

Peer Review File

Manuscript Title: Permafrost slows Arctic riverbank erosion

Reviewer Comments & Author Rebuttals

Reviewer Reports on the Initial Version:

Referee #1 (Remarks to the Author):

The paper examines the effect of permafrost on the pace of river migration. This has been a controversial issue in the literature as the authors acknowledge and discuss. Some studies have shown that permafrost acts to slow river migration, and some to accelerate it, with diverging consequences in predicting the fate of river migration in a warming arctic. The findings of this study predict that a warming climate will result in 30% to 100% increase in the migration rate of arctic rivers.

The paper is well written and all concepts and methods well explained, if not incredibly detailed at some places, such as for example when explaining the Fourier method for riverbank change detection in the main text (lines 74-81) and in the Methods, as well as some other passages. Although this might be distracting at times, I would not attempt to change the style of the authors. The figures are great and help to get the messages across.

A unique aspect of this study is that it approaches the problem at scales smaller than previous studies, i.e., sub-pixel scales at 3 m resolution using PlanetLab data, and at sub-seasonal and interannual time scales instead of decadal variations of Temperature and Discharge. This combination allows to differentiate the underlying mechanisms of erosion: entrainment-limited (where seasonal patterns of erosion should track discharge) vs. thaw-limited (where seasonal patterns of erosion should track water temperature) or a combination of the two. These hypotheses are tested by comparison of model predictions with observations.

Permafrost data is incredibly unreliable. Here they use the Pastick dataset (give the probability of permafrost within 1 m) and, as I understand, the Pastick permafrost map is trained on observations from 1999 to 2015 and their observations are mostly from 2018. Significant enough warming may have made the historically trained map a bit inaccurate, but they seem to control for this by identifying at what permafrost probability threshold the accuracy with respect to their 2018 field observations is maximized. So really good attention to detail to ensure the robustness of their analysis and results.

Commendable too that they not only acknowledge the issue of co-registration (i.e. high uncertainty in the location of the river) but also explicitly examine and controll for it by performing the analysis in reverse (section beginning on line 421 to line 500).

So overall, this is great work by the PhD student (first author) and the team.

Referee #2 (Remarks to the Author):

A. Summary of the key results

I have been pleased to review Permafrost slows Arctic riverbank erosion by Geyman et al. The authors have developed and applied a novel method for detecting sub-meter changes in river erosion rates for a study reach in Alaska. They have combined field, station, and modeled records of discharge, water temperature, and permafrost with these novel methods. Combining these observations with a clever study design, they find that erosion is above expected rates in river reaches with less permafrost.

This paper is a first step toward resolving a debate about the fate of Arctic rivers as they warm and is a potentially important contribution. I thoroughly enjoyed reading it and believe it is a model for research in the hard-to-study arctic. I have a few questions about the results and would be interested in seeing the authors push their approach past the relatively narrow hypothesis-test framework that is successful in this paper but leaves the reader with a few lingering questions.

--

B. Originality and significance: if not novel, please include reference/D. Appropriate use of statistics and treatment of uncertainties

To me, the clear strength of the paper is in methodological development and believe that it is an exciting approach that warrants testing around the world. This sub-pixel analysis has potentially high value even beyond the river meandering literature. The approach is well documented and many considerations that a reviewer might raise have already been addressed. Nice!

C. Data & methodology: validity of approach, quality of data, quality of presentation

The clever study design is also a great illustration of how to distill complex processes to a testable form. The authors also make a strong case for this type of local/regional study as at least a complement to the pan-arctic studies that have dominated this domain.

D. Appropriate use of statistics and treatment of uncertainties

Statistics and uncertainties were treated well. I think the authors can provide further comment on the potential statistical artifact/overfitting that they mention in the caption of Extended Data Table 1. "Note that the improvement in the goodness-of-fit (quantified through R^2) from Extended Data Fig. 4i to Fig. 4k could simply represent the increase in the number of tunable parameters from 1 (thaw-only) to 2 (entrainment-only) to 3 (combined thaw and entrainment)."

E. Conclusions: robustness, validity, reliability

F. Suggested improvements: experiments, data for possible revision

In tackling their hypotheses about erosion, the authors do a good job showing that permafrost reaches are eroding slower than non-permafrost reaches. However, as they point out, this introduces a chicken-or-egg question: are permafrost reaches eroding slower or are the fast-eroding reaches just reworking young-ish alluvial material?

To try to resolve this question, the authors turn to their state-of-the-art observations of sub-seasonal erosion and models that consider a) thaw, b) entrainment (river discharge), and c) thaw+entrainment. The authors do a good job of persuading the reader that neither the entrainment-only or the thaw-only model for erosion is a sufficient explanation for observed erosion in this reach of the Koyukuk River.

A model based on thaw+entrainment improves the results quite a bit, but I'm not necessarily convinced that this combined model for how things are working is a unique solution to their set of observations. Rather, it seems to be a model that a) makes a lot of sense, given the importance of discharge for sediment transport, b) has been proposed, and c) is convenient for hypothesis testing given the available data. It's certainly a logical place to start, but it would be more interesting to me if the authors had their observations guide a new hypothesis.

For instance, what's really amazing about the record is the drop-off in erosion in July, which is pretty stark and not necessarily captured by any of the models, including entrainment+thaw. For me, this immediately raises two questions not explicitly tested in this paper.

1. Is the seasonal pattern of erosion due to a form of hysteresis, i.e., sediment source exhaustion as the banks are stripped of the easy-to-transport chunks/slurries of eroding/eroded permafrost? It seems highly plausible to me that permafrost melting/slurrying at the end of the prior year summer could have built up a supply of sediment that was not transported (due to low late-summer flows). It's easy pickings for early season flows (maybe after some preliminary thaw), is subsequently all transported away, and then is regenerated by bank collapse as summer continues. I'm not sure if this is testable or if it was observed in the field, but it would be worth testing or, at the very least, commenting on.

2. Is vegetation an important story here: does erosion slow/halt as soon as the vegetation grows in for the season?

Particularly given the ongoing debate about vegetation and arctic riverbank erosion, the uncertain (but potentially rosy) fate of vegetation in a warming climate, and especially given that vegetation is the underpinning of their image analysis with NDVI, I would like to see the authors test vegetation as a control. For instance, the thaw-only model really does do a pretty good job on its own, at least in Figure 3a. I'm not sure why its performance is so strong there but so weak in Figure 3e, but I do wonder if the authors added an estimate for cohesion due to vegetation colonization during the summer, if that couldn't also combine with the thaw model to explain a lot of the variance in the interannual record.

Ideally, this would be informed by an analysis of banktop or bankslope NDVI, which I think is easily extracted from your record. Does vegetation grow in just as erosion wanes? It may be that vegetation is irrelevant. That would be great to show! It could be that vegetation gives a similar result to discharge. That would also be interesting, even though it might make it harder to arrive at a single interpretation.

In general, I believe that this paper has just about everything that we want in top-notch science: novel methods that advance the field and are transferrable beyond their application, an interesting and relevant science question being tested, and incorporation of models that allow for forward-looking predictions. I would be excited to see it published and think that the analyses I've suggested above would go a long way toward informing future research, while resolving some of the essential outstanding questions from past research.

G. References: appropriate credit to previous work?

References are clear and appropriate, to my knowledge.

Other comments/questions

H. Clarity and context: lucidity of abstract/summary, appropriateness of abstract, introduction and conclusions

The paper is well written and clear. I just have a few things here and there.

I wonder if the title could be softened just slightly – I'm not saying change it to "Permafrost slows riverbank erosion along a 450 km reach of the Koyukuk River, Alaska". But, given that Nature often publishes global-scale analyses, adding some element of regionality or case study to this could be good. Although I think this paper presents strong evidence and is well defended, it's not looking across the whole arctic. Maybe something like: "Permafrost slows riverbank erosion along an Arctic river"?

10: Add something about the cohesion or resistance of the banks?

12: Should probably be: "an individual river".

16: I don't think you need the "Yet". Perhaps just: "Observing bank erosion..."

144: One point of clarification — I think there are two R2 being calculated: one for the annual timeseries (at ~10 timesteps for five years?) and one from the relationship between X: annual estimated erosion from each model and Y: annual estimated erosion from the river migration satellite observations (n = 5 for each relationship)? I think making this clearer in the methods would be helpful. In addition, I think it would be helpful to add some commentary about why the R2 for the seasonal thaw-only model is pretty high (0.79), while for the interannual variation is pretty low (0.44).

377: May as well say "(see below)" or "(see next section)", since you expand on this below.

473–474: I suppose, although the accretion vs. erosion monotony is probably much less reliable, even over 6 years. I am not particularly concerned by this decision to only allow one direction of change, but I'm not sure that the comparison to this capital-L Law is necessary.

Extended Data Fig. 6: Unless I'm missing them, I believe these plots are missing letters.

Extended Data Fig. 7: In (g), along with "Stacked & bracked cumulative displacements", I would say from (f) to make it clear that this is processing of the real (noisy) data.

Photographs from the field of these eroding banks would be very instructive.

Referee #2 (Remarks on code availability):

I read the code but don't have Matlab, so I did not run the code. I think that it seems like a usable resource, more or less, although it doesn't have the original image files included. I think those are private (Planet data), so I'm not sure if they can be made available. The The commenting of the code could be more extensive but I think is appropriate.

Referee #3 (Remarks to the Author):

I have enjoyed reading and reviewing the manuscript by German et al., a manuscript that addresses an important and impactful subject: how does permafrost affect the rate of river migration in meandering rivers. As it is pointed out in the introduction, this is a question that - so far, and somewhat surprisingly - has contradictory answers; and I think the idea of focusing on a single river and pushing the limits of resolution as far as possible makes it more likely that significant insight is gained. The authors bring a wide range of sophisticated methods to get as much information as possible from the available datasets. If these methods (especially the sub-pixel detection of bank migration, regardless of river stage and georeferencing errors) are indeed as powerful as the manuscript suggests, they have the potential to revolutionize the way we estimate river bank migration. The manuscript is well written and nicely illustrated, with clear attention to all the important details. It is obvious that a lot of careful work went into this study. It is also great to see that our work on the relationship between curvature and migration rate and the application of this approach to satellite imagery is being put to good use and taken further, in innovative ways.

Note that my expertise primarily lies in detecting migration rates from time-lapse remote sensing data and linking them to channel geometry; so most of my comments focus on this aspect of the manuscript.

While the range and novelty of some of the methods and the strength of the correlations in the results make it quite tempting to fully trust the conclusions, I think there are a number of important aspects that require more clarity and/or discussion. I first list some of the broader and more important comments and suggestions, as follows:

1. The authors have compressed a lot of work and information into a relatively short manuscript, placing a lot of the details in the 'Methods' section, and the result is that it is not always easy to follow what exactly has been done. For example, I cannot locate much information on how channel bank locations were identified. Is this an automated process? What is the spacing between points where migration rates are estimated? How is curvature computed? How difficult would it be to scale this approach up to multiple rivers? I guess the answers to some of these questions could be found in the code, but I do not have access to the repositories that the authors mention and I am unable to assess how easy it would be to reproduce the results.

2. Bank migration rates that are lower than what is expected from channel curvature are likely due to low erodibility of the banks (e.g., Sylvester et al., 2019, especially Figs. 3b and 12 in Supplementary material). In the case of the Koyukuk, this low erodibility might be due to the presence of permafrost in the cutbank, as suggested by Geyman et al. However, it can be also due to a more erosion-resistant composition of the substrate; and this possibility is not discussed in the manuscript. Digital elevation models of the study area can be used to see where the river comes in contact with higher elevation, older, and likely more erosion resistant material - see figure 1 in the attachment (elevation data from Arctic DEM). I have highlighted places where low migration rates are probably due to reduced erodibility of the substrate, a reduced erodibility that is unlikely to result from the presence of permafrost alone. Some of the areas that are thought to be strongly affected by permafrost (e.g., dark red area in Fig. 2b) might also correspond to older and therefore

more erosion-resistant substrates. There is a danger of conflating the effect of permafrost with lithology-based variations in bank erodibility.

3. It looks to me like part of the bank migration observed in the two Sentinel images from 2016 and 2022 is due to stage variability, not actual deposition. The results shown in Fig. 1a and Extended Data Fig. 6 suggest a significant bias toward higher migration rates on the inner banks compared to the outer banks (see figure 2 in the attachment). If this reflects real bank movement, it also means that the river has systematically narrowed between 2016, because point bar accretion outpaced outer bank erosion. I think a more likely explanation is that the 2016 image is a high stage image in which the water tends to cover part of the vegetation (e.g., top of Extended Data Fig. 6); and the 2022 image is a low-stage image in which all the vegetation is exposed. Note that this kind of 'migration' matches well the pattern expected for actual migration when it comes to its position relative to curvature.

4. The ability to extract bank migration data at a temporal frequency so that an annual erosion time series can be constructed (Fig. 3d), despite the difficulties that arise from variable stage and coregistration errors is truly impressive. However, it would be helpful to see more data that supports this result and shows how it was obtained; e.g. a sequence of image chips from one summer, with the associated raw migration vectors, and how the final erosion time series looks like for that location. Much of the real data examples shown in the manuscript come from the Sentinel data and the arguably more interesting and innovative work that focuses on the higher resolution Planet Labs imagery is not illustrated.

5. Figure 3 caption - "Note that we only study the seasonal pattern of bank displacement on the eroding side of the river." - This seems like an important part of the methods but there is no other information about how and why this is done. Are erosional banks manually segmented for each time step? Or do you use an automated approach to identify these segments? Looks like the restriction to the eroding side only applies to the study based on the Planet Labs data — why not use the same approach for both the Sentinel and Planet Labs datasets? This approach would take care of the problems that come from 'enhanced' migration rates on inner banks in the Sentinel data.

6. Extracting robust erosion rates from noisy timeseries: it seems to be that this is one of the most interesting and novel contributions in the paper, certainly when it comes to the methods used, yet I don't fully follow what exactly is being done and why. How do 'stacking' and 'bracketing' affect the result? For example, it looks to me like the line in Extended Data Fig. 7d is the same as the first row of values is panel (e) of the same figure — so how do stacking and bracketing influence this? If the displacement is averaged along a row, you end up with a single average displacement for the entire time series - I doubt that this is what is happening. The cumulative displacement in the synthetic erosion data is already monotonic; I think that going through this exercise with actual data (like the data shown in Extended Data Fig. 7f) would be more insightful than relying on synthetic data that has no obvious 'backstepping'. Along the same lines, it is unclear to me why the MCMC approach is needed to get a monotonic displacement series; isn't it easier to just clip the 'backsteps', like it is often done when generating stratigraphy from topographic time series?

7. Removal of seasonal bias: this also seems like a great idea and useful trick but it is hard to follow

(for me at least) how exactly it is done. Is there a single mean value of E_{slow} computed for each time step? If so, I don't think it accounts for all the stage-related change along a gently dipping point bar. Also, is this method used to remove seasonal bias or coregistration errors, or both? The section title refers to seasonal biases but then the text says that these displacements most likely reflect coregistration noise. In addition, if 68% of the raw data consists of this noise, it would be good to see how some of the data looks like before and after this processing step.

8. Using NDVI means that the algorithm detects changes in the vegetation line, not erosion and deposition. This is likely to be fine on steep cutbanks, but things are more complicated on gently dipping depositional banks where, for example, young vegetation might be covered by the high-stage river, so that stage variability might still be an issue. This is discussed to some degree in the Methods section, but it should be mentioned in the 'image pre-processing' section as well (and maybe in the main text).

A few additional comments:

Line 15 - "not decades as in previous work" - some of the previous work used years (sometimes less than one year), not decades, especially in rapidly migrating rivers

Lines 18-19 - "We develop computational tools to resolve the spatial and temporal patterns of riverbank erosion at length scales $\sim 100\times$ finer than previous analyses" - this is a bit too broad brush, as part of the improvement in resolution comes from working with 3m-resolution data instead of 30m-resolution data

Lines 61-62 - "Existing methods (requiring whole-pixel change in 30 m Landsat satellite imagery) would require time intervals of at least 10 years to detect erosion." - this is not entirely true; e.g., we were able to detect changes of less than 30 m with our workflow (Sylvester et al., 2019; 2021).

Lines 81-83 - "A critical component of our methodology is a workflow to make the displacement estimates robust to georeferencing errors in the satellite imagery, since image-to-image co-registration errors can exceed the signal of riverbank erosion by $\geq 10\times$." - this is indeed a critical component and probably deserves a bit more detail in the main text. The same applies to dealing with stage variability / seasonal biases.

Line 92 - "we can ask" instead of "we we can ask"

Fig. 2 - How was the permafrost probability computed for each bend? E.g., do you only consider the pixels that are in contact with the eroding bank? Or do you consider a broader band? Looking at Fig. 2b, this can have a large impact on the permafrost probability values.

Fig. 3d - This is a remarkable result but it would be great to see example images and the derived data that show, for example, that erosion rates in June are significantly larger than in May.

Line 468 - it would be helpful to briefly explain somewhere that 'MCMC' stands for 'Markov-Chain Monte Carlo' and how it relates to the method described here.

Extended Data Fig. 6 - Panel labels (a, b, c, d) are missing and caption for upper part of figure (which I guess is supposed to be panel a?) seems to be missing.

Extended Data Fig. 6 - Sentinel images from 2016 and 2022 in the lower left corner of the figure are identical - I think both come from the 2016 data.

Extended Data Fig. 6 - If I am reading the text correctly, the upper part of this figure illustrates how a small displacement between two images can be recovered using the sub-pixel method. I think it would be useful to show an actual example of change detection as well where the difference between the images is more obvious.

Extended Data Fig. 6 - discussion of using NDVI - seems like this should be in the text, not the figure caption. Also, it seems to assume that the location of the vegetation line is immune to stage variability and it tracks deposition precisely.

I hope you find these comments useful.

Referee #3 (Remarks on code availability):

Access to the code repositories was not provided, so I could not review the code. Having access to it would have been useful as a lot of the details of the methods are (somewhat understandably, as not all the details can go into the text) missing or difficult to follow in the text.

Author Rebuttals to Initial Comments:

Reviewer 1

The paper examines the effect of permafrost on the pace of river migration. This has been a controversial issue in the literature as the authors acknowledge and discuss. Some studies have shown that permafrost acts to slow river migration, and some to accelerate it, with diverging consequences in predicting the fate of river migration in a warming arctic. The findings of this study predict that a warming climate will result in 30% to 100% increase in the migration rate of arctic rivers.

The paper is well written and all concepts and methods well explained, if not incredibly detailed at some places, such as for example when explaining the Fourier method for riverbank change detection in the main text (lines 74-81) and in the Methods, as well as some other passages. Although this might be distracting at times, I would not attempt to change the style of the authors. The figures are great and help to get the messages across.

A unique aspect of this study is that it approaches the problem at scales smaller than previous studies, i.e., sub-pixel scales at 3 m resolution using PlanetLab data, and at sub-seasonal and interannual time scales instead of decadal variations of Temperature and Discharge. This combination allows to differentiate the underlying mechanisms of erosion: entrainment-limited (where seasonal patterns of erosion should track discharge) vs. thaw-limited (where seasonal patterns of erosion should track water temperature) or a combination of the two. These hypotheses are tested by comparison of model predictions with observations.

Permafrost data is incredibly unreliable. Here they use the Pastick dataset (give the probability of permafrost within 1 m) and, as I understand, the Pastick permafrost map is trained on observations from 1999 to 2015 and their observations are mostly from 2018. Significant enough warming may have made the historically trained map a bit inaccurate, but they seem to control for this by identifying at what permafrost probability threshold the accuracy with respect to their 2018 field observations is maximized. So really good attention to detail to ensure the robustness of their analysis and results.

Commendable too that they not only acknowledge the issue of co-registration (i.e. high uncertainty in the location of the river) but also explicitly examine and control for it by performing the analysis in reverse (section beginning on line 421 to line 500).

So overall, this is great work by the PhD student (first author) and the team.

Thank you for the kind feedback.

Reviewer 2

A. Summary of the key results

I have been pleased to review Permafrost slows Arctic riverbank erosion by Geyman et al. The authors have developed and applied a novel method for detecting sub-meter changes in river erosion rates for a study reach in Alaska. They have combined field, station, and modeled records of discharge, water temperature, and permafrost with these novel methods. Combining these observations with a clever study design, they find that erosion is above expected rates in river reaches with less permafrost.

This paper is a first step toward resolving a debate about the fate of Arctic rivers as they warm and is a potentially important contribution. I thoroughly enjoyed reading it and believe it is a model for research in the hard-to-study arctic. I have a few questions about the results and would be interested in seeing the authors push their approach past the relatively narrow hypothesis-test framework that is successful in this paper but leaves the reader with a few lingering questions.

Thank you for the feedback. We address your remaining questions below.

B. Originality and significance

To me, the clear strength of the paper is in methodological development and believe that it is an exciting approach that warrants testing around the world. This sub-pixel analysis has potentially high value even beyond the river meandering literature. The approach is well documented and many considerations that a reviewer might raise have already been addressed. Nice!

Thank you.

C. Data & methodology: validity of approach, quality of data, quality of presentation

The clever study design is also a great illustration of how to distill complex processes to a testable form. The authors also make a strong case for this type of local/regional study as at least a complement to the pan-arctic studies that have dominated this domain.

Thank you.

D. Appropriate use of statistics and treatment of uncertainties

Statistics and uncertainties were treated well. I think the authors can provide further comment on the potential statistical artifact/overfitting that they mention in the caption of Extended Data Table 1. “Note that the improvement in the goodness-of-fit (quantified through R^2) from Extended Data Fig. 4i to Fig. 4k could simply represent the increase in the number of tunable parameters from 1 (thaw-only) to 2 (entrainment-only) to 3 (combined thaw and entrainment).”

Thank you for the feedback. In response to your comment about Extended Data Fig. 4i-k, we now account for the change in the number of tunable parameters from 1 (thaw-only) to 2 (entrainment-only) to 3 (combined thaw and entrainment). Specifically, we now compute the *adjusted* R^2 value, defined as¹⁶:

$$R_{adj}^2 = 1 - \frac{(1 - R^2)(n - 1)}{n - p - 1} \quad (1)$$

where n is the number of data points and p is the number of independent variables (parameters). The p in the denominator of equation (1) acts to penalize models with more tunable parameters (i.e., give them a lower R_{adj}^2 value). The adjusted R^2 is equivalent to the reduced Chi-square value, which is a standard quantity used to compare goodness-of-fit obtained from models with different numbers of parameters¹⁶.

For the bar charts in Extended Data Fig. 4i-k, $n = 5$, since we have 5 years of data. Following equation (1), the R_{adj}^2 values for the thaw-only, entrainment-only, and combined thaw and entrainment models are 0.25, 0.14, and 0.40, respectively. In other words, the R_{adj}^2 metric supports the conclusion that the combined thaw and entrainment model best explains the data.

We modified Extended Data Fig. 4i-k to show the R_{adj}^2 metrics in addition to the raw R^2 values. We also added a section to the Methods (***Evaluating model goodness-of-fit***) (lines 334-343) to explain our model evaluation and use of the R_{adj}^2 metric¹⁶.

F. Suggested improvements: experiments, data for possible revision

In tackling their hypotheses about erosion, the authors do a good job showing that permafrost reaches are eroding slower than non-permafrost reaches. However, as they point out, this introduces a chicken-or-egg question: are permafrost reaches eroding slower or are the fast-eroding reaches just reworking young-ish alluvial material? To try to resolve this question, the authors turn to their state-of-the-art observations of sub-seasonal erosion and models that consider a) thaw, b) entrainment (river discharge), and c)

thaw+entrainment. The authors do a good job of persuading the reader that neither the entrainment-only or the thaw-only model for erosion is a sufficient explanation for observed erosion in this reach of the Koyukuk River. A model based on thaw+entrainment improves the results quite a bit, but I'm not necessarily convinced that this combined model for how things are working is a unique solution to their set of observations. Rather, it seems to be a model that a) makes a lot of sense, given the importance of discharge for sediment transport, b) has been proposed, and c) is convenient for hypothesis testing given the available data. It's certainly a logical place to start, but it would be more interesting to me if the authors had their observations guide a new hypothesis. For instance, what's really amazing about the record is the drop-off in erosion in July, which is pretty stark and not necessarily captured by any of the models, including entrainment+thaw. For me, this immediately raises two questions not explicitly tested in this paper:

1. Is the seasonal pattern of erosion due to a form of hysteresis, i.e., sediment source exhaustion as the banks are stripped of the easy-to-transport chunks/slurries of eroding/eroded permafrost? It seems highly plausible to me that permafrost melting/slurrying at the end of the prior year summer could have built up a supply of sediment that was not transported (due to low late-summer flows). It's easy pickings for early season flows (maybe after some preliminary thaw), is subsequently all transported away, and then is regenerated by bank collapse as summer continues. I'm not sure if this is testable or if it was observed in the field, but it would be worth testing or, at the very least, commenting on.
2. Is vegetation an important story here: does erosion slow/halt as soon as the vegetation grows in for the season?

Particularly given the ongoing debate about vegetation and arctic riverbank erosion, the uncertain (but potentially rosy) fate of vegetation in a warming climate, and especially given that vegetation is the underpinning of their image analysis with NDVI, I would like to see the authors test vegetation as a control. For instance, the thaw-only model really does do a pretty good job on its own, at least in Figure 3a. I'm not sure why its performance is so strong there but so weak in Figure 3e, but I do wonder if the authors added an estimate for cohesion due to vegetation colonization during the summer, if that couldn't also combine with the thaw model to explain a lot of the variance in the interannual record. Ideally, this would be informed by an analysis of banktop or bankslope NDVI, which I think is easily extracted from your record. Does vegetation grow in just as erosion wanes? It may be that vegetation is irrelevant. That would be great to show! It could be that vegetation gives a similar result to discharge. That would also be interesting, even though it might make it harder to arrive at a single interpretation.

You raise two great questions that we now more clearly address in the main text and the *Supplementary Information*. First, is there an instability/hysteresis associated with the sediment supply or availability that makes the channel switch between states of high and low erosion? Alternatively (or in addition), does vegetation colonization in mid-summer cause a slow-down in erosion rates? We address each question below, and also in the revised manuscript.

1. The role of sediment supply and seasonal thaw layers.

See this recent publication that uses modeling to explore the role of a seasonally-varying thaw layer¹⁷:

Douglas, Madison M., and Michael P. Lamb. "A model for thaw and erosion of permafrost riverbanks." *Journal of Geophysical Research: Earth Surface* 129.4 (2024): e2023JF007452.

Late-season thaw does not necessarily lead to 'easy pickings' for early season flows because those thawed sediments can re-freeze during the winter^{18,17}. Thus, erosion and transport of that material during the next season requires the dual condition of thaw and entrainment^{17,19}, consistent with our model presented in Fig. 3 of the main text. It is true that the freezing front may not penetrate the entire e.g. ~1-meter

pile of slumped bank material over the course of a single winter¹⁷. Also, re-freezing likely does not occur below the water line (although the water level is very low in winter—see Extended Data Figs. 2-4). Even so, the numerical modeling¹⁷ and seasonal permafrost probe observations¹⁸ so far seem to support our assumption that, at least during the period when high erosion rates occur, the state of the bank is in local equilibrium with the erosion mode (entrainment- or thaw-limited)¹⁷. See our new language in **lines 290-294** of the revised manuscript for an explanation of this concept.

2. The role of vegetation.

We added a new section to the *Supplementary Information* document (section 3) to explain why seasonal changes in vegetation are unlikely to be driving the observed patterns in riverbank erosion rates. We summarize this evidence below.

We have repeat drone imagery of a stretch of Koyukuk River near Huslia (thanks to collaboration with a high-school science teacher in Huslia, Steve Swatling Jr.: <https://swatlingsjournal.com/>) from a full spring-summer-fall season (May, June, July, August, September, and October of 2023) (see *Supplementary Information*, Figure S4). From these drone surveys, we can observe the changes in vegetation that occur over the course of the summer. The eroding cutbanks typically represent old terrain ($\geq 1,000$ years^{20,21}) with a well-developed forest canopy that is dominated by black spruce in places where the floodplain is permafrost and by balsam poplar, birch, alder, and white spruce trees in places where the floodplain is unfrozen²². From spring to summer to fall, there is little to no change in biomass density (e.g., from shrubification¹³), because the forest ecosystem in this old floodplain terrain does not include a significant component from seasonal grasses or shrubs (*Supplementary Information*, Figure S4). It is true that there is vegetation colonization of the young point bar deposits over the course of the summer. However, this vegetation growth occurs only once the water recedes and the bar is sub-aerially exposed. At this point in the seasonal hydrograph (Extended Data Figs. 1-3), the river flow is not interacting with the vegetation²³ because the stage is too low. See section 3 of the *Supplementary Information* document (**lines 83-105**) for more details.

In general, I believe that this paper has just about everything that we want in top-notch science: novel methods that advance the field and are transferrable beyond their application, an interesting and relevant science question being tested, and incorporation of models that allow for forward-looking predictions. I would be excited to see it published and think that the analyses I've suggested above would go a long way toward informing future research, while resolving some of the essential outstanding questions from past research.

Thank you!

G. References: appropriate credit to previous work?

References are clear and appropriate, to my knowledge.

H. Clarity and context: lucidity of abstract/summary, appropriateness of abstract, introduction and conclusions

The paper is well written and clear. I just have a few things here and there.

I wonder if the title could be softened just slightly – I'm not saying change it to 'Permafrost slows riverbank erosion along a 450 km reach of the Koyukuk River, Alaska'. But, given that Nature often publishes global-scale analyses, adding some element of regionality or case study to this could be good. Although I think this paper presents strong evidence and is well defended, it's not looking across the whole arctic. Maybe something like: "Permafrost slows riverbank erosion along an Arctic river"?

You raise a good point that our data come from a single Arctic river rather than a global analysis. So, if the editor concurs with your suggestion, we are amenable to changing the title to ‘Permafrost slows riverbank erosion along an Arctic river.’ One argument in favor of keeping the current, more general title, however, is that our observations are used to calibrate and validate a numerical model that is designed to apply to any Arctic river. In our view, an important contribution of this study is our analysis of the numerical model to explore how changing the water temperature, discharge seasonality, and permafrost abundance influences the riverbank erosion rates (see the section **Modeling future erosion behavior** in the main text (**lines 148-157**) and Extended Data Fig. 5). This analysis reveals general insight into the role that permafrost plays in slowing Arctic riverbank erosion rates, and therefore gives us a tool to predict how such rates may change in the coming century²⁴. We think the title ‘Permafrost slows Arctic riverbank erosion’ most directly and succinctly communicates this message.

Line-specific comments

Line 10: Add something about the cohesion or resistance of the banks?

We added the phrase ‘cohesion of bank material²⁵’ to the list of factors that can vary from river to river.

Line 12: Should probably be: “an individual river”.

We changed ‘individual rivers’ to ‘an individual river’.

Line 16: I don’t think you need the “Yet”. Perhaps just: “Observing bank erosion...”

We removed the ‘Yet.’

Line 144: One point of clarification – I think there are two R^2 being calculated: one for the annual timeseries (at ~10 timesteps for five years?) and one from the relationship between X: annual estimated erosion from each model and Y: annual estimated erosion from the river migration satellite observations (n = 5 for each relationship)? I think making this clearer in the methods would be helpful. In addition, I think it would be helpful to add some commentary about why the R^2 for the seasonal thaw-only model is pretty high (0.79), while for the interannual variation is pretty low (0.44).

That’s exactly right—we evaluate the different models (thaw-only, entrainment-only, and combined thaw and entrainment) on two different datasets. The first dataset is the sub-seasonal pattern of riverbank erosion (i.e., how much erosion occurs in June vs. July vs. August), and the observations for that dataset are shown in Fig. 3d. The second dataset is the interannual pattern of riverbank erosion (i.e., how much erosion occurred in 2017 vs. 2018 vs. 2019), and the observations for that dataset are shown in Fig. 3e. To help clarify this point to the reader, we modified Fig. 3 so that the R^2 values say $R^2_{seasonal}$ in panels (a)-(c) and $R^2_{interannual}$ in panel (e).

The thaw-only model yields a relatively high R^2 value for the sub-seasonal dataset ($R^2 = 0.79$) because it does a reasonable job representing the shape of the seasonal cycle of erosion. However, the thaw-only model does a poor job explaining the interannual signal ($R^2 = 0.44$) because there is very little variability in the river water temperature from year to year (see Extended Data Fig. 4c for example temperature data from Pilot Station on the Yukon River). So, the thaw-only model cannot explain, for example, why there was so much less erosion in 2018 than 2019. With our use of the adjusted R^2 value (see our response above in the section **D. Appropriate use of statistics and treatment of uncertainties**), our analysis of both the seasonal and the interannual datasets lead to the same conclusion: the combined thaw and entrainment model best explains the observational data.

Lines 473-474: I suppose, although the accretion vs. erosion monotony is probably much less reliable, even over 6 years. I am not particularly concerned by this decision to only allow one direction of change, but I’m not sure that the comparison to this capital-L Law is necessary.

Thanks for the feedback. We now have softened the language in the Methods section *Markov-chain Monte Carlo (MCMC) sampling...* to both (1) remove the description of this capital-L ‘Law’, and (2) more clearly convey that our reference to stratigraphic age models is just an example where similar MCMC methodology is widely used^{26,27,28,29}, rather than a direct analogy or justification for our assumption of monotonicity in the erosion/accretion timeseries. See **lines 498-510**.

Extended Data Fig. 6: Unless I’m missing them, I believe these plots are missing letters.

We added letters to the sub-panels of Extended Data Fig. 6.

Photographs from the field of these eroding banks would be very instructive.

We agree. Since we already used 10 Extended Data figures in the original manuscript, we pushed the original Extended Data Fig. 8 figure to a Supplementary Information document and replaced that figure with a panel of field photos illustrating the eroding riverbanks. The new figure is shown in Figure 1 in this document. We also added additional field photos of eroding banks to the Supplementary Information document.




Figure 1: Representative field photos of the Koyukuk River near Huslia, Alaska (65.689°N, 156.381°W).

Reviewer 3

I have enjoyed reading and reviewing the manuscript by German et al., a manuscript that addresses an important and impactful subject: how does permafrost affect the rate of river migration in meandering rivers. As it is pointed out in the introduction, this is a question that - so far, and somewhat surprisingly - has contradictory answers; and I think the idea of focusing on a single river and pushing the limits of resolution as far as possible makes it more likely that significant insight is gained. The authors bring a wide range of sophisticated methods to get as much information as possible from the available datasets. If these methods (especially the sub-pixel detection of bank migration, regardless of river stage and georeferencing errors) are indeed as powerful as the manuscript suggests, they have the potential to revolutionize the way we estimate river bank migration. The manuscript is well written and nicely illustrated, with clear attention to all the important details. It is obvious that a lot of careful work went into this study. It is also great to see that our work on the relationship between curvature and migration rate and the application of this approach to satellite imagery is being put to good use and taken further, in innovative ways.

Thank you for the kind feedback.

Note that my expertise primarily lies in detecting migration rates from time-lapse remote sensing data and linking them to channel geometry; so most of my comments focus on this aspect of the manuscript.

While the range and novelty of some of the methods and the strength of the correlations in the results make it quite tempting to fully trust the conclusions, I think there are a number of important aspects that require more clarity and/or discussion. I first list some of the broader and more important comments and suggestions, as follows:

1. The authors have compressed a lot of work and information into a relatively short manuscript, placing a lot of the details in the Methods' section, and the result is that it is not always easy to follow what exactly has been done. For example, I cannot locate much information on how channel bank locations were identified. Is this an automated process? What is the spacing between points where migration rates are estimated? How is curvature computed? How difficult would it be to scale this approach up to multiple rivers? I guess the answers to some of these questions could be found in the code, but I do not have access to the repositories that the authors mention and I am unable to assess how easy it would be to reproduce the results.

We have added more details to the **Methods** section to clarify your questions above. For example:

1. **How were channel locations identified? Is this an automated process?** See the new paragraph called ***Channel identification*** that we added to the Methods section (**lines 434-441**). The text is reproduced here: *We developed an automated pipeline for channel extraction and the identification of riverbanks from satellite imagery. The pipeline uses a combination of thresholding, morphological opening operations, and skeletonization. It is implemented in MATLAB (see Code Availability). Briefly, the NDVI is calculated following equation (18), and then the image is thresholded to the 10th percentile of NDVI. The result is a binary mask that separates unvegetated pixels (principally rivers and lakes) from the other pixels. Next, we perform a morphological opening operation (with a disk structuring element of size 5 pixels = 50 meters) to eliminate the small channels from the binary water mask. Next, we identify the river from the other water features (e.g., lakes) in the water mask using a connected component analysis. Finally, we identify the river centerline by skeletonizing the river mask.*
2. **How is curvature computed?** See the new paragraph called ***Calculating river curvature*** that we added to the Methods section (**lines 442-449**). The text is reproduced here: *We follow*

ref.³⁰ and compute river curvature (which is the inverse of the radius of curvature, R) as:

$$\frac{1}{R} = \frac{x'y'' - y'x''}{(x'^2 + y'^2)^{3/2}} \quad (2)$$

where x' and x'' represent the first- and second-order derivatives of x . In equation (2), we use the x and y coordinates of the channel centerline, extracted at 10-m intervals and smoothed with a Savitzky-Golay finite impulse response (FIR) smoothing filter^{30, 31} with polynomial order 3 and frame length of 33. The frame length is chosen based the chip width, $n = 32$, plus one to ensure that the frame length is an odd integer³¹.

3. **How difficult would it be to scale this approach up to multiple rivers?** Our new methodology absolutely is scalable to multiple rivers. Even a global analysis of all large rivers is tractable. We see efficiency and scalability as one of the principle advantages of our 2D Fourier transform approach. Whereas a direct 2D cross-correlation analysis would become prohibitively expensive for larger rivers or higher-resolution datasets, the Fourier approach remains tractable at these large scales. See **lines 75-81** of the main text for a discussion of this idea.

We have worked to make all of the datasets in our study publicly available. They now are published online at the NSF Arctic Data Center. There are two repositories linked to this manuscript:

Resolving the spatial and seasonal pattern of riverbank erosion on the Koyukuk River, Alaska, 2016-2022. Arctic Data Center. doi:10.18739/A2HM52M6Q.

<https://doi.org/10.18739/A2HM52M6Q>

Quantifying river migration rates in the Yukon River Watershed from optical satellite imagery, Alaska, 2016-2022. Arctic Data Center. doi:10.18739/A2WW7719J.

<https://doi.org/10.18739/A2WW7719J>

The PlanetLabs satellite imagery (<https://www.planet.com/>) is not ours to share, but we have provided lists of the exact images used in our analysis (see Tables 1-2 of the Supplementary Information). Academic researchers can get free access to PlanetLabs imagery by registering for the Planet Education and Research Program: <https://www.planet.com/markets/education-and-research/>.

2. Bank migration rates that are lower than what is expected from channel curvature are likely due to low erodibility of the banks (e.g., Sylvester et al., 2019, especially Figs. 3b and 12 in Supplementary material). In the case of the Koyukuk, this low erodibility might be due to the presence of permafrost in the cutbank, as suggested by Geyman et al. However, it can be also due to a more erosion-resistant composition of the substrate; and this possibility is not discussed in the manuscript. Digital elevation models of the study area can be used to see where the river comes in contact with higher elevation, older, and likely more erosion resistant material - see figure 1 in the attachment (elevation data from Arctic DEM). I have highlighted places where low migration rates are probably due to reduced erodibility of the substrate, a reduced erodibility that is unlikely to result from the presence of permafrost alone. Some of the areas that are thought to be strongly affected by permafrost (e.g., dark red area in Fig. 2b) might also correspond to older and therefore more erosion-resistant substrates. There is a danger of conflating the effect of permafrost with lithology-based variations in bank erodibility.

Note: the figure provided by Reviewer 3 is shown in Figure 2 below.

We absolutely agree. We added a section to the Supplementary Information document (section 4) to discuss the multiple reasons why a particular section of riverbank may be more or less erodible, which we group into three categories: (1) variations in cutbank elevation, (2) pore ice, and (3) other bank material properties (lithology, cohesion of sediment grains, vegetation, etc.). The red areas in Figure 2 could be




Figure 2: **Image provided by Reviewer 3.** Digital elevation model (ArcticDEM) overlain with Extended Data Fig. 6, with likely locations of low erodibility highlighted in red.

low erodibility for reasons (1) or (3), both of which do not necessarily have to do with permafrost. In section 4 of the Supplementary Information, we explain how we isolate mechanism (2) (the role of pore ice) from other sources of low erodibility. These methods include (i) removing all areas where the river abuts high terrain (those regions shown in red in Figure 2) from our analysis for Fig. 2 of the main text, and (ii) using the *seasonal* signature of erosion to show that erosion rates are dependent on the water temperature, an observation that should not be true if low erodibility is caused only by different lithology rather than the presence of pore ice. See **lines 107-172** of the Supplementary Information for more details.

3. It looks to me like part of the bank migration observed in the two Sentinel images from 2016 and 2022 is due to stage variability, not actual deposition. The results shown in Fig. 1a and Extended Data Fig. 6 suggest a significant bias toward higher migration rates on the inner banks compared to the outer banks (see figure 2 in the attachment). If this reflects real bank movement, it also means that the river has systematically narrowed between 2016, because point bar accretion outpaced outer bank erosion. I think a more likely explanation is that the 2016 image is a high stage image in which the water tends to cover part of the vegetation (e.g., top of Extended Data Fig. 6); and the 2022 image is a low-stage image in which all the vegetation is exposed. Note that this kind of ‘migration’ matches well the pattern expected for actual migration when it comes to its position relative to curvature.

Note: the figure referenced by Reviewer 3 in the comment above is shown in Figure 3.

We agree that river stage variability can lead to apparent bank migration rates that are larger on the inner

bank than on the outer bank. This difference is due to the slope asymmetry in the channel cross-sectional profile. See Figure S3 of the Supplementary Information for a figure that illustrates this concept. Thus, we agree that the most likely explanation for why the bank migration vectors in Figure 3 tend to be larger for the inner bank than for the outer bank is that there was a drop in river stage between the 2016 and 2022 images, so the 2016 image had water covering most of the vegetation on the bars, whereas the 2022 image has vegetation exposed (Figure 3). ***There are two key reasons why this observation does not bias our results.*** First, we only analyze the displacements from the *eroding* side of the river (see the caption for Fig. 3 in the main text). We do not average the displacements on the eroding and accreting sides of the river, as would be the case for a centerline-tracking analysis method^{13,30}. Second, we correct for the possible biases that stage changes can impart by implementing the methodology illustrated in Figure S1 of the Supplementary Information document. We describe each reason in more detail below.

1. **We quantify riverbank erosion rates from the bank displacements on the eroding side of the river, rather than the motion of the channel centerline (i.e., the average of the eroding and accreting sides of the river).** Reviewer #3 astutely observed in Figure 3 that the bank displacements measured from the pair of 2016 and 2022 Sentinel-2 satellite images show consistently larger migration of the inner-banks than the outer-banks, an observation that is consistent with a drop in river stage from the 2016 to the 2022 image. However, as illustrated in Figure S3 of the Supplementary Information, the high-slope outer-banks are $\sim 5\times$ less sensitive to stage changes than the low-slope inner banks. So, for example, a 1-m water level drop on the Koyukuk River (a significant stage change, since the bankfull hydraulic radius of the Koyukuk (measured near Huslia, Alaska) is just ~ 9 m—see Figure S3 of the Supplementary Information), would result in an apparent bank migration of -2.7 m. Typical erosion rates at river bends on the Koyukuk River near Huslia are 5 m/yr (Main text, Fig. 1), corresponding to ~ 30 m of bank displacement from 2016–2022. Thus, even in the worst case scenario in which our bank-tracking method is fully sensitive to water level, the apparent migration from the stage change only represents about 10% of the total river migration. Additionally, note that our workflow uses the normalized difference vegetation index (NDVI), a spectral band ratio that makes water and unvegetated sediment appear virtually identical, and in sharp contrast to the vegetated floodplain (see Extended Data Fig. 6). For the accreting inner banks, this decision to use NDVI means that we are highly sensitive to whether vegetated low-elevation bars are submerged with water or not, as noted by Reviewer 3. However, the eroding outer banks typically are devoid of vegetation (see representative field photos in Extended Data Fig. 8 and Supplementary Information Figures S7–S9). Therefore, changes in water level simply substitute pixels representing water with pixels representing bare sediment, both of which yield similar NDVI values. Thus, we believe that our NDVI-tracking method is robust to water-level changes for the eroding side of the river, which is the side of the river that we use for our analysis. See **lines 413–424** and the **Fig. 3 caption** in the revised manuscript.
2. **We correct for biases caused by changes in river stage by measuring the apparent migration of stable river reaches.** Whereas the river bends on the Koyukuk River experience typical migration rates of ~ 5 m/yr (Main text, Fig. 1), many of the straight (zero-curvature) reaches have remained stable, experiencing less than ~ 1 m of total erosion or accretion over the 2016–2022 observational window. We use these stable straight reaches as observational controls to identify and remove seasonal biases in our river migration records that could be caused by changes in water level or other factors such as changes in the NDVI of the forest canopy on the floodplain. Specifically, we perform the same displacement estimates followed by the stacking, bracketing, and MCMC analysis shown in Extended Data Fig. 7 for the stable straight reaches, producing a seasonal curve of river migration analogous to Fig. 3d of the main text. The seasonal curve computed from the stable straight reaches represents a zero-erosion baseline signal that we subtract from the seasonal curve computed for the rest of the river (i.e., bends experiencing non-negligible erosion, ≥ 1 m) in order

to correct for the types of biases illustrated in Figure S1 of the Supplementary Information and to produce the final curve shown in Fig. 3d of the main text. See **lines 511-528** in the revised manuscript for a description of this procedure.




Figure 2. Red arrows highlight inner banks that show significantly higher migration rates than the corresponding outer banks.

Figure 3: Figure from Reviewer 3 illustrating how the inner banks appear to have higher migration rates than the outer banks, consistent with the expectations from the diagram in Figure S3 of the Supplementary Information.

4. The ability to extract bank migration data at a temporal frequency so that an annual erosion time series can be constructed (Fig. 3d), despite the difficulties that arise from variable stage and coregistration errors is truly impressive. However, it would be helpful to see more data that supports this result and shows how it was obtained; e.g. a sequence of image chips from one summer, with the associated raw migration vectors, and how the final erosion time series looks like for that location. Much of the real data examples shown in the manuscript come from the Sentinel data and the arguably more interesting and innovative work that focuses on the higher resolution Planet Labs imagery is not illustrated.

We have explored displaying the sequences of image chips, but found that this visualization is not very informative for the human eye because the total displacements are very small relative to the chip size. For example, a sequence of 10 images representing one year at a high-eroding river bend may have ~ 1 pixel of total change from start to finish. But that change is difficult to see amidst the 32×32 pixel image chips, especially in the face of the co-registration errors that exist in the raw satellite data. In some ways, the most exciting aspect of our new methodology (its ability to detect sub-pixel changes) makes it

particularly unexciting to see visually in the form of image data.

The visualization that we found to be most effective is the pairwise image displacement matrix shown in Extended Data Fig. 7f. We like this visualization because it conveys to the reader both the signal and the noise. In Extended Data Fig. 7f, the signal is the monotonic change from blue to red, and the noise is the non-monotonic speckle interrupting the monotonic trend. We discuss how to interpret this figure in the Extended Data Fig. 7 caption, and also in the Methods section ***Extracting robust erosion rates from noisy timeseries*** (lines 451-528).

5. Figure 3 caption - “Note that we only study the seasonal pattern of bank displacement on the eroding side of the river.” - This seems like an important part of the methods but there is no other information about how and why this is done. Are erosional banks manually segmented for each time step? Or do you use an automated approach to identify these segments? Looks like the restriction to the eroding side only applies to the study based on the Planet Labs data – why not use the same approach for both the Sentinel and Planet Labs datasets? This approach would take care of the problems that come from ‘enhanced’ migration rates on inner banks in the Sentinel data.

Yes, we agree that this approach helps to take care of the problems that come from ‘enhanced’ migration rates on inner banks (see our responses above). We use this approach (quantifying erosion rates by tracking only the erosional outer bends) in our analyses of both the Sentinel and Planet Labs datasets. Also, we use an automated pipeline for differentiating between the eroding outer bends and the accreting inner bends. We now describe these details in section 2 of the Supplementary Information (lines 12-83). Thank you for the feedback.

6. Extracting robust erosion rates from noisy timeseries: it seems to be that this is one of the most interesting and novel contributions in the paper, certainly when it comes to the methods used, yet I don’t fully follow what exactly is being done and why. How do ‘stacking’ and ‘bracketing’ affect the result? For example, it looks to me like the line in Extended Data Fig. 7d is the same as the first row of values in panel (e) of the same figure — so how do stacking and bracketing influence this? If the displacement is averaged along a row, you end up with a single average displacement for the entire time series - I doubt that this is what is happening. The cumulative displacement in the synthetic erosion data is already monotonic; I think that going through this exercise with actual data (like the data shown in Extended Data Fig. 7f) would be more insightful than relying on synthetic data that has no obvious ‘backstepping’. Along the same lines, it is unclear to me why the MCMC approach is needed to get a monotonic displacement series; isn’t it easier to just clip the ‘backsteps’, like it is often done when generating stratigraphy from topographic time series?

Great questions. We have revised the text in the section ***Extracting robust erosion rates from noisy timeseries*** (lines 451-528) to better explain the methodology and address your questions. Direct answers to your questions are given here:

- How do ‘stacking’ and ‘bracketing’ affect the result? For example, it looks to me like the line in Extended Data Fig. 7d is the same as the first row of values in panel (e) of the same figure — so how do stacking and bracketing influence this?

You interpreted the figure correctly. The first row (or, for that matter, any row) of the synthetic data matrix in Extended Data Fig. 7e shows the same data as in Extended Data Fig. 7d. However, the rows of the real data matrix in Extended Data Fig. 7f show the actual (and therefore noisy) realizations of the cumulative displacement record. So, ‘stacking’ and ‘bracketing’ have no impact on the result if the data are perfectly clean (noise-free), but they allow one to extract cleaner and more robust displacement records when the data are noisy (like in Extended Data Fig. 7f). See the revised caption for Extended Data Fig. 7.

- If the displacement is averaged along a row, you end up with a single average displacement for the

entire time series - I doubt that this is what is happening.

Good catch; stacking refers to the averaging of the differential displacements (the $\frac{d}{dx}$ values) along *columns* of the matrix. We have revised the caption of Extended Data Fig. 7.

- The cumulative displacement in the synthetic erosion data is already monotonic; I think that going through this exercise with actual data (like the data shown in Extended Data Fig. 7f) would be more insightful than relying on synthetic data that has no obvious ‘backstepping’.

It perhaps is difficult to see, but if you zoom in to Extended Data Fig. 7g, you will see that the data are not monotonic. There is some ‘backstepping’ interrupting the progressive cumulative displacement.

- Along the same lines, it is unclear to me why the MCMC approach is needed to get a monotonic displacement series; isn’t it easier to just clip the ‘backsteps’, like it is often done when generating stratigraphy from topographic time series?

Good question. There are many ways to produce monotonic displacement series from the noisy data, and indeed the most straightforward approach would be to ‘clip’ the backsteps (by marching through the time-series either forward or backward in time, and shaving down the observed (noisy) displacements in order to maintain monotonicity. The reason we opt for MCMC rather than the clipping approach is that, when we experimented with the clipping method, we found that it often resulted in putting too much weight on noisy data points. The MCMC approach we use is widely applied for generating age models for stratigraphic data^{26,27,28,29}, which is an analogous problem where the data points are geochronologic ages (e.g., from U-Pb or Ar-Ar dating) and the monotonic path comes from the notion that the age model must obey stratigraphic superposition (i.e., a sediment layer should be younger than the layer beneath it)^{26,27,28,29}. An advantage of the MCMC approach is that it treats each data point probabilistically (as a noisy realization of some underlying ‘truth’). Thus, the MCMC path is not forced to go through any individual data point. Instead, it seeks the monotonic path that best fits the entire ensemble of data. In more practical/applied terms, this means that the MCMC is good at using neighboring data points to inform the interpretation of any individual data point. If a noisy, non-monotonic data point is surrounded by data points that otherwise agree, the noisy data point largely will be neglected in the final MCMC path. See **lines 498-510** in the revised manuscript.

7. Removal of seasonal bias: this also seems like a great idea and useful trick but it is hard to follow (for me at least) how exactly it is done. Is there a single mean value of E_{slow} computed for each time step? If so, I don’t think it accounts for all the stage-related change along a gently dipping point bar. Also, is this method used to remove seasonal bias or coregistration errors, or both? The section title refers to seasonal biases but then the text says that these displacements most likely reflect coregistration noise. In addition, if 68% of the raw data consists of this noise, it would be good to see how some of the data looks like before and after this processing step.

We edited the section about the removal of seasonal biases (**lines 511-528**) to clarify how we view co-registration errors in the context of the seasonal biases. We write: *Displacements from co-registration errors (the $\{\Delta x, \Delta y\}$ shifts from image-to-image) largely cancel out by stacking the displacement observations from many banks.* That is, we see co-registration errors as an important source of noise, but not of seasonal bias. We agree that our approach does not account for all the stage-related change along a gently dipping point bar. However, recall that none of our results use accretion data from the point bar; we only analyze the erosion of the cutbank.

8. Using NDVI means that the algorithm detects changes in the vegetation line, not erosion and deposition. This is likely to be fine on steep cutbanks, but things are more complicated on gently dipping depositional banks where, for example, young vegetation might be covered by the high-stage river, so that stage variability might still be an issue. This is discussed to some degree in the Methods section,

but it should be mentioned in the image pre-processing' section as well (and maybe in the main text).

Yes, we absolutely agree. We have clarified the text to explain that the NDVI approach does not work as well on the accreting point bars, but that does not matter for our results because we only analyze data from the eroding cutbanks. See the **Fig. 3 caption** in the main text and **lines 12-82** of the Supplementary Information.

Line-specific comments

Line 15 - “not decades as in previous work” - some of the previous work used years (sometimes less than one year), not decades, especially in rapidly migrating rivers

We modified that sentence to say “...not years or decades as used in previous work”...

Lines 18-19 - “We develop computational tools to resolve the spatial and temporal patterns of riverbank erosion at length scales $\sim 100\times$ finer than previous analyses” - this is a bit too broad brush, as part of the improvement in resolution comes from working with 3m-resolution data instead of 30m-resolution data

Good point—the $100\times$ improvement in detection comes from a $10\times$ improvement in the spatial resolution of the satellite imagery (30 m to 3 m) and a $10\times$ improvement in the sensitivity of the detection method (using our sub-pixel algorithm that can detect change on lengthscales $\sim 10\times$ smaller than the pixel size). We have edited the text in **lines 18-20** to clarify this breakdown.

Lines 61-62 - ‘Existing methods (requiring whole-pixel change in 30 m Landsat satellite imagery) would require time intervals of at least 10 years to detect erosion.’ - this is not entirely true; e.g., we were able to detect changes of less than 30 m with our workflow (Sylvester et al., 2019; 2021).

Good point. We reference Langhorst & Pavelsky (2023)¹⁴ at the end of our sentence “*Existing methods (requiring whole-pixel change in 30 m Landsat satellite imagery¹⁴)*”. The Langhorst & Pavelsky (2023)¹⁴ method is an area-based framework³² rather than the centerline-based framework of Sylvester et al., 2019; 2021^{30,33}. The area-based approach, although more flexible for handling rivers with different morphologies (e.g., braided vs. single-threaded) and mobility patterns (e.g., avulsion vs. meandering)³⁴, is more beholden to whole-pixel change (i.e., it hinges on counting pixels that change from water to land or vice versa)^{14,32}.

Our goal in this sentence is simply to draw contrast between our method and existing methods that use as their first step the binarization of a satellite image into a mask of water vs. land (river vs. floodplain). To our knowledge, all existing methods use this preliminary thresholding step before calculating channel displacement^{14,30,33,35,34,32}. This first step removes much of the information at the river margins (i.e., the continuous pixel values that encode information about the fraction of a mixed pixel that represents river vs. land). You raise a good point, however, that sub-pixel change of the channel centerline^{30,33}—or of the average position of the bank^{14,35,34,32}—can be detected even when the river itself is identified with a whole-pixel binary classification. For example, Rowland et al. (2016)³² and Langhorst & Pavelsky (2023)¹⁴ both make the approximation/assumption that, if the river moves by 0.1 pixels between two images, and change is evaluated across a reach length of 200 pixels, then approximately 20 of those cross-sections will experience a whole-pixel binary change from 1 (land) to 0 (water)^{32,14}. In practice, the detection limit from using binary masks is not as good as this assumption would suggest^{14,32}. For example, one can not detect arbitrarily small changes (e.g., 0.01 pixels) by averaging across very large N . In summary, we agree with your feedback, and we modified the text to emphasize our distinction with methods that involve the analysis of binary masks rather than continuous pixel intensity values. We also softened the language associated with whether these existing methods can detect sub-pixel changes. See **lines 61-63** in the revised text.

Lines 81-83 - “A critical component of our methodology is a workflow to make the displacement estimates

robust to georeferencing errors in the satellite imagery, since image-to-image co-registration errors can exceed the signal of riverbank erosion by $\geq 10\times$.” - this is indeed a critical component and probably deserves a bit more detail in the main text. The same applies to dealing with stage variability / seasonal biases.

Thanks for the feedback. We currently are limited in manuscript length (see the “**LENGTH**” feedback in the **Editorial comments and formatting guidelines** section below), so we are restricted in how much detail we can add to the main text. However, we agree that the robustness of our method to co-registration errors is pivotal for being able to detect the small, sub-seasonal changes in the riverbank position. Therefore, we added more details to **lines 450-528** to explain how our methodology handles co-registration errors, stage variability, and seasonal biases.

Line 92 - “we can ask” instead of ‘we we can ask’

We corrected the typo.

Fig. 2 - How was the permafrost probability computed for each bend? E.g., do you only consider the pixels that are in contact with the eroding bank? Or do you consider a broader band? Looking at Fig. 2b, this can have a large impact on the permafrost probability values.

Yes, in the Pastick et al. (2015)³⁶ permafrost map, the river often is surrounded by a halo of lower permafrost probability values on both the eroding and the accreting sides of the river. First, we sought to constrain whether this observation is an artifact of the methodology used to produce the permafrost map³⁶, or a real phenomenon whereby the river water is conducting heat into the surrounding floodplain and thawing the permafrost (forming a ‘talik’^{13,37,38}). One of the ways we addressed this ambiguity is with field measurements of permafrost presence/absence using a permafrost probe that we inserted into the eroding cutbank (see Extended Data Fig. 10). These ground-truth observations revealed that, even in places where there is a halo in the Pastick et al. (2015)³⁶ permafrost map towards lower permafrost probabilities right next to the river, we still detected near-surface (≤ 1 m) permafrost in the eroding cutbank (Extended Data Fig. 10). Thus, we found that digitally ‘eroding’ the Pastick et al. (2015)³⁶ permafrost map away from the river reduced the bias associated with the halos and led to better agreement with our ground-truth observations. Thus, the two filtering steps we perform on the original Pastick et al. (2015)³⁶ permafrost map are:

1. Morphological erosion of the permafrost map by a distance of 1 pixel away from the river.
2. Moving median filter with a window size of 5 pixels.

The first step reduces the halo, and the second step averages pixel values in a 5×5 window. Since the river appears as NaN values in the permafrost map (and therefore river pixels are not included in the 2D median filter), the effect of the 5×5 median filter is to assign the near-bank pixels a value that represents the average (median) of the permafrost probability values over a distance of ~ 2.5 pixels away from the river. See Extended Data Fig. 10 in the revised manuscript.

Fig. 3d - This is a remarkable result but it would be great to see example images and the derived data that show, for example, that erosion rates in June are significantly larger than in May.

As we mentioned in a comment above, since the riverbank position changes are sub-pixel on the scale of individual months (e.g., May vs. June vs. July), the raw satellite image data do not illustrate this result to the reader. We think that the visualization in Fig. 3d best represents the signal (and the associated uncertainty envelope). We are working on a long-format methods manuscript with field data that includes drone imagery acquired over the course of the season (May, June, July, August, September, October). Since the drone imagery is of much higher spatial resolution (e.g., ~ 10 cm pixels), it can be used to directly illustrate the seasonal pattern of bank erosion.

Line 468 - it would be helpful to briefly explain somewhere that ‘MCMC’ stands for ‘Markov-Chain Monte Carlo’ and how it relates to the method described here.

We now spell out ‘Markov-Chain Monte Carlo’ and we briefly introduce the idea that we use MCMC as a way to identify monotonic paths through the bank position timeseries.

Extended Data Fig. 6 - Panel labels (a, b, c, d) are missing and caption for upper part of figure (which I guess is supposed to be panel a?) seems to be missing.

Thank you – we have added panel labels (a, b, c, d) to Extended Data Fig. 6.

Extended Data Fig. 6 - Sentinel images from 2016 and 2022 in the lower left corner of the figure are identical - I think both come from the 2016 data.

You’re right. We fixed the problem and now have different images for 2016 and 2022.

Extended Data Fig. 6 - If I am reading the text correctly, the upper part of this figure illustrates how a small displacement between two images can be recovered using the sub-pixel method. I think it would be useful to show an actual example of change detection as well where the difference between the images is more obvious.

Yes, your reading of the text is exactly right. The purpose of that figure panel is to illustrate how small changes in the riverbank position between two images can be recovered using our sub-pixel method. We also use that particular example because it shows how the stage change (i.e., the high water level in the 2016 image and low water level in the 2022 image) has a limited effect on the NDVI-transformed data. That is, the water pixels and the unvegetated sand bar pixels have very similar (low) NDVI values, which are easily distinguished from the vegetated floodplain (high NDVI values).

Extended Data Fig. 6 - discussion of using NDVI - seems like this should be in the text, not the figure caption. Also, it seems to assume that the location of the vegetation line is immune to stage variability and it tracks deposition precisely.

We discuss our use of the normalized difference vegetation index (NDVI) in the Methods (see the section *Image pre-processing*). Importantly, note that we do not rely on the assumption that the vegetation line tracks deposition (inner-bank accretion). This is because our measurements of riverbank erosion track the motion only of the eroding cutbank (not of the change in the channel centerline position). See, for example, the illustration in Extended Data Fig. 6 of the image chip that is centered on the bank edge on the eroding cutbank. Thus, our use of NDVI is based on the assumption that, when erosion occurs, vegetated floodplain terrain is replaced by water. We do not need to assume the reverse (e.g., that when accretion occurs, the newly-deposited river sands instantly become vegetated). Indeed, that second assumption would be flawed! See **lines 12-82** of the Supplementary Information for a more detailed discussion of this concept.

I hope you find these comments useful.

Thank you for the thorough and thoughtful feedback; your comments have helped us substantially improve the manuscript.

Referee #3 (Remarks on code availability)

Access to the code repositories was not provided, so I could not review the code. Having access to it ⁶¹¹ would have been useful as a lot of the details of the methods are (somewhat understandably, as not all ⁶¹² the details can go into the text) missing or difficult to follow in the text.

We published the code (and supporting data) through the NSF’s Arctic Data Center. The code now is available online at: <https://doi.org/10.18739/A2HM52M6Q>.

References

- [1]. Post, E. *et al.* The polar regions in a 2°C warmer world. *Science Advances* **5** (2019).
- [2]. Miner, K. R. *et al.* Permafrost carbon emissions in a changing Arctic. *Nature Reviews Earth & Environment* **3**, 55–67 (2022).
- [3]. Serreze, M. C. & Barry, R. G. Processes and impacts of Arctic amplification: A research synthesis. *Global and Planetary Change* **77**, 85–96 (2011).
- [4]. Schuur, E. A. *et al.* Climate change and the permafrost carbon feedback. *Nature* **520**, 171–179 (2015).
- [5]. MacDougall, A. H., Avis, C. A. & Weaver, A. J. Significant contribution to climate warming from the permafrost carbon feedback. *Nature Geoscience* **5**, 719–721 (2012).
- [6]. Schaefer, K., Lantuit, H., Romanovsky, V. E., Schuur, E. A. & Witt, R. The impact of the permafrost carbon feedback on global climate. *Environmental Research Letters* **9**, 085003 (2014).
- [7]. Treat, C. C. & Frolking, S. A permafrost carbon bomb? *Nature Climate Change* **3**, 865–867 (2013).
- [8]. Mascarelli, A. A sleeping giant? *Nature Climate Change* **1**, 46–49 (2009).
- [9]. Brown, A. Pandora’s freezer? *Nature Climate Change* **3**, 442–442 (2013).
- [10]. Jones, M. C. *et al.* Past permafrost dynamics can inform future permafrost carbon-climate feedbacks. *Communications Earth & Environment* **4**, 272 (2023).
- [11]. Torres, M. A. *et al.* Model predictions of long-lived storage of organic carbon in river deposits. *Earth Surface Dynamics* **5**, 711–730 (2017).
- [12]. Rowland, J. C. *et al.* Scale-dependent influence of permafrost on riverbank erosion rates. *Journal of Geophysical Research: Earth Surface* **128**, e2023JF007101 (2023).
- [13]. Ielpi, A., Lapôtte, M. G., Finotello, A. & Roy-Léveillé, P. Large sinuous rivers are slowing down in a warming Arctic. *Nature Climate Change* 1–7 (2023).
- [14]. Langhorst, T. & Pavelsky, T. Global observations of riverbank erosion and accretion from Landsat imagery. *Journal of Geophysical Research: Earth Surface* **128**, e2022JF006774 (2023).
- [15]. Chadburn, S. *et al.* An observation-based constraint on permafrost loss as a function of global warming. *Nature Climate Change* **7**, 340–344 (2017).
- [16]. Hughes, I. & Hase, T. *Measurements and their uncertainties: a practical guide to modern error analysis* (OUP Oxford, 2010).
- [17]. Douglas, M. M. & Lamb, M. P. A model for thaw and erosion of permafrost riverbanks. *Journal of Geophysical Research: Earth Surface* **129**, e2023JF007452 (2024).
- [18]. Douglas, M. *et al.* Geomorphic mapping and permafrost occurrence on the Koyukuk River floodplain near Huslia, Alaska (ESS-DIVE dataset) (2023).
- [19]. Douglas, M. M., Dunne, K. B. & Lamb, M. P. Sediment entrainment and slump blocks limit permafrost riverbank erosion. *Geophysical Research Letters* **50**, e2023GL102974 (2023).
- [20]. Douglas, M. M. *et al.* Discontinuous permafrost degradation and regrowth following meandering of the Koyukuk River, Alaska. *AGU Advances, in revision* (2024).
- [21]. Geyman, E. *et al.* Floodplain terrain age at three sites in the Yukon River Watershed: Huslia, Alakanuk, and Beaver (Alaska). *Arctic Data Center* (2024).
- [22]. Geyman, E. *et al.* Field and LiDAR-based estimates of above-ground biomass in the Yukon River Watershed, Alaska (2023). *Arctic Data Center* (2024).
- [23]. Kean, J. W. & Smith, J. D. Flow and boundary shear stress in channels with woody bank vegetation. *Riparian vegetation and fluvial geomorphology* **8**, 237–252 (2004).
- [24]. Smith, S. L., O’Neill, H. B., Isaksen, K., Noetzi, J. & Romanovsky, V. E. The changing thermal state of permafrost. *Nature Reviews Earth & Environment* **3**, 10–23 (2022).
- [25]. Constantine, C. R., Dunne, T. & Hanson, G. J. Examining the physical meaning of the bank erosion coefficient used in meander migration modeling. *Geomorphology* **106**, 242–252 (2009).
- [26]. Schoene, B. *et al.* U-Pb constraints on pulsed eruption of the Deccan Traps across the end-Cretaceous mass extinction. *Science* **363**, 862–866 (2019).
- [27]. Keller, C. Chron.jl: a Bayesian framework for integrated eruption age and age-depth modelling. *OSF (Open Science Framework)* **10** (2018).
- [28]. Schoene, B., Eddy, M. P., Keller, C. B. & Samperton, K. M. An evaluation of Deccan Traps eruption rates using geochronologic data. *Geochronology* **3**, 181–198 (2021).
- [29]. Zhang, T. *et al.* A Bayesian framework for subsidence modeling in sedimentary basins: A case study of the Tonian Akademikerbreen Group of Svalbard, Norway. *Earth and Planetary Science Letters* **620**, 118317 (2023).
- [30]. Sylvester, Z., Durkin, P. & Covault, J. A. High curvatures drive river meandering. *Geology* **47**, 263–266 (2019).
- [31]. Savitzky, A. & Golay, M. J. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry* **36**, 1627–1639 (1964).
- [32]. Rowland, J. C. *et al.* A morphology independent methodology for quantifying planview river change and characteristics from remotely sensed imagery. *Remote Sensing of Environment* **184**, 212–228 (2016).
- [33]. Sylvester, Z., Durkin, P., Hubbard, S. & Mohrig, D. Autogenic translation and counter point bar deposition in meandering rivers. *GSA Bulletin* **133**, 2439–2456 (2021).
- [34]. Greenberg, E., Chadwick, A. J. & Ganti, V. A generalized area-based framework to quantify river mobility from remotely sensed imagery. *Journal of Geophysical Research: Earth Surface* **128**, e2023JF007189 (2023).
- [35]. Chadwick, A. J., Greenberg, E. & Ganti, V. Remote sensing of riverbank migration using particle image velocimetry. *Journal of Geophysical Research: Earth Surface* **128**, e2023JF007177 (2023).
- [36]. Pastick, N. J. *et al.* Distribution of near-surface permafrost in Alaska: Estimates of present and future conditions. *Remote Sensing of Environment* **168**, 301–315 (2015).
- [37]. Roy-Leveille, P. & Burn, C. R. Near-shore talik development beneath shallow water in expanding thermokarst lakes, Old Crow Flats, Yukon. *Journal of Geophysical Research: Earth Surface* **122**, 1070–1089 (2017).
- [38]. Langer, M. *et al.* Rapid degradation of permafrost underneath waterbodies in tundra landscapes—toward a representation of thermokarst in land surface models. *Journal of Geophysical Research: Earth Surface* **121**, 2446–2470 (2016).

Reviewer Reports on the First Revision:

Referee #2 (Remarks to the Author):

To the authors -- particularly Dr-in-waiting Geyman:

Fantastic work on this study and addressing wide-ranging reviewer questions and comments.

I believe that this paper is a strong entry in the debate about permafrost and river migration in the arctic and is ready for publication. In addition to its main science finding, it gives us a tool to use in the arctic and elsewhere to measure sub-seasonal erosion rates.

Evan Dethier

Referee #2 (Remarks on code availability):

See previous review.

Referee #3 (Remarks to the Author):

I have gone through the revised manuscript and I think the authors have done a great job at addressing the reviewers' comments and suggestions. All of my key concerns have been dealt with and the paper is in a great shape.

I think there are two points where minor changes still should be considered (but not absolutely necessary for publication):

(1) Making it somehow more obvious in Fig. 1c, 1d, and 2 that only the erosional bank segments are used when comparing permafrost and non-permafrost migration rates - e.g., by highlighting the segments that provide the data for Fig. 2e-g.

(2) I am aware that it might be difficult to squeeze any new material in, and I understand that 32x32 pixel chips from the Planet Labs imagery do not illustrate well the intra-seasonal changes, but I still think it would be valuable to show some Planet imagery in the paper or at least in the supplementary material. Also, showing an unequivocal example of rapid bank retreat in the summer months with the help of the high-resolution drone imagery would strengthen the case that the variations in migration rate at this temporal scale are indeed due to melting permafrost.

Zoltan Sylvester

Referee #3 (Remarks on code availability):

Code and data has been made available through repositories and they are easy to access. However, I have not been able to run the code as I was unable to install Matlab on my computer.

Author Rebuttals to First Revision:

Reviewer 2

To the authors – particularly Dr-in-waiting Geyman:

Fantastic work on this study and addressing wide-ranging reviewer questions and comments.

I believe that this paper is a strong entry in the debate about permafrost and river migration in the arctic and is ready for publication. In addition to its main science finding, it gives us a tool to use in the arctic and elsewhere to measure sub-seasonal erosion rates.

Evan Dethier

Thank you for the kind words and for all of your detailed feedback from round 1 that helped us improve the manuscript.

Reviewer 3

I have gone through the revised manuscript and I think the authors have done a great job at addressing the reviewers’ comments and suggestions. All of my key concerns have been dealt with and the paper is in a great shape.

Thank you for the feedback.

I think there are two points where minor changes still should be considered (but not absolutely necessary for publication):

(1) Making it somehow more obvious in Fig. 1c, 1d, and 2 that only the erosional bank segments are used when comparing permafrost and non-permafrost migration rates - e.g., by highlighting the segments that provide the data for Fig. 2e-g.

In response to your suggestion, we revised Fig. 1 of the main text (see Fig. 1 below) to include dashed lines in panel (e) to highlight the erosional bank segments that are used for comparing permafrost and non-permafrost migration rates.

(2) I am aware that it might be difficult to squeeze any new material in, and I understand that 32x32 pixel chips from the Planet Labs imagery do not illustrate well the intra-seasonal changes, but I still think it would be valuable to show some Planet imagery in the paper or at least in the supplementary material. Also, showing an unequivocal example of rapid bank retreat in the summer months with the help of the high-resolution drone imagery would strengthen the case that the variations in migration rate at this temporal scale are indeed due to melting permafrost.

Thank you for the suggestion. We added two figures to the end of the Supplementary Information document (Figures S10-S11) to show example PlanetScope satellite image timeseries. The figures are included in this letter as Figures 2-3 below. The zoomed-in figure (Figure 3) helps the reader to understand and visualize the 32×32 pixel chips, since 32 pixels corresponds to 96 meters, which is approximately equal to the 100 m scale bar shown in that figure. We agree that this visualization is a useful addition to the supplementary materials. Thank you for your input.




Figure 1: A revised version of Fig. 1 of the main text. In response to the suggestion from Reviewer 3, we added dashed lines to highlight the erosional side of each river bend in panel (e). See Fig. 1 of the main text for the complete figure caption.




Figure 2: An example of a seasonal satellite image timeseries (using PlaneScope imagery) of an erosional river bend along the Koyukuk River. The image from May 19, 2023 (top left) captures the ‘ice-out’ period; blocks of floating river ice are traveling downstream with the current. The dashed white box shown in the image from July 19, 2023 (top right) shows the zoomed-in window for Figure 3. Coordinates are provided in WGS84 UTM zone 4N. This figure is included in the *Supplementary Information* as Figure S10.




Figure 3: A zoomed-in version of Figure 2. This figure is included in the *Supplementary Information* as Figure S11.