data

MDPI

Article

An Open-Data Framework for Screening Flood-Footprint Population Proxies and Potential Hospital Accessibility Disruption

Hossein Hassani [1,*], Leila Marvian Mashhad [2] and Nadejda Komendantova [1]

[1] International Institute for Applied Systems Analysis (IIASA), 2361 Laxenburg, Austria
[2] Big Data Lab, Imam Reza International University, Mashhad 91388-33186, Iran
* Correspondence: hassani.stat@gmail.com

Abstract

Urban disaster screening requires transparent methods that can integrate heterogeneous open datasets without implying unsupported causal or probabilistic relationships. This study presents an open-data framework for screening flood-footprint population proxies and potential disruption to hospital accessibility in Bucharest, Romania. The analysis covers 104 archived hexagonal spatial units and combines a 100-year flood-depth scenario from the Joint Research Centre, WorldPop 2020 population estimates, hospital-routing outputs derived from OpenStreetMap, and a publicly available seismic screening surface from the European Facilities for Earthquake Hazard and Risk. Flood and seismic information are retained as distinct screening dimensions because the available data do not support modelling their causal interaction, temporal sequence, joint probability, earthquake-related infrastructure damage, or hospital capacity. For spatial units that remain connected to a hospital, a transparent two-domain service-priority index is calculated using flood-footprint population proxy and the potential change in hospital accessibility under the flood scenario. Units for which no hospital route is available are reported separately as binary service-disconnection alerts rather than being assigned an arbitrary numerical penalty. The robustness and interpretability of the framework are examined through network monotonicity, score boundedness, Pareto dominance, penalty-free rank invariance, and rank-acceptability analysis. The results are communicated using a Pareto frontier, rank trajectories across the full weight simplex, indicator-contribution decomposition, and an exact hypergeometric assessment of class overlap. The proposed framework provides a transparent first-order planning tool for identifying locations that may require more detailed investigation. It should be interpreted as a screening approach rather than as a probabilistic risk, cascading-hazard, infrastructure-damage, or service-loss model.

Keywords: multi-hazard screening; potential accessibility disruption; flood footprint; hospital access; Pareto frontier; rank acceptability; Bucharest; open data

Check for updates

Academic Editor: Hu Wang

Received: 11 August 2026
Revised: 10 September 2026
Accepted: 19 September 2026
Published: 20 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Urban multi-hazard research increasingly distinguishes a spatial overlay from an actual compound, consecutive, or cascading process. A cascading model requires an explicit directional mechanism or temporal sequence through which one event changes the occurrence or consequence of another. In contrast, a composite map may be useful for planning while still remaining a screening product. This distinction is central to recent work

on multi-hazard indicators, spatial–temporal impact chains, and urban vulnerability [1–4]. It is also important for infrastructure applications, where a route-access model should not be re-labelled as observed service loss without facility capacity, traffic, dispatch, patient-behaviour, or observed-disruption evidence [5,6].

The Bucharest framework combines seismic screening, an RP100 flood scenario, population, facilities, and hospital accessibility while keeping their evidential roles distinct. It does not model earthquake-triggered flooding, common return periods across hazards, physical earthquake damage to roads or hospitals, hospital capacity, or a validated road-failure process. The framework is therefore defined as *multi-hazard screening*, not as cascading or compound risk. The seismic screening layer is kept separate from the flood-access priority score because the probability bases are not harmonised, and the continuous flood-service priority index is constructed from two auditable domains without an arbitrary disconnection penalty.

The research question is operational rather than causal: *How does adding potential hospital-accessibility deterioration to flood-footprint population proxy alter the planning-priority pattern obtained from simpler flood screens?* The contribution is an open-data, first-order decision-support workflow. It does not estimate realised loss, engineering damage, or joint earthquake–flood probability. This careful framing is consistent with recent emphasis on transparent risk terminology, uncertainty analysis, and interdependency-aware infrastructure research [7–11].

2. Materials and Methods

2.1. Study Area, Analytical Units, and Data Boundary

Bucharest is represented by 104 clipped hexagons generated from the municipal boundary. The archived analytical output combines WorldPop 2020 population totals, a Joint Research Centre Europe RP100 flood-depth result, a public EFEHR PGA screening field, OpenStreetMap-derived hospital-route outputs, and descriptive critical-facility counts. Coordinates are WGS 84 (EPSG:4326). The 1.5 km nominal grid defines the spatial resolution of the archived analytical dataset; it does not remove the modifiable areal unit problem. The analytical geometry should therefore be interpreted as a screening unit rather than a precise representation of resident locations. The data architecture and analytical role of each source are summarised in Table 1.

Table 1. Data architecture and analytical role of each source.

Input Family	Archived Source	Analytical Role
Population	WorldPop Romania, 2020	Population term in the flood-footprint population proxy.
Flood scenario	JRC Europe RP100 flood depth	Flood-footprint share and route-disruption scenario; not a joint probability.
Seismic context	Public EFEHR PGA screening extract	Separate contextual overlay; excluded from continuous service-priority index.
Hospital routes	OSM-derived baseline and disrupted shortest paths	Potential spatial access change for connected units; alerts for lost reachability.
Facilities	OSM hospitals, schools, and fire stations	Descriptive context only; not aggregated into a mixed exposure count.

The archived aggregate output does not retain the edge-level road-closure threshold, road-class speed table, bridge treatment, one-way implementation, node-snapping

parameters, OSM extraction date, hospital-capacity assumption, or EFEHR interpolation configuration used to produce the routing outputs. The archive also does not preserve the probability/return-period basis, classification threshold, or interpolation details used to define the high-seismic-screening class. These limitations are explicitly documented in the Zenodo repository metadata and provenance file. The route and seismic outputs are therefore treated as documented scenario inputs, and the analysis avoids claims of validated physical road damage, realised hospital service loss, or a harmonised seismic–flood probability.

2.2. Reproducibility Record and Processing Sequence

The analysis package reproduces the reported downstream screening calculations from the archived 104-unit input-derived result file. It contains the analytical data table, generated tables and non-map analytical figures, and a deterministic script for recalculating the score and sensitivity diagnostics. Because the two-domain weight space is one-dimensional, rank acceptability is evaluated exactly over the complete interval $w \in [0, 1]$ rather than approximated by Monte Carlo simulation. Reproducibility therefore applies to the documented post-processing workflow; unavailable upstream GIS and road-disruption settings are neither reconstructed nor inferred. The reproducibility boundary and repository contents are summarised in Table 2.

Table 2. Reproducibility boundary and repository contents.

Workflow Element	Status	Repository/Interpretation
Archived 104-unit analytical table	Available	Unit-level population, flood, route, seismic-screening and derived fields are included as CSV.
Flood-footprint proxy and SPI	Fully reproducible downstream	Recalculated from archived fields by `reproduce_analysis.py`.
Weight sensitivity/ rank acceptability	Fully reproducible downstream	Evaluated exactly over the complete $w \in [0, 1]$ interval using pairwise score-crossing partitions.
Tables and non-map analytical figures	Fully reproducible downstream	Generated from the archived analytical CSV and script.
Road-network disruption derivation	Not reconstructable	Closure threshold, speed table, bridge/ one-way treatment, snapping and OSM extraction date are unavailable.
Flooded-cell indicator derivation	Not reconstructable	The depth threshold used to define flooded raster cells is unavailable.
Seismic high-class derivation	Not reconstructable	Classification threshold and interpolation/probability details are unavailable; class membership is treated descriptively as archived input.
Hexagon GIS geometry	Not available	Unit identifiers support tabular downstream analysis but not independent reconstruction of the map geometry.

The sequential workflow is as follows. First, population, flood, accessibility, and seismic-screening fields are assigned to the hexagonal unit. Second, the flood-footprint population proxy $E_h = P_h F_h$ is calculated. Third, finite scenario routes define the connected scoring universe, whereas units without a finite hospital route are recorded as alerts. Fourth, connected units are normalised and scored. Fifth, all pairwise score-crossing

weights are identified and used to partition $w \in [0,1]$ into intervals on which the complete ranking is constant; rank-acceptability quantities are then integrated exactly over those intervals. Sixth, the index is compared with flood-only, flood-population-only, access-only, and percentile-normalised screens. Finally, the archived seismic class is used only for a descriptive co-location benchmark. This design follows recent calls for transparent separation of hazard, exposure, accessibility, and uncertainty dimensions in multi-hazard screening [2,4,12–17].

2.3. Baseline and Flood-Scenario Potential Accessibility

Let $G^0 = (V, E^0)$ be the directed baseline graph and $G^1 = (V, E^1)$ the flood-scenario graph. Let M be the set of mapped hospital destinations and o_h the road-network origin for hexagon h. The nearest-hospital travel time in scenario s is

$$T_h^{(s)} = \min_{m \in M} d_{G^s}(o_h, m), \tag{1}$$

where d_{G^s} is the shortest-path travel time. For connected units, potential accessibility disruption is

$$A_h = \frac{T_h^{(1)} - T_h^{(0)}}{T_h^{(0)}}. \tag{2}$$

If $T_h^{(1)}$ is not finite, the unit is labelled *disconnected after flood scenario*. It is a binary service alert and never receives a finite penalty. The metric captures potential spatial accessibility only: it does not represent ambulance dispatch, congestion, hospital capacity, patient choice, facility failure, or realised healthcare service loss. The scenario rules and interpretation boundaries are summarised in Table 3.

Table 3. Explicit scenario rules and interpretation boundaries.

Element	Retained Interpretation
Baseline condition	Shortest travel time to the nearest mapped hospital in the archived baseline graph.
Flood condition	Archived shortest-path result under the stated RP100 flood-routing scenario.
Hospital role	Hospital points are destinations; capacity, staffing, facility failure, and patient allocation are outside the model.
Road and bridge rules	Edge-level closure and bridge settings are unavailable in the archive and are not inferred.
Disconnection	A non-finite scenario route; it is a binary alert and is excluded from score normalisation.
Seismic layer	A contextual PGA screening field; it is not transformed into infrastructure damage or a joint probability.

Proposition 1 (Network monotonicity). *Assume that every retained edge $e \in E^1$ has positive flood-scenario cost $c^1(e) \geq c^0(e)$ and that $E^1 \subseteq E^0$. Then, $T_h^{(1)} \geq T_h^{(0)}$ for every hexagon connected in both scenarios.*

Proof. Every path feasible in G^1 is feasible in G^0, and its scenario cost is no lower than its baseline cost. Minimising over the smaller scenario path set therefore cannot yield a travel time below the minimum over the larger baseline path set. If no path remains, the travel time is infinite, and the unit is recorded as a separate alert. □

2.4. Flood-Footprint Population and Parsimonious Service-Priority Index

Let P_h denote the WorldPop population total and F_h the archived flooded-raster-cell share of hexagon h. The operational depth threshold used to create this flooded-cell indicator is not retained in the available provenance record. The flood-footprint population proxy is

$$E_h = P_h F_h. \tag{3}$$

This is a uniform-within-hexagon population proxy, not a pixel-exact count of residents in water. A direct WorldPop-raster overlay with the flood footprint cannot be reconstructed from the archived materials because the analytical polygon geometry and the full pixel-level processing state are unavailable. The uniform-within-hexagon assumption is therefore treated explicitly as a source of uncertainty. The proxy is used instead of entering both population and flood extent as separate positive components. Mean flood depth remains descriptive and is not independently weighted, avoiding partial double counting of the hydraulic dimension.

For each component x_{kh}, min–max normalisation is

$$Z(x_{kh}) = \frac{x_{kh} - \min_h(x_k)}{\max_h(x_k) - \min_h(x_k)}. \tag{4}$$

The general two-domain score is

$$SPI_h(w) = wZ(E_h) + (1-w)Z(A_h), \qquad 0 \le w \le 1. \tag{5}$$

The reported baseline uses $w = 0.50$. Facility counts, baseline travel time, mean flood depth, seismic PGA, and a disconnection penalty receive zero continuous-score weight. As a normalisation sensitivity check, each retained component is also transformed to its empirical percentile rank within the connected-unit universe, and the two percentile ranks are averaged with equal weights. Operational shortlist sensitivity is evaluated at top-10%, top-20%, top-25%, and top-33% thresholds; these are planning cut-offs rather than physical thresholds.

Proposition 2 (Boundedness and monotonicity)**.** *For every connected hexagon,* $0 \le SPI_h(w) \le 1$. *Holding the other component fixed, an increase in either normalised component cannot reduce* $SPI_h(w)$ *when its corresponding weight is positive.*

Proof. $SPI_h(w)$ is a convex combination of two values in $[0,1]$; hence, it lies in $[0,1]$. Moreover, $\partial SPI/\partial Z(E) = w \ge 0$, and $\partial SPI/\partial Z(A) = 1 - w \ge 0$. □

Proposition 3 (Pareto dominance)**.** *If* $Z(E_i) \ge Z(E_j)$ *and* $Z(A_i) \ge Z(A_j)$, *with at least one strict inequality, then* $SPI_i(w) > SPI_j(w)$ *for every* $0 < w < 1$.

Proof. Subtracting the two scores gives

$$SPI_i(w) - SPI_j(w) = w[Z(E_i) - Z(E_j)] + (1-w)[Z(A_i) - Z(A_j)]. \tag{6}$$

Both terms are non-negative and at least one is positive. □

Proposition 4 (Penalty-free rank invariance)**.** *When disconnected units are excluded from normalisation and continuous scoring, the ranks of connected units are invariant to any numerical coding of the disconnected state.*

Proof. For connected i and j, their order depends only on the sign of $SPI_i - SPI_j$. A binary alert outside the score cannot alter this sign. In contrast, inserting a penalty M into min–max normalisation changes a connected-unit difference to $(x_i - x_j)/(M - \min x)$, which converges to zero as M grows; a penalty can therefore compress connected-unit differentiation. □

2.5. Exact Full-Weight-Space Uncertainty and Descriptive Co-Location Benchmark

For operational threshold q, rank acceptability is defined as the fraction of the admissible weight interval for which unit h belongs to the leading q ranks,

$$\pi_h(q) = \int_0^1 \mathbb{I}\{\text{rank}[SPI_h(w)] \le q\}\, dw. \tag{7}$$

Because $SPI_h(w) = Z(A_h) + w[Z(E_h) - Z(A_h)]$ is linear in w, the ordering of any pair of connected units can change only at a pairwise score-crossing weight. Let $0 = b_0 < b_1 < \cdots < b_J = 1$ denote the sorted set of all such crossings plus the interval endpoints. The complete ranking is constant on every open interval (b_{j-1}, b_j), so

$$\pi_h(q) = \sum_{j=1}^{J} (b_j - b_{j-1})\, \mathbb{I}\{\text{rank}[SPI_h(w_j^*)] \le q\}, \tag{8}$$

where w_j^* is any point inside interval j. The same interval lengths provide exact weight-space means and quantiles of rank and rank-correlation summaries. In the archived dataset, the pairwise crossings generate 1514 constant-ranking intervals. Uniform measure over $w \in [0,1]$ is used only as a neutral sensitivity convention and is not interpreted as a probability model of stakeholder preferences.

The high-seismic-screening and high-service-priority sets are analysed separately. The seismic class is an archived descriptive input because its original classification threshold and interpolation/probability settings cannot be reconstructed. Conditional on the archived class sizes, their overlap X is compared with a hypergeometric reference distribution,

$$\Pr(X = x) = \frac{\binom{K_s}{x}\binom{n-K_s}{K_p-x}}{\binom{n}{K_p}}, \tag{9}$$

where n is the connected-unit universe, K_s is the archived seismic-class size, and K_p is the service-priority class size. The expected conditional overlap is $E[X] = K_s K_p / n$. This is a descriptive co-location benchmark only. It is not a test of earthquake–flood dependence, a joint return period, or causal cascading, and it does not adjust for spatial autocorrelation among neighbouring units.

3. Results

3.1. Priority Geography and Service Alerts

The continuous score is calculated for 100 connected hexagons. Four units, H006, H012, H059, and H073, have no finite hospital route in the flood scenario and are separately reported as service alerts. No artificial numerical score is assigned to these units because no disconnection penalty enters the index. The flood-footprint population proxy totals 933,777 across all units. The corresponding WorldPop population total of the 25 connected operational-priority units is 744,806; this quantity is not an estimate of residents affected by flooding. In the archived result, selected top-priority units have a 100% flood-footprint share; this means the proxy equals their hexagon population total and must not be read as an observed population in floodwater. The spatial pattern is shown in Figure 1.

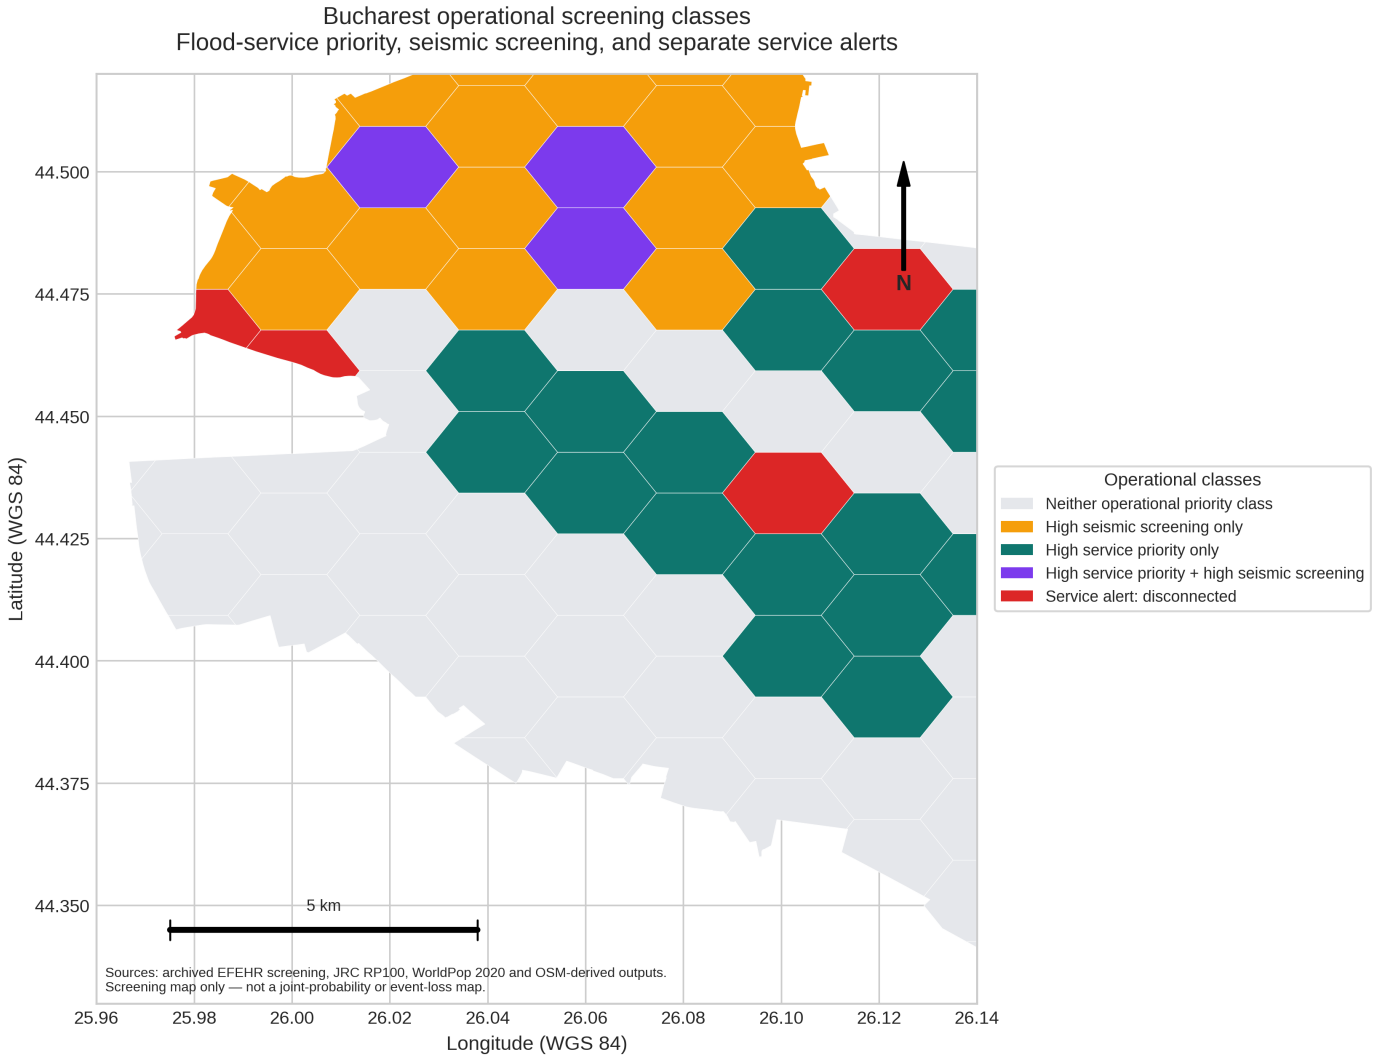

Figure 1. Two-axis operational screening map. Disconnected units are displayed as separate service alerts, while high seismic screening remains a distinct overlay rather than a component of the service-priority score.

H058 in Sector 4 is the highest-ranked connected unit, followed by H080 in Sector 3 and H069 in Sector 4. Sector 2 has the highest mean connected-unit priority score, 0.304, followed by Sector 1 at 0.236 and Sector 3 at 0.229. Sector means are calculated for connected units only, while service alerts are reported separately.

The four service alerts are not omitted from planning interpretation. They represent locations where the archived flood-routing output contains no finite route to a mapped hospital. Their separation from the continuous score prevents a small, low-population disconnected unit from automatically ranking above a highly populated unit with severe but finite accessibility deterioration. Planning therefore uses two outputs jointly: the ordered connected-unit priority score and the unordered list of binary reachability alerts. The leading connected units are listed in Table 4.

3.2. Comparison with Simpler Screening Approaches

The accessibility dimension changes the planning result rather than merely reproducing flood exposure. The service-priority rank has Spearman correlation 0.737 with flood footprint alone and a top-quartile Jaccard overlap of 0.471. It has correlation 0.813 with flood-footprint population proxy and Jaccard overlap 0.786. Access-change-only priorities are substantially

different, with correlation 0.598 and Jaccard overlap 0.163. These comparisons show the added value of the access signal but do not validate the index as observed service loss.

Table 4. Leading connected operational-priority hexagons.

Hexagon	Sector	Flood-Pop. Proxy	Access Change (%)	SPI	$\pi_h(25)$
H058	Sector 4	39,654.5	82.6	0.904	1.000
H080	Sector 3	45,282.6	49.7	0.753	1.000
H069	Sector 4	34,586.4	55.6	0.681	1.000
H048	Sector 1	22,095.9	64.4	0.609	1.000
H070	Sector 3	49,087.7	13.0	0.559	0.796
H082	Sector 2	43,373.9	19.9	0.544	0.846
H072	Sector 2	36,928.8	22.0	0.491	0.849
H095	Sector 3	42,228.4	13.4	0.491	0.772
H041	Sector 1	6524.7	69.9	0.486	0.758
H038	Sector 6	33,975.0	24.2	0.475	0.873

The comparison has a practical implication. Flood footprint alone selects locations largely according to the mapped hydraulic pattern. Flood-footprint population adds the distribution of residents, while the access-only result prioritises locations where routing is particularly affected. The two-domain score deliberately requires evidence from both retained domains. It therefore does not present the hospital-access component as an after-the-fact visual layer; it can change membership in the operational priority class. The method comparison is shown in Figure 2.

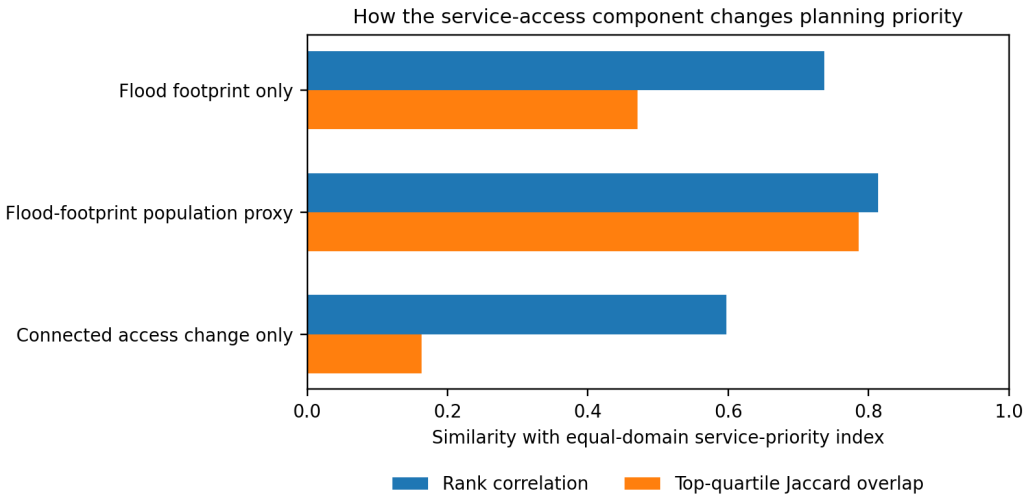

Figure 2. Similarity of the service-priority ranking with flood-only, flood-footprint-population, and connected-access-change screening approaches.

3.3. Rank-Acceptability Overview

Figure 3 provides a deliberately simplified two-panel sensitivity display. Panel A reports the exact fraction of the admissible weight interval for which each of the 15 leading connected units belongs to the top quartile. Panel B shows exact fifth–ninety-fifth percentile rank intervals and compares the equal-domain rank with the weight-space mean rank. The display uses ordinary axis labels rather than text drawn over data points, which avoids label overlap while retaining the same uncertainty information.

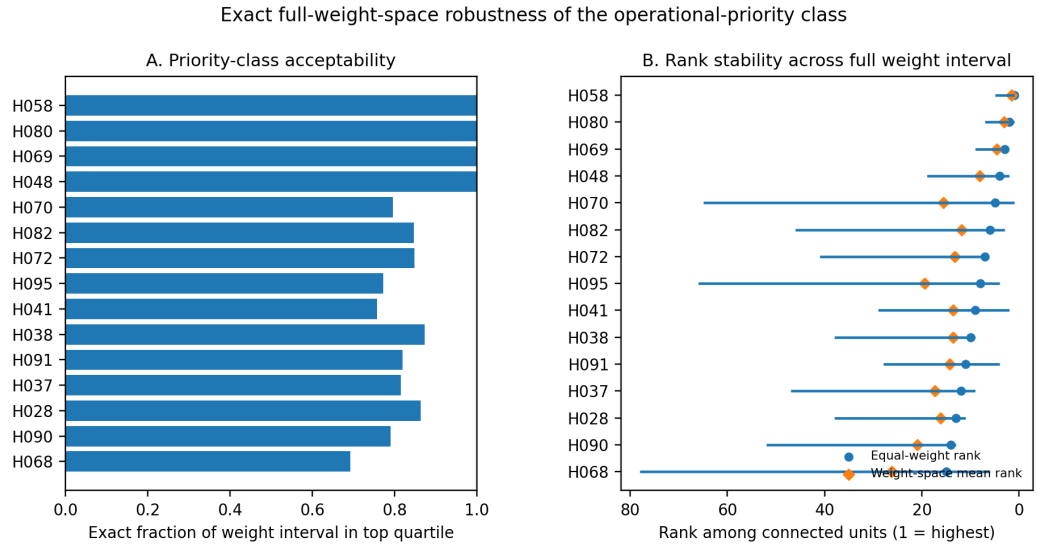

Figure 3. Exact full-weight-space robustness of the operational-priority class. Panel (**A**) reports the fraction of $w \in [0, 1]$ yielding top-quartile membership; Panel (**B**) reports fifth–ninety-fifth percentile rank intervals, equal-domain ranks, and weight-space mean ranks.

3.4. Advanced Diagnostics Embedded in the Results

Figure 4 explains the two-domain trade-off directly. Points on the Pareto frontier cannot be dominated simultaneously on flood-footprint population proxy and potential accessibility deterioration. This is a stronger statement than a single weighted ranking because Proposition 3 holds for every strictly positive weighting pair.

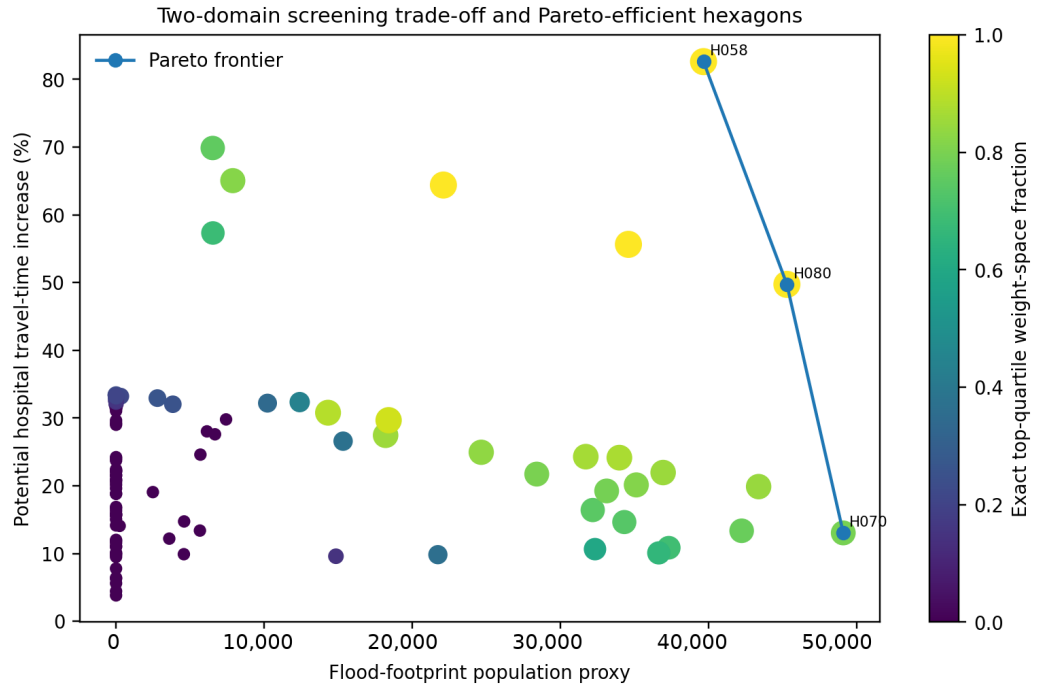

Figure 4. Pareto frontier for the two retained screening dimensions. Point size denotes the exact fraction of the admissible weight interval yielding top-quartile membership.

Figure 5 evaluates rank stability over the complete weight simplex, rather than relying on three hand-picked scenarios. The dashed line identifies the equal-domain baseline. Flat curves indicate stable priority; intersections reveal genuine value trade-offs between population in the flood footprint and access deterioration.

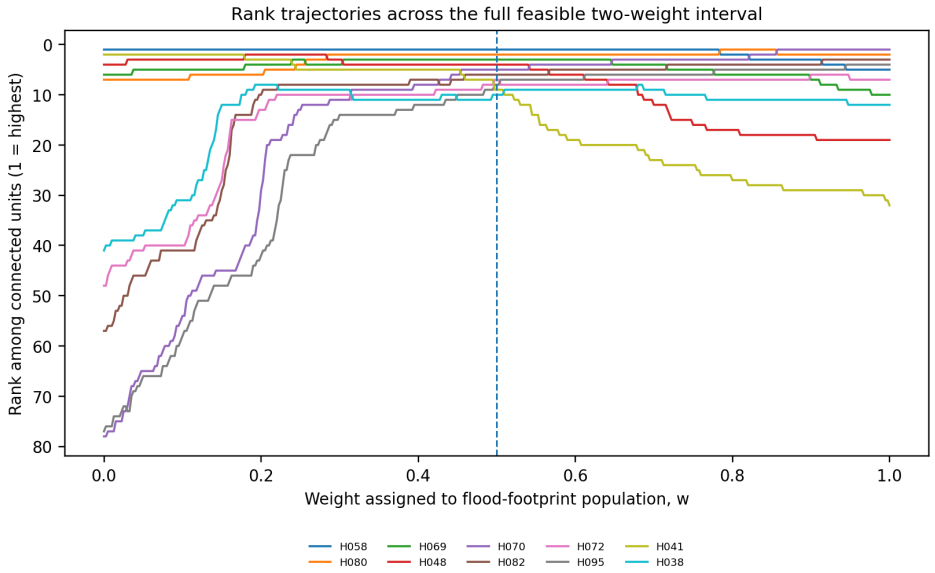

Figure 5. Rank trajectories of leading units over the full feasible weight simplex. The dashed line denotes the reported equal-domain baseline.

Figure 6 replaces circular correlations between a composite score and its own components. It displays the exact additive contribution of flood-footprint population proxy and accessibility change for each operational-priority unit. The graph therefore explains *why* each unit is prioritised under the stated baseline.

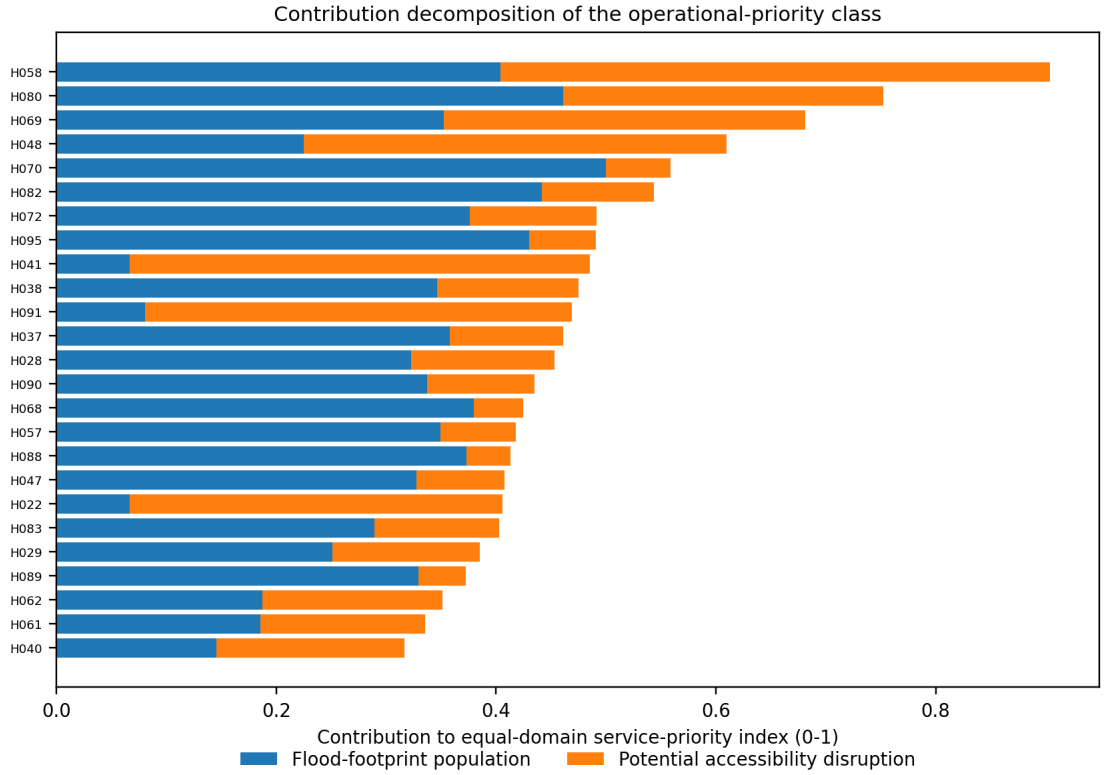

Figure 6. Additive contribution decomposition for the operational-priority class.

The equal-weight ranking is relatively stable across the complete weight interval. The exact weight-space median Spearman correlation with the equal-weight ranking is 0.968, while the fifth percentile is 0.683. Replacing min–max with percentile normalisation yields Spearman correlation 0.942 and Kendall correlation 0.846 with the equal-domain

baseline; the fixed-size top-quartile Jaccard overlap is 0.471 (16 shared units). Reporting these changes makes the value dependence of the shortlist explicit.

The top-quartile class is an operational planning threshold, not a physical discontinuity. Threshold sensitivity was also evaluated at top-10%, top-20%, and top-33%. The purpose is to show how the size of the planning shortlist affects membership, rather than to claim that exactly 25 units are objectively "hotspots". The project data directory contains the full threshold table and rank trajectories for the leading units. The leave-one-domain-out and normalisation sensitivity results are summarised in Table 5.

Table 5. Leave-one-domain-out and normalisation sensitivity results.

Scenario	Spearman ρ	Kendall τ	Top-Quartile Jaccard	Shared Units
Flood-footprint population only	0.813	0.670	0.786	22
Connected access change only	0.598	0.515	0.163	7
Equal-domain baseline	1.000	1.000	1.000	25
Percentile normalisation	0.942	0.846	0.471	16

Among connected units, 24 belong to the archived high-seismic-screening class, and 25 belong to the high service-priority class. Their observed overlap is three, compared with a conditional random-overlap expectation of six. Figure 7 displays the hypergeometric reference distribution. The lower-tail probability is 0.084; because the seismic classification rule is not reconstructable and spatial autocorrelation is not modelled, this value is reported only as a descriptive benchmark. The corresponding conditional overlap values are reported in Table 6.

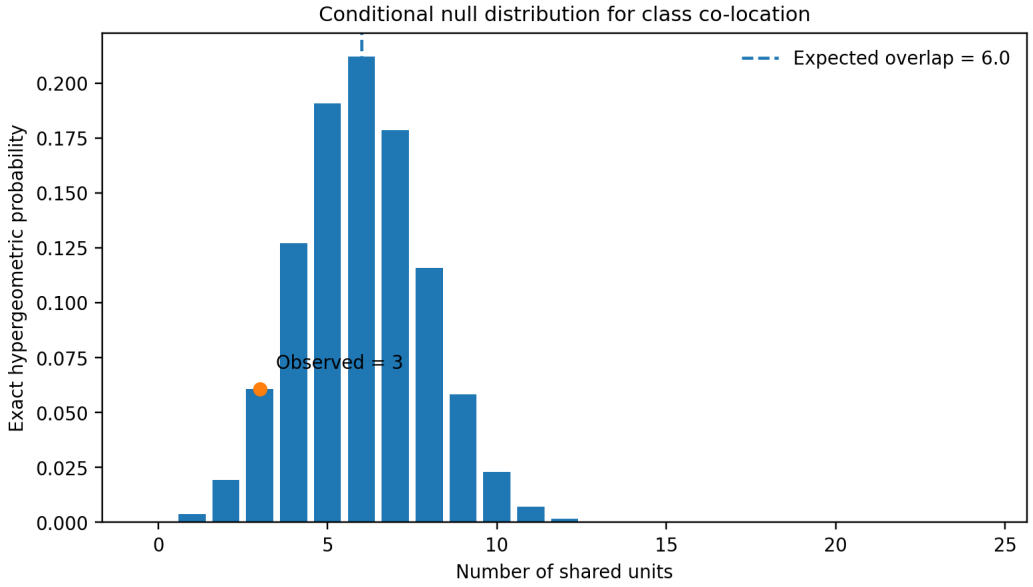

Figure 7. Conditional hypergeometric reference distribution for co-location of the archived high-seismic-screening class and the high service-priority class.

Table 6. Descriptive conditional class-overlap benchmark.

n	K_s	K_p	Observed	Expected	$P(X \leq x)$	$P(X \geq x)$
100	24	25	3	6.00	0.084	0.977

4. Internal Verification, Uncertainty, and Interpretation

The analytical framework distinguishes internal verification from external validation. Internal checks confirm that all 104 units have a baseline hospital travel-time result,

100 retain a finite scenario route, and disconnected units are excluded from continuous normalisation and ranking. These checks do not validate flood depth, edge closure, real ambulance travel time, hospital capacity, seismic interpolation, or observed service interruption. They are consistency checks of the archived analytical workflow rather than external validation of the underlying accessibility model. Historical closure records, emergency travel-time observations, hospital service data, and independent loss datasets were not available in the archived workflow. The internal checks and uncertainty boundaries are summarised in Table 7.

Table 7. Internal checks and explicit uncertainty boundaries.

Check/Boundary	Result	Interpretive Scope
Network-result completeness	104/104 baseline results; 100 finite flood-scenario routes; 4 disconnection alerts.	Internal consistency only; not validation of closures or travel times.
Route-change direction	All 100 connected units have positive scenario-minus-baseline travel-time change.	Checks archived scenario output, not real-world closure accuracy.
Disconnection handling	No numeric penalty enters normalisation or continuous scoring.	Avoids penalty-driven compression of connected-unit ranks.
Probability interpretation	RP100 flood and seismic screening are not harmonised probabilistically.	No combined return period, dependence, or cascading-event claim.
External validation	No observed closures, emergency travel times, hospital capacity, or event-loss records are archived.	Output remains a first-order planning screen.
Temporal alignment	WorldPop is 2020; other archived source vintages differ.	Temporal mismatch is a documented uncertainty.

Accordingly, the outputs should not be used for regulatory design, emergency dispatch, engineering loss modelling, or claims of population proxy exposure at the pixel level. A future operational study should archive raw OSM data and query date, specify edge-level flood and bridge rules, incorporate current hospital capacity and availability, harmonise probability interpretation where appropriate, validate against observed events, and evaluate alternative grid sizes and origins.

5. Discussion

A central feature of the framework is conceptual discipline. It does not transform a static earthquake layer plus an RP100 flood layer into a fictional joint probability. It separates a seismic context layer from a flood-service screen, assigns no mixed facility count, and avoids allowing a user-selected disconnection penalty to dominate the score. Formal results provide guarantees for the mathematical behaviour of the screening system, while advanced figures reveal trade-offs and uncertainty that a conventional choropleth cannot show.

Recent multi-hazard and urban-resilience research argues for the distinction between transparent planning indicators and causal or loss-based risk claims [2–4,9,11]. Flood-accessibility studies likewise show that road disruption can alter access to emergency and hospital services, while more operational studies require dynamic network, traffic, facility-capacity, and validation information that is outside the present archive [6,16–18]. The present contribution is therefore narrower: it demonstrates how archived flood-

footprint and route-change outputs can be combined transparently without converting them into unsupported service-loss or joint-hazard claims.

The exact weight-space analysis also clarifies where the planning result is robust and where it is value-dependent. With only two score domains, every ranking change occurs at a finite pairwise crossing, permitting exact integration over the full admissible interval rather than stochastic approximation. This strengthens interpretation of the shortlist and complements the Pareto frontier, which identifies units that remain non-dominated independently of any single interior weight choice. The percentile-normalisation and alternative shortlist thresholds further show that uncertainty arises not only from source data but also from defensible analytical choices.

The principal limitation for reuse is upstream rather than downstream. The archived analytical table and post-processing can be reproduced, but the road-network disruption derivation, flooded-cell threshold, seismic classification rule, and hexagon geometry cannot be independently reconstructed from the retained files. Consequently, the manuscript treats those fields as archived scenario inputs and avoids presenting the repository as a complete end-to-end GIS reconstruction. Future work should preserve raw OSM extracts and dates, explicit closure and speed rules, bridge and one-way handling, snapping parameters, flood-cell thresholds, seismic interpolation and classification settings, polygon geometry, and observed validation data.

6. Conclusions

This article provides an open-data first-order screen for flood-footprint population proxy and potential hospital accessibility disruption in Bucharest. It is not a cascading earthquake–flood risk model, a compound probability model, or a validated estimate of hospital service loss. The two-domain index, separate service alerts, Pareto analysis, exact full-weight-space rank acceptability, contribution decomposition, and descriptive overlap benchmark make the documented downstream workflow transparent and traceable within the archived scope.

Author Contributions: Conceptualisation, H.H., L.M.M. and N.K.; Methodology, H.H., L.M.M. and N.K.; Writing—review and editing, H.H., L.M.M. and N.K. All authors have read and agreed to the published version of the manuscript.

Funding: We would like to acknowledge two projects: SystR Project, grant number 101212761 as well as PARATUS project, grant number 101073954.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data used in this study are openly archived in Zenodo (https://doi.org/10.5281/zenodo.22685023). The archived materials support reproduction of the reported post-processing analyses from the hexagon-level derived table. The reproducibility package stores analytical tables, the variable dictionary, and the reproducibility-boundary record in CSV format; the post-processing workflow is supplied as Python 3.13.5 source code, generated analytical figures are supplied as PNG files, and provenance/workflow documentation is supplied as plain-text or Markdown files. The archive does not claim reconstruction of unavailable historical routing configurations, including undocumented road-closure, speed, bridge, one-way, snapping, or extraction-date settings, nor the original flooded-cell depth threshold or seismic-classification/interpolation settings. The analytical polygon geometry layer is not available; archived spatial identifiers should therefore be interpreted as references to analytical units rather than as a complete GIS layer. Reuse of archived materials is subject to the licence stated in the Zenodo record and to the original licences of the underlying source datasets. These sources include WorldPop population data for Romania (2020), European Commission Joint Research Centre (JRC) Europe 100-year return-period flood-depth data,

OpenStreetMap road-network and hospital data, and public seismic-hazard screening products from the European Facilities for Earthquake Hazard and Risk (EFEHR).

Acknowledgments: The authors thank the reviewers for their constructive comments and valuable suggestions. During manuscript preparation, ChatGPT (GPT-5.6, OpenAI; accessed September 2026) was used to assist with grammar and sentence structure.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kappes, M.S.; Keiler, M.; von Elverfeldt, K.; Glade, T. Challenges of Analyzing Multi-Hazard Risk: A Review. *Nat. Hazards* **2012**, *64*, 1925–1958. [CrossRef]
2. De Angeli, S.; Malamud, B.D.; Rossi, L.; Taylor, F.E.; Trasforini, E.; Rudari, R. A Multi-Hazard Framework for Spatial-Temporal Impact Analysis. *Int. J. Disaster Risk Reduct.* **2022**, *73*, 102829. [CrossRef]
3. Stolte, T.; Koks, E.E.; de Moel, H.; Reimann, L.; van Vliet, J.; de Ruiter, M.C.; Ward, P.J. VulneraCity—Drivers and Dynamics of Urban Vulnerability Based on a Global Systematic Literature Review. *Int. J. Disaster Risk Reduct.* **2024**, *108*, 104535. [CrossRef]
4. White, C.J.; Adnan, M.S.G.; Arosio, M.; Buller, S.; Cha, Y.; Ciurean, R.; Crummy, J.M.; Duncan, M.; Gill, J.; Kennedy, C.; et al. Towards Multi-Hazard and Multi-Risk Indicators: A Review and Recommendations for Development and Implementation. *Nat. Hazards Earth Syst. Sci.* **2025**, *25*, 4263–4281. [CrossRef]
5. Zhang, Y.; Li, X.; Kong, N.; Zhou, M.; Zhou, X. Spatial Accessibility Assessment of Emergency Response of Urban Public Services in the Context of Pluvial Flooding Scenarios: The Case of Jiaozuo Urban Area, China. *Sustainability* **2022**, *14*, 16332. [CrossRef]
6. Wassmer, J.; Bryant, S.; Schimansky, P.; Keegan, L.T.; Pregnolato, M.; Kurths, J.; Marwan, N.; Merz, B. Unveiling Hidden Risks in Healthcare from Flood-Induced Transportation Disruption in Germany. *Commun. Earth Environ.* **2025**, *6*, 676. Correction in *Commun. Earth Environ.* **2026**, *7*, 444. [CrossRef]
7. AghaKouchak, A.; Huning, L.S.; Sadegh, M.; Qin, Y.; Markonis, Y.; Vahedifard, F.; Love, C.A.; Mishra, A.; Mehran, A.; Obringer, R.; et al. Toward Impact-Based Monitoring of Drought and Its Cascading Hazards. *Nat. Rev. Earth Environ.* **2023**, *4*, 582–595. [CrossRef]
8. Arvin, M.; Beiki, P.; Hejazi, S.J.; Sharifi, A.; Atashafrooz, N. Assessment of Infrastructure Resilience in Multi-Hazard Regions: A Case Study of Khuzestan Province. *Int. J. Disaster Risk Reduct.* **2023**, *88*, 103601. [CrossRef]
9. Kapucu, N.; Ge, Y.; Rott, E.; Isgandar, H. Urban Resilience: Multidimensional Perspectives, Challenges and Prospects for Future Research. *Urban Gov.* **2024**, *4*, 162–179. [CrossRef]
10. Dunant, A.; Robinson, T.R.; Densmore, A.L.; Rosser, N.J.; Rajbhandari, R.M.; Kincey, M.; Li, S.; Awasthi, P.R.; Van Wyk de Vries, M.; Guragain, R.; et al. Impacts from Cascading Multi-Hazards Using Hypergraphs: A Case Study from the 2015 Gorkha Earthquake in Nepal. *Nat. Hazards Earth Syst. Sci.* **2025**, *25*, 267–285. [CrossRef]
11. Ferrario, D.M.; Sanò, M.; Maraschini, M.; Critto, A.; Torresan, S. Harnessing Data-Driven Methods for Climate Multi-Hazard and Multi-Risk Assessment. *Nat. Hazards Earth Syst. Sci.* **2026**, *26*, 2975–3024. [CrossRef]
12. WorldPop. *WorldPop Population Data for Romania, 2020*; WorldPop, University of Southampton: Southampton, UK, 2020.
13. OpenStreetMap Contributors. *OpenStreetMap*; OpenStreetMap Foundation: Cambridge, UK, 2025.
14. European Commission Joint Research Centre. *Europe 100-Year Return-Period Flood-Depth Raster*; European Commission, Joint Research Centre (JRC): Ispra, Italy, 2025.
15. European Facilities for Earthquake Hazard and Risk. *European Seismic Hazard Model 2020 Public Screening Products*; EFEHR Consortium/Swiss Seismological Service, ETH Zurich: Zurich, Switzerland, 2025.
16. Papilloud, T.; Steiner, A.; Zischg, A.; Keiler, M. Road Network Disruptions during Extreme Flooding Events and Their Impact on the Access to Emergency Medical Services: A Spatiotemporal Vulnerability Analysis. *Sci. Total Environ.* **2024**, *956*, 177140. [CrossRef]
17. Mitropoulos, L.; Karolemeas, C.; Tsigdinos, S. Road Network Accessibility Assessment during Flood Events by Using Infrastructure Facility-Based Indices. *Int. J. Disaster Risk Reduct.* **2025**, *122*, 105475. [CrossRef]
18. Abebe, Y.A.; Pregnolato, M.; Jonkman, S.N. Flood Impacts on Healthcare Facilities and Disaster Preparedness: A Systematic Review. *Int. J. Disaster Risk Reduct.* **2025**, *119*, 105340. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.