in this new setting.

I do believe that POD will be an important tool for glaciologists. I also believe that potential users should be well aware of the expected behavior of POD-type algorithms on different classes of dynamical signals. If the authors can sufficiently address and discuss these issues, then this work will indeed be a valuable contribution.

Thank-you for this assessment. In the submitted manuscript we tried to not present POD as a silver bullet to disentangling processes from observed spatio-temporal ice dynamics that can be applied without caution. We explicitly discuss the splitting of modes throughout and make recommendations for future work to focus on how modes evolve over specific glaciological features (L439-443). In our detailed responses and in a revised manuscript we address these issues and, in general, highlight more clearly the features of the POD approach, and how these must be considered during interpretation.

Major Comments:

1) Currently, my biggest concern is the application of POD to signals that propagate in space, which encompasses most of the signals related to ice dynamics. Taking the Sermeq Kujalleq (SK) data presented here as an example, the largest signal (Mode 1) is clearly related to the dynamic response of the glacier to stress perturbations at the terminus. However, previous work has shown that this response has a finite phase velocity, i.e. it takes some time for the signal to propagate upstream (~400 m/day; Riel et al., 2021, as a recent example). The fact that POD enforces every point in the modeling domain to have identical temporal eigenvectors (scaled by the spatial coefficients) means that every point has the same phase for Mode 1, which was stated by the authors. The consequence of this is that POD will end up splitting the spatially-propagating signal into multiple modes (with the same periodicity), just as the authors hypothesized for the conjugate Mode 2+3 pair. It's entirely possible that Modes 2 and 3 are partially "conjugates" for Mode 1, i.e. they are compensating for the error of modeling a finite phase velocity signal with an infinite phase velocity signal. Additionally, the spectral content of an upstream-propagating signal can change, e.g. pulse broadening. All in all, any interpretation of the weaker modes for subglacial hydrology is likely to be corrupted by the propagation of the compensation signal for Mode 1.

Please also refer to the top of this letter where we discuss that: (a) our POD domain is small (4 km long) relative to the 0.4 km/day travelling wave; (b) the image-pairs used to derived the velocity image are typically 11 day apart allowing enough time for the travelling wave to pass through, if present; (c) the POD domain is 11 km flow-line distance from the most retreated terminus position where the effect is less pronounced.

The Reviewer is indeed correct that Modes 2 and 3, that we hypothesis are dominated the glacier hydrologically processes, *may* contain some information on the kinematic wave. However, this is an inevitable ambiguity when assigning physical meaning to purely statistical features. We see no evidence for "leakage" from Modes 1 to Modes 2 and 3. In fact, we suggest the opposite, that the Mode 1 is "contaminated" by a hydrological signal at L308-310,

We are careful not to state that this pair of modes is exclusively due to glacier hydrology, rather that we present evidence that the modes are "dominated" by it. It is precisely that POD modes are purely statistical features that we attempt to build the case that they are dominated by glacier hydrological processes through comparisons to independent datasets and are very explicit about this in the manuscript (L249).

We agree that the manuscript should explicitly discuss the travelling wave phenomena recently described by Riel et al. (2021) in text. We particularly value the Reviewers suggestion “(c)” to add a discussion of how different “classes” of dynamic signals expected on glaciers may appear in POD datasets.

> *In general, the splitting of a spatially-propagating signal into multiple modes will be dependent on the phase velocity. When the phase velocity is very high, as would be the case for a dynamic wave resulting from redistribution of longitudinal stresses, then the splitting will be minimized. However, for slower moving waves, as would be the case for kinematic waves or subglacial hydrological effects, then one would expect the signal to be split into more modes, which would hinder any subsequent interpretation of these modes.*
>
> *With that said, the supporting data for the hydrological interpretation for Modes 2+3 are compelling (and very cool). If the authors can find a way to account for the leakage of Mode 1 into the other modes to isolate the hydrology signal, then the interpretation would be more convincing. Therefore, I would suggest the following initial diagnostic steps and improvements:*

These points are responded to the previous comments and the following comment “(b)” on the distribution of the eigenvalues.

> *a) Expand the spatial extent of the POD analysis (e.g., Figure 3) so the readers can see a wider picture of how the spatial coefficients change in space. I think the spatial extent shown in Figures 2b-d would be sufficient. This way, we can get a better sense of any mode splitting due to signal propagation.*

As outlined in the top of this letter, the issue with looking at a large area is that the signal then becomes difficult to interpret due to spatially moving signals being split over many modes. Furthermore, Sirovich (1987, Quarterly of Applied Maths) highlights that the number of pixels shouldn’t be significantly more than the number of timesteps, and so we can expect POD performance to degrade as domain size increases significantly relative to timesteps.

> *b) Include a plot of the singular values for the decomposition. I know these values were included as labels in Figure 3a, but I think a plot can give a better sense of any potential mode splitting. Slowly propagating signals will be split into a higher number of modes, and one would see a more gradual decay of the singular values. Rapidly propagating signals will have a much sharper spectral decay.*

This related to the point made by Reviewer 1 on the “percentage of variance” explained by each mode, as this is the common method of communicating the relative importance of modes in other POD studies. Briefly, in version 1 of this manuscript we discuss the relative importance of the singular values, not the eigenvalues. This resulted in some imprecise wording to be rectified in a revised manuscript. From the eigenvalues we can describe the “percentage of variance” explained by each mode: a common diagnostic metric for orthogonal decomposition work.

The eigenvalues show a strong dominance of Mode 1, explaining 98.1% of the variance. Despite the small *Var(%)* values of Modes 2 (0.87%) and 3 (0.45%) we do believe they contain useful information: indeed, the kink between 2 and 3 indicates that they are “paired” in some way. Arguably modes 2 and 3 cluster outside of normal, smooth pattern of decay: evidence that they contain meaningful information. Note also that the clustering of eigenvalues is only one of two common diagnostic criteria for identifying quasi-conjugate pairs, the other being the phase-shifted features between Modes 2 and 3 and the strong seasonality of their sum which we show in Figure 4.

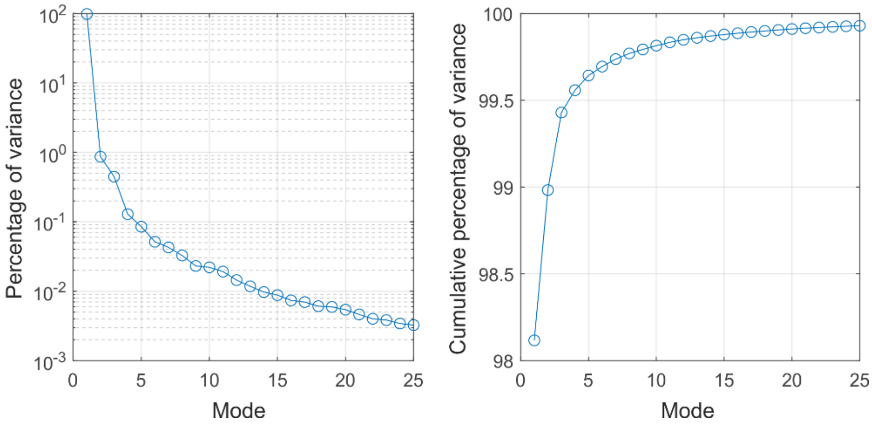

c) For the purposes of a methods paper, it would be hugely beneficial to compare the decomposition between SK and a glacier with known hydrological effects that dominate its dynamics. I would expect quite different spectral decay. Since that would entail a lot of additional work, alternatively, the authors should add some discussion on how different classes of dynamic signals will have different phase velocities and how those phase velocities would affect the POD results.

We do certainly take on board the idea that a short discussion about how one might expect different "classes of dynamic signals" and how they might exhibit in POD modes would be a worthwhile addition to the paper. In a revised paper this would be included in Section 2.1 and referred to in the Discussion and Interpretation. We are currently working with new collaborators on comparable SAR-derived velocity fields from other glaciers with different styles of hydrological behaviour.

2) The error analysis performed here is a bit too simplistic. Mainly, the errors are analyzed as if they are independent from pixel-to-pixel. However, it is possible for the errors to have some non-zero spatial correlation (off-diagonal terms in the scene-wide data covariance matrix), which could be quantified empirically via a variogram over areas that are known to have minimal ice velocity. A more robust analysis would be to generate random realizations of the velocity data (perhaps with covariance matrices with non-zero off-diagonal terms) and re-do the POD in a Monte Carlo sense. How would the spatial coefficients change with the different data realizations? Since POD is computationally efficient, this likely wouldn't take too much compute time.

A related point is also raised by Reviewer 1. The Reviewer is correct that there *may* be complex correlations in space and time in errors, however, it is also important to note that the published errors (10-20 ma^{-1}) are very small relative to the seasonal velocity fluctuation of -1000 to +1500 ma^{-1} and in the submitted manuscript we only assign glaciological significance to modes after comparison to independent datasets.

If I understand the variogram method proposed correctly the issue is that that in our POD domain we purposefully selected a dynamic area with no areas with minimal ice velocity. We do not take a Monte Carlo approach because the stated errors are variable from one image to the next, and so would certainly be "noise" in the data and accommodated in the low-variance Modes.

To address this comment, this we turn also to Reviewer 1's comment on providing an "error envelope". We run the POD again with the errors added, and then with them subtracted. In the below figure we find the POD Modes are insensitive to the published SAR-derived velocity errors and variations in modes are imperceptible when interpretated at our scale of interest and plotted on the same axis.

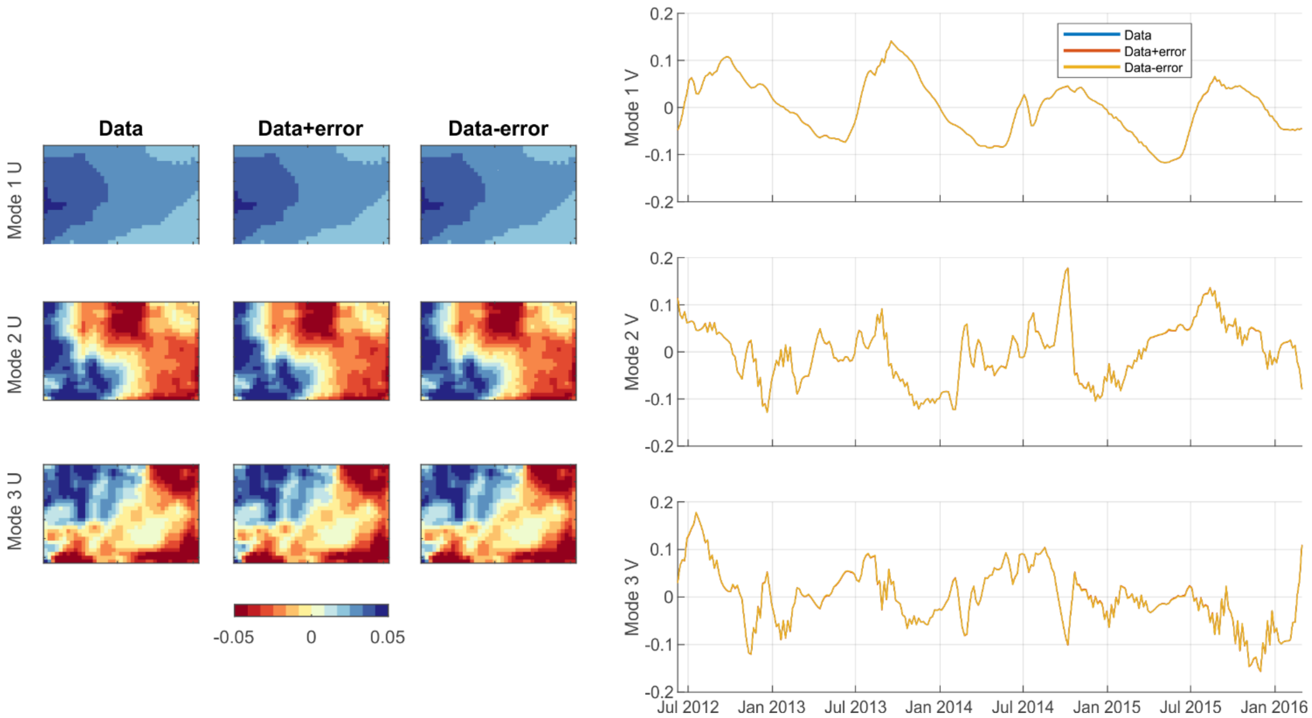

Line-by-line Comments:

- Line 39: Might be too soon to state the modes correspond to "distinct flow features".

Accepted, this statement comes from hydraulic applications of POD where this more appropriate. Will remove.

- Line 77: The distinction between eigenfunctions and eigenvectors can be a little misleading since an eigenfunction is instantiated as an (eigen)vector. In other words, they are both vectors, and my intuition is that the temporal vectors are closer to being instantiations of actual eigenfunctions. I would suggest to either replace the use of eigenfunctions with eigenvectors or to add a bit more detail into why the use of eigenfunctions is more appropriate.

Accepted, change to eigenvector.

- Line 79: What is meant by a "fully converged dataset"?

We use an autocorrelation approach (to ensure a sufficient number of integral length scales), this shows signals repeat sufficiently for POD analysis.

- Line 81: Orthogonal does not by definition mean independent. Achieving statistical independence (e.g., minimizing mutual information between modes) tends to be the objective of Independent Component Analysis and can result in modes with very different spatial and temporal eigenvectors. More importantly, orthogonality does not guarantee that independent forcing mechanisms will be expressed in the different modes.

Agreed, we can change this to spatially coherent and remove "by definition" which is potentially misleading.

- Line 99: Is there any indication of the numerical algorithm being performed by the "economy" SVD? Is it performing a full SVD and just removing a certain number of singular values? Or is it performing some form of randomized SVD? This information may be useful for reproducibility purposes.

We expect the economy SVD to have no effect on results. The approach is designed by Mathworks MATLAB, and from their documentation: *"The economy-size decomposition removes extra rows or columns of zeros from the diagonal matrix of singular values, S, along with the columns in either U or V that multiply those zeros in the expression A = U*S*V'. Removing these zeros and columns can improve execution time and reduce storage requirements without compromising the accuracy of the decomposition"*

> *- Line 124: I think it's a little confusing to state "the study area must be dynamic" since that may be interpreted as only dynamic time series can be decomposed with POD. I think what's being stated here is that dynamic glaciers tend to be well-observed by satellite platforms, so I would suggest some slight re-wording.*

Our argument is that to a POD is more likely to reveal statistical structure from an area with fast and variable flow, such as the trunk of a major outlet glacier (cf. a low-melt ice sheet interior). We will reword to make this clearer.

> *- Line 141: Is there really a requirement for the time series to be stationary in order to decompose them? What happens if non-stationary time series are used? I ask mainly because multi-year signals can be quite important for understanding the full spectrum of ice dynamics, and it would be interesting to see how POD handles those.*

The time series must be repeatable in term of its coherent signals, if not there is a strong risk of losing coherent structures in higher modes that are noisy and impossible to interpret.

> *- Line 145: A bit awkward to describe a slowdown with an "increased rate". It would be clearer to just state the decrease in velocity values after 2016.*

These are different statements; the velocity is already decreasing but does so at an increased rate. Will reword for clarity.

> *- Line 165-166: Are the formal data errors accounted for during the inpainting and interpolation procedure? If not, then every point is assigned equal weight, which could affect final interpolated signal.*

No, but the errors 10-20 ma^{-1} are very small for annual fluctuations of between -1000 ma^{-1} and +1500 ma^{-1}. As mentioned in text a maximum of 4 (of 232) measurements are interpolated over in a small area of the domain.

> *- Line 167: Is there any difference between post-smoothing the columns of V and pre-smoothing the time series of X?*

Thanks for this point. Yes, there is a difference, our approach follows the robust methods fully justified by Higham et al. (2016, Measurement Science and Technology). We will add this short justification in text.

> *- Line 195: What is the resolution for the "high-resolution" MODIS analysis?*

MODIS is daily and posted at 250 m pixel size. We can state this in-text.

> *- Line 220: Here's the main issue: the results are stating that the entire area moves in-phase, but previous results show the phase changes with upstream distance, with a finite phase velocity of ~400 m/day.*

In general, this point is rebutted in the detailed comment above. We state here that the Mode 1 velocity move in phase: this is essentially just a description of the POD algorithm output. This is an important distinction from saying the whole glacier *itself* moves in phase. Naturally, these travelling waves must be accommodated by the higher modes if resolved in our study area – but not Mode 1. While it is certainly important to acknowledge the limitations of the approach it is also important here to underline the stated aims of the paper: to apply technique to a new data type, and build a case through the comparison of POD to independent datasets that this technique shows something never seen before.

> *- Figure 3: It would be useful to include a length scale bar in the images.*

Okay, can include.

> *- Line 286: "we sum of Modes 2 and 3" -> "we sum Modes 2 and 3"*

Thanks, amended.

> *- Line 289: What are possible artefacts? If one were to interpret a splitting of a spatially-moving, seasonal signal into multiple modes as artefacts, then those would also exhibit strong seasonality.*

Okay, can be more explicit here in text. We would consider a "split" spatially moving, seasonal signal due to glaciological processes to not be an artefact as this is precisely the signal of interest in a study like this.

> *- Line 425: Conversely, it would also be useful to apply POD to ice flow model outputs (forced with distinct mechanisms) to determine how well the different signals can be separated.*

This is an interesting approach, and a test of POD capabilities with synthetic data: assuming the model itself was sufficiently accurate.

> *- Line 440: Again, this is a key point: spatial patterns of ice flow will split into multiple modes based on the propagation speed of the signal. Slower propagation speeds will be split into more modes.*

Discussed at length previously. We underline again here that in the Conclusion we made the point that POD users must be very careful what temporal and spatial scale you do the analyses, at as will depend on the particular process you wish to investigate.

> *- Line 460: Consider re-use of the word artefacts here as well. Mixing of signals into multiple modes could be considered an artefact but is definitely not a minor effect.*

We discuss at length the mixing of Modes 2 and 3, and will add sentences considering the potential for Mode 1 leakage.

> *- Line 467: My intuition is that hydrological effects will have slower spatial propagation speeds than redistribution of membrane stresses (such as the response to terminus forcing), which would cause more modal mixing.*

Okay, can note this at this point.