Response to Reviewer #2

Manuscript: *The Ice Sheet State and Parameter Estimator (ICESEE) Library (v1.0.0): Ensemble Kalman Filtering for Ice Sheet Models*

We thank the reviewer for carefully reading the manuscript and for the constructive comments. We appreciate the detailed feedback, particularly on the mathematical consistency of the EnKF formulation, the treatment of localization and stochastic perturbations, and the clarification of the hybrid assimilation-inversion workflow. These comments have helped us improve both the theoretical presentation and the description of the application examples.

Below we provide a point-by-point response. Reviewer comments are shown in blue, and our responses follow each comment below.

Main comments

Reviewer Comment:

1. The mathematical formulation of the EnKF algorithms needs careful review and clarification (Section 2). For example, the manuscript appears to assign the dimension $n \times N_e$ to $\mathbf{x}_i^f$, although $\mathbf{x}_i^f$ should represent a single ensemble member and therefore should have dimension n. The full ensemble matrix X^f would have dimension $n \times N_e$. This should be clarified throughout Section 2.

   The definition of the perturbed observation covariance is also ambiguous because D is described as containing columns $\mathbf{d}_i$, although $\mathbf{d}_i$ denotes perturbed observations rather than zero-mean perturbations. In addition, the replacement of $HX'(X')^T H^T + \Upsilon\Upsilon^T$ by $(HX' + \Upsilon)(HX' + \Upsilon)^T$ requires clarification because cross terms arise unless additional assumptions are made. Equation 10 is also unclear: if $X^* \in \mathbb{R}^{N_e \times N_e}$ is an ensemble-space transformation, it should multiply the full ensemble matrix X^f, not a single state vector $\mathbf{x}_i^f$. Similar dimensional issues appear in the deterministic filter equations. I recommend that the authors carefully revise Section 2, define the dimensions of all vectors and matrices, and ensure consistency with the cited EnKF formulations.

Response:
We thank the reviewer for this careful reading of Section 2, which identified several places where our notation was inconsistent or ambiguous. We have revised the section throughout to address each point.

Dimensions of X. We agree that our original notation did not clearly distinguish a single ensemble member from the full ensemble matrix. We have revised Section 2 to consistently use lowercase symbols for individual state/perturbation vectors (e.g., $x_i^f \in \mathbb{R}^n$, $x_i^a \in \mathbb{R}^n$) and uppercase symbols for the corresponding ensemble matrices (e.g., $X^f, X', X^a \in \mathbb{R}^{n \times N_e}$, with columns given by the respective vectors). This convention is now stated explicitly where X^f is first introduced and is applied consistently through the stochastic and deterministic filter formulations.

Definition of D. We revised the manuscript to distinguish the configured observation-error and analysis-factor modes. In the standard stochastic mode, $D = \bar{D} + \Upsilon$, where the columns of Υ are centered realizations of the prescribed observation error. ICESEE also supports spatially generated observation-error fields when spatially coherent measurement or representativeness errors are appropriate. The experiments reported here use the separately documented `legacy_prior_anomalies` analysis-factor mode. Motivated by the slow evolution of ice-sheet dynamics and the persistence of covariance structures inherited from the initialized ensemble, this mode sets $\Upsilon = Y'$ in the combined factor, $A = Y' + \Upsilon = 2Y'$, while leaving the innovation unperturbed, $D' = \bar{D} - Y^f$. Thus, Y' is not added to the observations and is not interpreted as an independent realization of the prescribed observation-error covariance R.

Cross terms. We agree that

$$Y'(Y')^T + \Upsilon\Upsilon^T = (Y' + \Upsilon)(Y' + \Upsilon)^T$$

holds exactly only when $Y'\Upsilon^T = \Upsilon(Y')^T = 0$. Independence gives zero cross-covariance only in expectation and does not impose exact sample orthogonality in a finite ensemble. In the legacy analysis-factor mode used here, the condition deliberately does not hold because $\Upsilon = Y'$, so $Y'\Upsilon^T = Y'(Y')^T$ and $A = 2Y'$. The implementation evaluates AA^T directly and therefore retains the resulting cross terms. The revised manuscript now identifies this construction as a legacy analysis-factor mode rather than presenting it as the standard perturbed-observation stochastic EnKF.

Equation 10 (X^* transformation). We agree with the reviewer that the original equation was dimensionally unclear. We have corrected this so that the $N_e \times N_e$ ensemble-space transformation matrix $\mathcal{X}^*$ multiplies the full ensemble matrix $X^f \in \mathbb{R}^{n \times N_e}$ rather than a single state vector, i.e., $X^a = X^f \mathcal{X}^*$, with $\mathcal{X}^* = I_{N_e} + K^* D'$ and $D' = D - Y^f$, all dimensions now defined explicitly in the text.

Deterministic filter equations. We have applied the same x/X (vector/matrix) convention throughout the EnSRF, ETKF, and DEnKF formulations, and corrected two related dimensional errors: (i) in the EnSRF update, the transformation is now applied to the full anomaly matrix X' rather than a single-column vector; (ii) in the ETKF update, the ensemble-space transformation $\sqrt{(N_e - 1)\Omega}$ is now correctly shown multiplying the anomaly matrix X', i.e., $X^a = \bar{X}^a + X'\sqrt{(N_e - 1)\Omega}$, resolving the earlier dimensional mismatch.

Nonlinear H. We removed the potentially misleading use of HX' for a nonlinear operator. The revised stochastic formulation evaluates $y_i^f = H(x_i^f)$ for every member, forms $Y^f = [y_1^f, \ldots, y_{N_e}^f]$ and its anomaly matrix Y', and computes the sample cross- and observation covariances from X' and Y'. The deterministic equations are displayed in their conventional linear-observation form, followed by an explicit rule that nonlinear applications use memberwise Y^f, Y', and the mean innovation $d^a - \bar{y}^f$. Thus no tangent-linear model or implicit matrix interpretation of nonlinear H is assumed.

Reviewer Comment:

2. Section 2.3 describes a noise modeling strategy based on spatially correlated pseudorandom perturbations and observation-space perturbations constructed either from ensemble anomalies or pseudo-random fields. This may help maintain ensemble spread and reduce some artificial correlations. However, it is not clear how this strategy removes the need for localization. Localization addresses sampling-error-induced spurious correlations in the ensemble covariance, whereas the proposed noise modeling seems mainly controls ensemble spread and perturbation structure. The manuscript should clarify whether the method truly replaces localization, approximates localization through prescribed spatial decorrelation of perturbations, or simply avoids explicit covariance localization in the presented experiments.

Response:

We thank the reviewer for raising this important distinction. We agree that the perturbation strategies (spatially correlated state perturbations and ensemble-consistent or pseudo-random observation perturbations) do not remove the sampling error that motivates covariance localization. They control ensemble spread and the spatial structure of the imposed perturbations, whereas localization explicitly suppresses spurious sampled state–observation correlations.

In direct response to this comment, we removed the claim that the perturbation strategy eliminates the need for localization and revised the manuscript to state that perturbation design and covariance localization address different problems. The former controls ensemble spread and the spatial structure imposed on perturbations, while the latter suppresses spurious correlations in the sampled forecast covariance.

Separately, while improving the revised ISSM experiment, we introduced a grouped patch-based local analysis so that sparse bed observations act locally on the bed-parameter block. This is not a covariance taper and was not introduced as a replacement for localization. In the MPI implementation, SciPy's `cKDTree` [Bentley, 1975, Virtanen et al., 2020] identifies observations inside each prescribed support. The tree is used only to accelerate the spatial search; it does not define the analysis. Nodes with the same active-observation subset are then grouped, and a single transform is computed per group, rather than computing an independent transform at every mesh node. This grouped implementation follows the local-analysis principle of Evensen [2003], but adapts it to distributed state-parameter vectors and avoids redundant transforms.

The Icepack demonstration uses the perturbation strategy without explicit covariance localization or local analysis. In the revised ISSM demonstration, sparse bed observations are handled with the added grouped patch-based local analysis, for which a 15 km radius selects the active observations. Only the bed block receives these local transforms; the globally observed surface-elevation and velocity fields remain in the global analysis. The framework retains configurable perturbation, inflation, covariance localization, and local analysis options for applications that require them; none is claimed to replace the others. The revised manuscript now explicitly reports both the added analysis option and the corresponding updated ISSM configuration.

Reviewer Comment:

3. The current text in section 4 (application examples) sometimes makes it difficult to follow which parameters are directly updated by the EnKF, which variables are observed, and how the hybrid workflow improves results relative to EnKF-only approach. For the ISSM example, the authors should specify exactly which variables are assimilated and which variables are updated by the filter. If the hybrid workflow is a major contribution, the manuscript should include a comparison against a baseline EnKF without the inversion step.

Response:

We have revised Section 4.2 to state the roles explicitly. The ISSM augmented vector is $X = [h, s, u, v, b, c]$ and serves as the common storage vector, while the active analysis blocks depend on the experiment. Surface elevation s and horizontal velocities u, v are directly observed and EnKF-updated. Bed topography b is directly observed only along sparse flight lines at ten survey times (years 2, 8, 14, 20, 24, 30, 36, 40, 46, and 50) and is updated by the grouped local bed analyses described above. Thickness h is not observed directly; its EnKF increment is obtained through ensemble cross-covariances, after which the geometry-consistency step recomputes base and surface rather than supplying an additional thickness update. Basal friction c is unobserved.

In the principal hybrid experiment, it is frozen during the EnKF analysis and inferred member-by-member with ISSM's regularized adjoint inversion after each applicable analysis; it remains an active EnKF parameter block only in the EBF comparison.

We also added the requested controlled baseline comparison. This is a separate 40-year experiment because preliminary WBF and EBF runs initialized directly with the 0.2-yr step used in the principal 100-year IBF experiment exhibited a sharp diagnostic-velocity initialization transient; in EBF, the analyzed friction changes did not propagate reliably into the next forecast. We therefore use a common four-year pre-assimilation stabilization at $\Delta t = 0.1$ yr and retain that step for all three comparison cases. A coupling preflight verified that EBF coefficient increments then altered the following dynamic forecast. Within this matched experiment, truth, initial ensemble, forcing, observations, stabilization, time step, and EnKF settings are identical, and only the basal-friction treatment changes: fixed wrong basal friction (WBF), EnKF-updated basal friction without inversion (EBF), and inversion-based basal friction within the hybrid framework (IBF). Over the free forecast from years 24–40, IBF has the lowest error in every tabulated state, grounding-line, and friction diagnostic. Its mean/final grounded-speed RMSE is 58.0/57.1 m yr^{-1}, compared with 102.9/101.5 for WBF and 99.7/89.3 for EBF. Its mean/final grounded-excluding-GL surface-elevation RMSE is 45.5/50.6 m, compared with 77.4/83.8 for WBF and 71.8/78.5 for EBF. Its mean/final centerline grounding-line error is 1.41/1.27 km, compared with 2.92/2.04 for WBF and 3.53/2.30 for EBF. Final grounded-friction RMSE is 641, compared with 777 for WBF and 789 for EBF (in $\mathrm{Pa\,m^{-1/3}\,yr^{-1/3}}$). The revised figure, table, and discussion therefore quantify the additional forecast benefit of the inversion-based treatment rather than merely asserting it.

Specific comments

Reviewer Comment:

Abstract "This matrix-free formulation eliminates the need for localization…". Please see the main comment above. I recommend revising this sentence. Avoiding explicit construction of the forecast covariance matrix is computationally valuable, but I am not convinced that it removes the effects of ensemble undersampling or spurious long-range correlations.

Response:

We agree that the original statement was incorrect. Avoiding explicit construction of the forecast covariance reduces memory and computational cost, but it does not remove ensemble undersampling or spurious long-range correlations. We deleted the claim that the matrix-free formulation "eliminates the need for localization" and revised the abstract and methods to state correctly what matrix-free algebra, perturbation design, inflation, and covariance localization each accomplish. The Icepack experiment uses neither explicit covariance localization nor local analysis. Separately, the revised ISSM configuration introduces grouped patch-based local analysis only for the bed block, with a 15 km observation-selection radius. We describe this added option as a hard spatial restriction on the observations entering each local bed analysis, not as a Gaspari–Cohn covariance taper or as something requested in this comment.

Reviewer Comment:

Page 2 "A further major bottleneck identified in the literature…" Please add references for the literature here.

Response:

We thank the reviewer for flagging this. In revising the manuscript, we replaced this passage with a more specific discussion naming the existing data assimilation frameworks relevant to this bottleneck, namely PDAF and DART, along with a description of the integration challenges they present for ice-sheet-specific applications. The corresponding references are provided in the revised manuscript.

Reviewer Comment:

Page 3 "Existing implementations are frequently tightly coupled to...". This sentence is not necessarily true since other DA software, such as DART, have already been coupled to many atmosphere/ocean models. Here, the contribution would be more specifically focused on ice sheet modeling and its coupling with inversion methods. Please revise this sentence.

Response:

This sentence overgeneralized: DART in particular has already been coupled to many atmosphere and ocean models, so it is not accurate to describe existing DA frameworks as tightly coupled to a single architecture. We have removed this claim and replaced it with the specific gaps ICESEE fills for ice-sheet applications: joint state-parameter estimation, bounded stochastic parameter updates, and the hybrid ensemble assimilation-inversion workflow, none of which are native to general-purpose frameworks like PDAF or DART.

Reviewer Comment:

Page 4 "Stochastic ensemble filter, such as the EnKF,..." The acronym "EnKF" appears to have been used earlier as a more general ensemble Kalman filter.

Response:

We have rectified this and now refer to the stochastic ensemble filter as "stochastic EnKF" and reserve "EnKF" for the general class of methods described as the ensemble Kalman filter.

Reviewer Comment:

Page 7. 2.2.2. Ensemble Transform Kalman Filter (EnTKF). I think the acronym "ETKF" is more commonly used.

Response:

We have rectified the acronym to ETKF

Reviewer Comment:

Page 10. "In such cases, ICESEE uses the assimilated state variables as inputs to physicbased inversion procedures.." I am not convinced that inversion is necessarily a more physical approach than EnKF since inversion is also based on Bayesian approach with regularization.

Response:

We agree that describing inversion as "physics-based" implied an inaccurate contrast with EnKF, since both methods operate within a Bayesian framework and inversion with Tikhonov regularization is itself a MAP estimate under a Gaussian prior. We have replaced "physics-based inversion procedures" with "deterministic inversion procedures" (Section 2.6).

Reviewer Comment:

Page 11. "testing infrastructure" I do not see a clear description of this testing infrastructure in the manuscript.

Response:
We have added a paragraph to Section 3.2.1 (Data Assimilation Core) describing ICESEE's continuous integration setup: on every commit, GitHub Actions installs ICESEE and runs both serial and parallel data assimilation experiments (the latter described in Section 3.2.2) using the reduced-order Lorenz-96 system across all supported EnKF variants, comparing the resulting trajectories against the true simulation to catch regressions before changes are merged. We also reference the ICESEE GitHub Wiki, which documents the full setup, deployment, and testing procedures.

Reviewer Comment:
Page 14. "observation functions": What do "observation functions" refer to here?

Response:
"Observation functions" which we now refer to as observation operators referred to the routines that map model output into observation space when observations are not defined on the model grid (interpolation or projection), which was otherwise only described later in Section 3.2.6. We have replaced the vague phrase with an explicit description and cross-reference to that section.

Reviewer Comment:
Page 16. "The figure also highlights how specific..." I recommend adding specific files and directory names associated with each workflow in Figure 3.

Response:
We have revised Figure 3 to include the specific files implementing each workflow stage directly within the figure, and expanded the caption and Section 3.2.6 to clarify how these map onto ICESEE's three execution modes (serial, partial parallel, full parallel), dispatched by run_models_da.py. Where a stage is implemented differently across modes (e.g., the forecast and assimilation steps use EnKF.py in serial mode but _mpi_forecast_functions.py/_mpi_analysis_functions.py in the parallel modes), both are now shown in the figure.

Reviewer Comment:
Page 17. ISSM - Ice-sheet and Sea-level System Model

Response:
We have corrected this so that ISSM is expanded consistently throughout the manuscript.

Reviewer Comment:
Page 22. "After about 60 years, the filter ... constrained parameters that are not directly observed": I think smb are observed according to the synthetic observation paragraph on page 21. Please see my main comment and specify which variables are observed and estimated.

Response:
We agree. We have corrected this sentence: smb is directly observed.

Reviewer Comment:
Figure 4: Could you also add changes in smb field?

Response:

Yes. The revised Figure 4 now includes the SMB field together with thickness and velocity, showing the reference field, initial error, and assimilated-minus-reference differences at the selected times. The accompanying text also distinguishes SMB as a directly observed, static forcing parameter that is updated during analysis but is not dynamically propagated like thickness.

Reviewer Comment:

Page 22. "The ensemble spread around the mean remains nearly constant...": In Figure 5, the ensemble spread appears to increase with time. It would also be helpful to include the observation uncertainty in Figure 5, so readers can compare the ensemble spread with the prescribed observational errors. I guess the initial spread is not enough or the initial mean is too far from the observation?

Response:

We thank the reviewer for catching this. We identified and fixed a bug in the parallel forecast routine that had introduced artificial oscillations in the ensemble trajectories, most visible in the SMB panel of the original Figure 5; all Icepack experiments have been re-run with the fix. We also agree the initial ensemble spread was too small relative to the initial bias, and have increased the initial standard deviations (Section 4.1.1). The revised Figure 5 now shows spread contracting for thickness and velocity, while SMB narrows more gradually as a static forcing parameter updated only during analysis, rather than a dynamically propagated state variable (Section 4.1.2). We have not added observation uncertainty directly to Figure 5, as it would visually compress the existing traces; the prescribed observation standard deviations are already provided in Section 4.1.1 for direct comparison.

Reviewer Comment:

Page 25. "to simulate the long-term evolution of a marine-terminating glacier in 3D": What approximation is used for this test case? Is it a 3D model or 2D shelfy shallow approximation for this test case?

Response:

This is a vertically integrated two-dimensional (plan-view) model, consistent with Choi et al. (2025); "3D" was a mislabel and has been corrected in the manuscript (Section 4.2.2).

Reviewer Comment:

Page 26. "exponential noise is added to the bed topography": what does "exponential noise" mean here?

Response:

This has been clarified — the bed perturbations are Gaussian random fields with an exponential covariance model, not a distinct "exponential noise" type. The manuscript text has been corrected accordingly (Section 4.2.2).

Reviewer Comment:

Page 26. "and a constant offset of 284 m is subtracted from": Why is this offset applied?

Response:

The historical 284-m offset is not used in the revised experiment. The revised heterogeneous prior instead combines a −80-m background bed displacement with a smoothly varying conditional spatial anomaly of nominal amplitude 120 m, bounded between −250 and 150 m. This produces a challenging but spatially heterogeneous prior without changing the bed-evolution equations. Bed and thickness are perturbed independently, after which base, surface, and the grounded/floating mask are recomputed consistently. The manuscript now states this construction explicitly (Section 4.2.2).

Reviewer Comment:

Page 26. "the bed topography conditioned on bed observations": Why are bed observations assimilated if the initial bed is already derived using bed observations? I am also not sure why bed observations are assimilated only once. What happens if bed observations are not assimilated, or if they are assimilated multiple times?

Response:

The reviewer is correct that reusing an identical observation realization would make the later analyses difficult to interpret. In the revised experiment, the conditional prior uses a separate synthetic baseline radar survey on the same flight-line geometry, whereas the sequential analyses use independently perturbed observation realizations. Thus, prior conditioning and sequential assimilation share sampling geometry but not noise realizations. The imposed −80-m displacement, bounded heterogeneous anomalies, and smooth downstream extension further ensure that the prior does not reproduce the reference bed. We also replaced the single-survey design with repeated but temporally sparse observations: radar-like flight-line measurements are assimilated at years 2, 8, 14, 20, 24, 30, 36, 40, 46, and 50. Each bed analysis uses a 15 km radius to select the observations entering its grouped local transform; bed increments are applied only at those survey times, and the bed is held fixed between them.

Reviewer Comment:

Page 26. "In addition, a uniform offset of 50 m . . . ": Why are these offsets applied?

Response:

This description no longer reflects the current implementation. The historical additional 50-m offset to base and surface has been removed. The revised prior retains the explicitly documented −80-m bed-background displacement and heterogeneous conditional anomalies; after perturbing bed and thickness, base, surface, and the grounded/floating mask are recomputed consistently.

Reviewer Comment:

Page 26. "aircraft-based radar flight lines oriented along the x-direction": is this the x-direction or the y-direction?

Response:

This was stated incorrectly. The flight lines are oriented along the y-direction, perpendicular to the along-flow (x) direction, with adjacent lines spaced 10 km apart along x. The manuscript text in Section 4.2.2 has been corrected accordingly.

Reviewer Comment:

Page 27. "the bed topography is observed directly only once at year 2": Again, why is the bed topography assimilated only once?

Response:

Bed topography is no longer assimilated only once: it is assimilated at ten sparse survey times (years 2, 8, 14, 20, 24, 30, 36, 40, 46, and 50). The revised geometry figure includes the reference bed, the initial unassimilated bed error, and assimilated bed-error maps at years 5, 15, 30, 50, 80, and 100. It also overlays the radar-flight-line sampling locations on the reference bed. The accompanying RMSE diagnostics quantify errors in the grounded, upstream-grounded, and whole-domain beds.

Reviewer Comment:

Page 28. "the grounding line position in the simulation with assimilation (solid line)". I think this should be "without" assimilation.

Response:

Confirmed, this was a typo — corrected to "without" assimilation in the manuscript.

Reviewer Comment:

Page 29. "The bed corrections remain relatively small.." could you add changes in the bed variable in Figure 8 to support this?

Response:

Yes. The revised geometry figure now shows the reference bed, the initial unassimilated bed error, and assimilated-minus-truth bed fields at years 5, 15, 30, 50, 80, and 100, with radar sampling locations overlaid on the reference bed. Companion RMSE diagnostics report grounded, upstream-grounded, and whole-domain bed errors. We also revised the discussion to describe the remaining upstream and marginal errors rather than characterizing the corrections only as "small."

Reviewer Comment:

Page 30. "The inferred friction differences between the true and assimilated fields…": is basal friction here inferred through inversion?

Response:

Yes. In the principal hybrid experiment, basal friction is excluded from the EnKF analysis and inferred member by member with ISSM's regularized adjoint inversion using the assimilated geometry and velocity together with the synthetic velocity observations. We now state this directly in the text and distinguish it from the EBF baseline, in which friction is updated only through EnKF cross-covariances.

Reviewer Comment:

Page 30. "The results show a large improvement…." This sentence is long and hard to follow. Please consider splitting it into two or three sentences. It would also be helpful to clarify whether the reduction in basal friction error is caused directly by the inversion step, while later improvements result from improved assimilated state variables being passed into subsequent inversions.

Response:

We split the sentence and clarified the causal sequence. The first large reduction in basal-friction error is produced directly by the inversion. Later cycles provide more gradual improvement because each new inversion receives geometry and velocity fields already corrected by the preceding EnKF analysis. The controlled WBF/EBF/IBF comparison further separates this inversion effect from improvement caused by the EnKF state update alone.

Reviewer Comment:

Page 30. "Friction is set to zero beneath..." I thought the inversion was performed only on grounded ice.

Response:

The reviewer is correct that only grounded-ice friction is physically meaningful. In the implementation, the inversion control is represented on the computational mesh, but the coefficient supplied at nodes currently classified as floating is set to zero before the inversion because basal traction is dynamically inactive there. We therefore interpret and evaluate the recovered coefficient only over grounded ice. In the revised friction maps, floating regions are shaded gray to make this lack of constraint explicit; the gray is visualization only and does not alter stored output. Section 4.2.3 and the figure caption now explain this distinction.

Reviewer Comment:

Page 32. Equation (20). Please check this equation. N_p is defined as "the number of ranks pinned to each ISSM instance," which seems inconsistent because that quantity is already represented by model_nprocs. It may also be clearer to write the left-hand side as "total MPI ranks" rather than "ICESEE ranks."

Response:

Confirmed, this was mislabeled. N_p denotes the number of concurrent ISSM instances (ensemble members run in parallel), not ranks per instance. The equation and left-hand side have been corrected to *total MPI ranks* $= N_p \times (\text{model\_nprocs} + 1)$.

Reviewer Comment:

Page 32. "This decline is primarily due to...": What about the number of nodes? Can the increase in number of nodes also contribute to the decline in speed?

Response:

Yes, this is a fair point. Beyond 80 ranks, the configuration also spans multiple nodes (Table 1), introducing inter-node communication over InfiniBand in addition to the growing MATLAB-instance overhead. We have revised the text to note both factors.

Reviewer Comment:

Page 35. "leveraging physic-based cost function..". Again, I do not think this method is more physical than EnKF DA.

Response:

Agreed. "Physics-based cost-function minimization" has been replaced with "traditional inversion" (or "regularized adjoint inversion," where specificity is useful) throughout the manuscript. We do not characterize inversion as intrinsically more physical than EnKF.

Summary

We thank the reviewer for these detailed and constructive comments. The manuscript now uses dimensionally consistent EnKF notation and an explicit memberwise formulation for nonlinear observation operators; corrects the relationship between perturbation design and covariance localization; documents the grouped local bed analysis added to the revised ISSM configuration; documents the testing infrastructure and observation mappings; identifies observed, EnKF-updated, and inversion-updated quantities in both applications; replaces the single bed survey with ten sparse local bed analyses using separately perturbed observation realizations and a deliberately displaced heterogeneous prior; adds SMB and bed-change diagnostics; and includes a controlled WBF/EBF/IBF comparison that quantifies the contribution of inversion-based basal-friction recovery. We also corrected the implementation issue affecting the original Icepack ensemble-spread behavior and revised the related figures and discussion.

References

J. L. Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, Sept. 1975. ISSN 0001-0782. doi: 10.1145/361002.361007. URL `https://doi.org/10.1145/361002.361007`.

G. Evensen. The ensemble kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4):343–367, 2003. doi: 10.1007/s10236-003-0036-9. URL `https://doi.org/10.1007/s10236-003-0036-9`.

P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, A. Vijaykumar, A. P. Bardelli, A. Rothberg, A. Hilboll, A. Kloeckner, A. Scopatz, A. Lee, A. Rokem, C. N. Woods, C. Fulton, C. Masson, C. Häggström, C. Fitzgerald, D. A. Nicholson, D. R. Hagen, D. V. Pasechnik, E. Olivetti, E. Martin, E. Wieser, F. Silva, F. Lenders, F. Wilhelm, G. Young, G. A. Price, G.-L. Ingold, G. E. Allen, G. R. Lee, H. Audren, I. Probst, J. P. Dietrich, J. Silterra, J. T. Webber, J. Slavič, J. Nothman, J. Buchner, J. Kulick, J. L. Schönberger, J. de Miranda Cardoso, J. Reimer, J. Harrington, J. L. C. Rodríguez, J. Nunez-Iglesias, J. Kuczynski, K. Tritz, M. Thoma, M. Newville, M. Kümmerer, M. Bolingbroke, M. Tartre, M. Pak, N. J. Smith, N. Nowaczyk, N. Shebanov, O. Pavlyk, P. A. Brodtkorb, P. Lee, R. T. McGibbon, R. Feldbauer, S. Lewis, S. Tygier, S. Sievert, S. Vigna, S. Peterson, S. More, T. Pudlik, T. Oshima, T. J. Pingel, T. P. Robitaille, T. Spura, T. R. Jones, T. Cera, T. Leslie, T. Zito, T. Krauss, U. Upadhyay, Y. O. Halchenko, Y. Vázquez-Baeza, and S. . . Contributors. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods*, 17(3):261–272, 2020. doi: 10.1038/s41592-019-0686-2. URL `https://doi.org/10.1038/s41592-019-0686-2`.