Probabilistic, Traffic-Weighted Prioritization of Post-Earthquake Bridge Inspections: A Reconstruction of Inspection Priorities for the 24 August 2014 South Napa Earthquake

Abstract

Immediately after a damaging earthquake, transportation agencies must deploy a finite number of inspectors across a large bridge inventory before the operational state of the network is known. We reconstruct a defensible, fully probabilistic inspection-prioritization scheme for the 24 August 2014 *M* 6.0 South Napa earthquake (U.S. Geological Survey event `nc72282711`) by coupling four authoritative data products: the reviewed USGS ShakeMap Atlas ground-motion grid (Atlas product `1624995941193`), the 2014 California National Bridge Inventory (NBI), the FEMA Hazus 6.1 highway-bridge fragility library, and Hazus restoration (downtime) functions. Peak ground acceleration (PGA), spectral acceleration at 0.3 s and 1.0 s, and their lognormal standard deviations were bilinearly interpolated from the ShakeMap grid to the coordinates of all 24,880 quality-controlled California bridges; 156 bridges fell within the Modified Mercalli Intensity (MMI) VI or greater footprint and were retained. Each bridge's NBI attributes (material, structure type, span count, structure length, year built, and skew) were programmatically mapped to one of the 28 Hazus bridge classes (HWB1–HWB28), honoring California-specific seismic design eras. A 10,000-trial Monte Carlo simulation per bridge propagated both ground-motion and fragility uncertainty to yield discrete damage-state probabilities, which were converted to an expected loss-of-function (downtime) duration through Hazus restoration curves and then weighted by average daily traffic (ADT) to produce expected vehicle-days of lost function. Ranking the 156 exposed bridges by this traffic-weighted downtime metric shows that risk is highly concentrated: the top 20 structures capture **81.29%** of the network's expected 19,025,725 vehicle-days of loss of function, while the top 5 alone capture 45.0%. We report the exact 20 NBI structure identifiers to inspect first, together with their Hazus class, expected downtime, and traffic-weighted impact; the highest-priority structure, `21 0098` in Napa County, contributes 3.58 million vehicle-days by itself. We deliberately omit two effects and state their direction: ground-failure/geotechnical hazards (liquefaction, lateral spreading, landslide), whose exclusion likely *under*-estimates damage at soft-ground crossings, and network rerouting/redundancy, whose exclusion *over*-states the relative importance of high-traffic bridges with good detours and *under*-states low-traffic but non-redundant links. The framework is transparent, reproducible, and executable within minutes of a ShakeMap release, offering a risk-informed alternative to epicentral-distance triage.

Keywords: post-earthquake bridge inspection; South Napa earthquake; USGS ShakeMap; Hazus fragility; Monte Carlo uncertainty propagation; loss of function; transportation resilience

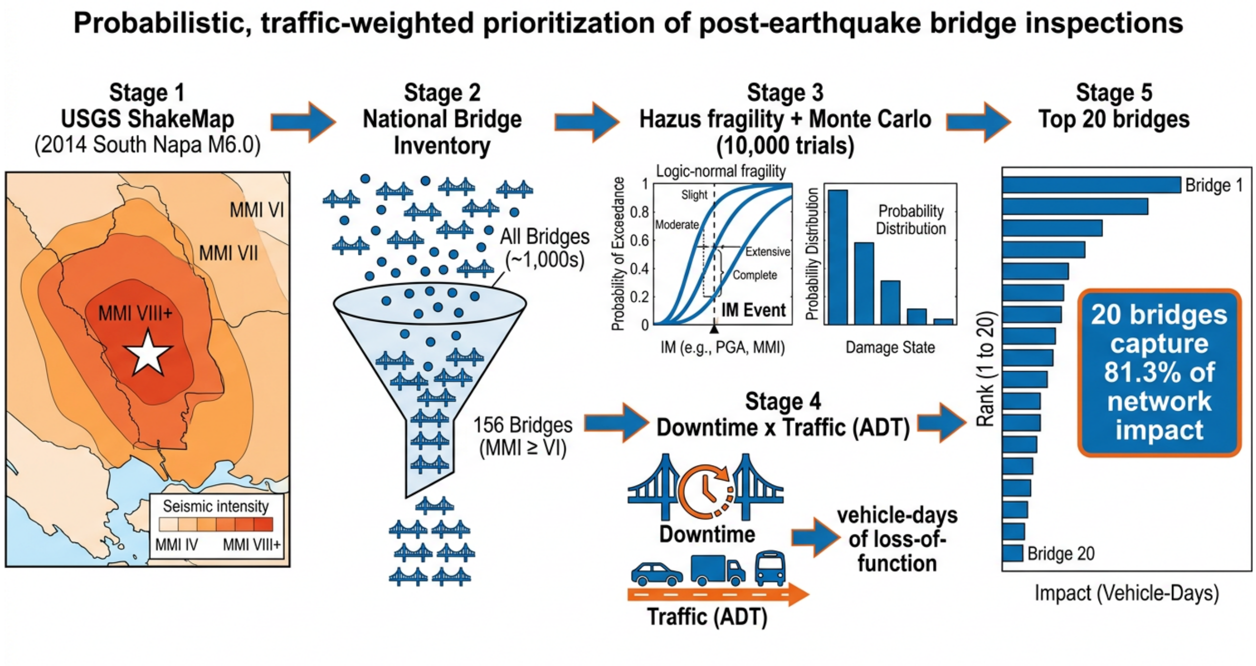

Figure 1: Graphical abstract. The analysis chain proceeds from the reviewed USGS ShakeMap Atlas ground-motion field for the 2014 South Napa earthquake, through spatial interpolation and MMI $\geq$ VI filtering of the 2014 California National Bridge Inventory, to Hazus fragility evaluation under Monte Carlo uncertainty, and finally to traffic-weighted expected loss-of-function days. Twenty bridges capture 81.3% of the network's expected vehicle-days of lost function.

1 Introduction

At 03:20:44 local time on 24 August 2014, a right-lateral strike-slip rupture on the West Napa fault system produced the $M\,6.0$ South Napa earthquake (USGS event `nc72282711`), the most damaging seismic event to strike the San Francisco Bay Area since the 1989 $M\,6.9$ Loma Prieta earthquake. The rupture propagated predominantly unilaterally toward the north and up-dip, focusing energy toward the City of Napa, where late-Holocene alluvial flood deposits amplified ground motion and concentrated damage [4, 16]. Near-fault stations recorded peak ground accelerations approaching 0.5–0.6 g and pulse-like velocity waveforms, and Modified Mercalli Intensities reached VII–VIII within the Napa Valley [4]. Although the event was of only moderate magnitude, it demonstrated a recurring lesson of earthquake engineering: a compact, shallow, directive rupture beneath a sediment-filled basin can impose near-design-level demands on the built environment over a spatially limited but densely developed footprint.

For the agencies responsible for the highway network—principally the California Department of Transportation and affected county public-works departments—the hours immediately following such an event pose an acute decision problem. Hundreds of bridges lie within the felt area, yet the number of qualified inspectors who can be dispatched in the first operational period is small. Bridges that have sustained damage severe enough to threaten public safety or to sever critical links must be identified and closed or restricted before the traveling public, emergency responders, and recovery logistics converge on them. The naive triage strategy—inspect the bridges nearest the epicenter first—is demonstrably inefficient, because seismic demand is governed not by epicentral distance alone but by rupture directivity, basin response, local site conditions, and, crucially, by the vulnerability of each individual structure [28, 34]. A modern, defensible triage must therefore fuse an estimate of the ground motion experienced at each bridge with an estimate of that bridge's fragility, and then weight the resulting damage estimate by the consequence of losing the structure.

The intellectual foundation for this fusion was laid by Ranf et al. [28], who combined USGS ShakeMaps, a bridge inventory, and bridge-specific fragility curves to rank Washington State bridges after the 2001 $M\,6.8$ Nisqually earthquake. Their central result—that a fragility-based ranking could identify roughly 80% of the moderately damaged bridges while inspecting only about 14% of the inventory, far outperforming an epicentral-distance rule—established risk-informed prioritization as the benchmark against which later methods are measured. The intervening two decades have seen the maturation of every ingredient of that pipeline. The USGS ShakeMap system has evolved from the original TriNet implementation [34] into a rigorously documented, uncertainty-aware interpolation framework [36, 37] that was explicitly designed as an operational tool for earthquake response and loss estimation [35], and the ShakeMap Atlas now provides consistently reprocessed, reviewed ground-motion fields for historical events [3], including the reviewed product for South Napa used here. In parallel, the analytical and empirical bridge-fragility literature has both underpinned and critiqued the default Hazus highway-bridge classes [5, 7, 23, 25], and probabilistic seismic risk assessment has embraced Monte Carlo simulation as the standard vehicle for propagating ground-motion and fragility uncertainty through to portfolio-level loss estimates [2, 15, 29].

Despite this maturity, a fully reproducible reconstruction that stitches the reviewed ShakeMap Atlas product, the contemporaneous National Bridge Inventory, the current Hazus 6.1 fragility and restoration library, and explicit Monte Carlo uncertainty propagation into a single traffic-weighted ranking—and that reports the exact structures to inspect—has not been documented for the South Napa event. Such a reconstruction is valuable for three reasons. First, South Napa is unusually well characterized: its ground motions, rupture, and (notably limited) ground failure have been extensively studied [4, 8, 16], making it an ideal validation case for a screening methodology. Second, the exercise exposes, in a concrete setting, exactly how inventory attributes propagate through the Hazus class taxonomy and how the resulting damage estimates combine with traffic exposure to concentrate network risk in a handful of structures. Third, and most importantly for responsible practice, a transparent reconstruction makes explicit what the screening tool does *not* model—here, geotechnical ground failure and network rerouting—so that the resulting priority list is used with appropriate caution.

This paper makes the following contributions. We (i) interpolate PGA, SA(0.3 s), and SA(1.0 s), together with their lognormal standard deviations, from the reviewed South Napa ShakeMap Atlas grid to every California bridge, and filter to the 156 bridges experiencing MMI $\geq$ VI; (ii) map the NBI attributes of those bridges to the 28 Hazus bridge classes with California seismic-design eras; (iii) propagate ground-motion and fragility uncertainty through a 10,000-trial per-bridge Monte Carlo simulation to obtain discrete damage-state probabilities; (iv) convert those probabilities to expected loss-of-function days via Hazus restoration curves and weight them by average daily traffic to obtain expected vehicle-days of lost function; and (v) rank the exposed bridges, deliver the exact top-20 structure identifiers with their expected damage and downtime, and quantify the share of the network's predicted vehicle-days that inspecting those 20 structures captures. Throughout, we state the direction of the bias introduced by omitting ground failure and network rerouting.

2 Methodology

The analysis pipeline comprises six sequential stages: data acquisition and quality control; spatial interpolation of the ground-motion field and MMI filtering; mapping of NBI attributes to Hazus bridge classes; Monte Carlo uncertainty propagation and damage-state estimation; downtime and loss-of-function modeling; and prioritization with network-impact synthesis. Each stage is implemented as a self-contained, seed-controlled Python module, and every stage writes a quality-control (QC) report so that intermediate results are auditable. We describe each stage in turn.

2.1 Data acquisition and quality control

Three primary data products were retrieved. The reviewed *ShakeMap Atlas* ground-motion grid for event `nc72282711` (Atlas product `1624995941193`) and its companion uncertainty grid were downloaded directly from the USGS product repository [31]. The Atlas is the systematically reprocessed, reviewed archive of ShakeMaps generated with a consistent, modern methodology [3, 36]; using the reviewed product rather than an early automatic map ensures that the ground-motion field reflects the best available reconciliation of instrumental recordings, macroseismic intensities, and ground-motion prediction equations. The grid contains 168,633 nodes (457 longitude by 369 latitude, at approximately 0.008,3°, or roughly 1 km, spacing) spanning longitude $[-124.17, -120.37]$ and latitude $[36.72, 39.78]$. Each node carries median MMI, PGA ($\%g$), peak ground velocity, and 5%-damped pseudo-spectral acceleration at 0.3 s, 1.0 s, and 3.0 s, together with the per-node lognormal standard deviations of the ground-motion parameters ($\sigma_{\ln \mathrm{PGA}}$, $\sigma_{\ln \mathrm{SA}}$) and the intensity standard deviation. Reported medians span MMI 2.1–8.2 and PGA up to $\sim$59 $\%g$, consistent with a moderate event with a Napa epicenter.

The *2014 California National Bridge Inventory* was obtained from the comma-delimited FHWA release [11], and the field semantics follow the FHWA Recording and Coding Guide [10]. Of 25,406 raw California records, 24,880 bridges survived quality control: records with invalid coordinates were dropped (522), records with non-positive average daily traffic were removed (4), and duplicate structure numbers were resolved. The NBI packs geographic coordinates as `DDMMSSss` integers (Item 16 latitude, Item 17 longitude); these were converted to decimal degrees, with longitudes assigned to the western hemisphere, and the conversion was verified against a bridge of known location. The cleaned coordinate ranges (latitude 32.54–42.00°N, longitude -124.40 to -114.18°W) lie within California, confirming the conversion. For each bridge we retained structure number, county code, year built, average daily traffic (ADT), structure kind (Item 43A), structure type (Item 43B), and number of spans (Item 45); structure length (Item 49) and degrees of skew (Item 34) were streamed separately from the raw file and joined on structure number where required downstream.

2.2 Spatial interpolation and MMI filtering

The ShakeMap grid was reconstructed from its one-dimensional tabular form into a regular two-dimensional lattice by identifying the 457 unique longitudes and 369 unique latitudes and reshaping each field into an `(nlat, nlon)` array via explicit index lookups, a procedure robust to the row order of the source file. The uniformity of the grid spacing (standard deviation of node spacing $\sim 5 \times 10^{-5}$ degrees) confirms a genuinely regular lattice, which justifies bilinear interpolation. For each ground-motion field we constructed a `scipy` regular-grid interpolator with bilinear (piecewise-linear) weights, no extrapolation, and `NaN` fill outside the footprint. We interpolated eight fields to the exact coordinates of every bridge: the four medians (MMI, PGA, SA(0.3 s), SA(1.0 s)) and their four uncertainty companions (σ_{MMI}, $\sigma_{\ln \mathrm{PGA}}$, $\sigma_{\ln \mathrm{SA}(0.3)}$, $\sigma_{\ln \mathrm{SA}(1.0)}$). Interpolation was performed in each field's native units—medians in the ShakeMap reporting units and uncertainties in natural-log units—so that the lognormal dispersion is preserved for the subsequent Monte Carlo step. Bilinear interpolation on the dense (∼1 km) ShakeMap lattice is appropriate because the reviewed field is already a smooth, spatially correlated surface; spatial-correlation structure at sub-kilometer scale, which matters for multi-site joint-loss estimation [14, 30, 32], is not re-introduced here because each bridge is treated as an independent inspection target rather than as part of a joint-loss integral. This choice is immaterial to the ranking and to the cumulative-capture statistic reported below: both depend only on the per-bridge marginal expectations, and the expected vehicle-days of lost function are linear in the marginal damage-state probabilities, so inter-site correlation would affect only the *variance* of the network-aggregate loss—which we do not report—and not its expectation or the induced ordering.

Of the 24,880 bridges, 15,889 fell outside the ShakeMap footprint (assigned `NaN` and dropped) and 8,991 lay within it. Applying the strong-shaking threshold of MMI $\geq$ 6.0 (intensity VI, the conventional onset of potentially damaging shaking in the empirical instrumental-intensity relationships that underlie ShakeMap [33]) yielded **156** exposed bridges, with zero missing values in any interpolated field. Across this filtered set, interpolated MMI ranges 6.01–8.11, PGA 10.3–55.8 %g, SA(0.3 s) 26.3–109.8 %g, and SA(1.0 s) 8.7–58.9 %g; the ground-motion standard deviations range σ_{MMI} 0.49–0.73 and $\sigma_{\ln \mathrm{PGA}}$ 0.28–0.42. As expected for a compact $M\,6$ source, the exposed bridges cluster tightly around the Napa epicenter near 38.2°N, −122.3°W (Figure 2).

2.3 Mapping NBI attributes to Hazus bridge classes

The Hazus earthquake model represents highway-bridge vulnerability through 28 bridge classes (HWB1–HWB28), each associated with a lognormal fragility curve for four damage states—slight, moderate, extensive, and complete [9, 21]. The lognormal parameterization and its statistical estimation from analytical or empirical damage data are well established in the fragility literature [19]. Assigning each NBI bridge to a Hazus class is the pivotal modeling translation, because the class governs both the fragility medians and the demand parameter used downstream. We implemented a deterministic, fully documented classification that mirrors the Hazus logic and preserves the intermediate reasoning for auditability.

First, structure length (NBI Item 49) was recovered for all 156 bridges from the raw NBI file (matched 156/156; lengths 6.7–1587.7 m, median 27.0 m). Bridges longer than 150 m with more than one span were assigned the *major bridge* class HWB1; single-span bridges (span count equal to one) were assigned HWB2. The remaining multi-span bridges were routed through thirteen structural groups defined by material family (concrete, prestressed concrete, steel, wood, masonry) and structural type (girder/beam, box, truss/arch, continuous), derived from NBI Items 43A and 43B. Within each structural group, the bridge was placed into one of a consecutive pair of Hazus classes according to its seismic-design vintage, using the year built (NBI Item 27) with a 1975 threshold that reflects the watershed in California bridge seismic design that followed the 1971 San Fernando earthquake—bridges built before 1975 are treated as conventionally (non-seismically) designed, and those built in 1975 or later as seismically designed. This year-of-construction stratification is directly supported by the empirical finding that California bridge fragility depends strongly on design era [27]. For transparency we retained a California seismic-design-era label with three bins (`pre_1975`, `1975-1990`, `post_1990`) alongside the HWB class and the intermediate material/type/seismic flags, so that canonical Hazus fragility parameters can be attached independently of the class label.

All 156 bridges mapped to a valid HWB class with no nulls or invalid codes. The resulting distribution is dominated by single-span bridges (HWB2, 44 bridges), continuous concrete construction (HWB9, 28), non-seismically designed concrete girder bridges (HWB3, 23), and major bridges (HWB1, 18), with the remainder spread across ten further classes (Table 2). By design era, 101 bridges are pre-1975, 26 are 1975–1990, and 29 are post-1990, so that a majority of the exposed inventory predates modern seismic detailing—an important context for the damage estimates that follow.

2.4 Uncertainty propagation and damage-state estimation

Damage-state probabilities were estimated by Monte Carlo simulation that jointly samples the two dominant sources of uncertainty: the ground motion experienced at the bridge (aleatory and epistemic variability encoded in the ShakeMap lognormal standard deviations) and the structural capacity at each damage state (fragility dispersion). We adopted Hazus 6.1 baseline lognormal fragility medians for all 28 classes [9], expressed in units of g. Following the Hazus convention, the demand parameter is PGA for single-span bridges (HWB2) and 1.0 s spectral acceleration for all multi-span classes (HWB1 and HWB3–HWB28); of the 156 bridges, 44 are driven by PGA and 112 by SA(1.0 s). This choice reflects the physical reality that the fundamental period of most multi-span highway bridges lies near 1 s, so that SA(1.0 s) is the more informative intensity measure for their seismic response [7, 25], a convention consistent with the component-based analytical fragility methodology developed for multi-span highway bridges and their retrofitted variants [24, 26].

Bridge skew, extracted from NBI Item 34 for all 156 bridges (mean 12.4°, capped at 60°), was incorporated through the Hazus skew-adjustment factor $K_{\text{skew}} = 1-\sin\alpha$ for multi-span bridges (with $K_{\text{skew}} = 1$ for single-span bridges), reducing the effective fragility median to $m_{ds} = K_{\text{skew}}\, m_{ds}^{\text{baseline}}$ and thereby increasing the damage probability of skewed structures; 87 of the 156 bridges carry non-zero skew. The inclusion of skew is motivated by evidence that geometric irregularity materially increases bridge seismic vulnerability and is imperfectly captured by generic fragility classes [22].

For each bridge, the simulation drew 10,000 trials (vectorized, fixed seed for reproducibility). In each trial, the demand intensity was sampled as $I = \exp(\ln m_I + Z_d\, \sigma_{\ln I})$, where m_I is the interpolated median intensity, $\sigma_{\ln I}$ the interpolated lognormal standard deviation, and Z_d a standard normal deviate; the capacity for each damage state was sampled as $C_{ds} = \exp(\ln m_{ds} + Z_c\, \beta)$ with a capacity dispersion $\beta = 0.40$. A single capacity deviate Z_c was shared across the four damage states within a trial, which preserves the ordinal nesting of the damage states ($C_{\text{slight}} \leq C_{\text{moderate}} \leq C_{\text{extensive}} \leq C_{\text{complete}}$) and thereby guarantees monotonically non-increasing exceedance probabilities and non-negative discrete probabilities. The exceedance probability of a damage state is the fraction of trials in which $I \geq C_{ds}$; discrete state probabilities follow by successive differencing. This nested-capacity Monte Carlo scheme is a standard and computationally efficient route to propagating combined demand and capacity uncertainty in portfolio seismic risk [2, 15, 29].

Validation confirmed that discrete probabilities summed to unity within floating-point tolerance (maximum absolute deviation 2.2×10^{-16}), that no discrete probability was negative, and that exceedance probabilities were monotonically non-increasing for every bridge. Averaged over the 156 bridges, the mean discrete probabilities are 0.806 (no damage), 0.097 (slight), 0.043 (moderate), 0.030 (extensive), and 0.024 (complete). Summing these probabilities across the inventory yields expected aggregate counts of roughly 125.7 undamaged, 15.1 slightly, 6.7 moderately, 4.7 extensively, and 3.8 completely damaged bridges—a distribution consistent with the observed, comparatively light bridge damage of the actual event.

2.5 Downtime and loss-of-function modeling

The consequence of interest for inspection triage is not damage *per se* but the loss of function that damage imposes on the traveling public. We mapped the discrete damage-state probabilities to an expected downtime (loss-of-function duration) using the Hazus 6.1 highway-bridge restoration times, which assign representative durations of 0.6, 2.5, 75.0, and 230.0 days to the slight, moderate, extensive, and complete damage states, respectively (zero for no damage) [9, 20]. The expected downtime of a bridge is the probability-weighted sum

$$\mathbb{E}[D] = 0.6\,P_{\text{slight}} + 2.5\,P_{\text{moderate}} + 75.0\,P_{\text{extensive}} + 230.0\,P_{\text{complete}}, \tag{1}$$

expressed in days. Because the extensive and complete states carry by far the largest restoration times, $\mathbb{E}[D]$ is dominated by the tail of the damage distribution, so that bridges with even modest probabilities of severe damage accrue large expected downtimes.

To translate downtime into network consequence, we weighted each bridge's expected downtime by its average daily traffic to obtain the expected traffic-weighted downtime, or expected vehicle-days of lost function:

$$\mathbb{E}[V] = \mathbb{E}[D] \times \text{ADT}. \tag{2}$$

This quantity estimates the aggregate travel demand that would encounter a non-functional structure over the course of its restoration, and it is the metric on which bridges are ranked. Across the 156 bridges, expected downtime ranges from 0.001 to 150.5 days (median 1.74 days, mean 7.96 days), and expected vehicle-days of lost function range from 0.62 to 3,582,373 (median 8,373, mean 121,960), summing to a network total of 19,025,725 vehicle-days. The distribution of both quantities is strongly right-skewed (Figure 3), foreshadowing the concentration of network impact documented in the Results.

2.6 Prioritization and network-impact synthesis

The 156 exposed bridges were ranked in descending order of expected vehicle-days of lost function (Equation 2). A deterministic secondary sort key (structure number) breaks any ties, so that every bridge receives a unique integer priority rank from 1 (highest priority) to 156, with no gaps or ties. We then computed the cumulative share of the network total captured as a function of the number of bridges inspected, and in particular the share captured by the top 20 structures, which constitute a plausible first-wave inspection assignment. Rank integrity was verified (ranks cover 1–156 without gaps; the metric is monotonically non-increasing along the rank order), and the metric contained no negative or missing values. The complete ranked table, the network-capture metrics, and the top-20 structure identifiers are provided as machine-readable outputs.

3 Results

3.1 Ground-motion exposure of the bridge inventory

The reviewed ShakeMap Atlas field and the 156 exposed bridges are shown together in Figure 2. The intensity field is strongly elongated along the northwest–southeast trend of the West Napa fault and is amplified within the Napa basin, so that the exposed bridges form a compact cluster around the City of Napa rather than a circular halo about the epicenter. This spatial pattern is the physical basis for preferring a shaking- and fragility-informed triage over an epicentral-distance rule: several bridges at comparable epicentral distances experience markedly different interpolated intensities depending on their position relative to the rupture and the basin, and the ranking must reflect the shaking each structure actually experienced.

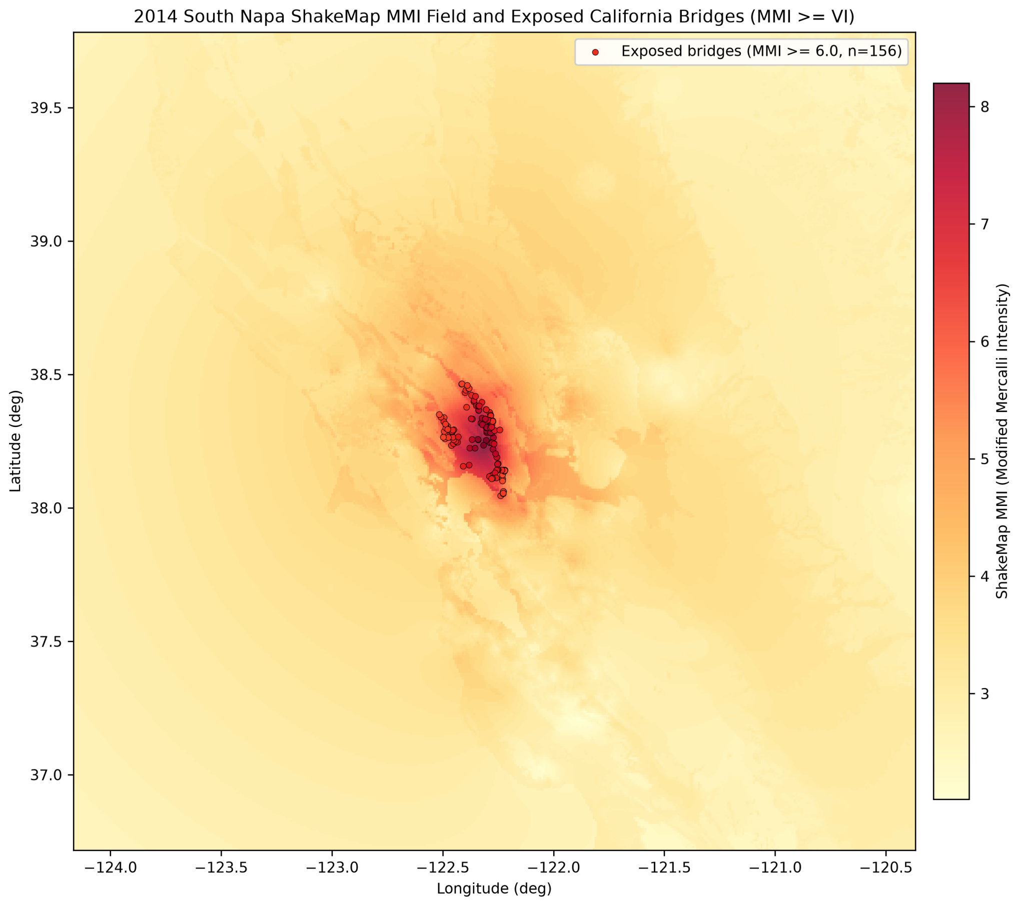

Figure 2: Reviewed USGS ShakeMap Atlas Modified Mercalli Intensity field for the 2014 South Napa earthquake (event `nc72282711`, Atlas product `1624995941193`) with the 156 exposed California bridges (MMI $\geq$ VI) overlaid as points. The intensity field is elongated along the fault strike and amplified within the Napa basin; exposed bridges cluster tightly around the City of Napa.

3.2 Distribution of expected downtime and traffic-weighted impact

Figure 3 summarizes the per-bridge distributions of expected downtime and expected vehicle-days of lost function. The expected-downtime histogram is dominated by bridges with sub-week expected downtimes, but carries a long tail of structures with expected downtimes of tens to more than one hundred days—these are bridges with appreciable probabilities of extensive or complete damage, for which the 75- and 230-day restoration times of Equation 1 dominate. When downtime is weighted by traffic, the distribution becomes even more skewed: the top bridges combine substantial expected downtime with high ADT, and the impact-concentration curve (Figure 3, right panel) rises steeply, indicating that a small number of structures account for a disproportionate share of the network's expected vehicle-days of lost function.

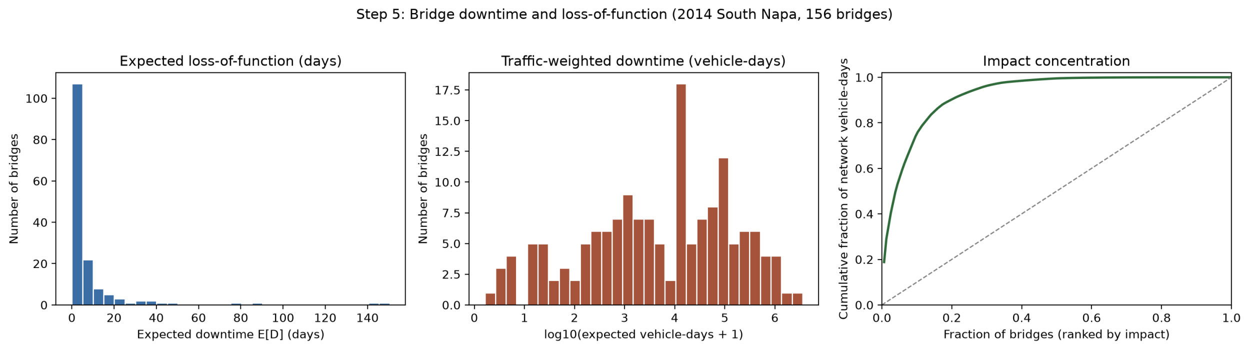

Figure 3: Per-bridge distributions of expected loss of function for the 156 exposed bridges. Left: histogram of expected downtime $\mathbb{E}[D]$ in days. Center: histogram of expected traffic-weighted downtime (vehicle-days, log scale). Right: the impact-concentration curve, showing the cumulative fraction of network vehicle-days as a function of the fraction of bridges ranked by impact. The steep concentration curve motivates a small first-wave inspection assignment.

3.3 Concentration of network impact and the top-20 capture

The central quantitative result of the study is the degree to which the network's expected loss of function concentrates in a small number of bridges. The 156 exposed bridges carry a total expected loss of function of 19,025,725 **vehicle-days**. The cumulative-capture curve (Figure 4) shows that the top 5 priority bridges capture 45.0% of this total, the top 10 capture 62.0%, the top 20 capture **81.29%** (15,465,162 vehicle-days), and the top 50 capture 97.0%. In other words, inspecting the 20 highest-priority structures—roughly 13% of the exposed inventory and a far smaller fraction of the full California inventory within the felt area—addresses more than four-fifths of the expected network travel disruption. This concentration is the quantitative analogue, for the South Napa event, of the efficiency gain reported by Ranf et al. [28] for Nisqually, and it is precisely the property that makes a first-wave inspection assignment both feasible and effective.

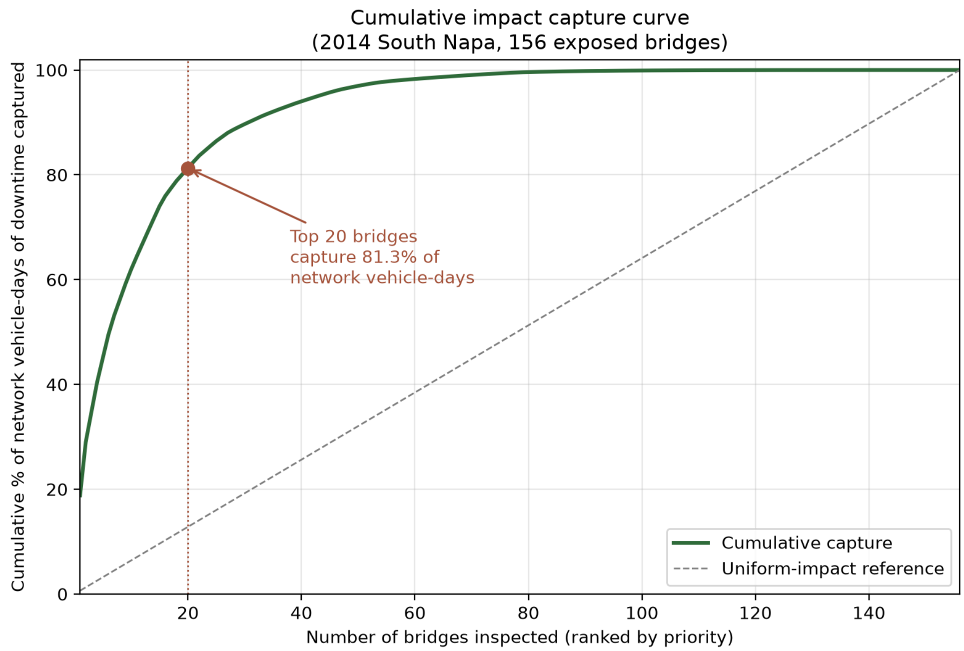

Figure 4: Cumulative share of the network's expected vehicle-days of lost function captured as a function of the number of bridges inspected, in priority order. The dashed line is the uniform-impact reference (equal impact per bridge). The top 20 priority bridges capture 81.3% of the network total, far above the ∼13% that a uniform assignment would capture.

3.4 The top 20 bridges to inspect first

Table 1 lists the exact 20 NBI structure identifiers that constitute the recommended first-wave inspection assignment, together with their county, year built, ADT, Hazus class, expected downtime, and expected traffic-weighted impact. The corresponding spatial distribution and impact bar chart are shown in Figure 5. County FIPS codes are 55 (Napa), 95 (Solano), and 97 (Sonoma); 18 of the top 20 bridges lie in Napa County, consistent with the concentration of intensity within the Napa basin.

Table 1: The 20 highest-priority bridges for post-earthquake inspection after the 2014 South Napa earthquake, ranked by expected traffic-weighted downtime (expected vehicle-days of lost function). County FIPS: 55 = Napa, 95 = Solano, 97 = Sonoma.

Rank	Structure ID	County	Year	ADT	Hazus	$\mathbb{E}[D]$ (days)	$\mathbb{E}[V]$ (veh-days)
1	21 0098	55	1981	45,000	HWB4	79.61	3,582,373
2	21 0049	55	1977	45,000	HWB1	43.05	1,937,442
3	21 0087	55	1961	7,260	HWB17	150.51	1,092,699
4	21 0101	55	2003	30,000	HWB22	35.96	1,078,684
5	23 0127	95	1997	133,000	HWB4	6.56	873,116
6	21 0048	55	1958	47,000	HWB9	18.30	859,999
7	21C0128	55	1991	8,002	HWB10	85.57	684,728
8	21 0088	55	1964	23,200	HWB15	24.84	576,218
9	21C0012	55	2002	12,050	HWB20	47.60	573,544
10	21C0126	55	2004	18,241	HWB22	28.95	528,112
11	21 0108L	55	2007	13,500	HWB1	35.32	476,870
12	21 0003	55	1918	22,850	HWB3	20.23	462,235
13	21 0105	55	1995	23,700	HWB10	19.30	457,363
14	21C0006	55	1972	3,211	HWB3	141.56	454,552
15	21 0108R	55	2007	13,500	HWB1	33.16	447,696
16	21 0014	55	1970	30,500	HWB9	11.90	362,948
17	21 0043R	55	1979	22,625	HWB10	12.40	280,477
18	21 0043L	55	1979	22,625	HWB10	12.09	273,555
19	23 0114	95	1958	25,425	HWB17	9.11	231,594
20	21 0022L	55	1964	23,500	HWB9	9.83	230,959
						Top-20 total	15,465,162
						Network total (156 bridges)	19,025,725
						Share captured by top 20	**81.29%**

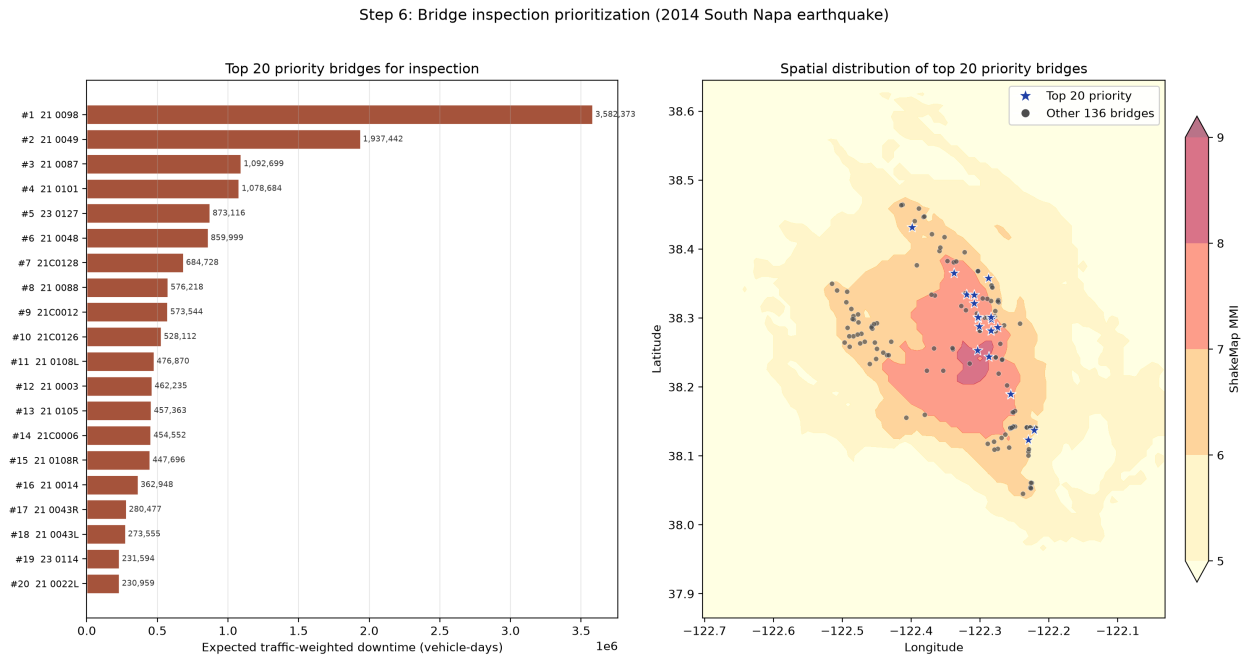

Figure 5: The 20 highest-priority bridges. Left: expected traffic-weighted downtime (vehicle-days) for each of the top 20 structures, labeled by priority rank and NBI structure identifier. Right: spatial distribution of the top 20 priority bridges (stars) relative to the other 136 exposed bridges (dots), over the ShakeMap MMI contours. The top-ranked structures concentrate within the high-intensity core near the City of Napa.

The ranking is driven jointly by structural vulnerability and traffic exposure, and the top 20 illustrate the several distinct ways a bridge can rise to the top of the list. The highest-priority structure, `21 0098` in Napa County (Hazus HWB4, a seismically designed continuous concrete bridge on a high-volume corridor), combines a large expected downtime (∼80 days) with a high ADT of 45,000 to contribute 3,582,373 vehicle-days—by itself 18.8% of the entire network total. Structure `23 0127` in Solano County enters the top five almost entirely on the strength of traffic exposure: despite an expected downtime of only 6.6 days, its ADT of 133,000 (the highest in the exposed set) produces 873,116 vehicle-days. At the opposite extreme, structures `21 0087` and `21C0006` rank in the top 15 despite modest ADT (7,260 and 3,211, respectively) because their expected downtimes exceed 140 days, reflecting high probabilities of severe damage in older, non-seismically designed construction. The presence of both mechanisms in the top 20 is a feature, not an artifact: a triage metric that captured only vulnerability would overlook heavily trafficked but relatively robust structures, while a metric that captured only traffic would overlook fragile structures on important-but-quieter routes.

The Hazus-class composition of the top 20 is also instructive. It includes both major bridges (HWB1) and a range of multi-span concrete and prestressed classes (HWB3, HWB4, HWB9, HWB10, HWB15, HWB17, HWB20, HWB22), spanning construction years from 1918 to 2007. Notably, several post-1990 (seismically designed) bridges appear in the list—not because their fragility is high, but because they carry very large traffic volumes; conversely, several pre-1975 structures appear because their fragility is high even though their traffic is moderate. This heterogeneity underscores that the traffic-weighted metric is doing genuine work in reconciling the two considerations.

4 Discussion

4.1 Interpretation: why risk concentrates

The dominant empirical finding—that 20 of 156 exposed bridges capture 81.29% of expected network loss of function—reflects the multiplicative structure of the ranking metric. Expected vehicle-days of lost function is the product of an expected downtime that is itself tail-dominated (because the extensive and complete damage states carry restoration times one to two orders of magnitude larger than the slight and moderate states) and a traffic volume that is heavy-tailed across the inventory (ADT spans four orders of magnitude, from single digits to 133,000). The product of two right-skewed quantities is extremely right-skewed, so that a handful of bridges—those unlucky enough to combine high fragility, strong shaking, and heavy traffic—dominate the total. This is the same phenomenon that makes portfolio seismic loss so sensitive to the upper tail of the exposure distribution [2, 15], and it is the property that a first-wave inspection assignment is designed to exploit.

The practical implication is favorable for emergency response: because impact is so concentrated, a small, well-chosen inspection team can address the great majority of the network's expected disruption in the first operational period. The cumulative-capture curve (Figure 4) quantifies the diminishing returns of additional inspections—the marginal bridge beyond rank 50 contributes very little to the captured total—which allows an agency to size its first wave against an explicit impact target rather than an arbitrary distance cutoff. Our reconstruction thus corroborates, for a California event with the current Hazus 6.1 library, the efficiency advantage of fragility-based triage first demonstrated for the Pacific Northwest by Ranf et al. [28] and later embedded in network-resilience and accelerated-repair frameworks [12, 38, 39].

At the same time, the concentration of impact means that the ranking is only as trustworthy as its treatment of the top structures. Structure `21 0098` alone accounts for nearly one-fifth of the network total, so any error in its interpolated ground motion, its Hazus class assignment, or its ADT propagates directly into the headline result. This sensitivity is a general feature of consequence-weighted triage and argues for treating the top of the list as a set of hypotheses to be confirmed by inspection, not as a settled damage assessment.

4.2 Relationship to the observed event

The reconstruction is a *prospective screening* calculation—what a risk-informed triage would have recommended given only the ShakeMap and the inventory—rather than a hindcast of observed damage. Its aggregate output is nonetheless physically reasonable: the expected damage-state counts (roughly 30 bridges with some level of damage, of which fewer than 10 extensively or completely damaged) are consistent with the comparatively light bridge damage actually sustained in the South Napa event, in which no major highway bridge collapsed and damage was concentrated in older, non-ductile construction and in non-structural and geotechnical elements [4, 8]. The value of the screening is not that it predicts precisely which bridges were damaged—no fragility-based method can—but that it orders the inventory so that the structures most likely to have been damaged *and* most consequential if damaged are inspected first.

4.3 Critical omission I: ground-failure and geotechnical hazards

The fragility model used here is driven entirely by ground shaking—PGA for single-span bridges and SA(1.0 s) for multi-span bridges—as interpolated from the ShakeMap. It does *not* model geotechnical ground-failure hazards: liquefaction, liquefaction-induced lateral spreading, and seismically induced landslide or slope failure. This is a material omission, because ground failure is historically one of the dominant causes of severe bridge damage. Case-history syntheses have repeatedly identified lateral spreading as the primary mechanism of liquefaction-induced bridge damage, physically displacing abutments and piers riverward and imposing large kinematic demands on pile foundations [6, 17]. Bridges are most exposed to these mechanisms precisely where they cross soft, saturated, or sloping ground—river crossings, approach embankments, and abutments on young alluvium—which are common settings for exactly the kinds of structures near the top of our list.

The direction of the resulting bias is clear: by omitting ground failure, the model *under-estimates* damage probability, and hence downtime, for the most geotechnically exposed structures. A bridge founded on liquefiable alluvium could be systematically *under-ranked* relative to its true risk, because its shaking-only fragility understates the demand it would actually experience. The South Napa event is an instructive and somewhat paradoxical case in this regard: despite mapped liquefaction susceptibility along the Napa River, field reconnaissance documented a notably limited extent of liquefaction and lateral spreading, attributed to drought-lowered groundwater and clay-rich Holocene soils, with only localized sand boils and minor lateral spreading (for example, near the Third Street bridge) [8]. That the omission happened to matter little for this particular event should not be mistaken for a general license to ignore ground failure; under wetter antecedent conditions, or for an event in a more liquefaction-prone setting, the same omission could grossly under-rank the most vulnerable crossings. A defensible operational deployment would therefore overlay the shaking-based ranking with a liquefaction- and landslide-susceptibility screen, elevating soft-ground and slope-crossing structures irrespective of their shaking-only rank.

4.4 Critical omission II: rerouting, detours, and network topology

The second deliberate omission concerns the consequence side of the metric. Expected vehicle-days of lost function (Equation 2) treats each bridge in isolation and uses the bridge's own ADT as a proxy for the disruption its closure would cause. This is a component-level, not a system-level, measure of consequence: it does not represent the topology of the transportation network, and therefore ignores traffic rerouting—the availability of detours, the added travel time they impose, and the downstream capacity constraints they trigger. The true operational and economic impact of a closure depends on network redundancy. Closing a bridge that is the sole crossing of a river or canyon produces far greater delay than closing an ADT-equivalent bridge that has parallel alternatives a short distance away, because in the redundant case traffic reroutes with modest added travel time while in the non-redundant case it must make a long detour or cannot complete the trip at all [1, 18, 39].

The direction of this bias is two-sided. For high-ADT bridges that possess good detour alternatives, the ADT-weighted metric *over-states* the true disruption, because much of the traffic would simply reroute at low marginal cost. For low-ADT bridges that are critical, non-redundant links—sole crossings, or the only access to a community or a hospital—the metric *under-states* the true disruption, because the loss of the link imposes delays or isolation out of all proportion to its raw traffic count. Capturing these system-level effects rigorously requires a network model with traffic assignment: total system travel time under equilibrium or dynamic assignment, evaluated with and without each candidate closure, is the theoretically appropriate consequence measure, and it is the basis of the resilience-optimization literature on post-earthquake restoration scheduling and inspection routing [18, 20, 38, 39]. Such models were beyond the scope of this ShakeMap-plus-inventory reconstruction, which was designed to be executable within minutes of a ShakeMap release using only nationally available data. In practice, the ADT-weighted ranking should be read as a fast, transparent first screen, to be refined—especially for the redundancy of the top-ranked structures—by a network-flow analysis where time and data permit.

4.5 Further limitations

Beyond the two headline omissions, several modeling choices bound the interpretation. The Hazus default fragilities are generic class curves; the literature documents that they can be biased—sometimes conservative, sometimes unconservative—for specific geometries, continuity conditions, retrofits, and aging states, and that bridge-specific or refined-class fragilities can differ materially [22, 23, 25, 27]. Aging and cumulative damage from prior events, which can shift fragilities over a structure's life, are not represented [13]. The ground-motion field is the reviewed ShakeMap median with its reported dispersion; while the reviewed Atlas product is the best available reconciliation of data and models [3, 37], residual epistemic uncertainty in the ground-motion model is not separately propagated, and inter-site spatial correlation—important if the objective were joint network loss rather than per-bridge triage—is not modeled [14, 30, 32]. The restoration times are Hazus point estimates rather than distributions, so downtime uncertainty is not carried through to the ranking. Finally, ADT is a static, pre-event average that neither reflects post-event demand changes (evacuation, emergency traffic, demand suppression) nor distinguishes vehicle classes (a lifeline route for emergency vehicles may warrant priority beyond its ADT). None of these limitations undermines the central, robust finding—that expected network loss of function is highly concentrated and that a small, risk-ranked first wave captures most of it—but each defines a direction in which a fuller analysis would refine the specific ordering.

5 Conclusion

We have reconstructed a fully probabilistic, traffic-weighted inspection-prioritization scheme for the 24 August 2014 M 6.0 South Napa earthquake by integrating the reviewed USGS ShakeMap Atlas ground-motion field, the 2014 California National Bridge Inventory, the FEMA Hazus 6.1 bridge fragility and restoration library, and per-bridge Monte Carlo uncertainty propagation. Interpolating PGA, SA(0.3 s), and SA(1.0 s) to bridge coordinates and filtering to MMI $\geq$ VI isolates 156 exposed bridges; mapping their NBI attributes to the 28 Hazus classes, propagating combined ground-motion and fragility uncertainty through 10,000 trials per bridge, and weighting the resulting expected downtime by average daily traffic yields an expected-loss-of-function ranking of the entire exposed set. The result is a concrete, actionable deliverable: the 20 NBI structure identifiers that should be inspected first (Table 1), with their expected damage, expected downtime, and traffic-weighted impact. These 20 bridges capture **81.29%** of the network's 19,025,725 expected vehicle-days of lost function, with the single highest-priority structure (`21 0098`, Napa) accounting for nearly one-fifth of the total on its own.

The broader contribution is methodological and operational. The framework is transparent—every translation from inventory attribute to Hazus class to damage probability to downtime is documented and auditable—reproducible under a fixed random seed, and fast enough to run within minutes of a ShakeMap release using only nationally available data. It offers emergency managers a defensible, quantitative alternative to epicentral-distance triage and a means of sizing a first-wave inspection assignment against an explicit impact-capture target. Equally important, we have been explicit about what the tool omits and in which direction each omission biases the result: excluding ground-failure and geotechnical hazards likely *under*-estimates damage at soft-ground and slope crossings, and excluding rerouting and network topology *over*-states the importance of high-traffic bridges with good detours while *under*-stating that of low-traffic but non-redundant links. The natural extensions—overlaying a liquefaction and landslide susceptibility screen, and refining the consequence metric with a network traffic-assignment model for the top-ranked structures—would tighten the ordering without disturbing its central, robust conclusion: after a damaging earthquake, expected network loss of function is concentrated in a small number of bridges, and a risk-informed first wave can find them.

Data and Code Availability

The reviewed ShakeMap Atlas product (event `nc72282711`, Atlas product `1624995941193`) is available from the USGS Earthquake Hazards Program, and the 2014 California NBI file is available from the FHWA. The processed datasets underlying this study—the interpolated exposed-bridge table, the Hazus class mapping, the Monte Carlo damage probabilities, the downtime and loss-of-function estimates, and the final ranked prioritization (`bridge_inspection_prioritization.csv`) with the top-20 summary (`prioritization_summary.json`)—together with the six seed-controlled analysis scripts and stage-by-stage QC reports, accompany this manuscript.

A Hazus bridge-class distribution of the exposed inventory

Table 2: Distribution of the 156 exposed bridges (MMI $\geq$ VI) across Hazus bridge classes and California seismic design eras. HWB1 denotes major bridges (length >150 m, multi-span); HWB2 denotes single-span bridges; HWB3–HWB28 denote multi-span bridges stratified by material family, structural type, and seismic design vintage.

Hazus class	Count	Design era	Count
HWB2 (single span)	44	pre-1975	101
HWB9 (concrete continuous)	28	1975–1990	26
HWB3 (concrete girder)	23	post-1990	29
HWB1 (major bridge)	18		
HWB22 (prestressed cont.)	13		
HWB10 (steel girder)	10		
HWB11 / HWB17	5 each		
HWB4	3		
HWB15 / HWB20	2 each		
HWB21 / HWB23 / HWB25	1 each		

References

[1] Mitsuyoshi Akiyama, Dan M. Frangopol, and Hiroki Ishibashi. Toward life-cycle reliability-, risk- and resilience-based design and assessment of bridges and bridge networks under independent

and interacting hazards: Emphasis on earthquake, tsunami and corrosion. *Structure and Infrastructure Engineering*, 16(1):26–50, 2019. doi: 10.1080/15732479.2019.1604770.

[2] Sinan Akkar and Yin Cheng. Application of a Monte-Carlo simulation approach for the probabilistic assessment of seismic hazard for geographically distributed portfolio. *Earthquake Engineering & Structural Dynamics*, 45(4):525–541, 2015. doi: 10.1002/eqe.2667.

[3] Trevor I. Allen and David J. Wald. *Evaluation of Ground-Motion Modeling Techniques for Use in Global ShakeMap*. Open-File Report 2009-1047. 2009. doi: 10.3133/ofr20091047.

[4] Annemarie S. Baltay and John Boatwright. Ground-motion observations of the 2014 south napa earthquake. *Seismological Research Letters*, 86(2A):355–360, 2015. doi: 10.1785/0220140232.

[5] Nesrin I. Basoz, Anne S. Kiremidjian, Stephanie A. King, and Kincho H. Law. Statistical analysis of bridge damage data from the 1994 Northridge, CA, earthquake. *Earthquake Spectra*, 15(1):25–54, 1999. doi: 10.1193/1.1586027.

[6] Ross W. Boulanger, Dongdong Chang, Scott J. Brandenberg, Richard J. Armstrong, and Bruce L. Kutter. Seismic design of pile foundations for liquefaction effects. In *Earthquake Geotechnical Engineering*, pages 277–302. Springer Netherlands, 2007. doi: 10.1007/978-1-4020-5893-6_12.

[7] Eunsoo Choi, Reginald DesRoches, and Bryant Nielson. Seismic fragility of typical bridges in moderate seismic zones. *Engineering Structures*, 26(2):187–199, 2004. doi: 10.1016/j.engstruct.2003.09.006.

[8] Kathleen M. Darby, Martin W. McIlroy, Ross W. Boulanger, and Jason T. DeJong. Analysis of liquefaction at a bridge site in the 2014 Napa earthquake. In *Geotechnical Frontiers 2017*, pages 277–289. American Society of Civil Engineers, 2017. doi: 10.1061/9780784480489.028.

[9] Federal Emergency Management Agency. Hazus earthquake model technical manual, Hazus 6.1. Technical report, Federal Emergency Management Agency, Washington, D.C., 2024. URL https://www.fema.gov/flood-maps/products-tools/hazus.

[10] Federal Highway Administration. Recording and coding guide for the structure inventory and appraisal of the nation's bridges. Technical Report FHWA-PD-96-001, U.S. Department of Transportation, Federal Highway Administration, Washington, D.C., 1995.

[11] Federal Highway Administration. National bridge inventory (nbi): 2014 California delimited data file. U.S. Department of Transportation, Federal Highway Administration, 2014. URL https://www.fhwa.dot.gov/bridge/nbi/2014/delimited/CA14.txt.

[12] Michael N. Furtado and Alice Alipour. Enhancement of post-earthquake resilience of bridge networks with prioritization and accelerated repair strategies. In *Structures Congress 2014*, pages 2014–2024. American Society of Civil Engineers, 2014. doi: 10.1061/9780784413357.177.

[13] Jayadipta Ghosh, Jamie E. Padgett, and Mauricio Sánchez-Silva. Seismic damage accumulation in highway bridges in earthquake-prone regions. *Earthquake Spectra*, 31(1):115–135, 2015. doi: 10.1193/120812EQS347M.

[14] Nirmal Jayaram and Jack W. Baker. Correlation model for spatially distributed ground-motion intensities. *Earthquake Engineering & Structural Dynamics*, 38(15):1687–1708, 2009. doi: 10.1002/eqe.922.

[15] Nirmal Jayaram and Jack W. Baker. Efficient sampling and data reduction techniques for probabilistic seismic lifeline risk assessment. *Earthquake Engineering & Structural Dynamics*, 39(10):1109–1131, 2010. doi: 10.1002/eqe.988.

[16] Chen Ji, Ralph J. Archuleta, and Cedric Twardzik. Rupture history of 2014 M_w 6.0 south napa earthquake inferred from near-fault strong motion data and its impact to the practice of ground strong motion prediction. *Geophysical Research Letters*, 42(7):2149–2156, 2015. doi: 10.1002/2015GL063335.

[17] Kemin Jia, Chengshun Xu, M. Hesham El Naggar, Xiaoling Zhang, Chunyi Cui, and Xiuli Du. Experimental study of pile group bridge failure mechanism caused by liquefaction-induced lateral spreading. *International Journal of Geomechanics*, 24(3), 2024. doi: 10.1061/IJGNAI.GMENG-8353.

[18] Ioannis Kilanitis and Anastasios Sextos. Integrated seismic risk and resilience assessment of roadway networks in earthquake prone areas. *Bulletin of Earthquake Engineering*, 17(1):181–210, 2018. doi: 10.1007/s10518-018-0457-y.

[19] David Lallemant, Anne Kiremidjian, and Henry Burton. Statistical procedures for developing earthquake damage fragility curves. *Earthquake Engineering & Structural Dynamics*, 44(9):1373–1389, 2015. doi: 10.1002/eqe.2522.

[20] Szu-Yun Lin and Sherif El-Tawil. Time-dependent resilience assessment of seismic damage and restoration of interdependent lifeline systems. *Journal of Infrastructure Systems*, 26(1), 2020. doi: 10.1061/(ASCE)IS.1943-555X.0000522.

[21] John B. Mander and Nesrin Basoz. Fragility curve development for assessing the seismic vulnerability of highway bridges. Technical report, Multidisciplinary Center for Earthquake Engineering Research (MCEER), University at Buffalo, 1999. Research progress and accomplishments report; technical basis for Hazus highway-bridge fragilities.

[22] Sujith Mangalathu, Jong-Su Jeon, Jamie E. Padgett, and Reginald DesRoches. Performance-based grouping methods of bridge classes for regional seismic risk assessment: Application of ANOVA, ANCOVA, and non-parametric approaches. *Earthquake Engineering & Structural Dynamics*, 46(14):2587–2602, 2017. doi: 10.1002/eqe.2919.

[23] Sujith Mangalathu, Farahnaz Soleimani, and Jong-Su Jeon. Bridge classes for regional seismic risk assessment: Improving HAZUS models. *Engineering Structures*, 148:755–766, 2017. doi: 10.1016/j.engstruct.2017.07.019.

[24] Bryant G. Nielson and Reginald DesRoches. Seismic fragility methodology for highway bridges. In *Structures Congress 2006*, pages 1–9. American Society of Civil Engineers, 2006. doi: 10.1061/40889(201)174.

[25] Bryant G. Nielson and Reginald DesRoches. Analytical seismic fragility curves for typical bridges in the central and southeastern United States. *Earthquake Spectra*, 23(3):615–633, 2007. doi: 10.1193/1.2756815.

[26] Jamie E. Padgett and Reginald DesRoches. Retrofitted bridge fragility analysis for typical classes of multispan bridges. *Earthquake Spectra*, 25(1):117–141, 2009. doi: 10.1193/1.3049405.

[27] Karthik Ramanathan, Jamie E. Padgett, and Reginald DesRoches. Temporal evolution of seismic fragility curves for concrete box-girder bridges in California. *Engineering Structures*, 97:29–46, 2015. doi: 10.1016/j.engstruct.2015.03.069.

[28] R. Tyler Ranf, Marc O. Eberhard, and Stephen Malone. Post-earthquake prioritization of bridge inspections. *Earthquake Spectra*, 23(1):131–146, 2007. doi: 10.1193/1.2428313.

[29] Archie Rudman, John Douglas, and Enrico Tubaldi. The assessment of probabilistic seismic risk using ground-motion simulations via a Monte Carlo approach. *Natural Hazards*, 120(7): 6833–6852, 2024. doi: 10.1007/s11069-024-06497-1.

[30] Erika Schiappapietra and John Douglas. Assessment of the uncertainty in spatial-correlation models for earthquake ground motion due to station layout and derivation method. *Bulletin of Earthquake Engineering*, 19(13):5415–5438, 2021. doi: 10.1007/s10518-021-01179-w.

[31] U.S. Geological Survey. ShakeMap atlas product for the 24 august 2014 M6.0 south napa earthquake (Event nc72282711, atlas product 1624995941193). USGS Earthquake Hazards Program, Advanced National Seismic System Comprehensive Catalog, 2021. URL `https://earthquake.usgs.gov/earthquakes/eventpage/nc72282711/shakemap/atlas`. Reviewed ShakeMap Atlas ground-motion grid (grid.xml and uncertainty.xml).

[32] Sarah A. Verros, David J. Wald, C. Bruce Worden, Mike Hearne, and Mahadevan Ganesh. Computing spatial correlation of ground motion intensities for ShakeMap. *Computers & Geosciences*, 99:145–154, 2017. doi: 10.1016/j.cageo.2016.11.004.

[33] David J. Wald, Vincent Quitoriano, Thomas H. Heaton, and Hiroo Kanamori. Relationships between peak ground acceleration, peak ground velocity, and modified mercalli intensity in california. *Earthquake Spectra*, 15(3):557–564, 1999. doi: 10.1193/1.1586058.

[34] David J. Wald, Vincent Quitoriano, Thomas H. Heaton, Hiroo Kanamori, Craig W. Scrivner, and C. Bruce Worden. TriNet ShakeMaps: Rapid generation of peak ground motion and intensity maps for earthquakes in southern california. *Earthquake Spectra*, 15(3):537–555, 1999. doi: 10.1193/1.1586057.

[35] David J. Wald, Lisa A. Wald, Bruce Worden, and Jim Goltz. *ShakeMap, a Tool for Earthquake Response*. Fact Sheet 087-03. 2003. doi: 10.3133/fs08703.

[36] David J. Wald, Bruce C. Worden, Vincent Quitoriano, and Kris L. Pankow. *ShakeMap Manual: Technical Manual, User's Guide, and Software Guide*. Techniques and Methods 12-A1. 2005. doi: 10.3133/tm12a1.

[37] C. B. Worden, D. J. Wald, T. I. Allen, K. Lin, D. Garcia, and G. Cua. A revised ground-motion and intensity interpolation scheme for ShakeMap. *Bulletin of the Seismological Society of America*, 100(6):3083–3096, 2010. doi: 10.1785/0120100101.

[38] Yangyang Wu, Guangyang Hou, and Suren Chen. Post-earthquake resilience assessment and long-term restoration prioritization of transportation network. *Reliability Engineering & System Safety*, 211:107612, 2021. doi: 10.1016/j.ress.2021.107612.

[39] Weili Zhang, Naiyu Wang, and Charles Nicholson. Resilience-based post-disaster recovery strategies for road-bridge networks. *Structure and Infrastructure Engineering*, 13(11):1404–1413, 2017. doi: 10.1080/15732479.2016.1271813.