

Reference #1:

1. The paper's central prescriptive recommendations, such as the use of Kahan's formula for edge length calculations, the preference for Eriksson's formula over L'Huilier and quadrature-based methods for face area computation, and the conclusion that k-d trees and ball trees exhibit broadly similar performance, are supported by benchmark results presented only in the supplementary material. Readers should not need to consult the supplement to assess whether the paper's core numerical claims are well supported.

I encourage the authors to incorporate the key quantitative findings from Supplement into the main text, including, at a minimum, representative error magnitudes and relative performance comparisons for the recommended approaches. Supplement can still serve an important role by presenting complete tables and additional comparisons.

We would like to thank the reviewer for careful consideration of our lengthy manuscript.

We agree that the key numerical claims supporting our prescriptive recommendations should be directly supported by quantitative evidence in the main text, without requiring readers to consult the supplementary material.

In the revised manuscript, we have incorporated representative error magnitudes and performance comparisons into the main text for all major recommendations, while retaining the supplementary material for complete benchmark tables and extended results.

Specifically:

- **Edge length computation (Kahan's formulation):**

We now report that Kahan's method achieves relative errors on the order of a small multiple of the machine epsilon u (approximately 10^{-15} in double precision), with no measurable computational overhead. In contrast, alternative formulations exhibit substantial degradation in ill-conditioned configurations, with errors increasing by several orders of magnitude (ranging from approximately 10^{-11} up to 10^{-3} , depending on the baseline angle).

- **Face area computation (Eriksson formulation):**

We now include explicit quantitative comparisons showing that the Eriksson formulation maintains relative errors on the order of a small multiple of u

(approximately 10^{-13} in double precision) across all tested resolutions. Quadrature-based methods achieve errors on the order of $10^{-3} u$ (approximately 10^{-12}), while the L'Huilier formulation can degrade to errors on the order of 10^{-5} in near-degenerate cases, reflecting significant precision loss.

- **Spatial indexing (k-d trees vs. ball trees):**

We now provide explicit performance comparisons indicating that KDTree construction is moderately slower (within a factor of ~ 1.5 – 2), while query performance is consistently faster by approximately one order of magnitude ($\sim 10\times$) across tested configurations.

These additions ensure that the main text contains sufficient quantitative evidence to support the paper's central recommendations, while the supplementary material continues to provide full benchmark data and additional comparisons.

2. The AccuCross and AccuXGCA algorithms represent key novel contributions of the paper. However, the manuscript does not currently include error bounds or empirical demonstrations of their accuracy, noting instead that such analysis is deferred to future work. To strengthen the paper, it would be helpful for the authors to include, at a minimum, either asymptotic error bounds or a concise empirical assessment of accuracy for AccuCross and AccuXGCA in the main text or appendix.

In the original submission, we provided empirical accuracy evaluations of AccuCross and AccuXGCA in the supplementary material, demonstrating that AccuXGCA consistently achieves errors near machine precision. In particular, its accuracy effectively matches that of evaluating the same expressions at approximately twice the working precision followed by rounding back to the target precision, without measurable performance overhead.

In the revised manuscript, we have now added a concise analytical error characterization to the main text, with full derivations provided in the supplementary material. Specifically, we now state that the relative error of AccuCross is bounded by $O(u^2)$, while the relative error

of AccuXGCA remains $O(u)$ in both asymptotic regimes considered. In contrast, the direct floating-point implementation exhibits an error bound of $O(\sqrt{u})$, corresponding to a loss of approximately half of the working precision.

3. Several contributions are described as being implemented in existing tools such as TempestRemap but documented here for the first time, including the extremal latitude formula and the sweep-line bounding rectangle algorithm. While documenting previously undescribed algorithms is valuable, the paper would benefit from a clearer and more explicit delineation, ideally in the introduction or in a dedicated summary, of which elements are algorithmically novel and which are being formally described for the first time.

We agree that clearly distinguishing between novel algorithmic contributions and existing methods that are formally documented in this work improves the clarity and positioning of the paper. In the revised manuscript, we have made this distinction explicit in the Introduction by introducing a notation to classify the status of each component in the regridding pipeline. Specifically, we now use a single asterisk to denote operations for which suitable algorithms exist in current codebases but are, to our knowledge, systematically documented here for the first time, and two asterisks to denote operations for which we propose novel methods.

and finally for the Minor Comment: The manuscript is well written, and the English is of a high standard throughout. Nevertheless, a careful final proofreading pass is recommended to catch isolated typographical errors before resubmission. As one example, Section 5.4 contains a duplicated word: "we we find from (3) that".

Thank you for catching this typographical error. We have performed a thorough proofreading of the manuscript and corrected this issue ("we we find from (3) that") along with a few others we found.

Reference #2:

I think the paper could be published in its present form, but the authors might consider attempting to make it more accessible and useful to the large number of analysts who are faced with a choice of remapping tools and would like to choose the one most appropriate for their research.

I can imagine that at the end of each section (or subsection) describing a particular operation, a summary could be provided indicating, for example, the relative merits of different approaches including 1) accuracy, 2) stability, and computational efficiency. Of course, in some cases, the differences might be minor, so a statement that all the approaches appear to be equally reliable and computationally efficient might be sufficient. The summaries provided could be supplemented by a table indicating for each of the (4?) remapping tools considered, which method is used for each of the required operations.

We would like to thank the reviewer for careful consideration of our lengthy manuscript.

We agree that summary-level guidance on the relative merits of different approaches can be valuable, particularly for readers seeking to understand practical trade-offs between methods. In the revised manuscript, we have strengthened the presentation within each subsection by more explicitly summarizing the accuracy, stability, and computational efficiency of the algorithms where appropriate. These additions are intended to improve accessibility for both practitioners and developers.

At the same time, we have chosen not to list the specific algorithms employed in particular regridding tools. This decision reflects the fact that major tools such as TempestRemap and YAC are developed in close collaboration with our team, and several of the algorithms presented in this paper are already being actively incorporated into these codebases. In many cases, implementations have already begun to evolve during the preparation of this manuscript. As a result, any static mapping between methods and tools would likely become inaccurate in the near term.

More broadly, our expectation is that these algorithms will be progressively adopted across TempestRemap and YAC, and related frameworks, and ultimately made available through higher-level interfaces such as the Python-based UXarray package. For this reason, we focus on presenting the algorithms and their numerical properties in an implementation-independent

manner, allowing both developers and users to make informed decisions based on the current capabilities of the tools they are using.

To further support practical adoption, we also provide vectorization-friendly APIs implementations of several key algorithms discussed in this work (including latitude–longitude bounding, point-in-polygon predicates, and intersection points construction) in the AccuSphGeom package.

We hope these revisions improve the accessibility and practical usefulness of the manuscript while maintaining its long-term relevance.