frontiers | Frontiers in Climate

TYPE Original Research
PUBLISHED 18 December 2024
DOI 10.3389/fclim.2024.1505268

Check for updates

OPEN ACCESS

EDITED BY
Michael C. Kruk,
National Oceanic and Atmospheric Administration (NOAA), United States

REVIEWED BY
Ramiro I. Saurral,
Barcelona Supercomputing Center, Spain
Pedro L. Fernández-Cabán,
Florida A&M University - Florida State University College of Engineering, United States

*CORRESPONDENCE
Irene Garcia-Marti
garciamarti@knmi.nl

RECEIVED 02 October 2024
ACCEPTED 02 December 2024
PUBLISHED 18 December 2024

CITATION
Garcia-Marti I, de Baar JHS, Noteboom JW, Sluijter R and van der Schrier G (2024) A data-driven impact-based analysis stemming from first responders reports to predict wind damage to urban trees.
Front. Clim. 6:1505268.
doi: 10.3389/fclim.2024.1505268

COPYRIGHT
© 2024 Garcia-Marti, de Baar, Noteboom, Sluijter and van der Schrier. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

A data-driven impact-based analysis stemming from first responders reports to predict wind damage to urban trees

Irene Garcia-Marti[1]*, Jouke H. S. de Baar[1], Jan Willem Noteboom[1], Rob Sluijter[2] and Gerard van der Schrier[1]

[1]Department of Observations and Data Technology, Royal Netherlands Meteorological Institute (KNMI), De Bilt, Utrecht, Netherlands, [2]Early Warning Centre, Royal Netherlands Meteorological Institute (KNMI), De Bilt, Utrecht, Netherlands

Transitioning from weather forecasts and warnings to impact-based forecast and warning services represents a paradigm shift in service delivery for many national hydrological and meteorological services (NHMS). NHMS typically excel at delivering information about hazardous weather, but are less experienced at inferring measures of risk of impact of extreme weather. Severe wind storms are high-impact weather phenomena that generally have a detrimental effect on distinct socio-economic sectors. In the Netherlands, the emergency services record locations where wind damage occurred to public or private property. In this work, we take 10 years of damage locations (2013–2023) provided by two safety regions in the Dutch province of Noord-Brabant. Each of the reports is enriched with an array of weather and environmental features, intended to describe the local conditions where wind damage was recorded. We model the wind reports using an ensemble of data-driven methods (i.e., One-Class Support Vector Machine) which are capable of learning from these hyper local conditions and predict for the rest of the study area. Results show how the ensemble of data-driven models are able to skillfully map locations where wind-induced damages are likely at spatial resolutions of 1 km and 5 km under high and low wind conditions scenarios. These results are encouraging for NHMS to strengthen national multi-hazard early warning systems by providing a new range of services at the urban scales in collaboration with external partners. As a consequence, the transition of scientific knowledge towards society would accelerate, hence helping at better protecting communities and livelihoods.

KEYWORDS

wind damage prediction, impact-based analysis, machine learning, climate action, climate adaptation

1 Introduction

The transition from weather forecasts and warnings to impact-based forecast and warning services represents a paradigm shift in service delivery for many national meteorological services (NHMS), including the Royal Netherlands Meteorological Institute (KNMI). The World Meteorological Organization (WMO) proposes a long-term 2020–2030 goal to strengthen national multi-hazard early warning systems and delivering user-oriented and fit-for-purpose services (WMO, 2021, 2014). NHMS typically excel at

delivering information about hazardous weather, but are less experienced at inferring measures of risk of impact, especially targeted at distinct socio-economic sectors. Risk is typically a result of the dynamic (and often non-linear) interactions between the hazard (e.g., the severity of a wind storm), the exposure (e.g., locations of urban trees), and vulnerability (e.g., health of individual trees) (Geiger et al., 2024; Gardiner, 2021); the last two firmly outside of the climate sciences and meteorology domain. Hence, it is pertinent that NHMS promote research lines in partnership with external agencies, who can provide the exposure and vulnerability components to devise new customized services (Uccellini and Ten Hoeve, 2019).

In the Netherlands, first responders (e.g., fire brigades, safety services, medical assistance and disaster crisis management) are organized in the so-called safety regions (in Dutch: *"veiligheidregio's"*; hereinafter: VRs). In collaboration with KNMI and Cobra Groeninzicht, a consulting company mapping urban nature and the living environment, an impact-based analysis was carried out, with the objective of modeling damage locations product of wind storms. The occurrence of wind damage to public (e.g., urban furniture, road pavement) or private (e.g., households) property triggers a cascade of events that are coordinated by the VRs. Hence, first responders are often handled together to minimize the time of public order. Windstorms can cause damage to urban trees, such as uprooting or toppling, with the subsequent risk for human life and damage to public or private property. In this context, VRs require detailed information on *where* and *when* wind damage might occur, so they can better coordinate their preparedness planning operations and mobilize resources efficiently for mitigation works (Potter et al., 2018; Taylor et al., 2018).

Extreme wind conditions generally have a detrimental effect on society and distinct economic sectors (e.g., energy, infrastructure) (Gliksman et al., 2023). Wind loading damages forests by causing uprooting or breakage (e.g., stem, branches) of trees, events that have the potential to negatively impact the livelihood of communities (e.g., human losses), disrupting communications (e.g., roads closed), or altering socio-economic activities (e.g., timber production) (Gardiner, 2021; Valta et al., 2019). The relationship between wind speed (or gust) and damage is complex and not fully understood, particularly during high-intensity events (Feuerstein et al., 2011). Wind damage estimates usually rely on post-disaster insurance data (Koks and Haer, 2020), which might not be publicly available. Researchers have modeled wind damage with sensitivity analyses (Koks and Haer, 2020), loss estimate models (Moemken et al., 2024), probability models (Suvanto et al., 2019), or data-driven methods (Hart et al., 2019; Pawlik and Harrison, 2021; Jahani and Saffariha, 2021). These works include different combinations of hazard, exposure, and vulnerability metrics, and range from small-scale regions like forests to continental scale. However, these works are not oriented at producing new user-oriented customized services, hence they remain relatively static studies.

In this project between the KNMI, the VRs and Cobra Groeninzicht, we focus our efforts at providing hourly wind damage estimates in the context of a risk assessment analysis for our study region in the Noord-Brabant province. To do so, we take ten years of wind damage reports provided by the VRs and we devise a set of hazard and exposure features that are subsequently modeled with an ensemble of machine learning methods. With this approach, we aim to predict the suitability of a grid cell location to develop wind damage and map the predictions into the geographic space. This work suggests there is untapped potential at including hyper-local reports collected by public services around the globe to enable impact-based analyses that might assist decision makers at better protecting communities and livelihoods.

2 Materials and methods

Severe windstorms hitting a region typically leave behind a trail of heavily damaged trees next to intact trees, which suggests that the wind damage phenomenon has a hyper-local nature. The physical characteristics of a tree, such as the crown volume or the branches strength, are fundamental to resist heavy wind loads, but eventually, prolonged severe wind conditions trigger tree damage (e.g., breakage, uprooting) (Gardiner, 2021). We see the set of damage reports as locations where hazard, exposure, and vulnerability non-linearly interact to exceed some threshold, ultimately triggering an instance of risk of impact. This section describes the methodology to characterize the damage reports collected by the VRs with a set of hazard and exposure metrics, that we subsequently model with an ensemble of data-driven methods to produce hourly maps of suitability of wind damage.

2.1 Building up a risk assessment dataset for storm damage locations

Risk assessment analysis typically conceptualizes risk (R) as a function of hazard (H), exposure (E), and vulnerability (V) (Reisinger et al., 2020), or $R = H \times E \times V$. This framework enables NHMS at better creating actionable impact information, by organizing new research lines and products by type and forecast range (Geiger et al., 2024). The goal of this work is to obtain models predicting storm damage, hence we focus at developing the H-E axis with a set of features (i.e., covariates) that we subsequently model with data-driven methods. Note that V metrics are not included in this work.

This subsection describes the process carried out to characterize each of the damage locations with a set of hazard and exposure metrics, summarized in Table 1. Features stem from weather and environmental datasets, in raster or vector formats, available in the context of this research. These datasets are processed using Python libraries and the resulting risk assessment dataset is subsequently modeled with data-driven methods in Section 2.2.

2.1.1 Description of the study area

The Netherlands is divided in 25 VRs, which are committed to the safety of citizens by providing services (e.g., fire brigades) protecting communities. In this project, we collaborated with two VRs, Brabant-Noord (BN) and Brabant-Zuidoost (BZO), both naturally providing services within the province of Noord-Brabant. Damage locations are thoroughly documented by the VRs, thus forming a hyper-local dataset on urban tree damage due to

TABLE 1 Features used in this work.

#	Short name	Description	Provider	Type
1	ws_loc	Wind speed at the hour of the report	ECMWF	H
2	wd_loc	Wind direction at the hour of the report	ECMWF	H
3	ws_loc_6h	Average wind speed in the 6 h prior to the report	ECMWF	H
4	wd_loc_6h	Average wind direction in the 6 h prior to the report	ECMWF	H
5	svf-25m	Proportion of visible sky at report location	KNMI	E
6	bomen-pct	Percentage of trees in pixel higher than 2.5 m	RIVM	E
7	bomen-hoogte	Height of trees in pixel	RIVM	E
8	spi1	Standard precipitation index in the past 30 days	KNMI	H
9	spi3	Standard precipitation index in the past 90 days	KNMI	H
10	spi6	Standard precipitation index in the past 180 days	KNMI	H
11	BG2017	Land use at the report's location	Cadaster Netherlands	E
12	top10nl_Hoogte	Terrain elevation	CBS	E
13	dist-top10nl_waterdeel_lijn	Distance to the closest line waterbody (e.g., canal, river)	CBS	E
14	dist-top10nl_waterdeel-vlak	Distance to the closest water surface (e.g., lake)	CBS	E
15	trees_pixel	Number of trees in each grid cell	Cobra Groeninzicht	E
16	trees_buff	Number of trees around the report location	Cobra Groeninzicht	E
17	closest_tree	Distance between the report location and the closest tree	Cobra Groeninzicht	E
18	avg_dist	Average distance between the reported location and all trees	Cobra Groeninzicht	E

Note they are classified in two groups: hazard (H) or exposure (E) metrics.

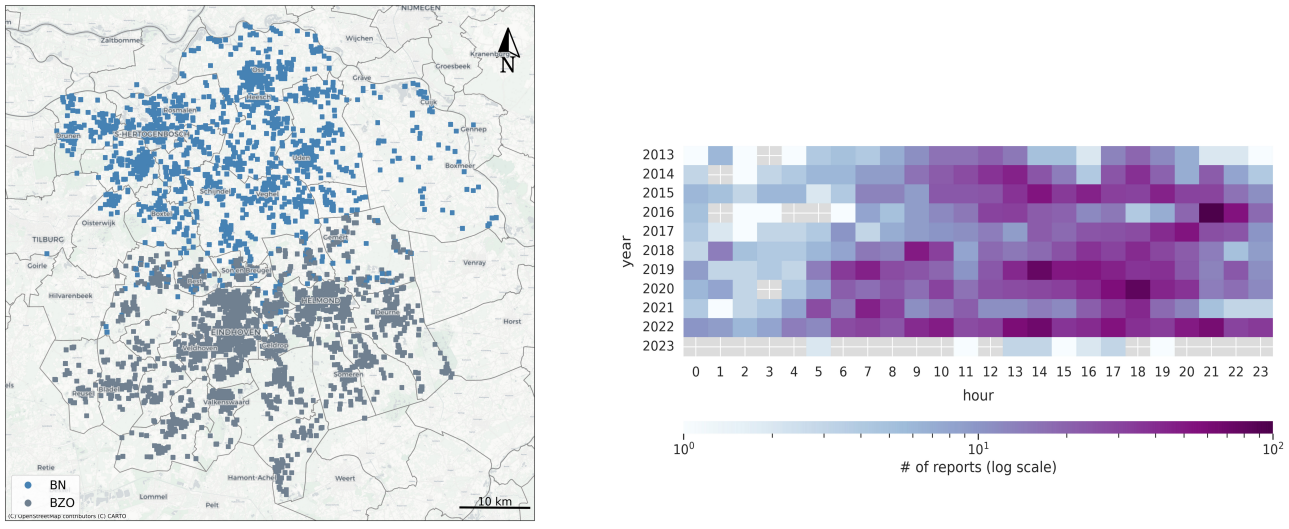

FIGURE 1
The left subplot is the spatial representation of the damage reports registered during the period 2013–2023 by the two participating VRs. Blue color is for the VR Brabant-Noord (BN), whereas gray color represents the VR Brabant-Zuidoost (BZO). The right plot shows the temporal dimension of the damage reports. The heatmap illustrates the number of reports per each hour of the day (i.e., *X*-axis) during the study period (i.e., *Y*-axis). As seen, during the small hours of the night citizens tend not to report damage to the VRs, these beginning around 5 a.m.–6 a.m.

strong winds. The dataset on tree damage contains roughly 5,000 locations in the period 2013–2023. Central to this study is the interpretation of each damage location. In this analysis, damage locations represent locations where the combination of weather and environmental characteristics triggered damage. Figure 1 (left) illustrates the study area and the dataset provided by the two participating VRs.

The study area is a region of (roughly) 80 × 80 km containing two major serviced cities (i.e., Eindhoven, in the south; 's-Hertogenbosch, in the north) along with several small cities (e.g.,

Helmond, Oss). In this work, we aim to create maps representing predicted wind damage suitability on an hourly basis, but the spatial resolution remains to be determined by the users. Hence, the study area and its ground-level properties and variables (i.e., see Table 1) will be sampled at 1 and 5 km of spatial resolution to reflect on the most suitable scale.

Regarding the damage reports, each of the symbols in the map (check Figure 1; BN in blue, BZO in gray) represents an address, where storm damage occurred, which is the basis of the analysis. However, from the original dataset, two aspects remain unknown (1) the severity of the damage; (2) the time at which the damage occurred (the metadata includes the time of reporting to the VRs). Figure 1 (right) shows the temporal dimension of the damage reports. The *X*-axis represents the hour of the day at which the damage was reported, whereas the *Y*-axis the year in the study period. As seen, the damage reporting signal remains very low during the night time and picks up in the early morning (5–6 h). It is unclear whether this is due to (a) human patterns, since during the small hours of the night citizens tend not to contact the VRs to report damage (except extreme episodes), or (b) the stability of the boundary layer, although during severe storms, the winds associated with these events might be sufficiently strong to suppress the build-up of the boundary layer by forceful turbulent mixing. These uncertainties inherent to the dataset are not a deterrent for data-driven methods, but they need to be taken into account at the time of devising the features in the H-E axis.

Note that the devised models will be tested over the study area on an hourly basis for the month of February 2022. This study month was selected because it contains the "Drielingstorm" (i.e., triplet storm) episode, in which storms Dudley, Eunice, and Franklin consecutively hit the Netherlands in a short period of time, concretely during 15–22 of February 2022.

2.1.2 Processing damage locations as a proxy for risk

In the risk assessment schema, each of the damage locations is an occurrence of risk (R) as product of particular combinations of H-E. The damage reports in their original form are characterized by metadata fields and internal codes from VRs, a timestamp and an address. Including these observations in a spatio-temporal data analysis with data-driven methods requires knowing the coordinates where damage happened. For this, we apply a process of reverse geocoding to obtain a pair of coordinates (i.e., latitude and longitude) for each damage location. Note that this process of translating a postal address into a geo-referenced location might introduce positional inaccuracies in the analysis, but it enables carrying out the envisioned analysis. Then the datasets are merged and standardized so that they contain the following columns: date & time, longitude, latitude, reporting VR, incident number, and geocoded address. After these steps, the risk is prepared to be combined with the features in Table 1.

2.1.3 Characterizing the damage locations with hazard and exposure metrics

In this work we define a set of features along the H-E axis that might be explanatory of the storm damage occurrence. In Table 1 we group these features in two categories: hazard or exposure. In this work, H metrics are intended to quantify a hydro-meteorological event that poses a level of threat to life, property or the environment, whereas E metrics provide a description of the assets exposed to such event (Intergovernmental Panel on Climate Change (IPCC), 2023). The computation of these features required gathering datasets on wind (i.e., ECMWF's ERA5 reanalysis), high-resolution tree data provided by Cobra Groeninzicht, KNMI's standard precipitation index and sky view factor datasets, and publicly available datasets on vegetation (i.e., RIVM's tree height and percentage of vegetation), topography (i.e., Dutch Cadaster's TOP10NL product) and land use (i.e., CBS's land use map). Note that the SVF is used in this work as a measure of urban morphology, since locations with a low SVF might correspond to dense urban areas, whereas locations with a high SVF to peri-urban residential areas. Also, the hyper local nature of the storm damage reports motivates including (very) high-resolution datasets, such as the location of individual trees, which help at understanding local conditions triggering damage. However, in this first stage of the research, we include wind data from ERA5-Land. The reason for this choice is that ERA5-Land provides average hourly estimates of 10 m wind components and grid cells present a horizontal resolution of 0.1°, roughly 17 km. In addition, the devised data-driven models will be tested at 1 and 5 km (check 2.1.1). Thus, all characteristics and requirements considered, ERA5-Land seemed a reasonable starting point to conduct this pilot impact-based analysis.

The selected datasets are either raster (e.g., NetCDF) or vector (e.g., Geopackage) geospatial formats. For each of the nearly 5,000 damage locations described in Section 2.1.1, we process the above-mentioned datasets to obtain an array of 18 values, corresponding to each of the features in Table 1. Note that the Type E features are all static through time, whereas the Type H features are dynamic, since they represent the potentially hazardous hydro-meteorological conditions. It is also important to remark that the wind features (i.e., Table 1, rows 1–4) are aggregated at two temporal resolutions: (a) closest time to the damage report; (b) average wind conditions of the previous 6 h. This is to accommodate the uncertainty described in Section 2.1.1, since we are unable to know the exact time at which the damage actually happened. Hence, we expect that providing the average wind conditions 6 h prior to the report might help the selected models. This procedure creates a training set for the data-driven methods.

The above procedure is then repeated at the two selected spatial resolutions and for each hour in February 2022 (check Section 2.1.1), hence producing two test datasets used in Section 2.2.3. Note that this is a computationally intensive task for a consumer-grade computer, but it is required to create maps with a spatio-temporal resolution matching the services at which the VRs operate, such as municipalities or neighborhoods.

2.2 Modeling storm damage reports with data-driven methods

Data-driven methods operate without any prior statistical assumption on the input dataset. In this case, the damage locations

conform a dataset which only contains occurrences of the "positive class", which represents the presence of a damage report. That is, we do not have a parallel dataset with locations where storm damage did not happen. This motivated using a machine learning method for one-class classification, concretely a One-Class Support Vector Machine (OCSVM) (Schölkopf et al., 2001). OCSVM is an algorithm typically used for anomaly detection. It projects the training samples into a kernel feature space in which they become separable using a maximum-margin n-dimensional plane (Manevitz and Yousef, 2002; Yang et al., 2016). By doing so the OCSVM classifier learns a boundary that wraps the positive class, while leaving outside potential outliers (Shin et al., 2005). In this way, OCSVM computes a binary function (i.e., "1" if inlier; "−1" if outlier) that delimits regions in the kernel feature space containing most of the probability density function (Schölkopf et al., 1999).

In this work, the modeling of the storm damage reports with OCSVM is carried out in three stages. First, we design four experiments intended to act as a sensitivity analysis for the inclusion or exclusion of high-resolution variables in the analysis (Figure 2). Second, we carry out a structured process of model selection to find out the optimal models fitting proportions of data. Third, we generate an ensemble of 200 OCSVM models with variable configurations that we apply to the geographic space. The remaining of this section describes in detail each of these stages.

2.2.1 Designing experiments

In this work, we are interested in exploring the effect of the high-resolution features (i.e., land use, Table 1 feature 11; high-resolution tree data, Table 1 feature 15–18) in the predictive power of OCSVM. The inclusion (or exclusion) of these features might be considered a sensitivity analysis. The idea is that if storm damage reports are indeed hyper local occurrences of a phenomenon, then the exclusion of these features should provide unrealistic predictions for wind damage. Thus, for each of the two selected spatial resolutions, 1 and 5 km, we run four experiments including or removing the high-resolution features listed in Table 1 and maintaining the rest. Figure 2 (top half) shows the four test beds designed for this work. In total, considering the sensitivity analysis and the two selected spatial resolutions, we run eight experiments that enable assessing how spatially consistent the predictions are and open the discussion on what are suitable spatial scales for the VRs.

2.2.2 Model selection process

The theoretical experiments design described in Figure 2 is followed by a model selection process. Model selection is a systematic procedure carrying out an in-depth exploration of the parameter space, simultaneously computing performance metrics. A fundamental reason to carry out this thorough model selection process is to control overfitting the damage reports. Overfitting happens when a trained model loses its generalization capabilities, hence misclassifying positive samples. In our case, overfitting is handled by controlling the complexity of the learned decision boundary.

The selected OCSVM model exposes three hyperparameters: "kernel", ν, and *gamma*. The complexity of the decision boundary is controlled by tuning the ν and *gamma* parameters. Basically, *gamma* controls how much pull over the decision boundary an individual sample has. Hence, high-values of *gamma* imply more complex decision boundaries that might not generalize well for unseen samples, hence yielding poor classification metrics. Then, ν controls the percentage of outliers in the training set. It is worth to mention that all the VR reports are real occurrences of damage. However, we do not know the severity of each intervention (e.g., broken branch vs. fallen tree). Hence, allowing the model to interpret some of the damage reports as outliers, might help at keeping the simplicity of the decision boundary, since outliers will be left out.

Therefore, two non-negative and incremental arrays are defined for ν and *gamma*, and the "kernel" parameter is set to use a radial basis function (RBF), since it has some desirable properties, such as high-generalization capabilities (Wang et al., 2004). For each pair of ν and *gamma* combinations, an OCSVM model is trained with an increasing fraction of the damage reports (i.e., characterized with the features of Table 1) and tested with the complementary fraction of data. Note that the training samples are randomly selected. This systematic process is carried out for each of the experiments in Figure 2. The metric used to evaluate the trained models is the percentage of correctly classified samples from the complementary fraction of data, which we know beforehand belong to the positive class.

2.2.3 Generating an ensemble of models to predict for the study area

The model selection process described in Section 2.2.2 illustrates the optimal range of parameters for the OCSVM models yielding a good percentage (i.e., >70%) of correctly classified samples. Rather than selecting an optimal single model that might not generalize properly, for example, due to overfitting, we create an ensemble of 200 OCSVM models. Generating the ensemble helps at mitigating the effects of overfitting. Each of the models in the ensemble is initialized with random and different parametrizations and will receive a random selection of 70% of the training samples. This double randomization is intended to generate an ensemble of diverse models with good generalization capabilities; some models will be more sensitive at identifying areas prone to receive wind damage, whereas others will have the contrary behavior. Hence, for each of the experiments in Figure 2, an ensemble with such characteristics is generated.

Each ensemble is applied to the two test datasets described at the end of Section 2.1.1 at 1 and 5 km spacing, and the four experiments in Figure 2 for the "Drielingstorm" episode (February 2022). The 200 ensemble predictions are averaged, so that for each hour and location a prediction value is produced. Averaged values range between [−1, 1], and this range is interpreted as how much the models of the ensemble agree (i.e., −1/1) or disagree (i.e., 0). Values closer to the limits of the interval imply that a large proportion of OCSVM models in the ensemble agree that a particular grid cell is suitable (i.e., close to upper limit "1") or unsuitable (i.e., close to lower limit "−1") to develop wind damage. Contrarily, averaged values around zero (e.g., [−0.25, 0.25]) mean that models disagree on whether a particular grid cell might develop wind damage. The averaged model predictions are then transformed into raster NetCDF files to ease the exploration of the wind damage suitability.

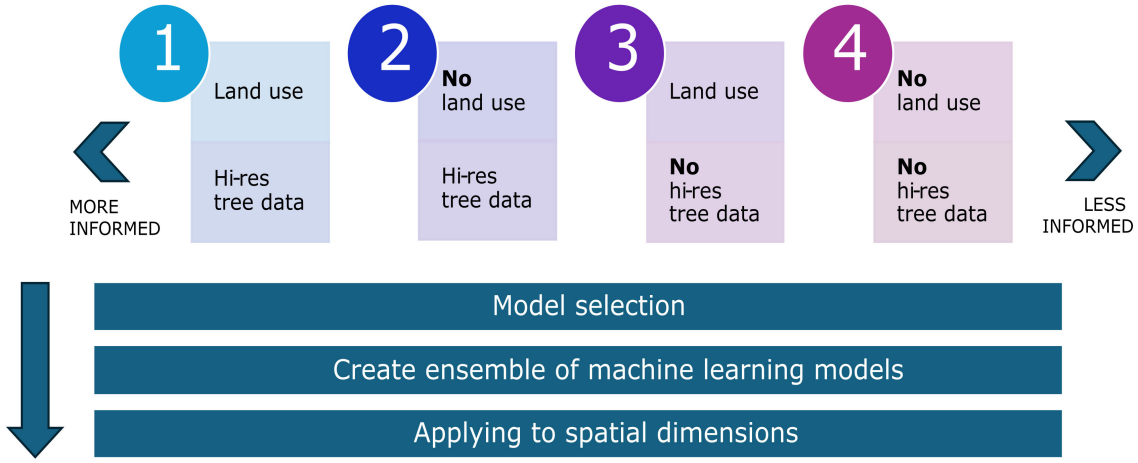

FIGURE 2
Four experiment types carried out in this work. They are designed to act as a sensitivity test to assess the effect of the inclusion (or exclusion) of high-resolution variables in Table 1 in the analysis. Note that each of these experiments is run at two spatial resolutions, yielding eight different combinations. The upper half is described in Section 2.2.1 and the bottom half in Sections 2.2.2 and 2.2.3.

3 Results

3.1 Model selection

Figure 3 shows the results of this model selection process. For simplicity, we only include the results with training data percentages $\geq$ 60%, organized per column. Hence, each of the 3D plots in the column shows the results for a fixed training data percentage and one of the experiments (i.e., with or without the high-resolution features). The X-axis of each 3D plot illustrates the values taken by the ν parameter, whereas the Y-axis shows the values taken by *gamma*. The Z-axis shows the percentage of test samples that are correctly classified as damage reports. Note the models are numerically labeled as in Figure 2, that is "1" for the most informed model; "4" for the least.

A general inspection of each of the 3D subplots shows dotted lines ranging from dark blue to yellow. Regions of the parameter space marked in dark blue represent complex decision boundaries (i.e., overfitting) with poor generalization capabilities. Regions marked in yellow represent models creating inflated decision boundaries that classify practically each test sample as "inlier". This effect is better appreciated visually navigating Figure 3 diagonally from the top left corner (i.e., more informed model with 60% of training data) to the bottom right corner (i.e., less informed model with 90% of training data). Note that for each subplot we have marked the hyperparameters of the model yielding the highest classification metric.

As seen, less-informed models, corresponding to rows "3" and "4", tend to present higher classification metrics (i.e., >90% of correctly classified samples) than the more-informed models, indicated in rows "1" and "2". Based on the percentage of correctly classified samples, the reader might be inclined to select the less-informed models. However, this might prove a sub optimal choice. OCSVM is a data-driven method that works by inflating the decision boundary to accommodate the maximum amount possible of training data. Models "3" and "4" do not contain high-resolution features, which are more complicated to fit, hence the decision boundary and its associated volume are too large. What happens in practice is that this model will mark every new unseen location as being within the boundary, which in the case of this project implies that a region will be marked as suitable for damage. Better informed models, trained with the high-resolution features, have a better view on "how damage locations look like at high-resolution", hence the decision boundary becomes more complex and constrained, which produces more classification errors.

The visual inspection of Figure 3 shows that balanced results across the four experiments are achieved when training data is 70%, the parameter ν ranges between [0.01–0.1] and *gamma* between [0.1–3]. After identifying these ranges, we can generate random values of ν and *gamma* to train an ensemble of OCSVM models.

3.2 Applying results

This section illustrates the ensemble predictions for the four experiments in two different situations: strong winds (i.e., 4 Beaufort; 5.5–7.9 m/s) during the passing of Storm Eunice over the Netherlands during the 18th of February of 2022 at 03:00 CET (Figure 4), and low wind conditions (i.e., 2 Beaufort; 1.6–43.3 m/s) on the 26th of February of 2022 at 03:00 CET (Figure 5). Note that for the date with low wind conditions, no damage reports were recorded by the VRs on that day. The color bar in both figures range between [−1, 1] and represents the average classification of the ensemble for damage to trees. Hence, values closer to "−1" (white) indicate agreement of the ensemble members at identifying a grid cell with low probability of developing wind damage, whereas values closer to "1" (purple) would identify cells more likely to develop damage. Each column corresponds to each

10.3389/fclim.2024.1505268

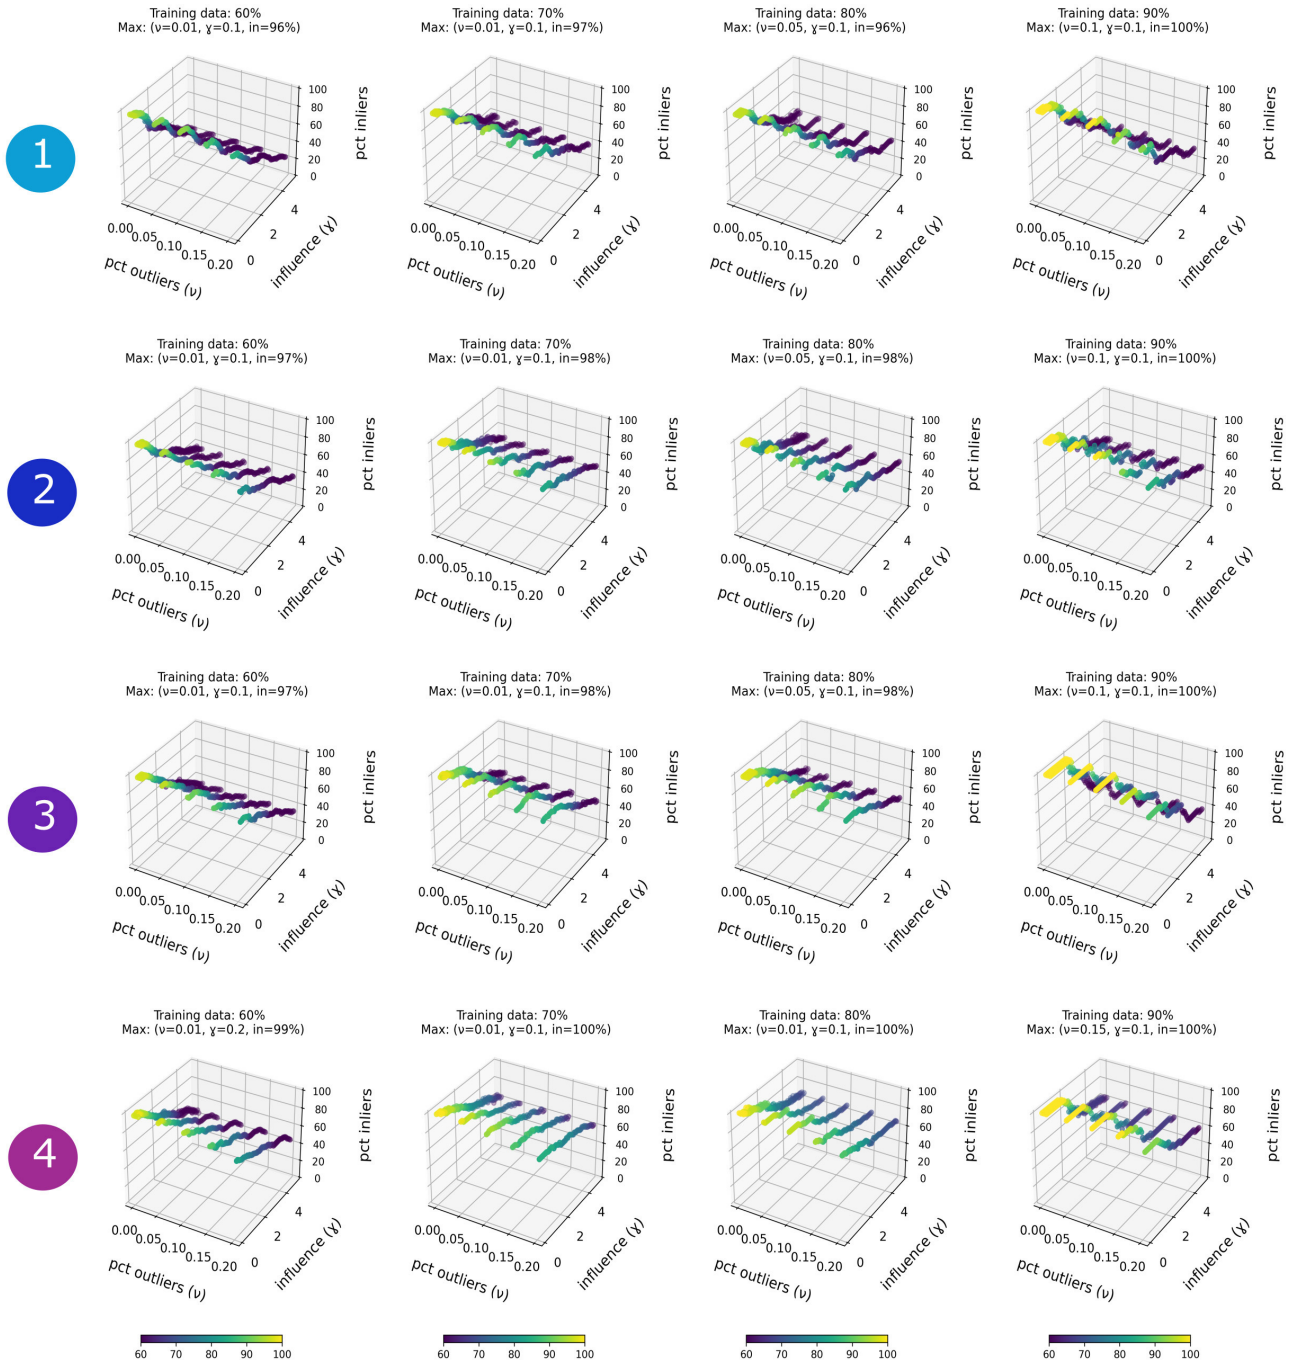

FIGURE 3
Visualization of the model selection stage carried out in this work. Each of the rows corresponds to one of the experiments described in Section 2.2.1, whereas the columns represent the percentage of randomly selected training samples used in the experiments (i.e., 10–90%). Hence, each of the subplots represent the parameter space exploration for a given experiment and training data. Subplots illustrate a 3D space in which the axes represent the model parameters (i.e., X: ν; Y: *gamma*) and the percentage of correctly classified samples (i.e., Z: inliers). To ease the interpretation of the figure, only results with $\geq$60% of the training samples are shown. High values of correctly classified samples are colored in green-to-yellow. Note that less-informed models yield high metrics, which might be misleading, as explained in Section 3.1.

of the experiments defined in Section 3.1 (i.e., labeled as 1–4) and each row the spatial resolution at which the model is applied. Figure 4 shows a stark contrast between the well-informed models (i.e., 1–2), and the less informed ones (i.e., 3–4). This relates with the complexity of the decision boundary explained in Section 3.2. Hence, the absence of high-resolution features prompts the models to yield predictions that might be too sensitive for a given moment. This can be better understood by inspecting Figure 5. In this hour, low speed conditions were monitored, yet the less informed models seem to identify most of the study area as prone to receive wind damage, which is not a realistic prediction. This seems to add critical mass to the hypothesis that damage due to severe winds

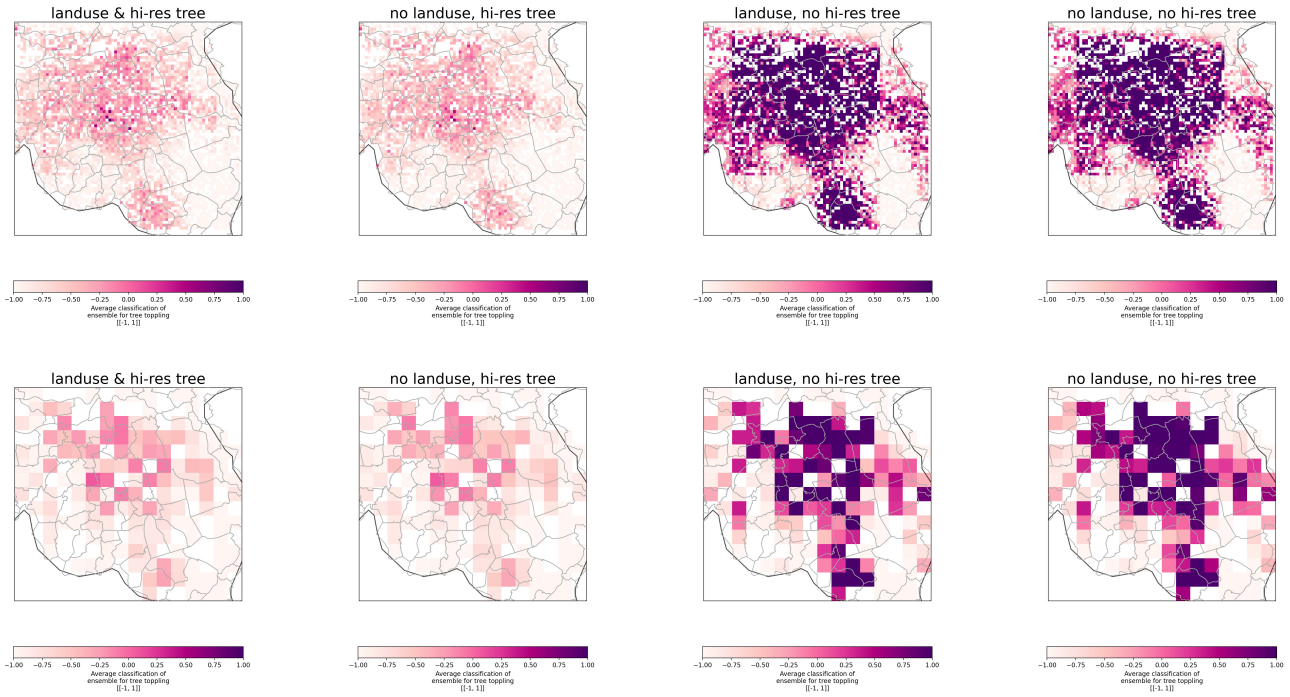

FIGURE 4
Ensemble predictions for the four experiment types during the 18th of February of 2022 at 03:00 CET, during the passing of Storm Eunice producing average winds of four Beaufort in that time slice. The color bar ranges between [−1, 1] and represents the average classification of the ensemble for damage to trees. Values closer to "−1" (white) indicate that most of the models in the ensemble coincided at identifying this location as not prone to be damaged, whereas values closer to "1" (purple) indicate the contrary case. As seen, less-informed models (i.e., columns "3" and "4") identify most of the grid cells in the study area as suitable to receive wind damage, whereas more-informed models (i.e., columns "1" and "2") are capable of narrowing down the risky grid cells to a small number. Note the stark contrast between including (or not) high-resolution features, better describing the hyper-local conditions, in the analysis on a windy hour. Also important to highlight how the cluster in the Eindhoven area (center) in the first column is visible at 1 km, but less prominent at 5 km.

is a hyperlocal phenomenon better captured with high-resolution variables.

The inspection of the maps produced at two spatial scales (i.e., 1, 5 km) for the well-informed models (i.e., 1–2) in Figure 4, shows that at 5 km the produced maps are very similar between them. However, the visual inspection of the 1 km maps reveals a cluster of purple grid cells close to the center of the map indicating that most of the models of the ensemble members are vulnerable to receive wind damage. This cluster is between the cities of Eindhoven, 's-Hertogenbosch, and Tilburg, and is visible at 1 km, but not so prominent at 5 km. This situation is repeated in Figure 5, in which the two clusters around the municipalities of Oss (north) and Gemert-Bakel (east), are less pronounced at 5 km. This suggests that the visualization of hyperlocal phenomena requires a high-resolution sampling of the study area, but this might add additional challenges to visualize and interpret the ensemble predictions.

4 Discussion

The modeling of the features along the H-E axis (Table 1) with an ensemble of calibrated OCSVM models yields hourly predictions of suitability for wind damage for the study area. This section comments on the results obtained and describes possible considerations for the future.

First, it is important to mention that the described methodology might be affected by the uncertainties inherent to the damage reports dataset. As mentioned in Section 2.1.1, each of the reports contains information about storm damage whose severity remains unknown, and also a non-quantifiable spatio-temporal inaccuracy. These limitations were considered during the development of the work, but might have some unaccounted effects in the predictive power of the ensemble.

Second, the features described in Section 2.1.2 conform a reasonable set of H and E features, that we think are sufficiently descriptive of the wind damage phenomenon. Nevertheless, the current work could benefit of including wind datasets at a higher spatial resolution than ERA5-Land. At Dutch latitudes, ERA5-Land grid cells are roughly 17 km, hence being too large for applications in small regions, such as the urban scales. In this work, the coarse resolution implies that many of the damage reports are characterized with the same wind conditions, which might hamper the learning of the OCSVM methods. For this reason, we recommend finding data sources on wind at finer

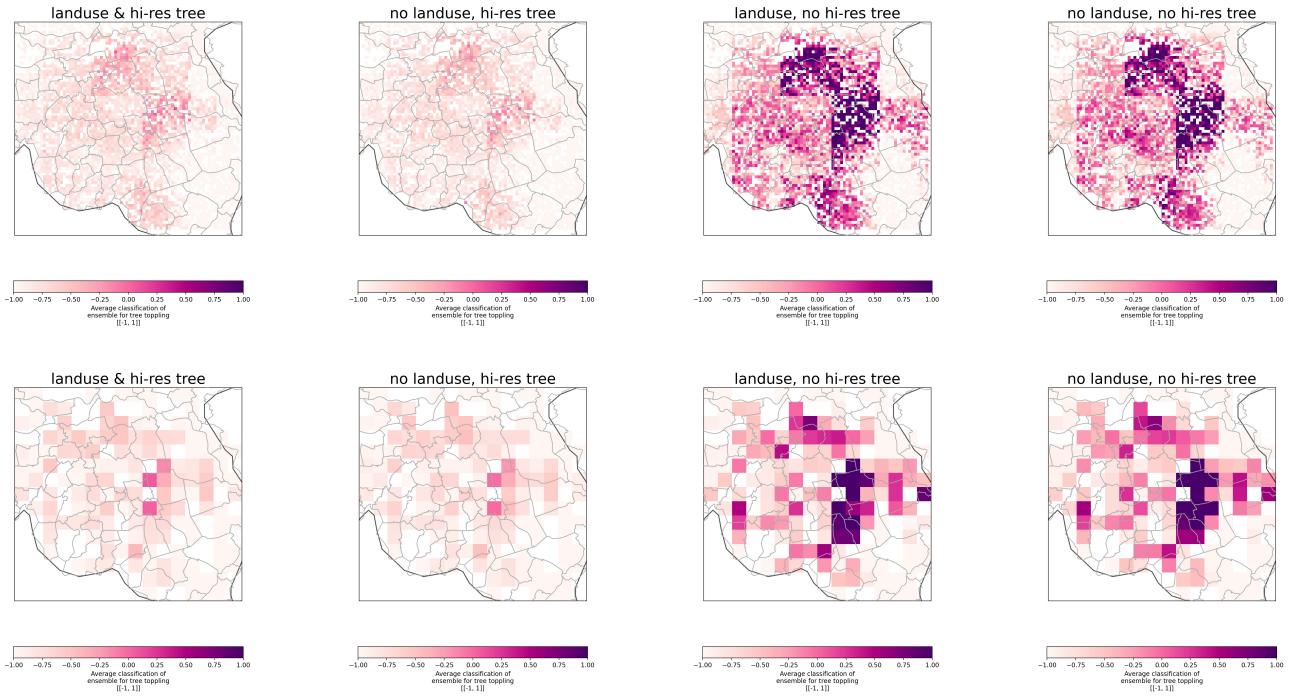

FIGURE 5
Ensemble predictions for the four experiment types during the 26th of February of 2022 at 03:00 CET, a night with low wind conditions (i.e., average winds 2 Beaufort). The color bar ranges between [−1, 1] and represents the average classification of the ensemble for damage to trees. Values closer to "−1" (white) indicate that most of the models in the ensemble coincided at identifying this location as not prone to be damaged, whereas values closer to "1" (purple) indicate the contrary case. Similarly to Figure 4, the less-informed models identify a large portion of the study area as suitable to receive wind damage, despite the non-hazardous weather conditions. The more-informed models are capable of narrowing down the risky grid cells to a small number. Note the clusters around the municipalities of Oss (north) and Gemert-Bakel (east) in the first column, are visible at 1 km but less prominent at 5 km of spatial resolution. This motivates further thinking on appropriate spatial scales for the VRs.

spatial resolutions. In addition, it might be interesting considering the inclusion of other H features, for example temperature or precipitation, or devising new E ones depicting city characteristics, such as the urban canyon effect.

Third, the approach described in Section 2.2 seems to require a relatively small amount of samples to produce hourly maps representing wind damage suitability. The models have been trained with double randomized conditions (i.e., model hyperparameters, training data) to ensure the creation of a very diverse ensemble of models. However, it remains unknown if the models have learned sufficiently from the study area to be applied to other regions in the Netherlands. Ideally, the collaboration with new VRs could provide new sources of damage reports, hence enabling more aspects of the validation process.

Fourth, the validation of the hourly layers is non-trivial. Roughly 5,000 locations are collected in a ten years period. This implies that for most of the hourly map predictions that the ensemble produces, there is not a substantial pool of damage reports to compare whether the model predicts high suitability for wind damage for the locations where there is an actual report. Future work could include comparing the hourly predictions with the stability of the boundary layer, so that it is possible to validate the low damage predictions and solve a part of the validation challenge. In addition, the reporting bias seen in Figure 1 (right), implies that for a given hourly prediction it is unclear what are the damage reports to compare with. The model selection process we carried out indicates that the OCSVM models can classify well the samples under certain hyper-parameter conditions, but with the current amount of damage reports the validation that we can carry out is limited.

Fifth, the sensitivity analysis reveals the effect of including high-resolution variables in the analysis. The inclusion of hyper local information seems to provide a detailed context for the ensemble of OCSVM, hence becoming less reactive to non-hazardous conditions. This is encouraging to pursue even higher-resolution models, reaching sub-km scales, such as neighborhoods. The current approach with the ensemble of OCSVM would be adequate to advance to such scales, but then it becomes fundamental what was mentioned above: including other H and E features at a high-resolution, particularly related to wind data.

Sixth, and related with the above point, there is an asymmetry between what is technically possible and what is required by the users. With the current approach, training for local conditions,

means that it is possible to predict for very high spatial resolutions (e.g., 100 m). However, this does not mean that the VRs require such a level of detail. In fact, our conversations with them revealed interest at getting maps aggregated at the municipality level. The current approach yields hourly maps at 1 and 5 km of spatial resolution, which is close to the municipality-level scales. Aggregating the gridded output of the model to the municipal level implies conditioning the predictions to the total area of the municipality, which might introduce new problems of interpretability. The general challenge here is to effectively bridge the gap between scientific outcomes and social agents, and this can only be achieved by sitting with users and understanding their needs.

5 Conclusion

In this work we illustrate how impact-based analyses might become fundamental for NMHS to progress toward the consolidation of early warning services. Establishing a collaboration with emergency services enables organizing the data analysis around a risk management framework. Our approach to map the risk of storm damage combines hazard (H) and exposure (E) metrics that are subsequently modeled with data-driven methods (i.e., OCSVM). The applied methodology seems to be able to recover a signal from the VRs damage reports, hence enabling predicting storm damage at high-resolution.

Our methodology contains a sensitivity analysis to study the effect of the high-resolution E metrics, followed by a thorough process of model selection. After identifying optimal parametrization ranges, we create an ensemble of OCSVM models that provide robust predictions for each grid cell in the study area. The ensemble is applied to two spatial resolutions (i.e., 1, 5 km) and its averaged predictions are mapped into the geographic space. We compare the maps for two days in February 2022 to illustrate the results of the modeling phase including the sensitivity analysis. Future work could consider the inclusion of other types of higher-resolution variables, especially wind, particularly if the scale of the analysis needs to be at the sub-km scale.

The results produced in this work are encouraging to pursue a new range of potential products and services at NHMS, at high spatio-temporal resolutions; even hyper-local in some cases. The seamless execution of impact-based analyses with external partners (e.g., disaster and crisis management agencies, safety and medical services) would accelerate the transition of scientific knowledge toward society, hence contributing to better protect communities and livelihoods. This requires that NHMS expand the current service delivery to strengthen national multi-hazard early warning systems and contributing at providing an effective response to hazardous weather.

Data availability statement

The code and data used in this work are publicly available via https://data.4tu.nl/datasets/142f8787-129c-4618-9621-05622698bd9f.

Author contributions

IG-M: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. JB: Investigation, Methodology, Validation, Writing – review & editing. JN: Funding acquisition, Project administration, Resources, Writing – review & editing. RS: Funding acquisition, Resources, Writing – review & editing. GS: Funding acquisition, Methodology, Resources, Validation, Writing – review & editing.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. This research has been funded by the Early Warning Center (EWC) programme of KNMI, which is part of KNMI's Multi-annual Strategic Plan (2019-2025).

Acknowledgments

We would like to express our gratitude to the two "veiligheidregio's" (i.e., safety regions), Brabant-Noord (BN) and Brabant-Zuidoost (BZO), participating in this research. Thanks to T. Trommelen and M. Oldenhof to be open to this initiative and providing the damage locations animating this research. Thanks to D. Otten (and the fire brigade dispatchers team) to show the scientists how an emergency service works; it was invaluable to understand the spatial and temporal scales at which first responders operate. Thanks to F. Groeneweg and K. de Zeeuw for helping with the coordination between organizations. Also, we would like to extend our gratitude to Cobra Groeninzicht, particularly to J. Verhagen, J. Verbeek, and T. Jak, for their readiness to collaborate and preparing the high-resolution tree dataset necessary for the analysis.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The authors declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Feuerstein, B., Groenemeijer, P., Dirksen, E., Hubrig, M., Holzer, A. M., and Dotzek, N. (2011). Towards an improved wind speed scale and damage description adapted for Central Europe. *Atmosph. Res*. 100, 547–564. doi: 10.1016/j.atmosres.2010.12.026

Gardiner, B. (2021). Wind damage to forests and trees: a review with an emphasis on planted and managed forests. *J. For. Res*. 26, 248–266. doi: 10.1080/13416979.2021.1940665

Geiger, T., Röösli, T., Bresch, D. N., Erhardt, B., Fischer, A. M., Imgrüth, D., et al. (2024). How to provide actionable information on weather and climate impacts?-a summary of strategic, methodological, and technical perspectives. *Front. Clim*. 6:1343993. doi: 10.3389/fclim.2024.1343993

Gliksman, D., Averbeck, P., Becker, N., Gardiner, B., Goldberg, V., Grieger, J., et al. (2023). A European perspective on wind and storm damage - from the meteorological background to index-based approaches to assess impacts. *Nat. Hazards Earth Syst. Sci*. 23, 2171–2201. doi: 10.5194/nhess-23-2171-2023

Hart, E., Sim, K., Kamimura, K., Meredieu, C., Guyon, D., and Gardiner, B. (2019). Use of machine learning techniques to model wind damage to forests. *Agric. For. Meteorol*. 265, 16–29. doi: 10.1016/j.agrformet.2018.10.022

Intergovernmental Panel on Climate Change (IPCC) (2023). *Annex II: Glossary. Climate Change 2022 - Impacts, Adaptation and Vulnerability*.

Jahani, A., and Saffariha, M. (2021). Modeling of trees failure under windstorm in harvested hyrcanian forests using machine learning techniques. *Sci. Rep*. 11, 1–13. doi: 10.1038/s41598-020-80426-7

Koks, E. E., and Haer, T. (2020). A high-resolution wind damage model for Europe. *Sci. Rep*. 10, 1–11. doi: 10.1038/s41598-020-63580-w

Manevitz, L. M., and Yousef, M. (2002). One-class SVMs for document classification. *J. Mach. Learn. Res*. 2, 139–154.

Moemken, J., Alifdini, I., Ramos, A. M., Georgiadis, A., Brocklehurst, A., Braun, L., et al. (2024). Insurance loss model vs meteorological loss index-How comparable are their loss estimates for European windstorms? *Nat. Hazards Earth Syst. Sci. Discuss*. 2024, 1–19. doi: 10.5194/nhess-2024-16

Pawlik, Ł., and Harrison, S. P. (2021). Modelling and prediction of wind damage in forest ecosystems of the Sudety Mountains, SW Poland. *Sci. Total Environ*. 815:151972. doi: 10.1016/j.scitotenv.2021.151972

Potter, S. H., Kreft, P. V., Milojev, P., Noble, C., Montz, B., Dhellemmes, A., et al. (2018). The influence of impact-based severe weather warnings on risk perceptions and intended protective actions. *Int. J. Disast. Risk Reduct*. 30, 34–43. doi: 10.1016/j.ijdrr.2018.03.031

Reisinger, A., Howden, M., Vera, C., Garschagen, M., Hurlbert, M., Kreibiehl, S., et al. (2020). *The Concept of Risk in the IPCC Sixth Assessment Report: A Summary of Cross- Working Group Discussions. Guidance for IPCC Authors*. Geneva: Intergovernmental Panel on Climate Change, 15. Technical Report September.

Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., and Williamson, R. C. (2001). Estimating the support of a high-dimensional distribution. *Neural Comput*. 13, 1443–1471. doi: 10.1162/089976601750264965

Schölkopf, B., Williamson, R. C., Smola, A., Shawe-Taylor, J., and Platt, J. (1999). "Support vector method for novelty detection," in *Advances in Neural Information Processing Systems, Vol. 12*, eds. S. Solla, T. Leen, and K. Müller (MIT Press).

Shin, H. J., Eom, D. H., and Kim, S. S. (2005). One-class support vector machines - An application in machine fault detection and classification. *Comp. Ind. Eng*. 48, 395–408. doi: 10.1016/j.cie.2005.01.009

Suvanto, S., Peltoniemi, M., Tuominen, S., Strandström, M., and Lehtonen, A. (2019). High-resolution mapping of forest vulnerability to wind for disturbance-aware forestry. *For. Ecol. Manage*. 453:117619. doi: 10.1016/j.foreco.2019.117619

Taylor, A. L., Kox, T., and Johnston, D. (2018). Communicating high impact weather: Improving warnings and decision making processes. *Int. J. Disast. Risk Reduct*. 30, 1–4. doi: 10.1016/j.ijdrr.2018.04.002

Uccellini, L. W., and Ten Hoeve, J. E. (2019). Evolving the national weather service to build a weather-ready nation. *Bull. Am. Meteorol. Soc*. 100, 1923–1942. doi: 10.1175/BAMS-D-18-0159.1

Valta, H., Lehtonen, I., Laurila, T. K., Venäläinen, A., Laapas, M., and Gregow, H. (2019). Communicating the amount of windstorm induced forest damage by the maximum wind gust speed in finland. *Adv. Sci. Res*. 16, 31–37. doi: 10.5194/asr-16-31-2019

Wang, J., Chen, Q., and Chen, Y. (2004). "RBF kernel based support vector machine with universal approximation and its application," in *Advances in Neural Networks*, eds. F.-L. Yin, J. Wang, and C. Guo (Berlin: Springer Berlin Heidelberg), 512–517.

WMO (2014). *The WMO Strategy for Service Delivery and Its Implementation Plan*.

WMO (2021). *The WMO Strategy for Service Delivery and Its Implementation Plan*.

Yang, J., Deng, T., and Sui, R. (2016). An adaptive weighted one-class SVM for robust outlier detection. *Lect. Notes Elect. Eng*. 359, 475–484. doi: 10.1007/978-3-662-48386-2_49