https://doi.org/10.5194/egusphere-2026-1625
Preprint. Discussion started: 19 May 2026

Pluvial flooding damage scenarios using remote sensing and deep learning models for urban object detection

Arianna Cauteruccio[1], Matteo Maragliano[2], Roozbeh Rajabi Toostani[3], Giorgio Boni[1], Gabriele Moser[2]

[1]Department of Civil, Chemical and Environmental Engineering (DICCA), University of Genova, Genoa, 16145, Italy
[2] Department of Electrical, Electronic, Telecommunications Engineering and Naval Architecture (DITEN), University of Genova, Genoa, 16145, Italy
[3]Geophysical Institute, University of Alaska Fairbanks, Fairbanks, USA

Correspondence to: Arianna Cauteruccio (arianna.cauteruccio@edu.unige.it)

Abstract

A quantitative assessment of the risk of pluvial flooding in an urban area is subject to various uncertainties. These include evaluating the intensity and patterns of rainfall across the area, and the exposure and specific vulnerability of different targets. In this study, the impact of rainfall pattern and intensity description is evaluated using rainfall information obtained from innovative opportunistic sensors based on satellite links and traditional rain gauges, located both within and outside the flood-prone area. Urban Flood Drifters (UFDs) are considered the primary targets, with a focus on motor vehicles (cars, vans, and motorbikes). Aerial images and deep learning–based object detection are used to evaluate the exposure of potential UFDs in terms of their numerosity and location within flood-prone urban areas. The specific vulnerability of these targets is considered using stability and damage curves from the literature, based on water depth and flow velocity. The magnitude and spatial distribution of these flood characteristics are obtained by running hydrodynamic simulations of two-dimensional pluvial flooding propagation in the urban area, based on rainfall patterns obtained from both opportunistic and traditional sensors. The impact is demonstrated in terms of vehicle safety, monetary damage experienced by flooded vehicles, and traffic disruption using an observed event. The results obtained from different rainfall data sources and aerial images with different spatial resolutions are compared and discussed in terms of their suitability for object detection and risk assessment. Flood hazard maps demonstrated that rain gauges positioned outside of the study area do not capture the actual magnitude of the event measured at the reference station. On the contrary, opportunistic sensors, despite significant underestimation of the peak rainfall, provide a good representation of the temporal evolution and total volume of the reference rainfall event. Damage scenarios demonstrated that the role of false negatives, in automated object detection, can be relevant to the safety analysis and the assessment of monetary damage, especially when spatial resolution is coarser, while it is quite negligible for the assessment of traffic disruption.

1 Introduction

The risk of pluvial flooding in cities can be quantified by assessing the vulnerability of potential targets to water depth and velocity, given a predetermined probability of occurrence. The pattern of flooding magnitude, in terms of water depth and flow velocity, over the affected area is typically derived from two-dimensional hydrodynamic modelling based on rain-on-grid simulations (see e.g., Chinchella et al., 2025). These models are highly sensitive to input rainfall information, and significant differences can be measured by rain sensors available in different locations that are only a few kilometres apart (Loglisci et al., 2024). Opportunistic sensors can be used to complement authoritative networks and improve the efficiency of the monitoring system, including earth-to-satellite microwave links (Colli et al., 2019; Giannetti and Lanza, 2023).

However, conducting vulnerability assessments in urban areas subject to pluvial flooding remains challenging due to the wide variety of potential targets in these environments. Pluvial flooding is characterized by limited water depth and flow velocity in comparison to fluvial floods, storm surges, or other types of floods. The probability of life losses in such conditions is rather low (Jonkman and Vrijling, 2008), although vulnerability may be high for goods, vehicles, economic activities and transportation. Beyond the damage that buildings and goods in ground-floor or basement premises can experience, the challenge lies in quantifying damage to mobile objects such as vehicles and street furniture. These are referred to as 'Urban Flood Drifters' (UFDs) in the literature, and only a few studies have examined their mobilisation and displacement in the event of flooding, or even their numerosity and exposure to flood hazard.

Valero et al. (2024) defined UFDs as objects that are not realized to remain stable in the event of flooding. . They investigated the onset of UFD mobilisation depending on flood characteristics for different types of loose objects, ranging from plastics to heavy vehicles. Although the authors recommended combining the stability curves resulting from their study with flood maps to evaluate the impact of UFDs during flooding, they did not use this information to quantify the vulnerability of the UFDs themselves or the overall effect of flooding on mobile assets within an urban area or during a specific event.

Remote sensing and artificial intelligence offer promising avenues for improving urban flood monitoring. The presence of sufficiently large UFDs in the urban environment can be suitably determined through the analysis of satellite or aerial images. Computer vision and deep learning techniques are used to contribute to the observation and characterization of flood-prone areas. By extracting detailed spatial information from aerial imagery, this approach supports the development of data-driven frameworks for flood risk assessment and urban resilience planning. Object detection algorithms can rapidly assess the number and position of vehicles, for example, at a specific time within areas prone to pluvial flooding.

Methods of object detection based on deep learning have made significant progress in recent years, demonstrating state-of-the-art performance in multiple areas. (Redmon et al., 2016; Terven et al., 2023). However, their direct application to aerial imagery remains challenging due to specific characteristics of aerial data — such as very small object size, dense and complex backgrounds, and significant viewpoint and scale variations caused by high-altitude imaging (Ammar et al., 2021). These factors often lead to substantially reduced detection accuracy when conventional object detectors developed for natural images are applied unchanged to aerial scenes. These challenges have been addressed by the proposal of several specialised frameworks tailored to aerial imagery. For instance, Yang et al. (2018) introduced the Double Focal Loss Convolutional Neural Network (DFL-CNN) for vehicle detection in aerial images. Their framework integrates three key components: skip-connections to fuse low-level spatial features with high-level semantic features (improving small object detection), a Region Proposal Network (RPN) employing a focal loss to mitigate class imbalance, and a classifier to refine detections. Experimental results show that DFL-CNN outperforms traditional detectors such as Faster R-CNN (Redmon et al., 2016) and classical feature-based methods on several aerial image datasets.

Similarly, Shen et al. (2022) proposed a lightweight, anchor-free, deep convolutional network for vehicle detection in aerial imagery. Instead of relying on predefined anchor boxes (which can be sub-optimal for small or arbitrarily oriented objects), their anchor-free detector directly predicts object centres and sizes. The network further incorporates two attention mechanisms — convolutional attention and triplet attention — to enhance feature representation by focusing on salient regions and modelling inter-channel and spatial dependencies. Their evaluation using performance indices like precision, recall, Sørensen-Dice coefficient – commonly known as the F_1 score – and mean average precision (mAP) confirms the effectiveness of this approach in densely packed aerial scenes.

The You Only Look Once (YOLO) model was first introduced by Redmon et al. (2016). The key concept behind YOLO is performing object detection in a single stage by processing the entire image at once. The YOLO loss function consists of three main components: localization loss for bounding box coordinates, confidence loss for predicting the presence or absence of an object, and classification loss for ensuring class accuracy (Terven et al., 2023).

A comprehensive review by Ammar et al. (2021) compared three widely adopted deep learning–based object detection frameworks—YOLOv4, YOLOv3, and Faster R-CNN—on aerial imaging datasets. Their analysis revealed that single-stage detectors, particularly those in the YOLO family, strike an effective balance between detection accuracy and computational efficiency... Thanks to their unified detection architecture, YOLO-based models enable real-time inference and scalable deployment, which makes them especially suitable for resource-constrained and time-critical applications (Redmon et al., 2016; Terven et al., 2023). While two-stage detectors such as Faster R-CNN may achieve higher accuracy in specific conditions, YOLO remains a widely adopted benchmark in object detection and classification due to its robustness, speed, and competitive performance across diverse scenarios.

YOLOv8, released in January 2023 by Ultralytics Inc., introduced several architectural and functional improvements. One major change is the introduction of the C2f module, which enhances feature representation by combining high-level semantic features with detailed spatial information, thereby improving detection accuracy. YOLOv8 employs Complete Intersection over Union (CIoU) and distributed focal loss functions for bounding box regression and uses binary cross-entropy for classification. The model is available in several variants based on parameter count, ranging from nano (3.2 millions of network parameters) to extra-large (68.2 millions of parameters). For efficient training and inference, the present study utilizes the lightweight nano variant (YOLOv8n).

YOLOv11 is the latest release in the Ultralytics YOLO series, published in September 2024 (Khanam and Hussain, 2024). The main architectural innovation in YOLOv11 is the introduction of the C3K2 block, which replaces the C2f block used in

YOLOv8. The C3K2 block offers a more computationally efficient implementation of the Cross Stage Partial (CSP) bottleneck using two smaller convolutional kernels rather than a single large one. Using the Cross Stage Partial with Spatial Attention (C2PSA) block enables the model to focus more effectively on relevant regions, thereby improving the detection of objects of different sizes and orientations. Multiple variants of YOLOv11 were released, ranging from nano (2.6 millions of parameters) to extra-large (56.9 millions of parameters). Like YOLOv8, this study employs the nano version (YOLOv11n), since it is often an effective balance between computational burden and accuracy.

In the present study, YOLOv8 and YOLOv11 models with oriented bounding box (OBB) detection capabilities were initially evaluated using pretrained weights from the DOTAv1 dataset (Dataset for Object Detection in Aerial Images). To improve the accuracy of the results, the training dataset was subsequently enhanced through manual identification and annotation of the exposed objects. DOTAv1 (Ding et al., 2021) is a large-scale benchmark for OBB detection in aerial imagery. It contains 2806 high-resolution images captured from different platforms, with object instances annotated using rotated bounding boxes. The dataset covers 15 object categories such as planes, ships, vehicles, and sports courts, featuring large variations in scale, orientation, and background. DOTAv1 is widely used to evaluate rotation-aware detectors under complex real-world aerial scenarios.

Following automated object detection, the exposure of mobile assets can be calculated by superimposing the location of each detected UFD onto flood hazard maps for various return periods and/or observed events. The number of assets at risk is an initial measure of the urban area's overall vulnerability. Assessing the specific vulnerability of UFDs — defined as the amount of damage suffered by a specific mobile asset once affected by pluvial flooding — is however more difficult. This depends on the object's specific nature, size, weight, geometry, monetary value, and susceptibility to physical damage. Traditionally, damage curves are sought for groups of targets when considering a specific natural hazard.

Depending on the resulting impact at an urban scale, displacement can be assumed to be the main cause of damage, and the number of items affected by flow conditions that exceed their mobilisation threshold can be used to estimate the expected amount of damage. Wherever possible, the cost of repairing or replacing the flooded item is estimated, for example in the case of flooded motor vehicles (cars, vans or motorbikes).

Stability (or mobilisation) and damage curves under flood conditions have been provided in the literature for various targets, including humans. Most such studies refer to water depth and flow velocity ranges that are more typical of fluvial floods than pluvial flooding events. For example, vehicular depth damage functions are formulated by FEMA (2015) as step functions, assuming flood damage to occur when critical components are immersed. A simplified approach was implemented, depending on whether the flood is below carpet, between carpet and dashboard, or above dashboard. Depth–incipient velocity relationships from theoretical analysis and experiments on large-scale model vehicles have been derived by various authors (e.g., Xia et al., 2014), with some variability observed as a function of the angle of attack of the water flow velocity and a clear dependence on the terrain slope.

However, there are few studies that focus specifically on urban areas and the characteristics associated with flooding in these areas. Martinez-Gomariz et al. (2019) developed an analysis of damages occurring to flooded vehicles using the approach formulated by FEMA (2015). Stability criteria for flooded vehicles, available in the literature, are summarized in Martinez-Gomariz et al. (2018). Kramer et al. (2016) formulated a stability criterion for the trafficability of inundated roads based on a constant specific energy (total head), which is obtained as the sum of the static head (water depth) and the kinetic head (velocity squared divided by twice the gravity acceleration). In that work, specific energy values between 0.3 and 0.4 are proposed as threshold curves for distinguishing between low- and high-risk conditions for passenger cars in the flow velocity versus water depth diagram (see their Figure 15).

Pregnolato et al. (2017) investigated the impact of flooding on the exposed roads in terms of traffic disruption. They proposed a relationship between the depth of standing water and vehicle velocity, derived from video analysis and supplemented by quantitative data extracted from existing literature. They obtained a relationship between the floodwater depth and the reduction of the vehicle speed by integrating the reviewed literature's data providing a best-fit quadratic function of the maximum acceptable velocity that ensures safe control of the vehicle given the water depth. Analogous information is provided by Hilly et al. (2018) for two-wheeled vehicles based on interviews of motorbike drivers involved in driving through flood-prone areas in Bangkok (Thailand).

In this study, we used object detection algorithms to determine the number and location of typical UFDs, focusing on motor vehicles (cars, minivans, and motorbikes) in an urban area prone to pluvial flooding. We used stability and damage curves to evaluate the vulnerability of the identified UFDs. Using different rainfall data sources covering the area, including conventional rain gauges located within and outside the flood-prone area, as well as innovative microwave sensors (satellite receivers), we performed hydrodynamic simulations to determine the flood hazard pattern in terms of local water depth and flow velocity. These simulations were performed for both a historical event and a hazard corresponding to a 10-year return period. Combining the localisation of UFDs from object detection algorithms, assumed as typical exposure scenarios for the area, with their specific vulnerability and the magnitude of flood characteristics acting on them during flood events, enabled us to assess the risk associated with the investigated UFDs when affected by pluvial flooding.

2 Materials and methods

Due to the limited size of urban catchments and the high density of impervious surfaces and buildings, cities typically experience pluvial flooding. In the present work, the input data availability issue, with specific focus on the spatial resolution of the rainfall data, is addressed by means of a comparison of pluvial flooding scenarios using rain gauges and opportunistic sensors for an urban case study. This study aims to extract actionable information from aerial imagery to enhance hydrologic models for pluvial flooding risk assessment. It focuses on detecting UFDs in urban environments using advanced computer vision techniques. These objects may encompass a diverse range of items, including two- and four-wheeled vehicles (Iqbal et al., 2022). The dataset utilized in this research originates from the Sampierdarena district in Genoa, Italy, an area historically prone to pluvial flooding events.

2.1 The case study

The case study of the present work is a densely built urban area within the Metropolitan area of Genoa (Italy), which, on September 24th, 2022, was affected by pluvial flooding associated with a rainfall event characterized by a low return period (T), between 1.5 and 3 years (see Figure 2b in Chinchella et al., 2025). The investigated urban area is in the Mediterranean region, in the western part of Genoa, named Sampierdarena district. It is characterised by a flat area measuring approximately 1 km^2 (see red area in Figure 1). The Mediterranean region is characterised by a rainfall climatology of short-duration, high-intensity events. These events typically have a limited spatial extent and evolve very rapidly, making them difficult to capture using traditional monitoring networks. Since urban storm water drainage is structured into a large number of small catchments, which are usually smaller than the typical spacing of rain gauge stations, urban catchments are often ungauged.

In the investigated area a traditional tipping-bucket rain gauge (named Arpal-FI), managed by the environmental protection agency of the Liguria region (ARPAL), is present. Two further rain gauges managed by ARPAL are considered in this study, the "Castellaccio" (Arpal-CA) and the "Centro Funzionale" (Arpal-CF), located 5.1 km and 4.5 km far from Arpal-FI, respectively.

One Smart Rainfall System (SRS), named SRS-SA, is also installed within the study area. SRS was developed at the University of Genova, as a collaboration between the Department of Electrical, Electronic, Telecommunication Engineering and Naval Architecture (DITEN) and the Department of Civil, Chemical, and Environmental Engineering (DICCA). This low-cost sensor estimates the rainfall intensity in real time by processing the attenuation of microwave satellite link signals (Federici et al., 2014).

The SRS provides an estimate of rainfall intensity averaged over the portion of the link between the antenna and the melting layer in the troposphere, using a simplified model of electromagnetic wave propagation. This model is based on the satellites that provide digital video broadcasting services. The method is based on the simplified assumption that rainfall intensity is uniform along the microwave link since the actual space distribution of the rainfall field is not known. The specific attenuation induced by rainfall is related to the rainfall intensity (RI) using a two-parameter power law equation, with the associated numerical exponent and coefficient being related to the polarization and frequency of the electromagnetic wave (Mercier et al., 2015).

Two further SRSs are available close to the investigated area: in Borzoli (named SRS-BO) and in Castelletto (named SRS-CA), positioned 2.7 km and 4.0 km far from Arpal-FI, respectively. The position of each instrument with respect to the study area is shown in Figure 1. The study area is also covered by meteorological radar, a C-band doppler polarimetric radar, installed on mount Settepani and managed by ARPAL. Radar maps have a pixel size of 1 km in both directions. This enables the real-time dynamic evolution of precipitation events to be observed. However, the ability to spatially detect and accurately assess the magnitude of the maximum rainfall intensity value in the study area is limited, as the latter has approximately the radar pixel size. Measurements from the ARPAL rain gauges and the meteorological radar are available at five-minute temporal resolution while SRSs provide measurements at one-minute resolution.

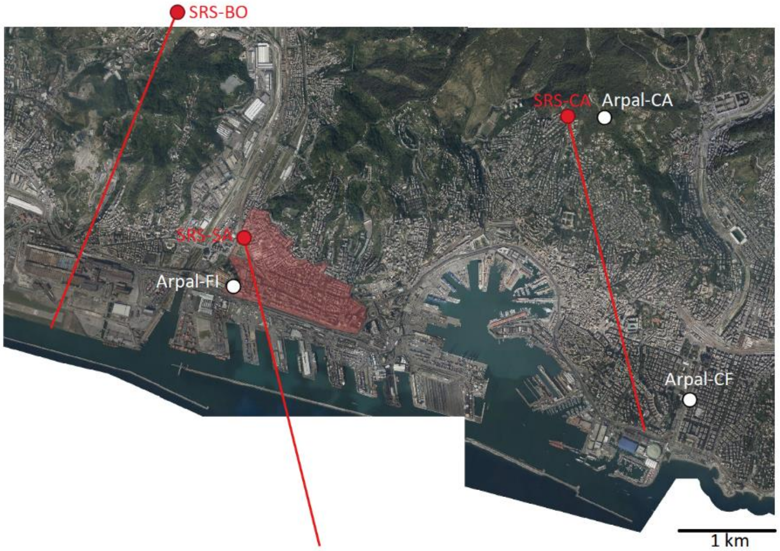

Figure 1: overview of the investigated urban area (red zone) overlaid with the aerial image of the city (source: Cartographic Repository of the Municipality of Genoa), along with the positions of the three rain gauges (white circles) and the three SRSs (red circles) with the associated satellite links (red lines).

The temporal evolution and magnitude of the sample event observed on September 24th, 2022, in the Sampierdarena district of Genoa are depicted in Figure 2 using two separate graphs. On the left-hand side, the rain depth evolution recorded by the three rain gauge stations managed by ARPAL is depicted, while the information inferred from the three available SRSs is reported in the right-hand graph. It is evident from the two graphs that the instruments located within the study area provide higher rainfall values, and that the rain gauges located outside the study area report delayed and lower rainfall.

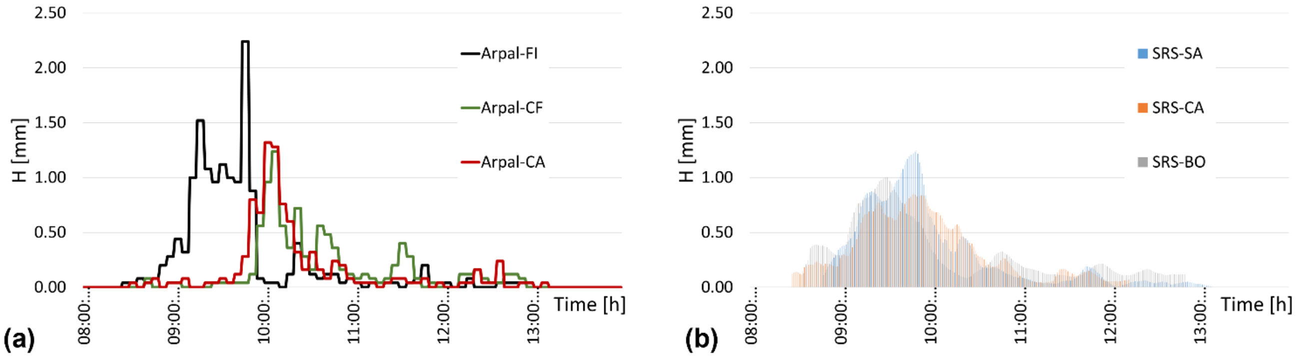

Figure 2: temporal evolution of the investigated rainfall event (September 24th, 2022) as measured by the three rain gauges (a) and the three SRSs (b).

2.2 Modelling of pluvial flooding

Pluvial flooding scenarios are selected considering the case in which a rain gauge station is available within the investigated catchment area, together with the possible alternatives in case of ungauged basins, including other contiguous rain gauges and opportunistic sensors. The SRS infers rainfall amounts or intensities by interpreting the additional attenuation caused by precipitation on the received signal level over a fixed trajectory (link) spanning a few kilometres. This may happen to improve the representativeness of the measurement as compared to point-scale rain gauges.

Recently, Chinchella et al. (2025) addressed the same case study, focusing specifically on the development of a mitigation solution. That paper simulated the overland propagation of water, using dedicated elements to reproduce the behaviour of stormwater inlets. Customised operating rules were applied to the hydraulic head and outflow rate. Simulations were performed for different return periods, revealing that the main causes of pluvial flooding in the Sampierdarena district are an insufficient number of inlets and their non-optimal distribution. Given the drainage system's inability to receive rainwater, the proposed solution was to limit the formation of stormwater runoff by implementing permeable pavements. For this reason, in this work, we neglect the role of the stormwater drainage network (including inlets), performing a purely two-dimensional simulation of the overland propagation of water.

Pluvial flooding scenarios were obtained for a return period of 10 years and for the observed event that occurred in September 2022. Hazard maps were obtained using the HEC-RAS 2D software, version 6.3.1, developed by the US Army Corps of Engineers (USACE, 2021). The propagation of stormwater volumes was modelled by solving numerically the shallow water equations (SWE-EM stricter momentum) at fixed time steps of 0.5 seconds, with a mesh size of 5 m and using a sub-grid bathymetry approach (Casulli, 2008). The rainfall input was set assuming only zenithal water using a Rain-on-Grid (RoG) approach (i.e. direct rainfall). The high-resolution digital terrain model (DTM), available at the Cartographic Repository of the Municipality of Genoa, with a resolution of 1 cm for elevation and a grid size of 1 m², was used. A vector layer reproducing buildings was implemented on a Geographic Information System (GIS) to locally raise the DTM elevation, ensuring that overland flow could not overtake them during flood propagation. This vector layer containing is available at the Cartographic Portal of the Liguria Region. During this process, the DTM was up sampled from 1 m to 0.25 m to minimise aliasing and a Gaussian filter with an influence radius of 1 m was applied to further smooth the transition between the terrain and the added buildings. The Manning coefficient was set equal to 0.018 s/m$^{1/3}$ (see Palla et al., 2018) and 0.030 s/m$^{1/3}$ for impervious ground surfaces and green areas, respectively. The boundary of the domain was set permeable. A dedicated infiltration layer with full infiltration characteristics was set for runoff volumes produced on the building roofs, as this is directly discharged into the underground drainage network. Based on the observed event duration (see Figure 2), simulations were run for six hours, with results saved at 15-minute intervals.

2.3 Deep learning-based object detection

Among UFDs, this work focuses on two- and four-wheeled vehicles, which represent the object to be identified and, hereinafter, will be named classes. These were detected using the YOLOv8n model applied to aerial images at the spatial resolutions of 5 and 20 cm. The two aerial image datasets are characterized by different daylight and traffic conditions, in particular rush traffic and off-peak hour traffic in the cases of the datasets at 5 and 20 cm, respectively. To efficiently manage the GPU memory, each original image in the dataset was partitioned into smaller patches of 1280 × 1280 pixels each.

Supervised machine learning models need well-annotated data to perform the training. This means that a subset of the objects of interest in the image patches selected for training must be annotated with their corresponding class. The annotation process can be time-consuming and requires a lot of attention to detail but is essential for creating a high-quality dataset.

The selected instances for training and validation purposes were identified in a georeferenced system and converted to the YOLO format using the LabelImg_OBB tool (Tkachenko and Team, 2020) to perform oriented bounding box annotation.

First, using the dataset at the resolution of 5 cm, four-wheeled vehicles that were completely visible (i.e. without partial occlusions due to trees or buildings) and with uniform lighting conditions (i.e. with no contrasting shaded and illuminated areas) were identified. Then, to improve the accuracy of the results, we also added vehicles with less clear visibility conditions, for both classes and for the datasets at both spatial resolutions. These two procedures of training data collection are referred to as “traditional annotation” and “detailed annotation,” respectively, in the following.

Both the YOLOv8n and YOLOv11n models were applied to quantify the associated performance and computational burden. For each object class, the performance of the model was expressed by calculating the True Positive (TP), the False Negative (FN), and the False Positive (FP) detections in the whole study area and by deriving the corresponding precision, recall, and F_1 scores. The precision index (P) measures the proportion of correctly predicted positive instances (true positives) among all predicted positives, indicating how often the detections obtained by the trained neural model are correct:

$$P = \frac{TP}{TP+FP}. \tag{1}$$

The recall index (R) quantifies the ability of the model to identify all relevant instances:

$$R = \frac{TP}{TP+FN}. \tag{2}$$

High precision implies that most detected objects correspond to actual instances in the ground truth, while high recall means that the detector successfully identifies most objects in the scene, although some may be missed. The F_1 score is used as an average measure of accuracy, for each object class, and is defined as the harmonic mean of precision and recall:

$$F_1 = \frac{2 \cdot P \cdot R}{P+R}. \tag{3}$$

2.4 Damage assessment

Damage scenarios were derived in terms of safety, monetary losses and traffic disruption, showing the capabilities and validity of the proposed approach in detecting exposed UFDs and the differences related to the characteristics of available rainfall measurements.

Using the safety criterion for the stability and roadworthiness of passenger cars reported in Kramer et al. (2016), we calculated the specific energy (E) of the flow for each four-wheeled vehicle detected in the study area and classified it as falling within the high, medium, or low class, assuming the following ranges for each class:

- Low class $E \in [0.1\text{m}, 0.3\text{m}]$
- Medium class $E \in (0.3\text{m}, 0.4\text{m}]$

https://doi.org/10.5194/egusphere-2026-1625
Preprint. Discussion started: 19 May 2026

- High class $E > 0.4$m.

Using the relationship between the water depth (h [m]) and the expected damage (D [%]) for city cars and minivans, as reported by Martínez-Gomariz et al. (2019) we derived the associated best fit curves (reported in Figure 3a) with the following formulation:

$$D\ [\%] = \frac{a \cdot h^b}{c^b + h^b} \tag{4}$$

The value of the coefficients and the associated R-squared are listed in Table 1.

The damage curve for two-wheeled vehicles was obtained in this work by considering the design of motorbikes and scooters typically used in urban areas and identifying the components of the engine and parts of the vehicle structure that are affected by flooding of increasing water depth at variable thresholds. The estimated cost of replacing or repairing the different parts was normalized with a reference total value of the vehicle obtained from commercial vendors. By fitting the calculated values, we obtain the percentage damage curve plotted in Figure 3a as a second-order polynomial in the form:

$$D\ [\%] = a \cdot h^2 + b \cdot h + c \tag{5}$$

whose coefficients, when h is expressed in meters, are listed in Table 1 with the associated R-squared.

Table 1: Coefficients of the best-fit curves used to calculate the percentage of damage of flooded vehicles for city cars, minivans, and motorcycles and to calculate the traffic disruption for four- and two-wheeled vehicles.

Vehicle model	a	b	c	R^2
City car – damage (Eq. (4))	117.7388	1.4856	0.7776	0.993
Minivan – damage (Eq. (4))	111.4042	2.0438	0.8344	0.992
Motorcycles – damage (Eq. (5))	62.1107	38.3051	0	0.993
Four-wheeled vehicles – velocity (Eq. (6))	0.0009	-0.5529	86.9448	0.950
Two-wheeled vehicles – velocity (Eq. (6))	0.0003	-0.2128	34.4674	0.994

We then used these curves to calculate the associated monetary losses by assuming an average typical cost of the specific vehicle model. A fixed share of minivans among the total number of four-wheeled vehicles was assumed equal to 30%, uniformly distributed across the study area. Both parked and moving vehicles are considered in the above-mentioned damage scenarios.

https://doi.org/10.5194/egusphere-2026-1625
Preprint. Discussion started: 19 May 2026

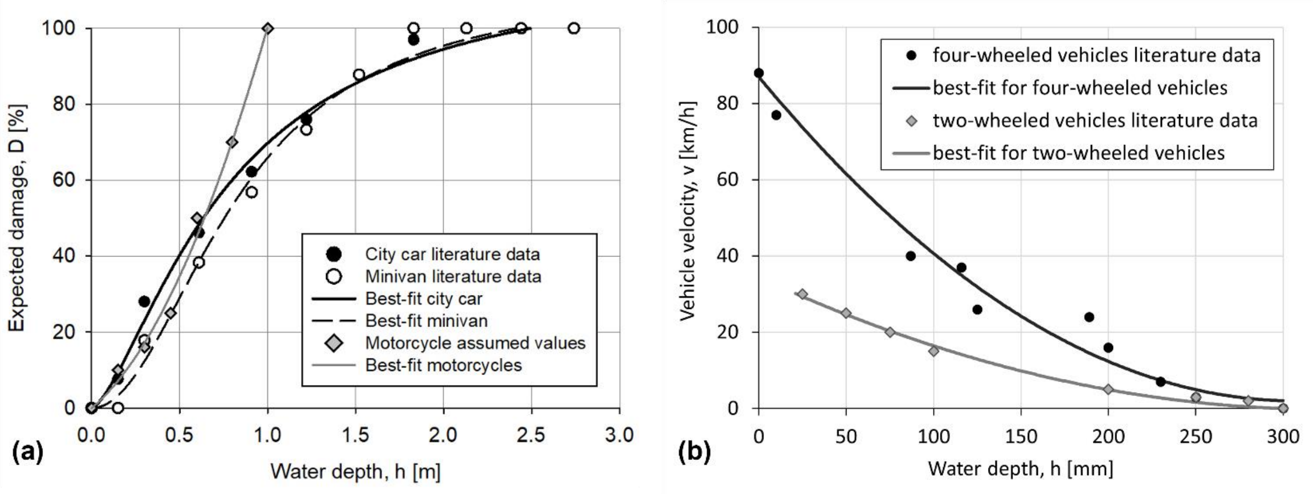

Figure 3: percentage expected damage curves (a) for city cars and minivans derived as best-fit curves from data reported by Martínez-Gomariz et al. (2019) and for motorcycles as assumed in this study. Depth-disruption functions (b) derived from Pregnolato et al. (2017) and Hilly et al. (2018) for four and two-wheeled vehicles, respectively.

We also quantified the traffic disruption by calculating the speed reduction of each moving vehicle present on the road as a function of water depth, using the second-order polynomial curves (see eq. 6) obtained by Pregnolato et al. (2017) and Hilly et al (2018) for four and two-wheeled vehicles, respectively (see Figure 3b). Data derive from experimental, observational and modelling studies. The coefficients of the two curves with the associated R-squared are reported in Table 1, when h is expressed in millimeters and the velocity of the vehicle (v) in km/h. The validity of the curves extends to a water depth of 30 cm, beyond which vehicles are completely blocked.

$$v(h) = a \cdot h^2 + b \cdot h + c \tag{6}$$

From the comparison of the investigated scenarios, we calculated the effect of the daylight and traffic conditions and the role undetected elements (FNs).

3 Results

This section is structured in three main parts. First, pluvial flooding hazard maps are compared by considering various rain measurement methods, with sensors positioned both within and outside the study area. The Arpal-FI rain gauge is assumed to be the reference instrument for the comparison, since it is located within the investigated urban area. Then, the performance of the adopted deep learning-based object detection model in identifying UFDs was synthesised and discussed using the aforementioned performance indices. Finally, different damage scenarios were developed, and the impact of having only approximate rainfall measurement data (different from the rain gauge located within the study area), as well as the influence of false negatives, was quantified. Moreover, the two available datasets of aerial images with resolutions of 5 and 20 cm were

https://doi.org/10.5194/egusphere-2026-1625
Preprint. Discussion started: 19 May 2026

used to reproduce different hourly periods and compare different traffic conditions, in the hypothesis that the aerial picture is representative of a specific exposure scenario.

3.1 Pluvial flooding hazard scenarios

The dynamic evolution of the investigated rainfall event was captured by the meteorological radar, as shown in Figure 4. The black square represents the location of the urban area under investigation in the Sampierdarena district. The analysis shows that the rainfall event passes over the area under investigation in approximately forty minutes. At 9:15, the rainfall intensity is around 20 mm/h. In just five minutes, this value doubles and remains at this level for a further fifteen minutes. At 09:45, the rainfall intensity value drops to approximately zero. This means that rainfall data with a time resolution of five minutes or less (e.g. one minute) can capture the main features of the rainfall event and contribute to flooding analysis. Nevertheless, although the observed dynamic evolution is consistent with that measured by the Arpal-FI rain gauge, the radar-measured RI values are less than half of those reported by the rain gauge due to the average field size that characterises radar measurements. This means that in terms of magnitude, the adoption of radar measurements leads to an underestimation of the hazard posed by pluvial flooding, which is reflected in the quantification of associated damage.

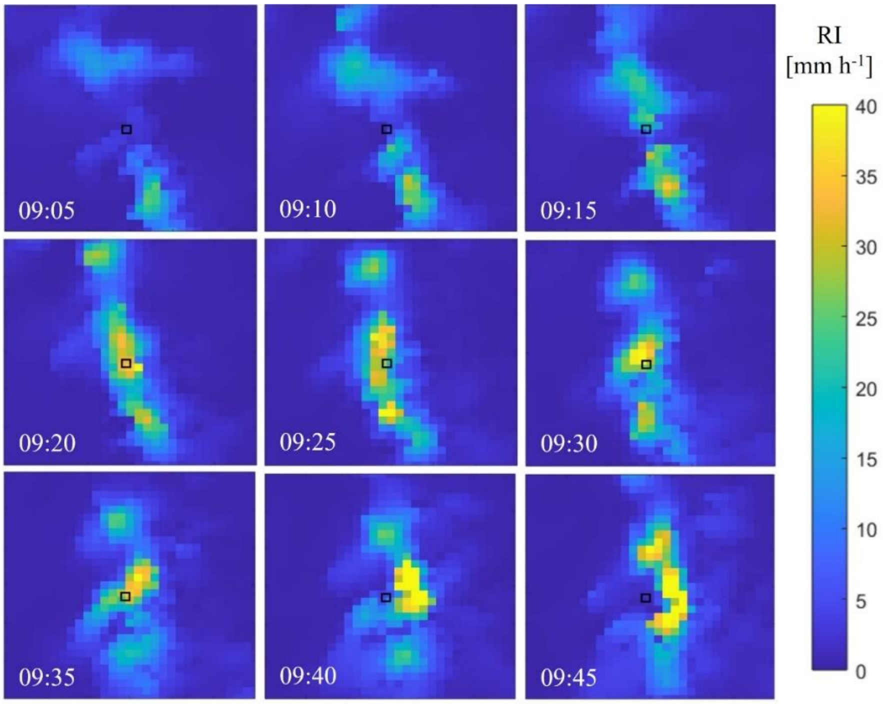

Figure 4: dynamic evolution of the rainfall event occurred on September 24th, 2022, as measured by the meteorological radar. The pixel size is equal to 1 km in both directions. The black square indicates the position of the Sampierdarena urban area.

The ratio between the rainfall amount measured by each instrument and the reference one (Arpal-FI), as well as the peak intensity ratio, is summarised in Table 2. It is evident that the two Arpal rain gauges positioned outside of the study area do not capture the magnitude and the duration of the rainfall event measured at the reference station. On the contrary, the SRSs, despite significant underestimation of the peak rainfall, provide a good representation of the temporal evolution and total volume of the reference rainfall event. This can be ascribed to the larger spatial representativeness of the SRS as compared to the traditional rain gauge. The satellite link may indeed cross the bulk of the rainfall event, even though the precise location of the sensor is not immersed in the rain field.

Table 2: Rainfall event and flooding parameters for each sensor as a comparison with the reference instrument (Arpal-FI).

Sensor	Rainfall amount ratio	Peak intensity ratio	Flood volume ratio conditional on h>5cm	Max flood depth ratio
Arpal-CF/Arpal-FI	0.64	0.55	0.50	0.91
Arpal-CA/Arpal-FI	0.62	0.59	0.54	0.91
SRS-SA/Arpal-FI	1.07	0.56	0.90	1.02
SRS-CA/Arpal-FI	1.02	0.38	0.84	1.01
SRS-BO/Arpal-FI	1.08	0.45	0.79	1.03

Hazard scenarios are reported in terms of maps of the maximum water depth in Figure 5. Table 2 reports the flooded volume (conditional on a minimum water depth h = 5 cm) and the maximum water depth obtained by using the various rain sensors and normalised with those obtained from the reference rainfall. The temporal evolution of the modelled flood volume (conditional on a minimum water depth h = 5 cm) is reported in Figure 6.

Both the maps and the ratios indicate that differences are limited when SRSs are employed, either when positioned in the same urban area or at a few kilometres from the area. The point nature of measurements taken at rain gauge stations outside of the study area reduces to a half the flooded volume and by 10% the maximum water depth. In all cases a reduction of the extension of the flooded area is shown indicating an underestimation of the flood hazard.

The temporal evolution of the modelled flood volume reflects the differences already observed in Figure 2 in the temporal evolution of the rainfall event as measured by the various sensors, including the shift in time of the core of the event in the case of the two rain gauge stations located outside of the study area. However, the significant underestimation of the rainfall peak experienced by the SRS measurements does not result in a corresponding reduction of the peak of the flood volume, probably because of the accumulation of flooded water in the depressed areas of the domain.

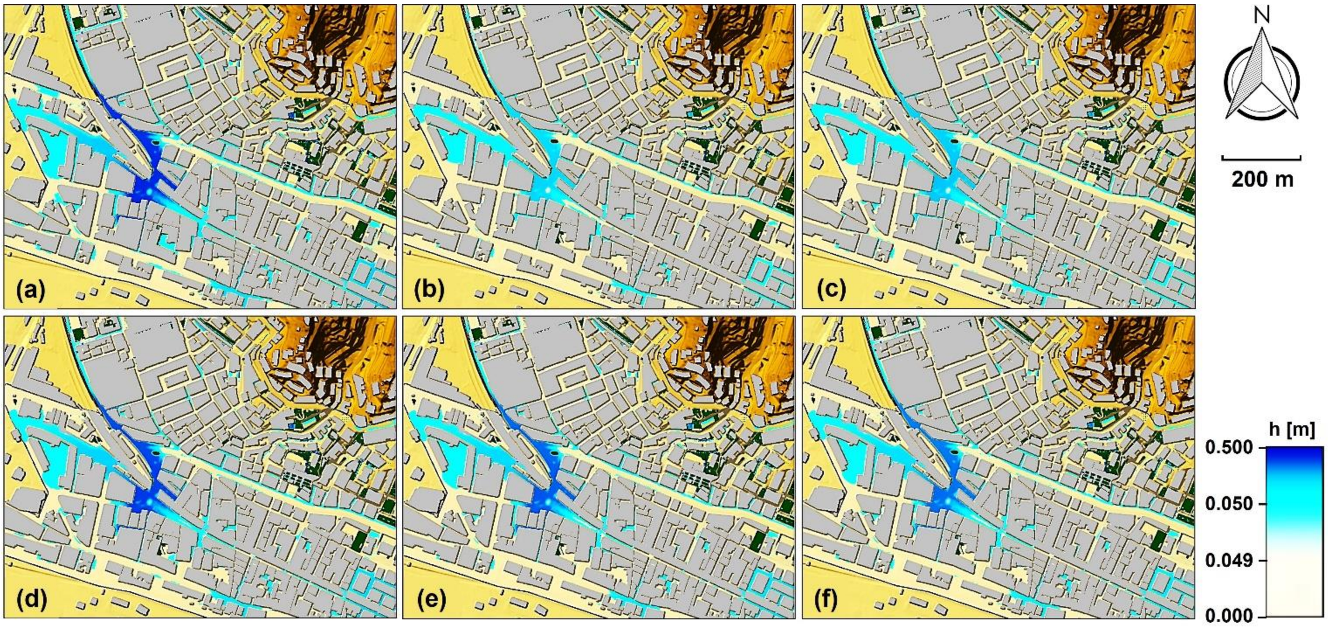

Figure 5: Maps of the maximum water depth during the investigated rainfall event using measurements from: (a) Arpal-FI, (b) Arpal-CF, (c) Arpal-CA, (d) SRS-SA, (e) SRS-CA and (f) SRS-BO. Background brown shades indicate terrain elevation.

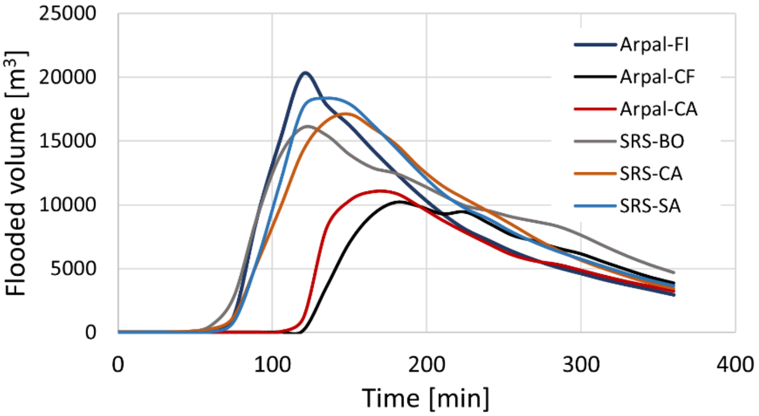

Figure 6: temporal evolution of the modelled flood volume (conditional on a minimum water depth h = 5 cm) obtained when using as input the rainfall measurements from various sensors.

3.2 Performance of the YOLO model

Preliminary analyses were conducted using thirty training patches and five validation patches. First, the model was run using the dataset at 5 cm resolution and the traditional annotation approach, with the goal to detect four-wheeled vehicles. Two recent versions of the YOLO model (v8n and v11n) were tested to determine an effective balance between performance and computational burden. As can be seen in Table 3, there is no improvement in performance, on the considered dataset, when

https://doi.org/10.5194/egusphere-2026-1625
Preprint. Discussion started: 19 May 2026

v11n is used. Despite newer models like YOLOv11n being more parameter-efficient, YOLOv8n provides a simpler and well-established lightweight baseline, ensuring lower implementation complexity. The latency on a GPU is lower for v8n (0.99ms) than for v11n (1.5ms). Therefore, based on the outcome of this preliminary analysis, the v8n version was adopted consistently to conduct the analysis on the entire study area.

A second analysis was performed using the detailed annotation approach, and the results were compared with those obtained using the traditional one. The comparison (see Table 3) shows that accurate results are obtained with the traditional annotation approach, although these are slightly lower than those obtained through manual annotation.

Table 3: Performance indices for the detection of four-wheeled vehicles using different annotation approaches and YOLO models for aerial images at the resolution of 5 cm, during the training (Train) and the validation (Val) stages.

YOLO model	Annotation approach	Data Splits	Instances	*P*	*R*	*F1*
v11n	traditional	Train	899	0.94	0.93	0.93
v11n	traditional	Val	179	0.92	0.87	0.89
v8n	traditional	Train	899	0.94	0.92	0.93
v8n	traditional	Val	179	0.93	0.88	0.90
v8n	detailed	Train	899	1.00	0.92	0.96
v8n	detailed	Val	179	0.93	0.97	0.95

An attempt was made to detect two classes in the same run. The observed reduction in the F_1 score was around 5% for the validation patches. As the aim of this study is to quantify expected damage under various pluvial flooding scenarios, the final object detection stage was conducted using the most effective procedure: detailed annotation with separate runs for each class.

The same approach involving detailed annotation was adopted for the dataset with a resolution of 20 cm to detect four-wheeled vehicles, for which twenty-three patches were used for training. A fine-tuning approach was adopted for this step: the pre-trained 5 cm model was used, and further training was then carried out to make it aware of the new image resolution.

Validation was performed manually on the entire study area in terms of TPs, FNs, and FPs as shown in Figure 7, and the associated performance indices are listed in Table 4 (indicating the performance of the adopted model for the test stage). By comparing the last line of Table 3 with the first line of Table 4, we can see that the performance index values are the same, with maximum differences of 0.01. The results suggest that the adopted network performs less well at detecting two-wheeled vehicles than four-wheeled vehicles. This is due to the limited size of these objects, and could only be achieved using the dataset with a resolution of 5 cm. The detection of four-wheeled vehicles using the dataset at a coarser resolution is obtained with rather high performance, yet lower by about 10% than in the case of the dataset at 5 cm resolution. These less accurate performances are expected and may also be affected by uncertainty relating to the operator who manually annotated vehicles across the entire study area. This task became more challenging as image resolution and object size decreased.

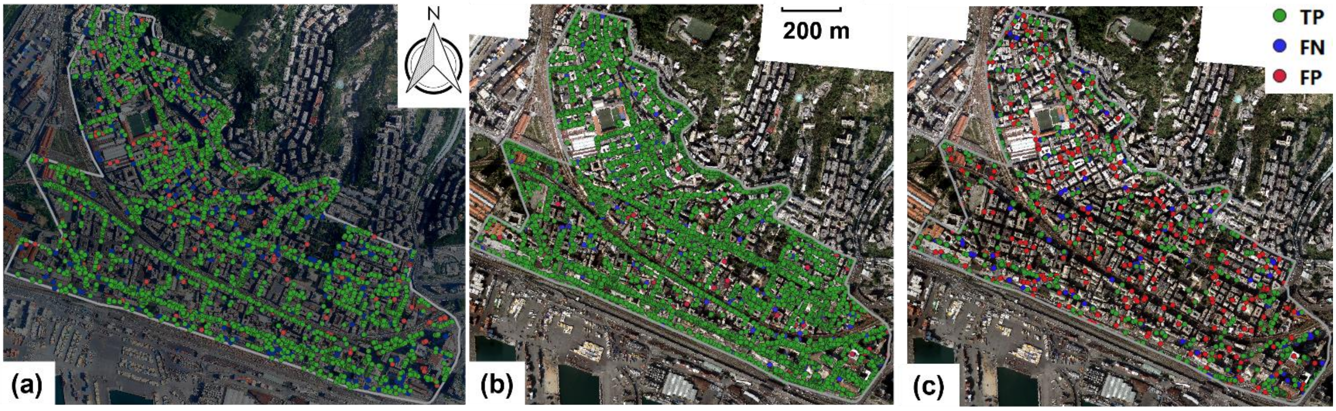

Figure 7: Results of the YOLOv8n model trained using detailed annotation for four-wheeled vehicles at the spatial resolution of 20 cm (a) and 5 cm (b), and for two-wheeled vehicles (c), as detected from the aerial images at 5 cm. The grey line indicates the boundaries of the study area. The aerial image at the spatial resolution of 20 cm is available at the Cartographic Portal of the Liguria Region, while the aerial image at the spatial resolution of 5 cm o is provided by the Cartographic Repository of the Municipality of Genoa.

A total of 1494 two-wheeled vehicles are present in the study area, while the dataset contains 5991 and 3447 four-wheeled vehicles at 5 cm and 20 cm resolution, respectively. These differences are due to the corresponding times of day: rush hour and off-peak, respectively. In the off-peak traffic scenario, most vehicles are parked, with only a limited number travelling along the roads.

Table 4: Performance indices for the detection of four-wheeled and two-wheeled vehicles, for the two datasets at spatial resolutions of 5 cm and 20 cm, across the entire study area

Class	TP[%]	FN[%]	FP[%]	*P*	*R*	F_1
Four-wheeled vehicles at 5 cm	94.3	5.7	2.3	0.94	0.98	0.96
Two-wheeled vehicles at 5 cm	64.9	35.1	29.6	0.65	0.70	0.68
Four-wheeled vehicles at 20 cm	82.0	18.0	10.2	0.82	0.90	0.86

3.3 Damage assessment

The stability and roadworthiness of passenger cars were quantified in specific energy classes across the entire study area using the maximum water depth and velocity. This assumption was verified by simulating the propagation of a flood with a return period of 10 years. Figure 8a shows the flooded areas coloured according to their associated energy class (low in green, medium in orange, and high in red). The values of specific energy obtained using the maximum water depth and velocity were compared with the associated temporal evolution by calculating the percentage of the flooded cell for each class. The comparison depicted in Figure 8b reveals that the energy obtained using the maximum values is very close to the maximum energy reached during the simulation, meaning that this simplification does not considerably affect the results. Therefore, specific energy maps for the observed event were derived, for each rain sensor, using the simplified procedure (see Figure 9).

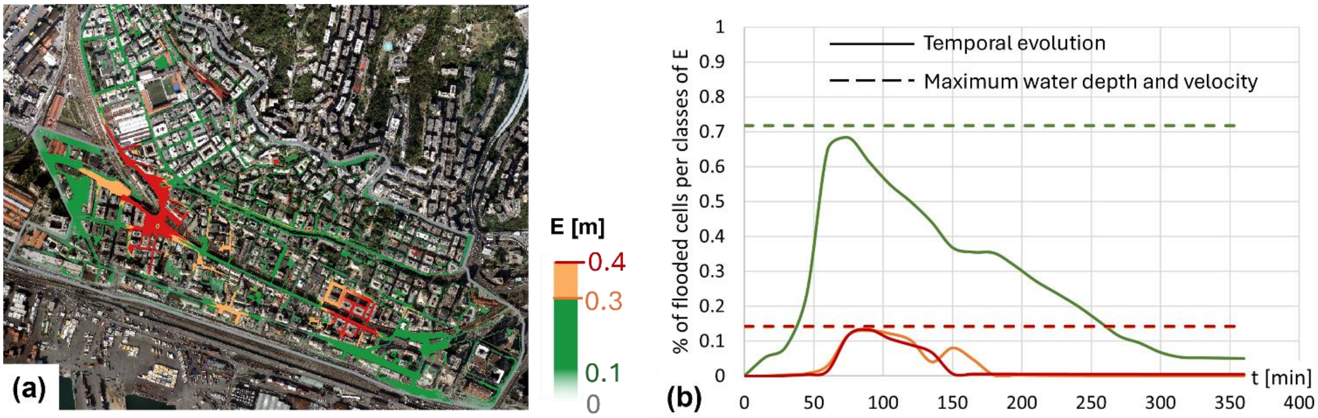

Figure 8: (a) Map of the specific energy of the flow overlaid with the aerial image of the study area (source: Cartographic Repository of the Municipality of Genoa) and (b) comparison between the simplified procedure and the temporal evolution of the percentage of flooded cells at T = 10 years and using the classification reported by Kramer et al. (2016)

Adopting the reference rain gauge (Arpal-FI), which is located within the study area, reveals that the highest class (shown in red) is reached in a significant portion of the territory. This red area is reduced when the SRS-SA is used and disappears entirely when SRS sensors located outside the study area are adopted. A pronounced underestimation is observed when rain gauges outside the study area are used, with the associated risk being completely overlooked (only green areas are present).

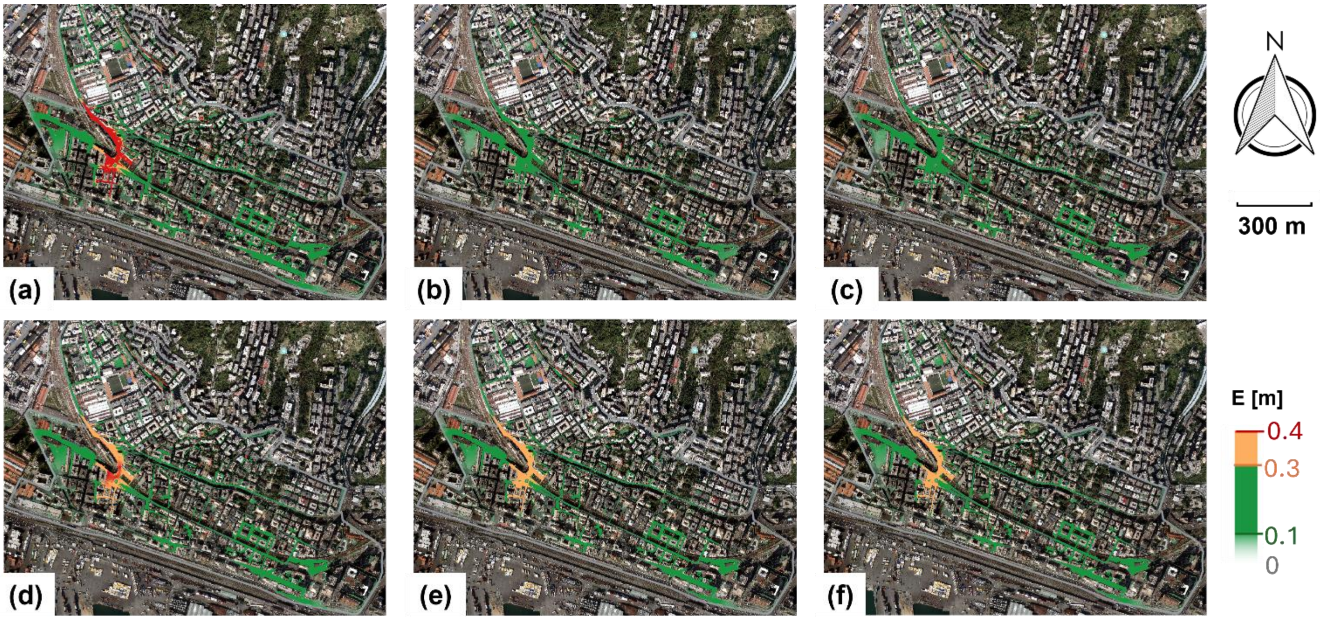

Figure 9: Maps of the specific energy of the flow, during the investigated rainfall event using measurements from: (a) Arpal-FI, (b) Arpal-CF, (c) Arpal-CA, (d) SRS-SA, (e) SRS-CA and (f) SRS-BO and the classification reported by Kramer et al. (2016), overlaid with the aerial image of the study area (source: Cartographic Repository of the Municipality of Genoa)

The number of flooded vehicles that fall in each class was quantified by sampling the maps of specific energy at the exact position of each detected vehicle (considering both parked and travelling four-wheeled vehicles), as shown in Figure 10.

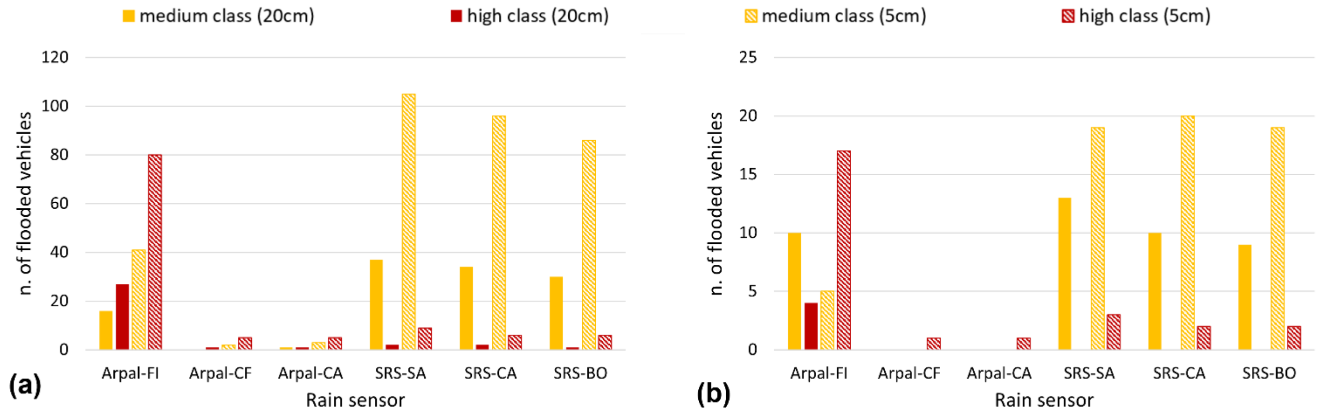

Figure 10: Number of flooded four-wheeled vehicles for each specific energy class as defined by Kramer et al. (2016) using different rain sensors. TPs (a) and FNs (b) are separately considered. Solid and dashed colour patterns are indicative of results obtained at the spatial resolution of 20 and 5 cm, respectively.

The impact of FNs is significant, accounting for 20% of vehicles in the high energy class. FPs, on the other hand, do not affect the analysis as they are typically positioned on roofs, and there is only one vehicle in the red zone. According to the energy maps, the use of SRS sensors shifts the number of vehicles from the high class to the medium class. Using rain gauges outside the study area incorrectly indicates that the risk is negligible.

The impact of using rainfall data different from the very local rain gauge positioned inside the study area is analysed by comparing the total monetary damage experienced by two- and four-wheeled vehicles. Results are presented in terms of the underestimation of the total damage deriving from using other rain sensors than Arpal-FI. The damage is calculated per each individual vehicle based on the damage curves presented in Figure 3a and the local hydraulic conditions, eventually summed over the total number of flooded vehicles. The reference total monetary damage is about four million Euros, including both two- and four-wheeled vehicles.

The underestimation due to the use of less representative rain sensors is reported in Figure 11 for the two datasets at 5 and 20 cm resolution, representing rush and off-peak hour traffic, respectively. The contribution of FNs is more pronounced at a resolution of 20cm. Considering both TPs and FNs the monetary damage fluctuations can reach 40%, depending on the time of day. As already mentioned above, the role of FPs is negligible.

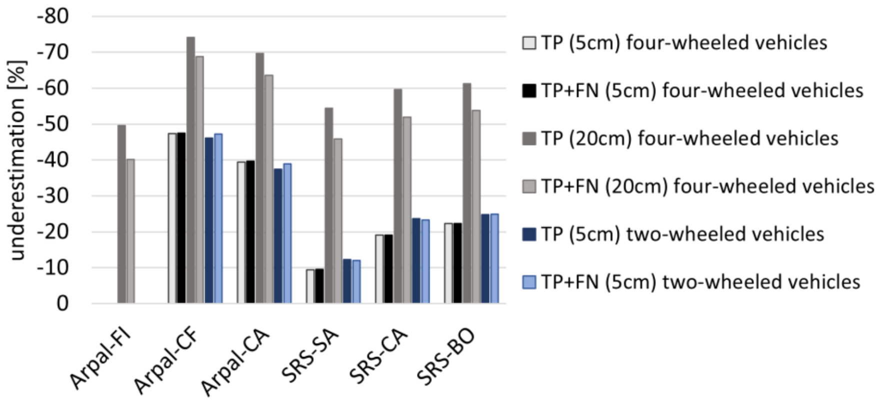

Figure 11: Underestimation of the total monetary damage with respect to using Arpal-FI at 5cm resolution and considering TPs only or the sum of TPs and FNs per each rainfall scenario.

To assess traffic disruption effects in the study area, we focused on the moving vehicles only, therefore excluding parked vehicles along the streets and in dedicated parking areas. This was obtained by filtering out vehicles detected within a strip of one to two meters (depending on the street dimensions) along the main streets. Using the literature depth-disruption functions reported in Figure 3b, we calculated the maximum velocity achievable by each detected two- and four-wheeled vehicle along the streets according to the local water depth.

Results are reported in Figure 12, where the various panels represent results obtained by using different rain sensors and dots indicate the detected vehicles, colour-coded according to their maximum velocity as constrained by the flooding conditions. Only TPs are reported in these pictures since we checked that the role of FNs in the analysis of traffic disruption is negligible.

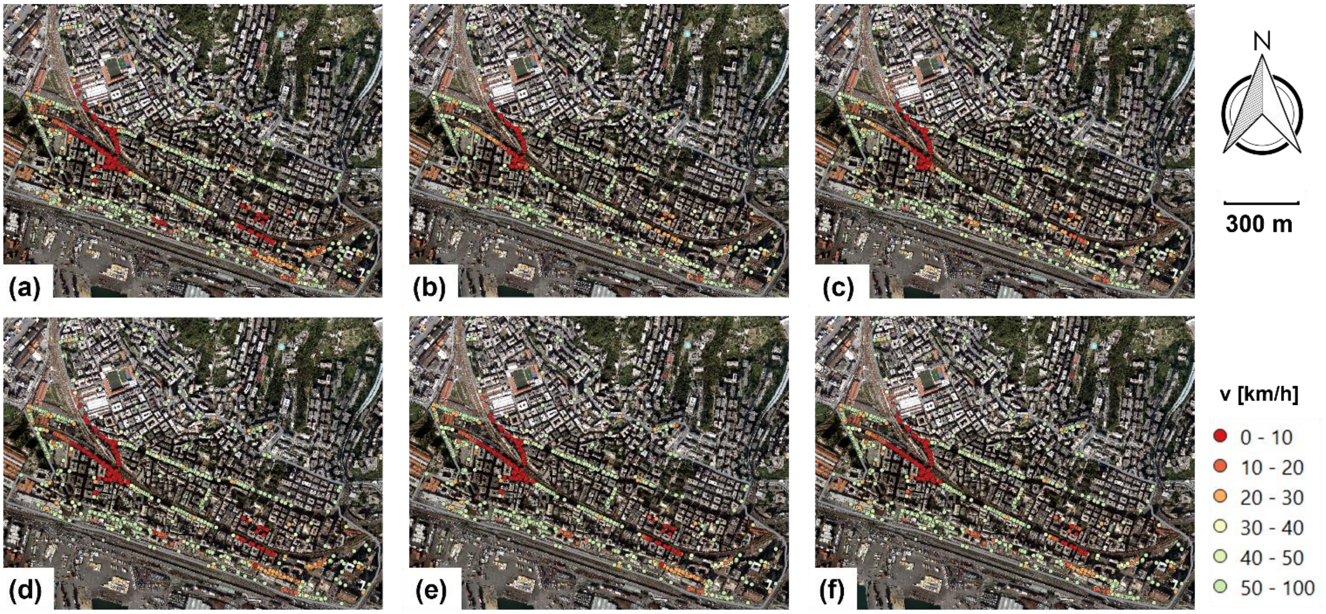

Figure 12: Traffic disruption maps showing the maximum velocity of detected vehicles according to the local flooding conditions using measurements from: (a) Arpal-FI, (b) Arpal-CF, (c) Arpal-CA, (d) SRS-SA, (e) SRS-CA and (f) SRS-BO. The depth-disruption functions were derived from Pregnolato et al. (2017) and Hilly et al. (2018) for four- and two-wheeled vehicles, respectively. The aerial image of the study area is provided by the Cartographic Repository of the Municipality of Genoa.

Compared to the reference scenario, when using rainfall observed by the SRS sensors, the resulting maps indicate that the three main roads are fully blocked (vehicles are about stopped due to the water depth), while if rain gauges outside of the study area are used, then the output result indicates the traffic in the road along the southern boundary of the domain is still regular.

Considering the results described above, the role of FNs in automated object detection can be deemed relevant to the safety analysis when the specific energy of the flow is considered, while it is quite negligible for the assessment of traffic disruption. For the assessment of monetary damage, FNs only affect the dataset at 20cm resolution.

https://doi.org/10.5194/egusphere-2026-1625
Preprint. Discussion started: 19 May 2026

Conclusions

Various pluvial flooding scenarios were simulated to highlight the impact of data availability and reliability over the assessment of the flood hazard using, as case study, a highly urbanised district of the city of Genoa (Italy). Significant differences were observed and quantified between flooding scenarios obtained by simulating the distribution of excess rainwater when various rainfall data were considered. A real-world event measured by rain gauge stations and opportunistic sensors is used to compare the results.

Although the overall flooding conditions are captured in all examined cases, the results show that significant differences in expected flood volumes and maximum water depths are obtained when using various rainfall information sources. In the case of the simulated event, the contribution of opportunistic sensors located within or in the proximity of the study area largely outperforms that of nearby rain gauge data, even when the gauges are located 5 km away from the study area. This was demonstrated for both the total flooded volume and the maximum water depth.

This work highlights the importance of input rainfall patterns in modelling pluvial flooding scenarios and quantifying associated damage for a specific case study involving a rapidly evolving event with limited spatial coverage. This suggests that input information plays a central role, even when detailed detection of UFDs is considered.

Results show that object detection through the application of a deep learning model to aerial images may provide a reliable tool to assess exposure in highly urbanized areas. The position and class of vehicles are retrieved with quite high accuracy, and the resulting maps can be used, in the hypothesis that the aerial picture is representative of a specific exposure scenario, in conjunction with hydrodynamic simulation results to assess the specific degree of hazard associated with each individual target. When the relevant damage curves are available, relating the local hydrodynamic characteristics to vehicle safety, monetary damage, and traffic disruption effect suffered by each class of targets, exposure and vulnerability can be combined with the hazard level (e.g. in terms of the return period of the investigated event) to quantify the risk. For this study, damage curves were obtained from the literature to illustrate the proposed procedure. However, different curves could be used without affecting the validity of the proposed approach.

The comparison of different scenarios performed in terms of both the rainfall input information and the object detection performance allowed us to assess the role of the various steps involved in the risk assessment chain. It was shown that the classification errors of the deep learning object detection model are less relevant than rainfall variability within a limited spatial scale for risk assessment.

The proposed procedure is replicable in urban and climatological contexts that may differ from those of the sample area investigated in this work. Further research is needed to assess whether additional UFDs can be suitably detected from aerial images at various pixel resolutions. Smaller size UFDs, like waste containers and debris, may be difficult to detect with sufficient accuracy. However, the overall damage associated with pluvial flooding is typically due to vehicle displacement and flooding and/or to traffic disruption. The presented approach is therefore able to quantify the most relevant contributions to the risk of pluvial flooding in urban areas.

Data availability.

The aerial images at the resolution of 5 cm and the DTM are available at the Cartographic Repository of the Municipality of Genoa (https://mappe.comune.genova.it/MapStore2/#/) while the aerial images at the resolution of 20 cm and the vector layer containing the outlines of buildings are available at the Cartographic Portal of the Liguria Region (https://geoportal.regione.liguria.it/) (last access: 21 January 2026). Georeferenced layers of manually detected vehicles and pluvial flooding hazard maps are available upon request by email to the first author (arianna.cauteruccio@edu.unige.it).

Author contributions.

AC, GB, and GM conceived the study and developed the methodology. AC, MM and RR modelled the investigated scenarios and wrote the original draft. AC prepared the validation dataset, analysed and synthesized the results and prepared the visualisation and presentation material. GB and GM supervised the work, reviewed the manuscript, and received funding.

Competing interests.

The authors declare that they have no conflict of interest.

Acknowledgements.

Rain gauge and meteorological radar measurements were kindly provided by the environmental protection agency of the Liguria region (ARPAL) while the measurements of the SRSs were kindly provided by Darts Engineering s.r.l.

Financial support

This study was carried out within the RETURN Extended Partnership and received funding from the European Union Next-GenerationEU (National Recovery and Resilience Plan – NRRP, Mission 4, Component 2, Investment 1.3 – D.D. 1243 2/8/2022, PE0000005).

References

Ammar, A., Koubaa, A., Ahmed, M., Saad, A., and Benjdira, B.: Vehicle detection from aerial images using deep learning: A comparative study, Electronics, 10(7), 820, doi:10.3390/electronics10070820, 2021.

Casulli, V.: A high-resolution wetting and drying algorithm for free-surface hydrodynamics. Int. J. Numer. Methods Fluids 60 (4), 391–408, doi:10.1002/fld.1896, 2008.

Chinchella, E., Cauteruccio, A. and Lanza, L.G.: Modelling the role of permeable pavements in mitigating pluvial flooding in Genoa, Italy, J. Hydrol. Reg. Stud., 60, 102596, doi:10.1016/j.ejrh.2025.102596, 2025.

Colli, M., Stagnaro, M., Caridi, A., Lanza, L.G., Randazzo, A., Pastorino, M., Caviglia, D.D. and Delucchi, A.: A field assessment of a rain estimation system based on earth-to-satellite microwave links. IEEE Trans. Geosci. Remote Sens, 57(5), 2864-2875, doi:10.1109/TGRS.2018.2878338, 2019.

Ding, J., Xue, N., Xia, G.-S., Bai, X., Yang, W., Yang, M., Belongie, S., Luo, J., Datcu, M., Pelillo, M., and Zhang, L.: Object Detection in Aerial Images: A Large-Scale Benchmark and Challenges, IEEE Transactions on Pattern Analysis and Machine Intelligence, doi: 10.1109/TPAMI.2021.3117983, 2021.

Federici, B., Gragnani, G.L., Parodi, G., Randazzo, A., Caviglia, D., Pastorino, M., Sguerso, D., Caridi, A. and Montecucco, C.: System and method for monitoring a territory, European Patent EP 2 688 223 B1, 2014.

FEMA - Federal Emergency Management Agency: Multi-hazard loss estimation methodology. Flood model. Hazus-MH MR5 technical manual. Washingtong, DC: Department of Homeland Security. Mitigation Division 449p, 2015.

Giannetti, F. and Lanza, L.G.: Rain Sensors (Editorial), Sensors, 23, 6934, doi:10.3390/s23156934, 2023.

Hilly, G., Vojinovic, Z., Weesakul, S., Sanchez, A., Hoang, D.N., Djordjevic, S., Chen, A.S. and Evans, B.: Methodological Framework for Analysing Cascading Effects from Flood Events: The Case of Sukhumvit Area, Bangkok, Thailand, Water, 10(1), 81, doi:10.3390/w10010081, 2018.

Iqbal, U., Riaz, M. Z. B., Barthelemy, J., Hutchison, N., and Perez, P.: Floodborne objects type recognition using computer vision to mitigate blockage originated floods, Water, 14(17), 2605, doi:10.3390/w14172605, 2022.

Jonkman, S.N. and Vrijling, J.K.: Loss of life due to floods, J. Flood Risk Manag., 1, 43–56, doi:10.1111/j.1753-318X.2008.00006.x, 2008.

Khanam, R., and Hussain, M.: Yolov11: An overview of the key architectural enhancements, arXiv, preprint, doi:10.48550/arXiv.2410.17725, 2024.

Kramer, M., Terheiden, K., and Wieprecht, S.: Safety criteria for the trafficability of inundated roads in urban floodings, Int. J. Disaster Risk Reduct., 17, 77-84, doi:10.1016/j.ijdrr.2016.04.003, 2016.

Loglisci, N., Boni, G., Cauteruccio, A., Faccini, F., Milelli, M., Paliaga, G., and Parodi, A.: The role of citizen science in assessing the spatiotemporal pattern of rainfall events in urban areas: a case study in the city of Genoa, Italy. Nat. Hazards Earth Syst. Sci., 24, 2495–2510, doi:10.5194/nhess-24-2495-2024, 2024.

Martínez-Gomariz, E., Gómez, M., Russo, B., and Djordjevic, S.: Stability criteria for flooded vehicles: a-state-of-the-art review, J. Flood Risk Manag., 11, S817–S826, doi:10.1111/jfr3.12262, 2018.

Martínez-Gomariz, E., Gómez, M., Russo, B., Sánchez, P., and Montes, J.A.: Methodology for the damage assessment of vehicles exposed to flooding in urban areas, J. Flood Risk Manag., 12, e12475, doi:10.1111/jfr3.12475, 2019.

Mercier, F., Barthès, L., and Mallet, C.: Estimation of finescale rainfall fields using broadcast TV satellite links and a 4DVAR assimilation method, J. Atmos. Ocean. Technol., 32, 1709–1728, doi:10.1175/JTECH-D-14-00125.1, 2015.

Palla, A., Colli, M., Candela, A., Aronica, G.T., and Lanza, L.G.: Pluvial flooding in urban areas: The role of surface drainage efficiency, J. Flood Risk Manag., 11, S663–S676, doi:10.1111/jfr3.12246, 2018.

Pregnolato, M., Ford, A., Wilkinson, Richard, S.M. and Dawson, R.J.: The impact of flooding on road transport: A depth-disruption function, Transp Res D Transp Environ, 55, 67-81, doi:10.1016/j.trd.2017.06.020, 2017.

Redmon, J., Divvala, S., Girshick, R., and Farhadi, A.: You Only Look Once: Unified, Real-Time Object Detection, 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 779-788, doi:10.1109/CVPR.2016.91, 2016.

Shen, J., Zhou, W., Liu, N., Sun, H., Li, D., and Zhang, Y.: An anchor-free lightweight deep convolutional network for vehicle detection in aerial images, IEEE Trans. Intell. Transp. Syst., 23(12), 24330–24342, doi:10.1109/TITS.2022.3203715, 2022.

Terven, J., Córdova-Esparza, D.-M., and Romero-González, J.-A.: A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas, Mach. Learn. Knowl. Extr. 2023, 5(4), 1680-1716, doi:10.3390/make5040083, 2023.

Tkachenko, M., Malyuk, M., Holmanyuk, A., and Liubimov, N.: Label Studio: Data labeling software, https://github.com/heartexlabs/label-studio, 2020.

USACE: HEC-RAS river analysis system, Hydraulic Reference Manual, US Army Corps of Engineering, Hydrologic Engineering Center version 6.0, 546, https://www.hec.usace.army.mil/confluence/rasdocs/ras1dtechref/latest, 2021.

Valero, D., Bayón, A., and Franca, M.J.: Urban Flood Drifters (UFDs): Onset of movement, Sci. Total Environ., 927, 171568, doi:10.1016/j.scitotenv.2024.171568, 2024.

Xia, J., Falconer, R.A., Xiao, X. and Wang, Y.: Criterion of vehicle stability in floodwaters based on theoretical and experimental studies, Nat Hazards, 70, 1619–1630, doi:10.1007/s11069-013-0889-2, 2014.

Yang, M.Y., Liao, W., Li, X., and Rosenhahn, B.: Vehicle detection in aerial images, Photogrammetric Engineering & Remote Sensing, 85(4), 297–304, doi:10.14358/PERS.85.4.297, 2018.