Revisions According to Referee and Community Comments

Assessment of Sentinel-3 Altimeter Performance over Antarctica using High Resolution Digital Elevation Models (https://doi.org/10.5194/egusphere-2024-3054)

Joe Phillips and Malcolm McMillan

Thank you to all reviewers for their comments and suggestions, and especially for their patience with this response. We have included our replies within the corresponding text provided by each referee, highlighting them for ease of review.

Anonymous Referee #1

General comments:

- This study evaluates the performance of the Sentinel-3 SAR altimeter over the Antarctic Ice Sheet using the REMA DEM. It finds that 90% of acquisitions capture the POCA within the Level-1b range window, but performance declines in complex terrain, missing 24% of topographic variance. The findings highlight limitations in capturing backscattered energy over steep or rough surfaces and provide insights to guide improvements for future missions like CRISTAL.
- The study is using REMA as a baseline for applying a SVD analysis to generate self-consistent surface slope and roughness datasets for the Antarctic Ice Sheet. The S3 altimeter performance is assessed by evaluating different topographic regimes by analysing the relationship between surface complexity and altimeter data quality. Furthermore, a waveform decorrelation analysis is made to assess the along-track impact of topography on waveform decorrelation to understand how varying terrain affects signal quality and backscattered energy capture.
- The methods and results concerning slope and roughness are well-executed and insightful.
- The research is valuable, addressing a gap in the literature. However, it's disappointing that the authors did not utilize the latest Sentinel-3 processing (B005), which has been available for over a year and addresses some of the highlighted issues. While the conclusions might remain valid with the reprocessed dataset, the omission weakens the study, as this remains uncertain. To make the impact the manuscript deserves, I suggest redoing the study using B005; at the very least, the new processed data should be mentioned, and what impacts the new processing has on this study.

Thanks for raising this point. We agree that an analysis of B005 would be worthwhile – this was not performed in the original manuscript, as the majority of our analysis was undertaken prior to the release of B005; hence our reporting of results for B004. Given the reviewer's comments, however, and the time that has elapsed since our original analysis, we have – as requested – rerun the analysis using the new S3 Hydro-Cryo Altimetry Thematic Baseline Collection 005 Land Ice product (available at https://dataspace.copernicus.eu/). As with the initial manuscript submission, our analysis was performed for cycle054 for S3A and cycle035 for S3B and we have updated the statistics accordingly. We have also included additional text highlighting the key changes between BC-004 and BC-005 over land ice. Additionally, we have also re-performed the same updated analysis (given the methodological changes made to address other comments) for the BC-004 product, which has allowed us to include a new cross-comparison between the 2 product versions. We have therefore added supporting text, analysis and statistics describing the differences between the two versions.

- The technical depth of the manuscript is inconsistent. For instance, the detailed reviews of retracking, REMA, and OCOG may not be necessary. In favor of the authors, they do present the topics in a clear and accessible manner, making them useful for newcomers.

We acknowledge that not all sections had a consistent level of technical depth; this was intentional in order to provide comprehensive descriptions for what we judged to be the most important and relevant technical elements for this study. In light of this, and the positive comments subsequently made by the reviewers relating to this level of detail, we prefer to retain this as we believe it will be useful to many readers who may not be familiar with the technical details.

Specific comments:

Principal of radar altimetry:

- Page 2-3, Lines 55-73: While textbook examples exist, the authors provide an exceptionally clear and convincing explanation. It could be highly beneficial for newcomers, though its relevance to this study should be carefully considered.

Thank you for your appreciation of the care we have put into writing these sections. Although we agree that the level of detail goes beyond what is commonly found in papers these days, we would prefer to keep this text because, as highlighted by all reviewers, it provides a valuable explanation for newcomers, which we found otherwise lacking in the literature.

- P3: Surface Tracking: For your information, Several places over Antarctica have been changed to open-loop acquisitions in 2021 and 2024. It would be nice to see if that has had a positive impact.

Thanks - we are aware of these sporadic acquisitions, but believe that an analysis of Sentinel-3's non-operational open loop mode is beyond the scope of the current work, which was intentionally focused on assessing the performance in nominal closed loop tracking. This choice was motivated by wanting to (1) assess the operational performance of the mission and (2) provide a comprehensive assessment at the ice sheet scale, which could not be done for open loop. That being said, we agree that - as part of future work - such an assessment would be worthwhile, e.g. within the context of CRISTAL, and we hope that our study will provide a framework for how this could be done. Within our revised manuscript we have included a more in-depth discussion of OLTC and recent developments in Section 3.1.

Specify

- P4, L48 Retrackers: You mention OCOG as part of S3 ground segment. True, but S3 products also include the UCL ice retracker, a beta-retracker or model-fit retracker (Wingham and Wallies, 2010).

We agree that this is important to mention, and so we have included additional reference to the UCL ice retracker (Wingham and Wallis, 2010).

Data:

- A fundamental weakness is choosing not to use the B005 Thematic Land Ice product, released over a year ago and specifically designed for land ice. This product uses extended window processing optimized for rough or steeply sloping surfaces. Further information is available in the B005 Validation Report: https://sentiwiki.copernicus.eu/__attachments/1681931/S3MPC-STM_RP_0114%20-%20Reprocessing%20Campaign%20BC005%20Validation%20Report%202023%20-%201.1.pdf?inst-v=bd8109db-ca17-4318-bc73-caad47319d7e and the data is currently in review in Nature Scientific Data by J. Aublanc.

Please see our prior response relating to the use of BC-005 - we acknowledge the value of this new product and have undertaken major new work to incorporate it into our revised analysis.

- Resampling REMA to 200 m, why? Have you considered using ICESat-2's gridded products?

This resampling was originally done in order to reduce computational overhead (P7, L191), however, within this revision, we have decided to use the REMA V2 mosaic at its native resolution of 100 m, and so we do not now perform this resampling. Notably, this has led to updated REMA-related statistics throughout the manuscript.

Regarding the ICESat-2 gridded product, we selected REMA in preference, because it is formed from stereoscopic techniques, and therefore provides an inherently 2D product for computing slope and roughness. In contrast ICESat-2 requires interpolation between tracks, especially as across-track spacing diverges away from the poles, and as such is therefore less well-suited to resolving slope and roughness. We also favour REMA, as it contains fewer data gaps than ATL14. For visual reference, we have included an example figure comparing REMA to the ICESat-2 ATl14 product (https://nsidc.org/data/atl14/versions/1) (Fig. 1).

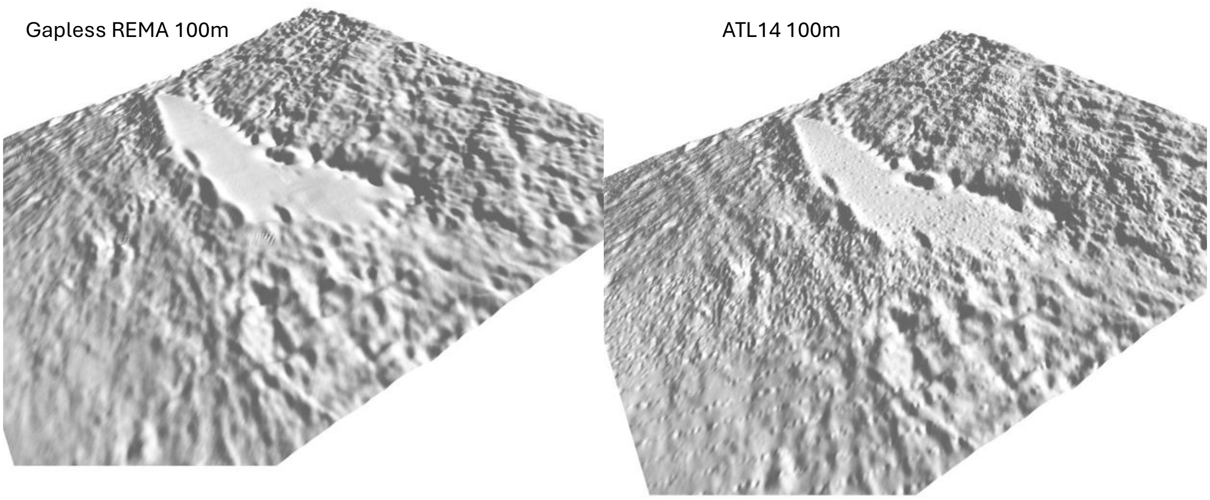

Figure 1.

We also refer to the ICESat-2 ATL14 ATBD (https://nsidc.org/sites/default/files/icesat2_atl14_atl15_atbd_r001_0.pdf), which states that "ATL14 and ATL15 are both limited as to the features they can resolve by (1) the spatial resolution of ICESat-2 tracks, (2) the temporal sampling of the tracks, and (3) the resolution of the grids chosen for the products."

We appreciate that our reasons for picking REMA were not clearly articulated, and so have also included additional clarification for this choice within the revised manuscript.

Methodology:

- P6, L238: Results interpreted as 68% of the roughness variance. Where does this number come from?

This value represented the proportion of points in each window that fall within ±1 standard deviation of the mean of the orthogonal residuals to the fitted slope, assuming a normal distribution. However, after discussion, we have decided to instead measure roughness using the difference between the minimum and maximum orthogonal residuals within each window. This is essentially an estimate of the peak-to-trough elevation difference, which we believe is more aligned with a reader's likely intuition related to this statistic.

- I am unfamiliar with the word seed waveform, but I assume it is a kind of a "first" waveform. How do you determine this seed waveform? Is it unclear if you calculate all waveforms as seed and then calculate the correlation to the 50 subsequent waveforms, or if the seed waveform is specially chosen?

We agree that this explanation was poor. Here, a seed waveform refers to the first waveform by which the correlation of 50 consecutive waveforms is compared against. This is computed for every single waveform, which is referred to as the "seed" waveform when it is the focus of these correlation computations. We have improved our explanation of how we calculated the along-track decorrelation of echoes, including replacing the use of the word "seed waveform" with easier to understand language.

Results:

- P13, L3556-358: 23.7% of echoes are not captured within 60m range window. Could you elaborate on the discussion? Could this be solved with the current S3 setup? Is it only a SWOT processing that would solve the problem? I am missing some consideration of the ability to change the on-board tracking modes for S3.

This statistic (now given as 57.4% of records) represents the proportion of records within the chosen cycles wherein the surface variability is such that, even with optimal placement (computed via maximising capture over 1 m incremental vertical displacements), the range window cannot fully capture the surface. This is in contrast to previously where we simply recorded the proportion of cases where the max–min difference of the elevation points within the beam-limited footprint exceeded 60 m. This change was made to account for the curvature of the range window, which introduces a vertical displacement of approximately 60 m at the edges compared to nadir. This result reflects a fundamental limitation imposed by the 60 m range window relative to the topographic variability, rather than an issue with onboard tracking.

- P14, L363: It would be nice to see how data processing baseline B005 performs. or at least discuss why it would still be an issue with the new dataset.

Please see our prior response relating to the use of BC-005 - we acknowledge the value of this and have incorporated it into our revised analysis.

- P15, L370: Please elaborate on what you mean about this statement

In order to successfully identify the leading edge of a waveform, it is important that there are a small number of bins proceeding it. Range window placement should consider this fact, allowing for a small buffer above the point-of-closest approach to the surface. We have revised this part of the text to make the point clearer.

- P16, L401-403: The results of LPCs are fine, but again this is not what was improved by S3 processing baseline B005?

We have included new, detailed discussion comparing the BC-004 and BC-005 products.

- Fig 5/6(b): It would be nice if the blue and purple colors were more distinct, but it is impossible to differentiate between the colors on the map.

The intention of these colour choices was to visually group classes that are above and below the range window as being equally unsuccessful cases of POCA capture. We agree that this colour choice makes the spatial figures difficult to read and so, as requested, we have revised our choices of colours for the POCA classes to make them more distinct.

- A along-track correlation study demonstrates that slope and roughness impact the result. We also see this in how the elevations are having trouble wrt. REMA or IceSat-2 results in the marginal zone. What do we gain from this method? Is this method better? It sounds like it takes a large amount of computational power.

We realise that the relevance of this section was not clearly conveyed in the original manuscript, and so we have revised this text to better clarify the purpose of this element of the study.

- How does S3 perform in this sense compared to other altimeter satellites?

We believe that the work involved in performing such a detailed assessment for other satellite would be considerable, and falls outside the scope of this current work, which was intentionally focused on developing novel analyses for understanding the performance of S3. We do agree that similar analysis for other satellites in the future would be worthwhile, and we hope that the approach developed and presented here for S3 can serve as a framework for similar assessments of other satellite altimetry missions in the future. We highlight this broader point within the revised manuscript, and we note that we do also now include a brief assessment of the impact of varying the size of the range window, which has relevance to other missions.

Technical comments:

- P6, L182: the dataset name appears to be incorrect ("S2_2_LAND__"); it should be "S3_2_LAN__".

Thank you for spotting this; we have corrected this, and updated the text to reflect the use of the new BC-005 thematic land ice product.

Anonymous Referee #2

The manuscript provides a very detailed description and history of radar altimetry, which is particularly beneficial for newcomers seeking to better understand the needs and various steps involved in processing the data. It also provides a good template for future analysis and mission instrument design.

General Comments:

In Section 2.3, a general description of different slope correction algorithms is presented, which is excellent for the reader. However, I find the section lacks a more definitive conclusion regarding how much better or worse the different algorithms perform. I recommend including more details here, such as why certain algorithms are better and by how much (e.g., statistical comparisons). Additionally, it would be helpful to describe how these algorithms improve elevation change estimates, particularly in areas such as outlet glaciers or regions with significant topographical variation.

We agree and have included additional text discussing and comparing the algorithms in more detail, as well as including quantitative statistics regarding their performance.

In Section 2.4, you discuss the geophysical corrections and provide ranges for the magnitudes of these corrections. It would be beneficial to add similar range information to the sections detailing retracking and slope corrections to maintain consistency and provide a more comprehensive overview.

We agree and have provided additional magnitude information relating to the retracking and slope correction sections.

In Section 4, there is significant discussion about using REMA, but it lacks detailed reasoning as to why REMA is particularly important. I've outlined specific comments on this below. Could you expand on why REMA is advantageous? Furthermore, if you intend to resample the DEM, I recommend avoiding NN (nearest neighbor) resampling when transitioning from higher to lower resolution. A more accurate approach would be to use the average of the pixels to avoid introducing bias.

We agree and - as outlined in our response to reviewer #1 - have included text to better justify our use of REMA. Within the revised manuscript we have also now opted to use REMA at 100 m resolution, thus avoiding any need for interpolation.

In Section 5.1, I understand the approach taken and the rationale behind it, but I would appreciate more details and statistics to support the arguments. For instance, if a more traditional approach (e.g., based on Horn or similar methods) were used, what would be the statistical differences compared to the method you propose? Additionally, what are the benefits of using SVD compared to a simple least-squares approach for generating slope and roughness parameters?

In terms of overall benefits, our motivation for using SVD is that it fits a plane by minimising orthogonal residuals, whereas the least-squares (LS) approach minimises residuals in the vertical direction. This distinction makes SVD particularly suitable for complex terrain, where minimising the true geometric (orthogonal) distance yields a more representative surface by reducing orientation bias. This is especially important for steep slopes, where vertical residuals tend to underestimate the true geometric deviation, leading to possible errors in slope estimation. Based on some preliminary tests comparing slope and roughness maps derived via SVD and least-squares, we found that the median difference and median absolute deviation difference in slope between the SVD and least-squares approaches was < 0.1 x 10^{-6} °. Roughness estimates showed a much larger median difference of 0.189 m and a median absolute difference of 0.130 m between the two methods. We expect this larger discrepancy in roughness is because for any given window, distances spanned in the horizontal (X and Y) directions span several orders of magnitude more than in the vertical (Z), so even minor differences in slope can lead to amplified variations in roughness. Although we found that the least-squares method was ~30% faster in terms of computational speed, we opted to use the SVD-derived maps due to their improved theoretical basis and fit quality over complex areas. This is quantitatively reflected in the coupling between slope and roughness, with the SVD approach producing a lower Pearson correlation coefficient (0.808) than the least-squares method (0.819), suggesting that the SVD solution was more effective at reducing algorithm-driven dependencies between the two variables.

To address this comment more explicitly, we have also included additional text in the manuscript to support our use of SVD over least-squares. We have also included additional text, statistics, and a new plot comparing our slope maps to those generated using Horn's method, and our roughness maps to those generated using the non-orthogonal Terrain Ruggedness Index.

Section 5.2 is somewhat unclear, particularly the last paragraph. Could you elaborate further on this section and include more details? As it stands, the explanation is too general to follow clearly.

We agree that this section would benefit from additional clarification and supporting text. The intention of this component of the analysis was to determine the proportion of the surface (as defined by REMA) illuminated by the beam footprint that falls above, within, and below the positioned range window for each record. We then aimed to identify the theoretical optimal positioning of the range window, which would maximise the proportion of the surface captured. By comparing the actual proportion of the surface captured to the proportion captured with an optimally placed range window (where the goal is maximal topographic capture), we obtain a performance metric that quantifies how well the current placement performs relative to the optimal case, whilst recognising the inherent limitations imposed by the range window size. We have revised the text to make these points clearer.

The font size in your figures needs to be increased, especially for the smaller inset figures, as they are difficult to read.

We agree and have increased the font sizes in the figures.

For Section 6.1, it would greatly enhance the manuscript to include a comparison to established slope and roughness algorithms or datasets, as you are presenting this as a novel approach. Consider creating a figure for traditional slope/roughness algorithms similar to Figure 3 to highlight the differences.

We agree and, as per our response to the related prior comment, have provided statistical comparisons of our new slope measurements relative to Horn's method, and of our roughness measurements with respect to the Terrain Ruggedness Index, together with an associated new figure.

Specific Comments:

- L48: Clarify why the use of REMA is advantageous for evaluating design choices. Is it due to its resolution, quality, or another factor? Be more specific.

As outlined in prior revisions comments, we have included additional text to better clarify our motivation for using REMA.

- L63: "Followed by gradual decay in the received power" – clarify the cause of this decay. What does it signify?

As requested, we have revised the text to better clarify the cause of this gradual decay in received power.

- L69: Backscatter is not the only parameter used to determine surface characteristics. Mention other commonly used parameters, such as leading edge width etc.

Agreed, we have added further text discussing the leading edge width, and trailing edge slope, along with additional citations.

- L85: Specify which DEM is used for Jason-2 or Sentinel-3A, including the resolution of the onboard DEM.

We have specified the onboard DEMs used for Jason-2, as well as the corresponding resolution. However, we could not find a source for the DEM used by Sentinel-3 over Antarctica during the commissioning phase. We are continuing to search for this information via our ESA contacts, and will add it in to a future version of the manuscript, if successfully identified. Related details are included alongside a more in-depth discussion of OLTC and recent developments in Section 3.1 as recommended by reviewer #1.

- L95: Provide additional references here; citing only (Quartly et al., 2020b) does not sufficiently encompass the breadth of prior work on retracking and scattering corrections.

Agreed, we have provided three additional references.

- L100: Add references to Curt Davis's work, particularly his "threshold retracker," which significantly advanced the field and built upon the ICE-1 tracker (OCOG).

Agreed, thanks for highlighting this omission. We have now referenced Curt Davis's threshold retracker.

- L105: Emphasize that "slope correction" is likely the largest source of uncertainty in the current processing pipeline, as corrections can vary by 1–100 m vertically and span kilometers across track.

Agreed, we have revised the text to emphasise that "slope correction" is likely the largest source of uncertainty in the current processing pipeline, including approximate magnitude values.

- L159: Elaborate on the decision to favor closed-loop tracking. What motivated this choice?

We have elaborated on this decision, as requested.

- L271: Are you evaluating the slope correction provided in the L2 product and comparing it with your own? If so, what are the differences, and why are they significant? Is the L2-provided correction insufficient for determining alignment within the 3dB beamwidth?

To clarify - the purpose of this work was not to develop or evaluate our own slope correction method, but rather to assess the performance of the window placement and L2 processing. Specifically, we first aimed to determine the extent to which the POCA location provided in the current L2 product falls within the beam-limited footprint and range window, and thus to assess whether the assumptions inherent to the L2 processing chain are self-consistent. In other words, we sought to evaluate whether the elevation measurement extracted via the leading edge of the waveform can reasonably be attributed to the provided POCA location. Secondly, we implemented an alternative REMA-derived approach, which was designed to allow us to assess the extent to which the window placement captured the 'true' POCA (as identified by REMA, and in contrast to the linear-slope POCA in the current L2 product). This allowed us to assess the extent to which the current window placement successfully tracked the closest point on the ice sheet surface. We have revised the text to better clarify the intention of this section of the paper.

- L292: Consider using interpolation for the comparison, as it accounts for gradients in the data. Although 200 m is smaller than the footprint, significant variations may still occur within a 1600x300 m altimeter footprint. Refer to Roemer et al. (2007), who suggest using a resolution of 10 m for improved location accuracy.

Whilst we agree that there are potential benefits to fine-scale interpolation when using this approach to estimate surface height with high precision (e.g. for resolving cm-scale elevation trends etc), we emphasise that this is not the purpose of our analysis here. Rather, here we are using REMA as the basis for a coarse-scale, discrete classification of the POCA location relative to the position of the range window in 3D space, and not as a relocation step. As such, we do not believe that interpolation to a higher resolution would significantly alter this classification, whilst it would introduce a substantial computational overhead. Therefore, due to the purpose of this element of the study, we have opted not to perform additional interpolation. As outlined in prior revision comments, we have, however, updated our analysis to use REMA at a resolution of 100 m. We have added text to the manuscript to better explain the purpose of this analysis.

- L308: Clarify what "removing any bins that lie beyond the bounds of the original waveform sample" means.

We agree that this wording is unclear and have revised accordingly. This refers to the process of removing the parts of the two waveforms that do not intersect when they are aligned according to their COG.

- L327: Do you mean the range from -1 to 0?

Thanks for spotting this; yes we do. We have changed this.

- L364: Which part of the waveform do you expect to record? If it is only a portion (e.g., the trailing edge), explain the value of recording just that part.

This varies, but our motivation here is that more data is better. Although the current literature is largely constrained to leading edge-based approaches in non-interferometric cases, it is likely that in the future, algorithm advances may make it possible to retrieve information from other parts of the waveform. As we are interested in the evolution of the ice sheet, it is important that as much information as is feasible is obtained at any one time, given the inherent limitations imposed by the size of the range window. We agree, however, that our explanation was not particularly clear, and that this would benefit from better clarification in the manuscript. We have therefore included additional motivating text to explain these elements of the analysis.

- L398: Why was the modal value used? Provide reasoning for this choice.

The modal value was used here as the data is discrete. We have revised the text to make this clearer.

- L499: Specify the algorithm used in the S3 product for slope correction.

We have specified the algorithm used in the S3 product for slope correction over land ice, which is a slope-based algorithm using a slope map derived by Helm et al. (2014).

Community Referee #1 *(Benjamin Smith)*

This is a quick and non-exhaustive review of "Assessment of Sentinel-3 Altimeter Performance over Antarctica using High Resolution Digital Elevation Models" by Phillips and McMillan, which uses the REMA mosaic to investigate how well range windows selected by the SENTINEL radar altimeters capture the surface in Antarctica. I found the manuscript to be nicely written, and thought that it gave a very good background discussion of how radar altimeters work, and that it presented its findings quite clearly.

I would like to raise one question about the scope of the study, and one about the presentation of the SVD analysis, and my only significant remaining concern relates to font sizes in the figures (hint: they're not too large).

Question 1: Scope of the study.

The study analyses the performance of the SENTINEL-3 altimeters over a range of Antarctic surfaces, and finds that in a lot of interesting places, the telemetered range window does not capture the POCA return from the surface and/or does not capture the full range of elevations illuminated by the radar beam pattern. This finding suggests that the SENTINEL missions and future missions should use a larger range window, and should consider implementing open-loop surface tracking to better position the range window relative to the surface. It would have been nice to see an explicit analysis of how these two options could be implemented- for example, the study could analyze how large the range window would need to be to consistently capture the surface, and could analyze the resolution of the on-board surface elevation model needed to consistently capture the surface.

We agree that an analysis focussed on determining the optimal range window size is an interesting question. As requested, we have therefore included an additional figure and more text to explore the impact of the size of the range window on the proportion of the surface captured. More specifically, we have calculated the maximum possible surface capture (iterating the range window placement vertically 1m at a time) for multiples of the original 60 m range window size; namely 60m, 120m, 180m, 240m, 300m, and 360m. Doing so allowed us to determine the point by which minimal capture gains are made when further increasing the range window size.

Regarding the resolution of the DEM used to form an OLTC - whilst we agree that this is an interesting question with regard to future satellites that will operate nominally with open loop tracking, we emphasise that the scope of this current work is to focus on an assessment of the closed loop tracking acquisitions routinely made by Sentinel-3, which do not utilise such an OLTC operationally. Thus we believe that this analysis, whilst interesting, is beyond the scope of the current study. That being said, we do hope that this work will serve as a framework for future assessment of altimeters that will operate in open loop tracking, such as CRISTAL, and also for defining the technical requirements of future radar altimeter missions. Accordingly, we have revised the text to emphasise that this work can act as a framework for future assessments, making reference to future possible studies related to the characteristics of the onboard DEM within the context of missions operating with open loop tracking.

I was left behind a bit by the discussion of topographic capture. My naïve assumption is that as long as the waveform captures the POCA point, the rest of the waveform structure is not generally interpreted. If this is the case, then perhaps section 6.2 is not needed in as much detail.

As outlined in a reply to anonymous reviewer #2, our motivation here was to optimise data volume retrieval, with a view to future methodological innovations that may allow retrieval of topographic information from beyond the leading edge; for example based on numerical waveform simulation or machine learning approaches. We agree, however, that this motivation was not clear and could have been better described. We have now revised the text to better explain our motivations for assessing the proportion of the surface return that is currently captured.

I was also unsure of the significance of the waveform decorrelation discussion. While this is interesting in the abstract, its importance for understanding the ice sheet is not as clear. Beyond this, I thought that section 6.4 would have benefitted from a little more discission of the mechanisms that determine waveform shape. I would assume that surface slope across the beam would be the most important driver of waveform shape, and that the correlation would then be more or less determined by the along-track consistency of surface slope and roughness. This seems testable using REMA.

We realise that the relevance of this section was not clearly conveyed and so have revised the text. We also agree that surface slope across the beam is likely the most important driver of waveform shape and so have refined or approach to calculate slope and roughness with respect to the across-track direction only. For each record, this then gives us a much more representative measurement of slope and roughness as seen by the altimeter, and impacting upon waveform morphology. We have included additional text to explain this process and the rationale behind it. As requested, we have also included additional text discussing the mechanisms that drive waveform shape.

Question 2: Presentation of the SVD.

Upon first reading, it was not at all clear to me why the SVD of the surface would give an estimate of the surface slope. I think section 5.1 would be much improved by a couple more equations describing the SVD approach. I'm not sure why the SVD is preferred over a simple least-squares calculation of the surface slope based on a the elevations within a small window on the ice sheet surface, which would in general be much easier to compute than the SVD because the matrix relating the surface elevations and the slope could be computed once and applied to every window on the ice sheet in the same way. It would also be good to define clearly the shape of the region to which the slope analysis was applied: it is not clear whether "The centre points in each region" (line 328) means the 5x5-pixel window, or some subset thereof.

We acknowledge that our explanation and justification for selecting this approach was not clear in the original manuscript. To clarify – as outlined in our response to anonymous reviewer #2, our motivation for using SVD was that it minimises the true geometric (orthogonal) distance from the fitted plane, whereas the least-squares (LS) approach minimises residuals in the vertical direction. This distinction makes SVD particularly suitable for complex terrain, as it reduces orientation bias by minimising the true geometric distance, avoiding slope underestimation. We see this quantitatively in a reduced slope-roughness Pearsons correlation coefficient (0.808), compared to SVD (0.819), signalling an improved algorithmic decoupling of slope and roughness. Furthermore, SVD directly provides the normal vector to the fitted plane via the right singular matrix, with SVD requiring a separate, additional calculation. This allows us to compute orthogonal residuals efficiently by taking the dot product of the centred elevation points with this normal vector. Although we agree with the reviewer that least-squares is a more computationally efficient approach (~30% faster according to the tests that we have performed), because the additional computational time is not prohibitive for this application, we prefer to stick with SVD, due to the improved theoretical plane fit with respect to least-squares. This is supported by the slightly lower Pearson correlation between slope and roughness observed with SVD (0.808), in comparison to least squares (0.819), indicating better decoupling of these metrics and, we believe, a more robust representation of local surface variability. As requested, we have also revised the text to better explain this methodological implementation, and included additional equations. Regarding the region used for the slope analysis, "the centre points in each region" refers to the pixels within the (now) 9x9 window after subtracting the means along the x, y, and z axis. We agree this was worded poorly, and have adjusted the text accordingly.

Specific editorial comments:

Line 43: It is the failure of the assumptions that leads to difficulties

We agree and have adjusted the text to make this clearer.

Line 77: low-> short

Agreed, we have made this change.

Line 145: specify "track-to-track spacing" rather than "across-track spacing"

We agree and have made this change.

Line 155: no hyphen between range and window

Thank you, we have made this change.

Line 221: "As such, values determined for slope and roughness calculated in this way encode each other via complex, non-linear interactions." I don't understand this sentence, and if I did, I think I would disagree with it.

The reasoning here for stating it in this way is that when roughness is measured using the residuals along the z-axis to a fitted plane, the resulting values are dependent on the slope of that plane. In other words, slope and roughness calculated this way become interdependent in a way that does not necessarily reflect their true geophysical relationship, making it more difficult to isolate their individual effects on a given dependent variable. The motivation for using orthogonal residuals is that they are far less sensitive to this issue. We agree that the original sentence was unclear, and that the use of "non-linear" in particular had no clear mathematical grounding, and so we have used alternative language and revised the text for clarity.

Section 5.4: please check the tense of the first and second paragraphs. The first paragraph should be in present tense.

We have changed the first paragraph to be in the present tense.

Lines 320-325:

Please define R

Here, R referrers to the Pearson correlation coefficient. We have now defined it explicitly within the text.

Please give an equation for the line approximating the correlation function. What are the units of the slope? It appears that they are (15 km)^(-1).

We have re-written this part of the text to be clearer, included a supporting equation, and clarified the units.

Figures:

These are nice figures, but I had to blow them up to the size of a large pizza to make out the text. The fonts need to be much larger!

We completely agree that the text is too small and have adjusted accordingly.