Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X
Volume 4, Issue 3, 2026

MULTIMODAL SENSOR FUSION AND OPERATIONS RESEARCH TECHNIQUES FOR REAL-TIME FLASH-FLOOD EARLY WARNING AND EMERGENCY RESPONSE: A REVIEW

Husnain Ali

Department of Mathematics, Govt. Zamindar Graduate College, Gujrat, University of Sargodha, Sargodha, Pakistan

hasnainali1427830@gmail.com

DOI: https://doi.org/10.5281/zenodo.23559382

Keywords
flash-flood early warning, multimodal data fusion, synthetic aperture radar, sensor-fault detection, kinematic wave routing, disaster operations research.

Article History
Received: 20 January 2026
Accepted: 05 March 2026
Published: 19 March 2026

Copyright @Author
Corresponding Author: *
Husnain Ali

Abstract
Flash floods are among the deadliest natural hazards because they typically develop within hours of the triggering event, frequently leaving disaster managers with lead times of only minutes to a few hours. This paper reviews the principal technical building blocks that underpin modern flash-flood early-warning and response platforms: satellite and radar-based flood mapping, in-situ hydrometric telemetry and sensor-fault detection, kinematic and hydrodynamic wave-routing models, and operations-research (OR) methods for emergency resource allocation. This review focuses in particular on a common weakness in operational systems, namely the tendency to interpret the sudden loss of a river-gauge signal as benign downtime rather than as possible evidence of station destruction during a flood peak. The paper also surveys evacuation-behaviour and loss-of-life models used to quantify the human benefit of earlier warning, and discusses the reproducibility problems that arise when historical disasters are used to benchmark such systems. The review concludes by outlining open research problems in cross-agency data standardisation, real-time computational tractability, and the coupling of physics-based and data-driven forecasting.

I. INTRODUCTION

Weather-, climate-, and water-related hazards accounted for close to 12,000 recorded disasters worldwide between 1970 and 2021, causing roughly two million deaths and reported economic losses in excess of US$4.3 trillion, with floods and storms responsible for a substantial share of both the casualty and cost totals [1]. Mortality has nonetheless fallen over the same period even as the number of events has risen, a decline that the WMO attributes largely to improved early-warning systems and better-coordinated emergency response [1].

Flash floods are harder to warn against than slower riverine floods for several reasons. Riverine floods build over days, whereas flash floods triggered by convective cloudbursts, dam or embankment breaches, or glacial and landslide-driven debris flows can inundate a catchment within tens of minutes to a few hours of the triggering rainfall or structural failure.

Within this compressed timeline, three difficulties recur across the literature surveyed in this paper. Data that could inform a warning, such as river-stage telemetry, rain-gauge readings, satellite imagery, seismic feeds, and citizen reports, are collected by different agencies at different spatial and temporal resolutions and are rarely fused automatically into a single situational picture (Sections II-A and II-F).

Physical monitoring stations can also be vulnerable to the hazard they are meant to detect, since a river gauge sited within the channel itself may be damaged or submerged as a flood peak arrives; whether such an event is correctly interpreted as a hazard signal rather than as routine downtime depends on how the monitoring system is designed (Section II-B). Finally, once a warning is issued, the allocation of rescue assets such as boats, ambulances, and helicopters is often performed through manual or population-proportional heuristics that do not account for which roads and bridges have already been cut by floodwater [11].

This paper surveys the technical literature relevant to each of these difficulties and to the broader question of how their solutions can be combined into a coherent decision-support pipeline. Section II reviews remote sensing, sensor-fault detection, hydrodynamic and kinematic flood routing, operations-research-based resource allocation, and evacuation/loss-of-life modelling in turn. Section III discusses the cross-cutting integration challenges that arise when these components are combined, including the reproducibility problems inherent in retrospective benchmarking. Section IV restates a quantitative relationship that recurs across the wave-routing literature reviewed in Section II-C. Section V concludes with open research directions.

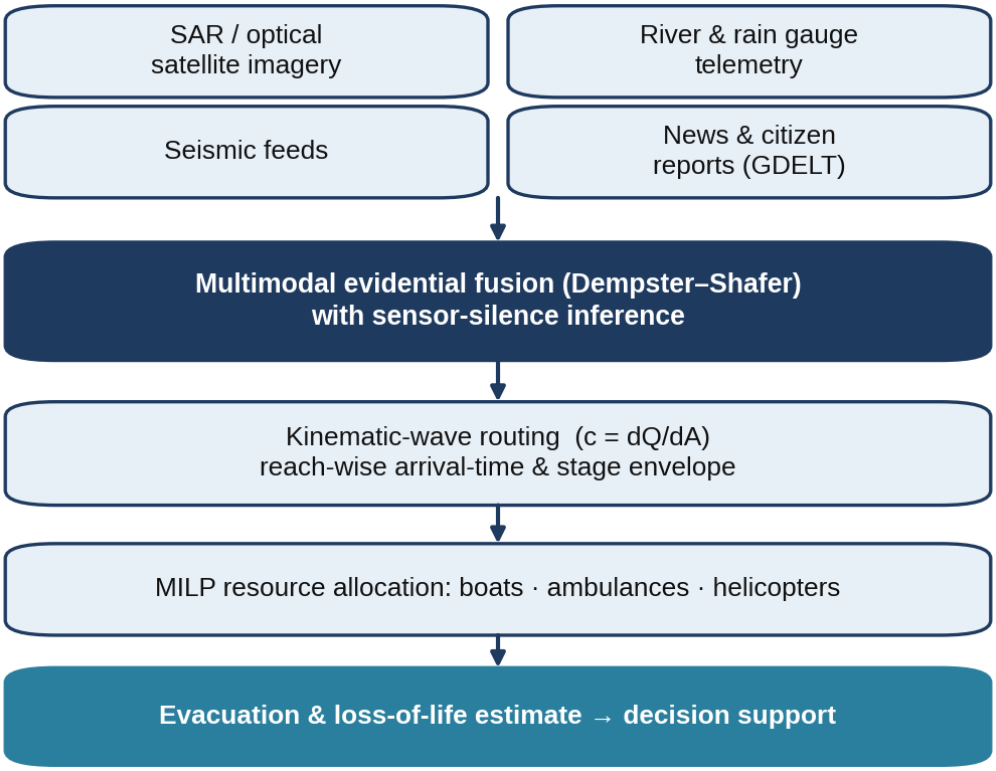


Fig. 1. Conceptual decision-support pipeline linking the evidence streams, fusion, routing, resource-allocation and evacuation components reviewed in Section II.

II. LITERATURE REVIEW

A. Remote Sensing and Satellite Flood Mapping

Spaceborne Synthetic Aperture Radar (SAR), most prominently the C-band Sentinel-1 constellation operated under the Copernicus programme, is widely used for flood delineation because active microwave imaging penetrates cloud cover and operates independently of daylight. Calm, flooded surfaces act as specular reflectors and return markedly lower radar backscatter than surrounding land; this contrast is commonly segmented using automatic bimodal thresholding methods such as Otsu's algorithm, though Otsu's method itself is a general image-segmentation technique rather than one developed specifically for flood mapping [2]. Martinis et al. demonstrated near-real-time flood detection from high-resolution X-band SAR using artificial neural networks [3],

while Chini et al. showed that interferometric SAR coherence can help detect floodwater between buildings in dense urban settings [4]. Optical multispectral sensors such as Sentinel-2 offer complementary, higher-fidelity delineation through indices such as the Normalized Difference Water Index, but are blinded by the very cloud cover that typically accompanies flood-producing storms [5].

A further limitation, particularly relevant to mountainous flash-flood corridors, is that steep terrain introduces radar layover, foreshortening, and shadow artefacts that can be mistaken for standing water unless cross-referenced against a digital elevation model [3]. Reported satellite revisit intervals of roughly one to several days and downlink-to-product latencies of a few hours further mean that SAR is generally considered better suited to post-onset confirmation and damage assessment than to the initial, sub-hourly detection of a fast-developing flash flood.

B. In-Situ Telemetry and Sensor-Fault or Silence Detection

Physical river- and rain-gauge networks operated by hydrological agencies remain the most direct source of ground-truth stage and discharge data, but they suffer from sparse spatial coverage and, critically, from vulnerability to the hazard itself. Evidential frameworks based on Dempster–Shafer belief theory [6], [7] provide one mathematical basis that can be used to combine measurements of differing reliability and to reason about missing or contradictory readings in a sensor network. To the extent that this can be judged from the systems and reviews examined here, most operational telemetry stacks appear to treat an unexpected gap in transmission as an undifferentiated fault state, rather than distinguishing between benign communication downtime and the specific, high-consequence possibility that the station has been destroyed or submerged during an active hazard. Framing sensor silence as a context-dependent inference problem, one that conditions the interpretation of a data gap on concurrent upstream rainfall, seismic, or discharge signals, offers a way to convert an apparent loss of information into a positive indicator of hazard severity rather than a neutral one.

IEER
Institute for Excellence in Education & Research

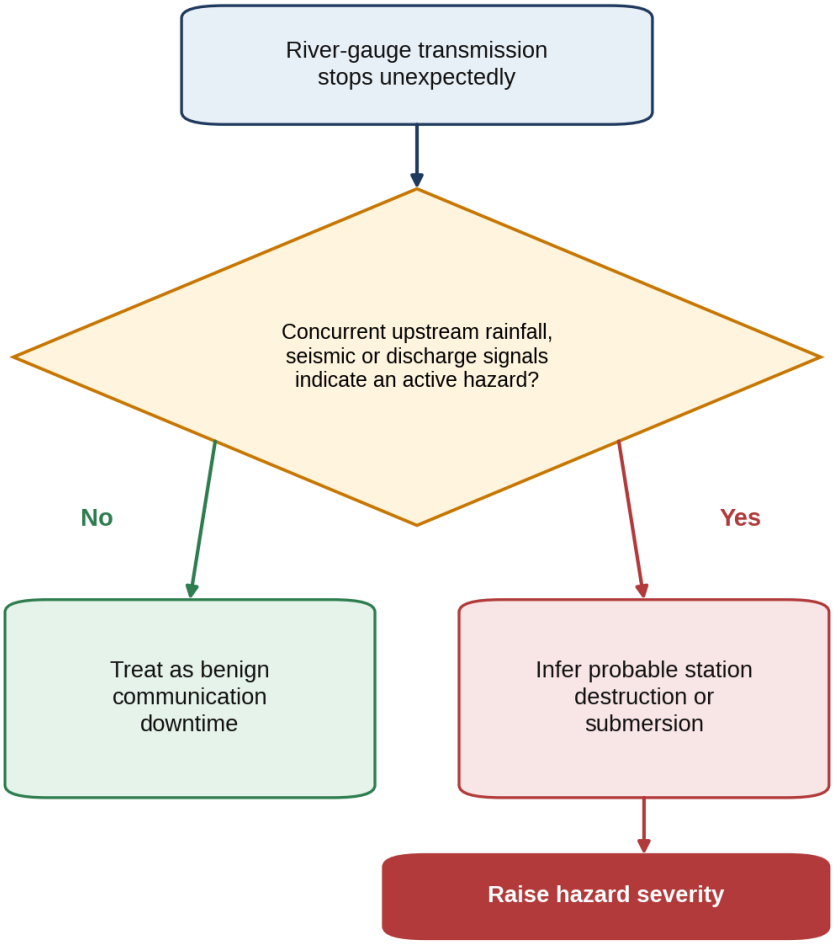


Fig. 2. Context-dependent interpretation of sensor silence: a transmission gap is treated as a hazard indicator only when concurrent upstream evidence supports it.

C. Kinematic and Hydrodynamic Flood-Wave Routing

Once a flood pulse has been detected or triggered, downstream arrival timing and peak-stage estimation are commonly addressed using either full two-dimensional shallow-water (Saint-Venant) solvers or simplified one-dimensional kinematic-wave approximations. Hunter et al. benchmarked several two-dimensional urban flood-inundation models and highlighted the substantial computational and calibration-data burden such models impose [8], a burden that is often incompatible with the minutes-scale decision cycles required during an unfolding emergency. The kinematic-wave approximation, formalised by Lighthill and Whitham for both traffic flow and flood-wave translation, instead expresses the wave celerity as the derivative of discharge with respect to cross-sectional flow area, $c = dQ/dA$, and allows a river corridor to be discretised into a sequence of reaches, each assigned a locally calibrated celerity [9]. This substantially reduces computational cost at the expense of spatial detail, and is well suited to producing rapid, corridor-wide arrival-time and stage envelopes for evacuation planning, provided that reach-specific celerities are recalibrated as channel geometry and bed roughness change during an extreme event.

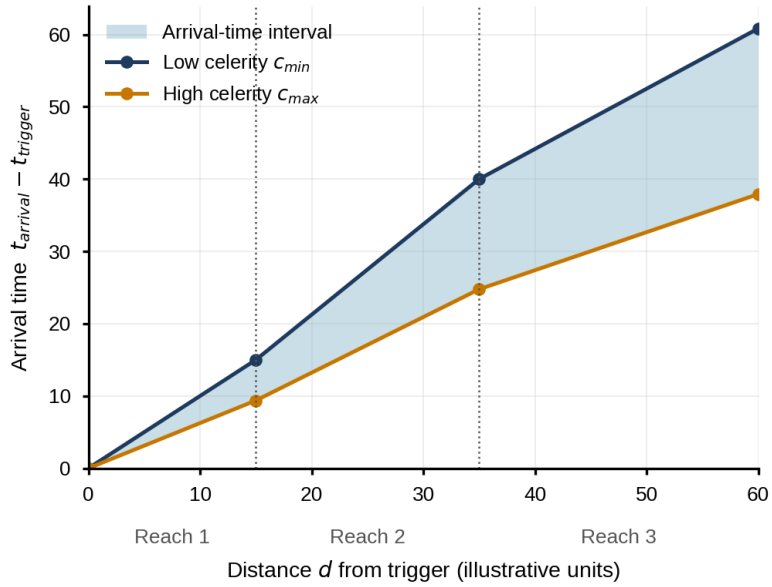


Fig. 3. Illustrative reach-wise kinematic routing: a locally calibrated celerity per reach yields an arrival-time interval rather than a single predicted time.

D. Operations Research for Emergency Resource Allocation

Humanitarian logistics has been a long-standing application area for operations research, encompassing facility location, vehicle routing, and resource-allocation problems under uncertainty and capacity constraints [10]. Perrigo, Altay, and Robinson's more recent review of OR applications in disaster response and evacuation highlights a persistent limitation: most published formulations assume a static, fully known road network and a fixed, pre-catalogued vehicle inventory [11], assumptions that rarely hold once floodwater has begun to sever bridges and low-lying highways. Mixed-Integer Linear Programming (MILP) formulations, solvable in practice with general-purpose solvers such as those in the Google OR-Tools suite, are well matched to this class of problem because they can encode discrete vehicle-type decision variables, per-resource capacity limits, and priority or isolation constraints within a single linear objective, typically some measure of expected lives protected. When the responding agency's exact fleet inventory is unknown, humanitarian minimum-standard frameworks such as the Sphere Handbook [12] and national disaster-management guidelines [13] provide reference coverage norms against which a minimum required fleet size can be estimated for a given at-risk population.

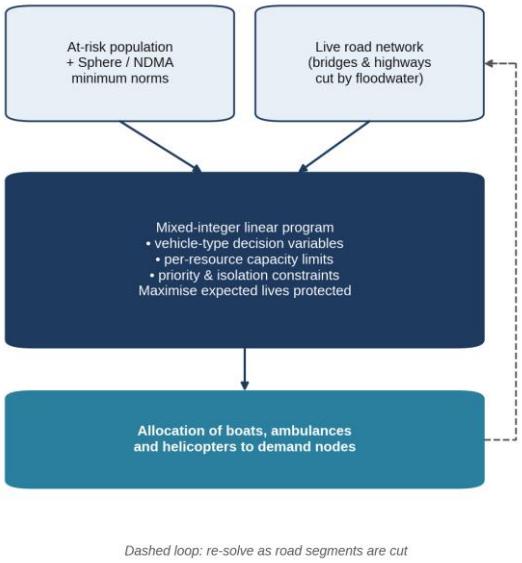


Fig. 4. Structure of the MILP-based emergency resource-allocation problem, re-solved as the feasible road network changes.

E. Evacuation Behaviour and Loss-of-Life Modelling

Translating additional warning lead time into an estimate of lives saved requires a behavioural model of how people respond to a warning and a physical model of flood lethality. Lindell and Perry's Protective Action Decision Model frames evacuation as a sequential process of warning receipt, risk perception, decision delay, and physical transit time [14], while Jonkman, Vrijling, and Vrouwenvelder's review of flood loss-of-life methods provides empirical mortality relationships as a function of flood depth, flow velocity, and available evacuation time [15]. Because both the compliance behaviour of the affected population and the physical parameters of the flood are uncertain, Monte Carlo simulation is commonly used to propagate this uncertainty into a distribution of avoidable-fatality estimates, allowing the marginal benefit of a warning-lead-time improvement to be expressed as a confidence interval rather than a single deterministic figure.

F. Open-Source and Crowd-Sourced Intelligence

Beyond physical sensors and satellites, translingual news-extraction platforms such as GDELT [16] and citizen-reported distress signals offer near-instantaneous, geographically broad situational awareness, particularly during the early minutes of a disaster before official telemetry or satellite tasking can respond. These sources, however, carry substantially higher noise and lower positional accuracy than instrumented feeds, and disaster-informatics surveys consistently recommend that they be weighted as corroborating rather than primary evidence within a fusion framework [17].

III. INTEGRATION CHALLENGES AND DISCUSSION

Combining the components surveyed in Section II into a single operational pipeline raises difficulties beyond those of any individual component. Data heterogeneity is the most immediate: hydrometric telemetry, satellite products, seismic feeds, and text-based reports differ in latency by several orders of magnitude, from minutes for a river gauge to hours for a satellite scene, so a fusion engine must reason coherently about hazard state even when its constituent evidence streams are asynchronous and incomplete at any given decision instant.

A second, often underappreciated difficulty concerns the scientific validity of retrospective evaluation. When a disaster-response algorithm is benchmarked against a historical event, it is essential that only the information genuinely available at each simulated decision time be used

in scoring the algorithm's performance; using a satellite scene or news report that was not yet published at the moment being evaluated constitutes a form of hindsight or data-leakage bias that can make an algorithm appear more effective than it would have been in real operational conditions [18]. Establishing and enforcing such chronological validity, commonly through timestamped, cryptographically verifiable provenance records, is therefore a prerequisite for any credible claim of warning-lead-time improvement.

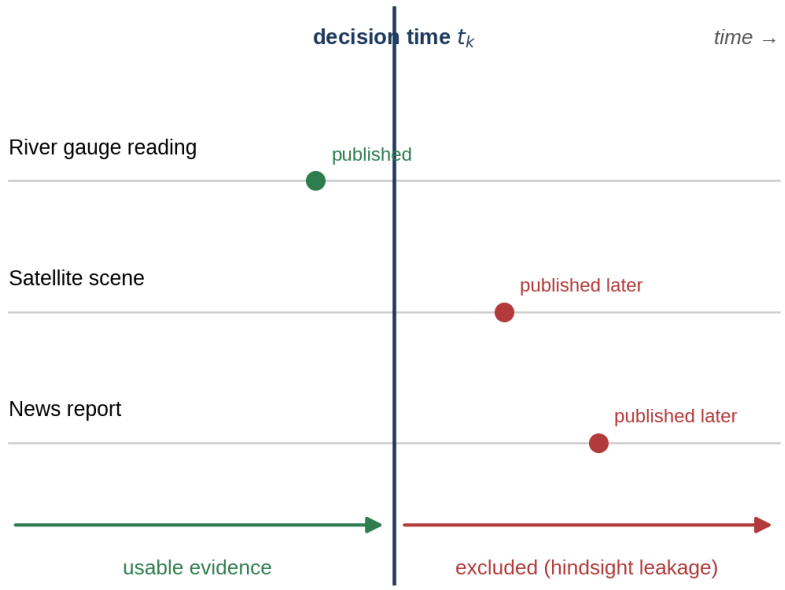


Fig. 5. Hindsight-free evaluation: only evidence published before the simulated decision time t_k may be used when scoring a retrospective benchmark.

A third challenge is computational tractability under real-time constraints. The MILP-based resource-allocation problems discussed in Section II-D are, in general, NP-hard integer programs; however, their practical instance sizes in a regional disaster (typically dozens to a few hundred demand nodes and a handful of resource categories) are generally considered within reach of modern branch-and-cut solvers for real-time use, though exact solve times depend on the solver, formulation, and hardware involved. The harder tractability problem lies instead in maintaining an up-to-date road-network graph as bridges and highway segments are progressively cut, since the allocation problem's feasible region itself changes during the response.

IV. A NOTE ON A RECURRING QUANTITATIVE RELATIONSHIP

One quantitative relationship recurs consistently across the wave-routing literature discussed in Section II-C and is worth restating explicitly. For a river reach located a distance d from a flood-triggering event occurring at time $t_{trigger}$, and for a locally calibrated wave celerity c, the expected arrival time of the flood crest follows the kinematic-wave relation

$$t_{arrival} = t_{trigger} + d / c . \quad (1)$$

Because the true celerity is uncertain and depends on channel geometry and bed roughness that can change during an extreme event, published applications of this relation typically evaluate c over a calibrated plausible range rather than at a single point estimate, yielding an arrival-time interval rather than a single predicted time [9].

V. CONCLUSION AND FUTURE DIRECTIONS

The preceding sections show that no single sensing or modelling technique addresses flash-flood early warning on its own. The most significant open problems that emerge from this review are the standardisation of data-sharing protocols across national hydrological, meteorological, and space agencies; the development of lightweight, edge-deployable models capable of operating during communication blackouts; the coupling of

physics-based kinematic routing with data-driven correction terms to improve accuracy in ungauged or morphologically unstable channels; and continued attention to reproducible, hindsight-free benchmarking so that reported improvements in lead time and life-safety impact can be trusted and compared across studies.

REFERENCES

[1] World Meteorological Organization, "WMO Atlas of Mortality and Economic Losses from Weather, Climate and Water Extremes (1970–2021)," WMO, Geneva, Switzerland, 2023.

[2] N. Otsu, "A threshold selection method from gray-level histograms," IEEE Trans. Syst., Man, Cybern., vol. 9, no. 1, pp. 62–66, 1979.

[3] S. Martinis, J. Kersten, and A. Twele, "Automatic near real-time flood detection in high resolution X-band synthetic aperture radar data using artificial neural networks," Remote Sens., vol. 7, no. 7, pp. 9360–9381, 2015.

[4] M. Chini, R. Pelich, L. Pulvirenti, N. Pierdicca, R. Hostache, and P. Matgen, "Sentinel-1 InSAR coherence for mapping floodwater in urban areas," Remote Sens. Environ., vol. 229, pp. 151–164, 2019.

[5] M. A. Clement, C. G. Kilsby, and P. Moore, "Multi-temporal Sentinel-1 and -2 data for rapid flood mapping and damage assessment," Int. J. Appl. Earth Obs. Geoinf., vol. 73, pp. 575–586, 2018.

[6] A. P. Dempster, "Upper and lower probabilities induced by a multivalued mapping," Ann. Math. Stat., vol. 38, no. 2, pp. 325–339, 1967.

[7] G. Shafer, A Mathematical Theory of Evidence. Princeton, NJ, USA: Princeton Univ. Press, 1976.

[8] N. M. Hunter, P. D. Bates, M. S. Horritt, and M. D. Wilson, "Benchmarking 2D hydraulic models for urban flooding," Proc. Inst. Civ. Eng.-Water Manage., vol. 161, no. 1, pp. 13–30, 2008.

[9] M. J. Lighthill and G. B. Whitham, "On kinematic waves II: A theory of traffic flow on long crowded roads," Proc. R. Soc. Lond. A, vol. 229, no. 1178, pp. 317–345, 1955.

[10] N. Altay and W. G. Green, "OR/MS research in disaster operations management," Eur. J. Oper. Res., vol. 175, no. 1, pp. 475–493, 2006.

[11] C. Perrigo, N. Altay, and E. P. Robinson, "Operations research applications in disaster response and emergency evacuation: A review," Eur. J. Oper. Res., vol. 292, no. 2, pp. 401–419, 2021.

[12] The Sphere Project, The Sphere Handbook: Humanitarian Charter and Minimum Standards in Humanitarian Response. Geneva, Switzerland: Practical Action Publishing, 2018.

[13] National Disaster Management Authority, National Disaster Management Guidelines: Management of Urban and Riverine Flooding. New Delhi, India: Govt. of India, 2020.

[14] M. K. Lindell and R. W. Perry, "The Protective Action Decision Model: Theoretical modifications and additional evidence," Risk Anal., vol. 32, no. 4, pp. 616–632, 2012.

[15] S. N. Jonkman, J. K. Vrijling, and A. C. W. M. Vrouwenvelder, "Loss of life caused by floods: An overview of methods and priorities for future research," Nat. Hazards, vol. 46, no. 1, pp. 35–56, 2008.

IEER Institute for Excellence in Education & Research