Anonymous Referee #2

I read with great interest the paper by Hasbini et al. and found that it provides useful and insightful links between scientific research and industrial applications, particularly in the (re-)insurance sector. Nevertheless, I have several queries, especially regarding the presentation, methodology, and motivation. More specifically, I found it difficult to fully understand the methodological approach, and I have a few concerns about the procedure used to attribute impacts to unique storms. Despite these issues, I find the content very interesting and would therefore recommend a major revision.

Since my concerns mainly relate to the methodological approach, in the following, I focus primarily on the first half of the paper. Addressing these comments could potentially lead to changes in the results or, at the very least, contribute to a better understanding of the conclusions.

We thank the referee for the review of our manuscript. We also thank the reviewer for the interesting discussion, which emphasizes the potential limitations and challenges of our analysis. The individual answers to each of the comments are in blue in the rest of the document.

1. Presentation

I found the use of language generally clear; however, I had some difficulty following the text overall. In my opinion, the introduction lacks a consistent narrative and does not make sufficiently meaningful use of technical language. I would strongly suggest revising it according to the comments below:

Line 21: Although I agree that tracking methods tend to show higher agreement on the tracks of well-developed storms, I would argue that tracking algorithms can still diverge in their results even for the most intense cyclones. Including a relatively brief review of tracking-method intercomparison studies would strengthen this point.

Indeed, the choice of the tracking method can highly condition the results (as highlighted later in the review). We will rephrase the section to refer more specifically to the studies on tracking-method intercomparison of (Flaounas et al., 2023; Neu et al., 2013).

Lines 27–29: The terms “explosive cyclogenesis” and “strong jet stream” cannot be straightforwardly characterized as physical characteristics. Please revise this phrasing. Additionally, the purpose of these phrases is unclear. In fact, the second half of this paragraph appears to consist of loosely connected statements rather than a coherent argument. I suggest restructuring it for clarity. Probably discussing cyclone dynamics is not of strong concern for this paper.

We included this section to give a dynamical background for damaging storms. Indeed, we do not specifically treat cyclone dynamics in the rest of the paper. We will clarify this section to refer more clearly to the dynamical mechanisms.

Lines 30–31: The type of distribution may indeed vary depending on the tracking method used. However, I am not sure I understand how the distribution type justifies the varying number of ETCs “observed at a given location and for a given period.” Please clarify.

We will clarify this sentence.

Lines 33–34: The statement is vague. Please be more precise.

The definition of serial clustering given by (Dacre & Pinto, 2020) is deliberately vague, as no unique definition of the clusters exist. Following this, the clustering used in this paper was defined as an abnormal frequency of storms over a restricted time period. This period was fixed to 96hours, in order to align with the length of reinsurance events. We will clarify this section to emphasize more clearly that no unique definition of cyclone clustering exists, and we chose a “user-based” definition.

Line 34: RWB typically occurs due to external forcing on the waveguide — for example, from the outflow of warm conveyor belts. I assume the authors intended to express a different idea here. Please revise, clarity or omit.

We will revise this in order to refer more precisely to the arguments raised by (Pinto et al., 2014; Priestley et al., 2017).

Lines 38–40: Dominate what? Which “dynamics” are being referred to? It seems that the term dynamics is being used as a generalization for “all the above.” Please be more specific and precise when using scientific terminology.

In this sentence, “dynamics” refers specifically to the secondary cyclogenesis. We will clarify this section to refer more specifically to the mechanisms involved. The sentence will be rephrased as follows: “Recent studies (e.g. Pinto et al., 2014; Priestley et al., 2020) have demonstrated that cyclones formed by secondary cyclogenesis are more numerous during clustered periods over Western Europe. Better understanding secondary cyclogenesis formation is therefore key for improving our knowledge of clustering mechanisms. Such analyses are beyond the scope of the present study.”

Lines 46–48: Why is a 4-day window necessary? Please provide justification for this choice (seems that this is done in the discussions section and thus it comes too late).

The choice of temporal windows depends on the research or practical objectives. From an event and impact perspective, an absolute definition with a 96-hour window, corresponding to the length of reinsurance contracts for Generali, is particularly relevant for capturing the temporal clustering of extreme events. We will add this justification earlier in the text.

Lines 49–53: This section seems more like an introductory paragraph to clustering. Consider integrating it more clearly with the broader discussion.

In this section we want to highlight that understanding cyclone clustering is of great interest not only because of the physical perspective but also from the impact perspective, as seasons with intense clustering were also association with substantial losses. We will rephrase this to make the connection with the rest of the section more direct.

Line 61: The meaning of “driving characteristics” is unclear. Please define the term more precisely.

In this section, we emphasize that characterizing storm events, solely by aggregated surface wind is too simplistic to determine whether their physical characteristics increase the likelihood of damage. Such an approach overlooks important aspects of storm evolution (e.g., secondary cyclogenesis, explosive deepening, successive events) that can strongly influence both the occurrence and severity of damage. We will clarify this point more explicitly and describe more clearly what “driving characteristics” refers to.

Lines 62–63: I agree with the general point, but in some cases, a strict one-to-one attribution of impacts to individual storms is difficult, if not impossible. For example, if a new storm develops within the frontal region of a mature storm, then attributing damages exclusively to one or the other becomes completely arbitrary.

We agree with this point. However, from an insurance and reimbursement perspective, claims still need to be grouped around specific storm event.

What we also aim to highlight in this paper is that directly attributing damage to a single storm is often complex. As you mentioned, the development of successive storms over frontal zones can further complicate attribution. One objective of the manuscript is also to understand how frequently such situations occur, particularly when it is unclear which storm was the primary driver (for instance during clustered episodes. This will be discussed in the revised manuscript.

Lines 83–85: Reading the objectives, it seems that resolution-related issues are not thoroughly discussed earlier in the text. If necessary please do so. Additionally, while you mention the limitations of aggregated wind speed data in attributing impacts to a single storm, the importance of distinguishing between storms is not sufficiently developed (see also comments on methodological approach). I overall recommend formulating the introduction to better align it with the study's stated objectives.

In fact, the resolution-related issues are more a discussion point, as we did not test results obtained with wind-speed of different resolutions. We will reformulate the first research question as follow: "How can insurance claims data be reliably linked to ETC?".
We will also highlight in the introduction and in the discussion about damage dataset, that aggregated losses (at country and department levels) make it impossible to distinguish between successive events.

We will consider the limitations you raised and rephrase part of the introduction to make it align better with the results presented in the rest of the manuscript.

2a. Storm tracks approach

My impression is that cyclone tracking is naively taken for granted considering every track as "a storm". It is quite clear from many cyclone tracking papers that the number of "storms" can be easily tuned and lead to much different results. In the case of this paper, with some very simplistic calculations and supposing a rather constant number of storms per year, then the authors seem to conclude to 100 storms per 6-month season (4439 storms for 45 years), i.e. ~16 storms per month, or more naively, one storm every other day close to France. The potential conclusion of the important role of clusters in provoking damages is thus "predecided" by the storm tracking approach. I feel that the methodological approach should be better discussed and the reader should be provided with more insights about the tracked features (see specific comments below).

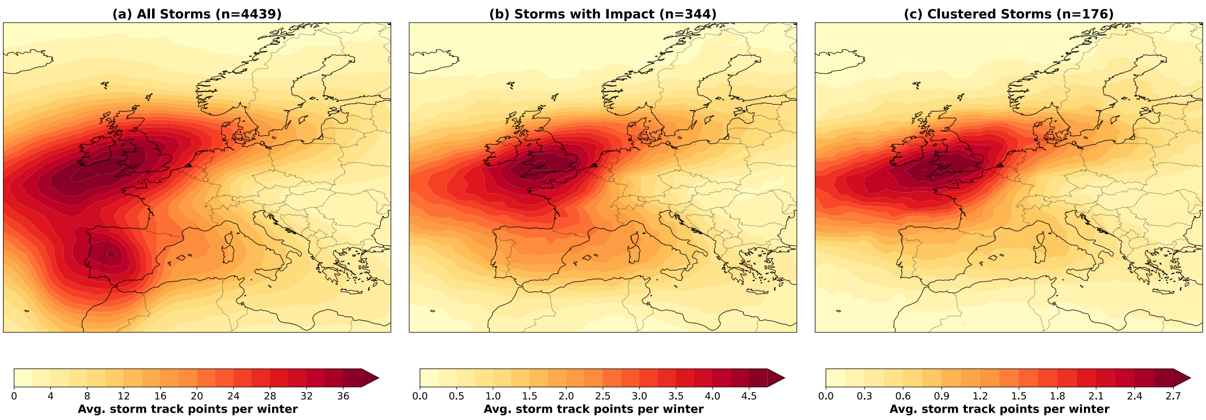

Figure 1- Distribution of storm tracks with an impact radius of 700km for the set of all storms (a), storms resulting in impact for Generali (b) and storm resulting in clustered impact (c)

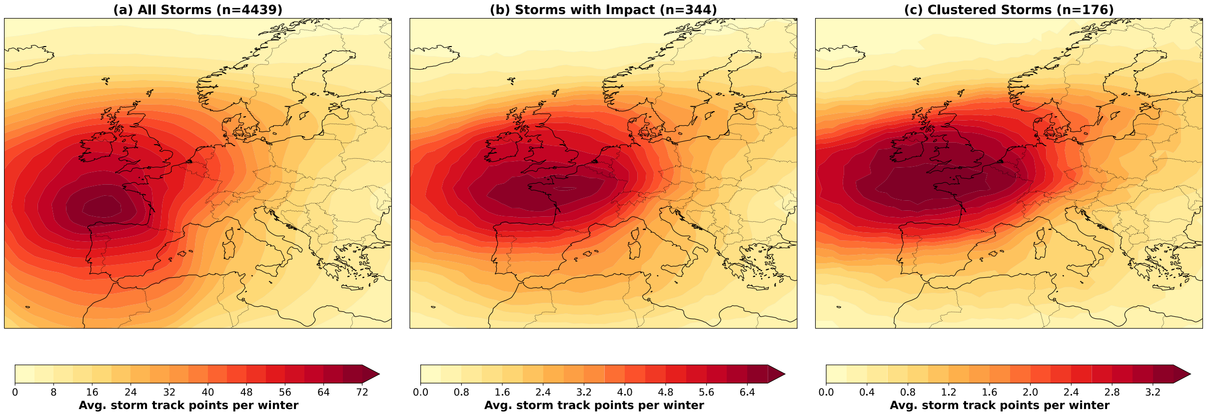

Figure 2- Same as Figure 1 with an impact radius of 1300km

Indeed, this could be interpreted as a large number but all the atmospheric depressions are, at this stage, the only potential candidates for storm impact. As mentioned later, only a few tracked depressions have led to some impact for the insurance company. Most of the storms counted in this climatology are either too far from France, or not deep enough and not associated with intense wind.

The reason for which we kept all the depressions in the initial set of tracks is because we realized that adding more constraints to the tracking algorithms (such as increasing the threshold over minimal lifetime, or decreasing the radius of 1300km) would discard some major storms with impacts. With a smaller radius, some important wind gust associated to storm Lothar were discarded. Also, some storms occurring at the beginning of January 2018, which were known to be associated with important impact (Vautard et al., 2019), were not present in the set of tracks. Therefore, we acknowledge that trial-and-error tuning on observed and documented events has to be performed to achieve a balance between the number of tracked storms and their potential impacts.
Figures 1 and 2 represent the number of storms detected with impact areas of respectively 700 and 1300km. We see in these two illustrations that the number of storms is drastically reduced after filtering over storms associated with impact for Generali. The more exposed regions are the North of France and Brittany, known to be particularly exposed to storm events, and for which we find on average 3 to 4 storms per winter, with a radius of 700km, or 6 with a radius of 1300km. The number of storms associated with impact for Generali per winter as well as the most exposed areas is consistent with the two selected radii.

Specific comments

Line 109. The storms seem to be treated as "all of the same". For instance, the same radius of influence is used (1300 km) for all tracked features, but I would argue that this is a rather unrealistic assumption. From the perspectives of a morphological approach, the impacts are expected to take place close to the center and along the fronts. So a circular area with a radius of 1300 km seems to be overwhelming. Maybe not so overwhelming for capturing the extent of the fronts but for a small storm this would be certainly the case. This radius actually compares "rather too big" with the surface of France. This seems to further favor the collocation of storms and the conclusion of clusters' high importance. I would advise the authors either to use a dynamically varying effective area for the storms, e.g. change according to intensity, or better, to detect and attribute specific wind patterns to storm track points.

Thank you for the suggestion: better characterizing the footprints could indeed improve the method. The radius was chosen at 1300km following some literature using radius of 12° (Hawcroft et al., 2012;

Sinclair & Catto, 2023). Smaller radius of 6° can also be found in the literature (Gramcianinov et al., 2020; Zappa et al., 2013) but some case studies over our set of identified storms reveals that this radius was not sufficient to capture all the wind-related impacts.

We believe automatic association of wind to frontal systems and storms tracks would be an interesting research project, but this is beyond the scope of our paper. As the aim of the proposed manuscript was more to get the physical insight from an impact dataset, we decided not to include it. Nonetheless, adding frontal detection to this approach could be a really interesting perspective as it would also give insight on the structure of fronts most likely to result in damage.

Finally, extensive literature also showed that applying spatial radius filtering over storm tracks is usually enough to "decontaminate" the footprints (Copernicus C3S, 2025; Lockwood et al., 2022).

Line 96. I am not sure what is meant by "potential noise". As far as I know, the TRACK algorithm identifies and tracks local maxima of relative vorticity. Relative vorticity is a high frequency field. Even a relatively coarse resolution of 0.25x0.25 would pose a computational challenge due to numerous identified local maxima that need to be tracked in time. So smoothing is necessary (I would actually argue that noise is not even appropriate here). How much smoothing is applied and what is the native horizontal resolution of the input field matters greatly for the eventual number of tracked features. These aspects, plus the arbitrary choice of cyclones duration (line 101), are all important to conclude to a high number of tracks. My impression is that in the trascks dataset might be included features (i.e. local maxima of vorticity) that could be rather small, or nested within greater cyclonic structures (e.g. local maxima of vorticity nested in a cold front). Therefore, several (many?) of the tracked features could be hardly considered as "distinct storms".

This is not a critical comment, because bottom line, there is no right or wrong when it comes to the number of tracks, albeit the reader needs to have a good statistical knowledge of the characteristics and nature of tracked features. For instance, it is not rare for tracking algorithms to perceive a "large scale cyclone" with 2 or 3 distinct centers as a set of that many different storm tracks. Whether there is one "big" cyclone or 2-3 "smaller" ones is an arbitrary choice and depends on the tuning of the cyclone tracking method. However, if there are large heterogeneities in the intensities of tracked cyclones, or if many of these cyclones overlap in time, then one needs to adjust their approach when attributing specific impacts to "storms" (e.g. adjusting the radius of influence). Lines 113-115 suggest that cyclone tracking approach was designed to include heterogeneous storms. This needs to be shown, discussed and linked with an accordingly adjusted attribution methodology (see also below).

We can justify our method as follow:

- Filtering for intense tracks is essential when working with Eulerian statistics or aiming to derive global measures. In contrast, for an event-based approach, the absence of the tracks (whether unique or multiple) of the event in question becomes a major limitation. While diving into such analysis, we realized that the usual storm criteria can be overly, excluding certain some high-impact storms (for example storms of January 2018).
- The potential division of one "big" cyclone into smaller subset is a particular case of interest. With the choice of the storm tracks we could expect this to happen. On top of this the association method might also associate the damage to each "member" of this bigger depression. And as you underscore, we would call a cluster this big depression which was divided into 2 to 3 smaller storms. The main concerns are then: "Can we really call a cluster of storms an event for which the several centers where part of bigger system?", "Shouldn't they be combined ?"

  We argue that the more we are able to specify and divide the loss, the better it is. Studying the type of objects encompassed in this cluster definition and whether they were subdivisions of the same global represent correspond to a different research question.

- Finally, we emphasize that storms are also filtered through the association process, which is specifically designed to retain only storms with measurable impact. This filtering reduces the number of storm events based not only on their physical characteristics (such as the maximum wind gust within their footprint) but also on the number of claims that can be linked to them.

The comments you raised underlines the complexity of t working with physical storms events. We will clarify this in the revised version of the manuscript.

Section 2.2

except if I missed it, Figure 1 is not referenced in the text. Please also provide dates in the caption or within the figure (if these are the tracks of named storms, that would be also nice to know). Also it would be useful to provide a reference in the map for 7.5 degrees West.

We will add the name of the storms are well as the 7.5W reference line (see Figure below). We will also add a reference to this figure when introducing the concept of cluster of storms.
Note that the cluster used for the illustration (Cluster 67) shows an example with 2 “storms” associated with few claims. In fact, they are respectively association with 120 and 70, however it’s clear from the location of the claims that they are coming from 2 separate events. In fact, the claims associated to the depression of the 2006-12-30 are located in the North of France, which corresponds to the area were intense wind have been recorded.

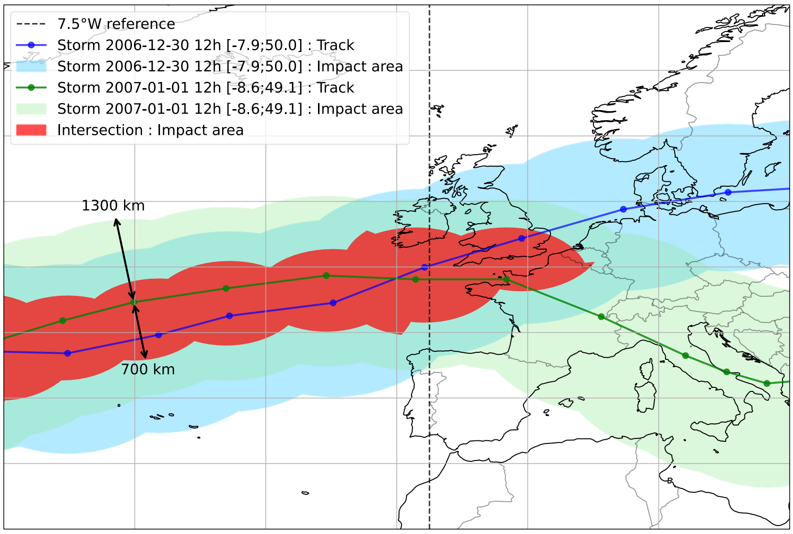

Figure 4 - Example of clustering with two storms. The green and blue lines represent the storm tracks of storms 1 and 2, with 6-hourly time increments. The green and blue shadings represent, respectively, the impact area of storms 1 and 2, defined with a radius of 1300 km around the centre of the track. The red shading illustrates the intersection of "high-impact" areas, different by a radius of 700 km around the storm tracks.

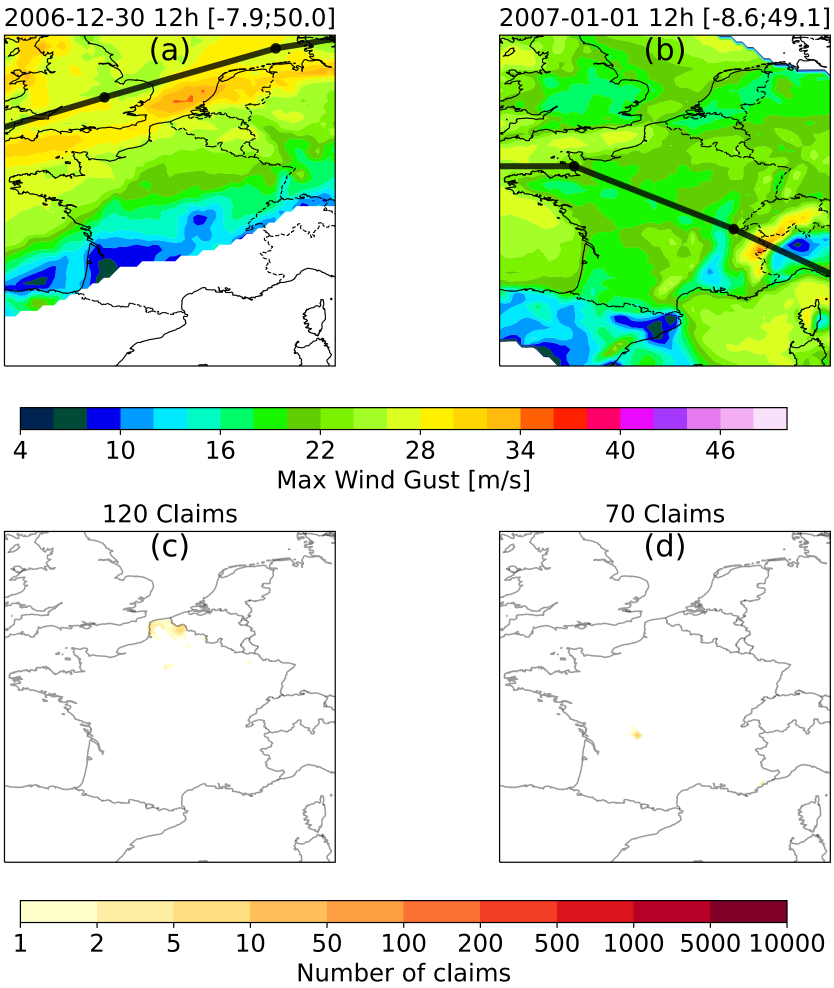

Figure 3- Association of claims to storms. Each column contains the maximum wind gust speed and the number of claims associated with a storm in Dec. 1999. Panels (a) and (c) show maps obtained for storm "2006-12-30 12h [-7.9;50.0]"and (b) and (d) for storm "2007-01-01 12h [-8.6;49.1]". Panels (a) and (b) show wind-gust footprints and the storm trajectories (thick black lines), panels (c) and(d) show the spatial distribution of the number of claims associated with each storm, while titles indicate the total number of claims for the whole event.

Lines 127-128: Distance is 420 km (assuming that 70 km refers to 1 degree..). But is 96 hours a realistic timescale for a cluster of storms hitting the same place? Please elaborate. The 96h criterion is explained in section 5 but this comes too late.

We will explain the meaning of the 96-hour criterion in the introduction, highlighting that it is intended to align with existing reinsurance policies. The remarkable winter of 1999, marked by the successive storms Lothar and Martin, illustrates the realistic possibility of multiple intense storms striking the same region within a 96-hour period.

Lines 130-134. While I appreciate the numbers provided here, could we have a visual impression of the frequency of areas affected by storm clusters in France. Also, in connection with the above, could we have a figure with statistical information about the storms included in every cluster?

We will add in the annexes of the revised manuscript, the Figure 1 of this review-answer to illustrate the area most impacted by the clustered events. We can see in this figure that the storms being clustered and resulting in an impact for Generali are similar to the set of storms associated with impact.

2b. Attribution procedure

Section 3.1 was rather difficult to understand. Probably I get the big picture but the text is rather complex, dense and in some points rather incomprehensible (at least to me). I would certainly appreciate having an illustrative example (maybe also a flow diagram?) where each step of the procedure is shown in a bit more detail. Maybe the examples in section 4 could serve this cause.

We will add the following flow-diagram to clarify the different steps of the association procedure. The diagram was done for the case study of mid-January 2025 which contain the cluster of storm Klaus.

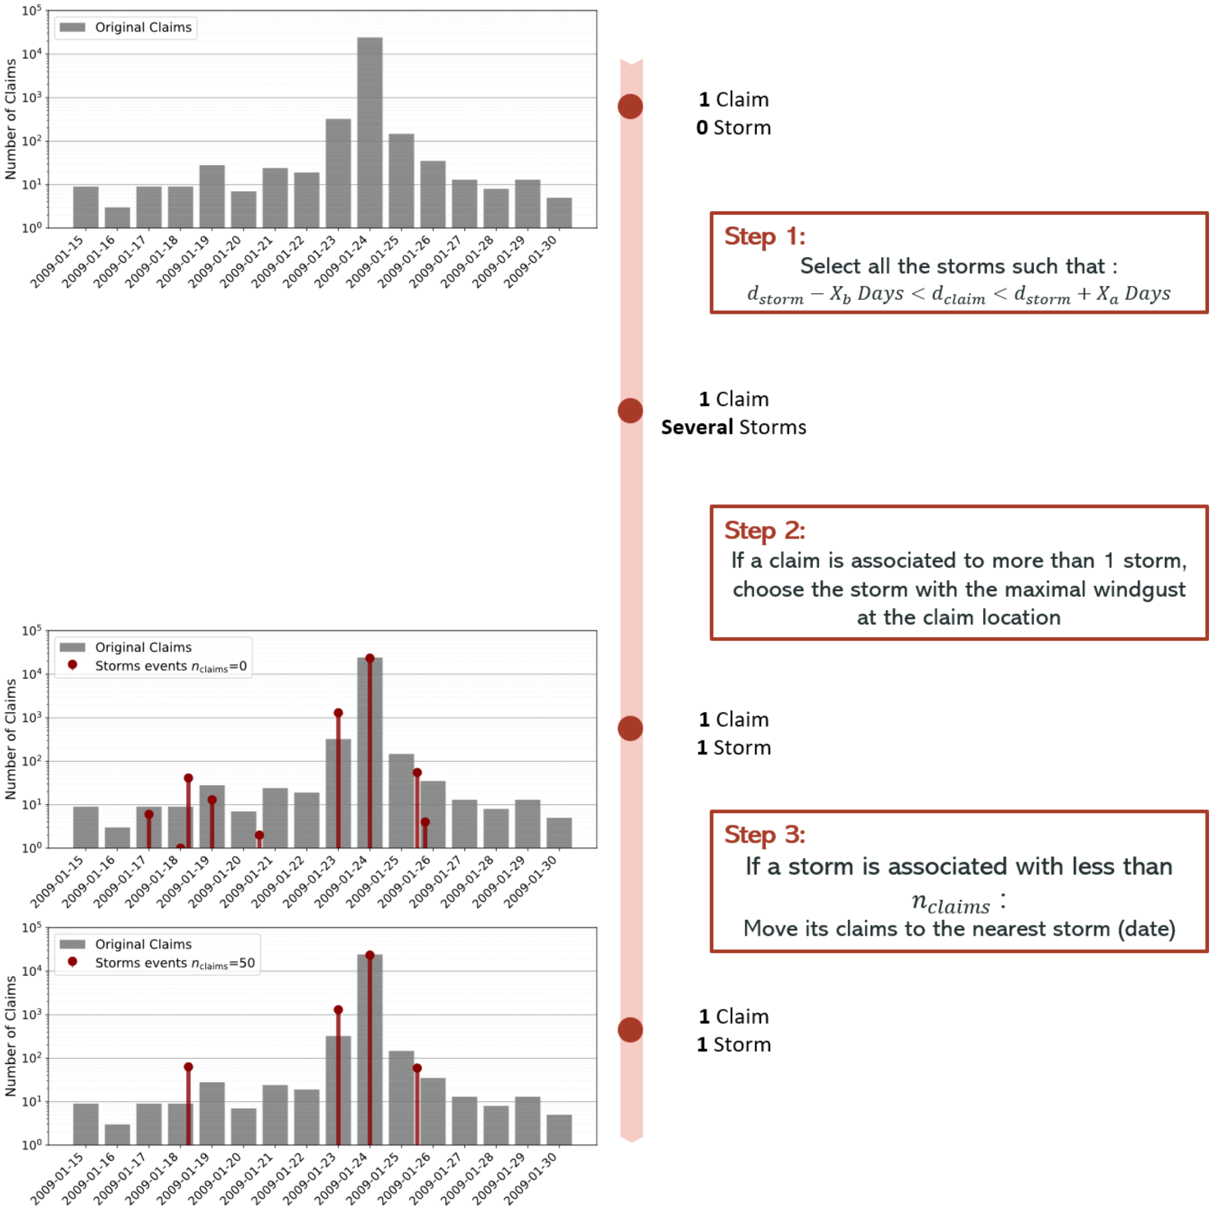

Figure 5 - Flowchart of the association method for the case study of mid-January 2009.

Lines 173-174 are quite cryptic to me. What does it mean "understanding of storm damage and Generali's exposure"? And what would it mean "few claims"? Please elaborate. How does this affect the results in this study. How representative is the exposure of Generali for the total of assets in France? Can your results be generalized about storm impacts in France? Actually, why the minimum number of claims is important here? Maybe one claim alone in an area where Generalli has a low exposure still insinuates an impactful event(?).

What we meant here is that storms are known to be large physical events leading to impact over large areas (as opposed to hail of flood hazards). When such events are detected, we can expect more than 10 houses impacted. This can be put in perspective with Generali's exposure. Generali France owes approximately 1million contracts over France, which represent 3% of the market share in France. The percentage of the market share can seem relatively small but in France, property damage insurance, is mandatory for renters, co-ownerships and for owner if they need a loan. Hence we can assume that almost 100% of properties are insured. We will clarify this in the text.

Generali does not hold a homogeneous risk over France and is more exposed in some area, such as the South-Est or North of France. Nonetheless it's highly unlikely that a storm event will result in some damage in less that 50 properties.
Additionally, for reinsurance perspectives, storms event a reasonable number of claims. While this number is arbitrary, we decided to set it to 50 as it was minimizing the best the cost function defined in section 3.2.

I actually failed to understand several concepts in section 3.1. In a rather naive approach, I would say that since you already have the coordinates of the circular area influenced by every track point and since you already have the information on a claim, then you could simply say that if the claim spatially and temporally coincides with a storm-affected area (+/- a certain amount of time), then the claim is attributed to that track point. Of course this way, the claim could be attributed to more than one tracked features. In this case, the claim could be certainly attributed to a cluster of tracks (or storms). But I do not see the feasibility of the "..ultimate goal to associate each claim with a single storm" as stated in lines 171-172. Supposing that two track points are very close to each other and correspond to two different storms, or maybe they correspond to the tracks of a "bigger" storm with two centers. Then, the high wind gusts somewhere between two track points will be due to the interaction of these two storm centers (e.g. two distinct storm centers close to each other may result in higher pressure gradients and thus higher wind speeds). There is no really a reason to say that the high gusts are due to one or the other track point. Even in such a case, one could e.g. attribute the claim to the closest track point of either storm. Given that claims and track points are pinpointed and temporally well-defined, I am not very sure I follow the complexity of the approach here, the meaningfulness of the "dstorm" variable and of the cost function. As stated before, I would appreciate illustrative examples to better understand their necessity and use.

We first stress that claims are well pinpointed but not temporally well-defined. It's visible on the added flowchart as well as on Figure 7 and 8 of the manuscript. The figures of the flowchart show that although, some claims peaks can be identified, claims are recorded almost every day. This makes the association to the correct storm complicated, in particular when the dataset of tracks contains numerous members. The 2nd figure of the flow chart underlines that if no constraint is added on the number of claims a storm should be associated to, we end up with numerous storms members, in particular, some are associated with less than 10 claims. Given the number of policies held by Generali, this is extremely unlikely. Figure 7 also underlines the incorrect dating of claims. We can see claims declared on the 27, 28 (not shown be some claims were shown until the 30) over the area impacted by storm Martin.

Additionally, as you already mentioned earlier, the windstorm dataset contains numerous tracks, some of them being only weak low-pressure systems. As you also correctly pointed out, the radius of 1300km could be too big for such small events which would lead to contaminated footprints. In such cases, the footprints of small and weak depressions could include wind intensity from other storms. The final step of the association, which will group the claims around events with at least 50 claims will, in a sense, correct this contamination. In fact, final claims will be associated with the storm with the largest impact area. Tuning the n_{claims} parameters will correct this potential contamination of storms footprint. A too small value would keep some potential overlapping of the footprint while a too large value will only identify the major events.

The implementation of the cost function ensures the robustness of the association. As explained in section 3.1, the association rely on 3 parameters. Two of them corresponds to the length of the day windows (X_a and X_b), the last one is the intensity of the grouping over the major storm events (n_{claims}). This one can be tuned to aligned the best to the observed input claims, while keeping in mind that these are a bias representation of the actual damage caused by the storms. This was first done visually over specific winter with known successive cluster events (Anatol, Lothar, Martin in December 1999; Klaus in January 2009 but also Carmen and Eleanor in January 2018 which was not shown in the article). To asses more quantitatively and automatically the quality of the association, we designed 3 performances metrics. Again, we could tune them visually to see which combination of (X_a, X_b, n_{claims}) would be the best but might not lead to robust and reproducible result. The construction of the cost function ensures the finding of the best parameters without extensive visual arbitrary decisions. Also, we showed in Appendix A that the optimal parameters obtained are not too sensible to the constructed performance metrics. I agree that the whole process is technically heavy but it was needed to ensure the reproducibility and to justify the robustness of the association method.

We will clarify all of this in the revised manuscript.

Few additional comments

In the examples of Dec 1999, Line 350 states that misattribution compromises the vulnerability curves. I am not an expert but if I understand correctly the field, the vulnerability curve links wind speed with losses. So this is indifferent of a storm attribution procedure(?). If not, could you please explain a bit more this point?

Exactly: vulnerability curves associate potential loss ratio to wind speed intensities. They are usually built using observed damage but also building vulnerability features. A correct association of wind speed to loss ratios is crucial. Let's take a claim declared on the wrong day which was associated to a storm A but it was caused a storm B. The corresponding point for the calibration of the vulnerability curve would be to link the wind speed B to the loss ratios but if the claims are not attributed to the correct storm, the loss ratio would then be linked to the wind speed A, which might not be the correct one.
This could be illustrated with the case of Lothar and Martin. Let's take a claim declared around Bordeaux. If, based only on the date, the claim is associated to storm Lothar, it would be associated to a wind of approximately 30m/s. Alternatively, as a result of the association method, the claim is more likely to be associated to storm Martin and to wind values of about 38m/s. In the construction of the vulnerability curve, forecasting a given loss ratio for 30 or 38 m/s is drastically different.

The case above is of course idealized because we never know which of the two storms caused the damage and it could also be the combination of both which lead to such losses. The study of how the association procedure change the functions of the vulnerability curves, and how clustered events, with the succession of intense storms, should be treated in such curve, is the scope of another research.

In your methodological approach, if I understood correctly, "dstorm" is defined by coordinates which are located over the Ocean (7.5W). So choosing Xb = 3 days means that the claim could take place while a storm is really far from France (I guess still within 1300 kms of radius). Could you show examples where evidently high wind speed is relevant to a storm with a far-reached center? Maybe I misunderstood this part. Could you please clarify.

In fact, it might be counter intuitive to allow the window to take some days before the storm actually crossed the line of 7.5W but this is necessary due to the inherent bias in claim dates. As illustrated in the added flowchart (Figure 5 of this review-answer), claims are declared every day. When filling a claim request, the insured person can fill any date in the past. When the person does not know the exact date of the damage, which is usually the issue with such declaration, the person is as likely to fill a date before and after the actual damage. Hence this number of days before the storm crossed the line of 7.5W should not be viewed as an anticipation of a damage but more as a correction of the mistake potentially made during the declaration of the claim date. We will clarify this in the text.

As an additional note, please avoid using the verb "land" for stating that a track point is over continental areas (or a cyclone influences such). First of all, the coordinate of 7.5 W used here is over the sea, second, I presume that the use of this terminology is inspired by the field of tropical cyclones where the usual term is "landfall". But even so, "landfall" has a special weight because the highest wind speed is really close to the cyclone center. Therefore, time and area of landfalling cyclones is directly relevant to impacts, which is not -always- the case in ETCs.

We will rephrase this to “storm impact date”. We will also clarify in the introduction that it’s impossible to give a unique date to a storm event as it is a moving object. However, for comparison purposed, we decided to define the storm impact date as the date at which it crossed to longitude of 7.5W

Section 2.2 seems more like a "methods" rather than "Data" as stated in the title of section 2.

We will revise the name of this section

Section 3.1 Lines 161-167. This paragraph seems more adequate for the introduction.

We will consider adding this paragraph to the introduction

Copernicus C3S. (2025). *Windstorm tracks and footprints derived from reanalysis over Europe between 1940 to present* [Jeu de données]. ECMWF. https://doi.org/10.24381/BF1F06A9

Dacre, H. F., & Pinto, J. G. (2020). Serial clustering of extratropical cyclones : A review of where, when and why it occurs. *Npj Climate and Atmospheric Science*, *3*(1), 48. https://doi.org/10.1038/s41612-020-00152-9

Flaounas, E., Aragão, L., Bernini, L., Dafis, S., Doiteau, B., Flocas, H., Gray, S. L., Karwat, A., Kouroutzoglou, J., Lionello, P., Miglietta, M. M., Pantillon, F., Pasquero, C., Patlakas, P., Picornell, M. Á., Porcù, F., Priestley, M. D. K., Reale, M., Roberts, M. J., … Ziv, B. (2023). A composite approach to produce reference datasets for extratropical cyclone tracks : Application to Mediterranean cyclones. *Weather and Climate Dynamics*, *4*(3), 639‑661. https://doi.org/10.5194/wcd-4-639-2023

Gramcianinov, C. B., Campos, R. M., De Camargo, R., Hodges, K. I., Guedes Soares, C., & Da Silva Dias, P. L. (2020). Analysis of Atlantic extratropical storm tracks characteristics in 41 years of ERA5 and CFSR/CFSv2 databases. *Ocean Engineering*, *216*, 108111. https://doi.org/10.1016/j.oceaneng.2020.108111

Hawcroft, M. K., Shaffrey, L. C., Hodges, K. I., & Dacre, H. F. (2012). How much Northern Hemisphere precipitation is associated with extratropical cyclones? *Geophysical Research Letters*, *39*(24), 2012GL053866. https://doi.org/10.1029/2012GL053866

Lockwood, J. F., Guentchev, G. S., Alabaster, A., Brown, S. J., Palin, E. J., Roberts, M. J., & Thornton, H. E. (2022). Using high-resolution global climate models from the PRIMAVERA project to create a European winter windstorm event set. *Natural Hazards and Earth System Sciences*, *22*(11), 3585-3606. https://doi.org/10.5194/nhess-22-3585-2022

Neu, U., Akperov, M. G., Bellenbaum, N., Benestad, R., Blender, R., Caballero, R., Cocozza, A., Dacre, H. F., Feng, Y., Fraedrich, K., Grieger, J., Gulev, S., Hanley, J., Hewson, T., Inatsu, M., Keay, K., Kew, S. F., Kindem, I., Leckebusch, G. C., … Wernli, H. (2013). IMILAST : A Community Effort to Intercompare Extratropical Cyclone Detection and Tracking Algorithms. *Bulletin of the American Meteorological Society*, *94*(4), 529-547. https://doi.org/10.1175/BAMS-D-11-00154.1

Parfitt, R., Czaja, A., & Seo, H. (2017). A simple diagnostic for the detection of atmospheric fronts. *Geophysical Research Letters*, *44*(9), 4351-4358. https://doi.org/10.1002/2017GL073662

Pinto, J. G., Gómara, I., Masato, G., Dacre, H. F., Woollings, T., & Caballero, R. (2014). Large-scale dynamics associated with clustering of extratropical cyclones affecting Western Europe. *Journal of Geophysical Research: Atmospheres*, *119*(24). https://doi.org/10.1002/2014JD022305

Priestley, M. D. K., Dacre, H. F., Shaffrey, L. C., Schemm, S., & Pinto, J. G. (2020). The role of secondary cyclones and cyclone families for the North Atlantic storm track and clustering over western Europe. *Quarterly Journal of the Royal Meteorological Society*, *146*(728), 1184-1205. https://doi.org/10.1002/qj.3733

Priestley, M. D. K., Pinto, J. G., Dacre, H. F., & Shaffrey, L. C. (2017). Rossby wave breaking, the upper level jet, and serial clustering of extratropical cyclones in western Europe. *Geophysical Research Letters*, *44*(1), 514-521. https://doi.org/10.1002/2016GL071277

Renard, R. J., & Clarke, L. C. (1965). EXPERIMENTS IN NUMERICAL OBJECTIVE FRONTAL ANALYSIS[1]. *Monthly Weather Review*, *93*(9), 547-556. https://doi.org/10.1175/1520-0493(1965)093<0547:EINOFA>2.3.CO;2

Schemm, S., Rudeva, I., & Simmonds, I. (2015). Extratropical fronts in the lower troposphere–global perspectives obtained from two automated methods. *Quarterly Journal of the Royal Meteorological Society*, *141*(690), 1686-1698. https://doi.org/10.1002/qj.2471

Sinclair, V. A., & Catto, J. L. (2023). The relationship between extra-tropical cyclone intensity and precipitation in idealised current and future climates. *Weather and Climate Dynamics*, *4*(3), 567-589. https://doi.org/10.5194/wcd-4-567-2023

Vautard, R., Van Oldenborgh, G. J., Otto, F. E. L., Yiou, P., De Vries, H., Van Meijgaard, E., Stepek, A., Soubeyroux, J.-M., Philip, S., Kew, S. F., Costella, C., Singh, R., & Tebaldi, C. (2019). Human influence on European winter wind storms such as those of January 2018. *Earth System Dynamics*, *10*(2), 271-286. https://doi.org/10.5194/esd-10-271-2019

Zappa, G., Shaffrey, L. C., & Hodges, K. I. (2013). The Ability of CMIP5 Models to Simulate North Atlantic Extratropical Cyclones*. *Journal of Climate*, *26*(15), 5379-5396. https://doi.org/10.1175/JCLI-D-12-00501.1