Response to Reviewers - Second Round

Towards robust fracture mapping: benchmarking automatic fracture mapping in 2D outcrop imagery

egusphere-2026-1097

Reply to second comment of Thomas Dewez, 01 Sep 2026

Dear Dr Thomas Dewez,

Thank you for your positive assessment and for the helpful consistency observation. Please find our response below.

Best regards,
Ayoub Fatihi, Jefter Caldeira, Tom Beucler, Sam Thiele, and Anindita Samsu

Why are figures 10 and 11 laid out as lattice plots with datasets per row, while figure 13 has a lattice plot with dataset per column. Both sets of figures seem to represent the same concept. For consistency in comparison, wouldn't it be better to have all three figures layed out identically?

Agree: We transposed Figure 13 so that, like Figs. 10 and 11, each row now corresponds to one dataset/site, with the original orthophotograph, U-Net prediction, and SegFormer prediction shown as columns within each row. The caption was updated accordingly: *"Deep-learning fracture-trace predictions for four representative orthophotographs, one per row. Each row shows the original orthophotograph (a, d, g, j) alongside the corresponding U-Net (b, e, h, k) and SegFormer (c, f, i, l) predictions."* **(see Fig. 13)**.

Reply to second comment of Billy Andrews, 11 Sep 2026

Dear Dr. Billy Andrews,

Thank you for your positive assessment of the revised manuscript and for the constructive minor comments. We are pleased the revision addressed your major concerns. Below we respond to each remaining point; the manuscript has been updated accordingly. Please note that page and line numbers in your comments refer to the manuscript sent for this round of review, while page and line numbers in our responses refer to our newly revised manuscript.

Best regards,
Ayoub Fatihi, Jefter Caldeira, Tom Beucler, Sam Thiele, and Anindita Samsu

Introduction

I suggest using "e.g.," where there are a limited number of references where there are a large number of studies in this area.

Agree: We reviewed the citation groups in the Introduction; representative examples are already marked with *"e.g.,"* where a group stands in for a broader body of work (e.g., p. 2, *"related to image quality (e.g., lighting, surface weathering, vegetation cover, and image resolution)"*), while shorter, exhaustive lists of directly relevant studies are left unmarked. We did not identify additional instances requiring this treatment, but remain happy to revisit specific cases the reviewer has in mind.

L2: "of a fractured rock mass" or "fractured rock masses"

Agree: Corrected. The abstract now reads: *"...to characterise the mechanical and hydraulic properties of a fractured rock mass and to interpret its deformation history."* **(see L2)**.

L6: should it be segmentation of?

Agree: Corrected to *"...for pixel-wise fracture segmentation of combined high-resolution RGB orthophotographs and digital elevation models (DEMs)..."* **(see L5)**.

L33: "Remarkable" should be replaced with mm to cm resolution or similar.

Agree: Replaced with a quantitative statement consistent with the resolution range already given in the abstract: *"...that capture rock surfaces at millimetre-to-centimetre resolution, including outcrops..."* **(L22)**.

L61: Bond et al., 2007 do not discuss fracture mapping- I do not believe this reference is applicable here.

Agree: Agreed; the citation was not specific to fracture mapping. We removed it from both instances in this sentence, which now cites only fracture-mapping-specific studies of interpreter variability: *"...and to the experience and preferences of the interpreter (Scheiber et al., 2015; Peacock et al., 2019; Andrews et al., 2019). As a result, fracture maps produced by different interpreters...may vary in fracture identification, segmentation, and connectivity (Scheiber et al., 2015)."* **(L41-43)**.

L74-77: Please split this sentence for readability.

Agree: Split into two sentences: *"Wavelet and multi-scale transforms (Mallat and Hwang, 1992; Candès and Guo, 2002; Candès and Donoho, 2005; Do and Vetterli, 2005; Prabhakaran et al., 2019) improve the handling of curvilinear, multi-scale features, although classical wavelets lack directionality. Dedicated line-segment detection and tracking algorithms (Rahnama and Gloaguen, 2014; Masoud and Koike, 2017; Saint Jean Patrick Coulibaly et al., 2021; Fan and Ni, 2023; Gaikwad et al., 2023) offer a complementary approach for capturing such features."* **(L62-64)**.

Figure 1: Something has gone wrong with your figure caption, a and b are c and d and vice versa.

Agree: Thank you for catching this; the figure caption had a misplaced closing brace that truncated the caption text after item (b), leaving items (c) and (d) as stray, unformatted text. This has been fixed; the caption now reads as a single block with all four items (a)-(d) correctly matching the in-text references and the panel order in the figure **(Fig. 1)**.

Methods

The limitation of the three datasets is addressed. It would be good to discuss how these limitations could impact the results briefly in your discussion. I worry that some authors could use this approach without thinking how lithology or weathering specific effects may impact their interpretations. Adding a sentence in your discussion would be enough to solve this concern.

Agree: Added to the Discussion (Sect. 4.1, Model Performance): *"We caution, however, that lithology- and weathering-specific effects (e.g., grain size, surface staining, or vegetation cover characteristic of a given site) can influence both fracture appearance and model performance; users applying this approach to new settings should validate predictions against site-specific conditions rather than assuming uniform performance across lithologies."* **(L499)**.

Fig2: Consider adding a scalebar to panel (c). Additionally, please add a north arrow to at least one of the panels.

Agree: A scale bar has been added to panel (c), and a north arrow has been added to panel (a) of Fig. 2 **(Fig. 2)**.

Fig3 is much improved. consider adding "Benchmark dataset preparation" or similar above the first part of the figure to match the other panels. Also, the numbers do not matter for this figure, and are not legible in the histograms. Consider removing them and just having blank axis labels. In the figure caption, I would add after ": not all datasets shown here" that you have added the datasets in detail in appendix A1-A3.*

Agree: We added a *"Benchmark dataset preparation"* heading above the dataset panels and removed the illegible histogram tick labels (blank axes retained). The caption now also reads: *"...*: not all outcrops used shown here. Detailed dataset overviews are provided in Appendix Figs. A1-A3."* **(Fig. 3)**.

Figure 6 - I am still struggling with the readability of this figure, with the numbers being very small. I wonder if it is required in the main text as I find your expanded caption describes the process well. Maybe consider moving Figure 6 to the supplementary and having it as a larger landscape figure.

Agree: We would prefer to keep this figure in the main text, as the U-Net architecture is central to our methodology. Instead, we have enlarged the figure and increased the font size of all internal labels (dimensions, channel counts) to substantially improve legibility. We hope this addresses the readability concern while keeping the architecture visible alongside its description **(Fig. 6)**.

Discussion

The discussion is improved but I still believe it would benefit from a small sub-section regarding the usefulness of the method for fracture network analysis studies. Your point about time in the results is very powerful, this has direct implications both within industry and academia. This would only need an additional paragraph in the discussion and I think would increase the impact of this work.

Agree: We added a new subsection, *"Implications for Fracture Network Analysis*" **(Sect. 4.4, L525)**, building on the timing result (Sect. 3.2.5): *"Beyond raw segmentation accuracy, the speed of the automated workflow has direct implications for fracture network analysis. Our timing comparison shows that generating a fracture probability map takes seconds rather than the tens of minutes to hours required for manual or semi-automated interpretation of an equivalent area (Thiele et al., 2017b). This speed-up changes what is practically feasible: instead of interpreter time limiting studies to a handful of representative outcrops, automated mapping makes it tractable to sample fracture networks over larger areas, across more outcrops, or repeatedly through time (e.g., before and after a rockfall or excavation, or across a monitoring interval), producing larger and more statistically robust datasets for connectivity, spacing, and orientation analyses. In academic settings this lowers the barrier to regional-scale fracture network studies that are currently limited by manual digitisation effort, while in industry contexts (e.g., reservoir characterisation, slope stability, or geo-energy site screening) it enables faster turnaround between data acquisition and first-pass structural interpretation. We emphasise that automated segmentation alone does not yet replace the post-processing step needed to obtain topologically consistent fracture networks; rather, it removes the main time bottleneck upstream of that step, making fracture network analysis at scale substantially more accessible."*.