Response to Reviewer 1

The manuscript is well prepared and presents a valuable contribution to the field. I have a few minor suggestions that should be incorporated into the revised version to further enhance its clarity, rigor, and overall impact.

We thank the reviewer for the careful evaluation of our manuscript and for the constructive comments. We appreciate the positive assessment of the relevance of the topic and the Antarctic-ice-shelf-wide application. In the revised manuscript, we have addressed the reviewer's suggestions regarding methodological clarity, interpretation, and presentation. Please find below our point-by-point responses. All line numbers cited below refer to the clean revised manuscript, not to the track-changes version.

1. The manuscript would benefit from stating even more consistently, particularly in the Abstract, Introduction, and Conclusions, that the retrieved grain size represents an effective ASCAT-conditioned parameter. Likewise, the FAC-related analysis should be explicitly framed as a proof of concept rather than an operational retrieval product.
   In the revised manuscript, we have made this framing more consistent in the most visible parts of the paper. In the Abstract, we now state that the study aims to "infer an effective firn grain-size parameter across Antarctic ice shelves, constrained by 15 years (2007--2021) of active microwave observations from the C-band Advanced Scatterometer (ASCAT)" (lines 4--8). We also frame the backscatter--FAC analysis as a proof of concept, stating that "by statistically reducing grain-size-driven scatter after inversion, we present a proof of concept for a more interpretable backscatter--FAC relationship" that may support future development and evaluation of FAC-oriented diagnostic retrieval frameworks (lines 16--20). In the Introduction, we similarly clarify that the study derives "a more interpretable backscatter--FAC relationship by statistically reducing grain-size-driven scatter, as a proof of concept for FAC-oriented interpretation of ASCAT" and that "the retrieved grain size is interpreted as an effective parameter conditioned by ASCAT observations within the adopted FDM--SMRT framework" (lines 60--64). In the Conclusions, we now describe the retrieved quantity as "an ASCAT-constrained grain-size retrieval parameter" and state that the optimized grain size should be interpreted as an "effective, ASCAT-conditioned parameter" whose confidence is regime-dependent and whose independent validation remains limited (lines 649--669).

2. Section 3 contains all the necessary methodological details; however, the overall workflow could be presented more clearly. In particular, it would be helpful to state explicitly, ideally at the beginning of the Methods section or in a brief schematic

overview, that firn stratigraphy (including density, temperature, liquid water content, and layering) is prescribed from IMAU-FDM; that the inversion estimates only a single column-integrated grain-size parameter; and that the ANOVA and grain-size standardization are applied as post-inversion diagnostic analyses. This clarification would greatly improve the transparency of the methodological framework without altering the underlying approach.

In the revised manuscript, we have added a short overview paragraph at the beginning of Sect. 3 to clarify the methodological workflow. The new text states that "firn stratigraphy is prescribed from IMAU-FDM, including layer thickness, density, temperature, and liquid water content" and that "the inversion estimates only a single column-wide grain-size parameter representative of the upper firn column sampled by ASCAT" (lines 127--130). The same paragraph also clarifies that "the ANOVA and grain-size standardization are then applied as post-inversion diagnostic analyses to interpret interannual backscatter variability and the backscatter--FAC relationship" (lines 130--132). We also revised the roadmap paragraph immediately afterward to clarify the purpose of each methodological step. Specifically, the revised text now describes the sequence from the baseline comparison, to targeted FDM--SMRT sensitivity experiments, to the core grain-size inversion, and finally to the two post-inversion diagnostic analyses: ANOVA variance partitioning and grain-size standardization (lines 133--140). These changes improve the transparency of the workflow without altering the underlying approach.

3. The manuscript appropriately notes that the excellent agreement between the optimized FDM-SMRT backscatter and ASCAT should be interpreted as internal consistency within the adopted modelling framework, rather than as an independent validation. Given the strength of the reported fit, however, this important caveat should be highlighted more prominently in the main Results section and/or in the caption of Appendix Figure A1, so that readers encounter this interpretive context alongside the performance metrics.

   In the revised manuscript, we have made this caveat more visible alongside the performance metrics. In Sect. 4.3, we now state that "this agreement should be interpreted as internal consistency within the adopted framework and not as independent validation of the retrieval" (lines 341--343). We also clarify that the comparison between IMAU-FDM grain size and the optimized grain-size parameter should not be interpreted as "a direct comparison between modeled grain size and physical truth," but rather as framework-dependent differences between the IMAU-FDM-prescribed firn state, the FDM--SMRT simulated scattering response, and the ASCAT-conditioned effective parameter retrieved by the inversion (lines 341--346).

The same clarification has been added to the caption of Appendix Fig. A1, where we state that the agreement reflects internal consistency within the adopted framework, not independent validation (Appendix A, Fig. A1 caption, around lines 680--683). These changes ensure that readers encounter this interpretive context directly alongside the fit statistics and comparison figure.

4. The manuscript correctly acknowledges that the ANOVA should be interpreted as a conditional diagnostic within the adopted framework, since the optimized grain size is itself derived from ASCAT observations. However, some instances of the term “explained variance” may inadvertently imply a stronger degree of independence or causality than intended. To enhance interpretive precision, I recommend consistently adopting terminology such as “variance partitioning within the adopted framework” in Section 3.4, Section 4.4, and Appendix B, and, where appropriate, in the corresponding figure captions. This refinement would strengthen the rigor of interpretation without requiring additional analyses.
   In the revised manuscript, we have revised the ANOVA terminology throughout the relevant sections to describe the analysis as conditional variance partitioning within the adopted framework. In Sect. 3.4, the subsection is now titled “Variance partitioning of interannual ASCAT variability within the adopted framework”, and the text states that ANOVA is used “as a variance-partitioning tool” to estimate “fractional contributions associated with FAC, grain size, their interaction, and a residual term” (lines 238--247). The same section further states that this analysis should be interpreted as “a conditional diagnostic within the adopted framework rather than as an independent attribution of ASCAT variability” (lines 245--247). In Sect. 4.4 and the Fig. 5 caption, we similarly describe the results in terms of “variance partitioning” and “fractional contributions” associated with grain size, FAC, their interaction, and residual variability (lines 381--400). Appendix B has also been revised to use the same terminology, including “ANOVA-based variance partitioning”, “fractional variance contribution”, and “partitioned ASCAT variability within the adopted framework” (lines 684--708).

5. One of the manuscript’s key strengths is its clear recognition of the conditions under which the inversion is more or less informative. For example, grain-size identifiability appears stronger in certain regimes and weaker over shallow-slope plateau regions, while unmodelled near-surface processes may become increasingly important in dry, high-FAC environments. Similarly, the FAC-related interpretation is appropriately presented as a proof of concept. To further enhance clarity for readers, I recommend adding a concise synthesis paragraph, either at the end of the Results section or at

the beginning of the Limitations section, summarizing where confidence in the retrieval is relatively high, where it is lower, and the physical reasons underlying these differences. This would improve accessibility and interpretability without altering the scientific content.

In the revised manuscript, we have added a concise synthesis paragraph at the beginning of Sect. 5.1, now titled "Methodological scope and future development", to summarize where the retrieval is relatively more or less informative. The new text states that confidence is highest "where the sensitivity experiments indicate sufficient grain-size leverage, particularly where the σ^0-r response remains on the steep part of the curve and where the 2-D response surface is steep in grain size relative to density/FAC (Fig. 2b--c)" (lines 535--540). We further state that this includes many dry-snow and intermediate-FAC regions for which the optimized grain size remains physically consistent with the adopted FDM--SMRT framework (lines 539--540). The same paragraph also clarifies the lower-confidence regimes: "Confidence is lower in shallow-slope ('plateau') regimes, particularly in strongly depleted firn, where Fig. 2b--c shows that the C-band response becomes less sensitive to grain size and more strongly governed by density/FAC" (lines 540--542). Finally, we note that some dry, high-FAC regions may also be influenced by unmodeled near-surface processes and reiterate that the optimized grain size is best interpreted as a regime-dependent effective parameter, while the FAC-related interpretation remains a proof of concept rather than a standalone operational product (lines 542--546). These additions directly address the request for a concise synthesis of where the inversion is more or less informative.

Response to Reviewer 2

I debated between major and minor for this paper but ultimately decided on minor. I believe the research objective is extremely important. Using ASCAT data to constrain grain size evolution on Antarctic ice shelves addresses a real gap in firn modeling. The knowledge gained about how grain size influences microwave backscatter would genuinely help progress related research on firn permeability and meltwater retention.

That said, there are substantive issues to address. The single-layer grain-size assumption needs more discussion given vertical processes like refreezing. The FAC retrieval ambiguity at low backscatter values requires clarification. Several methodological aspects such as temporal aggregation choices, computational feasibility, and some interpretive framing etc can be strengthened.

The basic idea of the conducted research is good and the contributions are valuable. However, I strongly recommend addressing the reviews in the subsequent revisions before accepting it for the publication.

We thank the reviewer for the careful reading of our manuscript and for the constructive comments. We appreciate the positive assessment of the study objective and the recognition that using ASCAT observations to constrain firn grain-size evolution addresses an important gap in firn modelling. In the revised manuscript, we have addressed the reviewer's comments by clarifying the methodological assumptions, refining the interpretation of the ASCAT-conditioned optimized grain-size parameter, and better defining the scope of the FAC-oriented analysis. Please find below our point-by-point responses. All line numbers cited below refer to the clean revised manuscript, not to the track-changes version.

Please find attached my detailed comments.

6. **Column-wise grain-size assumption may mask physically meaningful heterogeneity.**
   The optimization treats the entire ~20 m ASCAT-sampled column as having a single uniform grain size. While this is justified by under-representation concerns (line 513-514), it creates a conceptual tension: if you are optimizing to match observations, why not allow some vertical structure? The authors acknowledge depth-dependent grain size matters (lines 520-522) but dismiss multi-layer retrieval as causing "optimizer collapse". This is a significant limitation that deserves more than brief mention.

- What kinds of physical processes are likely hidden by this simplification (e.g., refreezing layers, metamorphism gradients)?
- Whether the retrieved parameter should be interpreted as an effective value weighted toward certain depths based on scattering physics?
- What if we completely ignore the vertical structure, how much confidence will you put on your spatial patterns?

In the revised manuscript, we have clarified the interpretation and limitations of the single column-wide grain-size parameter. In Sect. 3.3, we now state that "the inversion targets a single, column-wide, uniform-with-depth grain-size parameter representative of the upper ~20 m firn column sampled by ASCAT" and that this parameter "should be interpreted as an ASCAT-constrained retrieval parameter within the adopted FDM--SMRT framework, rather than as a direct microphysical grain-size measurement, a depth-resolved grain-size profile, or a classical effective-medium property derived independently of the observations" (lines 219--222). We also explicitly define "effective" as referring to "the ASCAT-constrained parameter inferred through the inversion" (lines 223--225).

We have also expanded Sect. 5.1 to discuss the physical implications of this simplification. The revised text states that "a single column-wide grain-size parameter cannot represent vertical heterogeneity in firn microstructure, including near-surface metamorphic gradients, buried refrozen layers, ice lenses, or depth-dependent grain growth associated with firn densification and melt--refreeze events" (lines 597--604). To address the reviewer's question about scattering physics, we further clarify that the optimized value is "likely weighted toward the parts of the upper firn column that contribute most strongly to C-band volume scattering" because the scattering physics depend on the interaction between microwave frequency, grain size, density, temperature, and layer structure (lines 597--604). Consequently, we place greatest confidence in large-scale, footprint-scale spatial patterns, particularly where the forward sensitivity analysis indicates stronger identifiability, while avoiding interpretation of the retrieval as local or vertically resolved grain-size truth (lines 598--600).

Finally, we clarify that even simple depth-varying alternatives would remain underconstrained with the present single-frequency ASCAT setup unless additional assumptions or external constraints were introduced. We therefore state that relaxing this limitation would require expanding the observation space through multi-sensor, multi-frequency, and multi-angular constraints, which could support future joint inversion frameworks to better constrain vertical grain-size heterogeneity (lines 608--615).

7. **ANOVA variance partitioning assumes additivity that may not hold.**
   The manuscript partitions backscatter variance into FAC, grain size, their interaction, and residual components (Section 3.4). The assumption of additive main effects with a small interaction term (Fig. 5a) is presented as finding, but this has important implications. If the relationship between backscatter and these drivers is not truly additive, then partitioning may be misleading regardless of whether it is "conditional." Consider explicitly stating that variance partitioning assumes linear(izable) relationships and discuss what non-additive physics might look like.
   In the revised manuscript, we have clarified that the ANOVA is a first-order statistical diagnostic and should not be interpreted as a complete physical decomposition of firn--microwave scattering processes. In Sect. 3.4, we now state that "ANOVA is used here as a variance-partitioning tool" to decompose year-to-year backscatter variability into "fractional contributions associated with FAC, grain size, their interaction, and a residual term" (lines 240--242). We further clarify that "this decomposition is statistical and assumes that the variability can be represented, to first order, by approximately additive contributions from FAC, grain size, their interaction, and a residual term" and that it "should therefore not be interpreted as implying that the underlying firn--microwave scattering physics are strictly additive" (lines 242--245). Because the optimized grain size is inferred by fitting FDM--SMRT simulations to ASCAT observations, we also state that the ANOVA should be interpreted as "a conditional diagnostic within the adopted framework rather than as an independent attribution of ASCAT variability" (lines 245--247).
   We have also expanded Sect. 5.1 to clarify what non-additive behaviour could involve. The revised text states that strongly non-additive behaviour may not be fully captured by the ANOVA and gives examples: "when grain size and FAC co-evolve during melt--refreezing transitions", "when ice saturation causes threshold-like scattering changes", or "when layering and near-surface structure make the sensitivity to one variable depend on the state of another" (lines 568--575). We further clarify that "the variance fractions should not be interpreted as uniquely separable physical controls" and that "some degree of confounding may remain when FAC and grain size co-vary, particularly during melt--refreezing transitions" (lines 575--579).

8. **The "proof of concept" FAC retrieval curve has an unresolved ambiguity problem.**
   The adapted inverted U-curve (Fig. 6) is presented as a way to infer FAC from backscatter once grain size is controlled. But the two-branch ambiguity at low backscatter values remains fundamentally problematic.
   The manuscript mentions using "simple, independent priors" (lines 470-472) but does not clarify how often these would be needed or whether they are actually

available across the Antarctic domain. In addition, are there any other ways to deal with this ambiguity?

In the revised manuscript, we have clarified that the adapted inverted-U relationship should be interpreted as an indicative proof of concept rather than an operational FAC retrieval curve. We now state that it "remains conditioned on the adopted FDM--SMRT framework, the statistical standardization of grain size, and the assumptions used to distinguish firn regimes", and that "additional validation and treatment of ambiguities would be required before such a relationship could be used for standalone FAC retrieval" (lines 504--508). We also clarify that converting observed backscatter to FAC would require "explicit handling of the two-branch ambiguity at low returns with independent priors" (lines 508--510).

We have also clarified where independent priors are needed and how they would help resolve the ambiguity. The revised text states that, "at low observed backscatter, however, the inverted-U shape yields two plausible states: a depleted, low-FAC branch and a dry, high-FAC branch" (lines 514--515). We further state that "independent priors are therefore mainly required in this low-backscatter regime, whereas the intermediate, locally monotonic part of the curve is less ambiguous" (lines 515--516).

To make these priors more concrete, we now list examples including "melt-pond presence from optical imagery", "melt-day information from ASCAT or climate-model analyses", "MoA climatology", field-based grain size where available, and guidance from firn models (lines 516--519). We also clarify that these constraints are "not uniformly or consistently available across the Antarctic domain", and that their absence limits the use of the adapted curve as a standalone FAC retrieval (lines 519--520). Finally, we state that once the depleted- versus dry-side branch is independently identified, FAC can be read consistently from the adapted curve (lines 520--521).

We further note that unmodeled surface processes, including roughness, snowdrift, and accumulation variability, can shift ASCAT independently of grain size or FAC, and that these effects are beyond the scope of the present proof-of-concept analysis but indicate how such an assessment could be developed and evaluated in future work (lines 521--524).

9. **Comparison between IMAU-FDM and optimized grain size may conflate model error with physical reality.**

   The manuscript presents discrepancies between IMAU-FDM grain sizes and optimized values as evidence that the model needs improvement (e.g., lines 540-548). But this comparison assumes:

The optimization retrieved a "true" value, yet we have established it is an effective, ASCAT-conditioned parameter. IMAU-FDM errors are solely due to grain-size physics when they could also reflect errors in temperature history, accumulation, or other drivers that feed into the model.

Consider reframing: instead of "model underestimates/overestimates," discuss whether discrepancies indicate inconsistency between modeled and observed scattering behavior, which could stem from multiple sources.

In the revised manuscript, we have reframed the comparison between IMAU-FDM grain size and the ASCAT-conditioned optimized grain-size parameter more carefully. In Sect. 4.3, we now state that because "the optimized grain-size parameter is itself inferred by fitting the model to ASCAT", the close agreement between optimized FDM--SMRT simulations and ASCAT "should be interpreted as internal consistency within the adopted framework and not as an independent validation of the retrieval" (lines 341--343).

We also clarify that "the comparison between IMAU-FDM grain size and the optimized grain-size parameter should not be interpreted as a direct comparison between modeled grain size and physical truth" (lines 343--344). Instead, we state that deviations between the two indicate "framework-dependent differences between the IMAU-FDM-prescribed firn state, the FDM--SMRT simulated scattering response, and the ASCAT-conditioned effective parameter retrieved by the inversion" (lines 344--346).

To avoid implying that the optimized values represent true physical grain size, we have also revised the language throughout Sect. 4.3. Rather than stating that IMAU-FDM underestimates or overestimates grain size, we describe where "IMAU-FDM grain-size estimates are generally lower than the ASCAT-conditioned optimized effective values" (lines 347--348), where they are "substantially higher than the ASCAT-conditioned optimized effective values" in ice-saturated areas of Amery (lines 355--357), and where they are "substantially lower than the ASCAT-conditioned optimized values" in relatively drier and colder parts of Amery and Larsen-C (lines 357--359). We further state that the consistency between IMAU-FDM estimates and ASCAT-conditioned optimized values is "strongly regime-dependent, particularly with respect to the degree of ice saturation and associated firn processes" (lines 378--380).

In the Discussion, we also clarify that the observed differences "should not be interpreted as direct evidence that the optimized retrieval represents physical truth or that IMAU-FDM errors arise solely from grain-size physics" (lines 445--447). Instead, we state that they indicate framework-dependent differences between modeled firn-state/scattering behaviour and the ASCAT-conditioned effective parameter within the adopted FDM--SMRT framework, and that such differences may arise from the IMAU-FDM grain-size parameterization affecting the simulated scattering response (lines 447--449). We then discuss possible contributors to these differences, including the limited effect of refreezing on grain size in IMAU-FDM and limited calibration in depleted conditions (lines 450--459).

10. **The backscatter-FAC relationship analysis does not account for frequency dependence.**
    C-band ASCAT is used throughout, but microwave scattering physics is inherently frequency-dependent (as also noted in lines 523-525). While you are not expected to do multifrequency inversion now, consider:
    Briefly acknowledging that FAC interpretation may differ at other frequencies. I would also clarify whether C-band's particular sensitivity profile limits or biases certain conclusions about grain-size versus FAC control.
    In the revised manuscript, we have clarified that the FAC-oriented interpretation developed here is specific to C-band ASCAT. We now state that it "should not be assumed to transfer directly to other microwave frequencies" and that "the relative sensitivity of backscatter to grain size, density, and FAC depends on frequency-dependent penetration depth and scattering behaviour" (lines 526--528). To make the link to our own sensitivity analysis explicit, we further state that "the C-band sensitivity pattern shown in Fig. 2b--c therefore reflects the ASCAT-specific balance between grain-size-related volume scattering and density/FAC-related saturation effects", and that this balance should not be assumed to transfer directly to other microwave frequencies (lines 528--530). Finally, we identify multi-frequency observations as an important future direction for testing whether the grain-size and FAC sensitivities identified here are consistent across different penetration depths and scattering regimes (lines 530--532).

11. **Temporal aggregation choices lack justification.**
    The study uses winter-mean values throughout. This choice has implications:
    - Why winter specifically? Does this avoid melt-related physics that might be more relevant to hydrofracture concerns?
    - What seasonal variability is being lost by aggregating to single values per year/season?

    In the revised manuscript, we have clarified the reason for using winter-mean ASCAT backscatter. In Sect. 2.1, we now state that "to minimize the direct influence of active melt, transient liquid water, and short-lived seasonal effects on the ASCAT signal, we focus on winter (June, July, and August, or JJA) mean backscatter values" for each year between 2007 and 2021 (lines 89--90). We further state that this choice is intended "to isolate a more stable interannual signal related to firn microstructure", yielding 15 annual values per pixel for the temporal analysis (lines 90--91). This aggregation means that melt-season timing, short-lived melt--refreezing events, seasonal transitions, and sub-seasonal backscatter variability are not resolved in the present analysis. The framework therefore focuses on stable interannual firn-structure contrasts rather than seasonal firn evolution.

12. The "effective grain-size" parameter definition is circular.

The manuscript defines the optimized value as an "ASCAT-conditioned effective parameterization" (line 436).

This framing is honest but creates interpretive issues:

- An "effective" parameter in radiative transfer typically means something derived from upscaling microphysical properties, but here it's being inverted from observations.
- Consider clarifying whether this should be thought of as a retrieval constrained by ASCAT, or an effective medium property that happens to match ASCAT when plugged into the forward model.

In the revised manuscript, we have clarified that the optimized quantity should be understood as a retrieval constrained by ASCAT observations, not as an independently derived effective-medium property. In Sect. 3.3, we now state that the retrieved parameter "should be interpreted as an ASCAT-constrained retrieval parameter within the adopted FDM--SMRT framework, rather than as a direct microphysical grain-size measurement, a depth-resolved grain-size profile, or a classical effective-medium property derived independently of the observations" (lines 219--222). We also explicitly define the meaning of "effective" in this study. The revised manuscript states that "effective" refers to "the ASCAT-constrained parameter inferred through the inversion" (line 225). This clarification distinguishes the retrieved parameter from a classical radiative-transfer effective-medium property and frames it as an ASCAT-conditioned parameter within the adopted forward-model assumptions.

13. Comparison with previous work lacks depth.

The manuscript cites Alley et al. (2018) and Dattler et al. (2024) but doesn't systematically compare:

- Whether their FAC-inference approach would have produced similar results if applied to the same domain/data.

- What advantages / disadvantages the current grain-size-constrained approach has over purely empirical backscatter - FAC relationships.
- How this work advances beyond simply confirming that ASCAT is sensitive to both variables (which was already known).

In the revised manuscript, we have strengthened the comparison with previous work in Sect. 5. We now clarify how our framework differs from Dattler et al. (2024): while both studies couple firn-model output with SMRT to interpret microwave signals, our study uses "active C-band ASCAT backscatter", whereas Dattler et al. used "passive AMSR-2 brightness temperatures for physics-based melt detection" (lines 417--420). We also clarify that our analysis performs "an ice-shelf-wide, per-pixel, per-year inversion to optimize grain size", whereas Dattler et al. focused on 13 sites on the Larsen-C Ice Shelf (lines 420--422).

We also compare our approach with more empirical or observational ASCAT-based interpretations of firn state, such as the backscatter--melt relationships explored by Alley et al. (2018). The revised manuscript states that the present study "complements such approaches by explicitly treating grain-size variability within a physics-based FDM--SMRT framework before examining the backscatter--FAC relationship" (lines 422--425).

Finally, we have added a clearer statement of the advantages and limitations of our approach relative to empirical backscatter--FAC relationships. We state that the advantage is "greater physical interpretability by identifying how grain-size-related variability affects ASCAT backscatter and FAC-oriented interpretation", while the corresponding limitation is that "the results depend more strongly on the adopted forward-model assumptions" (lines 425--428). We further clarify that the main advance is not simply confirming that ASCAT is sensitive to both grain size and FAC, but showing "how a physics-based inversion can constrain grain-size-related variability and thereby support more interpretable FAC-oriented analyses across Antarctic ice shelves" (lines 428--430).

14. **Surface roughness effects are acknowledged but not tested.**

The manuscript mentions unmodeled surface processes contribute to residual variance in dry, high-FAC zones (lines 497-502) and proposes adding a rough-surface scattering module as future work. But, without testing whether including this would actually reduce residuals, the attribution remains speculative. Can you clarify this part?

In the revised manuscript, we have softened the interpretation of surface roughness and related near-surface processes. In Sect. 5, we first flag that "unmodeled surface processes, including roughness, snowdrift, and accumulation variability, can shift

ASCAT independently of grain size or FAC" and that these effects are beyond the scope of the present proof-of-concept analysis (lines 521--524).

In Sect. 5.1, we now state more cautiously that some dry shelves at the very high-FAC end "exhibit elevated residual variability, which may be consistent with unmodeled surface-process controls" such as roughness, snowdrift, and accumulation variability, and that the residual "may indicate unrepresented near-surface processes" rather than a demonstrated mechanism (lines 564--567).

We further clarify that "the residual contribution in dry, high-FAC regimes may partly reflect surface processes rather than microstructure or FAC alone" and that "this interpretation remains a hypothesis within the present framework, because surface roughness, snowdrift, and accumulation variability are not explicitly tested here" (lines 579--582). We also state that these processes are plausible contributors because they can modulate ASCAT backscatter independently of grain size or FAC, but that "their relative importance cannot be quantified from the present analysis alone" (lines 582--584).

Finally, we outline how this hypothesis could be tested in future work. The revised manuscript states that independent constraints on surface roughness and drift/accumulation patterns, such as in-situ roughness surveys, azimuthal anisotropy proxies for sastrugi and microtopography, or snow-transport diagnostics, could be combined with a rough-surface scattering module in SMRT to test whether explicitly representing surface scattering reduces the residual contribution in dry, high-FAC regimes (lines 584--589).

15. The ASCAT footprint size creates spatial averaging issues.

At 27 km resolution, each grid cell represents significant spatial heterogeneity:

- Could different parts of the footprint be in different regimes, making the single grain-size assumption even less valid?
- Is there any literature of whether regime transitions within footprints affect results?

In the revised manuscript, we have clarified that the ASCAT-constrained grain-size parameter should be interpreted at the footprint scale. In Sect. 5.1, we now state that "the inversion is performed at the 27 km RACMO/IMAU-FDM grid scale", and that the retrieved quantity should therefore be interpreted as "a spatially averaged, footprint-scale effective parameter" (lines 617--620).

We also explicitly acknowledge that a single ASCAT footprint may include multiple firn regimes. The revised manuscript states that "different parts of a single ASCAT footprint may span different firn regimes, especially near transitions among dry, intermediate, and depleted firn" (lines 620--622). We further clarify that such "subfootprint regime mixing can make the single grain-size assumption less representative of local variability", although this effect is not explicitly quantified in the present study (lines 622--623).

Finally, we have added a pathway for testing this limitation in future work. We now state that heterogeneity in grain size, density, stratigraphy, anisotropy, and near-surface structure can influence microwave scattering and its interpretation, and that future work could test the importance of footprint-scale mixing by applying the framework to higher-resolution active microwave C-band observations, such as Sentinel-1, and comparing the resulting effective parameters with those retrieved from ASCAT (lines 623--627).

16. **The manuscript does not discuss potential aliasing between FAC and grain size.** While ANOVA suggests small interaction effects, this is based on a specific dataset and framework. If both variables are changing simultaneously (e.g., during melt seasons), could their effects be partially aliased even if the forward model treats them as separate?

    In the revised manuscript, we have clarified that the ANOVA should be interpreted as a conditional diagnostic within the adopted framework. In Sect. 3.4, we now state that ANOVA is used "as a variance-partitioning tool" to decompose year-to-year backscatter variability into "fractional contributions associated with FAC, grain size, their interaction, and a residual term" (lines 240--242). We also state that this decomposition is statistical and assumes that variability can be represented, to first order, by approximately additive contributions from FAC, grain size, their interaction, and a residual term, and therefore "should not be interpreted as implying that the underlying firn--microwave scattering physics are strictly additive" (lines 242--245). Because the optimized grain size is itself inferred by fitting FDM--SMRT simulations to ASCAT observations, we further state that the analysis should be interpreted as "a conditional diagnostic within the adopted framework rather than as an independent attribution of ASCAT variability" (lines 245--247).

    We have also expanded Sect. 5.1 to address the possibility of aliasing or confounding between FAC and grain size. We now state that the ANOVA-based variance partitioning should be interpreted "as a first-order statistical diagnostic rather than as a complete physical decomposition of the firn--microwave system" (lines 568--571). We further clarify that strongly non-additive behaviour may not be fully captured, for example "when grain size and FAC co-evolve during melt--refreezing transitions", "when ice saturation causes threshold-like scattering changes", or when layering and near-surface structure make the sensitivity to one variable depend on the state of another (lines 571--575).

Finally, we explicitly state that "the variance fractions should not be interpreted as uniquely separable physical controls", and that "the small interaction term should not be taken as proof that aliasing between FAC and grain size is absent" (lines 575--577). We also note that some degree of confounding may remain when FAC and grain size co-vary, particularly during melt--refreezing transitions, and that the partitioning should therefore be interpreted within the sampled interannual variability and the adopted inversion framework (lines 577--579).

17. **The manuscript does not clearly distinguish between diagnostic and prognostic capabilities.**

The work is presented partly as showing how ASCAT could contribute to FAC-oriented monitoring (line 551).
However:

- Diagnostic capability (describing current state) does not necessarily imply prognostic capability (predicting future changes or responding to perturbations).
- Consider clarifying what kinds of applications the approach is actually suited for versus what would require additional development.

In the revised manuscript, we have clarified the distinction between diagnostic and prognostic capabilities throughout the paper. In the Abstract, we now describe the FAC-related component as supporting future development and evaluation of "FAC-oriented diagnostic retrieval frameworks" and frame the broader contribution as supporting "large-scale diagnostic monitoring of firn-state variability" (lines 17--20). In the Introduction, we similarly describe the adapted backscatter--FAC relationship as "a proof of concept for FAC-oriented interpretation of ASCAT" that informs future development of FAC-oriented diagnostic retrieval approaches, and we clarify that the retrieved grain size is interpreted as an effective parameter conditioned by ASCAT observations within the adopted FDM--SMRT framework (lines 60--64).

We have also added an explicit statement in Sect. 5.1 that the present framework should be interpreted as diagnostic rather than prognostic. The revised manuscript states that "the present framework should be interpreted as diagnostic rather than prognostic" and that it is suited to "interpreting present firn-state variability", identifying regimes where ASCAT provides stronger or weaker grain-size constraints, and supporting FAC-oriented interpretation once grain-size-driven scatter is reduced (lines 628--630). We further clarify that the framework "does not by itself predict future firn evolution or the response of firn conditions to future climate perturbations" (lines 630--631).

Finally, we now state that prognostic applications would require additional development, including stronger external validation, explicit coupling to firn-evolution modelling, and treatment of the unresolved processes and ambiguities discussed in the manuscript (lines 631--633). In the Conclusions, we also refer to ASCAT as a complementary diagnostic and monitoring metric and state that the framework supports large-scale diagnostic monitoring of firn-state variability (lines 668--672).

18. **The optimized grain sizes are presented without assessing whether they represent physically realistic firn states.**
    - In ice-saturated regions (FAC ≈ 0), what range of grain sizes should be expected? Are the retrieved values consistent with ice crystal growth physics under those conditions?
    - The IMAU-FDM model "lacks a maximum grain size constraint" in depleted areas (line 547). Does this imply optimized values could also become unphysically large, or does the ASCAT inversion naturally constrain them?
    - In high-FAC, dry-snow regions where FAC is >20 m and temperature is very low, are the retrieved small grain sizes (<0.5 mm) consistent with observed metamorphism rates in such conditions?

In the revised manuscript, we have expanded the discussion of physical plausibility by firn regime. For high-FAC, dry-snow regions, we now state that IMAU-FDM grain size and the ASCAT-conditioned optimized effective parameter show broad agreement, with both indicating relatively small grain sizes of approximately <0.5 mm. We further state that these small values are "qualitatively plausible for cold, dry firn with limited melt influence and slower metamorphic evolution", although they should not be interpreted as independently validated grain-size benchmarks (lines 440--443).

For melt-affected and depleted regions, we now discuss why the physical interpretation is more uncertain. We state that the larger differences between IMAU-FDM grain-size estimates and ASCAT-conditioned optimized effective values "should not be interpreted as direct evidence that the optimized retrieval represents physical truth or that IMAU-FDM errors arise solely from grain-size physics" (lines 445--447). Instead, we clarify that they indicate framework-dependent differences between modeled firn-state/scattering behaviour and the ASCAT-conditioned effective parameter within the adopted FDM--SMRT framework (lines 445--450).

We also explain two mechanisms that may contribute to these differences in melt-affected and depleted firn: the minimal effect of refreezing on grain size in IMAU-FDM and the lack of calibration in depleted conditions (lines 451--452). Specifically, we now state that the IMAU-FDM grain-growth parameterization assumes that refreezing affects grain size only below 0.25 mm, and that this threshold "is therefore not a physical upper limit on grain growth, but a simplifying parameterization" (lines 452--456). We further note that refrozen layers may produce an ASCAT scattering response consistent with a coarser effective microstructure than represented by the IMAU-FDM grain-size parameterization, particularly in melt-affected regions such as Larsen-C (lines 457--460).

We have also clarified the implication of the missing maximum grain-size constraint in IMAU-FDM. The revised manuscript states that, in old, depleted firn such as the Amery Ice Shelf, grains can continue to grow indefinitely in the model because "there is no prescribed maximum grain size in IMAU-FDM" (lines 465--467). However, we explicitly state that this absence of an upper bound applies to the IMAU-FDM parameterization and "does not imply that the ASCAT-conditioned optimized parameter can grow freely in the same way", because the inversion is constrained by the observed ASCAT backscatter and the FDM--SMRT forward-model sensitivity (lines 467--471). We therefore clarify that the optimized parameter is not unconstrained in the same way as the model parameterization.

Finally, we have strengthened the discussion of external plausibility checks. We now state that additional calibration data from depleted regions would be required to assess the physical plausibility of both modeled grain-size evolution and the ASCAT-conditioned optimized effective parameter in these regimes (lines 475--477). Where limited field observations are available, we compare the retrieved effective grain size with in-situ measurements near Dumont d'Urville, but explicitly treat this comparison as a plausibility check rather than strict validation because of differences in seasonality and spatial scale between summer point measurements and winter 27 km ASCAT-scale retrievals (lines 477--493).

19. **Computational feasibility of scaling the methodology not discussed.**

The paper mentions using the DelftBlue supercomputer with 16 cores allocated per grid point (line 508) but doesn't address whether this approach is computationally feasible for operational or large-scale applications. This is relevant because if the framework is intended as a "complementary monitoring metric" (line 557), it will be useful to know whether it's practically deployable or remains primarily a research tool.

In the revised manuscript, we have added a practical consideration in Sect. 5.1 to distinguish the present implementation from an operational product. We now state that "the current inversion was designed as a research-scale framework to test the physical consistency and interpretability of ASCAT-constrained grain-size retrievals across Antarctic ice shelves, rather than as an operationally deployable monitoring system" (lines 633--637).

We also clarify that, although the workflow is feasible on high-performance computing infrastructure, "routine large-scale or near-real-time application would require additional methodological development" (lines 637--638). To make the scalability options more concrete, we now identify possible routes including "reduced-order formulations, pre-computed lookup tables, surrogate or emulator-based forward models, or selective regional and temporal updating focused on regimes where ASCAT provides the strongest retrieval leverage" (lines 638--640).

Finally, we state that "the present results demonstrate research-scale feasibility, while operational deployment would require further optimization of the inversion framework" (lines 640--641).

20. Typographical Issues.

- In Section 4.2, at line 305, the word "in" is duplicated into "in in".
- Inconsistent use of "Antarctic-ice-shelf-wide" vs "AIS-wide".
- In some instances, "FAC × GS" is written as "FAC×GS" (without spaces around '×').
- Invalid URL of Alley et al., 2018 in the References. It would be worth double checking and fix similar errors (if any).

We have corrected the duplicated word in Sect. 4.2, standardized the terminology throughout the manuscript to use "Antarctic-ice-shelf-wide" consistently, and checked the formatting of the FAC×GS notation in the variance-partitioning text and figure caption (around lines 381--384 and 392--394). We have also checked the reference list and corrected the Alley et al. (2018) DOI/URL entry.