

Urban floods disrupt healthcare before, during, and after peak flooding

Mikhail Sirenko

`m.sirenko@tudelft.nl`

Delft University of Technology

Tina Comes

Deutsches Zentrum für Luft- und Raumfahrt e. V. (DLR)

Yared Abebe

Delft University of Technology

Nikola Stanić

Gemeente Rotterdam

Koen Tromp

Gemeente Rotterdam

Carlo Sobral de Vito

Gemeente Rotterdam

Tije Bakker

Gemeente Rotterdam

Farhan Meke

Gemeente Rotterdam

Kevin Su

Gemeente Rotterdam

Maria Pregnolato

Delft University of Technology

Article

Keywords:

Posted Date: July 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10243607/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Urban floods disrupt healthcare before, during, and after peak flooding

Mikhail Sirenko^{1,*}, Tina Comes², Yared Abayneh Abebe¹, Nikola Stanić³, Koen N. Tromp³, Carlo Sobral de Vito³, Tije M. Bakker³, Farhan Meke³, Kevin Su³, and Maria Pregnolato¹

^{*}*m.sirenko@tudelft.nl*

Abstract

Urban floods can increase demand for care while making functioning healthcare facilities harder to reach or creating a need for evacuation. Protecting urban health systems, therefore, depends not only on preventing facilities from flooding, but also on maintaining access to care as roads, services and population needs change throughout an event. Yet flood risk assessments commonly rely on maps of maximum flood extent or depth, which cannot show when healthcare disruption begins, how long it lasts, or which population groups and facilities face the greatest continuity of care challenge. Here we develop a spatio-temporal framework that links hourly flood dynamics, road network passability, healthcare facility exposure, population presence and neighbourhood health vulnerability. We apply it to a one-in-400-year rainfall scenario, resembling a real event that caused the Limburg 2021 flood, to Rotterdam, the Netherlands. We show that healthcare disruption follows a different timeline from flood extent: hospital access loss begins almost a day before maximum inundation, the first facility crosses an operational flood threshold 10 hours before that, and severe road and facility flooding peaks later. Neighbourhoods form a typology of access loss profiles, ranging from early, short, scattered losses to persistent, high losses that continued through the end of the 10-day simulation. At the facility scale, impacts form distinct timing-duration profiles, separating early-hit facilities requiring anticipatory action, a cascade-period group facing simultaneous evacuation and coordination challenges, and later-hit facilities requiring continued support during recovery. Our findings show that urban flood resilience requires maintaining equitable continuity of healthcare services throughout the flood onset, peak, and recovery phases.

Background

Floods create a triple challenge for urban health systems. They can increase demand for emergency care while simultaneously limiting residents' ability to reach functioning healthcare facilities and necessitating the evacuation of affected facilities (Aggarwal et al., 2025; Tariverdi et al., 2023; Abebe et al., 2025). Recent European flood events illustrate these coupled pressures. For example, the July 2021 floods in Limburg, the Netherlands, exposed gaps in evacuation decision-making (Jong et al., 2023; Wal et al., 2025). The 2024 floods in Valencia, Spain, disrupted access to critical healthcare services, particularly for vulnerable populations that required intensive care, while widespread flooding left many residents trapped and unable to reach healthcare facilities (Mas-Coma et al., 2025; Sancho et al., 2026). These events highlight that the resilience of urban health systems depends not only on protecting healthcare facilities but also on service continuity as flood conditions change over time (WHO, 2021).

Flood risk assessment has traditionally represented a flood event through a static map: often, maximum water depth or flood extent overlaid with exposed people, buildings, and infrastructure (Cea and Costabile, 2022). Such maps remain useful for identifying direct physical exposure and potential damage. For healthcare, however, they commonly translate into counts of flooded facilities or residents within the inundation extent (Antony et al., 2021). These types of assessment answer an important question, "where is the (maximum) floodwater?", but not the questions that matter for healthcare continuity and emergency decision-making: which residents can still reach care, which facilities remain reachable and functional, when disruption begins, how long it persists, and who is present when a facility is flooded. This is particularly important in urban floods because disruption is produced by interacting systems rather than by the water depth alone (Melnychuk et al., 2022). Flood impacts on healthcare must therefore be understood as a dynamic interaction among hourly flood hazard, flood-related changes in road network passability, facility-specific flood-exposure triggers, dynamic population presence, and neighbourhood health vulnerability. (Petricola et al., 2022; Alam et al., 2024; Lin et al., 2026).

Urban flood resilience is increasingly understood as the ability to preserve essential functions during disruption and recovery, rather than simply reducing the area or depth of inundation (Hammond et al., 2015; McClymont et al., 2020; Rözer et al., 2022). Dynamic flood models show that inundation evolves over time rather than appearing as a single maximum-depth flood extent (Thrysøe et al., 2021). Flood accessibility studies demonstrate that road disruption can substantially increase travel times to emergency and healthcare services, often with spatially uneven consequences (Yin et al., 2021; Gangwal and Dong, 2022; Papilloud et al., 2024). Equity-oriented

analyses further show that the same disruption can impose different losses across communities because populations differ in their resources, health needs, and ability to tolerate service loss (Dong et al., 2020). More recently, large-scale modelling has shown that flood-induced transport disruption can also create hidden pressure on hospitals located far outside the inundated area by redirecting patients towards the remaining accessible facilities (Wassmer et al., 2025).

Yet these elements remain only partly integrated. Accessibility studies have generally focused on travel time or emergency service reachability, without addressing the timing and duration of facility exposure or the number and types of people present when a facility gets flooded. In contrast, studies on healthcare facility evacuation show that patient condition, transport availability, receiving capacity, and coordination are critical but are usually based on retrospective interviews or optimisation scenarios rather than hourly flood simulations and occupancy estimates (Zhou et al., 2023; Kim et al., 2024). Elderly homes are especially important in this respect (Abebe et al., 2025). Their residents may require assisted movement, medication, medical equipment, and uninterrupted personal care, suggesting that even a short disruption can create substantial operational demands (Landeg et al., 2019; Abebe et al., 2025).

To address this gap, we propose a multi-scale spatio-temporal framework for analysing healthcare service continuity during urban floods caused by extreme rainfall. The framework links five layers: hourly flood hazard, flood-related changes in road network passability, facility-specific flood exposure triggers, dynamic population presence, and neighbourhood health vulnerability. We demonstrate the framework by simulating an extreme rainfall scenario with a 400-year recurrence interval in Rotterdam, the Netherlands, over a 240-hour simulation period.

The contribution of this study is threefold. First, we move beyond static flood exposure by tracing how flood extent, road disruption, healthcare access, and facility exposure evolve. Second, we quantify both neighbourhood-level vulnerability-weighted hospital access loss and facility-level potential evacuation demand. Third, we identify distinct temporal profiles of neighbourhood access loss and facility exposure, showing how the timing and duration of disruption can inform emergency planning and response. Together, these contributions shift attention from static flood exposure towards spatio-temporal health service continuity, focusing not only on where floodwaters are deepest, but also on when, where, and for whom access to functioning healthcare is most at risk.

Results

We examine the impacts of an extreme rainfall event on access to healthcare facilities as a dynamic urban disruption unfolding across three connected scales. At the city scale, we trace the temporal evolution of flood extent, road network disruption, and healthcare facility exposure. At the neighbourhood scale, we show how these disruptions translate into distinct neighbourhood-level profiles of hospital access loss. At the facility scale, we estimate potential evacuation demand by accounting for the timing of inundation, expected occupancy at the onset of flooding, and the duration of flood exposure.

Flood extent, road disruption, and healthcare facilities exposure peak at different times

The simulated flood event does not produce a single peak across the urban system. Flood extent, road disruption, healthcare accessibility loss, and healthcare facility exposure follow different temporal trajectories (Figure 1). Therefore, the implications of the event cannot be inferred solely from the maximum flooded area.

The flood extent expands in distinct phases. We see the first accumulation of water within the first 5-7 hours, after which the inundated area increases rapidly. The maximum flood extent, 12.8 km², occurs at hour 30 (Figure 1A). The footprint then remains close to this maximum until approximately hour 42, despite rainfall continuing until hour 48, and then gradually recedes over the following days. This peak represents the maximum in spatial extent, but it does not coincide with the peak in road disruption, loss of healthcare access, or healthcare facility exposure.

Mobility impacts emerge much earlier, while the total flooded area is still small. At hour 5, one neighbourhood already experiences an increase in travel time to its nearest functioning hospital. This occurs 25 hours before the maximum flood footprint. Such an early disruption may arise when a small number of critical road links are affected, even when the total inundated area remains small. Figure 1B shows that the length of roads affected by any water increased rapidly, while the length of roads with water depths exceeding the vehicle passability threshold of 0.30 m (Pregolato et al., 2017) reaches its peak later. At hour 36, 249 km of the drive network (approximately 10%) exceeds this threshold, which is six hours after the flood extent reached its maximum.

Healthcare facility exposure follows a third trajectory. Under our modelled operational evacuation trigger (i.e., when more than 20% of the 5 m facility ring flooded above 0.30 m), the first trigger occurs at an elderly home at hour 20 with an estimated 92 people present. This is 10 hours before the maximum flood footprint. Healthcare

facility exposure intensifies afterwards. The 90th percentile of water depth surrounding healthcare facilities reaches 0.94 m at hour 38 (Figure 1C), indicating a substantial temporal lag between peak flood extent and peak facility exposure

Taken together, these trajectories show that urban flood impacts are temporally staggered rather than concentrated around a single moment of peak flooding. Loss in healthcare accessibility begins at hour 5; the first facility is directly affected at hour 20; the flood extent peaks at hour 30; and severe road and facility flooding peaks only later. This sequencing creates a window for targeted response, such as warning affected neighbourhoods, protecting critical road links, or preparing facilities for evacuation, before the event reaches its most extensive stage.

At the same time, receding water should not be equated with immediate recovery of healthcare access. Physical inundation may subside faster than the roads do; affected facilities are likely to remain closed longer than the event lasts; evacuated patients and residents will still require care elsewhere (Hertelendy et al., 2025; Arefi et al., 2025). Flood response and recovery should therefore be organised around health service continuity rather than a single snapshot of maximum depth or inundation.




Figure 1: Temporal evolution of flood dynamics and its impact. (A) Flood footprint, measured as the area affected by any floodwater. (B) Length of the drivable road network intersecting any water (grey) and exposed to water depths of at least 0.30 m (red), the threshold used to classify roads as inaccessible to ordinary cars. (C) The 90th percentile of simulated water depth within 5 m rings around healthcare facilities. A facility crosses the operational flood trigger when more than 20% of its ring is inundated to a depth of at least 0.30 m. Vertical markers indicate key timings, and the shaded area denotes the rainfall period (hours 0-48). All times are measured from flood onset.

Neighbourhood access loss follows distinct spatio-temporal profiles

The staggered disruption of roads and health care facilities leads to different patterns of hospital access loss across neighbourhoods. At any given hour, a neighbourhood’s loss depends on the combined effect of road disruption, which alters the available travel routes; hospital availability, which changes the set of facilities that can be reached; and the neighbourhood’s population and underlying health vulnerability. Therefore, the resulting loss varies in severity, timing, duration, and concentration.

We characterise these differences using hourly vulnerability-weighted hospital access loss:

$$VWAL_{n,t} = P_n \times V_n \times \Delta\tau_{n,t}$$

where P_n is the neighbourhood population, V_n is the composite neighbourhood health vulnerability score ranging from 0 to 1, and $\Delta\tau_{n,t}$ is the additional travel time, in hours, required to reach the nearest open hospital relative to normal conditions. And, as a result, $VWAL$ units are vulnerability-weighted person-hours. $VWAL$ measures population-level deterioration in potential access to healthcare, rather than the number of realised hospital trips.

Clustering the hourly $VWAL$ trajectories results in a neighbourhood typology of access loss (Figure 2). While the resulting four profiles are descriptive, neighbourhoods assigned to the same group still differ in population, vulnerability, and absolute increase in travel time. Their value is distinct in temporal and spatial patterns through which hospital access deteriorated over the course of the event.




Figure 2: Spatio-temporal profiles of neighbourhood-level vulnerability-weighted hospital access loss. Lines show the aggregate hourly loss for each group, while heatmaps show hourly losses for the neighbourhoods displayed in each panel. The vertical dotted line marks the citywide peak in hospital access loss at hour 42. Colours use a shared scale and are clipped at the 95th percentile; colour differences among the darkest cells should therefore not be interpreted as proportional differences in loss. For readability, the panels display the neighbourhoods with the largest cumulative losses in each group; group totals and trajectories nevertheless include all neighbourhoods assigned to that group.

Profile A contains 24 neighbourhoods with *early, short, scattered losses*. Loss begins at hour 6, peaks at hour 30, and largely disappears by hour 53. Despite its broad geographical distribution, this profile contributes only 41,716 vulnerability-weighted person-hours of cumulative loss. Its mean vulnerability is moderate 0.497 (Table 1), but the range is wide, indicating that the profile includes both high- and low-vulnerability neighbourhoods. The defining feature is therefore not low vulnerability, but the short duration and fragmented geography of the disruption. These neighbourhoods experienced an early accessibility shock, primarily due to the temporary deterioration of

surrounding road access.

Profile B has six neighbourhoods with *low-intensity, sustained losses*. Losses emerge at hour 18, peak at hour 41, and remain until hour 140. Yet their total cumulative burden is only 2,257 vulnerability-weighted person-hours. This is partly consistent with the profile’s comparatively low vulnerability profile, with a mean score of 0.319 and no neighbourhood exceeding 0.455. Profile B demonstrates that persistence alone does not determine impact: access can remain degraded for several days while producing only modest additional travel times and affecting comparatively fewer vulnerable residents.

Profile C represents *peak-centred, concentrated losses*. Across 23 neighbourhoods, losses rise rapidly from hour 18, peak at hour 42, coinciding with the citywide maximum in hospital access loss, and remain until approximately hour 183. The profile contributes 1.21 million vulnerability-weighted person-hours, more than half of the cumulative loss across all four profiles. Its mean vulnerability is moderate-to-high (0.538), although the broad range shows that high losses are not restricted to the most vulnerable neighbourhoods. Instead, the pattern reflects the convergence of substantial road disruption, large exposed populations, and elevated vulnerability around the citywide peak in access loss. In several neighbourhoods, the disruption extends beyond a longer detour to an alternative hospital and includes periods when no hospital could be reached via the road network.

Profile D forms a *persistent high-loss tail*, concentrated primarily in the southern part of the study area. Loss begins at hour 29, peaks almost immediately at hour 30, and remains substantial for the rest of the 240-hour simulation. This profile has the highest average vulnerability score (0.639), and even its least vulnerable neighbourhood has a score of 0.304. It produces 1.02 million vulnerability-weighted person-hours of cumulative loss. Unlike Profile C, Profile D does not primarily represent a short, citywide shock. Rather, it reflects a prolonged service-continuity problem in which residents remain dependent on more distant alternatives after nearby hospital capacity is lost or inaccessible (see Figure 3). Because the simulation ends while these losses remain active, the duration of this loss might be understated rather than fully observed.

Profile	n	Population (thousands)	Timing (h)			Vulnerability Mean [min–max]	Total VWAL (millions)
			Onset	Peak	Recovery		
A — Early, short, scattered	24	212.2	6	30	53	0.497 [0.164–0.915]	0.042
B — Low-intensity, sustained	6	16.4	18	41	140	0.319 [0.189–0.455]	0.002
C — Peak-centred, concentrated	23	224.5	18	42	183	0.538 [0.064–0.864]	1.214
D — Persistent high-loss tail	13	98.8	29	30	240	0.639 [0.304–0.897]	1.024

Note: Total VWAL denotes cumulative vulnerability-weighted access loss of the profile’s neighbourhoods. Vulnerability is reported as the mean across neighbourhoods in each profile, with minimum and maximum values in brackets.

Table 1: **Timing, population, vulnerability, and total vulnerability-weighted access loss across neighbourhood loss profiles.**




Figure 3: **Spatial distribution of neighbourhood level vulnerability-weighted hospital access loss profiles.** Colours indicate the profile. Hatched neighbourhoods experienced no loss of hospital access during the simulation period. Hospital markers distinguish facilities that were flooded from those that remained operational.

Potential evacuation demand depends on timing, duration, and occupancy

Facility-level consequences depend not only on whether a healthcare facility is exposed to flooding, but also on the timing and duration of that exposure, as well as the number of people present when flood conditions become operationally critical. We estimate potential first-hit evacuation demand as the number of people (i.e., patients or residents, staff, and visitors) present when a facility first exceeds the operational flood threshold, defined as 20% of the facility's 5 m ring being flooded with water more than 0.30 m. This metric represents a hypothetical planning scenario designed to identify the scale and timing of potential evacuation needs. It should not be interpreted as a prediction that every exposed facility would close or require evacuation.

Under the default timing assumption, with rainfall beginning on Monday at 00:00, the first facility crosses this threshold at hour 20 or Monday 20:00 (see Figure 4). This is an elderly home with an estimated 92 people present. Two other elderly homes are first affected at hour 23, bringing the cumulative potential evacuation demand to approximately 251 people before midnight.

The main escalation occurs between hours 27 and 32, corresponding to Tuesday 03:00-08:00 under the assumed event timing. During this six-hour window, 11 of the 15 affected facilities first cross the flood threshold. Together, they account for approximately 1,507 of the 1,881 people estimated to require evacuation during the event, or about 80% of the total demand. The largest single step occurs when a hospital becomes flooded at hour 30, adding an estimated 436 people.

During the early morning, hospitals and elderly homes are occupied primarily by patients and residents rather than visitors. Of the 1,881 people estimated to be present at the time of impact, 1,308 are patients or residents of elderly homes likely to require assisted movement, while 573 are staff and visitors.

The final facility is flooded at hour 39, when another elderly home crosses the threshold, adding approximately 123 people to the potential evacuation demand. This later impact illustrates that the evacuation problem does not end when the main impact window has passed.




Figure 4: **Cascade of modelled evacuation demand as healthcare facilities cross the flood trigger.** The upper panel shows the cumulative estimated number of people present at affected facilities, separated into patients/residents likely to require assistance and staff/visitors. The lower panel shows the new evacuation demand entering each hour. The shaded period, hours 27-32, accounts for approximately 80% of the total estimated demand. The vertical line marks the first impact on the hospital at hour 30. The building is considered flooded when 20% of its building ring is flooded to a water level of 0.30 m or more.

Figure 5 translates the sequence of impacts on facilities into a typology of affected facilities. We classify each facility according to two dimensions: the timing of first flood exposure and the flood duration exceeding the threshold. Facilities are grouped into an early-hit phase (hours 20-27), a cascade phase (hours 28-31), and a late-hit phase (hours 32-39). These are combined with three persistence classes: brief exposure (≤ 24 h), sustained exposure (25-48 h), and prolonged exposure (49-72 h).

This typology makes explicit the potential need for different preparedness and response actions depending on the type of facility (Table 2). *Early-prolonged* facilities may be the highest priority for anticipatory action because they are affected before the main cascade and may require extended service substitution or relocation. *Cascade-period* facilities generate the most acute coordination challenge because several sites may require transport, staffing, and receiving capacity at nearly the same time. *Late-brief* facilities may require targeted short-term support, whereas *late-intermediate* facilities can form a second wave of disruption that outlasts the acute response phase. The spatial distribution of facilities also shows that these profiles are spatially dispersed across the city, meaning that response capacity cannot be concentrated in a single part of Rotterdam.




Figure 5: **Timing, duration, and potential evacuation demand at flood-affected healthcare facilities.** Circle size represents the estimated number of people present when each facility first crosses the flood trigger. Colours classify facilities jointly by the timing of first trigger crossing (early: hours 20-27; cascade: 28-31; late: 32-39) and the duration of threshold-exceeding flooding (brief: ≤ 24 h; sustained: 25-48 h; prolonged: 49-72 h).

First hit	Duration	Facilities	Hospitals	Elderly homes	People at first hit
Early (20–27 h)	Brief (≤ 24 h)	0	0	0	0
	Intermediate (25–48 h)	2	0	2	171
	Prolonged (49–72 h)	2	0	2	157
Cascade (28–31 h)	Brief (≤ 24 h)	4	1	3	795
	Intermediate (25–48 h)	5	0	5	589
	Prolonged (49–72 h)	0	0	0	0
Late (32–39 h)	Brief (≤ 24 h)	1	0	1	47
	Intermediate (25–48 h)	1	0	1	123
	Prolonged (49–72 h)	0	0	0	0

Note: People at first hit denotes the estimated number of people present when facilities in each category first cross the flood trigger.

Table 2: **Healthcare-facility exposure typology by first-hit timing and flood duration.**

The timing-sensitivity analysis reinforces the importance of hourly occupancy patterns (Table 3). Shifting the flood timeline and daily occupancy by only four hours increases potential first-hit evacuation demand from 1,881 to 2,518 people (+34%). Nearly all of this increase comes from staff and visitors, whose estimated numbers more than double, while the numbers of patients and residents remain broadly unchanged. Event timing, therefore, changes not only the total number of people involved but also the composition of the potential evacuation operation: a daytime event may require much greater crowd management and staff coordination, whereas an overnight event concentrates the challenge on assisted evacuation and continuity of care.

Indicator	Current model timing	Busiest-hour ment	align-	Difference
<i>Timing profile</i>				
Clock shift	+0 h	+4 h		+4 h
First trigger profile hour	20:00	00:00		—
Peak profile hour	07:00	10:00		+3 h
80% surge window	03:00–08:00	07:00–11:00		Later and shorter
<i>Evacuation demand</i>				
People needing evacuation	1,881	2,518		+637 (+34%)
Patients/residents needing assistance	1,308	1,294		−14 (−1%)
Staff/visitors	573	1,225		+652 (+114%)
Largest one-hour demand	679	1,051		+372 (+55%)

Note: The busiest-hour alignment shifts the flood timeline by four hours to coincide more closely with periods of high staff and visitor presence. Demand denotes the number of people who require evacuation when a facility is first affected.

Table 3: **Sensitivity of evacuation demand to alternative event timing.**

These estimates remain scenario-based: we assume facilities to keep their baseline occupancy until they first cross the flood trigger. They should therefore be interpreted as potential demand for planning, rather than a prediction that every affected facility would necessarily be closed or would need to be evacuated. The central result is that healthcare flood risk produces a sequence of facility-specific decisions about evacuation, staffing, receiving capacity, and continuity of care, rather than a single peak-time emergency.

Discussion

This study extends the emerging healthcare continuity research by jointly tracing hourly flood dynamics, road network disruption, hospital accessibility, neighbourhood health vulnerability, and facility exposure and occupancy over space and time. By doing so, we address a gap between urban flood modelling, transport accessibility analysis, and health emergency preparedness and response, which are often disconnected (Dorfer et al., 2025; Abebe et al., 2025).

The first main finding is that healthcare disruption does not peak at the same time as flood extent. In the simulated event, hospital access losses begin at hour 5, the first facility crosses the flood trigger at hour 20, and the maximum flood extent occurs at hour 30, with more serious road disruption and facility exposure intensifying

later. This temporal mismatch shows why maximum inundation maps alone cannot identify the timing of the most consequential healthcare disruptions. A relatively small flood footprint can already disrupt a critical road link and lengthen travel to a hospital, while facility exposure and the loss of nearby destinations can create access loss that persists even after floodwater recedes.

This result is consistent with the broader flood resilience and transport literature, which shows that disruption unfolds through multiple temporal trajectories rather than as a single-peak effect (Hammond et al., 2015; Rebally et al., 2021; Long et al., 2025). However, existing work has rarely linked these trajectories directly to changing access to functioning healthcare facilities at fine urban resolution. Transport studies have often focused on travel time disruption or network performance, whereas healthcare flood studies have focused more on direct facility impacts, evacuation, or broader health system pressures (Rebally et al., 2021; Abebe et al., 2025; Dorfer et al., 2025). Our analysis connects these perspectives by showing that floodwater, road passability, hospital accessibility, and facility exposure can follow different trajectories within the same event. The practical implication is that the response should be organised around a timeline of deteriorating service continuity rather than a single snapshot of the flood peak.

This temporal perspective also identifies an opportunity for anticipatory action. Because road access deteriorated before the maximum flood extent and before most facilities crossed the flood trigger, the period before visible peak flooding may be the most useful time for intervention. Forecast-informed preparedness could include monitoring critical access routes, warning affected neighbourhoods, restricting non-essential admissions and visits, reinforcing staff coverage, preparing alternative transport, and confirming receiving capacity at less exposed facilities. Such measures do not eliminate flood risk, but they may reduce the likelihood that a manageable disruption becomes a healthcare access or evacuation crisis. This supports calls in resilience research to focus not only on flood resistance, but also on the capacity to adapt and maintain essential functions as conditions evolve (McClymont et al., 2020; Rözer et al., 2022).

The second main finding is that losses in healthcare access are not spatially or temporally uniform. Instead, they form 4 descriptive profiles of vulnerability-weighted hospital access loss: early and short, scattered disruption; low-intensity but sustained disruption; peak-centred, concentrated disruption; and a persistent, high-loss tail. In contrast to analysing changes in travel time or vulnerability alone, these profiles allow us to examine the combined effects of urban floods on neighbourhoods. They combine the number of residents, underlying health vulnerability, road network disruption, resulting changes in travel time, and hospital availability. They show that the timing, duration, and concentration of disruption matter as much as the maximum additional travel time.

This finding reinforces existing evidence that flood-related mobility losses and underlying vulnerability are spatially heterogeneous (Moreira et al., 2021; Hoover and Tate, 2025; Ermagun et al., 2026). Previous research has shown that vulnerability can vary across preparedness, response, and recovery phases, and that the same hazard may impose different burdens across places and population groups (Moreira et al., 2021; Stimpson et al., 2025). Our contribution is to demonstrate how this heterogeneity appears in spatio-temporal healthcare access during urban floods caused by extreme rainfall. In particular, the persistent high-loss profile shows that the most challenging service-continuity problem is not necessarily the neighbourhood experiencing the largest short-lived delay at the citywide peak. It may instead be the area where access remains degraded long after the most visible phase of flooding has passed.

This distinction has implications for emergency preparedness and response. Neighbourhoods facing acute, peak-centred disruption may require protected emergency routes, traffic management, temporary transport support, and targeted communication for residents temporarily cut off from hospital care. Persistent high-loss areas require a different response: sustained access to alternative facilities, continuity of outpatient and chronic care, support for residents with limited mobility, and restoration of local service availability. A single citywide intervention is therefore unlikely to address the different forms of access loss. More broadly, the results suggest that equitable flood resilience requires attention not only to who is exposed but also to who experiences more prolonged difficulties in accessing care.

The third main finding covers healthcare facilities themselves. Existing reviews show that flooding can disrupt facilities in multiple ways, including inundation, loss of utilities, interrupted transport access, staff shortages, and changes in care-seeking (Abebe et al., 2025; Dorfer et al., 2025; Valente et al., 2026). Our analysis adds a temporal dimension to this literature by estimating when facilities first cross an operational flood trigger, how long trigger-exceeding conditions persist, and how many people may be present at that moment. The results indicate that potential evacuation demand is not only a question of how many facilities are affected, but also of when impacts occur, how long they last, and who is present.

From an emergency planning perspective, the concentration of first impacts during a six-hour cascade period is particularly relevant. Approximately 80% of the estimated potential first-hit evacuation demand occurs during this short window. Moreover, changing the assumed event timing by 4 hours increased the estimated demand by 34%, largely due to increased staff and visitor presence. We do not predict actual evacuation demand, facility closure, or patient outcomes. Instead, we demonstrate that the operational burden associated with a flood can vary substantially depending on the time of day facilities are affected. A daytime event may require greater crowd

management and staff coordination, whereas an overnight event may concentrate the challenge on assisted transfers of patients and residents.

These results support framing facility emergency preparedness as service continuity rather than evacuation alone. Crossing the flood trigger should not automatically be interpreted as requiring immediate closure or full evacuation. Rather, it should guide a structured assessment of whether admissions should be reduced, visits restricted, staffing reinforced, resources pre-positioned, or patients and residents transferred before routes and receiving capacity become constrained. Our facility timing-duration typology can support this process. Early-prolonged facilities may warrant anticipatory action because they are affected before the main cascade and may face extended disruption. Cascade-period facilities create the most acute coordination challenge because several sites may require transport, staffing, and receiving capacity at nearly the same time. Late-hit facilities may require targeted support after the initial response surge, particularly where disruption extends into the recovery period.

Our findings also draw attention to elderly homes. 14 of the 15 facilities that crossed the modelled trigger are elderly homes. This does not mean that elderly homes are universally more flood-prone than hospitals or that all affected homes would require evacuation. It does, however, identify long-term care as a major planning concern. Residents of elderly homes may require assisted movement, medication or equipment, continuous supervision, and continuity of personal care. Previous reviews of flood impacts on health systems identified care continuity for older populations and for non-hospital facilities as underrepresented areas of research (Abebe et al., 2025; Dorfer et al., 2025; Stimpson et al., 2025). Our analysis addresses this gap by demonstrating how dynamic facility-level exposure analysis can help identify where and when these continuity challenges may emerge.

Several limitations should guide interpretation. First, the study examines a modelled extreme rainfall scenario rather than a historical event in Rotterdam. It is, however, based on a historical event in Limburg, the Netherlands, in 2021. The sequence and duration of disruption can differ with shorter rainfall events, riverine floods, compound flooding, or events occurring at different times of day. Second, vulnerability-weighted hospital access loss represents a modelled deterioration in potential access, rather than realised hospital trips or unmet need. The metric is intended to identify where access conditions deteriorate most severely, not to estimate how many people would actually seek hospital care during the event. Third, the analysis treats Rotterdam as a closed system. In practice, some residents may seek care in surrounding municipalities, although these alternatives may be more distant, difficult to reach during a flood, or have capacity constraints of their own. The analysis also focuses on access to hospitals with cars. Other modes, including walking, cycling, public transport, or ambulances, may alter accessibility. Public transport may be particularly vulnerable to widespread disruption, while ambulances may retain access under conditions that prevent regular car travel. Fourth, estimates of potential evacuation demand represent the number of people present when a facility first crosses the modelled flood trigger; they do not account for patient condition, transfer needs, ambulance availability, willingness and ability to travel, receiving facility capacity, or actual healthcare building functionality. Fifth, the results depend on assumptions about flood depths, road passability, speed reductions, the 0.30 m vehicle threshold, the flood trigger based on flooding within a 5 m ring, and the timing of modelled extreme rainfall scenario onset. These assumptions provide a transparent way to examine relative patterns of disruption, but they should not be interpreted as direct evidence of road closures, building flooding, or mandatory evacuations.

In conclusion, we argue for positioning urban flood resilience of healthcare as a dynamic service continuity problem. Urban flood resilience cannot be assessed solely by analysing maximum flood extent maps, counts of exposed facilities, or peak travel times. It must consider whether people can maintain equitable access to functioning care as floodwater, road passability, facility availability, and population needs and presence change over time. By integrating these processes across city, neighbourhood, and facility scales, our framework identifies when access losses begin, where they become acute or persistent, which facilities may require early action, and how potential evacuation needs may cascade. This provides a basis for moving health emergency preparedness and response beyond using static exposure towards maintaining equitable access to essential services, such as healthcare, throughout the onset, peak, and recovery phases.

Methods

Analytical framework

We analyse healthcare continuity during an urban flood as a linked sequence of hazard, mobility, facility exposure, population presense and health vulnerability. Our framework has three scales. At the city scale, we combine hourly flood-depth rasters with the road network and healthcare facility footprints to identify when roads become slower or impassable and when facilities cross an operational flood trigger. At the neighbourhood scale, we use changes in the road network and the set of open hospitals to estimate how much longer residents need to travel to access care, and whether this loss falls on neighbourhoods with higher health vulnerability. At the facility scale, we combine the time at which each healthcare building first crosses the flood trigger with expected hourly occupancy

to estimate potential first-hit evacuation demand.

Study area

We model Rotterdam, the Netherlands, focusing on the municipality's central, most densely populated part. The study area contains 75 of Rotterdam's 94 neighbourhoods and had a combined population of 604,340 residents in January 2026. The flood model is organised into 11 hydraulic domains, which together cover 71.9 km². These domains represent the main drainage and surface-water areas used in the municipal hydrodynamic model.

The healthcare system included in the analysis contains 58 facilities: 4 hospitals and 54 elderly care homes (Figure 6). Following calls from previous studies (Abebe et al., 2025; Dorfer et al., 2025; Stimpson et al., 2025), we include elderly homes in the analysis, as flood disruptions in long-term care can create substantial continuity of care needs even when hospitals remain open. Residents of elderly homes may require assisted movement, medication or equipment, continuous supervision, and therefore cannot be treated as equivalent to the general population.




Figure 6: **Study area, hydraulic domains, neighbourhoods and healthcare facilities.** The analysis covers 75 Rotterdam neighbourhoods located within 11 hydraulic domains. Blue crosses show the four hospitals, and red circles show the 54 elderly homes.

Input data

The analysis combines 5 datasets: flood model outputs; a routable drive network; facility locations and their building footprints; neighbourhood population and health vulnerability indicators; and facility occupancy estimates

(Table 4).

Dataset	Spatial/temporal resolution	Role in analysis
Rainfall scenario	48 h event profile	Flood-model forcing for the Limburg-analogue extreme rainfall event
InfoWorks ICM depth rasters	2 m; hourly; 0–240 h	Hourly flood hazard for hydraulic areas, roads and facility rings
OpenStreetMap road network	Directed road edges	Depth-adjusted shortest-path routing from neighbourhoods to facilities
Healthcare facilities and footprints	Building level	Facility location, type and flood-trigger sampling around each building
Population and health indicators	Neighbourhood level	Resident exposure and composite Health Vulnerability Index
Facility capacity and visitation data	168 h weekly profiles	Hourly occupancy and potential first-hit evacuation demand

Table 4: **Datasets and their roles in the spatio-temporal healthcare continuity analysis.**

Extreme rainfall scenario and hydrodynamic model

This study uses a single modelled extreme rainfall event as input to the integrated urban water model developed in InfoWorks ICM. The event replicates a return period of approximately 400-500 years, similar to the extreme Limburg flooding event in July 2021 (Kok et al., 2023). The Limburg event lasted about 48 hours and produced a cumulative rainfall depth of roughly 200 mm. Unlike short cloudburst events, however, it was characterised by relatively low peak rainfall intensities, reaching approximately 14 mm/hr, but long duration (Figure 7).




Figure 7: **Rainfall event profile: July 2021 Limburg extreme rainfall analogue applied to Rotterdam.** Blue bars show hourly rainfall intensity; the red line shows cumulative depth. The event is characterised by low peak intensity (approximately 14 mm/h) sustained over 48 hours, producing roughly 200 mm of total rainfall. Estimated return period: 400-500 years.

To simulate the hydrodynamic response of Rotterdam’s urban water system to the modelled extreme rainfall, we develop flood models using the InfoWorks Integrated Catchment Models (InfoWorks ICM). The models capture the urban hydrological system, including surface water bodies, sewer networks, topography, land use, and simplified groundwater processes 8. The software employs a coupled one- and two-dimensional (1D–2D) approach, in which 1D modelling simulates flows in pipes and channels, and 2D modelling simulates the propagation of surface water across the terrain. This enables the integrated simulation of above- and below-ground drainage systems within a single modelling framework. For a more complete overview of the flood model, please see the Supplementary Information.




Figure 8: **Schematic representation of the coupled urban flood model.**

The model discretises the city into 11 computational domains (hydraulic areas) that together cover the whole study area. For each domain, the model produces two-dimensional water-depth rasters at a horizontal resolution of 2 m and at hourly intervals. We assemble a single seamless depth surface for each hour by mosaicking the per-domain rasters. The resulting time series spans 241 hourly snapshots, from the onset of rainfall ($t = 0$) through to $t = 240$ h (10 days), capturing both the inundation and the subsequent recession of water.

Road network and healthcare facilities

We use Rotterdam’s drivable road network from OpenStreetMap, processed with the OSMnx Python package (Boeing, 2025) on 15 May 2026. After clipping to the flood-model footprint, the routable network contains 13,652 nodes and 30,324 directed edges. We assign a single urban free-flow speed of 30 km/h to every edge as a modelling assumption based on Ligterink (2016). This avoids inconsistent or missing OpenStreetMap maxspeed tags and ensures that baseline and flooded travel times differ only because of flood-depth effects.

We collect facility locations from two openly available national data portals: Volksgezondheid en Zorg (VZ-info.nl), maintained by the National Institute for Public Health and the Environment (RIVM, <https://www.vzinfo.nl/>), and Zorgkaart Nederland (<https://www.zorgkaartnederland.nl/>), maintained by the Netherlands Patients Federation (Patiëntenfederatie Nederland). Building footprints are taken from Overture Maps using the 15 April 2026 data dump.

Population and neighbourhood health vulnerability

We obtain the neighbourhood population and age data from Statistics Netherlands’ Wijk- en buurtkaart 2025. For health indicators, we use the 2024 RIVM Gezondheid per Wijk en Buurt survey. We select seven indicators that capture conditions likely to reduce people’s ability to cope with healthcare access disruption: share of residents aged 65 and over, chronic conditions, severe overweight, psychological complaints, risk of anxiety or depression, low personal resilience or limited sense of control, and severe or very severe loneliness 9.




Figure 9: **Neighbourhood indicators used in the Health Vulnerability Index.** Histograms show the distribution of the seven input indicators across the 75 study neighbourhoods. Indicators are grouped into demographic, physical health, mental health, and social domains.

Facility occupancy data and its estimation

We estimate facility occupancy as expected hourly presence over an average 168-hour week. The aim is not to reconstruct real-time occupancy, but to estimate how many patients, residents, staff and visitors are likely to be present if a facility first becomes operationally disrupted at a given hour.

Baseline anchors for hospitals include bed capacity, staff headcount and annual or daily patient activity. We collect these from public national healthcare portals and provider sources, including VZinfo/RIVM, Ziekenhuischeck, hospital annual reports and hospital websites. For elderly homes, we collect resident counts from provider websites. All four hospitals used in the accessibility analysis have facility-specific capacity and activity anchors. Facility-specific resident counts are available for 18 of the 54 elderly homes; for the remaining 36, we use the city-level average.




Figure 10: **Weekly occupancy profiles used to estimate hourly presence.** Lines show unit-mean multipliers for hospitals and elderly homes. Static populations remain close to one throughout the week, while staff and visitor-related populations vary strongly by hour and weekday.

We derive weekly activity shapes for hospitals from parking data, and use Google Popular Times data for elderly homes. Healthcare shift assumptions for staff, inpatients and elderly home residents are near constant.

This presence model separates relatively static groups from groups that vary strongly by hour. Hospitals are represented by inpatients, staff and outpatients/visitors. Elderly homes are represented by residents, staff and visitors. Each group has a deterministic, unit-mean weekly profile and a stochastic scale factor for its average headcount. We run 1,000 Monte Carlo realisations with a fixed seed and summarise each facility-hour using the 10th, 50th and 90th percentiles. The full estimation algorithm, dwell-time assumptions, priors and checks are reported in the Supplementary Information Section. The median estimate (p50) is used in all downstream calculations.

Spatio-temporal modelling of flood-related healthcare service continuity

Road passability and travel times

For every hourly flood snapshot, we sample water depth along each road edge. We densify edge geometries at 2 m intervals, and query the depth raster at each sample point using a nearest-cell lookup. We record summary statistics for every edge-hour, including maximum, mean and 90th-percentile depth, the number of wet sample points, and the wet fraction.

Because pluvial flood rasters report depth at ground level, bridges and tunnels require special handling. For bridge edges, we sample the two endpoints of the edge and treat the bridge as flooded only if both approaches are wet; otherwise, we treat the bridge deck as dry. For tunnel edges, we assign water depth to the tunnel only if both portals are wet; if both are wet, we set the effective tunnel depth to the greater of the two portal depths.

We classify roads into depth bands: dry (< 0.01 m), pedestrian-passable (0.01-0.15 m), car-passable (0.15-0.30 m), emergency-vehicle only (0.30-0.60 m) and impassable (≥ 0.60 m). For routing, however, we use the continuous depth-speed relationship from Pregnolato et al. (2017). The curve is evaluated per edge using the effective maximum edge depth, capped at the urban free-flow speed and expressed as a speed multiplier between 0 and 1. Passenger-car routing treats edges as impassable at 0.30 m and assigns them infinite travel time for that hour (Figure 11).




Figure 11: **Depth-speed functions for passenger-car and ambulance routing.** Passenger-car speeds follow the continuous depth-disruption function from Pregolato et al. (2017), capped at 30 km h^{-1} , and roads become impassable at 300 mm.

We assume that roads become passable again as soon as the modelled water depth falls below the relevant threshold. This is optimistic because mud, debris, vehicle blockage or inspection requirements could keep roads closed after floodwater recedes. We did not add a separate flow-velocity term to the routing model; road disruption is represented through depth-dependent passability and speed reduction.

We define accessibility as the travel time from each neighbourhood to the nearest operational facility of a given type. Each neighbourhood is represented by all road-network nodes within its polygon, rather than by a single centroid (Figure 12). We compute travel times from all sample nodes to all candidate facilities and aggregate by the median to represent a typical resident. If more than half of a neighbourhood’s sampled nodes cannot reach any facility, the neighbourhood is recorded as cut off for that facility type. The flooded accessibility analysis is run both for a variant in which facilities that have crossed the flood trigger are removed from the candidate set.

— Neighbourhood boundary ● Sampled intersections (n=46) ★ Representative-point fallback



Figure 12: **Neighbourhood origin sampling for accessibility modelling.** Red points show all road-network intersections inside an example neighbourhood polygon. The black star shows the representative point fallback used only when a polygon contains no internal network node. The median travel time across all sampled intersections is used as the neighbourhood travel-time estimate.

Healthcare facility flood exposure and operational flood trigger

To determine when a healthcare facility is directly affected by flooding, we sample flood depth within a 5 m radius around each building footprint at every hourly step. The radius is used because building entrances and immediately adjacent access points are often more relevant to operational disruption than the interior of the footprint in a raster-based model. For each facility-hour, we record maximum, mean and 90th-percentile ring depth, along with the fraction of ring cells that are wet, defined as water depth greater than 0.01 m.

A facility crosses the operational flood trigger at the first hour when the maximum depth in the 5 m ring is at least 0.30 m and more than 20% of ring cells are wet. Once this threshold is crossed, the facility is treated as unavailable for the remainder of the simulated event in accessibility and evacuation routing. This is a conservative service-continuity assumption: the trigger is an operational disruption proxy, not proof that a facility would necessarily close or evacuate in practice. Supplementary Information reports the distinction between any ring wetting and trigger exceedance, as well as threshold sensitivity.

Neighbourhood health vulnerability index

Not all residents are equally able to cope with a loss of healthcare access. We summarise this susceptibility at the neighbourhood scale using a composite Health Vulnerability Index (HVI). We select seven indicators spanning four domains: demographic (share of residents aged 65 and over), physical health (share with chronic conditions; share with severe overweight), mental health (share reporting psychological complaints; share at risk of anxiety or depression) and social (share with low personal resilience / limited sense of control; share reporting severe or very severe loneliness).

We orient each indicator so that a higher value indicates greater vulnerability, normalise it to the unit interval, and combine them into a single 0-1 score. Our primary index is an equal-weighted composite of the indicators' percentile ranks across neighbourhoods, which is robust to outliers and produces an intuitive, comparable ranking. To test the influence of the weighting and normalisation choices, we constructed four alternative composites: equal-weighted min-max, equal-weighted winsorised min-max, an age-emphasised variant and a health-burden-emphasised variant and report them in the Supplementary Information.

Vulnerability-weighted access loss

We summarise the event's impact on the healthcare system with metrics that combine hazard, exposure and vulnerability. We accumulate person-hours of lost access: for each neighbourhood-hour, we multiply the extra travel time (in hours) by the resident population, then sum across neighbourhoods and time. From the resulting city-wide time series, we derive recovery metrics: the hours at which 50 %, 90 % and 99 % of the total disruption has accrued, and the first hour after the peak at which the hourly loss falls below 5 % of its peak and stays there (a back to normal marker).

To weight disruption by who is affected, we define the Access Resilience Loss (ARL) per neighbourhood and hour as

$$VWAL_{n,t} = P_n \times V_n \times \Delta\tau_{n,t}$$

where P_n is the neighbourhood population, V_n is the composite neighbourhood health vulnerability score ranging from 0 to 1, and $\Delta\tau_{n,t}$ is the additional travel time, in hours, required to reach the nearest open hospital relative to normal conditions. And, as a result, $VWAL$ units are vulnerability-weighted person-hours. $VWAL$ measures population-level deterioration in potential access to healthcare, rather than the number of realised hospital trips.

Potential first-hit evacuation demand

We estimate potential first-hit evacuation demand only for facilities that cross the operational flood trigger defined above: maximum depth of at least 0.30 m in the 5 m facility ring and more than 20 % wet ring coverage. The first-hit hour is the first hour at which both criteria are met. This metric is a planning proxy for the number of people who may need evacuation-related support when flood conditions first become operationally critical; it is not an observed evacuation, closure or transfer decision.

For each first-hit facility, we join the first-hit hour to the median (p50) occupancy estimate for the corresponding hour of the weekly occupancy profile. Hospitals contribute inpatients, staff and outpatients/visitors; elderly homes contribute residents, staff and visitors.

We assess sensitivity to event timing by recomputing the same first-hit occupancy counts after shifting the mapping between flood hours and occupancy-profile hours by four hours. The flood-trigger sequence, affected facility set, and occupancy model remain fixed; only the hour of day at which each first hit samples the occupancy profile changes. We use this sensitivity to compare total potential demand, patients or residents likely to require assistance, staff and visitors, and the greatest one-hour demand.

Neighbourhood access-loss profiles

We identify neighbourhood access-loss profiles by clustering the hourly hospital $VWAL$ trajectories. Neighbourhoods with no hospital $VWAL$ during the event are excluded from the clustering and mapped separately as having no modelled hospital-access loss. For affected neighbourhoods, the raw hourly $VWAL$ time series is first scaled by each neighbourhood's own maximum value. This means that clustering groups neighbourhoods by the timing and shape of access loss rather than by population size alone.

We then apply hierarchical agglomerative clustering with Ward linkage and Euclidean distance to the scaled trajectories, and cut them into four clusters. The number of clusters was chosen to provide a small, interpretable set of response profiles for the Results rather than by clustering performance metrics or to optimise future label predictors. For each profile, we report the number of neighbourhoods, onset hour, peak hour, recovery hour, duration, cumulative hospital $VWAL$, population and HVI summary statistics. Onset and recovery are defined from the group-average raw $VWAL$ trajectory as the first and last hours above 10% of that group's peak.

Facility exposure timing-duration profiles

We classify flood-affected facilities along two dimensions: the timing of first trigger crossing and the duration for which trigger-exceeding conditions persist. First-hit timing is divided into three event phases: early (hours 20-27), cascade (hours 28-31) and late (hours 32-39). These phase boundaries are descriptive and were chosen from the observed sequence of first trigger crossings: a small early set, a concentrated surge around the hospital impact, and a final later set.

Duration is calculated as the number of hours during which both facility-trigger criteria are met. Facilities are classified as brief (≤ 24 h), sustained (25-48 h) or prolonged (49-72 h). Combining first-hit timing and duration forms the timing-duration typology used in the Results. For each cell of the typology, we report the number of facilities, the number of hospitals and elderly homes, and the estimated number of people present at first hit.

Author contribution

M.S. conceived and designed the study, developed the analytical framework, prepared and integrated the input data, performed the analyses, produced the figures, interpreted the results, and wrote the original draft. N.S., K.N.T., C.S.d.V., T.M.B., and F.M. developed and provided the Rotterdam flood model. K.S. provided mobility data. Y.A.A. supported data collection and contributed to manuscript writing and revision. T.C. and M.P. contributed to conceptual development, interpretation of results, and manuscript writing and revision. All authors reviewed and edited the manuscript and approved the final version.

Funding declaration

This work was carried out within the Frontrunner 3 project, "Pandemic lessons for flood disaster preparedness", and was funded by the Pandemic and Disaster Preparedness Center. The funder had no role in the study design, data collection and analysis, interpretation of results, or preparation of the manuscript.

References

- Abebe, Y. A., Pregnolato, M., and Jonkman, S. N. (2025). Flood impacts on healthcare facilities and disaster preparedness—a systematic review. *International Journal of Disaster Risk Reduction*, 119:105340.
- Aggarwal, S., Hu, J. K., Sullivan, J. A., Parks, R. M., and Nethery, R. C. (2025). Severe flooding and cause-specific hospitalisation among older adults in the USA: a retrospective matched cohort analysis. *The Lancet Planetary Health*, 9(7):101268.
- Alam, M. S., Kim, K., Horner, M. W., Alisan, O., Antwi, R., and Ozguven, E. E. (2024). Large-scale modeling of hurricane flooding and disrupted infrastructure impacts on accessibility to critical facilities. *Journal of Transport Geography*, 116:103852.
- Antony, R., Rahiman, K. U. A., and Vishnudas, S. (2021). Flood Hazard Assessment and Flood Inundation Mapping—A Review. In Thomas, J., Jayalekshmi, B., and Nagarajan, P., editors, *Current Trends in Civil Engineering*, pages 209–218, Singapore. Springer.
- Arefi, F., Tavan, A., Moradi, S. M., Daneshi, S., and Farahmandnia, H. (2025). Identifying challenges and future directions of flood hazards mitigation strategies in health facilities: a systematic literature review. *BMC Emergency Medicine*, 25(1):174.
- Boeing, G. (2025). Modeling and Analyzing Urban Networks and Amenities With OSMnx. *Geographical Analysis*, 57(4):567–577. [eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/gean.70009](https://onlinelibrary.wiley.com/doi/pdf/10.1111/gean.70009).
- Cea, L. and Costabile, P. (2022). Flood Risk in Urban Areas: Modelling, Management and Adaptation to Climate Change. A Review. *Hydrology*, 9(3):50.
- Dong, S., Esmalian, A., Farahmand, H., and Mostafavi, A. (2020). An integrated physical-social analysis of disrupted access to critical facilities and community service-loss tolerance in urban flooding. *Computers, Environment and Urban Systems*, 80:101443.
- Dorfer, N., Bailie, J., Ahern, C., McNaught, R., Scott, K., Matthews, V., Morgan, G. G., Ekanayake, K., and Bailie, R. (2025). How floods impact health systems: a scoping review of Australian research. *BMJ Open*, 15(12):e108856.

- Ermagun, A., Janatabadi, F., Safarloo, Z., Kremer, P., and Lindley, S. (2026). Urban flooding restructures mobility through coupled behavioral and network disruption: A systematic review of evidence. *Sustainable Cities and Society*, 141:107285.
- Gangwal, U. and Dong, S. (2022). Critical facility accessibility rapid failure early-warning detection and redundancy mapping in urban flooding. *Reliability Engineering & System Safety*, 224:108555.
- Hammond, M., Chen, A., Djordjević, S., Butler, D., and Mark, O. (2015). Urban flood impact assessment: A state-of-the-art review. *Urban Water Journal*, 12(1):14–29. eprint: <https://doi.org/10.1080/1573062X.2013.857421>.
- Hertelendy, A. J., Dresser, C., Gorgens, S., Hertelendy, A. P., Biddinger, P. D., and Ciottone, G. (2025). Strengthening healthcare system resilience: a comprehensive framework for tropical cyclone preparedness and response. *The Lancet Regional Health – Americas*, 48.
- Hoover, S. and Tate, E. (2025). Spatial heterogeneity in social vulnerability to flood exposure. *Natural Hazards*, 121(11):12695–12719.
- Jong, A. d., Beek, J. v., Korving, H., Koopmans, M., and Boelee, E. (2023). Health effects of the 2021 flooding in Limburg. *Journal of Coastal and Riverine Flood Risk*, 2.
- Kim, K. Y., Toplu-Tutay, G., Kutanoğlu, E., and Hasenbein, J. J. (2024). A stochastic optimization model for patient evacuation from health care facilities during hurricanes. *International Journal of Disaster Risk Reduction*, 108:104518.
- Kok, M., Slager, K., Moel, H. d., Botzen, W., Bruijn, K. d., Wagenaar, D., Rikkert, S., Koks, E., and Ginkel, K. v. (2023). Rapid damage assessment caused by the flooding event 2021 in Limburg, Netherlands. *Journal of Coastal and Riverine Flood Risk*, 2.
- Landeg, O., Whitman, G., Walker-Springett, K., Butler, C., Bone, A., and Kovats, S. (2019). Coastal flooding and frontline health care services: challenges for flood risk resilience in the English health care system. *Journal of Health Services Research & Policy*, 24(4):219–228.
- Ligterink, N. (2016). On-road determination of average Dutch driving behaviour for vehicle emissions. Technical report.
- Lin, B., Cai, H., Tian, H., Fu, D., Das, P. C., Sumiya, N. N., and Yin, Z. (2026). Quantifying healthcare visitation disruptions using mobile phone location data during disasters: A case study of hurricane Beryl. *Progress in Disaster Science*, 31:100623.
- Long, L., Junjia, Y., Jiao, W., Alias, A. H., Haron, N. A., and Abu Bakar, N. (2025). Urban flood resilience evaluation in China: a systematic review of frameworks, methods, and limitations. *Geomatics, Natural Hazards and Risk*, 16(1):2445631. eprint: <https://doi.org/10.1080/19475705.2024.2445631>.
- Mas-Coma, S., Artigas, P., Cuervo, P. F., De Elías-Escribano, A., Fantozzi, M. C., Colangeli, G., Córdoba, A., Marquez-Guzman, D. J., Mas-Bargues, C., Borrás, C., Pérez-Pérez, P., Bethencourt-Estrella, C. J., Rodríguez-Expósito, R. L., Peña-Prunell, M. D., Chao-Pellicer, J., García-Pérez, O., Domínguez De Barros, A. T., García-Ramos, A., Sirvent-Blanco, C., Gajeta-Arenas, M., Córdoba-Lanús, A. E., Piñero, J. E., Valero, M. A., Lorenzo-Morales, J., and Bargues, M. D. (2025). Infectious disease risk after the October 2024 flash flood in Valencia, Spain: Disaster evolution, strategic scenario analysis, and extrapolative baseline for a One Health assessment. *One Health*, 21:101093.
- McClymont, K., Morrison, D., Beevers, L., and Carmen, E. (2020). Flood resilience: a systematic review. *Journal of Environmental Planning and Management*, 63(7):1151–1176. eprint: <https://doi.org/10.1080/09640568.2019.1641474>.
- Melnychuk, E., Sallade, T. D., and Kraus, C. K. (2022). Hospitals as disaster victims: Lessons not learned? *JACEP Open*, 3(1):e12632.
- Moreira, L. L., de Brito, M. M., and Kobiyama, M. (2021). Review article: A systematic review and future prospects of flood vulnerability indices. *Natural Hazards and Earth System Sciences*, 21(5):1513–1530.
- Papilloud, T., Steiner, A., Zischg, A., and Keiler, M. (2024). Road network disruptions during extreme flooding events and their impact on the access to emergency medical services: A spatiotemporal vulnerability analysis. *Science of The Total Environment*, 956:177140.

- Petricola, S., Reinmuth, M., Lautenbach, S., Hatfield, C., and Zipf, A. (2022). Assessing road criticality and loss of healthcare accessibility during floods: the case of Cyclone Idai, Mozambique 2019. *International Journal of Health Geographics*, 21(1):14.
- Pregolato, M., Ford, A., Wilkinson, S. M., and Dawson, R. J. (2017). The impact of flooding on road transport: A depth-disruption function. *Transportation Research Part D: Transport and Environment*, 55:67–81.
- Rebally, A., Valeo, C., He, J., and Saidi, S. (2021). Flood Impact Assessments on Transportation Networks: A Review of Methods and Associated Temporal and Spatial Scales. *Frontiers in Sustainable Cities*, 3.
- Rözer, V., Mehryar, S., and Surminski, S. (2022). From managing risk to increasing resilience: a review on the development of urban flood resilience, its assessment and the implications for decision making. *Environmental Research Letters*, 17(12):123006.
- Sancho, E., Martín, R., Abejaro, B., Zaragoza, R., Carrasco, C., and Ramírez, P. (2026). Critical care during natural disasters: ICU insights from Valencia’s October 2024 floods. *Medicina Intensiva (English Edition)*, 50(6):502395.
- Stimpson, J. P., Rashed, A. L., Pandya, J., Baudot, E. C., Whitfill, J., and Ortega, A. N. (2025). Health equity in the wake of disasters and extreme weather: evidence from an umbrella review. *Health Affairs Scholar*, 3(11):qxaf207.
- Tariverdi, M., Nunez-del Prado, M., Leonova, N., and Rentschler, J. (2023). Measuring accessibility to public services and infrastructure criticality for disasters risk management. *Scientific Reports*, 13(1):1569. Publisher: Nature Publishing Group.
- Thrysøe, C., Balstrøm, T., Borup, M., Löwe, R., Jamali, B., and Arnbjerg-Nielsen, K. (2021). FloodStroem: A fast dynamic GIS-based urban flood and damage model. *Journal of Hydrology*, 600:126521.
- Valente, M., Del Prete, C., Facci, G., Musso, F., Cenati, S., Calligaro, S., Ragazzoni, L., and Barone-Adesi, F. (2026). The impacts of extreme weather events on health services and systems: A systematic review of reviews. *Public Health*, 250:106049.
- Wal, W. v., Barten, D. G., Ketelings, L., van Osch, F. H. M., Rao, M., Mortelmans, L., and Bierens, J. (2025). Preemptive hospital evacuation during the 2021 European floods: A qualitative study on factors influencing decision-making. *The Journal of Climate Change and Health*, 24:100460.
- Wassmer, J., Bryant, S., Schimansky, P., Keegan, L. T., Pregolato, M., Kurths, J., Marwan, N., and Merz, B. (2025). Unveiling hidden risks in healthcare from flood-induced transportation disruption in Germany. *Communications Earth & Environment*, 6(1):676.
- WHO (2021). *Health service continuity planning for public health emergencies*. World Health Organization.
- Yin, J., Yu, D., and Liao, B. (2021). A city-scale assessment of emergency response accessibility to vulnerable populations and facilities under normal and pluvial flood conditions for Shanghai, China. *Environment and Planning B: Urban Analytics and City Science*, 48(8):2239–2253.
- Zhou, Z., Zhang, Y., Gou, J., Kan, N., Zhao, F., Yao, L., and Hou, S. (2023). Successful Large Hospital Evacuation With 11 350 Patients Transferred in the 2021 Zhengzhou Flood. *Disaster Medicine and Public Health Preparedness*, 17:e434.