

Response to Reviewer 1

We thank the editor, referees, and community contributors for their careful reading of our manuscript and for their constructive and insightful comments. We have revised the manuscript in response. In particular, we added a new intercomparison with published surface elevation change products with an evaluation against ICESat-2 ATL15 (new Section 4.1.2 and Table 3), strengthened the framing in the Abstract and Introduction, and expanded the Discussion to clarify modeling assumptions and limitations. We also clarified methodological details, and addressed all specific comments throughout the manuscript. We believe that these revisions have significantly strengthened both the technical clarity and the scientific context of the study.

Reviewer comments in black

[Our response in blue](#)

Response to Reviewer 1

Reviewer comment:

The manuscript "A State-Space Model for Monitoring Greenland Ice Sheet Surface Elevation Change from CryoSat-2" presents a method to derive surface elevation changes of the Greenland ice sheet at a monthly temporal resolution and at 5 km spatial resolution. For estimating the surface elevation changes a state-space filtering technique is applied to CryoSat-2 data (radar altimetry). A comparison is made with satellite and airborne laser altimetry data. The data is validated with data from automatic weather stations at a total of three locations. Further development of surface elevation change products is highly appreciated, as they could provide valuable insights into the ongoing processes related to the Greenland Ice Sheet. Advances in the processing of radar altimetry data to derive spatially and temporally coherent SEC are a very important part of this, which makes the topic of this study significant.

In its current form, the article essentially presents a data set and the underlying method, a state-space model.

However, much remains unclear to me. In particular, after reading the manuscript, I am left with the following questions:

- Is this method as good as or better than other methods?
- Which limitations of existing methods does the presented method overcome, and what limitations does it have itself?
- Does the final SEC product agree better with independent data than existing products that have already been published?

As long as these questions remain unanswered, the relevance of the study remains unclear to me. The application of a method alone does not justify this study.

The article's technical focus also makes it difficult to identify the relevance of the study. The authors should state more clearly what the actual problem is that they are working on. If I understand correctly, the challenge lies in the fact that the retracted heights (POCAs) are highly irregular in both space and time. The task now is to generate regular monthly grids from these irregularly sampled observations. I consider the choice of a state space filter to be extremely suitable for this task. It would be helpful to provide a clearer justification for this choice and demonstrate to the reader what the spatial and temporal sampling of the initial data looks like. Additionally, a review of alternative methods for this task and their limitations is missing, which are apparently to be resolved here.

Essentially, I can think of two ways of how to proceed with the manuscript.

1. The development and analysis of the methodology could be more rigorous. It would be interesting to see a comparison of the state space model with other methods of creating temporally and spatially coherent grids from POCAs. If applicable, the manuscript might then be better suited to a remote sensing journal specializing in methodology or a journal focusing on geodata analysis methods. The authors may also consider to focus on publishing a dataset and aiming for an article in a data journal like ESSD. The authors should consider which readers they want to reach, bearing in mind that readers of *The Cryosphere* are potentially more interested in insights into cryospheric processes than in data analysis methods. The choice of title makes it already clear that this is a technical and methodological study, rather than one focusing on conclusions about the Greenland ice sheet.
2. The study could be expanded by demonstrating that signals can be detected with a higher accuracy compared to other SEC products, or maybe even compared to other remote sensing methods. I agree with Jichang Shen's comment that a comparison with other SEC products should be included. While this would not be an independent validation, it would be a helpful comparison. Ideally, a validation dataset would be used to determine which SEC product performs best. I think it is problematic that the current validation is supposed to be meaningful for the entire SEC product across the entire GIS, based on only three validation points. Ideally, a validation is designed to allow statements to be made about the entire GIS as extensively as possible. Furthermore, I think it is necessary to expand the analysis of the signals contained in the determined SEC. The abstract promises that it should be possible to "detect seasonal cycles, long-term trends and abrupt changes". Accordingly, the following questions could be answered after some more analysis: What can we learn about the seasonal cycles of the Greenland ice sheet in space and time? How is summer ablation changing over time and in which regions? How significant is glacial thinning over the observation period? How does the long-term trend and its uncertainty compare to that obtained using other methods? The authors may also be aware of questions that could not be answered using previous processing strategies, but which can be answered using the developed "State Space Model".

Response:

We thank the reviewer for the detailed and constructive assessment of the manuscript and for recognizing the relevance of improving surface elevation change (SEC) products for the Greenland Ice Sheet. We agree that the original version of the manuscript did not sufficiently clarify the motivation of the study, the methodological contribution of the proposed approach, and its relation to existing SEC products. In the revised manuscript, we have therefore clarified the objectives of the study, expanded the validation and intercomparison analysis, and improved the description of the methodological framework.

We have revised the Introduction to more clearly explain the core problem addressed in this study. CryoSat-2 radar altimetry provides highly irregular point-of-closest-approach (POCA) elevation observations in both space and time. Reconstructing spatially and temporally coherent monthly SEC fields from such irregular observations therefore represents a non-trivial spatio-temporal estimation problem. The state-space model is introduced specifically to address this reconstruction step by jointly modeling spatial dependence through a Gaussian Markov Random Field and temporal evolution through an autoregressive process. This formulation allows the method to reconstruct continuous monthly SEC fields while explicitly accounting for irregular sampling and observation uncertainty.

Following the reviewer's suggestion, we have added an explicit intercomparison with two published Greenland-wide SEC datasets: the monthly 5 km radar-altimetry SEC product of Zhang et al. (2022) and the smoothed monthly Greenland SEC product of Khan et al. (2025). All datasets were brought to a common grid and evaluated against the independent ICESat-2 ATL15 dataset. The comparison shows that the proposed dSEC product exhibits the lowest overall misfit relative to ATL15 among the evaluated CryoSat-2-based products, with reduced RMSE and reduced spread of elevation-change differences. This comparison is now presented in Section 4.1.2 and summarized in Table 3.

To further improve transparency and reproducibility of the methodology, we now provide a simplified code example together with input data for a single processing tile in the supplementary material. This example allows readers to reproduce the state-space reconstruction workflow and test the method on a representative subset of the CryoSat-2 observations. The example demonstrates how irregular CryoSat-2 POCA observations are transformed into monthly SEC fields within the state-space framework. We believe this addition improves the accessibility of the methodology and facilitates reproducibility of the results.

Finally, we have clarified the scope of the study. The primary aim of this work is to present and evaluate a framework for reconstructing spatially and temporally coherent monthly SEC fields from irregular radar-altimetry observations. While the manuscript focuses primarily on the methodology and resulting dataset rather than a detailed glaciological interpretation of elevation-change signals, the resulting dataset enables analyses of seasonal variability and long-term trends across the Greenland Ice Sheet and provides a basis for future studies of cryospheric processes.

We thank the reviewer for these helpful suggestions, which we believe have improved the clarity, motivation, and completeness of the manuscript.

General Comments

Reviewer comment:

The abstract is very technical and does not provide enough contextualization of the relevance of the study presented. What is the motivation for the article? Why are monthly surface elevation changes needed, and what is the problem with existing products? Why is the SSM needed? It would also be desirable to quantify the statement "... consistent elevation change records that capture both seasonal variability and long-term trends of the Greenland Ice Sheet". Can you state the "seasonal variability" and "long-term trends" with a value and an uncertainty?

Response:

We thank the reviewer for this helpful suggestion. The abstract has been revised to better emphasize the scientific motivation and context of the study. In particular, we now clarify that monthly surface elevation change estimates are needed to resolve seasonal variability while simultaneously capturing long-term elevation trends. We also explicitly describe the challenge posed by irregular satellite sampling and data gaps, which motivates the use of a spatio-temporal state-space reconstruction framework.

Furthermore, the abstract now quantifies key signals identified in the dataset. The Greenland Ice Sheet shows an average elevation change rate of -0.069 ± 0.002 m yr⁻¹ over the period 2011–2025, with stronger thinning in the ablation zone (-0.314 ± 0.009 m yr⁻¹) and near-stable conditions in the interior (-0.014 ± 0.002 m yr⁻¹). Seasonal variability is evaluated using independent AWS and SERAC observations, where correlations between dSEC and AWS surface-height records range from 0.58 to 0.75.

These additions provide clearer context for the motivation of the study and better quantify the seasonal and long-term signals captured by the dataset.

Reviewer comment:

The introduction does not mention any research questions or hypothesis derived from the current state of knowledge. What are the research questions being investigated here? What are the hypotheses being tested here? Could it be said that the existing methods of altimetry processing listed in I37 are not capable of quantifying the processes mentioned in I32 with sufficient accuracy? Currently, it sounds as if the existing methods are very good ("mark a major step forward"). Why is the SSM needed then?

Response:

We thank the reviewer for this important point. We do think that the existing methods cited in the Introduction represent major progress for deriving surface elevation change from radar altimetry. Our intention is not to suggest that these approaches are inadequate, but rather to clarify that reconstructing monthly, spatially coherent SEC fields from irregularly sampled CryoSat-2 observations remains a challenging problem.

The motivation for introducing the state-space model is therefore not that previous methods fail, but that they make different methodological choices. In particular, existing approaches may rely on predefined seasonal formulations.

The research objective of this study is to investigate whether a spatio-temporal state-space framework can reconstruct monthly Greenland Ice Sheet SEC fields that are spatially coherent, temporally continuous, and quantitatively consistent with independent observations. Our working hypothesis is that modeling the SEC field as a latent spatio-temporal process, with spatial dependence represented by a GMRF and temporal persistence by an AR(1) process, provides a robust way to estimate monthly elevation-change fields from irregular radar-altimetry sampling without imposing a predefined seasonal cycle.

Reviewer comment:

The introduction does not clearly state which specific step of altimetry processing is being addressed here. This should be clarified in advance for readers who are not familiar with all the details of altimetry processing. State-space filtering is used here as a spatial-temporal interpolation method (or extrapolation method), is that correct? Is the step then similar to, e.g., IDW, collocation/kriging or similar?

Response:

State-space modeling in our framework indeed performs a spatio-temporal reconstruction of the elevation field from irregular observations. In that sense, it serves a role similar to interpolation methods such as inverse distance weighting or kriging, in that it estimates values on a regular grid from scattered measurements. However, the approach used here differs fundamentally from classical interpolation techniques.

Rather than applying a deterministic interpolation scheme, the state-space model formulates the elevation field as a latent spatio-temporal stochastic process. Spatial dependencies are represented using a Gaussian Markov Random Field (GMRF), while temporal evolution is modeled using an autoregressive AR(1) process. The model parameters and the latent elevation field are estimated jointly through maximum likelihood using the Laplace approximation implemented in RTMB.

This formulation allows the method to account explicitly for observation uncertainty, spatial dependence between neighboring grid cells, and temporal correlations between successive epochs. As a result, the reconstruction can be interpreted as a statistically consistent estimate of the underlying elevation field rather than a purely geometric interpolation.

Reviewer comment:

The question that keeps coming to mind as I read, and where I make comments in several places: How compares your result to existing published SEC products (e.g. 10.5194/essd-14-973-2022, 10.5194/essd-2024-311, 10.5194/essd-17-3047-2025) in space and time? Can we conclude that this SEC product represents an improvement on existing ones? Or more general: Why is there no quantitative comparison with other studies that did similar things?

Response:

We thank the reviewer for raising this important point. In the revised manuscript, we now explicitly address this by adding a quantitative intercomparison with two published SEC products (Zhang et al., 2022; Khan et al., 2025), together with validation against ICESat-2 ATL15 as a reference (new Section 4.1.2 and Table 3). This analysis shows that dSEC achieves lower RMSE and reduced spread of differences relative to ATL15 compared to the evaluated community products, providing quantitative evidence that the proposed state-space framework yields improved agreement with independent laser altimetry.

Reviewer comment:

Am I correct in understanding that validation of short-term signals only takes place at three individual AWS stations? The caption for Figure 1 does not state that these are shown there. This should be added. Section 2.4 contains a lot of general information about AWS measurements. Please clearly state at the outset which data is used, and justify the selection of data. Furthermore, I do not understand why it is justified not to use data from all the other AWS stations (e.g. 10.5194/essd-13-3819-2021). I'm wondering whether the authors only selected stations where the comparison with the SEC works well. In my opinion, a comparison with all available AWS stations is necessary in order to assess the validity of the study. How do you validate short-term signals from ice flow changes? Maybe firn thickness changes from publicly available firn models are useful to get further insights in some parts of the GIS. How do the seasonal signals from the SEC product align with results from regional climate modelling?

Response:

We agree that point-based AWS observations alone cannot characterize the performance of a Greenland-wide SEC product. For this reason, the AWS comparisons are intended specifically to evaluate seasonal surface-height variability at selected sites, rather than to provide comprehensive spatial validation. To assess large-scale performance, we rely primarily on spatially extensive independent datasets, namely ICESat-2 ATL15 and Operation IceBridge, which provide Greenland-wide coverage and form the basis of the trend validation (Section 4.1.1).

The three AWS sites were selected as case-study locations where several complementary datasets are available. In particular, these sites provide long and continuous AWS records, are located away from grid boundaries and persistent CryoSat-2 data gaps, and coincide with locations where the high-temporal-resolution SERAC time series are available. This allows a consistent comparison between dSEC, SERAC, and in situ observations, which is useful for evaluating seasonal variability in the reconstructed elevation-change signal.

The AWS/SERAC analysis is therefore intended as a localized illustration of seasonal variability rather than a Greenland-wide validation of the SEC product. Ice-sheet-scale validation is instead performed using spatially extensive datasets (ICESat-2 ATL15 and Operation IceBridge). We have clarified this distinction and the rationale for the station selection in Section 2.4 and in the caption of Figure 1.

Reviewer comment:

If I understand correctly, one objective is also to determine SEC in glacial valleys where other data sets encounter difficulties. Is laser altimetry data even suitable for comparison or validation there, given that these glacial valleys are often cloud-covered? It would be also very useful to know what the spatial and temporal sampling of the laser altimetry data looks like in space and time. Are there regions in Greenland where you could first demonstrate how well CryoSat-2 and ICESat-2 agree before generating the SEC grids?

Response:

We thank the reviewer for this point. Our primary objective is not to resolve SEC specifically in glacial valleys, but rather to reconstruct Greenland-wide monthly SEC fields. The ICESat-2 data are used primarily as an independent reference dataset for validation of the reconstructed SEC fields rather than to resolve elevation change in specific glacial valleys. The comparison with ICESat-2 ATL15 is therefore performed at the grid scale and over regions where both datasets overlap spatially and temporally. The purpose of this comparison is to assess the overall consistency between the CryoSat-2-derived monthly SEC fields and independent laser-altimetry-based elevation change estimates.

ICESat-2 ATL15 is derived from repeat-track laser altimetry measurements that are aggregated into spatial grids. Although individual photon returns can be affected by cloud cover, the ATL15 processing combines multiple observations to produce spatially continuous elevation-change fields. The comparison therefore reflects aggregated observations rather than individual passes.

The validation is restricted to regions where both datasets overlap spatially and temporally, ensuring that the comparison reflects areas with sufficient sampling from both radar and laser altimetry.

Reviewer comment:

A somewhat more detailed explanation of the motivation for the stochastic description of SEC in space and time would be desirable. Are there any statistical analyses of SEC that could be cited as a basis here? How do you justify to use the GMRF in space and the AR(1) process in time?

Response:

We thank the reviewer for highlighting the need for clearer justification of the stochastic formulation. Surface elevation change evolves continuously in both space and time and is therefore well suited to a spatio-temporal statistical description. In this framework, the Gaussian Markov Random Field (GMRF) represents local spatial dependence between neighboring grid cells, while the AR(1) process describes the temporal persistence of elevation change from one time step to the next. This combination allows the model to capture spatial coherence and temporal continuity while remaining computationally efficient at ice-sheet scale.

The use of GMRF-based spatial models and autoregressive temporal processes is common in geostatistical and environmental time-series modeling, particularly when reconstructing spatial fields from irregularly sampled observations. We clarify this motivation more explicitly in the revised manuscript and include references to related statistical approaches used in cryospheric and geophysical applications.

Reviewer comment:

It is not clear to me from the description of the method how a distinction can be made in the approach between uncorrelated neighbours ($\phi = 0$) and measurement noise (ϵ_i). How are temporal and spatial error correlations taken into account, which can be particularly significant in radar altimetry? Is it justified to describe observational errors with a normal distribution (Eq. 1)? I think it is very necessary to clearly state and justify the assumptions made in the methodology.

Response:

We thank the reviewer for raising this point. In the model, spatial correlations between neighboring grid cells are represented through the Gaussian Markov Random Field (GMRF), which describes the spatial structure of the underlying elevation-change field. In contrast, the observational error term (ϵ_i) represents measurement noise associated with individual altimetry observations and is treated as independent noise. In this way, the model separates spatially correlated signal from random measurement noise.

Temporal correlations are represented through the AR(1) process that describes how the elevation-change field evolves in time. The assumption of Gaussian observational errors is a common approximation for radar altimetry data after retracking and quality filtering.

To improve transparency, we will also provide the R code for the model implementation and an example input data in the supplementary material.

Reviewer comment:

Assumptions made about stochastic processes are not discussed or critically evaluated within the discussion.

&

The discussion lacks a critical assessment of the method presented. Why are the results initially so noisy and why is a post-processing step necessary? This suggests that the SSM method has limitations that need to be explored and discussed.

Response:

We appreciate this important point. Although the state-space model separates signal and noise under the assumed statistical structure, the reconstructed monthly fields may still exhibit small-scale variability related to sampling heterogeneity and residual measurement noise. The spatial median filtering and temporal Savitzky–Golay smoothing are therefore applied to improve robustness, enhance interpretability, and facilitate comparison with independent datasets. The post-processing step does not modify the large-scale signals

recovered by the state-space model but primarily reduces small-scale variability arising from heterogeneous sampling density and residual observational noise.

In the revised manuscript, uncertainties are now propagated through this post-processing step using a Monte Carlo approach. Perturbed realizations of the reconstructed field are filtered in the same way as the mean solution, and the final uncertainty is estimated as the ensemble standard deviation. The reported uncertainties therefore correspond to the fully post-processed dSEC product, ensuring consistency between the values and their associated error estimates.

Specific comments

114: Please specify "high-resolution" & 114: What do you mean by "temporally consistent"?
We have revised the Abstract to explicitly specify the spatial and temporal resolution of the product. The term "high-resolution" has been replaced with "monthly temporal resolution and 5 km spatial resolution" to avoid ambiguity.

121: Please provide a number to contextualize "significant"
We have rephrased the sentence to remove the term "significant" and instead state that the Greenland Ice Sheet is contributing to sea-level rise, supported by Otosaka et al. (2023).

133: Please provide references for the "annual SEC products traditionally derived ...".
We have included references to that sentence (e.g., Simonsen et al., 2017; Sørensen et al., 2018)

137: You may add Helm et al. (2024, 10.5194/tc-18-3933-2024) when you provide references for "new processing methods".
We thank the reviewer for this helpful suggestion and have added Helm et al. (2024) to the list of recent processing approaches.

139: What do you mean by "integration" in "better integration with regional climate models and in situ observation"?
We thank the reviewer for this comment and have clarified the wording by replacing "integration" with "direct comparison and joint analysis with regional climate model outputs and in situ observations."

140: Please provide an introduction of state-space models and please argue which limitations you aim to resolve with them. In the paragraph before you mention "These efforts mark a major step forward ..." But where are the limitations? What justifies the need of a new method? (See also my general comments given above). & 140: "detect subtle patterns" That would indeed be interesting, but as far as I can see, it plays no role at all in the rest of the manuscript.
We have revised this part of the Introduction to briefly introduce state-space models and clarify that their primary role here is to reconstruct spatially and temporally coherent elevation-change fields from irregular and noisy altimetry observations while accounting for measurement uncertainty and spatio-temporal dependence. The phrase "detect subtle patterns" has been removed, as it was not sufficiently substantiated elsewhere in the manuscript.

1155: You may mention that you use level-2 data at the beginning of Section 2.

We thank the reviewer for this suggestion and have added a statement at the beginning of Section 2 to clarify that CryoSat-2 Level-2 data are used as input to the analysis.

1110: Please argue why the comparison with SERAC is useful as there is already a comparison with laser altimetry (Sect. 2.2 and 2.3).

We thank the reviewer for this comment. We have clarified this in Section 2.5 by adding the following sentence: “*SERAC provides localized, high-temporal-resolution elevation-change time series that complement the ATL15 and OIB validations by enabling assessment of seasonal and short-term variability.*”

1112: Why is the SERAC time series only created for the three AWS points? What are the arguments for not doing this for all points on the ice sheet?

The SERAC–FDM fusion time series are computationally intensive to generate and are primarily intended here for localized, process-level evaluation of seasonal variability. We therefore limited this analysis to three sites where high-quality AWS records are also available, allowing a consistent three-way comparison between dSEC, SERAC, and in situ observations. Greenland-wide validation is instead performed using spatially extensive datasets (ICESat-2 ATL15 and OIB), which are more suitable for assessing large-scale performance.

1145: I don't understand what was inefficient about previous approaches. What is the problem?

We have revised the text to clarify that the inefficiency refers to computational scalability and the limited ability of some approaches to jointly model spatial and temporal dependencies when reconstructing monthly fields from irregularly sampled altimetry data. We now explicitly state that “*because CryoSat-2 POCA elevation anomalies are irregular in space and time, we adopt a framework that jointly models spatial and temporal correlations to reconstruct coherent surface elevation anomaly fields.*”

1154: Please clarify up to what correlation lengths processes can be captured temporally and spatially with the GMRF and AR(1), or does 1174 mean that only direct neighbours are considered?

We thank the reviewer for this important clarification. The model does not impose an explicit spatial or temporal correlation length scale. Instead, spatial dependence is defined through a 4-nearest-neighbour GMRF sparse precision structure on the 5 km grid. While only adjacent grid cells are directly connected in the precision matrix, this structure induces marginal correlations that propagate through the field and decay with distance. Similarly, the AR(1) temporal formulation implies exponentially decaying correlations across multiple time steps.

1159: Some motivation would be desirable for the chosen spatial and temporal resolution.

We have clarified the motivation for the chosen 5 km spatial resolution and monthly temporal resolution. The 5 km grid represents a balance between resolving spatial variability and ensuring sufficient data density for stable estimation given the irregular CryoSat-2 POCA sampling. The monthly resolution is selected to capture seasonal variability while maintaining adequate observational support and numerical stability in the spatio-temporal reconstruction. We have added the following sentence in the Method section: “*The 5 km spatial resolution and monthly temporal resolution represent a compromise between resolving spatial and seasonal variability and ensuring sufficient data density for spatio-temporal*

reconstruction given the CryoSat-2 sampling that has a 369 day repeat cycle and a 30-day subcycle (Bouzinac, 2019)."

1187: A more mathematical explanation would be desirable, detailing how the problem is solved and how it is implemented numerically. I do not see how one would be able to reproduce the results based on the description. The authors could also consider making the routines used publicly available alongside detailed documentation. Am I correct in understanding that no Kalman filter or smoother is used here to estimate the states?

We thank the reviewer for this important comment. The state estimation is performed using Template Model Builder (RTMB), where the marginal likelihood is obtained via Laplace approximation of the joint likelihood of the latent field and model parameters. In this framework, the high-dimensional latent spatio-temporal field is treated as a random effect, and optimization is carried out numerically using automatic differentiation and using a sparse matrix.

No, that is correct, the approach does not rely on a Kalman filter or smoother. Instead, the technical solution here is the Laplace approximation.

1191: Please explain "generalized delta-method".

We have revised the manuscript to briefly clarify that *"the generalized delta-method approximates the uncertainty of the reconstructed field by propagating parameter uncertainty through a first-order Taylor expansion around the optimum."*

1196: Where does the unrealistic high-frequent variability comes from? Are these errors? In 146 you promised SSMs explicitly separate signal from noise.

The state-space model does separate signal and noise under the assumed Gaussian observation model and specified spatio-temporal structure. However, radar altimetry measurements can exhibit residual variability related to spatially correlated measurement errors, and surface scattering changes, that are not fully captured by the model.

The subsequent spatial median filtering and temporal Savitzky–Golay smoothing are therefore applied as a post-processing step to reduce residual high-frequency variability and enhance interpretability, for trend and seasonal analyses. We have clarified this distinction in the manuscript with the following change of wording: *"to reduce residual high-frequency variability not fully captured by the statistical model."*

1200: How do you estimate the seasonal component and the trend? As far I can see both are not separate parameters in your model.

Yes, the state-space model does not explicitly parameterize seasonal or trend components. Instead, it reconstructs the full latent spatio-temporal elevation anomaly field, within which seasonal variability emerges directly from the data through the temporal AR(1) structure and spatial regularization. The seasonal signal is therefore not predefined but inferred from the observations and therefore datadriven.

l205: Is the uncertainty then representative? Why is it not propagated to the post-processed SEC?

We thank the reviewer for this important question. In the revised manuscript, the uncertainty is now propagated through the post-processing. Starting from the posterior standard deviations provided by the state-space model, we generate an ensemble of perturbed dSEC realizations and apply the same spatial median filtering and temporal Savitzky–Golay smoothing to each realization. The uncertainty of the final product is then estimated as the ensemble standard deviation (1σ) of the post-processed fields. The reported uncertainties therefore represent the propagated uncertainty of the fully post-processed dSEC product, ensuring consistency between the values and their associated error estimates.

l207: I don't quite understand that. Apparently, the estimated parameters still contain too much noise, so the estimated model is smoothed spatially and temporally in a post-processing step. Shouldn't this noise, which is filtered out in the post-processing step, also be included in the uncertainties? Having read section 3, I don't really get the impression that the methodology can sufficiently separate signal and noise. Ideally, the SSM would be designed so that the post-processing step would not be necessary. Or, to put it another way: How does this variant perform compared to other approaches for generating spatially and temporally consistent SEC grids from radar altimetry observations?

We thank the reviewer for this thoughtful and important comment. The state-space model estimates the latent spatio-temporal elevation anomaly field under the assumed Gaussian observation model and GMRF–AR(1) structure, providing a statistical separation of signal and noise within that framework. However, radar altimetry data can exhibit residual variability related to spatially correlated measurement errors, irregular sampling, and localized artifacts that are not fully represented in the observation model. Consequently, some small-scale or short-term variability remains in the reconstructed fields.

The spatial median filtering and temporal Savitzky–Golay smoothing are therefore applied as a pragmatic step to enhance robustness and interpretability, particularly for trend and seasonal analyses. In the revised framework, uncertainties are propagated through this post-processing using a Monte Carlo approach, as described above. The reported uncertainties thus correspond to the fully post-processed product.

The performance of the proposed framework is evaluated through comparisons with independent datasets (ICESat-2 ATL15 and Operation IceBridge) and existing SEC products, as presented in Section 4. These comparisons demonstrate improved spatial coherence and competitive agreement, supporting the robustness of the joint spatio-temporal modeling approach.

l220: But this does not account for spatial correlations across grid cells, is that right? This explanation should be included in the Methods, not the Results.

We thank the reviewer for this important observation. The regional uncertainties are derived by propagating the grid-cell uncertainties through the area-weighted spatial averaging used to compute the Greenland-wide and regional mean elevation-change time series. This calculation does not explicitly propagate the full spatial covariance between grid cells. As

suggested, we have moved and clarified this methodological description in the Methods section.

Figure 1: I suggest showing the Greenland map (Panel D) as a separate, much larger figure. It would be helpful to plot metres per year, as in Figure 3.

We thank the reviewer for this helpful suggestion. The Greenland map previously shown as Panel D in Figure 1 has now been separated and presented as a standalone figure (Figure 1) with an increased size to improve readability and spatial detail. The surface elevation change rate is now expressed in m yr^{-1} , consistent with Figure 3.

To maintain clarity, the six elevation-change time series previously shown alongside the map are now presented in a separate figure (Figure 2). The figure captions and references in the manuscript have been updated accordingly.

Figure 1: The uncertainties are constant over time, am I right? Is that realistic? Other products identify significant uncertainty of the trend (e.g. Fig. 16 in 10.5194/essd-17-3047-2025), and I would have expected time-correlated errors to be reflected in the SSM result. Are time-correlated errors taken into account?

The figure referred to in the ESSD paper shows cumulative mass change, where the associated uncertainty increases over time due to the propagation of errors in the cumulative estimate. In contrast, our figure shows surface elevation change time series at individual grid cells, rather than cumulative quantities, and therefore the uncertainty behaviour differs.

In the state-space model (SSM), time-correlated errors are accounted for through the temporal structure of the model. However, the magnitude of the uncertainty at each epoch is largely controlled by the local data density. Because the CryoSat-2 observation density within a given grid cell tends to remain relatively similar over time, the estimated uncertainty at that location can appear approximately constant throughout the time series.

Importantly, the uncertainties vary between grid cells, reflecting spatial differences in measurement density and coverage across the ice sheet. Thus, while the uncertainty may appear relatively constant over time at a specific location, it still reflects both the temporal correlation structure of the model and the spatial variability in the observations.

Figure 2: It would be helpful to show an integrated volume time series here, so that interior + ablation = whole Greenland. Here, too, the total uncertainty does not seem to include any uncertainties in the trend, which I consider unrealistic.

Figure 2 (Now Figure 3 in the revised manuscript) shows area-averaged elevation change time series rather than a cumulative or integrated volume change. Consequently, the interior and ablation zone curves are not expected to sum directly to the whole GrIS curve, as it is computed independently as the spatial mean over all grid cells within the ice mask.

Regarding the uncertainties, the shaded regions represent the uncertainty of the mean elevation change at each timestep, derived from the grid-cell uncertainties estimated by the state-space model and propagated through the spatial averaging. Because the time series

represent elevation anomalies at each timestep rather than cumulative changes, the uncertainties do not accumulate over time. This clarification has been added to the revised manuscript.

Section 4.1.1: The name of the section is "Validation of trend". Essentially, this section deals with how the post-processing step minimizes noise in the differences between the result and ICESat-2 and OIB. How does the linear trend rate of ICESat-2 and the presented product compares over the period from 2019 to 2025? It would be helpful to see this as a spatial grid. Furthermore it would be interesting to a time series comparison between the SEC product and ICESat-2.

This was initially done however, including all corresponding figures resulted in a large number of panels. We therefore previously summarized the results using the table and a histogram. In the revised manuscript, we have added a spatial map of the difference between the dSEC and ATL15 trends over the full overlapping period. This complements the histogram by illustrating the spatial distribution of the differences across the ice sheet.

Regarding the suggestion of a full GrIS time series comparison, we have chosen not to include an additional time series analysis here. A similar temporal comparison is already presented in Section 4.1.2 for the three case-study sites, where the dSEC time series are evaluated against the independent high resolution SERAC product that is also based on ICESat-2 data.

Section 4.1.2: I find it highly problematic to refer to "validation of seasonal amplitude" when a comparison is only carried out at three points, because the section title suggests this is representative for the whole GIS. Why is the phase not taken into account?

We agree that the original section title could suggest a broader validation than intended. The title has therefore been revised to clarify that the analysis represents a comparison at selected validation sites rather than across the entire Greenland Ice Sheet. In addition, we note that the phase of the seasonal signal is implicitly considered through the comparison of the mean seasonal cycles and the correlation analysis between the detrended time series.

l298: The conclusion is not clear to me. "Reducing noise" is only achieved through the post-processing step.

We thank the reviewer for pointing this out. We agree that the reduction of high-frequency noise also comes from the post-processing step rather than just from the state-space model itself. The sentence has been revised to the following: *"The combination of a GMRF representation of spatial dependence with an autoregressive temporal component allows the model to isolate meaningful elevation-change signals, while additional post-processing of the radar data further reduces high-frequency noise."*

Figure 6: How were the average seasonal cycles calculated?

The average seasonal cycles are calculated from the detrended time series by grouping the data by calendar month and computing the mean value for each month over the common overlap period (2011–2024). Specifically, each time series is first linearly detrended to remove the long-term signal. The detrended observations are then assigned to their respective calendar month, and the monthly means are computed across all available years. A sentence

explaining this have been added to the revised manuscript line 331.

l323: "Independent validation against ICESat-2 ATL15 and OIB confirms the 3D-ECM method's effectiveness." Effectiveness sounds quite vague. What does it mean? Can it be quantified?

We thank the reviewer for this comment. We agree that the term "effectiveness" was too vague. We have revised the text to quantify performance more explicitly. The validation against ICESat-2 ATL15 and Operation IceBridge demonstrates that post-processing improves quantitative agreement: compared to the raw reconstruction, the smoothed dSEC product reduces the annual standard deviation of differences by 40–45 % and lowers the median absolute deviation.

l343: Please show and quantify these patterns that were not revealed before. & l344: What do you mean by "great detail"? Which integration do you refer to? Please explain how the SEC product would be valuable for climate models?

We thank the reviewer for these helpful comments. We agree that the previous wording was too broad and could be interpreted as claiming detection of previously unresolved patterns or direct application to climate model development, which is beyond the scope of the present study. We have therefore revised the text and clarified the statement. The revised sentence now reads:

"The 3D-ECM framework integrates radar altimetry observations within a coherent spatio-temporal reconstruction, enabling the generation of spatially and temporally consistent, gap-filled SEC time series across the Greenland Ice Sheet."

l347: Please explain more extensively why "entirely data-driven" is a great advantage? What are the caveats of the other studies you are citing? Here it becomes clear once again that a comparison with these products would be very useful.

We thank the reviewer for this helpful comment. By "entirely data-driven," we mean that the seasonal signal is not prescribed using a fixed sinusoidal (sine/cosine) formulation, as is common in several other approaches. In many SEC products, seasonality is represented by fitting predefined harmonic functions, which implicitly assume a stationary, symmetric seasonal cycle with constant amplitude and phase.

In contrast, our framework does not impose a parametric seasonal model. The seasonal variability emerges directly from the observations through the joint spatio-temporal estimation, allowing for non-sinusoidal shapes, temporal changes in amplitude, and shifts in phase. This is particularly relevant for Greenland, where seasonal signals can vary regionally and interannually due to evolving melt and accumulation conditions.

We have clarified this distinction in the discussion (line 402) in the revised manuscript and expanded the comparison with existing SEC products to better highlight methodological differences and their implications.

1350: The authors could consider expanding the study to include a comparison with swath-processed data. It would certainly be very interesting to see how these data sets compare in the outflow glacier regions where the sampling is sparse.

We agree that a comparison with swath-processed data would indeed be very informative, particularly in fast-flowing outlet glacier regions where sampling is sparse and elevation gradients are strong. Such an analysis could help assess how well the spatio-temporal reconstruction captures variability in dynamically complex areas.

However, we consider this extension beyond the scope of the present study, which focuses on the methodological development, and validation of the 3D-ECM framework. A dedicated evaluation against swath-processed products in selected outlet glacier regions would be a valuable direction for future work.

1356: Please clarify "ice mask". Do you mean the mask of the Greenland ice sheet without peripheral glaciers? Please provide a reference to the mask you are applied.

To clarify, the "ice mask" refers to the Greenland Ice Sheet mask used within the IMBIE framework, which excludes peripheral glaciers and ice caps not dynamically connected to the main ice sheet. In the revised manuscript, we have clarified this explicitly and added the appropriate reference to the IMBIE Greenland Ice Sheet mask (IMBIE Team, 2020).

1364: Please explain what do you mean by scalability and statistical rigour.

We have clarified that by *scalability* we mean the computational efficiency of the framework, which allows it to be applied to large, ice-sheet-scale datasets, and by *statistical rigor* we refer to the use of a consistent statistical modeling approach with explicit uncertainty estimation. The sentence has been revised accordingly to make this clearer and less technical.

1367: It's not entirely clear to me how the approach enhances the temporal resolution. Maybe you could use "unifies" instead of "enhances".

We agree that the term "enhances" may be misleading, as the framework does not increase the native sampling frequency of radar altimetry missions. We have therefore revised the wording accordingly.

1368: You conclude here that the approach "supports improved assessments of ice-sheet mass balance and dynamics". On what findings is this conclusion based? The manuscript contains no analyses of mass balances or ice sheet dynamics.

We agree that the manuscript does not include a direct analysis of ice-sheet mass balance or ice dynamics. Our intention was to indicate that improved spatio-temporal consistency in surface elevation change products can support such assessments, rather than to claim that this study itself performs them. To avoid overstatement, we have revised the wording accordingly.

Technical Corrections

You introduced SSM as an abbreviation for State Space Model. Please use throughout.

We have checked that the abbreviation “SSM” is used consistently throughout the manuscript.

When referring to uncertainties, please indicate the sigma level.

In the revised manuscript, the uncertainty level is stated when uncertainties are reported.

Figures: Please increase labels to a sufficient font size as they are hard to read at some places (e.g. Figure 1, panel C Figure 5 etc.)

The font sizes of the labels and annotations have been increased throughout the figures to improve readability. In particular, Figures 1 and 5 in the original manuscript (now Figures 1–2 and Figure 6 in the revised version) have been updated with larger labels and clearer annotations.

156: Please introduce the abbreviation as usual: Surface Elevation Reconstruction And Change detection (SERAC)

We have checked that SERAC is now defined in the usual manner.

1195: "Thorson and Kristensen, 2024" is not in the reference list.

We have corrected the reference to (Kristensen et al., 2016)

Figure 5: Please provide time periods.

This has now been added to the Figure caption in the revised manuscript.

1337: "Taken together, these comparisons indicate that dSEC generally agrees well with both the AWS records and the altimetry-based SERAC product, underscoring the ability of the 3D-ECM framework to extract physically meaningful elevation change signals that are consistent with in situ observations and the altimetry/model-based SERAC." This is redundant and could be formulated more concisely.

The sentence has been revised to make the statement more concise.