Technical Note: Resolution Enhancement of Flood Inundation Grids

Response to Reviewer #1

Seth Bryant, Guy Schumann, Heiko Apel, Heidi Kreibich, Bruno Merz

Friday, October 20, 2023

We thank the editor and the reviewers for their comments. These have been very helpful in improving the manuscript. Responses to the major comments are provided below (along with the original comment in italics). Some responses to minor comments are also provided where necessary (other minor comments will simply be incorporated into the manuscript).

Reviewer #1

> *In my view, the strongest result, and one that best showcases the superiority of the new algorithm, is in Figure 6. It shows that the new algorithm can much better represent small-scale inundation extent. This is very important and, clearly, cannot be well captured by standard evaluation metrics. I encourage the author to emphasize this result (including in the abstract and conclusions) and provide additional examples (using the same flood case study). Do not shy away from qualitative comparison.*

This is a nice point which strengthens the manuscript that we had not considered. To incorporate this comment, the following sentence will be added to the abstract:

> Qualitatively, the algorithm generates more physically coherent flood maps in some hydraulically challenging regions compared to the state-of-the-art.

Similarly, the following sentence will be added to the Results and Discussion section along with the accompanying figure in the supplement:

> Qualitatively however, *CostGrow* generates more physically coherent depth grids in some fringe areas as shown in Figure 8 and S8.

And finally, the following sentence in the conclusions will be modified:

> This algorithm outperforms the state-of-the-art, with a six-fold improvement in runtime for our case study, a slight improvement in standard performance metrics, and improvements in some fringe areas using qualitative evaluation.

Adding building impact may also help showcase the improved capability at small scales.

We considered this and agree that this may also better demonstrate the improvement provided by the new algorithm; however, we felt the added complexity and length required for this comparison were not justified in a Technical Note.

The new algorithm's description needs to be extended.

This section will be expanded and the sub-routines elaborated as follows:

> The novel CostGrow algorithm employs the four phases summarized in Fig. 1: 1) grid resampling; 2) least-cost mapping; 3) filtering high-and-dry sub-grid cells; and finally 4) an isolated-cell filter all of which are parameterless. In the first grid resampling phase, various techniques have been developed by others for applications in image analysis and spatial analysis (Bierkens et al., 2000) with bilinear being the most common for terrain manipulations in hydraulic applications (Heritage et al., 2009; Muthusamy et al., 2021) as it provides a smooth result while preserving centroid values. For downscaling, bilinear resampling computes the s1 value from the four adjacent s2 centroid values weighted by distance as seen in Fig. 1a (notice the s2 values are preserved by the center s1 cell). CostGrow implements bilinear resampling from the popular spatial analysis package GDAL (GDAL/OGR contributors, 2022). In the second phase, the resampled grid is extrapolated using a cost-distance analysis, a common GIS algorithm for computing the path of least cost, determined by weighting distance and some cost map to obtain the effective distance from source cells to sink cells Foltête et al. (2008). For this study CostGrow implements a cost-distance routine with N8(P) adjacency and a neutral cost surface (Lindsay, 2014, CostAllocation). This first maps the dry portion of the domain in terms of catchment areas for each boundary W SEs1,i from the previous phase, then maps corresponding boundary W SEs1,i value to each of its catchments. In effect, this grows each W SEs1,i boundary cell value outwards, filling the dry domain with the W SEs1,i values that are closest in distance. For the toy example shown in Fig. 1b, this is a simple extrapolation onto the dry right-side of the domain. Future implementations could employ a non-neutral cost surface to incorporate levees or some other flood obstructions into the analysis. In the third phase, high-and-dry cells are filtered from this cost-distance map by comparing cell-by-cell to the terrain values (where DEMs1 > W SEs1,i set W SEs1,i = NULL) as shown by the blank cells in Fig. 1c. This often results in many isolated pockets of flooding in low-lying areas shown beyond the initial contiguous W SEs2 flood which are addressed below (see Fig. 1c red circle). In the final phase, isolated or disconnected groups of flooded cells are filtered from the result such that only the largest or main flooded water body remains. To accomplish this, the filtered grid is converted to a binary inundation grid, from which each contiguous clump is identified and ranked according to size (Lindsay, 2014, Clump) (see Fig. 1d). From the largest clump, an inverted mask is generated and applied to the water level grid to remove isolated flooding cells from the result.

Line 177 – what was the performance of s2?

This section discusses the inundation performance metrics. Line 174 states '…the results for the Resample downscaling algorithm (Fig. 6a0 and b0) show the performance of the coarse (s2) parameterization'. The original figure is pasted below.

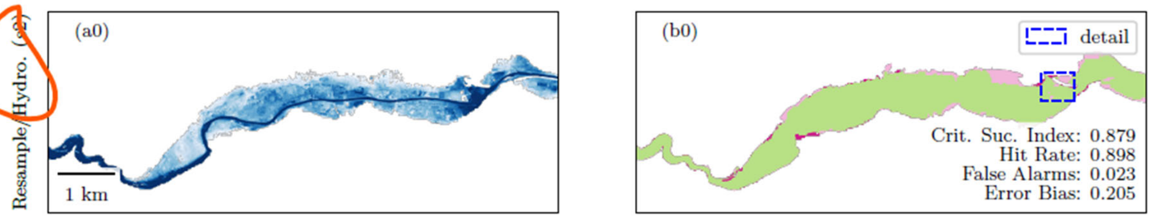

Figure 5 – 'rvalue' – is this R or R2? R2 should be reported. Explain the observations with values close to 0.

The figure will be revised to show r^2. The simulations close to zero are for HWM locations where the simulated extents predict dry (False Negative). This explanation will also be provided in the manuscript.
