

Manuscript gc-2026-4130 - "Sea level rise in Los Santos and Los Angeles: a comparative flood analysis using Grand Theft Auto V"

Response to Reviewer 1 (Chris Skinner)

Reviewer comments are reproduced in black; our replies follow in green. Line and section numbers refer to the revised manuscript unless stated otherwise.

We are grateful for a review that was both generous and specific. The recommendation to drop the classroom proposal was the most consequential comment we received, and we have followed it in full. Acting on the review also led us to re-examine the figures, where we found two errors of our own that we report below.

General remarks

From my own experience of working with game developers, it is likely that the GTAV devs will have incorporated freely available elevation data in the development of the terrain mesh for example SRTM. This would strengthen your arguments but without speaking to the developers we cannot know.

We agree this is plausible, and it would indeed strengthen the argument. We have chosen not to assert it in the manuscript, because it is exactly the kind of inference about developer intent that the reviewer asks us to mark as interpretation in the line-331 comment below. We have instead applied that standard consistently: the one place where we did make such an inference, about the editorial compression of the Ballona Creek area, is now explicitly flagged as our reading rather than established fact.

1. Remove research question 3 and Section 5.3

Although there is educational value in this work, the approach suggested is not the right one and has not been fully justified. It does not really add anything to the research. [The four supporting points concern the 18+ age rating in Germany; the comparable effort of building a real-city environment, e.g. CoastCraft; speculative and unreferenced argument at lines 369-372; and the mismatch between the lesson plan and the 'international conflict and conflict resolution' curriculum strand.] My recommendation is the authors replace 5.3 with a shorter (one or two paragraphs) discussion on possible (and appropriate) applications of the game in education and outreach. I suggest the authors consider non-formal routes, for example, the projections might be of interest to game-based and climate creators, including Any Austin or Simon Clark.

Done in full, and we accept the reasoning. The age-rating objection in particular is not one we can answer: the fact that students below the rating play the game anyway is not a basis on which to build a lesson, and we should not have treated it as one.

Specifically: research question 3 has been removed from the abstract, the introduction, the opening of Sect. 5 and the conclusions, and the paper now states that it addresses two questions. Table 4 and the entire six-lesson sequence have been deleted. Sect. 5.3 is retitled 'Applications in communication and outreach' and reduced to two paragraphs. The first states plainly why we no longer propose a classroom application, and points to CoastCraft as the more appropriate vehicle where a classroom exercise is the goal, with the added advantage that it can depict a real city rather than a fictional analogue. The second sets out the non-formal routes the reviewer suggests: games and climate creators, museum and science-centre installations, and media literacy about

flood maps themselves.

Six references cited only by the deleted section (Dickey, Klafki, Nicholson-Cole, Spence, Squire and Wdowinski) have been removed. We have retained the citation to the underlying thesis, now placed in the introduction, since the study derives from it.

On one point the new text is more explicit than the review required. Removing the objectionable content is technically straightforward, and Sect. 5.3 now says so: GTA V has a large and active modding community, mods that strip out weapons and violence are commonplace, and an age-appropriate build is therefore achievable, which is what the underlying thesis proposed. We think the objection survives that, for two reasons the section now states. The familiarity such a lesson would trade on comes from the unmodified game, which the intended age group is not supposed to have played, so the pedagogical premise and the legal position pull against each other. And building, distributing and maintaining a classroom-safe build across a school's machines costs about as much effort as producing a purpose-made environment, which is the reviewer's own second point.

2. Tighten language around simulation throughout

Much of the description and discussion of the modelled results falls into the model land problem (see Thompson, 2022). In essence, the authors discuss them as if they were reality and instead it should be clear they are discussing only possible outputs. [...] This shift in language should be made throughout the manuscript.

Done. We have adopted the reviewer's suggested rewording of the abstract almost verbatim, and applied the same shift to the four other places where results were stated as fact: the abstract sentence on RCP 2.6, and three sentences in the conclusions. Throughout, 'would inundate' is now 'with the model showing ... inundated', and 'rendering the port ... non-functional' is now 'placing the port ... at risk'. Thompson (2022) has been added to the reference list. We had not registered how consistently the original text slipped into that register, and the correction improves the paper.

3. Increase the critique of bath tub models

In Section 3.4 the authors discuss the bath tub model used for the flood simulation, stating they are appropriate for first-order assessment of coastal flood risk. It would be helpful here to point to recent critiques of these methods for use in assessing coastal flood risk, for example [Ramirez et al., 2016; Sanders et al., 2024]. For additional context [...] there has been heavy criticism for bath tub (or similar) model derived flood maps as part of public engagement with climate change, with them being a potential source of misinformation if used inappropriately. [...] This should be a consideration for any educational application of this work.

Done. Sect. 3.4 now states that static approaches are known to overestimate flood extent relative to hydrodynamic models, particularly where flow is constrained by topography or where an event is too short for the inundated volume to fill, citing Ramirez et al. (2016) and Sanders et al. (2024); both have been added to the reference list. The same paragraph now says that the maps should be read as showing land lying below a given still-water level, not as predictions of where water would reach in a specific event.

The misinformation point is taken seriously rather than acknowledged in passing: it is the third and, as we now argue, the most valuable of the outreach applications in the rewritten Sect. 5.3. A recognisable game world is an unusually good place to make the distinction between 'land below a water level' and 'where water will go' visible, because the audience can navigate the terrain themselves. The section closes by stating that this limitation has to travel with the figures whichever route is taken.

Line-by-line corrections

Line 46: Geo-games and geo-gaming already have a specific meaning, being games are based on a player's current location (e.g. Pokemon Go). I'd avoid using the term here. I'd suggest game-based approaches in geoscience instead (not very catchy...).

Corrected. The phrase 'this emerging "geo-gaming" literature' now reads 'work on game-based approaches in geoscience'. The remaining occurrences of the term were inside the deleted Sect. 5.3 and have gone with it.

Line 70: The heightmap was not extracted from the game it was reproduced. This should be clear here.

Corrected in both places where the claim appeared, in the abstract and in the closing paragraph of the introduction: 'terrain data extracted from' now reads 'terrain data reproduced from'. The distinction matters for the data-availability question raised by Reviewer 2 as well, and the statement there has been rewritten accordingly.

Table 1 (and elsewhere): State the indicative present-day magnitude of the storm surge applied here.

Added to the Table 1 caption: the 1.591 m value is the highest still-water level in the Los Angeles gauge record and is applied as an upper-bound present-day extreme rather than a surge of any particular return period.

Figure 4: Consider the accessibility of the shading used here. Green and red is not the best combination.

Corrected, and the reviewer is right by a wider margin than we expected. We tested the published green/yellow/red scheme against 40,000 pixels sampled from the basemap under the land mask, simulating protanopia, deuteranopia and tritanopia. Under deuteranopia the three classes separate by only dE 10.7, and under tritanopia the overlay separates from the terrain by only dE 6.0. We compared ten alternatives, including ColorBrewer Blues, PuBu, YlGnBu, OrRd and PuRd, viridis and cividis, and adopted the Okabe-Ito palette in lightness order (#56B4E9, #E69F00, #D55E00), which gives dE 18.6 between classes and dE 22.7 against the basemap in the worst case across all three dichromacies. All flood figures and the supplement figures have been redrawn.

Two further figure problems surfaced while doing this, which we report because they are ours rather than the reviewer's to find. First, the scale bars in the submitted figures were wrong: the bar labelled 1000 m in Figs. S1-S4 spans about 1555 m, and the 200 m bar in Fig. 4 spans about 220 m. This was confirmed three independent ways, the clearest being that the manuscript's own stated total map area of 74.3 km² matches the raster extent at 1.0524 m per cell to 0.1 %. All bars have been recomputed. Second, Sect. 3.1 stated a 1 m x 1 m cell size, which is inconsistent with the same area calibration; it now states 1.05 m.

All four flood-figure captions have been rewritten to match the redrawn figures, naming the colour classes explicitly, and Fig. 3 has been rebuilt as well: the previous version carried the Maribus attribution inside the plotted image, while the new one shows the full AR5 Table 13.8 model spread with the upper bounds we used marked on it, and adds the AR6 projections for the Los Angeles gauge at 2100 and 2150. The inline figures in the manuscript have been replaced, not only the separate figure files supplied for typesetting. Figure 7 needed the same treatment and turned up two more problems. Its three Los Santos panels have been re-rendered with correct scale bars; the Los Angeles panels are unchanged, since those are CoSMoS output for the real city and are correctly georeferenced. In doing so we found that the caption described panel (f) as 'RCP 8.5 + storm surge, 3.59 m (2300)' while the panel's own legend read 2.61 m, and 3.59 m is in any case the value without storm surge. The text in Sect. 5.2 refers to the port remaining above water at 2.61 m, so the legend was right and the caption wrong; all three Los Santos panels are at 2.61 m and the caption now says so. We also relabelled panel (d), which the caption gave as Chumash although the district shown, and labelled in the imagery, is Pacific Bluffs. The caption now states that the shading is the modelled water surface, which includes water already present at mean sea level, and is therefore not a map of newly inundated land.

Line 331: You are speculating about an editorial compression of the landscape - you would need to speak to the devs to know this. Make it clear this is your interpretation of what has happened (I don't disagree btw).

Corrected. The sentence now reads 'which we interpret as a deliberate editorial compression of the landscape for game design purposes, although this is our reading of the map rather than something confirmed by the developers'.

Three further corrections we made ourselves

Checking our own scenario inputs against the primary source, prompted by the reviewer's call for methodological sharpness, turned up three errors that no reviewer had raised.

First, Table 1 gave 0.83 m as the RCP 2.6 value for 2300, while the combined value in the same row implied 0.85 m and the figure legends used 0.85 m. IPCC (2013), Table 13.8 gives total sea level change for the low scenario category as 0.41 to 0.85 m at 2300, and every value in our Table 1 is the upper bound of that table. The correct figure is 0.85 m; Table 1 contained a typographical error. The simulations themselves used 0.85 m, so no result changes, and we have corrected the table cell and the Fig. 4 caption.

Second, the manuscript attributed the sea level values to IPCC (2013, Fig. 13.14). That figure shows the multi-millennial sea level commitment per degree of warming, not scenario projections. The correct sources are Table 13.8 and Fig. 13.13, and Sect. 3.3 now gives the full assessed spread for both scenario categories rather than only the upper bounds we used, together with AR5's own low-confidence caveat on the ice-sheet contribution.

Third, the 2500 values of 1.02 m and 6.63 m were attributed to Maribus (2017) and described as illustrative extremes carrying very low confidence. They are in fact the AR5 upper bounds at 2500 in the same Table 13.8, which extends to that horizon. We have reattributed them to the primary source and removed the Maribus citation. This strengthens rather than weakens the 2500 results, which are AR5-assessed rather than drawn from a popular-science synthesis, although AR5's low confidence in the ice-sheet component still applies.