

Review 1 of Antwerpen et al 2026, TC Discussion

Note to editor and referees: we decided to revise our previous replies, please use the following responses.

The manuscript presents a Machine Learning approach to calculate the spatial and temporal evolution of albedo in the bare ice zone on the SW part of the Greenland Ice Sheet. PIXAL is trained on MODIS albedo and compared to the method for calculating albedo in MAR, after calibrating the MAR method to the same MODIS data. While I found the interpretation of the SHAP a little difficult to follow and I was missing some discussion on the incident angle/slope dependence on MODIS data, the manuscript is clear and well presented and I can recommend publishing after some minor revisions. I hope the authors will take this work further in the future and implement PIXAL or something similar in MAR. We would like to thank Referee #1 for their thorough comments and helpful suggestions. We appreciate the effort it took to review our manuscript and are confident the manuscript will be improved after we address these comments.

I understand and agree with your concern regarding the clarity of the SHAP method. I addressed this in your comment about line 338 and changed that sentence to better clarify how SHAP works in our workflow. Further, I agree with your concern about the incident angle/slope dependence on MODIS data. I address this point in your comment about line 380-386 and added text to the manuscript to discuss this.

Here are my line by line comments:

48: maybe add to the sentence something like: [...] resulting in a net mass loss *and exposure of the bare ice surface every year*.

Thank you for the suggestion, I added “[...] and the potential exposure of the bare ice surface.”

67: I think Cryoconite is a collective term for dust and algae, and not a type of particle in itself?

You're right, I removed cryoconite from the list.

88: Consider adding a *thus*: [...] which can lead to underestimates of surface melting and *thus* sea level rise [...]

Done.

225-227: Maybe I missed it somehow, but for the albedo prediction in the XGBoost you use all the MAR output listed in line 140-146, except Albedo and cloud cover. But what about: surface melt and shortwave upward radiation. These must also be dependent on the MAR

albedo. But I am unsure how this affects the results, could you maybe add a sentence about this?

Shortwave upward radiation is included in the list in 140-146. I added the following sentences to clarify why surface melt is not included in the list: “While surface melt is available as an output variable of MAR, we did not include this variable in this list because it strongly correlates with runoff. We chose to include runoff over surface melt because we want to capture the potential redistribution processes of LACs through runoff and the consequent impacts on ice albedo.”

Figure 2 c): Why do you think there is this spatial difference in the performance? I wonder if the difference you see is due to slope? There must be an issue with MODIS seeing albedo differently depending on slope. I think it would be easy to compare that here – although maybe out of the scope for this study.

That’s an interesting point. The largest XGBoost-MODIS discrepancies are found in areas with a low slope, so this would likely not be the cause. The sun and satellite viewing angles could be a culprit at these high latitudes though. Testing this is unfortunately out of scope for this study, but I added the following sentences to section 5.2:

“Moreover, some spatial variability is present in the performance of XGBoost compared to MODIS (Fig. 2c). This discrepancy may be partly due to spatial variability in the number of valid MODIS data points, which can lead to a bias in the mean observed ice albedo in areas with fewer valid data points (Fig. 2a). Lastly, MODIS exhibits a positive albedo bias above 70° N due to sun and satellite viewing angles at high latitudes (Alexander et al., 2014). The largest differences between XGBoost and MODIS albedo are found near this latitude, potentially explaining the XGBoost-MODIS discrepancy as a positive bias from MODIS rather than an underestimate from XGBoost.”

338: I am not sure I understand this: “In other words, the SHAP value shows how much the ice albedo prediction increases or decreases due to each individual feature relative to the mean ice albedo”. When you say ice albedo prediction increases do you then mean that it is the performance of the albedo prediction that gets better?

With an increase/decrease in ice albedo prediction I mean that the absolute value of the predicted albedo is higher/lower. In this sentence I do not relate it to the performance, just the actual albedo value. I changed the sentence to: “In other words, a feature’s SHAP value represents the extent to which that variable pushes the prediction above or below that mean.”

340-342: I don’t understand why you determine the surface height and slope to be the primary and then the climatic to be the secondary. To me it looks like temperature, shortwave incoming and wind is a better predictor than slope and surface height. Temperature in particular looks unambiguous. But maybe it relates to the fact that I have not really understood what you are comparing in the SHAP? I think a few sentences should be added to clarify this.

I agree that this section is somewhat confusing. According to the results in Fig. 5, surface height, temperature and slope are the most impactful drivers, since the features are ordered by the mean absolute SHAP values (how much each variable changes the albedo prediction). The other variables also impact the albedo prediction, but less significantly. We describe this in more detail in Section 5.1. I've adjusted the wording to:

“The topographic features of surface height and slope, as well as temperature, are the key drivers of ice albedo. Radiative and atmospheric features are important additional drivers. The significance and impact of each variable listed in Fig. 5 is described in more detail in Section 5.1.”

I have also adjusted the abstract and the conclusion to account for this new wording.

367-374: This is all correct, but maybe add some discussion here on what does that mean for results from PIXAL if used for future runs?

Thank you for the great suggestion. We address this issue in the limitations section in L462-469. I added the following sentence to point to this discussion:

“The impact of these effects on the performance of PIXAL is further discussed in Section 5.2.”

380-386: As mentioned above, I think that MODIS albedo is likely affected by the slope, giving a somehow skewed picture of actual albedo e.g. Wang and Zender (2010) in “MODIS snow albedo bias at high solar zenith angles relative to theory and to *in situ* observations in Greenland” (<https://doi.org/10.1016/j.rse.2009.10.014>). I am missing bit of discussion on this.

I added this sentence: “Note that the MODIS-derived albedo may be affected by high solar zenith angles (SZA) $>55^\circ$ (Wang and Zender, 2010; Alexander et al., 2014). This may introduce a negative albedo bias, especially in areas with a high ice surface slope.”

426-230: I was looking forward to hearing your thoughts about the albedo bias as mentioned above. To me this seems like an obvious next step and I am looking forward to see your results! In this study it would be interesting to see how often MAR albedo based on PIXAL would have to be recalibrated.

This is a great suggestion and we are excited about this next step too! Unfortunately, this approach is outside of the scope of this manuscript. We cannot confidently comment on the impact or number of iterations required to achieve a desired accuracy without testing the approach. I added the following sentence to this section to address the question of number of iterations required for this approach:

“While this iterate-update approach is likely required before permanent embedding of PIXAL in MAR or other models can be considered, testing this is unfortunately outside of the scope of this study. We therefore cannot confidently comment on the impact of this approach, nor on the number of iterations required to reach a desired accuracy. We look forward to applying this method in a future study.”

440-450: Again I am missing some discussion on MAR and incident angle / slope
This is addressed in the comments above about lines 380-386 and figure 2c.

Review 2 of Antwerpen et al 2026, TC Discussion

Note to editor and referees: we decided to revise our previous replies, please use the following responses.

1 Summary

In this study the authors introduce an ML based model, PIXAL. It implements an XGBoost-based architecture that maps 14 environmental features from the MAR climate model to Greenland ice albedo from MODIS observations. The architecture significantly outperforms traditional parameterizations ($R^2 = 0.563$ vs. 0.062). The strength of the study lies in the application of SHAP, an interpretability tool, to PIXAL. The SHAP analysis yields surface height and temperature as the primary drivers of the change of ice albedo changes. While providing a robust data-driven alternative to MAR formulations, the model operates as a feature-based regressor. Although suitable for hindcasts and forecasts, PIXAL can struggle with out-of-distribution samples in a changing climate. I think this is a valuable study and hope that my comments and technical critiques are useful to the authors in further strengthening the manuscript.

We would like to thank Referee #2 for their thorough comments and helpful suggestions. We appreciate the effort it took to review our manuscript and are confident the manuscript will be improved after we address their comments.

2 Reviewer Comments

2.1 Major

The physics-informed aspect of PIXAL

As a general comment, I would encourage the authors to clarify exactly the "physics informed" aspect of the ML model developed in the study. Physics-informed models typically use knowledge from physics to introduce constraints or loss functions to nudge neural networks to produce outputs that are consistent with the physics of the problem at hand. Could the authors perhaps clarify their claim and/or sources as to what exactly makes their model physics-informed?

Thank you for your comment. I agree with your statement as we did not include any physical laws into the model or loss function besides setting a constraint for the predicted ice albedo to be in the range 0-1. I changed the title, abstract, and introduction to reflect this with "Physics-Inspired" instead of "Physics-Informed".

Use of a tree-based model architecture

Generalizing capabilities: PIXAL is based on XGBoost, which is a treebased ML model. It is widely understood that tree-based methods can have issues with generalizing outside of their training dataset. Do the authors think using a neural network based model architecture be more

useful for generalization? Is there a particular reason/advantage for choosing a tree-based model in this case?

Thank you for raising this point. Our initial tests for ML model selection included neural networks (NN) as well as tree-based approaches. We found a lower performance for NNs and decided to continue with the best-performing models and omit the NNs from further tests and analyses. In L462-469 we describe that our dataset has a $\sim 6\sigma$ spread in values, ensuring some robustness against future out-of-distribution values. Please also see our response to the next comment, where we discuss the model architecture choices (including NNs) and add additional text to the limitations in the discussion section of the manuscript.

Spatial structure: The ice albedo is a spatial feature. XGBoost does not understand the relationship between a pixel and its neighbors and from my understanding, in this study each grid is treated as an independent data point. Can the use of a convolution based method help capture the gradients in the dark zone more accurately?

Thank you for this suggestion. Similar to the comment above, our initial tests for ML model selection included a convolutional neural network (CNN) and a CNN+long short-term memory (CNN-LSTM) model. These spatially-aware architectures showed better performance than a standard NN, though still not as good as XGBoost. We acknowledge that SHAP can be applied to tree-based models, CNNs and NNs. However, implementations of SHAP for CNNs and NNs rely on approximations of SHAP values, while SHAP for tree-based models computes the exact SHAP values. Given the importance of the explainability of the model in this manuscript, and the superior performance of XGBoost over NNs and CNNs, we opted for XGBoost as the basis for PIXAL.

We agree that the generalization limitations (compared to NNs) and spatial limitations (compared to CNNs) of our tree-based approach may affect the results and that this requires acknowledgement in the manuscript. We therefore added the following paragraph to the end of the limitations section in the discussion:

“Beyond these distributional considerations, we make two architectural trade-offs by opting for XGBoost as the basis of PIXAL. First, neural networks can in principle generalize outside of the training distribution, whereas tree-based models will assign out-of-distribution inputs the same prediction that is associated with the nearest boundary of the training data, effectively limiting extrapolation. However, as described above, the $\sim 6\sigma$ spread of our training dataset partially mitigates this limitation. Moreover, preliminary experiments showed a lower performance for NNs than tree-based approaches in modeling ice albedo. Second, tree-based models treat each grid cell as an independent data point and do not encode spatial relationships between neighboring pixels. This may limit their ability to capture spatial gradients such as those in the dark ice zone. Convolutional neural networks (CNNs) are theoretically better suited to leverage the information surrounding each pixel due to the spatially-aware nature of their convolutional kernels. However, our preliminary tests showed a superior performance of tree-based models over a CNN and a CNN-long short-term memory (CNN-LSTM) model in modeling ice albedo. Lastly, while SHAP can be applied to NNs and CNNs, these implementations rely on

approximations of the SHAP values, whereas the SHAP implementation for tree-based models computes the exact SHAP values. Given the importance of model explainability to identify drivers of ice albedo variability, we opted for XGBoost as the basis for PIXAL.”

Scalability: The authors recognize that the unexplained variance in the ice albedo prediction can also be attributed to a lack of information about several physical processes. For XGBoost, adding more features (such as information about algae and dust, even if contrary to reality, they were available at least as proxies) would significantly increase the computational cost. Have the authors considered other methods of representing this information without coming up against the difficulty of computational scalability of tree-based models?

Regarding computational scalability, training runtime does increase with additional input variables. However, in principle, training only needs to be performed once. Future iterations of PIXAL will likely prune non-informative variables, further reducing computational costs.

Regarding your point on adding additional LAC-related datasets or LAC proxies, we considered adding additional datasets to the input data collection. However, we found that no datasets are available on the same spatiotemporal range and resolution as the rest of our data (MAR and MODIS). We included the 2-year annual algal bloom dataset from Wang et al, (2020) in initial tests and found that the uneven input data availability resulted in biases and artefacts in our model output. We are unaware of other, more comprehensive datasets or proxies that are able to add sufficient information to improve the performance of the XGBoost model. To address this point, I added the following text to Section 5.2:

“Some datasets are available that could represent these LAC-driven darkening processes, e.g. annual algal bloom data (Wang et al., 2020). However, this data is currently not available with the same spatiotemporal range and resolution as the MAR and MODIS datasets used in this study. Initial experiments showed that this 2-year annual algal bloom dataset created uneven input data availability resulting in biases and unwanted artefacts in the model output. We therefore omitted this data from our input data collection. We are unaware of other, more comprehensive datasets that are able to add sufficient information to improve the performance of the XGBoost model. Beyond adding direct observational inputs, alternative approaches to representing LAC-driven darkening processes could include proxy variables that co-vary with LAC concentration, or explicitly accounting for the contribution of unknown processes through uncertainty quantification. However, identifying reliable proxies at the required spatiotemporal resolution is an open challenge. We consider both directions valuable for future development of PIXAL.”

2.2 Minor

Line 37: "hindrance to" instead of "hindrance for"

Done.

Line 88: "underestimation" instead of "underestimates" fits better

Done.

Line 213: more appropriate in ML terminology would be "generalize to data"

Done.

Line 229: "predictand" instead of "predictant"

Done.

Line 236-240: Could these be put into a table? If the authors wish to leave the hyperparameter information in a paragraph, some of the text here seems to be in a different font.

We prefer the text to be in a paragraph. I changed the font to match the rest of the manuscript.

Line 272: "insufficient variability" instead of "too-low variability"

Done.