Panakkal, P., & Padgett, J. E. (2024). More eyes on the road: Sensing flooded roads by fusing real-time observations from public data sources. Reliability Engineering & System Safety, 110368.
https://doi.org/10.1016/j.ress.2024.110368

Graphical Abstract

More Eyes on the Road: Sensing Flooded Roads by Fusing Real-Time Observations from Public Data Sources

Pranavesh Panakkal, Jamie Ellen Padgett

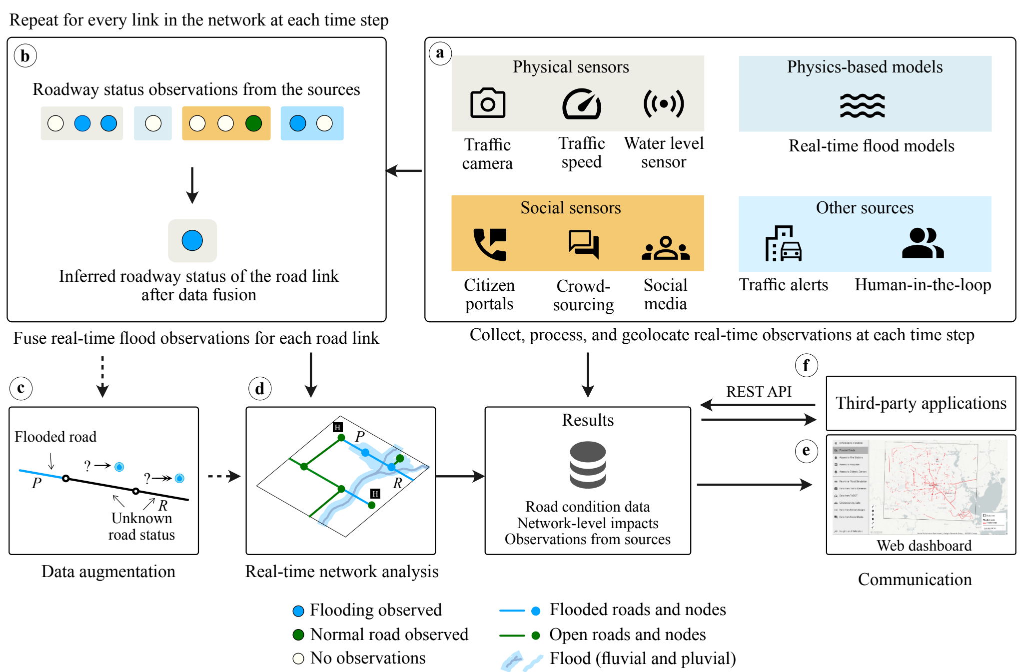

Highlights

More Eyes on the Road: Sensing Flooded Roads by Fusing Real-Time Observations from Public Data Sources

Pranavesh Panakkal, Jamie Ellen Padgett

- A new situational awareness framework for real-time sensing of flooded roads
- Poses methods to infer road condition by fusing observations from public data sources
- Offers communities a pathway to improve situational awareness using existing sources

More Eyes on the Road: Sensing Flooded Roads by Fusing Real-Time Observations from Public Data Sources

Pranavesh Panakkal[a], Jamie Ellen Padgett[a,b]

[a]*Department of Civil and Environmental Engineering, Rice University, 6100 Main St., Houston, 77005, Texas, USA*
[b]*Corresponding author, Email address: jamie.padgett@rice.edu*

Abstract

Reliable sensing of road conditions during flooding can facilitate safe and efficient emergency response, reduce vehicle-related fatalities, and enhance community resilience. Existing situational awareness tools typically depend on limited data sources or simplified models, rendering them inadequate for sensing dynamically evolving roadway conditions. Consequently, roadway-related incidents are a leading cause of flood fatalities (40%-60%) in many developed countries. While an extensive network of physical sensors could improve situational awareness, they are expensive to operate at scale. This study proposes an alternative—a framework that leverages existing data sources, including physical, social, and visual sensors and physics-based models, to sense road conditions. It uses source-specific data collection and processing, data fusion and augmentation, and network and spatial analyses workflows to infer flood impacts at link and network levels. A limited case study application of the framework in Houston, Texas, indicates that repurposing existing data sources can improve roadway situational awareness. This framework offers a paradigm shift for improving mobility-centric situational awareness using open-source tools, existing data sources, and modern algorithms, thus offering a practical solution for communities. The paper's contributions are timely: it provides an equitable framework to improve situational awareness in an epoch of climate change and exacerbating urban flood risk.

Keywords: Urban flooding, Roadway flooding, Situational awareness, Data fusion, Roadway safety, Emergency response, Smart resilience

1. Introduction

Flooding poses a significant risk to urban mobility: While inundated roadways and overtopped bridges isolate communities and limit roadway mobility, the paucity of reliable real-time road condition data causes delays and detours, reduces emergency response efficiency, and poses safety risks [1–10]. Further, existing situational awareness tools are often limited in their ability to accurately sense dynamically evolving road conditions [11, 12], thus limiting communities' ability to respond to flood events. Consequently, mobility-related incidents are linked to 40%-60% of flood fatalities in many developed economies [3–5, 13]. Although structural changes are necessary to reduce flood risk, improving situational awareness could, in the short term, enhance our ability to sense and respond to flooding, reduce flood casualties, and strengthen community resilience. Reliable situational awareness tools are especially essential considering climate-exacerbated flood risk to urban mobility [14, 15], aging or inadequate stormwater infrastructure [16], and the scale of emergency response in major urban centers (for example, first responders evacuated more than 122,300 people during Hurricane Harvey [17]). Situational awareness is defined here as the ability to timely and accurately sense flood impacts on road transportation networks at the link and network levels.

Most existing situational awareness tools for detecting flooded roads, or flooding in general, depend on a limited number of sources and consequently inherit their limitations, biases, and inaccuracies. For example, though physical sensors [18–22] deployed along streets can detect road conditions reliably, deploying, maintaining, and securing sensors at scale is prohibitively expensive. Similarly, although social sensors (social media platforms [23] or custom crowdsourcing tools [24, 25]) can offer enhanced situational awareness, they are often replete with bias, misinformation, noise, or model errors [26–29]—thus limiting their application as the sole source of situational awareness data for emergency response applications. Further, studies [30–33] have also successfully used remote sensing techniques (satellites, UAVs, and other aerial platforms) to infer road or flood conditions. While capable of observing large areas, time delays due to satellite revisit times and unavailability of aerial platforms during inclement weather conditions, such as hurricanes, limit their application for emergency response applications requiring limited time lag. With recent advances in deep learning [34, 35], automated image processing models [36–38] can infer roadway flood conditions from traffic camera images;

however, camera data are often only limited to select watchpoints along major highways. Similarly, authoritative data from the Departments of Transportation [39, 40] are usually limited to major highways or arterial roads, limiting data availability for minor roads and residential streets. Recently, studies [41–44] have shown successful applications of machine learning models to predict flooding and roadway status. Often trained on limited historical or simulated data, these models have unknown reliability and generalizability for unseen future events. Moreover, the data-driven models inherit biases and uncertainties associated with the training data, limiting their application. Studies [45–50] have also used physics-based models to predict roadway conditions at select watchpoints as well as at watershed levels. While more reliable than surrogate models for unseen storms, physics-based models are computationally expensive to run in real-time, and simplifications such as the inability to model storm drainage networks could lead to model errors. Some studies have attempted to use precompiled maps [51] to overcome the computational burden of real-time models at the cost of accuracy. Similarly, studies have also attempted to correlate road conditions to nearby gages [52] or rainfall sensors [39] with varying levels of accuracy. However, such simplified or empirical methods are often insufficient for large-scale emergency response and high-risk applications. While these frameworks have advantages and work reliably for limited case study applications, they often fail to provide comprehensive mobility-centric situational awareness solutions at scale.

The shortcomings of current mobility-centric situational awareness frameworks are primarily due to limited real-time data, as they rely solely on a small number of sources. An alternative is to fuse information from multiple sources using data fusion techniques. When data from compatible sources are combined, their collective observations can overcome their individual limitations. Concurrently, data fusion also engenders the challenge of combining information from disparate sources with varying spatial and temporal resolution, reliability, robustness, and modality. Although real-time mobility-centric applications are limited, examples of data fusion-based methods are available for flood monitoring and hindcasting. For example, Wang et al. [53] used social media data with crowdsourcing data for flood monitoring. Rosser et al. [54] fused remote sensing data with social media data and topographical data for flood inundation mapping. Ahmad et al. [55] used remote sensing and social media to detect passable roads after floods. Frey et al. [56–58] used a digital elevation model and remote sensing images to identify trafficable routes. Albuquerque et al. [59] used social media and authoritative data for filtering reliable social media messages. Bischke et al. [60] used social multimedia and satellite imagery for detecting flooding. Werneck et al. [61] proposed a graph-based fusion framework for flood detection from social media images. These methods showcase the application of the data fusion approach for situational awareness or hindcasting, albeit with a very limited number of data sources. Fusing observations from limited sources (especially leveraging social or remote sensors) might not effectively provide reliable situational awareness data for emergency response applications requiring high reliability and limited time lag. In summary, a comprehensive mobility-centric situational awareness framework that can sense roadway conditions at link and network levels is still lacking in the literature. Such a framework should ideally (a) observe a majority of roads, including residential streets, with limited time lag through all stages of flooding; (b) yield reliable and accurate predictions devoid of spatial, temporal, and social bias or inequity; (c) be robust to provide reliable data even with failure of some dependent data sources; (d) quantify link- and network-level impacts on flooding to facilitate a holistic view of flooding; and (e) be accessible to a majority of communities. This study addresses this need for improved roadway sensing and proposes a mobility-centric real-time situational awareness framework leveraging data fusion.

While a data fusion approach can potentially revolutionize situational awareness, a key challenge remains unaddressed—data sources directly reporting flood road conditions are scarce. In contrast, urban centers are replete with data sources that may either directly or indirectly infer flooding or road conditions. Some common data sources include citizen service portals from the city or utility provider, water level sensors located along streams, and traffic cameras, to name a few. Often, these sources are not primarily designed for sensing flood conditions on roads, although they may provide indirect observations of flooding or flood impacts on roads. For example, live video data offers visual evidence of roadway flooding, and water level sensors provide insights on roads colocated with streams. The value of such data sources was evident during Hurricane Harvey in Houston: many people—including emergency responders—resorted to manually examining data sources to infer probable road conditions to overcome the dearth of reliable real-time road condition data [11]. While manual examination of multiple data sources provided temporary relief, they also could result in information scatter, cognitive overload, increased likelihood of misinterpretation, and the risk of using outdated data. An alternative is to leverage observations from multiple public data sources in an automated data fusion framework to sense current flood conditions. Such a framework could significantly improve situational awareness: they can enhance data availability; reduce information scatter; improve accuracy, robustness, and reliability of road condition data; and reduce the cognitive overload of first responders. Moreover, such a data fusion-centric approach might be more affordable to communities than deploying, maintaining, and securing physical sensors at scale.

This study addresses the need for reliable mobility-centric situational awareness and presents a new framework called Open Source Situational Awareness Framework for Mobility using Data Fusion (OpenSafe Fusion). OpenSafe Fusion leverages data collection and processing, data fusion and augmentation, and spatial and network analyses to infer link- and network-level impacts of flooding by fusing observations from real-time data sources that observed flooding or roadway conditions. Any new situational awareness framework should ideally address the needs of stakeholders; consequently, the design of this framework is informed by insights from extensive stakeholder interviews (n = 24) and needs assessment following the tenets of a user-centered design process [62], a detailed description of which is available in Panakkal et al. [11]. This paper primarily focuses on the methodological underpinning of the OpenSafe Fusion methodology and its components. The remainder of the paper is arranged in three sections. A brief overview of the OpenSafe Fusion methodology is provided in the next section, followed by a case study application of the framework in Houston, Texas. The final section presents key insights from the experiments in the context of mobility-centric situational awareness.

2. Proposed Architecture and Methods

OpenSafe Fusion (Fig. 1) is a modular framework composed of five steps: data acquisition and processing, data fusion, data augmentation, impact assessment, and communication. During the data acquisition step (Fig. 1a), real-time data from select sources are acquired, processed to infer road conditions, and geolocated. During the data fusion step (Fig. 1b), road conditions inferred from the selected sources in the data acquisition step are fused at the road link level to estimate road flood conditions while explicitly accounting for the characteristics of the data sources. Similarly, during the optional

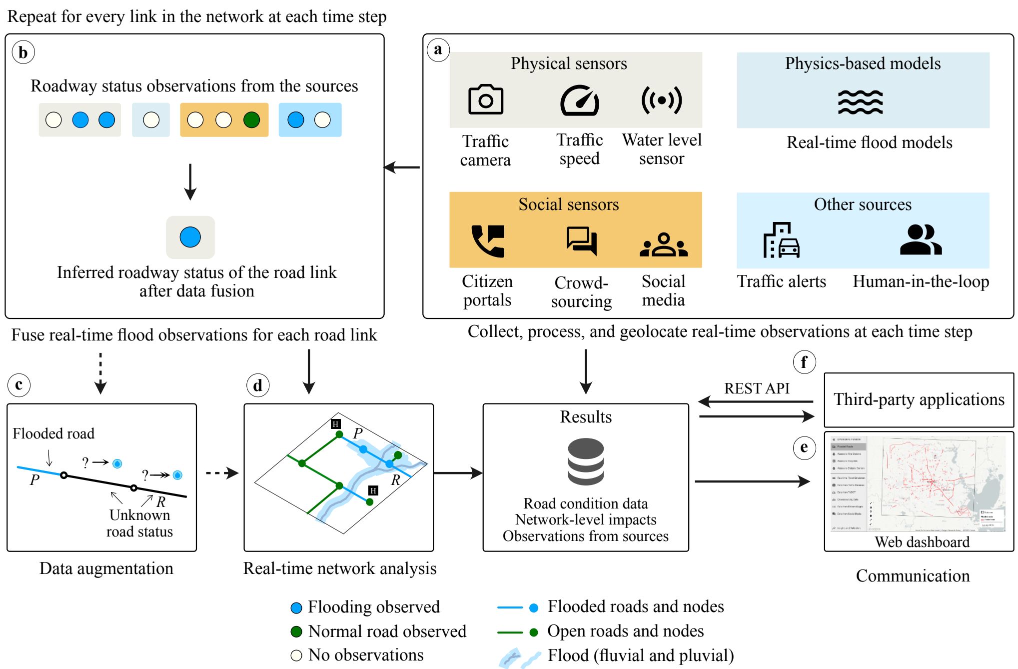

Figure 1: Overview of the OpenSafe Fusion methodology: (a) real-time observations from diverse sources are collected and processed; (b) observations from sources are fused for each road link in the study area to infer the roadway status; (c) data augmentation techniques infer the conditions of roads for which direct observations are unavailable; (d) real-time network analysis quantifies the network-level impacts of flooding; and finally (e) observations and road condition data are communicated to stakeholders via a web dashboard and REST API.

data augmentation step (Fig. 1c), observed roadways status in the current time step are used to infer the state of roads for which direct observations are unavailable. Next, the impact assessment step (Fig. 1d) estimates the network-level impacts of roadway flooding on access to select facilities. Finally, the results are communicated to stakeholders using a web dashboard (Fig. 1e) and REST API (Fig. 1f).

2.1. Data sources

Before deploying the OpenSafe Fusion framework in a region, real-time data sources that can observe flooding or road conditions—either directly or indirectly—should be identified. Some example sources include authoritative sources (e.g., Department of Transportation alerts), social sensors (e.g., crowdsourcing, social media, and citizen service portals), physical sensors (e.g., traffic speed sensors and water level sensors), remote sensors (e.g., UAVs, satellite imagery), and physics-based or hybrid models (e.g., flood alert systems built upon hydrologic and hydraulic models). Once data sources are identified, their historical performance and characteristics are studied. Some example data source characteristics include modality (text from Tweets vs. images from traffic cameras), accuracy, availability, and time lag. Characterization of data sources is necessary to fuse real-time multi-modal data while explicitly accounting for data type heterogeneity, spatial and temporal resolution mismatch, and time lag. Once the data sources are identified and characterized, automated source-specific workflows are developed to extract road condition data from the sources. These data sources and proposed data processing workflows are presented in Section 2.3 after introducing the methodological core of OpenSafe Fusion: the data fusion method.

2.2. Data Fusion

This section presents the methodology proposed to fuse observations from diverse sources and infer the current status of road links. Let the variable $\mathcal{X}_t$ represent the state of a road link at time t and x represent the specific value that $\mathcal{X}_t$ might assume at a time step. A street link could be either impassable (f) or open (o) (i.e., $x \in \{f, o\}$). $p(\mathcal{X}_t = f)$ or simply $p(f)$ denotes the probability that the road link is impassable at a time step.

Consider that time is discretized over a time step δt. The distribution of trajectories of road condition sampled over time $t = 1, ..., T$ is $P(\mathcal{X}_1, ..., \mathcal{X}_T)$ or its abbreviated form $P(\mathcal{X}_{1:T})$. The state of the road at a time is not directly known ($\mathcal{X}_t$ is a hidden variable) but can be observed through sensors with varying characteristics, availability, and noise. $U = \{u^1 ..., u^k\}$ is a set of k sensors available in the study area. A sensor in the context of OpenSafe Fusion is any real-time data source that observes flooding, flood impacts, or road conditions.

As a road link evolves through states $\mathcal{X}_1, ..., \mathcal{X}_T$ under the influence of external actors $e_1, ..., e_T$, the state of the link is observed by sensors in U as $z_1, ..., z_T$. Here, e_t represents the environmental factors ($\{a^1, .., a^p\}$) in the time interval between $t-1$ and t (i.e., in the $(t-1, t]$ time window) that drive the transition of roadway condition from $\mathcal{X}_{t-1}$ to $\mathcal{X}_t$. These environmental factors are often hard to quantify as they include complex factors (rainfall, topography, and built environment) and their interactions at various timescales. To elaborate, transition from $\mathcal{X}_{t-1}$ to $\mathcal{X}_t$ is influenced by

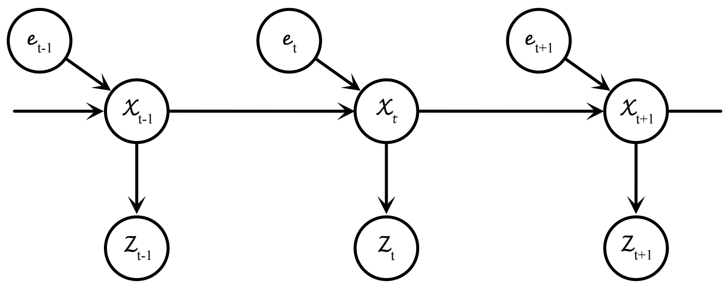

Figure 2: Overview of the dynamic Bayes network for modeling roadway condition.

the actors at time $(t-1, t]$ (i.e., e_t; e.g., rainfall since $t-1$), actions in the short-term (i.e., $e_{t-10:t}$; e.g., delayed peak flow) and actions in the long-term (i.e., $e_{1:t}$; e.g., influence of soil moisture).

Since the actors affecting the transition from $\mathcal{X}_{t-1}$ to $\mathcal{X}_t$ are hard to characterize and the state itself is hidden, an observer is only left with imperfect observations ($z_t = \{z_t^1, ..., z_t^k\}$) by sensors in U at time t to infer the current road link condition $\mathcal{X}_t$. Here, z_t^1 is the observation from sensor u^1 at time t. Figure 2 shows a simplified representation of the transition of road conditions, external actors affecting the transition between time steps, and observations by the sensors at the end of each time step.

OpenSafe Fusion uses Bayes' theorem to fuse observations from diverse sources. Specifically, it uses the discrete form of the Bayes Filter [63] to sense current flood conditions from multi-sensory observations. The formulation presented here is adapted after Thrun et al. [63]. Following Bayes' theorem, the probability of a road link assuming a state at time t (i.e., x_t) given past observations ($z_{1:t}$) and external actions ($e_{1:t}$) is given as:

$$p(x_t|z_{1:t}, e_{1:t}) = \frac{p(z_t|x_t, z_{1:t-1}, e_{1:t}).p(x_t|z_{1:t-1}, e_{1:t})}{p(z_t|z_{1:t-1}, e_{1:t})} \tag{1}$$

Equation 1 can be simplified using a normalizing constant η as:

$$p(x_t|z_{1:t}, e_{1:t}) = \eta.p(z_t|x_t, z_{1:t-1}, e_{1:t}).p(x_t|z_{1:t-1}, e_{1:t}) \tag{2}$$

2.2.1. Prediction Step

In Equation 2, $p(x_t|z_{1:t-1}, e_{1:t})$ represents the Prediction step which (Eq. 3) predicts the current road condition (x_t) from historical records of external actions ($e_{1:t}$) and sensor measurements ($z_{1:t-1}$). Note that the prediction step happens after the external actions in time $(1, t]$ (i.e., $e_{1:t}$) and before receiving the sensor measurements at time t (i.e., z_t is not available).

$$pred(x_t) = p(x_t|z_{1:t-1}, e_{1:t}) \tag{3}$$

The prediction stage can be modeled in its most complete form by employing a surrogate model (e.g., a neural network) that infers the current condition from external actions and sensor data records. To model the intricate relationships it attempts to capture, such a model requires substantial historical data, which is often unavailable, necessitating a simpler formulation for the prediction step. Following the chain rule, Equation 3 can be expressed as:

$$pred(x_t) = \int p(x_t|x_{t-1}, z_{1:t-1}, e_{1:t}).p(x_{t-1}|z_{1:t-1}, e_{1:t})dx_{t-1} \tag{4}$$

Assuming that once the state x_{t-1} is observed, no additional data prior to the time step $t-1$ is required to infer the road condition x_t at t. To elaborate, if a road link is known to be flooded at time $t-1$, only information on the external actions acting on the system between $t-1$ and t is sufficient to predict the state of the road at t. Thus, Equation 4 can be further simplified as:

$$p(x_t|x_{t-1}, z_{1:t-1}, e_{1:t}) = p(x_t|x_{t-1}, e_t) \tag{5}$$

$$pred(x_t) = \int p(x_t|x_{t-1}, e_t).p(x_{t-1}|x_{t-2}, e_{t-1})dx_{t-1} \tag{6}$$

$$pred(x_t) = \int p(x_t|x_{t-1}, e_t).pred(x_{t-1})dx_{t-1} \tag{7}$$

As previously stated, external actions in the present and previous time steps play an active part in the transition of flood conditions in the current time step. Neglecting external actors beyond the current time step may impair the Prediction step's capacity to accurately capture the state transition of road links. The effects of such errors will be more prominent if limited sensor measurements are available at each time step to correct the predicted road condition. Thus, for regions with limited real-time data sources, it is crucial to model the Prediction step accurately without invoking the Markov assumption.

Equation 6 expresses the Prediction step as a recursive update equation. $p(x_t|x_{t-1}, e_t)$ can be modeled using a surrogate model that considers the road condition at time step $t-1$ and external actors e_t to predict the road condition x_t at time t.

It is often impractical to identify the external factors that drive the complex flood process, model their interactions, and sense them in real time. As a result, external actors are not always observable. This necessitates further simplification of Equation 6 as:

$$pred(x_t) = \int p(x_t|x_{t-1}).p(x_{t-1}|x_{t-2})dx_{t-1} \tag{8}$$

$$pred(x_t) = \int p(x_t|x_{t-1}).pred(x_{t-1})dx_{t-1} \tag{9}$$

Equation 9 represents the simplest form of the prediction step. Here, a transition function is used to predict the next state of a road, given the current state of the road (i.e., $p(x_t|x_{t-1})$). Please note that the selected time step will impact the transition function and the influence of environmental factors. Moreover, for a simple two state system (i.e., $x \in \{f, o\}$), a state transition matrix can be used to model the transition function [64].

Finally, mathematical functions describing the Prediction step should ideally be learned from extensive historical data. In the absence of such observations, the function form of the Prediction step can be based on prior knowledge (i.e., expert judgment) for the initial deployment. With additional data available after each storm, such functions should be updated to reflect the most recent information.

2.2.2. Measurement and Update Steps

While the Prediction step predicts the current condition from past observations, the current state of the road is hidden and only observable through imperfect sensors. In the Measurement Step, $p(z_t|x_t, z_{1:t-1}, e_{1:t})$ from Equation 2 is estimated. $p(z_t|x_t, z_{1:t-1}, e_{1:t})$ estimates the probability of observing z_t at time t given the road is at x_t state, past sensor observations are $z_{1:t-1}$ and historical external actions are $e_{1:t}$. Since z_t primarily depends on x_t, it is reasonable to believe that no prior measurements or external actions will yield any additional insights if x_t is known. Thus, Equation 2 reduces to Equation 10.

$$p(x_t|z_t) = \eta'.p(z_t|x_t).pred(x_t) \tag{10}$$

Assuming that multiple sensors will report the road condition at time t, and the sensors independently observe flooding, $p(x_t|z_t)$ can be rewritten as:

$$p(x_t|z_t) = \eta''. \prod_{i=1}^{k} p(z_t^i|x_t).pred(x_t) \tag{11}$$

Here, $p(z_t^i|x_t)$ is the likelihood of observing a sensor measurement z_t^i for sensor u^i at time t give the state of the road x_t. Similar to the Prediction step, surrogate functions can be developed to model $p(z_t^i|x_t)$ either from historical data or expert judgement. Data sources in OpenSafe Fusion often observe flooding independently of other data sources. For example, traffic cameras sense flooding independently of physics-based flood models. However, not all data sources observe flooding independently; dependency on other sources is common in social sensors, where people will report flooding based on data from other sources (e.g., traffic cameras). Sources with extensive interdependencies might disproportionately affect model predictions if Eq. 11 is adopted. While the impacts of such interdependencies on model accuracy are generally limited (as they represent a confirming observation), with extensive historical data, better models capturing the $p(x_t|z_t)$ can be developed that also consider interdependencies in the data sources.

OpenSafe Fusion uses several data sources as sensors. The performance of the sensors and consequently $p(z_t^i|x_t)$ vary both spatially and temporally. For example, observations from the flood model used in OpenSafe Mobility are more reliable near a bayou than in other areas. Similarly, flood models are less accurate for small floods (or in the early stages of the flood) than for severe floods (or in the later stages). Further, environmental and sociodemographic factors may influence sensor performance. For example, camera data are more reliable under sufficient illumination. Hence, automated flood detection from camera data might be more reliable during a bright day. Likewise, it is more likely to acquire better social media data for urban regions with more active users compared to sparsely populated regions. While quantifying the influence of different factors is difficult, it is necessary to reliably estimate current flood conditions from diverse data sources. Finally, observations from different data sources may be available at different rates; the OpenSafe Fusion uses the latest available data from the sources for each link for fusion. In scenarios with significant delay in receiving the data, OpenSafe Fusion reruns all affected timesteps for the reported road link. It is important to carefully choose the time step (δt) after considering data availability and frequency, accuracy, and computational resources.

2.3. Data Processing Workflows

This subsection provides nine examples of data processing workflows for deriving input data to the fusion method modeled after the data available in Houston, TX. These workflows also serve as templates for transferring the framework to other study regions.

2.3.1. Department of Transportation Alerts

Departments of Transportation (DOT), such as the Texas Department of Transportation (TxDOT), operate traffic information systems (TIS) to alert road users on real-time road conditions. For example, DriveTexas [40] is an online traffic information system developed and operated by TxDOT to provide real-time information on highway conditions in Texas. In DriveTexas, road conditions are reported by reliable sources such as law enforcement and are then verified by TxDOT employees or contractors (Fig 3a). The reported road conditions include the location of incidents such as accidents, construction, damage, flooding, and snow (Fig 3b). Users can access roadway status using a variety of mediums, including web dashboards [40] and APIs [65].

During operation, OpenSafe Fusion utilizes the API functionality offered by DOTs to collect real-time information at regular intervals. DOT road condition data are often geocoded and can be used directly in OpenSafe Fusion. Rarely, minor geometry differences in the reported road geometry may occur due to disagreements between the road databases used by OpenSafe Fusion and DOT. In such cases, mapping functions are used to locate roads from the OpenSafe Fusion road network that correspond to the roads in the official road condition reports. Example mapping functions might consider proximity, orientation, and road description to perform the mapping.

2.3.2. Traffic Speed

Real-time traffic speed data (e.g., Houston TranStar [39], Waze [24]) can be used to monitor highway performance. Typical traffic speeds could indicate the normal functioning of roads, and any abnormally low traffic speed could imply adverse or atypical conditions. OpenSafe Fusion leverages real-time traffic speed data to sense the opening of flooded roads. To elaborate, OpenSafe Fusion assumes that if the traffic speed is near normal (as defined using a threshold value or the posted speed limit), it is likely that the road is open to traffic—either partially or fully. OpenSafe Fusion does not use real-time speed data to identify flooded roads, as various factors, including

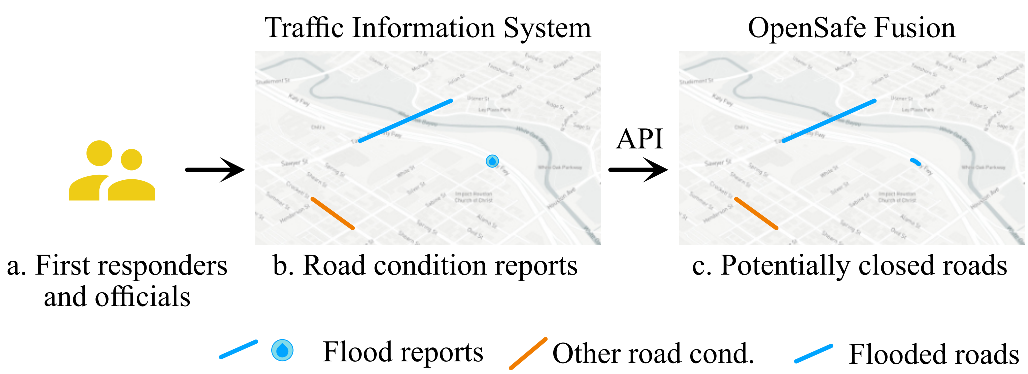

Figure 3: OpenSafe Fusion uses API calls to collect road condition data from DOT alerts. Typically, DOT alerts contain geolocated data on roadway conditions which can be used directly in OpenSafe Fusion with minimal or no processing. (Maps © Mapbox)

flooded roads, traffic congestion, accidents, faulty equipment, stagnant traffic, or special events, could also cause speed reduction. Consequently, relying on traffic speed to detect flooded roads could result in erroneous detection.

To demonstrate the OpenSafe Fusion methodology, Fig 4 shows real-time traffic speed data and OpenSafe Fusion road conditions for two time-steps—5 am and 7 pm. At 5 am, OpenSafe Fusion reports two flooded roads (c and d). While slow traffic speed at links a and b might suggest flooding, OpenSafe Fusion did not consider this observation in its calculation. At 7 pm, the traffic speed at road links a, b, and d returned to normal, indicating a transition to normal condition. Accordingly, OpenSafe Fusion now reports links a, b, and d as likely open to traffic.

2.3.3. Sensors

Sensors deployed along streams and roads provide point estimates of water level at the deployed location. Many gages operated by public agencies such as the United States Geological Survey (USGS) are easily accessible via API or web dashboards. For sensors located along roads, the water level estimates can be directly used to infer the road condition. For sensors situated away from roads, such as water level sensors deployed along rivers, sensing the state of nearby streets requires additional processing. Fig. 5 and Equation 12 illustrates the methodology used by OpenSafe Fusion to convert point estimates at sensor locations to areal estimates to facilitate the identification of roadway conditions. The sensor data processing workflow presented here is inspired from bathtub flood models [66].

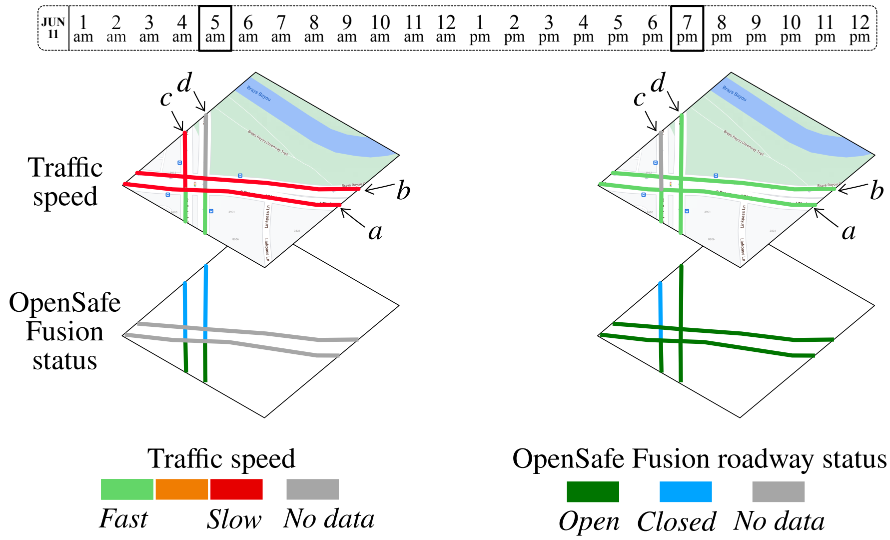

Figure 4: OpenSafe Fusion uses real-time highway speed data to sense the opening of flooded roads. (Maps © Google LLC)

First, the digital surface model (DSM) for the region around the sensor location is collected. DSM is a digital representation of the terrain and contains elevation data of infrastructure elements such as roads and bridges. Water level data from the sensor is gathered during real-time operation and used to construct a constant water surface elevation raster (WSE) in the same datum as the DSM data. A new raster depth map is produced by subtracting the DSM from the WSE map; any places with positive depth values are likely to be flooded. Fig 5b shows an example illustration of the water depth map corresponding to water level 1 in Fig 5a.

All cells with a positive depth value might not be flooded, as indicated by Fig 5c. Here, the presence of a levee protects the right bank from inundation. To account for such situations, OpenSafe Fusion only considers cells with positive water depths that are also contiguous with the location of the water level sensor. The proposed methodology yielded reliable results in our limited testing, especially for inferring the water depth for regions closer to the sensor location. As we move away from the sensor location, the ability of the model to predict water depth reduces. The reduction in predictive ability depends on factors such as water depth and topography. Consequently, this approximate method should only be applied to regions close to the sensor location. Fig 5 uses the four distances R_r, R_l, R_u, and R_d to describe this region. Here, R_r and R_l are the offset towards the right and left banks, and R_u and R_d are the buffers towards the upstream and downstream sides of the sensor location. Historical flood inundation data or results from flood models can be used to estimate the optimal buffer distances for each gage location. This method is only used to detect flooded road ($D_d^l > 0$) and is not used to identify open roads (i.e., $D_d^l = 0$ is neglected).

$$D_d^l = \begin{cases} d - d_s^l, & \text{if } d - d_s^l \in \mathbb{R}^+ \text{and } l \in C^* \text{and } l \in S_{rlud} \\ 0, & \text{otherwise} \end{cases} \tag{12}$$

where:

l	= a raster cell location defined by latitude and longitude
d	= water level reading at the sensor
D_d^l	= water depth at location l due to water level d
d_s^l	= elevation at location l from digital surface model
$\mathbb{R}^+$	= positive real number
C^*	= region contiguous with the sensor location
S_{rlud}	= region bounded by R_r, R_l, R_u, R_d distances from the sensor

2.3.4. Social Media

Past studies have shown that social media analytics can detect flooding, track flood impacts, and sense community response to flooding [67–69]. Several automated workflows [68] exist in the literature to process social media data to sense urban flooding. Following existing literature, OpenSafe Fusion adopts a five-step workflow to glean information on flood conditions in the study area. First, OpenSafe Fusion collects relevant tweets from Twitter using Twitter API. Search queries include flood impacts keywords (e.g., flood, road flooded), event-specific keywords (e.g., Harvey, Ike), location-specific keywords (e.g., Houston, Bayou City), and location constraints (e.g., latitude and longitude of Houston). All collected tweets are then passed through a deep learning-based natural language processing classifier trained to filter relevant tweets. A relevant tweet is a text that contains information on flooding or flood impacts on communities suitable for informing situational awareness. Filtered tweets are then passed through a deep learning model trained to identify entities. For this study, entities are primarily real-world geographical features (e.g., addresses, roads, places). Tweets with

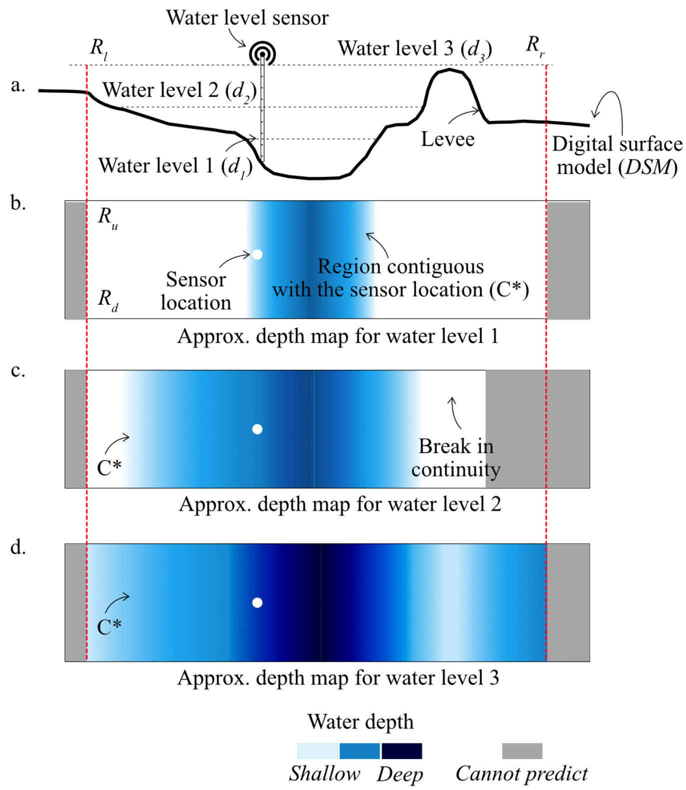

Figure 5: OpenSafe Fusion methodology for identifying flooded regions from sensor data.

identified entities are then geolocated using geocoding techniques [70, 71]. Finally, geocoded tweets are passed through another suite of models that extracts relevant attributes from the text. Relevant attributes include the intensity of flood impacts, time of flood report, and flood depth data. The extracted attributes are then assigned to the corresponding geolocated tweets and mapped on a web interface.

Existing datasets and models are primarily suited to identify entities such as standardized street addresses. Consequently, current models have limited skill in extracting information related to roads. Limited skill in identifying flooded roads necessitates deploying approximate methods to sense road conditions from geolocated flood condition reports. For example, if the following conditions are met, OpenSafe Fusion will mark a road flooded: 1) the report is within a buffer distance of the road; 2) the roadway is at a lower elevation

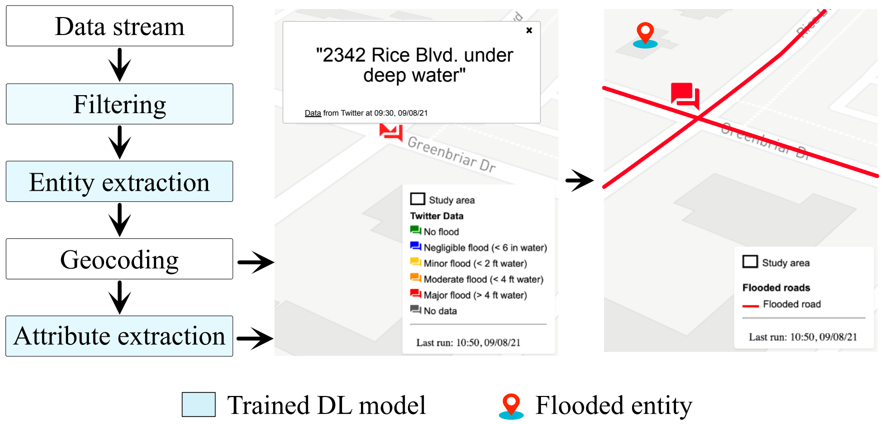

Figure 6: OpenSafe Fusion methodology for collecting and processing social media data to identify flooded roads. (Maps © Mapbox)

than the reported location; and 3) the flooding at the reported location is severe. Similarly, OpenSafe Fusion uses geolocated tweets to identify open roads if conditions 1 and 2 are met, and the tweet reports dry conditions at the location. While automated pipelines that use natural language processing are often noisy and prone to misinformation from malicious or misinformed actors, they serve as an inexpensive source with high availability in urban regions with high social media activity. The precision and dependability of flood mapping using social media can be improved by combining social media data with human-in-the-loop frameworks (see Section 2.3.9).

2.3.5. Traffic Cameras

Many urban areas have live traffic cameras along major highways and busy intersections. Live video or image feeds from these cameras enable traffic management agencies to monitor highway conditions. Such cameras are often in the public domain and can be accessed via a website or API. For example, Houston TranStar [39] operates and publishes data from more than 700 cameras in the Houston region. As observed during past events in Houston, manual inspection of live camera feeds can sense road conditions. While manual sensing of flooding from cameras might be accurate, it is often not practical or scalable. OpenSafe Fusion proposes a framework for automated sensing of flooded roads from camera images using deep learning models. A new dataset especially annotated to sense roadway flooding is developed and deep learning architectures are used to create a robust image

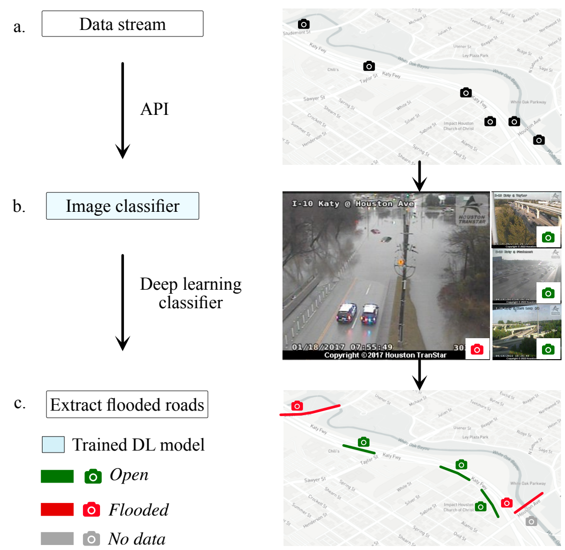

Figure 7: OpenSafe Fusion methodology for identifying flooded roads from traffic camera data (image courtesy of Houston TranStar). (Maps © Mapbox)

classifier capable of predicting flood conditions from camera images. During real-time operation, live traffic camera data is collected at regular intervals (e.g., 10 min). The images are then processed by a deep learning-based image classifier trained to infer the flood condition captured in the image. Flood conditions from the images are then used to identify the status of roads linked to the traffic camera. For example, detecting a severe flood condition on the camera data in Fig. 7b might suggest flooding on I-10 at Houston Ave.

2.3.6. Physics-Based Models

Real-time analysis using physics-based flood models can enable reliable road condition sensing. For example, in regions with radar or rain gage coverage, the OpenSafe Mobility framework [72, 73] (Fig. 8) can provide real-time estimates of flood depth at roads. OpenSafe Mobility collects real-time rainfall radar data from reliable sources (Fig. 8a) such as NEXRAD at frequent intervals. The radar data is then processed to identify flood-inducing rainfall conditions. A flood-inducing rainfall [73, 74] is a rainfall event that could initiate flooding in the study region. Once the rainfall exceeds any flood-inducing rainfall thresholds, radar data at discrete time steps within a maximum considered duration (d_{max}) are concatenated to

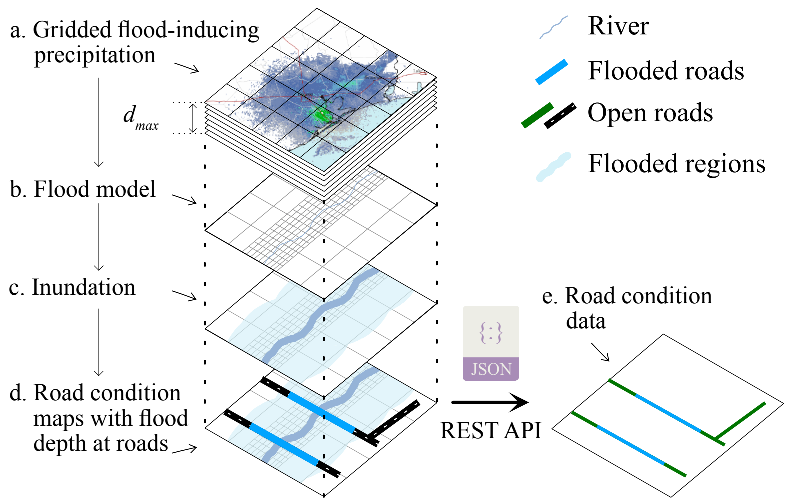

Figure 8: OpenSafe Mobility methodology for identifying flooded roads.

generate a rainfall event. The maximum considered duration is selected after accounting for factors such as the model runtime, acceptable time lag, and available computational resources. The rainfall event is then simulated in a calibrated and validated flood model (Fig. 8b), which routes the rainfall over a digital representation of the study region and estimates the current water surface elevation (WSE) (Fig. 8c). The WSE map and roadway elevation from LiDAR data are then used to estimate the flood depth at road links (Fig. 8d). Flood depth and flow velocity at roads can then be used to assess the trafficability of a road link considering vehicle characteristics such as the safe wading height or stability requirements. Finally, the road conditions are communicated to stakeholders via a website or through REST API.

2.3.7. Crowdsourcing

Several recent studies [24, 25] have demonstrated the effectiveness of crowdsourcing as a medium for collecting real-time flood observations, particularly during severe flood events in urban areas. For example, many ad hoc crowdsourcing platforms [25, 75] were active during Hurricane Harvey in Houston to address the unmet need for situational awareness data. OpenSafe Fusion leverages crowdsourcing as one of the data sources for three reasons: it provides an alternative data source in urban regions; it facilitates communication between users (e.g., first responders active in the field); and it enables stakeholders to overwrite inaccurate predictions from the model.

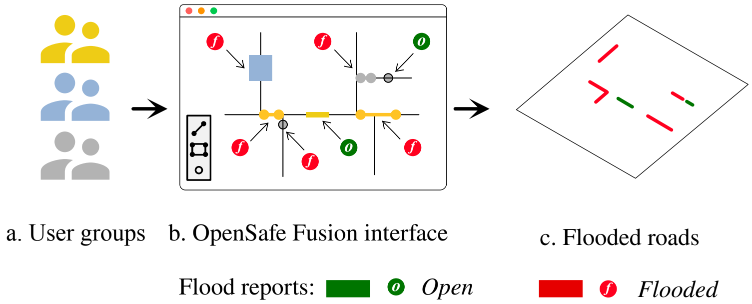

Figure 9: OpenSafe Fusion methodology for collecting and processing crowdsourcing data.

Figure 9 shows an example workflow adopted by OpenSafe Fusion to collect and process crowdsourcing data. To ensure data trustworthiness and prevent misinformation from malicious or misinformed actors, OpenSafe Fusion divides its user group into three different credibility categories: high, medium, and unknown. The high credibility group comprises known first responders (e.g., police officers and FEMA search and rescue team) and officials from organizations responsible for managing flood response (e.g., Houston TranStar). The medium credibility group comprises registered and verified platform users (e.g., city officials and community stakeholders) with a track record of reliable reporting during past events. The unknown credibility group comprises all other users not covered in the first two categories. During data fusion, observations from the high credibility group are assigned the highest importance, followed by the medium and unknown credibility groups. During operation, users can mark the current condition of roads or regions by drawing shapes on the map using interactive draw tools. Example geometry includes points (e.g., flooded intersections), lines (e.g., open roads), and polygons (e.g., flooded neighborhoods). Further, users could also provide auxiliary data describing each report. The auxiliary data could include information such as flood conditions (flooded or open), flood depth, and comments from users. Finally, OpenSafe Fusion uses the user-generated shapes to infer road conditions.

2.3.8. Citizen Service Portals

Many urban regions are equipped with citizen service portals (e.g., the City of Houston 311 system [76]), where residents can report problems such as flooding. The citizen service portal reports are usually associated with the issue report time, closed time, a brief description of the problem, and

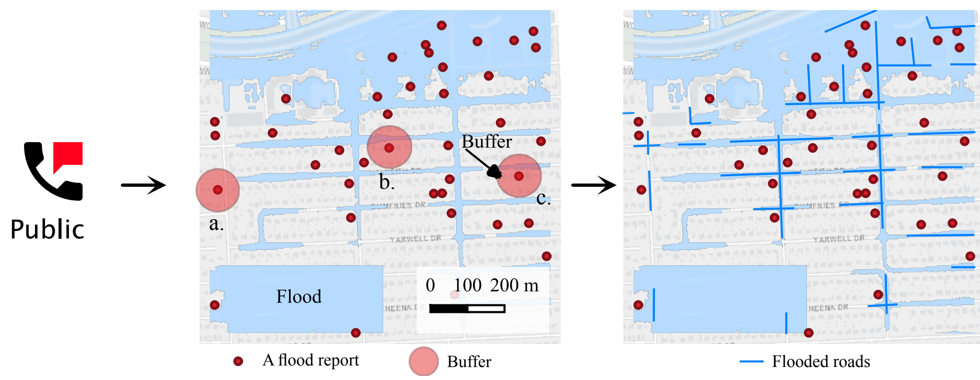

Figure 10: OpenSafe Fusion methodology for collecting and processing data from citizen service portals. (Maps © ESRI)

the required service location. The service locations are most often encoded using a standard street address. Comparing past reports with flood hindcast inundation map indicate that the flooding was often localized to the adjacent streets, and the encoded residential property was not flooded at any point during the storm. For example, Figure 10 compares CoH 311 flood reports to an inundation map for Hurricane Harvey. Here, many reported parcel locations were often not flooded, but the adjacent roads were flooded primarily due to their lower elevation compared to the adjoining parcels.

Figure 10 illustrates OpenSafe Fusion methodology for identifying flooded roads from citizen service portal reports. OpenSafe Fusion marks all streets within a buffer distance (e.g., points a, b, c) of a flood report flooded. To acknowledge uncertainty, OpenSafe Fusion assigns a confidence value to these observations. For example, the probability of a road link flooding given a flood observation within a predefined buffer distance of 100 m is 85 percent. Historical flood reports and hindcast flood maps can be used to determine the buffer distance and the corresponding confidence value. While flood sensing using citizen service requests lacks specificity, the reports in the presence of observations from other sources might provide better sensing of flooded entities in a data fusion framework.

2.3.9. Human-in-the-Loop

Real-time automated data processing for sensing, mapping, and tracking floods to guide emergency response decision-making is a high-risk application. Any mistakes in model prediction will expose first responders and evacuees to possible safety risks and cause delays and detours that limit emergency response efficiency. In the long term, model errors will impact stakeholder trust in the framework leading to reduced use and continued mistrust. The unproven generalizability of machine learning and automated models—often trained on limited historical data—on unseen new events in high-risk scenarios necessitates substantial safety measures to limit risk to stakeholders. In the short term, while visible disclaimers and acknowledgment of uncertainty in model predictions might improve stakeholder trust, they might increase the cognitive overload of first responders in stressful conditions.

To partially address the need to ensure prediction quality, OpenSafe Fusion adopts a human-in-the-loop strategy (Fig. 11). Here, a group of trained human agents monitors the performance of different data processing workflows. The OpenSafe Fusion framework assigns a confidence score to observations from data processing workflows to facilitate review prioritization. The confidence score ranges from 0 to 1, with higher values indicating more reliable predictions. Three methodologies are used by OpenSafe Fusion to assign confidence scores. First, physics-based constraints imposed by the study region's topography are employed to detect potentially inaccurate observations (see Sec. 2.4 for more details). Consider two adjacent and connected roads on sloping terrain. If the road at a higher elevation is observed flooded, the road at a lower elevation is most likely be flooded. If observations from data sources contradict physical constraints imposed by terrain, OpenSafe Fusion will automatically assign low confidence scores for the observations and tag the observation for review. Second, performance metrics inherent to mathematical models are used to assign confidence scores. Example metrics include model accuracy or F1-score for classification models (for deep learning framework used to identify flooded roads from live camera images) and RMSE or MAE for models estimating water depth. Third, the historical performance of the data processing workflows (e.g., flood models are more accurate near bayous compared to regions away from bayous) is used to assign confidence scores. In summary, the assigned confidence score depends on the expected model performance considering environmental, technical, and other factors influencing model predictions. To further facilitate review prioritization, high-impact observations are identified by considering both confidence scores and the population density of the report location.

Reviewers can rectify any inaccurate predictions by using the crowdsourcing capabilities offered by OpenSafe Fusion. Additionally, human oversight can monitor the model's performance in real-time and disable or modify the confidence of data processing workflows whose accuracy is subpar. It is

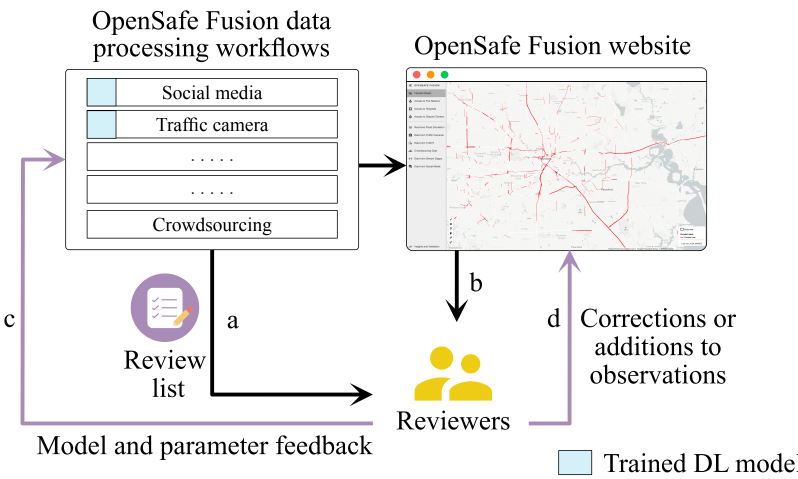

Figure 11: Conceptual human-in-the-loop framework for enhancing the accuracy of OpenSafe Fusion. (Maps © Mapbox)

crucial to highlight that OpenSafe Fusion already considers the accuracy of observations during the fusion process (see Sec. 2.2). The human-in-the-loop strategy provides an additional opportunity to augment existing data for better predictions. Further, the human-in-the-loop component is intended to be operated by emergency response managers and coordinators at command and control centers and not by field personnel to prevent cognitive overload. Finally, the human-in-the-loop is optional; OpenSafe Fusion can sense current conditions without human supervision.

2.4. Data Augmentation

Direct flood observations may be sparse. Depending only on sparse observations may limit the efficacy of OpenSafe Fusion. A possible strategy to augment data availability is to leverage existing observations in the context of the region's topography to infer the status of roads with no direct road condition data. Figure 12 illustrates some example scenarios. In scenario s-1, road link a is observed flooded while conditions of roads b and c are unknown. Given the topography (mean elevation and slope) of the connected roads, link b is likely to be inundated as link a is flooded (one-step logical deduction). While link c lacks observations for its surrounding roads, once the state of link b is inferred, the possible state of link c can be deducted (two-step logical deduction). Similarly, iterative logical reasoning can be used to infer the states of additional road links, frequently at the expense of accuracy. It is ideal to limit data augmentation to only one step to ensure accuracy.

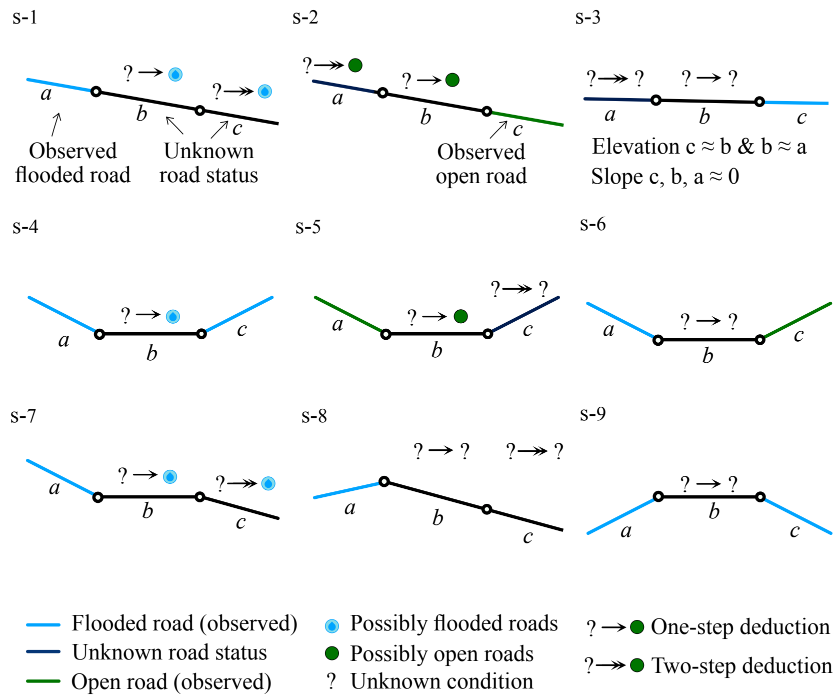

Figure 12: Example data augmentation scenarios for select roadway profiles.

Using logical deduction is not always possible for all road links. Consider scenario s-8; though link a is flooded, the status of links b and c cannot be reliably inferred due to the presence of a ridge. Similarly, in s-3, the status of links a and b can only be reliably estimated if significant flooding is reported at link c (to account for any localized flooding of link c). Further, data augmentation via deduction could occasionally lead to contradictions. For example, in s-6, link b is both flooded (as determined by the condition of link a) and open (based on link c). This contradiction could imply the failure of logical deduction for link b or point to inaccuracy in existing observations for either link a or b. OpenSafe Fusion will tag these roads for further review by a human agent. The data augmentation methodology is summarized in Equation 13. It is critical that the data augmentation approach presented here is not employed for scenarios involving long road links or roads in flat terrain. Additionally, a road is only deemed open if its full stretch is dry;

otherwise, errors could occur in instances such as s-5. Finally, while DSM data are used for inferring road conditions from sensor data (Section 2.3.3) and for data augmentation, the data processing workflows, input data needs, and application criteria differ (see Equations 12 and 13).

$$C^k(\mathcal{R}, \delta\mathcal{L}, \delta\mathcal{D}) = \begin{cases} likely\ flooded, & \text{if } \mathcal{R}^+_\searrow \notin \{o\}\ \&|\mathcal{R}^+_\searrow \in f| > 0\ \&\mathcal{R}^-_\searrow \notin \{o\} \\ likely\ open, & \text{if } \mathcal{R}^+_\searrow \notin \{f\}\ \&|\mathcal{R}^+_\searrow \in o| > 0 \\ unknown, & \text{otherwise} \end{cases} \tag{13}$$

where:

C^k = condition of the road link k

$\mathcal{R}$ = a set of all links connected to the link k. The links must have an elevation difference of at least δD and a maximum length of δL.

$\mathcal{R}^+_\searrow$ = a set of links at a higher elevation than the link k and sloping towards the link k. $\mathcal{R}^+_\searrow \in \mathcal{R}$.

$\mathcal{R}^-_\searrow$ = a set of links with lower elevation and sloping away from the link k. $\mathcal{R}^-_\searrow \in \mathcal{R}$.

f = a set of all roads flooded in the current time step.

o = a set of all opened roads in the current time step.

2.5. Network Analysis

Information on flooded roadways alone does not provide a comprehensive view of flood impacts. Factors such as network topology and the location of facilities could influence network robustness (defined here as the ability to maintain connectivity between communities and critical facilities). Consequently, quantifying the network-level impacts of flooding via real-time network analysis is essential to provide a holistic view of flood impacts to support decision-making and to prioritize emergency response.

OpenSafe Fusion represents the topology of a road network as graph $G = (V, E)$. Here, V is a set of nodes modeling points of interest, such as access locations or roadway intersections, and E is a set of road links connecting nodes. For a specific critical facility group k (e.g., all hospitals), baseline connection between every node in the network and the nearest facility is assessed. $D^n_{x\to k}$ denotes the shortest distance (measured in route length) in the original road network between a node x and the nearest facility in k (e.g., the nearest hospital). During operation, OpenSafe Fusion identifies impassable links (v_t^f) and inundated nodes (e_t^f) at every time step. The flooded entities are then removed to create an updated road network $G_t^f = (V_t, E_t)$, where $V_t = (V - v_t^f)$ and $E_t = (E - e_t^f)$ at time t. The shortest distance ($D_{x \to k}^t$) between node x to the nearest facility in k at time t is then estimated. Further, the connectivity loss ($CL_{x \to k}^t$) ratio [77], defined as $1 - D_{x \to k}^n / D_{x \to k}^t$ for facility k and node x at time t, is utilized to quantify flood impacts on access to the facility group k. $CL_{x \to k}^t$ ratio varies between 0 (no impact of flooding on the network access) and 1 (complete loss of access). Finally, the node-level results can be aggregated at a geographical unit level, such as Census Tracts, to visualize the spatial distribution of flood impacts on access to each facility type. Connectivity loss maps can be generated for various critical facilities (e.g., fire stations, pharmacies, and dialysis centers) to enhance situational awareness and aid decision-making.

2.6. Publishing

Stakeholders have access to four categories of data through the OpenSafe Fusion framework: observations from data sources, road condition data after data fusion, road condition data after data augmentation, and network-level flooding impacts. Observations from individual data sources enable stakeholders to verify OpenSafe Fusion results. Road condition data can be used for routing. Network-level flood impacts help identify isolated neighborhoods, prioritize emergency response, and support decision-making. The OpenSafe Fusion results could be published via web-based tools built following the tenets of user-centered design [11] to address the needs and preferences of stakeholders. Further, OpenSafe Fusion results should also be made available via REST API to facilitate interoperability with existing situational awareness and decision-making tools.

3. Case Study Evaluation

This section presents results from case study experiments designed to evaluate the OpenSafe Fusion framework for its strengths and limitations. A limited case study deployment of the framework is developed for Houston, Texas. Data sources in the study region are analyzed, and OpenSafe Fusion workflows are created. The OpenSafe Fusion framework is evaluated by reenacting Hurricane Harvey (2017). OpenSafe Fusion model predictions are compared to ground observations during enactment to quantify model performance. The following subsections describe the experiments in detail.

3.1. Study Area

Houston, Texas, (Fig 13) is the fourth most populous city in the United States. Houston is prone to recurring urban flooding due to several factors, including its location in the hurricane-prone Gulf of Mexico, flat topography with few relief features, urban sprawl, lack of zoning laws, limited stormwater drainage capacity, and soil conditions [78]. High flood hazard was evident during recent storm events such as Memorial Day Flood (2015), Tax Day Flood (2016), Memorial Day Flood (2016), Hurricane Harvey (2017), Tropical Storm Imelda (2019), and Tropical Storm Beta (2020). Dong et al. [79] demonstrated that even minor flooding in Houston could trigger network-wide catastrophic capacity reduction due to compound failures. While flooding causes network failures, its impacts are exacerbated by the limited information on road conditions during a flood event. Flooding and a lack of situational awareness reduce safety and efficiency during emergency response and mobility during flooding. For example, 21 of the 57 drowning fatalities during Hurricane Harvey in Houston are linked to vehicle use [80].

While flood mitigation studies are required to reduce Houston's flood hazard, increased availability of situational awareness data can improve roadway safety and emergency response efficiency in Houston. Although Houston has several real-time data sources, they are not organized in a unified framework to enhance situational awareness. This study evaluates the OpenSafe Fusion framework's capacity to monitor flood impacts on roads by leveraging data sources varying in data types, accuracy, and reliability. Any improvement in situational awareness could help responders identify flooded roads and affected communities improving the safety and efficiency of emergency response. Recurring flooding and the availability of real-time data sources make Houston an ideal testbed for OpenSafe Fusion.

3.2. Hurricane Harvey

Hurricane Harvey (2017) is reenacted in OpenSafe Fusion to critically assess its effectiveness. Hurricane Harvey (25 August to 2 September 2017) brought record-breaking rainfall to Harris County. The Houston metro area saw rainfall amounts totaling 36-48 inches. As a result of this slow-moving storm, more than 122,000 people were rescued by emergency responders [17]. Additionally, roadways throughout Houston were flooded, including major highways such as I-10, I-45, and US-59. NOAA estimates damages from Harvey at around $125 billion, making it the second costliest tropical cyclone in

Figure 13: Houston, Texas is used to demonstrate OpenSafe Fusion. (Maps © ESRI)

the United States, next to Hurricane Katrina (2005). The lack of real-time information about roadway conditions was especially detrimental to emergency response efficiency and safety. For example, two ad hoc projects [25, 75] implemented by community members to share roadway status had more than a million map views. Experiences during Hurricane Harvey further highlight the need for reliable mobility-centric situational awareness tools in Houston.

3.3. Data Sources and Data Processing Workflows

This study identifies eight public data sources that observe floods in real time, either directly or indirectly. The identified data sources are: (1) Texas Department of Transportation DriveTexas [40]; (2) Houston 311 database [76]; (3) OpenSafe Mobility [73]; (4) U-Flood crowdsourcing [25]; (5) Gage data from USGS [81]; (6) Houston TranStar traffic camera network [39]; (7) Real-time traffic speed data from Houston TranStar [39], and (8) Twitter data [23, 82]. A majority of these data sources were active during Hurricane Harvey. An exception is the OpenSafe Mobility framework, which was created in response to the need for better mobility-centric situational awareness tools. It is included here to demonstrate its capability and compare it to other data sources. A summary of the characteristics of different data sources selected for this case study application is provided in Table 1. Screenshots from select data sources used in this study are shown in Fig 14.

After identifying the data sources, automated source-specific data processing procedures are developed for each data source. These data processing algorithms use a variety of approaches, including deep learning and spatial analysis, to determine present flood conditions and, consequently, flood impacts on roads. The remainder of this subsection presents an overview of the data sources and data processing workflows.

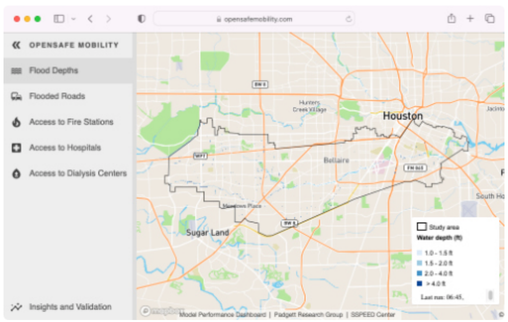

a. OpenSafe Mobility website

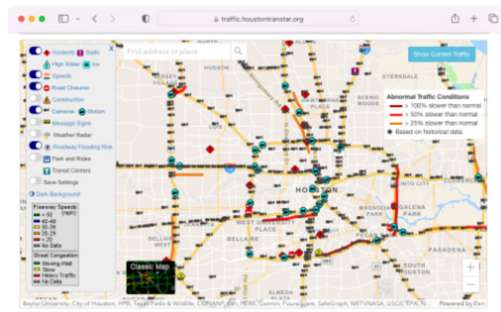

b. Houston TranStar website

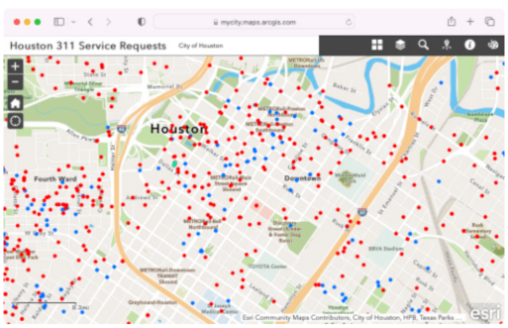

c. CoH 311 Citizen Service Portal

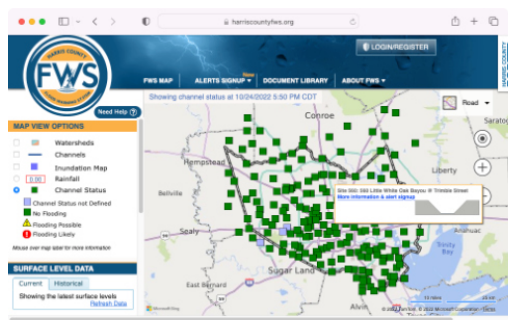

d. Gage data from HCFCD

Figure 14: Screenshots from select data sources used in this study. (Images courtesy of © Houston TranStar, City of Houston, Harris County Flood Control District, Mapbox)

Table 1: List of data sources used in this case study and their characteristics. All data sources except OpenSafe Mobility were available during Hurricane Harvey. OpenSafe Mobility was created in response to the need for a mobility-centric situational awareness framework during Hurricane Harvey.

Data source	**Obs. type**[1]	**Data type**[2]	**Transition**[3]	**Availability**[4]	**Delay**[5]	**Accuracy**[6]	**Bias**[7]	**Cost**[8]	**Sensor type**[9]
DriveTexas	■■■	○	□□	⋆	⋄⋄	✓✓✓	-	$	Authoritative
Houston 311	■	○	□□	⋆⋆	⋄⋄⋄	✓	§	$	Crowdsourcing
U-Flood	■■■	○	□	⋆⋆⋆⋆	⋄⋄⋄	✓	§	$$	Crowdsourcing
Gage data	■	○○○	□□	⋆	⋄	✓✓	-	$	Physical sensor
OpenSafe Mobility	■■■	○○○○	□□	⋆⋆⋆⋆	⋄⋄	✓✓	-	$	Physics-based
Traffic data	■■■	○	□□	⋆⋆⋆	⋄	✓	-	$	Physical sensor
Traffic cameras	■■	○○○○	□□	⋆	⋄	✓✓	-	$	Physical sensor
Twitter	■■	○○	□	⋆⋆⋆⋆	⋄⋄	✓	§	$$	Crowdsourcing

[1] **Observation type**: ■ = No direct observation of road conditions is usually available, and post-processing is required to infer road conditions. ;■■= Road conditions are usually inferred, but direct observations are sometimes available. ;■■■ = These sources directly observe flooding on roads. [2] **Data type**: ○ = Data sources usually report binary roadway status (flooded/open); ○○ = Usually, binary status is reported, but sometimes flood depth at roads is available; ○○○ = Water depth at roads is always available; ○○○○ = Flood depth and flow velocity at road links are available. [3] **Transition**: □ = Typically transition from open to flooding is reported, but the change from flooded to normal condition is not reported; □□ = State transitions from open to flooded and from flooded to open are continuously reported. [4] **Spatial availability**: ⋆ = Low spatial availability (usually only for major roadways or limited by the availability of sensors); ⋆⋆ = Moderate data availability (usually available for arterial links); ⋆⋆⋆ = Hight data availability (usually available for collector streets); ⋆⋆⋆⋆ = Highest spatial availability (usually available for even residential/local streets).
[5] **Time delay**: ⋄ = Data are typically available instantaneously; ⋄⋄ = Data are usually available without much delay; ⋄⋄⋄ = Data could be delayed significantly. [6] **Accuracy**: ✓ = Low accuracy reports are possible due to several factors; data could contain noise or errors; ✓✓ = Reports are usually accurate but could contain errors; ✓✓✓ = Reports are accurate and validated. [7] **Bias**: §= Data might be biased (e.g., usually available for densely populated regions thus could misrepresent the spatial distribution of flood impacts).
[8] **Cost to acquire data for a future implementation in Houston or a similar region**: $ = Free; $$ = Low; $$$ = Moderate; $$$$ = High. [9] OpenSafe Mobility uses rainfall radar data and a physics-based flood model to infer current flood conditions.

3.3.1. Texas Department of Transportation DriveTexas

In Houston, the TxDOT DriveTexas website provides real-time information on road conditions via the DriveTexas website [83] and through API [65]. Historical road closure data from DriveTexas was collected for Hurricane Harvey and used in this study. A closer examination of TxDOT data reveals that all roads marked closed due to flooding are not flooded. Many roads, such as Interstate-610 loop around Houston, were partially open but marked closed to the public. Further, the DriveTexas platform only reports road conditions for TxDOT-maintained roads. This limits the data availability to major roads such as Interstates, US and State Highways, and Fram-to- and Ranch-to-Market roads. Non-TxDOT maintained roads include roads maintained by the city or county, including frontage roads and several arterial, collector and local streets. Thus DriveTexas will not report the road conditions of several roads essential for urban mobility. TxDOT DriveTexas API provides georeferenced road condition data. While OpenSafe Fusion uses OpenStreetMap road data, DriveTexas uses a different road dataset, thus necessitating a mapping function. This study maps DriveTexas condition data to OpenSafe Fusion data by matching location (within a 30m margin), road name, and orientation. In limited testing, this mapping logic identified the correct mapping in most cases.

3.3.2. Crowdsourcing

During Hurricane Harvey, multiple citizen-led crowdsourcing tools were deployed to address the unmet need for situational awareness data. Of the ad hoc tools, U-Flood [25] was focused on real-time information on flooded streets. U-Flood enabled the public to share information on flooded roads by marking roadway status on a web dashboard built using Mapbox and OpenStreetMap. During its operation, U-Flood saw more than 1 million map loads. User-generated content from U-Flood during Hurricane Harvey is used here to model crowdsourcing data. A closer look at the data reveals two significant findings. First, data on local roads and residential streets are overrepresented, complementing sources that primarily report on the status of main highways. Second, while many individuals report flooded roads, the number of reports indicating the transition from flooded to open state is rare. Hence, flood reports quickly become untrustworthy in dynamic flooding scenarios where road conditions rapidly evolve.

Past studies have also highlighted that social sensors, such as crowdsourcing data, are prone to misinformation due to malicious or misinformed actors. For example, Sebastian et al. [78] observed the presence of fake flood reports in social sensors during Hurricane Harvey. Similarly, Praharaj et al. [84] reported that only 71.7% of the crowdsourced Waze flood incident data was trustworthy in a Norfolk, Virginia case study. Thus, additional measures such as verifying crowdsourcing observations using a human-in-the-loop strategy and dividing user groups into trust categories might help improve the reliability of crowdsourcing data.

3.3.3. Traffic Speed Data

The Anonymous Wireless Address Matching (AWAM) system of Houston TranStar [39] employs multiple roadside AWAM readers. These readers sense the MAC address from Bluetooth-enabled devices such as cellular phones, mobile GPS systems, and in-vehicle navigation systems as they pass the reader station. The report times of a device at successive AWAM readers are used to estimate the roadway segment's average travel time and speed. The Houston TranStar Speed Map archive was used to acquire historical traffic data for this study. Houston TranStar has maintained a database of 15-min average speeds for 485 freeway links in Houston since January 2009. Houston TranStar also provides API access to the traffic speed data for real-time applications.

3.3.4. Sensors

Houston is amongst the most extensively gaged region in the US, with more than 50 gages in the study region. The USGS and the Harris County Flood Control District (HCFCD) are the primary operators of these gages. USGS offers API access to real-time and historical data, whereas HCFCD data is only available through a web dashboard, necessitating web scraping. Data from 40 USGS-operated gages were used in this investigation due to their ease of access. Following Section 2.3.3, historical gage data for selected gages are collected and processed to estimate flood extents. Flood extents are then used to estimate water depth at roads; roads with a depth of greater than 50 cm are considered flooded in this study.

3.3.5. Citizen Service Portals

This study uses historical reports from the City of Houston (CoH) 311 citizen service portal to identify flooded regions. Flood reports from Hurricane Harvey are collected and geolocated. As described in Section 2.3.8, flood reports are encoded using the standard street address in CoH 311 data, thus preventing the accurate localization of the reported condition. At each time step, all roads located within a buffer of 30 m (100 ft) of an active flood report are considered flooded in this study.

3.3.6. OpenSafe Mobility

OpenSafe Mobility [73] is a mobility-centric situational awareness system that uses real-time radar data and a physics-based flood model to identify flooded roads. A version of the OpenSafe Mobility framework has been operational since September 2021 for the Brays Bayou Watershed area in Houston, Texas. For this study, OpenSafe Mobility is expanded to include other watersheds in the Houston region. The newly considered regions include a) Greens and Hunting Bayou Watersheds; b) Sims and Vince Bayou Watersheds; c) White Oak Bayou Watershed; and d) Buffalo Bayou Watershed. New physics-based flood models are developed and calibrated for each region using historical rainfall from Tax Day Flood (2016). Together the five models (one pre-existing and four newly developed models) cover most of the study area, thereby significantly improving the data availability. Historical rainfall radar data are used in this study to reenact model outputs for Hurricane Harvey.

3.3.7. Traffic Cameras

Houston TranStar [39] operates more than 700 live traffic cameras. An automated deep learning model that can sense road conditions from traffic cameras can significantly improve data availability, especially for major roadways. Existing labeled image datasets are either limited in size or unsuitable for inferring road conditions from low-resolution traffic cameras. The lack of relevant annotated data necessitated the development of an image classifier from scratch. This study collected and labeled 2300 images related to roadway flood conditions. Flooded images are collected from various sources, including traffic camera images, Flickr, Bing, Google search, Twitter and others. The collected images are then manually inspected to filter images featuring roads—either flooded or open. The shortlisted images are then annotated using Supervise.ly annotation platform. Two classes are considered while annotating images. The considered classes are a) roads either not flooded or with minor flood and passable to most vehicles and b) flooded roads that could pose unsafe road conditions. The annotated images are then manually cross-checked to ensure quality. The images are then used to train deep-learning-based image classifiers using transfer learning. The best among the trained models (based on ResNet-34 [85]) can detect open and impassable roads using traffic camera data with 83% accuracy. For this case study, historical traffic camera data are collected for the study region. Due to the delay in data collection and the absence of archived data, data from all Houston TranStar cameras through Hurricane Harvey are not available. The limited images collected (n=15) are used here to demonstrate the application of automated deep learning workflow to sense flooding on roads.

3.3.8. Social Media

Despite recent advances in annotated datasets [86–88] and reliable geocoding tools (e.g., Google Geocoding API), limited testing during this study reveals that more research is required to enable automated identification and mapping of flooded roads and entities from tweets. Specifically, adding social media to OpenSafe Fusion did not significantly improve its accuracy but introduced noise to observations due to the lack of specificity in observations derived from tweets. To elaborate, existing annotated datasets [86–88] can identify informative tweets, classify relevant tweets into preidentified humanitarian categories, and estimate infrastructure damage severity from tweets. However, the datasets cannot estimate flood depth or severity from tweets. Thus, new datasets that can estimate flood depth or severity from tweets are necessary. Further, existing annotated datasets for geographic feature extraction (and geocoding tools) focus on standard street addresses and place names, thus, failing to identify roads as entities reliably. Hence, an entity extraction dataset that can identify roads and other geographic features are necessary. Finally, existing annotated datasets focus on either classification or entity extraction and are not suited for mapping the identified flood impacts to the affected entity. To elaborate, consider the tweet, "Brompton St. South of Holcombe Blvd. is Flooded." While processing this tweet, an entity extractor can identify two entities: Brompton St. and Holcombe Blvd. A tweet classifier can identify that the tweet is related to flooding. However, models trained on existing datasets might not help identify the flooded road section from the two identified entities. Thus, a new joint entity and relation extraction dataset that maps the flood condition to entities is required to facilitate an accurate mapping of flood impacts. Such a dataset should map flood conditions to entities (e.g., *entity::*Brompton St.—*relation::*attribute—*condition::*Flooded) and also help identify the affected portion of the entity (e.g., *entity::*Brompton St.—*relation::*South of—*entity::*Holcombe Blvd.). In summary, a new dataset that can estimate flood depth or severity, identify roads and other entities, and map the relation between entities and flood severity are necessary for leveraging social media data. Since OpenSafe Fusion is intended for emergency response applications, it was decided not to leverage social media data in this case study and initiate the development of datasets that can accurately identify flooded roads from tweets.

3.4. Validation Results

This section reenacts Hurricane Harvey in OpenSafe Fusion to critically evaluate its performance. The main stages of OpenSafe Fusion are illustrated in Figure 15. First, the OpenSafe Fusion model is activated when flood-inducing conditions are detected in the study area. Once activated, OpenSafe Fusion uses the road transportation network of the study region to begin analysis. The road transportation network used in this example is extracted from OpenStreetMap and contains more than 62,000 road links. All major highways and arterial roads are covered, while some residential streets are not considered for this case study. In the beginning, all road links are assigned an initial probability of flooding. In this example, the initial probability of flooding is set at 50% to encode the model's lack of knowledge about the initial state of the roads. Once initialized, OpenSafe Fusion will collect, process, and fuse data at regular intervals. The time interval between runs is set to one hour for this demonstration. For a real-time application, shorter time steps could be used to ensure the recency of model predictions.

During a new time step, previous states of the road, past observations, and external actors can be used to predict the state of the road link in the next time step. Figure 16 shows the average transition probability for roads in Houston during Halloween Day Flood (2015), Memorial Day flood (2015), and Tax Day Flood (2016). Here, OpenSafe Fusion road network and physics-based flood models are used to track link states and estimate the state transition for each time step (Fig 16). In all three cases, the transition probability of an open road remaining open ($P(X_{t+1} = Open|X_t = Open)$) in the next time step (1 hour) is 0.99. The transition probability of flooded roads remaining flooded ($P(X_{t+1} = flooded|X_t = flooded)$) hovers between 0.90 and 0.99 (mean transition probability is 0.97 for all events). While some fluctuations can be observed for transition probability for flooded to flooded transitions in the early stages of flooding, the value quickly converges to 0.97. Insights from the three past events indicate that the Prediction step can be approximately modeled as a Markov Process, especially for Hurricane Harvey, as it was a slow-moving flood event. This study uses two Prediction

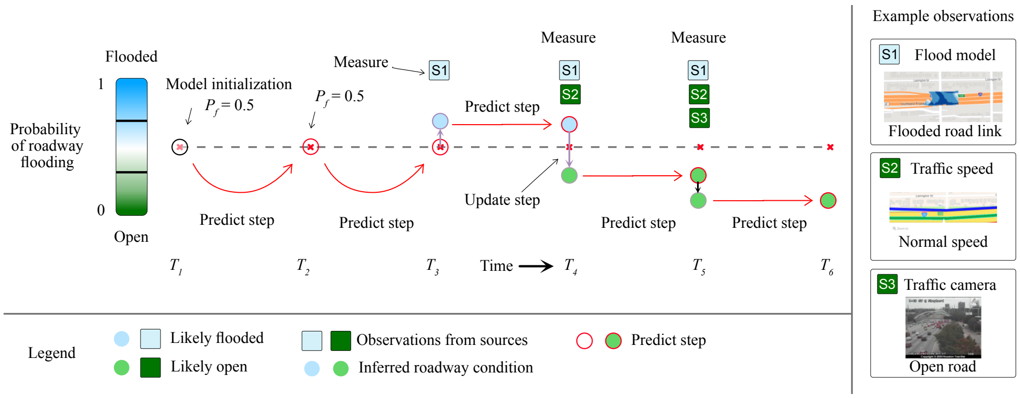

Figure 15: Prediction, measurement, and update steps for a road link in OpenSafe Fusion. The model is initialized at time step T1 with an initial probability of the road link flooding set at 50%, encoding the lack of information on roadway status. At T2, the model maintains the initial belief since no observation was received. After obtaining a flood observation from the OpenSafe Mobility flood model, the model believes the link may be flooded at step T3. OpenSafe Fusion sees typical traffic speeds at the link at T4, and it now updates its belief to a likely open road. At T5, OpenSafe Fusion receives more evidence from a traffic camera that the road is open, leading to an updated belief that the link is probably open. (Images courtesy of © Houston TranStar, Google LLC, Mapbox)

models (Table 2): P1 and P2. A road link is initialized with the P1 model as it holds the assigned initial probability of flooding. Once the link is observed, OpenSafe Fusion switches the prediction model to P2. With each time step, Prediction model P2 will move the state of the road closer to the open state.

Next, observations from data sources are collected and processed using the data processing workflows described above. Only the Prediction step is executed if no observations are available during a time step (see time step T2 in Fig 15). If observations are available, data fusion is initiated using the formulation presented in Equation 11. Equation 11 disregards data source interdependencies, overemphasizing simultaneous observations from interdependent sources. In this initial study, sufficient historical data was unavailable to model and study the interdependencies among data sources and their impacts on data fusion accuracy. Future research should investigate interdependencies among data sources and model them if it improves model accuracy. For this case study, four sources (OpenSafe Mobility, Sensors, Traffic Camera, and Citizen Portals) independently observe flooding, while three sources (UFlood, TxDOT, and Twitter) might have dependencies on other sources. Consider, for example, a TxDOT employee reporting flooding

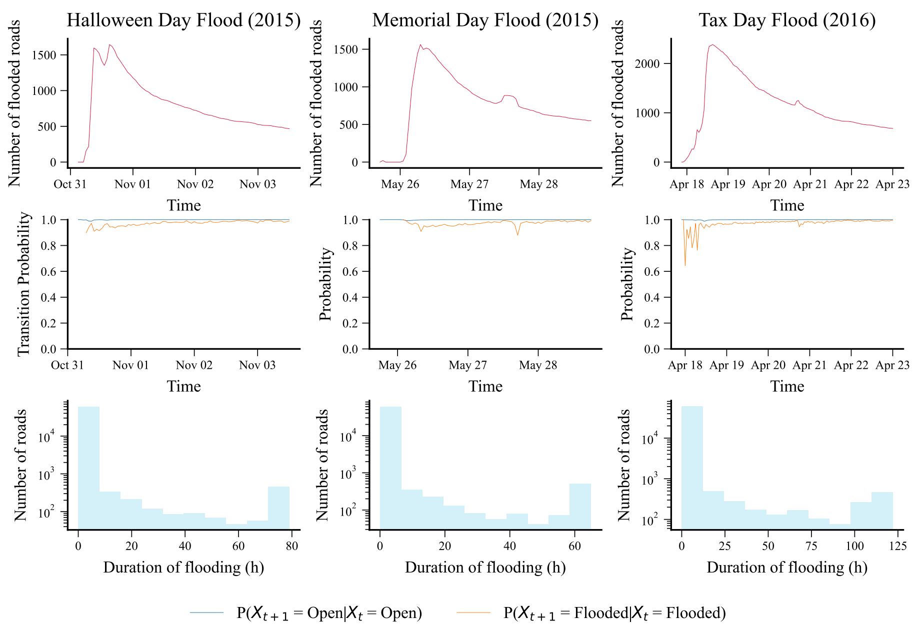

Figure 16: Figures showing the evolution of flood impacts on roads during three recent floods in the study region. Similarity can be observed in the distribution of flooded duration and the temporal evolution of flood impacts on roads (i.e., the number of flooded roads). More importantly, consistent transition probability between flooding states observed in the modeled flood events indicates that a Markov model can be used to model the Prediction step.

Table 2: Model parameters for OpenSafe Fusion Hurricane Harvey case study. Only OpenSafe Mobility and Traffic Camera reports both open and flooded status. While Traffic Speed data only reports open status, the remaining sources only observe flooding.

Model	Model ID	Description
Transition Model	P1	$P(X_{t+1}=f\|X_t=f)=0.99$; $P(X_{t+1}=o\|X_t=o)=0.99$
	P2	$P(X_{t+1}=f\|X_t=f)=0.97$; $P(X_{t+1}=o\|X_t=o)=0.99$
OpenSafe Mobility	OSM-1	$P(z=o\|X=o)=0.90$
		$P(z=f\|X=f)=1/(1+e^{-c1*(wd-c2)})$; c2=2, c1=2
Traffic Camera	CAM-1	$P(z=f\|X=f)=0.83$; $P(z=o\|X=o)=0.83$
Traffic Speed	SPEED-1	$P(z=f\|X=f)=0.95$; $P(z=o\|X=o)=0.95$
TxDOT	TXDOT-1	$P(z=f\|X=f)=0.95$; $P(z=o\|X=o)=0.95$
UFlood	UFLOOD-1	$P(z=f\|X=f)=0.70$; $P(z=o\|X=o)=0.70$
Citizen Portal	COH-1	$P(z=f\|X=f)=0.85$; $P(z=o\|X=o)=0.85$
Sensors	USGS-1	$P(z=f\|X=f)=0.85$; $P(z=o\|X=o)=0.85$
Twitter	TW-0	—

after observing a flooded road from a traffic camera.

Table 2 reports $p(z|x)$ (see Equation 11) for the considered data sources. These models are based on historical data (for Citizen Portals and Sensors), model performance (for camera data), insights from similar studies [84](for U-Flood), design considerations (for TxDOT), or a preliminary informed assumption (for OpenSafe Mobility). For OpenSafe Mobility, the sigmoid function with two parameters is used to model $p(z=f|x=f)$. Leveraging the sigmoid function enables OpenSafe Fusion to dynamically change model confidence based on the predicted flood depth (wd in feet) at roads. Further, the sigmoid formulation also facilitates road-link-specific flood threshold selection to consider potential ponding effects due to numerical errors. After measurement and update, OpenSafe Fusion pauses until the next time step is initiated. The process of prediction, measurement, and update continues with each time step until the stopping criteria is reached (e.g., OpenSafe Fusion detects no flooded road in the study area).

Figure 17 shows the spatial distribution of road condition observations from select sources and OpenSafe Fusion. OpenSafe Mobility, U-Flood, and TxDOT are the three sources that provided the majority of flood observations. While TxDOT and traffic speed observations are primarily for major highways, other sources also offer data on minor streets, thus addressing the need for detecting local road conditions. The reports from CoH 311 data are mainly focused on residential streets, whereas data from gages is centered close to bayous. Since U-Flood was an ad hoc situational awareness tool deployed during Hurricane Harvey, the data is only available starting August 31, 2017. Contrasting OpenSafe Fusion data availability with individual sources indicates that it successfully improved data availability throughout the event, even for minor roads—thus achieving one of the main goals of OpenSafe Fusion. Better data availability can translate to better situational awareness and improved roadway safety.

The effectiveness of data fusion in achieving just situational awareness and overcoming data inequities depends primarily on the availability of reliable observations from multiple data sources. Fig. 17 indicates that OpenSafe Fusion observations are available throughout urban Houston, while other sources exhibit clustering around select neighborhoods (U-Flood; Fig. 17c) or sparse availability (Fig. 17b, e-g) outside major highways or bayous. While fusion can help reduce situational awareness data inequity, it cannot eliminate them entirely (data-rich regions will always have better situational awareness). However, any reduction in situational awareness bias will promote equitable emergency response. With only U-Flood reports, responders might prioritize the observed areas, leading to unjust resource allocation and reduced emergency response efficiency in other communities. In contrast, OpenSafe Fusion enables better sensing for all regions, thus promoting just resource allocation and safer and efficient emergency response navigation. Finally, better characterization of data sources and enhancing the accuracy of OpenSafe Fusion workflows could also enable the framework to offer just situational awareness.

Figure 18 evaluates OpenSafe Fusion performance using ground truth data collected from images showing road conditions (both flooded and open). These images are collected from diverse sources, including TranStar, Twitter, and ESRI [89]. The impacted roads are located, and water depth over roads are estimated by contrasting collected images with terrain data from Google Map. Additionally, this study only considers pictures whose time of capture is known. TranStar camera data are used to increase the validation data availability; consequently, OpenSafe Fusion model results are generated without considering the traffic camera data source. For each observation, flood depth obtained from the image is compared to the OpenSafe Fusion predicted probability of flooding (Fig 18). Next, OpenSafe Fusion model per-

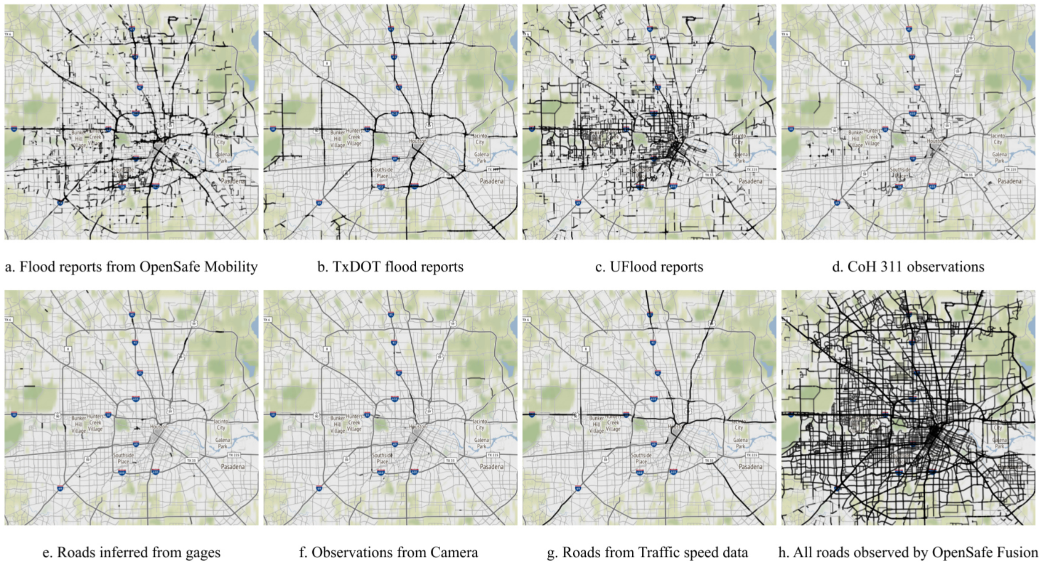

Figure 17: Spatial distribution of data availability from various sources and OpenSafe Fusion during Hurricane Harvey. All roads with observations are marked using black lines. For OpenSafe Mobility, roads without flood depth data can be considered open.

formance is quantified using the following five metrics: AUC (0.84), Weighted F1-Score (0.87), Balanced accuracy (0.88), Weighted Precision (0.88), and Weighted Recall (0.875). For developing these metrics, roads with a probability of flooding higher than 0.5 are classified as flooded. Further, Figure 18 also reports the Confusion Matrix and ROC curve. The findings show that in 87 percent of cases, OpenSafe Fusion can detect the state of roads accurately. OpenSafe Fusion, in particular, has a low false negative rate (1/14 or 7.14%; Fig 18). For situational awareness, a low false negative rate is vital since incorrectly designating roads open can pose safety risks and result in detours and delays.

A closer examination of wrongly predicted roads indicates that lack of real-time observations and terrain with a predisposition for ponding are the two main reasons for incorrect classification. A significant source of data for OpenSafe Fusion is OpenSafe Mobility. OpenSafe Mobility's flood models are currently unable to simulate stormwater networks; as a result, low-lying areas that are predominantly drained by the stormwater network will be misclassified as flooded. Such regions are easily discernible from the digital terrain model. It is possible to ignore OpenSafe Mobility observations from

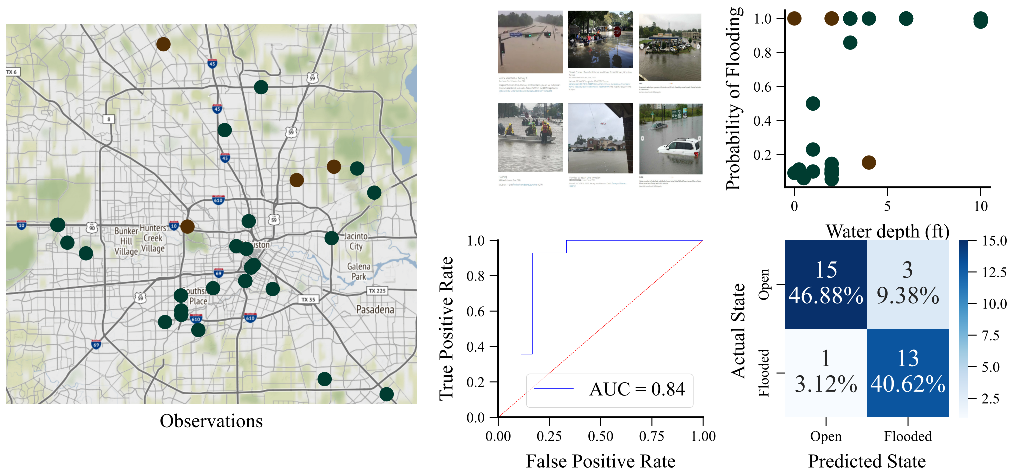

Figure 18: Validation of OpenSafe Fusion using geolocated images during Hurricane Harvey.

these regions or establish a higher bar for declaring a road to be flooded.

Ablation studies (Fig. 19) are performed to examine OpenSafe Fusion further. Specifically, six experiments are run to offer insights into the performance, data availability, accuracy, and robustness of OpenSafe Fusion. In each experiment, one data source is held back and used as the "ground truth," while the remaining data sources are used to run OpenSafe Fusion. Next, OpenSafe Fusion predictions are then compared to the held-back data set, and performance metrics (AUC and Weighted F1) are estimated for each time step. While extensive validation studies are essential before adopting OpenSafe Fusion, the ablation study presented here offers initial insights into the characteristics of the OpenSafe Fusion framework. Figure 19 reports the temporal distribution of data availability and model performance for each scenario. With the exception of OpenSafe Mobility, OpenSafe Fusion outperforms all other data sources in terms of data availability. Out of the network's 62,000 roadways, OpenSafe Fusion continuously monitors around 37,000 of them. Most highways without observations are found near the periphery of Houston (Fig.17).

Further, caution should be exercised when interpreting temporal variation of AUC and F1 scores. While estimating these measures, the held-back data source is considered the ground truth, which often is not true. For data sources that use physical sensors (cameras, speed data, and gages), Open-

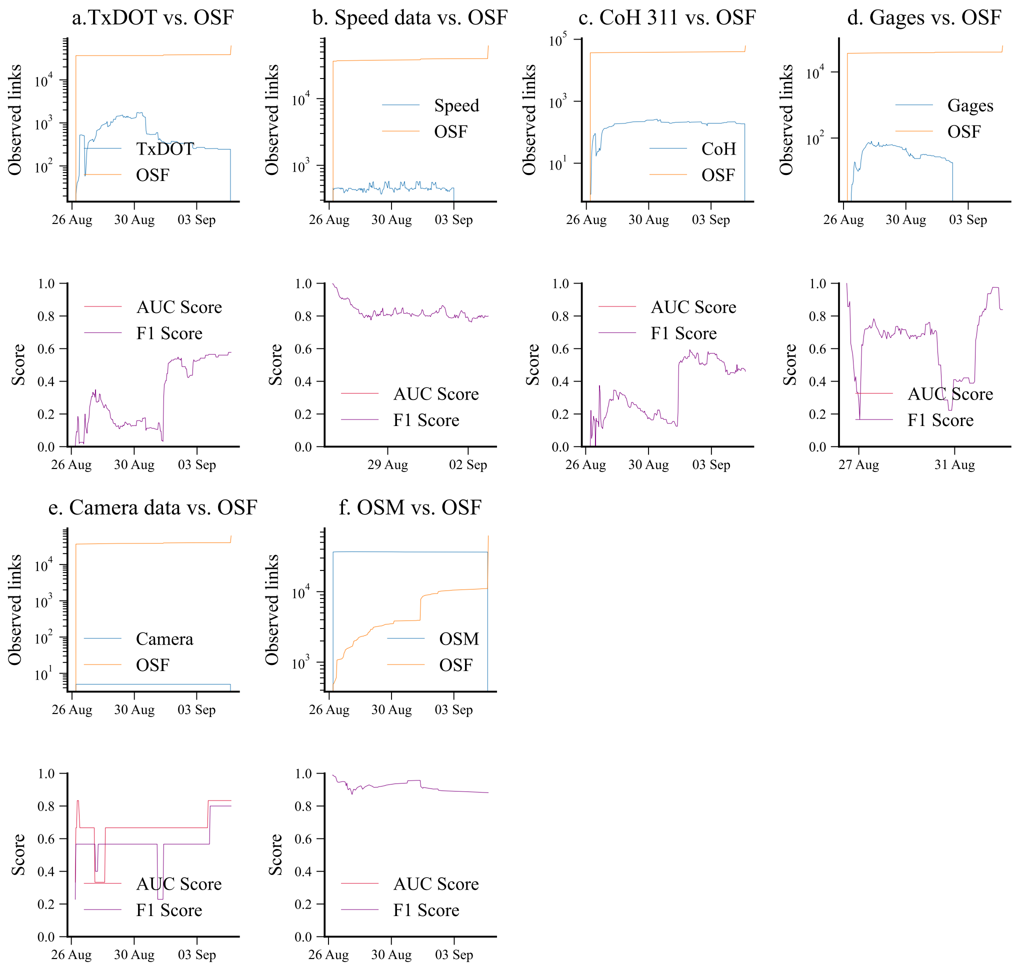

Figure 19: Results from ablation studies. Comparison of data availability (top) and temporal variation in F1 and AUC scores (bottom) between individual data sources and OpenSafe Fusion (OSF).

Safe fusions predictions show good temporal performance. For other sources (TxDOT, CoH 311), OpenSafe Fusion performance is low during the initial phases of flooding. On closer examination, some inherent characteristics of these data sources might have contributed to the low OpenSafe Fusion model performance. To elaborate, all TxDOT flood reports are not flooded. Entire stretches of highways are often marked flooded proactively due to partial closure of a link or flooding of access roads. In some cases, traversable roads are marked flooded to caution drivers about the presence of water. Similarly, for COH-311 data, many initial reports might be related to nuisance flooding. Ablation studies indicate that, for the selected case study, (a) OpenSafe Fusion observes more road links than all sources except OpenSafe Mobility. It also highlights OpenSafe Fusions ability to observe road status during the initial stages of flooding; (b) OpenSafe Fusion provides acceptable accuracy when compared to other sources, particularly considering physical sensors; and (c) OpenSafe Fusion exhibits robustness by accurately monitoring roads even if a specific data source becomes unavailable (a common occurrence during major flood events).

Finally, Fig. 20a shows the predicted roadway status on 28 August 2017 at 5 AM. From the figure, it is evident that a majority of roads in the urban centers of Houston are observed. Moreover, the unobserved roads are primarily located in the suburban regions—primarily because of the limited data generation from this region. Deploying additional data in the suburban regions could further enhance data availability. Similarly, Fig. 20b shows the network-level impact of flooding on hospital access. Specifically, it identifies regions with significant loss of connectivity to hospitals; such regions are more vulnerable due to the lack of hospital access. OpenSafe Fusion results are finally communicated via a web dashboard and REST API. OpenSafe Fusion and the accompanying web tool are designed after extensive user feedback following the tenets of user-centered design. For additional details, please refer to Panakkal et al. [11].

Completeness of OpenSafe Fusion predictions can be assessed through four key dimensions: availability, timeliness, certainty, and accuracy. Availability, measured as the percentage of road links observed, provides insight into spatial data availability (Figs. 17 and 19). In this case study, OpenSafe Fusion typically observed 60% of roads, except when OpenSafe Mobility data was not included (Fig. 19; Parts a-d and e-f). Further, Fig. 20 indicates that urban Houston has more complete observations for flooded roads than suburban areas in the periphery. Timeliness, measured as the time elapsed since the last observation from data sources for each road link, can identify regions with potentially outdated data. However, timeliness was not examined in this case study as archived data was used, and the time of data reporting was unavailable. Certainty, gauged through the predicted probabilities (Fig. 20), offers stakeholders a sense of OpenSafe Fusion's confidence in the estimated roadway status. For instance, OpenSafe Fusion is more confident in its assessment when it estimates a 98% probability of flooding than 60% for a link. Real-time accuracy can be calculated by comparing OpenSafe Fusion predictions (such as in ablation studies) to a reliable, independent source uniformly distributed through the study region. Ideally, the independent source should be selected such that excluding it from the data fusion process should not diminish the overall performance and data availability of OpenSafe Fusion. Finally, while ablation and validation studies offer insights on model performance, a comprehensive assessment of OpenSafe Fusion performance is still lacking; especially, a detailed comparison study with other tools and frameworks under diverse conditions is required. Ideally, OpenSafe Fusion should be evaluated holistically, considering model performance on five dimensions: availability, timeliness, uncertainty, fairness [90], and accuracy.

4. Discussions and Conclusions

This paper presents the methodological underpinning of the OpenSafe Fusion framework. OpenSafe Fusion addresses a key impediment to improving situational awareness—the lack of reliable real-time data on road conditions during flooding—and offers a real-time mobility-centric situational awareness framework. While additional research is required, the presented case study show that fusing multi-modal observations from existing data sources can significantly improve our ability to sense flood impacts at the link and network levels in real time. Specifically, (a) this study demonstrated that carefully designed source-specific workflows considering data source characteristics enable the extraction of road condition data from diverse sources, even sources that do not directly observe flooded roads—thus significantly increasing data availability; (b) this study also addressed the methodological challenges in fusing observations from sources diverse in characteristics and reliability to estimate the probability of roadway flooding. The presented link-level data fusion approach is adaptable, modular, and efficient and can effectively model the spatiotemporal variation in source characteristics; (c) this study illustrated that a data fusion-based approach can offer a real-time

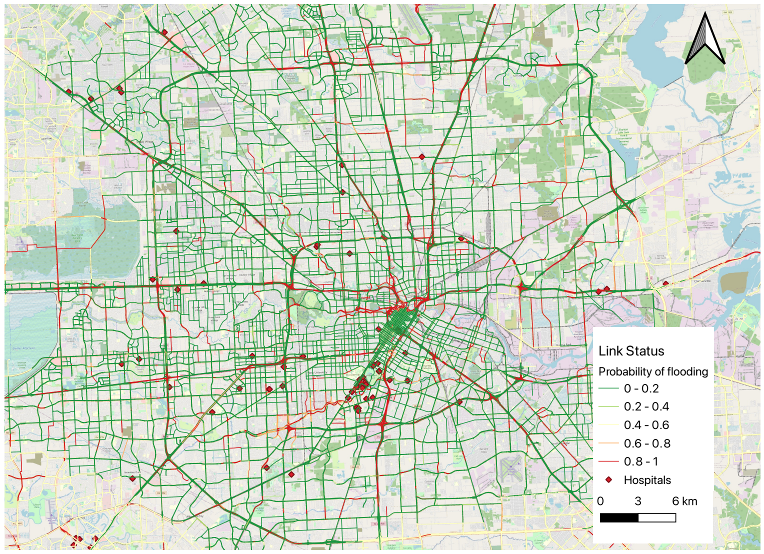

(a)

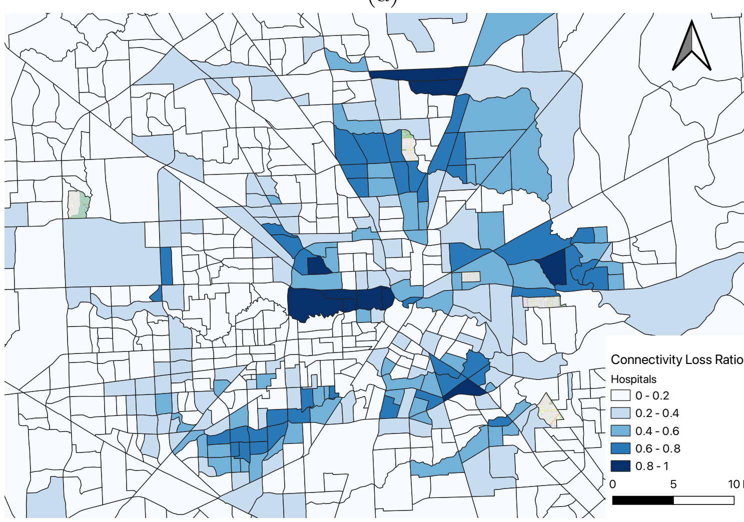

(b)

Figure 20: OpenSafe Fusion predicted roadway status (a) and connectivity loss (b) to hospitals at a time step during Hurricane Harvey.

situational awareness framework capable of monitoring road conditions of a majority of roadways and yield comprehensive and credible estimates of flood impacts at the road link and network levels. Moreover, such a data fusion-centric approach also has the potential to be more robust and equitable; finally, (d) the study offers tools, methods, and insights to enable real-time data processing, data fusion, data augmentation, and network analysis. Communities can tailor the framework to their region and available data sources to enhance roadway situational awareness—thus promoting community resilience.

OpenSafe Fusion advances the current state-of-the-art in mobility-centric flood situational awareness. Specifically, it is the first open-source framework designed following the tenets of the user-centered design process [11] and adhering to responsible design principles [91–95] that offer interpretable and grounded real-time probabilistic estimates of flood impacts on road transportation infrastructure. OpenSafe Fusion framework can significantly improve data availability and accuracy compared to existing situational awareness models depending on limited data sources (e.g., physical sensors, physics-based models, alerts). Compared to machine learning methods, OpenSafe Fusion offers interpretable, transparent, and grounded predictions; for each road link, users can identify the real-time observations used by OpenSafe Fusion to make predictions. Machine learning and physics-based models often remain static in their initial configuration and parameters, thereby failing to adapt to the changing conditions (e.g., new pumps, terrain changes, new detention basins), resulting in diminishing performance, which could often go unnoticed until significant errors occur. OpenSafe Fusion, on the other hand, will constantly adapt to changing ground conditions as it primarily leverages ground observation; in addition, the degrading performance of any source-specific workflow is easier to notice in the context of other observations. OpenSafe Fusion can promote situational awareness data equity by combining observations from multiple reliable urban sources. Compared to existing data fusion-based situational awareness tools, OpenSafe Fusion stands apart in its ability to leverage diverse urban sources that directly or indirectly observe roadway status. Finally, the OpenSafe Fusion is human-centered, contestable, and tenable to human oversight, thus promoting user trust, adhering to responsible design principles, and offering guardrails against significant model errors.

While the limited case study presented here precludes generalization, the presented proof-of-concept alludes to several advantages of the proposed framework. First, by leveraging existing data sources, communities could improve situational awareness without deploying and maintaining physical sensors at scale. Repurposing existing sources leveraging open-source tools is especially advantageous to communities without significant resources. Second, as demonstrated in the case study and ablation experiments, OpenSafe Fusion can improve data availability—spatially (throughout the watershed for both pluvial and fluvial floods) and temporally (through all stages of flooding). The improvement in data availability is especially prominent for regions with multiple data sources. Enhanced spatial and temporal data availability could translate to enhanced safety and efficiency of emergency response. Third, based on the limited case study presented here and in the context of situational awareness tools used in Houston, OpenSafe Fusion is robust and fault-tolerant as it uses multiple data sources. While sensor errors or unavailability of data sources could reduce the model performance, OpenSafe Fusion might still provide reliable results if other sources observe flooding. Deploying replicas of OpenSafe Fusion on multiple computers that are not co-located can ensure the availability of OpenSafe Fusion during power outages that frequently accompany flooding. Fourth, OpenSafe Fusion can produce reliable results by leveraging data from multiple data sources. The reliability of OpenSafe Fusion will depend on several factors, including data availability and the accuracy of data collection, processing, fusion, and augmentation workflows. Moreover, understanding the data characteristics (e.g., accuracy, bias) and factors influencing them under diverse conditions is essential for effectively fusing observations. Fifth, OpenSafe Fusion can help reduce inequities in situational awareness data availability. Many frameworks rely on limited data sources and, consequently, carry biases in the availability and accuracy of the relying sources. For example, social sensors might be concentrated near urban regions, and physical sensors are more affordable for affluent communities. Inequities in data sources could translate to inequities in situational awareness. By combining diverse sources and leveraging data augmentation, OpenSafe Fusion might be able to reduce inequity. Although OpenSafe Fusion might help ameliorate inequity in situational awareness data availability and accuracy, it cannot eliminate it—model results might be more accurate in regions with reliable and abundant data than in regions with sparse or unreliable data.

The advantages of OpenSafe Fusion should be considered in the context of its limitations. First, OpenSafe Fusion requires reliable data sources; limited, incomplete, or biased data will affect model performance. Second, OpenSafe Fusion used the discrete formulation of the Bayes Filter to fuse observations from sources. Consequently, the likelihoods, prior, and posterior are all discrete, and the model produces a deterministic estimate for the probability of a road link flooding. Additional data fusion strategies could be adopted to characterize the probability of roadway closure and associated uncertainties in the continuous domain. Third, a Markov model is sufficient for modeling OpenSafe Fusion's prediction step in Houston since reliable data is available at regular intervals. A Markov-based prediction step might not be appropriate for applications in data-scarce regions. It might be beneficial to develop generative or time series models that can predict the potential state of the system (and the uncertainty bounds) over multiple time steps without frequent observations. Fourth, since sufficient historical data was unavailable to learn interdependencies, the data fusion model adopted here neglected the dependencies between sources. Neglecting data source interdependencies may result in errors, and once data is available, more refined fusion models that can account for sensor interdependencies can be developed. Fifth, exhaustive testing and validation studies are required to validate OpenSafe Fusion and its components before a widespread deployment. Ideally, the OpenSafe Fusion framework should be deployed, and model performance should be validated over diverse storm types, including flash floods, compound floods, severe storms, and multi-peak events. Additionally, the framework's transferability and scalability should be assessed by implementing it in communities of various sizes, ranging from megacities to small towns. Sixth, it might be challenging for communities without sufficient resources to develop, deploy, and maintain OpenSafe Fusion. To facilitate faster adoption and application, the authors envision national agencies (e.g., FEMA) or non-profit organizations developing, validating, maintaining, and updating OpenSafe Fusion components and making them available to communities through API calls and easily usable modular tools. A service-based approach might allow communities with limited resources to leverage state-of-the-art situational awareness tools and overcome technological and financial accessibility and affordability barriers—thus promoting social equity and community resilience. Finally, this study used distance-based metrics to measure network-level flood impacts; future implementations could also use real-time traffic speed data to estimate travel time-based metrics to better inform situational awareness and emergency response decision-making.

Our future work will continue to improve the OpenSafe Fusion framework and its components. A prototype OpenSafe Fusion web tool is currently being tested for usability following the tenets of user-centered design [11]. Once deployed, OpenSafe Fusion will be supported by extensive data collection, processing, and archiving workflows to develop a rich dataset of sensor observations. While a wealth of literature exists on data fusion [96, 97], it predominantly deals with physical sensors or sensors with known or stationary characteristics. OpenSafe Fusion, in contrast, employs sensors whose characteristics are non-stationary, frequently unknown, and affected by various complex variables, such as location, socioeconomic and environmental factors. The gathered dataset will help characterize data sources accurately, evaluate and enhance data processing workflows, and facilitate the development of data fusion models that can capture the complex interdependencies among the data sources. Further, each component of OpenSafe Fusion can be improved. Additional sources, such as data from connected cars and the Internet of Things, could be considered. Similarly, improved data processing workflow will be developed and tested. For example, Panakkal et al. [73] report the development and performance of OpenSafe Mobility. Additional data labeling and model development are underway to accurately and precisely extract roadway status from text data (e.g., tweets) and estimate flood depth from traffic camera images. While the current version of OpenSafe Fusion offers the probability of road link flooding, future versions should offer flood hazard (depth and velocity) and vehicle-specific stability at the road links, leveraging data from relevant sources (as outlined in Table 1). Further, opportunities exist to improve the data augmentation model to consider short- and long-range spatial correlation in flooding and roadway status. Historical or simulated flood or road condition data will be used to develop spatial correlation models to support data augmentation. Better data augmentation models can improve data availability in data-scarce regions, detect outdated data, and provide a check against malicious or misinformed data from social sensors when combined with the human-in-the-loop strategy. Likewise, while human-in-the-loop strategy offers potential benefits such as enabling human supervision, enhancing transparency, contestability, and user trust, concerns arise regarding its practicality and usability in high-pressure emergency response situations with limited resources. Extensive validation studies, testing, and refinement might be required to operationalize an effective human-in-the-loop workflow. Finally, the performance of OpenSafe Fusion will be reviewed after major storm events, and the insights gathered will be used to improve the framework and its components further.

In summary, this paper addresses the need for reliable real-time mobility-centric situational awareness data—a long-standing problem with societal significance. The proposed framework offers tools and methods to sense flood impacts at the link- and network levels. The OpenSafe Fusion architecture is simple, practical, and modular, allowing communities to reuse existing data sources to improve situational awareness and upgrade the framework when more data or better models become available. While extensive additional validation studies are required, OpenSafe Fusion offers communities a potential pathway to improved situational awareness—a vital contribution to community resilience in an epoch of climate-exacerbated flood risk.

Acknowledgments

The authors gratefully acknowledge the support of this research by the National Science Foundation (award numbers 1951821 and 2227467) and the National Academy of Science, Engineering, and Medicine, Gulf Research Program (award number 2000013194). The authors thank Dr. Devika Subramanian for her general guidance and support in designing the image classifier. The authors also thank Allison Wyderka for sharing flood model results for the case study and Johnathan Roberts for annotating training data for the image classifier. The authors acknowledge the fruitful discussions with Dr. Philip Bedient, Dr. Danielle King, Dr. Andrew Juan, and Elisa Fattoracci throughout the project. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsors.

References

[1] A. Gori, I. Gidaris, J. R. Elliott, J. Padgett, K. Loughran, P. Bedient, P. Panakkal, A. Juan, Accessibility and Recovery Assessment of Houston's Roadway Network due to Fluvial Flooding during Hurricane Harvey, Natural Hazards Review 21 (2) (2020) 04020005. doi:10.1061/(ASCE)NH.1527-6996.0000355.

[2] C. Bernier, I. Gidaris, G. P. Balomenos, J. E. Padgett, Assessing the accessibility of petrochemical facilities during storm surge events, Reliability Engineering & System Safety 188 (2019) 155–167. doi:10.1016/j.ress.2019.03.021.

[3] Z. Han, H. O. Sharif, Vehicle-Related Flood Fatalities in Texas, 1959–2019, Water 12 (10) (2020) 2884. doi:10.3390/w12102884.

[4] S. Drobot, C. Benight, E. Gruntfest, Risk factors for driving into flooded roads, Environmental Hazards 7 (3) (2007) 227–234. doi:10.1016/j.envhaz.2007.07.003.

[5] S. T. Ashley, W. S. Ashley, Flood Fatalities in the United States, Journal of Applied Meteorology and Climatology 47 (3) (2008) 805–818. doi:10.1175/2007JAMC1611.1.

[6] R. C. B. Bucar, Y. M. Hayeri, Quantitative assessment of the impacts of disruptive precipitation on surface transportation, Reliability Engineering & System Safety 203 (2020) 107105. doi:10.1016/j.ress.2020.107105.

[7] U. Gangwal, S. Dong, Critical facility accessibility rapid failure early-warning detection and redundancy mapping in urban flooding, Reliability Engineering & System Safety 224 (2022) 108555. doi:10.1016/j.ress.2022.108555.

[8] S. Dong, X. Gao, A. Mostafavi, J. Gao, U. Gangwal, Characterizing resilience of flood-disrupted dynamic transportation network through the lens of link reliability and stability, Reliability Engineering & System Safety 232 (2023) 109071. doi:10.1016/j.ress.2022.109071.

[9] K. Amini, J. E. Padgett, Probabilistic risk assessment of hurricane-induced debris impacts on coastal transportation infrastructure, Reliability Engineering & System Safety 240 (2023) 109579. doi:10.1016/j.ress.2023.109579.

[10] T. M. Logan, M. J. Anderson, A. C. Reilly, Risk of isolation increases the expected burden from sea-level rise, Nature Climate Change 13 (4) (2023) 397–402. doi:10.1038/s41558-023-01642-3.

[11] P. Panakkal, E. S. M. Fattoracci, J. E. Padgett, D. D. King, T. Yoo, Sensing flooded roads to support roadway mobility during flooding: A web-based tool and insights from needs assessment interviews, Natural Hazards Review 24 (4) (2023) 04023039. doi:10.1061/NHREFO.NHENG-1753.
URL `https://ascelibrary.org/doi/abs/10.1061/NHREFO.NHENG-1753`

[12] K. C. Dey, A. Mishra, M. Chowdhury, Potential of Intelligent Transportation Systems in Mitigating Adverse Weather Impacts on Road Mobility: A Review, IEEE Transactions on Intelligent Transportation Systems 16 (3) (2015) 1107–1119. doi:10.1109/TITS.2014.2371455.

[13] A. Gissing, S. Opper, M. Tofa, L. Coates, J. McAneney, Influence of road characteristics on flood fatalities in Australia, Environmental Hazards 18 (5) (2019) 434–445. doi:10.1080/17477891.2019.1609407.

[14] M. Pregnolato, A. Ford, V. Glenis, S. Wilkinson, R. Dawson, Impact of Climate Change on Disruption to Urban Transport Networks from Pluvial Flooding, Journal of Infrastructure Systems 23 (4) (2017) 04017015. doi:10.1061/(ASCE)IS.1943-555X.0000372.

[15] E. Mühlhofer, E. E. Koks, C. M. Kropf, G. Sansavini, D. N. Bresch, A generalized natural hazard risk modelling framework for infrastructure failure cascades, Reliability Engineering & System Safety 234 (2023) 109194. doi:10.1016/j.ress.2023.109194.

[16] C. Haddock, S. Kanwar, Infrastructure Design Manual 2021, Tech. rep., City of Houston, Houston Public Works, Houston (Jul. 2021).

[17] FEMA, Historic Disaster Response to Hurricane Harvey in Texas — FEMA.gov, accessed July 03, 2022 (2017).

[18] N.-B. Chang, D.-H. Guo, Urban flash flood monitoring, mapping and forecasting via a tailored sensor network system, in: 2006 IEEE International Conference on Networking, Sensing and Control, IEEE, 2006, pp. 757–761.

[19] M. A. Islam, T. Islam, M. A. Syrus, N. Ahmed, Implementation of flash flood monitoring system based on wireless sensor network in bangladesh, in: 2014 international conference on informatics, electronics & vision (ICIEV), IEEE, 2014, pp. 1–6.

[20] B. Arshad, R. Ogie, J. Barthelemy, B. Pradhan, N. Verstaevel, P. Perez, Computer vision and iot-based sensors in flood monitoring and mapping: A systematic review, Sensors 19 (22) (2019) 5012.

[21] D. Loftis, D. Forrest, S. Katragadda, K. Spencer, T. Organski, C. Nguyen, S. Rhee, Stormsense: A new integrated network of iot water level sensors in the smart cities of hampton roads, va, Marine Technology Society Journal 52.

[22] H. Farahmand, X. Liu, S. Dong, A. Mostafavi, J. Gao, A Network Observability Framework for Sensor Placement in Flood Control Networks to Improve Flood Situational Awareness and Risk Management, Reliability Engineering & System Safety 221 (2022) 108366. doi:10.1016/j.ress.2022.108366.

[23] Twitter, Inc., Twitter, accessed March 26, 2022 (2022).

[24] Google LLC, Driving directions, live traffic & road conditions updates, accessed March 18, 2022 (2022).
URL `https://www.waze.com/live-map/`

[25] N. McIntyre, H. Needham, Flood-map: A simple web map to visualize flood information on maps, accessed August 04, 2022 (2017).
URL `https://github.com/nickmcintyre/flood-map`

[26] C. Fan, M. Esparza, J. Dargin, F. Wu, B. Oztekin, A. Mostafavi, Spatial biases in crowdsourced data: Social media content attention concentrates on populous areas in disasters, Computers, Environment and Urban Systems 83 (2020) 101514, publisher: Elsevier.

[27] X. He, D. Lu, D. Margolin, M. Wang, S. E. Idrissi, Y.-R. Lin, The signals and noise: actionable information in improvised social media channels during a disaster, in: Proceedings of the 2017 ACM on web science conference, 2017, pp. 33–42.

[28] S. Praharaj, F. T. Zahura, T. D. Chen, Y. Shen, L. Zeng, J. L. Goodall, Assessing Trustworthiness of Crowdsourced Flood Incident Reports Using Waze Data: A Norfolk, Virginia Case Study, Transportation Research Record 2675 (12) (2021) 650–662, publisher: SAGE Publications Inc. doi:10.1177/03611981211031212.
URL `https://doi.org/10.1177/03611981211031212`

[29] F. Jin, W. Wang, L. Zhao, E. Dougherty, Y. Cao, C.-T. Lu, N. Ramakrishnan, Misinformation propagation in the age of twitter, Computer 47 (12) (2014) 90–94.

[30] K. Ahmad, K. Pogorelov, M. Riegler, O. Ostroukhova, P. Halvorsen, N. Conci, R. Dahyot, Automatic detection of passable roads after floods in remote sensed and social media data, Signal Processing: Image Communication 74 (2019) 110–118. doi:10.1016/j.image.2019.02.002.

[31] M. Wieland, S. Martinis, A modular processing chain for automated flood monitoring from multi-spectral satellite data, Remote Sensing 11 (19) (2019) 2330.

[32] P. Matgen, S. Martinis, W. Wagner, V. Freeman, P. Zeil, N. McCormick, et al., Feasibility assessment of an automated, global, satellite-based flood-monitoring product for the copernicus emergency management service, Luxembourg: Publications Office of the European Union.

[33] M. T. Perks, A. J. Russell, A. R. Large, Advances in flash flood monitoring using unmanned aerial vehicles (UAVs), Hydrology and Earth System Sciences 20 (10) (2016) 4005–4015, publisher: Copernicus GmbH.

[34] L. Alzubaidi, J. Zhang, A. J. Humaidi, A. Al-Dujaili, Y. Duan, O. Al-Shamma, J. Santamaría, M. A. Fadhel, M. Al-Amidie, L. Farhan, Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions, Journal of Big Data 8 (1) (2021) 53. doi:10.1186/s40537-021-00444-8.

[35] A. Shrestha, A. Mahmood, Review of Deep Learning Algorithms and Architectures, IEEE Access 7 (2019) 53040–53065. doi:10.1109/ACCESS.2019.2912200.

[36] M. Geetha, M. Manoj, A. S. Sarika, M. Mohan, S. N. Rao, Detection and estimation of the extent of flood from crowd sourced images, in: 2017 International Conference on Communication and Signal Processing (ICCSP), IEEE, 2017, pp. 0603–0608.

[37] J. Jiang, J. Liu, C.-Z. Qin, D. Wang, Extraction of Urban Waterlogging Depth from Video Images Using Transfer Learning, Water 10 (10) (2018) 1485, number: 10 Publisher: Multidisciplinary Digital Publishing Institute. doi:10.3390/w10101485.

[38] P. Chaudhary, S. DAronco, M. Moy de Vitry, J. P. Leito, J. D. Wegner, Flood-Water Level Estimation from Social Media Images, ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences IV-2/W5 (2019) 5–12. doi:10.5194/isprs-annals-IV-2-W5-5-2019.

[39] H. TranStar, Houston TranStar - Traffic Map, (Last Accessed: 2022-03-27) (2022).
URL https://traffic.houstontranstar.org/layers/

[40] Texas Department of Transportation, Texas Department of Transportation - DriveTexas, (Last Accessed: 2022-03-15) (2022).
URL `https://drivetexas.org`

[41] A. Mosavi, P. Ozturk, K.-w. Chau, Flood prediction using machine learning models: Literature review, Water 10 (11) (2018) 1536.

[42] F. T. Zahura, J. L. Goodall, J. M. Sadler, Y. Shen, M. M. Morsy, M. Behl, Training machine learning surrogate models from a high-fidelity physics-based model: Application for real-time street-scale flood prediction in an urban coastal community, Water Resources Research 56 (10) (2020) e2019WR027038.

[43] F. Yuan, C.-C. Lee, W. Mobley, H. Farahmand, Y. Xu, R. Blessing, S. Dong, A. Mostafavi, S. D. Brody, Predicting road flooding risk with crowdsourced reports and fine-grained traffic data, Computational Urban Science 3 (1) (2023) 15.

[44] R. C. B. Bucar, Y. M. Hayeri, Quantitative flood risk evaluation to improve drivers' route choice decisions during disruptive precipitation, Reliability Engineering & System Safety 219 (2022) 108202. doi:10.1016/j.ress.2021.108202.

[45] P.-A. Versini, E. Gaume, H. Andrieu, Application of a distributed hydrological model to the design of a road inundation warning system for flash flood prone areas, Natural Hazards and Earth System Sciences 10 (4) (2010) 805–817, publisher: Copernicus GmbH. doi:10.5194/nhess-10-805-2010.

[46] J.-P. Naulin, O. Payrastre, E. Gaume, Spatially distributed flood forecasting in flash flood prone areas: Application to road network supervision in Southern France, Journal of Hydrology 486 (2013) 88–99, publisher: Elsevier.

[47] D. Mioc, J. Nkhwanana, K. Moreiri, B. Nickerson, M. Santos, E. McGillivray, A. Morton, F. Anton, A. Ahmad, M. Mezouaghi, L. Mofford, P. Tang, Natural and man-made flood risk mapping and warning for socially vulnerable populations, International Journal of Safety and Security Engineering 5 (3) (2015) 183–202. doi:10.2495/SAFE-V5-N3-183-202.

[48] M. M. Morsy, J. L. Goodall, G. L. O'Neil, J. M. Sadler, D. Voce, G. Hassan, C. Huxley, A cloud-based flood warning system for forecasting impacts to transportation infrastructure systems, Environmental Modelling & Software 107 (2018) 231–244. doi:10.1016/j.envsoft.2018.05.007.

[49] X. Ming, Q. Liang, X. Xia, D. Li, H. J. Fowler, Real-Time Flood Forecasting Based on a High-Performance 2-D Hydrodynamic Model and Numerical Weather Predictions, Water Resources Research 56 (7) (2020) e2019WR025583, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2019WR025583. doi:10.1029/2019WR025583.

[50] J. M. Johnson, J. M. Coll, P. J. Ruess, J. T. Hastings, Challenges and Opportunities for Creating Intelligent Hazard Alerts: The FloodHippo Prototype, Journal of the American Water Resources Association 54 (4) (2018) 872–881. doi:10.1111/1752-1688.12645.
URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/1752-1688.12645`

[51] P. Panakkal, A. Juan, M. Garcia, J. E. Padgett, P. Bedient, Towards enhanced response: Integration of a flood alert system with road infrastructure performance models, in: Structures Congress 2019: Buildings and Natural Disasters, American Society of Civil Engineers Reston, VA, 2019, pp. 294–305.

[52] S. Dong, T. Yu, H. Farahmand, A. Mostafavi, Probabilistic modeling of cascading failure risk in interdependent channel and road networks in urban flooding, Sustainable Cities and Society 62 (2020) 102398. doi:10.1016/j.scs.2020.102398.

[53] R.-Q. Wang, H. Mao, Y. Wang, C. Rae, W. Shaw, Hyper-resolution monitoring of urban flooding with social media and crowdsourcing data, Computers & Geosciences 111. doi:10.1016/j.cageo.2017.11.008.

[54] J. F. Rosser, D. G. Leibovici, M. J. Jackson, Rapid flood inundation mapping using social media, remote sensing and topographic data, Natural Hazards 87 (1) (2017) 103–120. doi:10.1007/s11069-017-2755-0.
URL `10.1007/s11069-017-2755-0`

[55] K. Ahmad, K. Pogorelov, M. Riegler, O. Ostroukhova, P. Halvorsen, N. Conci, R. Dahyot, Automatic detection of passable roads after floods

in remote sensed and social media data, Signal Processing: Image Communication 74 (2019) 110118. doi:10.1016/j.image.2019.02.002.

[56] D. Frey, M. Butenuth, D. Straub, Probabilistic graphical models for flood state detection of roads combining imagery and dem, IEEE Geoscience and Remote Sensing Letters 9 (6) (2012) 10511055. doi:10.1109/LGRS.2012.2188881.

[57] D. Frey, M. Butenuth, Multi-temporal damage assessment of linear infrastructural objects using dynamic bayesian networks, in: 2011 6th International Workshop on the Analysis of Multi-temporal Remote Sensing Images (Multi-Temp), 2011, p. 6164. doi:10.1109/Multi-Temp.2011.6005048.
URL `10.1109/Multi-Temp.2011.6005048`

[58] D. Frey, M. Butenuth, Trafficability analysis after flooding in urban areas using probabilistic graphical models, in: 2011 Joint Urban Remote Sensing Event, 2011, p. 345348. doi:10.1109/JURSE.2011.5764790.
URL `10.1109/JURSE.2011.5764790`

[59] J. P. d. Albuquerque, B. Herfort, A. Brenning, A. Zipf, A geographic approach for combining social media and authoritative data towards identifying useful information for disaster management, International Journal of Geographical Information Science 29 (4) (2015) 667689. doi:10.1080/13658816.2014.996567.

[60] B. Bischke, P. Bhardwaj, A. Gautam, P. Helber, D. Borth, A. Dengel, Detection of flooding events in social multimedia and satellite imagery using deep neural networks 3.

[61] R. De O. Werneck, I. C. Dourado, S. G. Fadel, S. Tabbone, R. Da S. Torres, Graph-based early-fusion for flood detection, in: 2018 25th IEEE International Conference on Image Processing (ICIP), 2018, p. 10481052. doi:10.1109/ICIP.2018.8451011.
URL `10.1109/ICIP.2018.8451011`

[62] A. C. Robinson, J. Chen, E. J. Lengerich, H. G. Meyer, A. M. MacEachren, Combining Usability Techniques to Design Geovisualization Tools for Epidemiology, Cartography and geographic information

science 32 (4) (2005) 243–255. doi:10.1559/152304005775194700.
URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2786201/

[63] S. Thrun, W. Burgard, D. Fox, Probabilistic robotics (intelligent robotics and autonomous agents series), Intelligent robotics and autonomous agents, The MIT Press.

[64] K. P. Murphy, Probabilistic Machine Learning: An introduction, MIT Press, 2022.
URL probml.ai

[65] Texas Department of Transportation, DriveTexas API, https://api.drivetexas.org/ (2023).

[66] L. L. Williams, M. Lück-Vogel, Comparative assessment of the gis based bathtub model and an enhanced bathtub model for coastal inundation, Journal of Coastal Conservation 24 (2) (2020) 1–15.

[67] F. A. Azlan, A. Ahmad, S. Yussof, A. A. Ghapar, Analyzing Algorithms to Detect Disaster Events using Social Media, in: 2020 8th International Conference on Information Technology and Multimedia (ICIMU), 2020, pp. 384–389. doi:10.1109/ICIMU49871.2020.9243599.

[68] C. Fan, F. Wu, A. Mostafavi, A Hybrid Machine Learning Pipeline for Automated Mapping of Events and Locations From Social Media in Disasters, IEEE Access 8 (2020) 10478–10490. doi:10.1109/ACCESS.2020.2965550.

[69] C. Zhang, C. Fan, W. Yao, X. Hu, A. Mostafavi, Social media for intelligent public information and warning in disasters: An interdisciplinary review, International Journal of Information Management 49 (2019) 190–207.

[70] P. A. Zandbergen, A comparison of address point, parcel and street geocoding techniques, Computers, Environment and Urban Systems 32 (3) (2008) 214–232.

[71] K. Lai, J. R. Porter, M. Amodeo, D. Miller, M. Marston, S. Armal, A natural language processing approach to understanding context in the extraction and geocoding of historical floods, storms, and adaptation measures, Information Processing & Management 59 (1) (2022) 102735.

[72] P. Panakkal, A. Price, J. Padgett, P. Bedient, Opensafe mobility, accessed October 04, 2022 (2022).

[73] P. Panakkal, A. M. Wyderka, J. E. Padgett, P. B. Bedient, Safer this way: Identifying flooded roads for facilitating mobility during floods, Journal of Hydrology (2023) 130100doi:https://doi.org/10.1016/j.jhydrol.2023.130100.

[74] D. A. Dao, D. Kim, S. Kim, J. Park, Determination of flood-inducing rainfall and runoff for highly urbanized area based on high-resolution radar-gauge composite rainfall data and flooded area GIS data, Journal of Hydrology 584 (2020) 124704. doi:10.1016/j.jhydrol.2020.124704.

[75] M. Dempsey, Rachael Dempsey, Carolyn Edds, Cheryl Phillips, Anonymous Contributors, Flooded Streets due to #Harvey (No longer updated), accessed August 04, 2022 (2017).

[76] City of Houston, Houston, Texas 3-1-1, `https://www.houstontx.gov/311/` (2022).
URL `https://www.houstontx.gov/311/`

[77] A. Gori, I. Gidaris, J. R. Elliott, J. Padgett, K. Loughran, P. Bedient, P. Panakkal, A. Juan, Accessibility and recovery assessment of houstons roadway network due to fluvial flooding during hurricane harvey, Natural hazards review 21 (2) (2020) 04020005.

[78] A. Sebastian, K. T. Lendering, B. L. M. Kothuis, A. D. Brand, S. N. Jonkman, P. H. A. J. M. van Gelder, M. Godfroij, B. Kolen, M. Comes, S. L. M. Lhermitte, K. Meesters, B. A. van de Walle, A. Ebrahimi Fard, S. Cunningham, N. Khakzad, V. Nespeca, Hurricane Harvey Report, Tech. rep., Delft University Publishers, Delft (2017).
URL `http://resolver.tudelft.nl/uuid:54c24519-c366-4f2f-a3b9-0807db26f69c`

[79] S. Dong, X. Gao, A. Mostafavi, J. Gao, Modest flooding can trigger catastrophic road network collapse due to compound failure, Communications Earth & Environment 3 (1) (2022) 1–10. doi:10.1038/s43247-022-00366-0.

[80] S. N. Jonkman, M. Godfroy, A. Sebastian, B. Kolen, Brief communication: Loss of life due to Hurricane Harvey, Natural Hazards and Earth

System Sciences 18 (4) (2018) 1073–1078. doi:10.5194/nhess-18-1073-2018.

[81] USGS, USGS Water Data for the Nation, https://waterdata.usgs.gov/nwis (2022).
URL https://waterdata.usgs.gov/nwis

[82] M. E. Phillips, Hurricane harvey twitter dataset, University of North Texas, Denton, TX, USA,.
URL https://digital.library.unt.edu/ark:/67531/metadc993940/

[83] Texas Department of Transportation, Texas Department of Transportation - DriveTexas, https://drivetexas.org, accessed on: 2022-03-15 (2022).

[84] S. Praharaj, F. T. Zahura, T. D. Chen, Y. Shen, L. Zeng, J. L. Goodall, Assessing Trustworthiness of Crowdsourced Flood Incident Reports Using Waze Data: A Norfolk, Virginia Case Study, Transportation Research Record (2021) 03611981211031212Publisher: SAGE Publications Inc. doi:10.1177/03611981211031212.
URL https://doi.org/10.1177/03611981211031212

[85] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.

[86] M. Imran, P. Mitra, C. Castillo, Twitter as a lifeline: Human-annotated twitter corpora for nlp of crisis-related messages, in: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016), European Language Resources Association (ELRA), Paris, France, 2016.

[87] F. Alam, F. Ofli, M. Imran, Crisismmd: Multimodal twitter datasets from natural disasters, in: Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM), 2018.

[88] CrisisNLP, https://crisisnlp.qcri.org/, accessed November 01, 2022 (2022).

[89] Story Map Series, publisher: ESRI.
URL https://cuahsi.maps.arcgis.com/apps/MapSeries/index.html?appid=e201458acb

[90] A. Castelnovo, R. Crupi, G. Greco, D. Regoli, I. G. Penco, A. C. Cosentini, A clarification of the nuances in the fairness metrics landscape, Scientific Reports 12 (1) (2022) 4209. doi:10.1038/s41598-022-07939-1.

[91] A. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, F. Herrera, Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI, Information Fusion 58 (2020) 82–115. doi:10.1016/j.inffus.2019.12.012.

[92] Tobias Clement, Nils Kemmerzell, Mohamed Abdelaal, Michael Amberg, XAIR: A Systematic Metareview of Explainable AI (XAI) Aligned to the Software Development Process 5 (1) (2023) 78–108. doi:10.3390/make5010006.

[93] X. Wu, L. Xiao, Y. Sun, J. Zhang, T. Ma, L. He, A Survey of Human-in-the-loop for Machine Learning, Future Generation Computer Systems 135 (2022) 364–381. arXiv:2108.00941, doi:10.1016/j.future.2022.05.014.

[94] B. Shneiderman, Bridging the Gap Between Ethics and Practice: Guidelines for Reliable, Safe, and Trustworthy Human-centered AI Systems, ACM Transactions on Interactive Intelligent Systems 10 (4) (2020) 26:1–26:31. doi:10.1145/3419764.

[95] H. Lyons, E. Velloso, T. Miller, Fair and Responsible AI: A Focus on the Ability to Contest (Feb. 2021). arXiv:2102.10787, doi:10.48550/arXiv.2102.10787.

[96] B. Khaleghi, A. Khamis, F. O. Karray, S. N. Razavi, Multisensor data fusion: A review of the state-of-the-art, Information Fusion 14 (1) (2013) 28–44. doi:10.1016/j.inffus.2011.08.001.

[97] F. Castanedo, A Review of Data Fusion Techniques, The Scientific World Journal 2013 (2013) 1–19. doi:10.1155/2013/704504.