

Nat. Hazards Earth Syst. Sci. Discuss., author comment AC1 https://doi.org/10.5194/nhess-2021-148-AC1, 2021 © Author(s) 2021. This work is distributed under the Creative Commons Attribution 4.0 License.

Reply on RC1

Katy Burrows et al.

Author comment on "Integrating empirical models and satellite radar can improve landslide detection for emergency response" by Katy Burrows et al., Nat. Hazards Earth Syst. Sci. Discuss., https://doi.org/10.5194/nhess-2021-148-AC1, 2021

We thank the reviewer for their helpful comments on our manuscript. Our responses are below.

The Authors investigate the advantages of adding coherence-based landslide features detection in earthquake-induced landslide empirical modelling.

The paper is well written but difficult to review because this research makes use of multiple expertise including machine learning, remote sensing, and geomorphology. As an average reader, I acknowledge that some of my notes/questions are more like clarification requests than criticisms.

The topics (there are many) are interesting and quite current, in particular the use of SAR for rapid detection/mapping, and earthquake-induced landslide susceptibility modelling (probably the use of the term susceptibility here is `controversial' because these models exploit `dynamic variables'. I'm not particularly expert in this so I will not delve into). Some of the results are interesting and encouraging.

My concerns include:

Data re-sampling consequences seem to be underestimated or not fully considered: re-sampling is here used to overlap layers with different resolutions but, when data are re-sampled, at least two points/questions must be taken into account: (I) is my process still represented/sampled at the correct scale? (II) Does re-sampling introduce artifacts (e.g. aliasing) that in quantitative analysis can impact the results?

Apart from the re-sampling, the original scales/resolutions of the data are not reported but they look to be quite different. Are they all adequate to sample the different processes and to be combined in a machine learning model (no ghost relationships)?

We address this point in more detail in response to the reviewer's specific questions. It is true that by coarsening the resolution of our input products, we have lost information. However, the resolution at which we have applied our model is similar to that used by

other studies including the landslide model of Nowicki Jessee et al. (2018) that is published online by the USGS following an earthquake for use in emergency response.

The results are (quite well) numerically discussed but the geomorphological part is missing. So we know where and when there are numerical advantages in the combination of the two products, but we don't know why (e.g. geomorphological characterization of the true/false positives/negatives), this makes it difficult to evaluate the real possibility to export the framework in different contexts (e.g. Par 4.4)

We believe that the geomorphic implications of the models are beyond the scope of the paper, which is aimed at application during an emergency response. However, this is an excellent suggestion for further work on the topic, and would be appropriate to include in a follow-on paper aimed at wider evaluation of our approach across a larger global dataset of events.

There is a point that I admit, I still have problems to unravel. I feel that the input variable ICF is somehow a proxy of what the models want to find (the output variable). Furthermore, it sounds to me as if the ground truth was used twice: the first time when landslides (Y - independent variable) are exploited to label the units, and then, when the ICFs say (again, but as feature variable - X) to the model where landslides occurred. Can a regression/classification model make use of variables that somehow are what the model wants to predict/classify (no spurious correlations?)? If yes, can their importance be evaluated in the same way the importance of the other variables is evaluated? My doubt can be originated by the difficulties I'm having to interpret the output of the model as landslide mapping instead of landslide (spatial) forecasting (see my several comments later.

We have responded to the individual comments later in the text. But the key point is that the ICFs are indeed a proxy for landslide distribution (as mapped from optical imagery), just as the various other model inputs are (whether they are static such as topographic slope, or dynamic such as peak ground velocity). And just as the other model inputs have limitations and are imperfect proxies/predictors, so do the ICFs – i.e. while InSAR coherence is sensitive to landslides, it is not purely a measure of landslides, since other changes to the ground surface e.g. building damage, soil moisture changes, vegetation growth can result in the same signal as a landslide. So essentially, what we aim to predict is the landslide distribution (as mapped from optical satellite imagery), from a number of imperfect proxies. Therefore, of course we hope that the ICFs and optical maps are correlated, just as we hope that PGV and the optical maps are correlated. But neither of these are spurious correlations because they are comparisons between imperfect proxies distribution (at the optical maps are correlated. But neither of these are spurious correlations because they are comparisons between imperfect proxies and independent 'ground truth' of landslide distribution obtained from optical datasets.

It would also be helpful to evaluate how ICFs false positive and negative (random, systematic, reproducible?) impact the model.

We observed that false positives that were present in the InSAR coherence products of Burrows et al (2020) that were attributed to vegetation damage or to built-up areas are dampened in the models we present here.

New text: "We are also able to make a visual comparison between the InSAR coherence products of Burrows et al. (2020) and the 2-week models shown here in Figures 3 and 4. Burrows et al. (2020) observed false positives south of the landslides triggered by the Hokkaido earthquake, which they attribute to wind damage to vegetation associated with Typhoon Jebi. They also observed false positives in built-up areas close to the coast in

Lombok. These areas of false positives are visibly dampened in both our same-event and our global 2-week models." (Lines 405-409)

More details can be found later. I hope this helps.

Title: not so sure the title is correct: the empirical models don't detect landslides but they model the spatial probability of occurrence.

The choice of word is tricky since we are combining a method of landslide detection with a method of landslide prediction. InSAR coherence can be used in landslide detection but the empirical model is a prediction of where landslides are most likely to have happened. However, we agree that since we do not detect individual landslides, the title was misleading.

Suggested new title "Integrating empirical models and satellite radar improves landslide prediction for emergency response"

20-25: "generating information on the intensity and spatial the extent of landslides in the immediate aftermath of a large earthquake": what does intensity mean here?

By intensity, we meant the area affected by landslides, but since this is covered by "spatial extent", we will simplify to: New text "there are two options for generating information on the spatial extent of triggered landsliding in the immediate aftermath of a large earthquake" (line 22-23)

25-44: The discussion about whether to consider driving factors or not into susceptibility maps is quite hot and open. I suggest anticipating here the comment at par 2.3.2 to avoid possible misunderstandings.

In this paragraph, we are discussing methods of generating landslide information immediately after an event. Empirical models that do not incorporate causative factors do not fit into this category since they are not event specific. To make it clearer the type and application of empirical model we are talking about here, we make the following change:

New text "To generate an empirical model of triggered landslides following an earthquake, a training dataset of landslides is analysed alongside maps of 'static' factors known to influence landslide likelihood e.g. slope and landcover, as well as 'dynamic' causative factors e.g. ground shaking estimates, and a model is produced that predicts landslide likelihood based on these inputs." (line 29-32)

45: maybe helpful: https://doi.org/10.1016/j.earscirev.2021.103574

Thank you, this paper is highly relevant and has been added to the reference list. (referenced at lines 28 and 549)

60-65 'Here we aim to establish which of these three options is best': susceptibility and detection provide different types of results (LAD probability, and LAD measure), and what is best is not absolute but it depends on how they are used. I suggest explaining what 'best' means here.

As the paper is oriented towards emergency response, we have assumed that the three methods will all be used in the same way: to indicate areas which have been heavily damaged by triggered landsliding following an earthquake. To make this clearer, we make the following change:

New text: "which of these three options provides the best indication of areas strongly affected by triggered landslides after an earthquake" (lines 62-63)

70-75 "We chose to model LAD rather than individual landslide locations as both empirical models and SAR-based methods perform best at relatively coarse spatial resolutions": It does not sound very clear. I suggest rephrasing the sentence and try to explicit more the relationship between spatial resolution and the models (I guess susceptibility and LAD from coherence and not random forest). Another possibility can be to demand this concept to the discussion.

Yes, since SAR methods which use polarimetric data, or those which are based on SAR backscatter are often applied at a higher resolution to detect individual landslides, this sentence was too vague. We have rewritten it for improved clarity, and also note that the improvement with coarser resolution is supported by previous work on the ICFs. Empirical models are usually generated at the resolution of their coarsest input feature, which requires them to model landslide intensity rather than individual landslides.

"Rather than attempting to delineate individual landslides, we chose to model LAD as both empirical models and detection methods based on InSAR coherence perform well at relatively coarse spatial resolutions (within the range 0.01-1 km² (Burrows et al., 2019; Nowicki Jessee et al., 2018; Robinson et al., 2017))." (lines 70-73)

109-115: 2 questions:

1: 'a small number of landslides', small, but I guess the number must be representative from a statistical point of view and a geomorphological point of view, in the sense that it should be statistically significant and well representing the density of landslides in the area, correct?

This is correct. The number of mapped landslides can be small but the samples must be dispersed throughout the affected area in order to provide a representative sample of the affected area

New text: "Providing these landslides are dispersed throughout the study area to constitute a representative training dataset, their distribution can then be used to predict the landslide distribution across the whole affected area..." (lines 110-112)

2: `we randomly selected 250 landslides...': if I understand well, you labeled as yes, those pixels in which LAD > 0.01, and no elsewhere, correct?

Yes that is correct. New text to improve clarity: "we randomly selected 250 landslide (LAD >= 0.01) and 250 non-landslide (LAD < 0.01) pixels for use as training data." (line 113)

115-120 'we randomly undersampled': I guess you mean that you randomly selected the pixels, correct?

Yes that is correct. New text to improve clarity: "We randomly selected 1000 landslide $(LAD \ge 0.01)$ and 1000 non-landslide (LAD < 0.01) from each event." (lines 117-118)

115-120 'resulting model was therefore trained on equal numbers of pixels from the two training events': interesting point. The choice is probably dictated by the model, I'm wondering what implications this choice can bring. Don't you 'decorrelate' the strength of the earthquake (represented somehow by the ground shaking data) and the number of landslides (the consequence of the strength) (and other covariates)? Forgive me if I'm using some terms inappropriately. It is true that a larger earthquake will generally trigger more landslides, but it is not necessarily the case that the larger earthquake will have higher LAD values. For example, out of the cases we look at here, the Hokkaido earthquake is Mw 6.6, while Nepal is Mw 7.8, and indeed the Nepal earthquake triggered many more landslides across a much larger area, but Hokkaido has more cells with high LAD values than Nepal.

Furthermore, if we used more samples from large earthquakes compared to small earthquakes, the model could be biased to recreate landsliding more similar to one event than the other. This is particularly important here, since we are only using two events as training data. We therefore think it better to use equal numbers of landslides from each event.

New text: "The resulting model was therefore trained on equal numbers of pixels from the two training events, which prevents it from being dominated by the larger training event." (Lines 118-119)

120-125 `each model within the ensemble is trained on a different set of cells': do you estimate/suppose that the inventories have locally the same quality?

We do not expect all landslide inventories to be of equal quality. It is for this reason that we repeat the sampling process 100 times and take the median of the ensemble prediction. This should reduce the influence of any single training dataset that could be of poorer quality by chance (i.e. one that is not representative or is biased in some way). Using different sets of cells allows us to produce a more diverse ensemble, resulting in a more robust prediction.

Coarsening the resolution from individual landslide polygons to LAD should also help with this, since it should overcome differences in e.g. detail (the number of points used to delineate the landslide) or how precisely the landslides are located.

120-125 'This process reduces variability': in what sense?

It is necessary to take equal numbers of landslide and non-landslide pixels for every event, which requires us to undersample the data, but exactly which pixels are selected affects the model output. By 'variability' we meant that if you run the model several times, its prediction and performance changes slightly each time because it is using different sets of pixels. However if you run the ensemble of models several times, the median output does not vary as much. This was not very clear in the text, we have updated this.

New text: "This process allows the model to use more of the available training data and improves robustness" (Line 122)

125-150 'Masato...': I suggest to better contextualise: the model tuned using rainfall does not use (a polite guess...) shaking data...

Masato model uses only SAR and topography, it does not use any causative input. To make this clearer, we have altered this to

"Ohki et al (2020) demonstrated that when using fully polarimetric SAR data alongside topographic input features, it was possible..." (Lines 127-128)

(Reviewer 2 indicated that Ohki, rather than Masato was the family name of the author – we are referring to the same paper here as before)

130-135 'Instead, a high performance on at least one predicted event would be considered a success': I would be a bit more cautious, one result is not

statistically significant, I suggest to cancel the sentence and say (correctly) that would encourage further investigation.

We have altered this to "Instead a high performance on at least one predicted event would suggest that our approach is worthy of further investigation and is likely to improve when trained on a number of case-studies more comparable to those used by current "gold-standard" models" (Lines 132-134)

2.3 Input Features

Two general comments related to the fact that this is a quantitative (and not qualitative or based on interpretation) analysis:

- resampling to a higher resolution can be problematic, what type of resampling did you use?
- how did you aggregate here? In this case, have you considered aliasing problems? Is the result still sampling the variable at the right scale?

First, the resolution of these input data was finer than or similar to the resolution of our model, so we did not resample to a higher resolution. For our resampling, we use linear interpolation for continuous input features (e.g. slope) and nearest-neighbour for categoric input features (e.g. lithology). Information on resampling, aggregation and on the resolution of the original data is now included for each input feature.

It is true that perhaps small-scale features are lost as we coarsen the resolution, but there is no way to avoid this, since we need to obtain a set of input features in the same geometry that are derived from different datasets with different spatial resolutions. For topography, the number of different derived surfaces goes some way to address this. For example, if the mean slope of a cell does not represent the fact that a small area of the aggregate cell is very steep while the rest is flat, this will be picked up by the standard deviation of slopes within the aggregate cell, or by the maximum slope within the aggregate cell.

150-155 'a static proxy for soil moisture': did you find it relevant (from a geomorphological point of view) for earthquake-induced landslides? Or just from a numerical point of view?

The proxy is used by Nowicki Jessee (2018) and Tanyas et al. (2019). It is relevant because it is a metric for spatial variation in phreatic surface depth and thus pore pressure (e.g. Montgomery and Dietrich, 1994), which in turn influences material strength (and thus stability) by reducing effective normal stress on the failure surface.

New text: "the compound topographic index (CTI), a static proxy for pore water pressure, which alters the effective normal stress on the failure surface and thus the material strength." (line 151-152)

2.3.2 Ground shaking estimates

General comment: I suggest adding the scale/resolution of the map, and eventually its uncertainty.

Since this layer is also quite different from the others (sort of causing factor) I also suggest commenting (maybe not here) how this can be an adequate sampling of the measure in relation to the product you want to obtain and its resolution. Is the map enough resolute to influence/characterise the LAD? Is it compatible with the other products? The grid spacing is 0.00083 degrees (~92 m at equator). We will add this information to the text. This is a similar resolution (slightly finer) than the resolution of our model (~200 m) and so we expect that the product should be well suited to characterising LAD. The USGS ground failure products, which are generated at a very similar resolution to the models we present here, also make use of this shaking intensityproduct and find it is well-suited to landslide hazard modelling

With regards to the uncertainty of the ground shaking models, this varies depending on the local seismic network and is not currently incorporated into empirical models so we do not consider it here.

Additional text "Gridded PGV estimates are available with pixel spacing 0.00083° as part of the USGS ShakeMap product. We resampled this onto the 200 x 220 m geometry used here using a linear interpolation. The uncertainty of the ShakeMap product varies depending on the local seismic network used to gather the shaking data on which the shaking estimates are based (Wald et al. 2020) but is not currently taken into account when using these estimates in empirical models (Nowicki Jessee et al, 2018)." (Lines 165-169)

160-165 'is what distinguishes': so, is there any way to call it differently?

New text: "is what differentiates"

170-176 'Our initial susceptibility model may therefore perform better than one that uses the data made available within the first few hours of the earthquake...study': this should be part of the discussion. In any case, I still strongly recommend testing and compare the two products, since you are evaluating a progressive improvement of the performances of the system devoted to working in an emergency phase, and SAR images might be available just after the event. (I guess initial is a typo).

There are more than two products to consider here – the ShakeMap product is updated multiple times after an earthquake. We feel it would make the manuscript overly complicated to incorporate successive changes to the model that are a result of updates to the ground shaking estimates as well as the changes we can obtain by incorporating ICFs. We have rewritten this paragraph to better explain why we use only the final ShakeMap product in our models.

New text: "An initial estimate of ground shaking is generally available within hours of an earthquake from the USGS ShakeMap webpage and is then refined as more data become available. The difference that these updates to the ground shaking estimates can make to empirical models of landsliding has been explored by Allstadt et al. (2018) and Thompson et al. (2020). Here our aim was to investigate changes to modelled landsliding due to the incorporation of SAR data and so for simplicity we chose to keep all other input features constant through time. We used the final version of PGV published for each event in all our models. In reality, model performance would evolve both due to updates to the estimated PGV and the availability of SAR data. But SAR data are also used to improve PGV estimates, and the final 'best' PGV estimate is often achieved when the first post-event SAR data are incorporated. Therefore, in practice, SAR-based landslide indicators will only be added to final or close-to-final versions of these empirical models." (Lines 170-178)

2.3.2 Lithology

General comment: I suggest adding the scale/resolution of the map

The GliM is provided as vector data, with a resolution that varies depending on exact location since it was generated through combining national geological maps from a range of sources. However, these data were used by Nowicki Jessee (2018) who found them suitable for generation of an empirical model at the same scale as ours.

New text: "These data are provided in vector format. We rasterized these data onto the 20 x 22 m grid at which the SAR data were processed. Then for each 200 x 220 m aggregate pixel, we took the dominant basic lithological class..." (lines 183-185)

180-185 `One advantage of Random Forests is their ability...': I suggest moving this elsewhere

We placed this here to justify this processing step, but will shorten to

"Random forests can accept both categoric and continuous input features..." (Line 186)

2.3.5 USGS ground failure products

General comment: I suggest removing/move the first part of the paragraph. Here only the method should be described and not the reasons for which others failed. In the discussion, you can say why you did not use the same model...

We feel that the first paragraph of this section is necessary as it justifies using this product as an input feature for the global models instead of building them from the same input features as our same-event models. The paragraph has been rewritten as follows:

"For our global model, instead of the set of individual input feature maps described above, we use one single feature map that encapsulates the likely best results from a global model trained across a large number of events: the output USGS Ground Failure product of Nowicki Jessee et al. (2018). This was necessary since we found that models generated from individual input feature maps had limited skill in Hokkaido and no skill in Lombok or Nepal, which is not representative of the performance of existing global models (e.g. Nowicki Jessee et al. 2018; Tanyas et al. 2019). This poor performance is likely to be due to the limited number of case studies used as training data in our global models and would make it difficult to draw conclusions on the benefits of adding SAR to a global empirical model. We found that by replacing the individual input features with the model output of Nowicki Jessee et al. (2018) we improved model performance in Nepal and Lombok, and obtained a more consistent result across the three events." (Lines 201-209)

And again, I suggest finding another way to say that the model would have probably worked worse if you had used another product... this sounds like 'incomplete'.

New text: "USGS estimates of PGV are updated multiple times after an earthquake and the ground failure products based on these are also updated. Again, we use final published version of the ground failure product (based on the final PGV estimates) in all our models." (lines 212-215)

2.3.6 InSAR coherence features (ICFs)

230-235 'This volume of SAR data..' for sure for S1, also true for ALOS-2?

In most cases, we expect at least one ALOS-2 image to be available within two weeks of an earthquake, since the satellite is tasked with acquiring imagery following an earthquake for emergency response and has a 14 day repeat time. 240 – 245 'Burrows et al (2020) ... and 2019, and the whole paragraph: I had a look at the papers, unfortunately, I'm not sure I got correctly a point: the different methods are based on differences (generally speaking) and to decide what is the right threshold, a ROC analysis is required. A benchmark is then required, so you need to have the landslide map already prepared. I see a kind of loop. Can you please clarify? Or, If I'm wrong, can you please make explicit when 'lower' or higher' is low enough or high enough to say that the changes were caused by landslides?

In the papers, a threshold was not applied, since the tolerance for false positives and false negatives is dependent on factors such as the available resources during the emergency response phase and this was beyond the scope of the paper. Therefore, the output from these papers is a continuous surface between 0 and 1 with 1 indicating a strong likelihood of landsliding and 0 indicating the opposite. ROC analysis was used to evaluate the different continuous outputs independent of the choice of threshold, not to choose a specific threshold. It is these continuous surfaces which are used as inputs in the model. This is clarified in the new version of the document.

New text: "The output of each of these methods is a continuous surface which can be interpreted as a proxy for landslide intensity" (lines 258-259)

255 – 260 'showed some level of landslide predictive skill': not sure I would use predictive here, I suggest detection capacity or similar

Agreed, we now use "detection skill"

2.4 Random forest theory and implementation

General comment: not sure that Fig 2 and some parts of the description of RF are really helpful. My suggestion is to remove fig 2, shorten the paragraph, and eventually add some references.

We have shortened this section by removing the figure and altering the description in some places. The figure is now included in the supplementary material.

325 – 330 'by comparing the test areas of these predicted surfaces with the mapped LAD calculated in Section 2.1.': LAD obtained from the external inventories, correct?

Yes correct. New text to improve clarity "We assessed model performance by comparing the test areas of these predicted surfaces with the mapped LAD calculated from the inventories of Ferrario (2019); Roback et al. (2018) and Zhang et al. (2019) in Section 2.1" (lines 326-329)

330 - 335: I understand the need of choosing a threshold, in fact, I think the most appropriate sentence to define what is tested here is "The ROC AUC values calculated here, therefore, represent the ability of the model to identify pixels with LAD > 0.1" but, according to the fact that the percentages are so different for Hokkaido and for Lombok, I'm not so sure that the value can have the same weight in the testing phase (different densities related to different events/geosettings).

It is true that the proportion of landslides from the three optical satellite images that are represented by this 0.1 threshold is different. However, in general for emergency response it is more important to detect bigger landslides or areas of more intense landsliding, and so we choose to keep this value of 0.1 constant. The proportion of small landslides

mapped in each inventory may also be a consequence of the resolution of the optical satellite imagery or how easily visible small landslides were in optical imagery in each case, and not necessarily reflect real differences in the size distribution of landslides. In contrast, there is likely to be less sampling bias between different events for larger landslides and more intensely affected areas.

340 – 345 'The second method...': this sounds to me like an indirect evaluation because in between the real numbers and the results of the RF there is a further model to obtain the interpolation. If correct, fine for me but I suggest to better motivate the choice.

High R2 values do not necessarily indicate a good model LAD, but low R2 values, such as those seen in this paper mean that the level of noise in the model is too high to obtain a useful model. Since the R2 values are so low anyway, we don't feel it is necessary to add in the gradient or intercept of the linear regression. Even if these were perfect, the low R2 value would mean our results would not change.

New text: "A high r2 coefficient value (up to a maximum of 1.0) indicates low levels of random errors, while an r2 close to zero indicates that levels of random errors in the model are too high to obtain a good fit to observed LAD" (Lines 346-347)

3 results

3.2 Global models

General comment: I would have preferred to see the original ROCs and not only the differences.

The ROC AUC values are visible on the figure and are also included in the text. At line 397, where we had previously described an increase of 0.2 in AUC, we now give the values: "the model has a consistently higher AUC than the individual Sentinel-1 coherence surfaces (approximately 0.8 compared to 0.6)"

3.3 Do these models outperform individual InSAR coherence methods?

General comment: I'm not so sure that I got the point. Susceptibility and mapping are two different things and when you compare using the LAD benchmark you are comparing different results. In the first case the capacity of predicting spatial landslide occurrence (if this is a real susceptibility ... what is more, including a sort of ground truth obtained from the coherence-based LAD), in the second you measure the capacity of mapping using a technique.

InSAR coherence is not, in fact, a direct measurement of landslides and so cannot really be thought of as a ground truth. The ICFs used are indicative of changes to the ground surface, which include landslides but may also include changes to vegetation/ building damage/ ground surface rupture. The empirical model predicts where landslides might be expected to have happened following an earthquake of the observed characteristics. InSAR coherence methods and empirical models both provide imperfect information on the distribution of landslides immediately after an earthquake, and neither can be considered as ground-truth. We therefore think it is useful to compare these two different approaches to estimating the distribution of earthquake-triggered landslides (as measured from mapping of optical satellite imagery).

3.4 Feature importances

410 'where these become the most important feature in the model': Here there

are at least 3 different 'categories' of features: (I) those features that take a static picture of the pre-event situation (lithology, slope...) (II) the velocity, which is strongly contingent upon the event, and (III) the Insar features, that map the event. An RF feature importance evaluation is only numerical, I suggest providing a geomorphological interpretation, and considering whether the 3 classes can be really evaluated together using the same criterion.

This evaluation is only numerical, and it is limited in that there are dependencies between some of the categories (e.g. coherence is influenced by slope and land cover). However, it clearly illustrates the growing role for ICFs over time after the earthquake and that is the primary purpose of this analysis. Geomorphological interpretation is not the focus of the paper but we do take this opportunity to clarify why these different data sources can be evaluated together using the same criterion.

Additional text in Section 3.4, Lines 414-416

"As described in Section 2.4.1, these measures of importance are limited since our input features are not fully independent, but it is clear that the importance of the ICFs increases through time in each case."

4 Discussion

General comment: I in principle agree with all the topics chosen for the discussion, but I think that a real discussion on the geomorphological interpretation of the results is missing.

We believe that a geomorphological interpretation, while of interest for future studies, is beyond the scope of this study, since we are considering methods of generating information for use in an emergency response context and not for use in understanding the changes caused by triggered landslides to the landscape.

4.1 Model interpretation based on ROC AUC and r2

420 – 425 'how the current generation of empirical models at this spatial scale should be interpreted.?: any hint, personal opinion, on this?

Our opinion was given in the following sentences, we have corrected this to make it clearer.

New text: "This result directly impacts how the current generation of empirical models at this spatial scale should be interpreted: the low r2 values we have observed indicate that the ability of the models to predict LAD as a continuous variable is poor, while the more encouraging AUC values indicate that the models are well-suited to discriminating between affected and unaffected pixels." (Lines 432-435)

4.5 Current recommendations for best practice

General comment on L-C bands: probably one of the things in common among the 3 events is the presence of vegetation in the affected areas before the event. Would your suggestion change if Ic was different? See recent earthquakes in Iran that triggered landslides.

It is true that our recommendations apply only to vegetated areas. We now specify this at the beginning of this section. We also already explicitly address the case of more sparsely vegetated areas in the discussion – indeed this is an important area for further research

525 – 535: I reiterate a previous comment: can the two products be directly compared? One is mapping, the other is 'probabilistic spatial forecasting'. I agree with the suggestion about having a single product to update.

We believe the comparison is useful since both products are developed for the same application and can be assessed using ROC analysis. See previous response for changes to text.

4.6 Future possible SAR inputs

545 'is thus not usually reliable in landslide detection...' I suggest relax this sentence, actually, it is, it just needs more controls that are not highlighted in the cited studies.

Random Forests rely on establishing thresholds when splitting the data. Since landslides can cause both increases and decreases in amplitude and background amplitude is dependent on numerous factors including surface roughness, soil moisture content and particularly on slope orientation relative to sensor, we expect this method would struggle.

New text: "However, other studies have found amplitude to be unreliable in landslide detection (Czuzhlewski et al., 2003; Park et al., 2019). SAR amplitude depends on several factors including surface roughness, soil moisture content and slope orientation relative to the satellite sensor. The amplitude of non-landslide pixels is therefore likely to vary significantly across the study area. Landslides can result in both increases and decreases in radar amplitude. Since the Random Forests technique relies on establishing thresholds on which to divide data, input features based on amplitude methods may not be suitable." (Lines 553-557)

560 – 565: See my previous comment on C-L bands, I suggest mentioning these parts earlier in the recommendations for best practices.

We mention them here as a motivation for possibly using data from the planned NiSAR Lband mission in the future.

4.7 Possible applications to rainfall-triggered landslides

General comment, which can be somehow applied to the next paragraph as well: it seems to me too generic, and I'm not so sure that this topic deserves a dedicated paragraph. I suggest shortening this part.

4.8 Possible application in arid environments

see my previous comment

We will combine these two sections into a single section "application in other settings"