PNAS RESEARCH ARTICLE | SOCIAL SCIENCES — SUSTAINABILITY SCIENCE

Where Global Travel-Time Maps Miss Seasonal Access: A Route-Level Risk Dashboard for Rural Zambia

Ryuichi Ikeda[a,1], Hanah Habash[a], Noah Syrdal[a], Nanzheng Xie[a], and Joey Zhang[a]

This manuscript was compiled on May 15, 2026

Weiss et al. (2020) produced a widely used global map of travel time to healthcare facilities, but the model assumes stable travel conditions year-round. In rural Zambia, communities that appear accessible on a static map can become effectively isolated during the rainy season. We test whether the Weiss model's predictive relationship with DHS-reported access barriers varies across provinces and seasons, using two independent DHS waves. The key finding, replicated across both waves, is that model performance differs sharply by province. To close this gap, we train a route-level model that augments Weiss travel time with route-level precipitation, route geometry, and accessibility features, and predicts the cluster-level share of women reporting distance to a health facility as a big problem. Its per-cluster predicted value is delivered as a seasonal risk score through an interactive dashboard, shifting the dashboard output from travel time to reported access risk.

Healthcare accessibility | Seasonal flooding | Zambia | DHS | Spatial analysis

Significance

Global maps of travel time to healthcare are used to decide where clinics are needed, where outreach should go, and where health workers are under-supplied. They are built once and assume the landscape does not change with the seasons. In rainy-season geographies, this assumption hides a problem that does not show up in national averages: some communities look accessible on the map but are effectively cut off for months of the year. We show this mismatch concentrates where rivers and wetlands dominate the landscape, and deliver a dashboard that assigns each community a predicted seasonal risk score—an estimate of how many of its women will report distance to care as a serious problem—so that planners can stage supplies, community health workers, and infrastructure investment before the rainy season rather than after.

Rural Zambia, and more specifically Luapula Province, sits at the intersection of two compounding disadvantages compared to the rest of the country. Relative to other provinces, it has historically lagged in infrastructure investment and healthcare coverage, leaving residents with fewer facilities and greater travel burdens. This structural gap is further strained by Luapula's geography: the province encompasses the Bangweulu Wetlands and the Luapula River system, and receives among the highest rainfall in the country during the wet season. When the rains arrive, roads become impassable, river crossings are cut off, and communities that appear accessible on a static map become effectively isolated.

Luapula is widely known across Zambia as the land of water. The province is defined by lakes, rivers, and wetlands, a geography that proves both a blessing in times of drought and a crisis during the wet season, which typically runs from November through March. Each year, seasonal flooding inundates villages, damages bridges, and severs road connections that communities depend on for even basic services. For residents, the rains do not simply bring inconvenience; they bring isolation. Flooding is damaging on its own, but it becomes a compounding crisis when it intersects with healthcare need. Maternal healthcare in Zambia is already constrained by structural factors, including high rates of malaria, nutritional deficiencies, and the vast distances separating rural communities from the nearest clinic. Seasonal flooding intensifies these barriers, creating situations where women who have successfully navigated pregnancy face devastating outcomes simply because they cannot physically reach care when complications arise at delivery. This is not a hypothetical concern; it is a lived reality

Author affiliations: [a]School of Information, University of California, Berkeley, CA, USA

R.I. led the project, designed the routing pipeline, and built the dashboard. H.H. contributed Luapula field context and maternal health framing. N.S. designed and implemented the predictive model. N.X. contributed dashboard visualization. J.Z. implemented OSM-based routing for cross-validation. All authors contributed to writing and reviewed the manuscript.

The authors declare no competing interests.

[1]To whom correspondence should be addressed. E-mail: i.ryuichi@berkeley.edu

across Luapula and similar flood-affected provinces such as Western Province.

This project was designed with that ground-level reality in mind. Its aim is to identify and map the specific access breakdowns caused by seasonal flooding, not as an academic exercise, but to keep the stakeholders who can act on this information meaningfully informed. Ambulance drivers need to know the fastest viable routes under flood conditions. Government planning officers need to understand where targeted infrastructure investments (whether a new bridge or an elevated road) would yield the greatest health impact. NGOs and mobile health units operating where clinics are seasonally unreachable need to know precisely where and when their deployments are most urgently needed. This project aims to equip each of these actors with actionable, spatially grounded information.

This matters because the tools that healthcare planners and NGOs rely on assume stable travel conditions year-round. Our baseline analysis shows that while the Weiss model predicts perceived access barriers reasonably well in provinces like Central, its predictive accuracy in Luapula collapses during the wet season, with the correlation between modeled travel time and reported access barriers dropping to near zero. This project investigates why that breakdown occurs and how it can be addressed. By incorporating seasonal environmental variables along travel routes, including precipitation, surface water extent, and river crossings, we aim to build a more accurate predictor of when and where access to maternal healthcare is most at risk, delivered through a public-facing dashboard for planners and policymakers.

Related Work

Spatial Accessibility as a Determinant of Healthcare Access. The structural barriers limiting healthcare access in low- and middle-income countries are well established. Peters et al. (1) synthesize evidence from multiple developing-country health systems to show that rural populations face substantially higher access barriers due to longer travel distances, weaker transportation infrastructure, and seasonal environmental disruptions. Their conceptual framework integrates spatial accessibility with social and economic determinants, motivating the use of geographic accessibility metrics as predictors of healthcare access outcomes. Operationalizing this insight at scale, however, requires accessibility metrics that can be computed across countries and updated over time—a need that motivated the development of global travel-time models.

Static Global Travel-Time Models. The Weiss et al. global travel-time map (2) is one of the most widely used inputs to healthcare planning in low- and middle-income countries. It combines healthcare facility locations with friction-surface modeling and a least-cost-path algorithm, producing a 1 km global raster of travel time to the nearest hospital or clinic. The methodology directly extends an earlier travel-time-to-cities product (3) that integrated ten global-scale datasets—roads, railways, rivers, land cover, elevation, and slope—into a 1 km friction surface, validated against 53,091 settlement point pairs ($R^2 = 0.66$, MAE = 20.7 minutes). Both products assume optimal motorized movement under stable conditions; the authors themselves identify seasonal variability as a key limitation (3), prompting a second line of work that explicitly models seasonal access disruptions.

Seasonal Accessibility Models. Mroz et al. (4) developed GIS-based models incorporating dynamic floodwater variables for the Barotse Floodplain in western Zambia, applying both raster-based walking access and vector-based vehicular referral models. Their results showed that 55% of women of reproductive age had two-hour walking access to the nearest delivery site at the dry-season peak, but this dropped to 29% at flood peak. Sehi et al. (5) took a different approach in Côte d'Ivoire, using AccessMod with the assumption that all travel speeds decrease by half during the wet season; they found that average travel time increased from 1.8 hours (dry) to 3.4 hours (wet), with the population within two hours of a facility dropping from approximately 80% to 66%.

Both studies rely on a prospective physical simulation or an assumed speed reduction. Neither directly tests whether existing static models like Weiss accurately predict the access barriers that residents themselves perceive. This study takes a complementary empirical approach. We first use resident-reported access barriers from the Demographic and Health Survey to identify where the static Weiss model fails, then build a route-enhanced regression model trained directly on those reported barriers—rather than constructing a physical simulation—and deliver the resulting predictions through a dashboard.

Results

For each DHS wave we compute Spearman's ρ between cluster-level mean V467D (the share of women in a cluster reporting distance to a facility as a "big problem") and Weiss travel time at the cluster's coordinates, stratified by interview month and by province. Absolute ρ values are not directly comparable across waves: the 2015 Weiss product measures travel time to cities, the 2019 product to healthcare facilities, producing systematically different travel-time distributions. What is comparable, and what we focus on, is the *within-wave pattern* by province and season. If the same pattern holds in both waves, we treat it as robust.

The National-Level Signal. At the national level the Weiss model captures a moderate signal of perceived access in both waves (2013–14: monthly $\rho \approx 0.35$–0.59; 2018–19: $\rho \approx 0.32$–0.61, Fig. 1). The question is whether this signal degrades in wet-season months, which would be the simplest form of the "seasonal blind spot" hypothesis. It does not, consistently.

In 2013–14 the monthly ρ declined from around 0.59 in September to 0.35 in November–December, the period of peak precipitation, and partially recovered in subsequent wet months. In 2018–19 the pattern was different: ρ peaked at 0.61 in November and fell to 0.41 in December. Monthly sample sizes range from $N = 18$ to $N = 141$, so single-month values are noisy. The two waves do not tell a consistent seasonal story at the national level. We take this as evidence that a uniform nationwide "wet-season degradation" is not the right hypothesis—and begin looking for structure below the national aggregate.

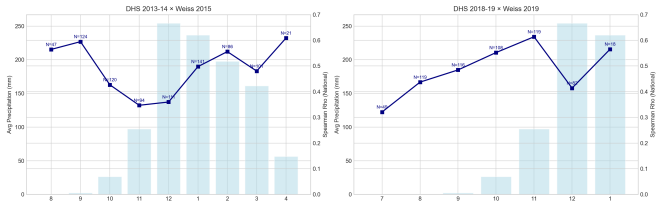

Fig. 1. National Spearman ρ between Weiss travel time and cluster-level mean V467D, by interview month, overlaid with average monthly precipitation. Left: DHS 2013–14 × Weiss 2015. Right: DHS 2018–19 × Weiss 2019. Sample size N annotated at each point.

The Provincial Signal. Stratifying by province reveals a structure that the national aggregate obscures, and the structure replicates across both waves (Fig. 2). Central Province consistently shows positive correlations across most months, in both 2013–14 and 2018–19. Luapula Province consistently shows weak or null correlations across most months, in both waves. Other provinces sit between these extremes with varying degrees of stability.

Many cells in the heatmap have small sample sizes ($N < 15$) and individual cell values should not be over-interpreted on their own; we draw our conclusions from the consistent *pattern* that holds across both waves, not from any single cell. The Central–Luapula contrast persists across two survey years, two Weiss products, and two analytically independent samples, providing stronger evidence than either wave alone.

Central and Luapula. The provincial signal is clearest in the Central–Luapula contrast (Fig. 3). Central Province wet-season correlations are strong and statistically significant in both waves: $\rho = 0.77$, $p < 0.001$, $N = 44$ (2013–14) and $\rho = 0.67$, $p < 0.001$, $N = 25$ (2018–19). Dry-season correlations in Central are positive but weaker: $\rho = 0.39$, $p = 0.074$, $N = 22$ (2013–14) and $\rho = 0.45$, $p = 0.018$, $N = 27$ (2018–19). In Central, Weiss travel time tracks perceived access across seasons, with the signal strongest in the wet season when the access gradient is largest.

Luapula's pattern is different. The 2013–14 wet-season correlation is flat: $\rho = 0.00$, $p = 0.976$, $N = 53$ (Fig. 3, top right). Weiss travel time carries essentially no information about perceived access in Luapula during the wet season of 2013–14. In 2018–19 a positive trend reappears ($\rho = 0.40$, $N = 19$) but does not reach significance ($p = 0.092$), and sample size is smaller. Luapula's dry-season correlations are also weak ($\rho = 0.24$, $N = 22$ in 2013–14; $\rho = 0.19$, $p = 0.276$, $N = 34$ in 2018–19), suggesting the Luapula result is not purely seasonal but reflects a broader mismatch between static travel time and local access conditions.

Weiss travel time is a *necessary but insufficient* predictor of perceived access in rural Zambia. It captures real signal, and in provinces where terrain conditions are relatively stable that signal is enough on its own. Where rivers, wetlands, and heavy seasonal rainfall dominate the landscape, it is not. Closing the gap requires features that describe the route, not just its endpoints—precipitation along the path, route geometry, and broader accessibility context. These features feed a route-enhanced regression model that predicts cluster-level V467D directly. The model's per-cluster prediction—bounded in $[0, 1]$—is the seasonal risk score the dashboard is built around.

Route-Enhanced Model. We next test whether route-level context improves prediction of perceived access barriers beyond Weiss travel time alone. Using the 2018–19 DHS wave (535 clusters), we train a model that predicts the cluster-level share of women reporting distance to a health facility as a big problem.

The model augments Weiss accessibility with route distance, route shape, CHIRPS rainfall summaries, and season, using features extracted from Google road routes and Weiss least-cost paths to the rank-1 matched facility. OSM routes are retained as a Luapula-specific visual cross-check and dashboard layer. Details of route recovery and feature extraction are described in Materials and Methods.

The best-performing model under the strictest evaluation scheme is the weighted ensemble. Under leave-one-province-out validation—training on nine provinces and testing on the tenth—it achieves $R^2 = 0.293$, MAE $= 0.191$, AUC $= 0.695$, and cluster-level Spearman $\rho = 0.543$ between observed and predicted V467D. Under the easier cluster-grouped 5-fold setting, the same model reaches $R^2 = 0.357$ and AUC $= 0.730$. For comparison, the cluster-level Spearman correlations between Weiss travel time and observed V467D shown in Fig. 2 reach only $\rho = 0.00$–0.40 in Luapula's wet-season cells. The route-enhanced model's province-holdout Spearman $\rho = 0.543$, a stricter generalization test, suggests that route-level features carry information about perceived access beyond what Weiss travel time captures alone.

These results suggest that the model is most useful for ranking communities by relative risk, but not for precise forecasting of each cluster's barrier rate. Figure 4 shows this directly: predicted and observed cluster-level V467D move together overall, but with substantial dispersion. V467D is a self-reported access measure shaped not only by travel friction, rainfall, and route geometry, but also by household resources, transport availability, and facility choice—only partly captured in the predictors. We therefore interpret the model's per-cluster predicted V467D as a seasonal accessibility risk score for prioritization rather than as a causal estimate of unmet access.

Dashboard. The dashboard presents the route-enhanced model's output for people who plan healthcare delivery, such as District Health Officers and NGO program managers working on maternal health in Luapula. Its core is a per-cluster *predicted seasonal risk score*—the route-enhanced model's prediction of cluster-level V467D, bounded in $[0, 1]$ and interpretable as the expected share of women in the cluster who will report distance as a big problem in a given season—displayed on an interactive map (Fig. 5). The dashboard is a browser-based Folium application delivered as a standalone HTML file that requires no server, hosted as a web-based interactive dashboard.

Clusters are color-coded by predicted risk score on a sequential yellow-to-red scale calibrated to the observed range in our predictions. A season toggle (top right) flips every cluster's color between the dry-season and wet-season prediction, which is the dashboard's primary lens for asking "where does access deteriorate the most when the rains arrive?" A complementary panel surfaces the largest seasonal gaps directly. For the focused province, it lists the top five clusters by $|\text{wet} - \text{dry}|$ and zooms to any

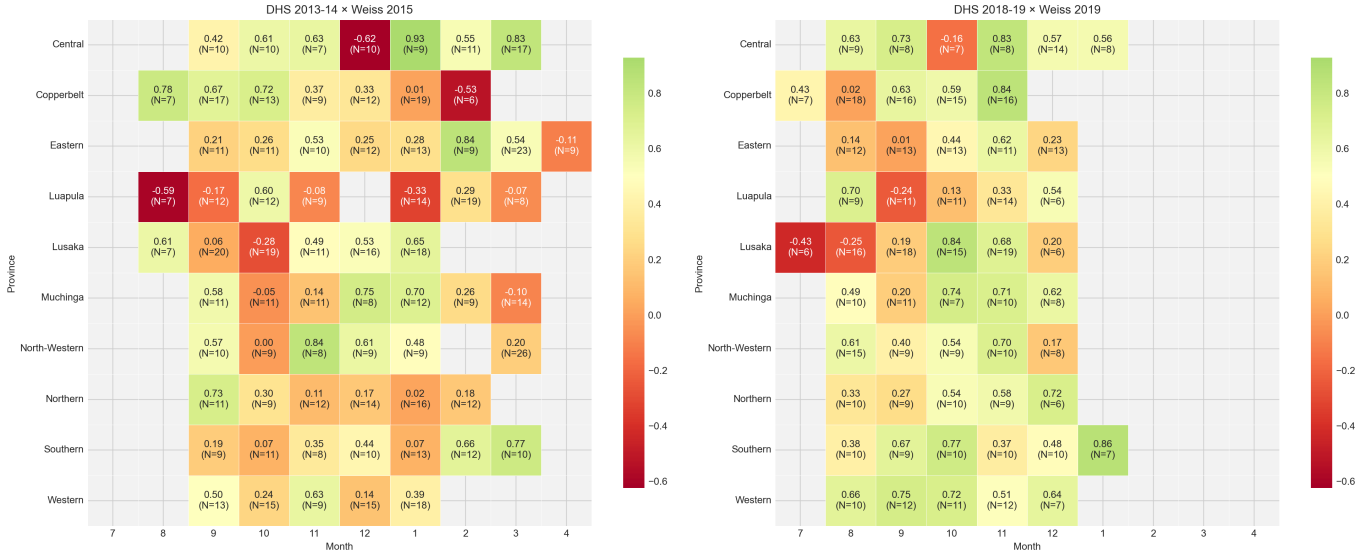

Fig. 2. Spearman ρ between Weiss travel time and cluster-level mean V467D, by province and interview month. Left: DHS 2013–14 $\times$ Weiss 2015. Right: DHS 2018–19 $\times$ Weiss 2019. Cell annotations report ρ and N (clusters). Luapula Province shows persistently low ρ; Central Province remains consistently high across both waves.

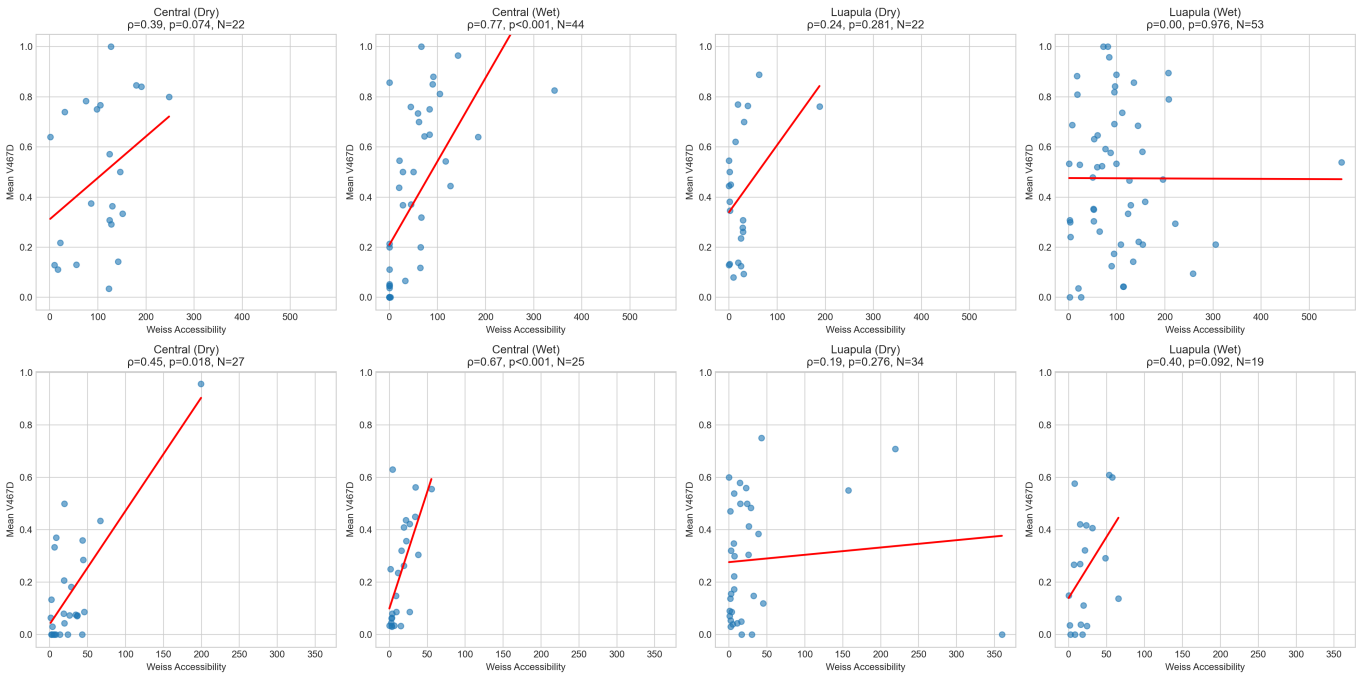

Fig. 3. Scatter plots of cluster-level mean V467D against Weiss travel time, by province and season. Top row: DHS 2013–14 $\times$ Weiss 2015. Bottom row: DHS 2018–19 $\times$ Weiss 2019. Red lines show linear fits for visualization; reported ρ values are Spearman. Central wet-season relationships are strong and significant; Luapula wet-season 2013–14 is flat.

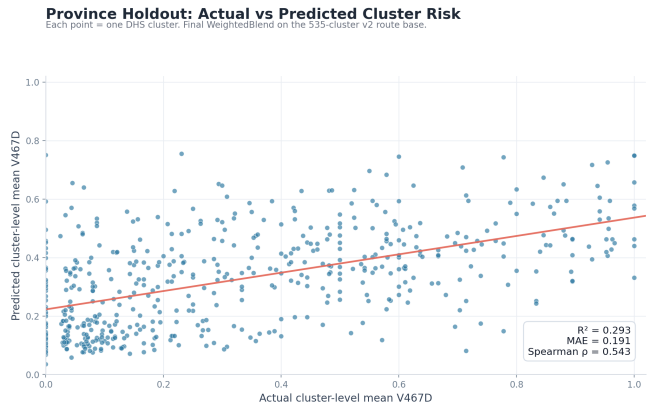

Fig. 4. Province-holdout performance of the final WeightedBlend model. Each point is one DHS cluster; the x-axis is observed cluster-level mean V467D and the y-axis is the corresponding predicted value. The fitted trend line highlights the model's positive ranking signal, while the spread around the line shows substantial residual error.

cluster on click, so a planner can locate the most seasonally volatile communities without scanning the map.

Clicking any cluster opens a popup with the predicted scores for both seasons, the underlying Weiss travel time, and the cluster's three nearest facilities. The same click highlights the routes connecting the cluster to those three facilities: the Google road route, the Weiss least-cost-path route, and (in Luapula) the OSM route, layered together so the planner can see whether the formal road network and the friction-grid optimum disagree about how the journey is made. A risk-distribution histogram in the lower-left summarizes the distribution of predicted scores in the focused province; clicking a histogram bar zooms to the clusters in that bin.

To support extension beyond the 535 DHS clusters, the dashboard exposes a "Custom Point Prediction" control that accepts user-supplied coordinates, runs the same routing pipeline (three nearest facilities from the MFL, Google + Weiss routes), and returns the model's predicted dry- and wet-season scores.

Supporting layers include the full Master Facility List of health facilities, OpenStreetMap and satellite basemaps, and the route layers themselves. They are independently toggleable. The default basemap is the standard OpenStreetMap tiles; satellite imagery is available for visual sanity-checking against terrain features.

The dashboard connects three groups of actors (Fig. 6): supporting stakeholders who provide the data and institutional context behind the risk score (Ministry of Health, public geospatial data providers, local authorities, and the Disaster Management and Mitigation Unit); decision makers and users who consult the dashboard directly; and frontline health workers and rural communities in Luapula, whose access outcomes the tool is ultimately designed to improve. The Use Cases that follow in the Discussion describe how the five decision-making roles translate the same per-cluster risk score into different operational decisions.

Discussion

Below we describe how planners use the seasonal risk score, what the analysis cannot tell us, and what evidence would be needed to take the same approach elsewhere.

Use Cases. Throughout this subsection, "predicted risk score" refers to the route-enhanced model's predicted cluster-level V467D: a number between 0 and 1 that estimates the share of women in a cluster who will report distance to a facility as a "big problem" in a given season. All five use cases below share a common entry point: the planner opens the dashboard, selects a cluster of interest, and reads its predicted risk score. What the planner does next differs according to the planner's goal. Different supporting information from the dashboard turns the scores into different decision-making support tools.

Dry-season supply staging. District health offices finalize supply plans before the rainy season begins in November. The planner sorts clusters by predicted wet-season risk score and pre-positions essential supplies (antimalarials, oral rehydration salts, maternal health kits, etc.) in the highest-risk communities before roads deteriorate and access breaks down. Luapula's combination of high rainfall, dense river networks, and limited road infrastructure means that windows for resupply can close quickly and without warning once the rains arrive. This is the use case closest to the dashboard's core: the risk score alone drives the decision, and acting before the season begins is the difference between communities that have what they need and those that go without.

Community Health Worker (CHW) deployment. NGOs allocating limited CHW capacity need to find communities where facility-based care fails seasonally even though a facility is nearby. These are clusters with a high predicted risk score but short Weiss travel time: the score indicates that access breaks down during the wet season, yet the facility is not geographically distant under stable conditions. In Luapula, this combination is precisely the signature of flooding, a community that can reach a clinic in forty minutes in July may face a washed-out road or a submerged crossing in January, making the same journey impossible. The dashboard surfaces that combination directly: a cluster click shows both numbers side by side, making the CHW-deployment candidates easy to spot and allowing NGOs to concentrate their limited capacity where the seasonal gap is largest.

Infrastructure prioritization. Road and bridge investment budgets are limited and must be targeted. Starting from the highest-risk clusters, the planner opens the GeoJSON path layer and reads the per-route features in the popup to see which physical elements drove the risk: which paths cross rivers, which traverse flood-prone stretches, or which accumulate high precipitation. These are the segments where a bridge or a raised road would translate most directly into year-round access. Luapula's geography makes this targeting especially consequential, a single river crossing can render an entire community unreachable for months, as the 2020 bridge destruction in Samfya District illustrates. Infrastructure investment guided by seasonal risk scores rather than aggregate travel time alone is investment directed at the disruptions that actually cut communities off from care.

Ambulance and emergency vehicle routing. Ambulance drivers and dispatchers responding to medical emergencies face a problem that is distinct from the planning problems above: they need the right answer now, not next month. Which route to the nearest facility is passable under current conditions? What is the viable alternative if the primary route is flooded? Every minute spent

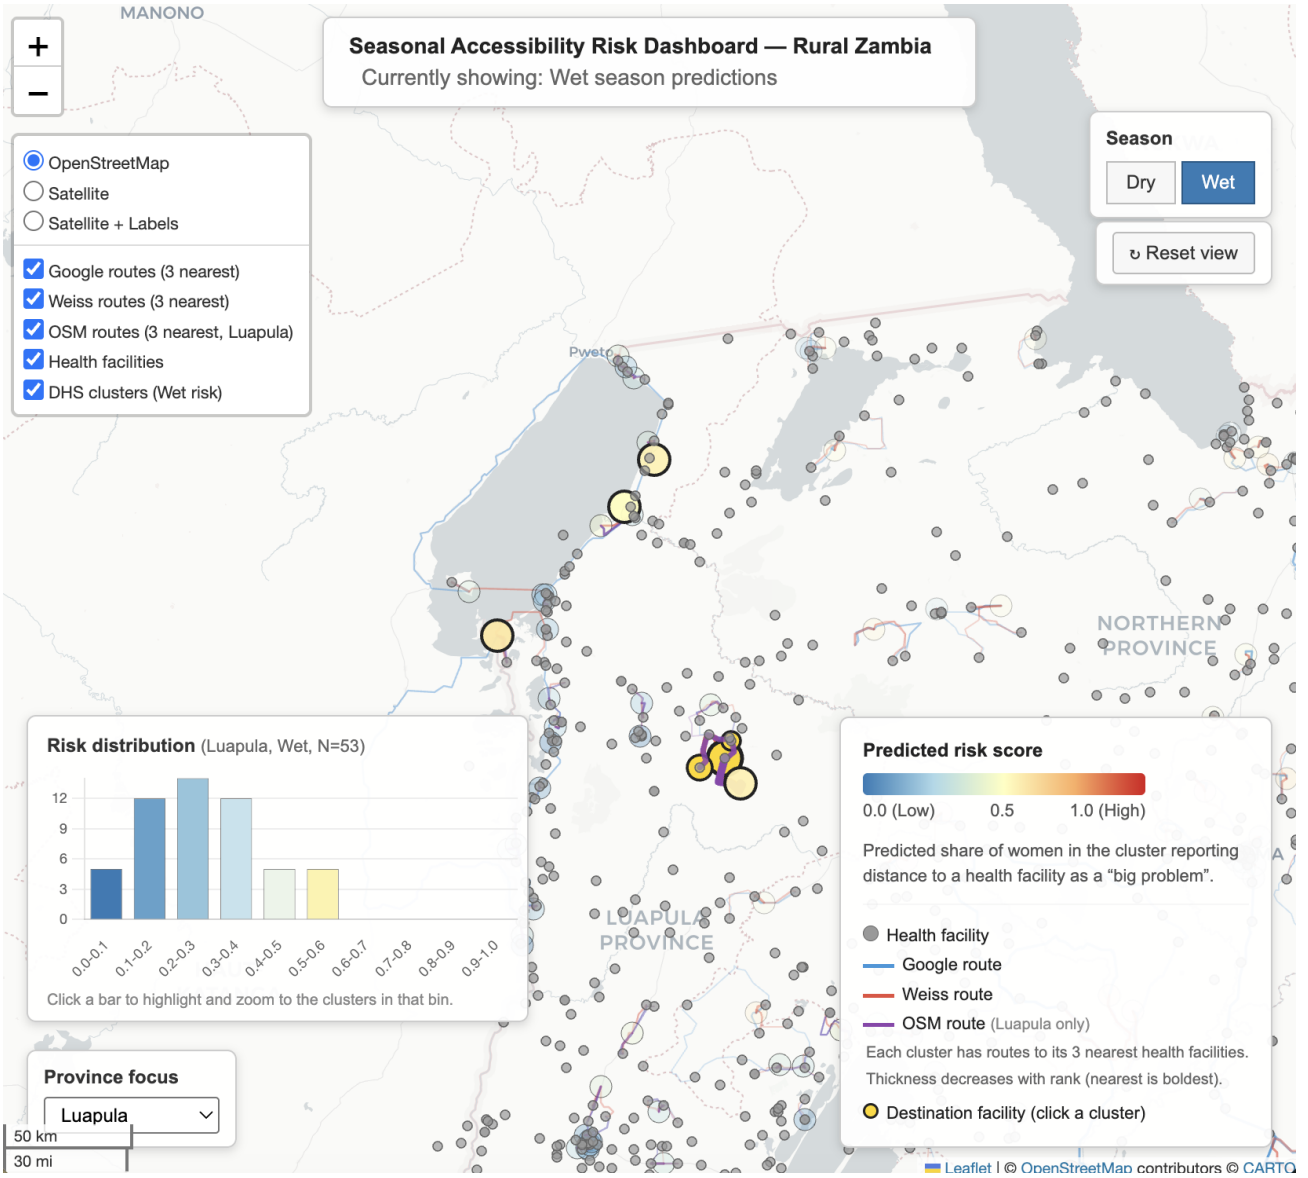

Fig. 5. Interactive dashboard for Luapula Province. Clusters are color-coded by predicted seasonal risk score from the route-enhanced model—the predicted share of women in the cluster who will report distance as a big problem. Cluster popups surface the score alongside the Weiss travel time and route-level features that produced it; a GeoJSON path layer shows the underlying cluster-to-facility route on demand. A season toggle flips between dry- and wet-season predictions, while a side panel lists the largest seasonal gaps in the focused province.

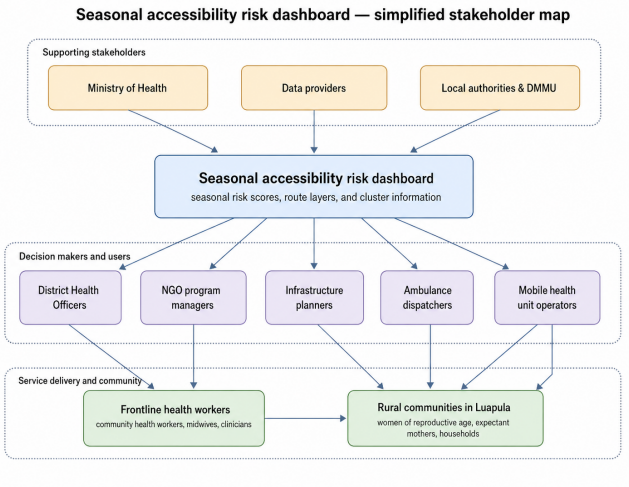

Fig. 6. Stakeholder map for the seasonal accessibility risk dashboard.

deliberating at a washed-out crossing is a minute of care delayed for a woman in labour or haemorrhage. The dashboard addresses this not by providing real-time conditions but by reducing the decision burden before emergencies occur. A driver or dispatcher familiar with the dashboard can consult it at the start of the rainy season to pre-identify primary and alternative routes for each community they serve. For any cluster, the route layer surfaces the three candidate paths to nearby facilities alongside the per-route features (river crossings, flood-prone stretches, precipitation accumulation) that determine seasonal passability. The Google road route, the Weiss least-cost-path route, and the OSM route are layered together, allowing drivers to see where the formal road network and the friction-grid optimum diverge, which is itself a signal of where seasonal conditions are most likely to produce unexpected blockages. This use case differs from the others in its user and its time horizon. The primary user is operational rather than strategic, and the goal is not to allocate resources across communities but to move a single patient safely under conditions that may not match what the map assumed. Using the dashboard before the rainy season could help drivers identify likely problem routes in advance, reducing the need to make routing decisions during an emergency.

Mobile health unit scheduling. Mobile health units operating in Luapula serve a population that facility-based care cannot reliably reach during the wet season. But mobile units face the same seasonal access constraints as the women they serve: the roads that become impassable for a patient travelling to a clinic are often the same roads that a mobile unit would need to reach her. Deploying mobile capacity effectively therefore requires knowing not only which communities have the greatest unmet need, but which communities remain reachable for a vehicle during the months when facility access collapses. The dashboard's season toggle makes this targeting direct. A program manager can identify clusters whose predicted risk score increases sharply from the dry season to the wet season; communities that are manageable in July but effectively cut off by January, and schedule mobile unit visits to coincide with those high-risk windows, while the routes remain open enough to allow access. The complementary panel listing the top clusters by seasonal gap provides a ranked starting point for this scheduling exercise. Unlike CHW deployment, which targets communities where the facility is nearby but access breaks down, mobile unit scheduling must also account for the unit's own route: the route layer and its per-segment features allow planners to identify not only where to go but whether getting there is feasible during the months that matter most.

Data Quality Lessons. Building the tool required integrating four public geospatial datasets, which surfaced issues generic enough to be worth naming.

Coordinate integrity. The Zambia Master Facility List contained records with swapped lat/lon, missing decimal points, and coordinates outside the country. We cleaned these manually, retaining 97% of records. We recommend bounding-box sanity checks before use.

Facility typologies. "Hospital," "clinic," and "health post" are classified differently across sources. Our analysis avoids this problem because V467D does not specify facility type, but any type-filtered analysis must confront it.

Temporal alignment. Weiss 2015 and 2019 use different destination sets (cities vs. healthcare facilities). Treating them interchangeably would obscure replication claims. The rule is to compare patterns across waves, not magnitudes.

Routing-source divergence. The three routing methods we tested (Google Routes v2, Weiss LCP, OSM) produce visibly different geometries in flood-prone Luapula. This divergence is itself useful—it flags where the formal road network and a friction-grid optimum disagree—but consumers of any single source should know that a "route" is a function of the source's assumptions, not an observed fact about how people travel.

Ethical Analysis. *Community Dignity and Representation.* Any project centered on vulnerable populations carries an obligation to represent those communities with accuracy and respect. Luapula's residents are not simply data points or passive subjects of analysis; they are people navigating a genuinely difficult structural reality. Our data draws on DHS survey responses that represent real hardship and, in the context of maternal complications, potentially life-and-death experiences. This project has a responsibility to frame those experiences honestly: not to sensationalize suffering, and not to flatten the complexity of the problem into a single narrative. The framing of Luapula as a "failure case" for the Weiss model, for example, is analytically precise, but care must be taken throughout the project's communication materials to ensure that the province is not cast more broadly as simply deficient, obscuring the structural inequities in infrastructure investment that drive these outcomes.

Data Privacy and Consent. The DHS data used in this project was collected under established informed consent protocols, and cluster-level GPS coordinates have been randomly displaced by the DHS program (up to 5 km in rural areas) specifically to prevent re-identification of individual respondents. Our analysis operates exclusively at the cluster level and does not attempt to identify or profile specific individuals. All spatial datasets used, including CHIRPS precipitation, HydroRIVERS, and the

Zambia health facilities dataset, are publicly available and do not contain personally identifiable information. The public-facing dashboard will similarly present aggregated risk scores at the cluster level rather than individual data, and will be designed to minimize any potential for misuse or stigmatization of specific communities. *Gender and Health Equity.* At its core, this project is about a gendered health equity problem. The burdens of inaccessible maternal healthcare fall entirely on women, and they fall most heavily on women who are poor, rural, and geographically marginalized. Seasonal flooding compounds an already unjust distribution of healthcare risk. By making this disparity visible and spatially legible, our project aims to contribute to the broader effort to close that gap, not merely to describe it.

Limitations. *V467D measures perceived, not actual, access.* The outcome captures whether a woman reports distance as a "big problem," not whether she reached care. We read correlations with Weiss travel time as evidence about model-to-perception alignment, not ground-truth accuracy.

DHS spatial jitter. Cluster coordinates are displaced up to 5 km (rural) and 2 km (urban). Weiss et al. used the same jittered coordinates in their extended analysis, establishing precedent for cluster-level work. Nonetheless, jitter bounds the precision of route-level features and is one reason we route to the three nearest facilities rather than only the single nearest.

Small cell sizes. Many province–month cells in Fig. 2 have $N < 15$. Our claims rest on the consistent pattern across both waves, particularly Central and Luapula where key wet-season cells have adequate N (44 and 53 in 2013–14).

V467D seasonal differences may reflect sampling bias. Mean V467D is higher in wet-season months (2013–14: $0.36 \rightarrow 0.43$; 2018–19: $0.32 \rightarrow 0.38$), but wet- and dry-season respondents are not random subsamples of each other. We do not claim this demonstrates seasonal worsening of actual access.

Routing assumes drivable roads. All three routing methods we tested are biased toward formal road networks. In rural Luapula, foot paths, and informal tracks carry substantial real-world traffic but do not appear on any of our base layers. The route geometries the model sees are therefore a partial approximation of the journeys women actually make. A related gap is behavioral: the model has no visibility into which facility a woman chooses to travel to, or whether she travels at all.

Taken together, these gaps (unmapped roads, unobserved behavior, and jittered coordinates) share a common implication: ground-level validation is necessary before the dashboard can be recommended for operational use, and cannot be replaced by data improvements alone.

Future Work. Four directions extend naturally from this work. First, the cluster-level diagnostic approach is portable to any country where DHS provides geolocated clusters with a V467D-equivalent variable. Replicating the analysis in two or three hydrologically comparable settings (for example, Bangladesh's delta region or Mozambique's Zambezi floodplain) would test whether the provincial heterogeneity pattern we observe in Zambia generalizes. Second, the risk layer currently relies on historical precipitation from CHIRPS. Integrating near-real-time or forecast precipitation data would allow planners to anticipate access disruptions days or weeks ahead rather than relying on seasonal averages. Third, the use cases described in the Discussion remain untested with actual end users. Field validation with District Health Officers and NGO program managers in Luapula, observing whether and how the tool changes planning decisions, is necessary before the dashboard can be recommended for operational use. Ground-level knowledge of which roads are actually traveled, and which facilities communities actually use, cannot be recovered from the data alone and would directly inform both the routing assumptions and the behavioral gaps identified in Limitations. Fourth, the routing pipeline is currently limited to the formal road network. Higher-resolution satellite image segmentation could improve detection of informal tracks and dirt roads in rural settings. A segmentation model trained on labeled multi-temporal satellite imagery could learn the visual signature of dirt tracks in rural landscapes and flag road-like features absent from OSM or Google's network.

Materials and Methods

Data sources. Our primary outcome signal is from two waves of the Zambia Demographic and Health Survey: the 2013–14 Women's Recode (ZMIR61FL, $N = 16{,}411$ women across 721 clusters) and the 2018–19 Women's Recode (ZMIR71FL, $N = 13{,}421$ women across 545 clusters; 535 retained after route recovery). We use cluster-level geographic coordinates, the self-reported access barrier variable V467D ("distance to a health facility: big problem"), and the interview month variable V006. DHS applies random displacement of up to 5 km for rural clusters and 2 km for urban clusters to the coordinates published to researchers.

Our primary accessibility signal is from two Weiss rasters paired temporally with the DHS waves: the 2015 *travel-time-to-cities* product (3) with the 2013–14 DHS, and the 2019 *travel-time-to-healthcare* product (2) with the 2018–19 DHS. Both rasters are sampled at the (jittered) cluster coordinates and enter the analysis as minutes of travel time.

Contextual data used in the route-level pipeline and dashboard include CHIRPS monthly precipitation at approximately 5 km resolution from the UC Santa Barbara Climate Hazards Center; HydroRIVERS river centerlines from the HydroSHEDS framework; and ...health facility locations from Zambia's Master Facility List maintained by the Ministry of Health, which catalogs the country's three-tier facility structure of health posts, district hospitals, and central hospitals (6). Facility coordinates were sanity-checked for bounding-box and decimal-place integrity before use, as described in Discussion.

Outcome construction. V467D is a binary individual-level variable (1 = "big problem," 0 = otherwise). We aggregate to the cluster level as the arithmetic mean across all women interviewed in a cluster, producing a continuous outcome bounded by $[0, 1]$ that represents the share of women in the cluster reporting distance as a big problem. This aggregated outcome is our response variable throughout.

Seasonality definition. We classify each DHS observation as dry-season or wet-season using V006 (interview month), with dry = July–October and wet = November–April. This follows conventional Zambian seasonal definitions and aligns with CHIRPS monthly precipitation. The 2013–14 wave spans August 2013 to April 2014; the 2018–19 wave spans July 2018 to January 2019. Both waves include both dry and wet interview months, enabling within-wave seasonal comparison.

Correlation analysis. The primary analysis computes Spearman's rank correlation coefficient ρ between cluster-level mean V467D and Weiss travel time at the cluster coordinate. Correlations are computed separately for each DHS wave and stratified by interview month and by province. Reported p-values are two-sided; we do not apply a multiple-comparison correction because our claims rest on patterns that replicate across waves rather than on any single cell. Province–month cells with fewer than a handful of clusters are shown in figures for completeness but are not used individually in the argument.

Route recovery and feature extraction. For each of the 535 DHS clusters in the 2018–19 wave, we begin by constructing a candidate set of health facilities from Zambia's Master Facility List (MFL). Specifically, we identify the three nearest candidate facilities by haversine distance from the DHS cluster coordinate. This candidate step addresses two practical issues: DHS cluster coordinates are randomly displaced by up to 5 km for rural clusters, so the truly nearest facility in any underlying geometry may differ from the truly nearest under the jittered coordinates we observe; and V467D asks about the perceived distance barrier overall rather than about a specific facility, so a route-level signal aggregated over the three nearest captures more of women's real choice set than a single best route. The MFL contained records with swapped lat/lon, missing decimal points, and coordinates outside the country, which were corrected manually; 97% of records were retained after cleaning.

We compute three route representations from each cluster to each candidate facility. The first is a Google Routes v2 query (7) with `routingPreference=TRAFFIC_AWARE` and `travelMode=DRIVE`, returning a polyline along the formal road network with distance and duration estimates. The second is a least-cost path (LCP) computed on the Weiss 2019 friction surface (2) using `skimage.graph.route_through_array` on the same one-kilometer raster the Weiss travel-time map is built on; this is the route the Weiss model would itself favor between two points. The third, used as a visual cross-check in Luapula where commercial road data may be sparse, is an OpenStreetMap-based route on a NetworkX graph: each Linestring is decomposed into coordinate vertices, every consecutive pair forms a weighted edge with travel time estimated from OSM functional class (`fclass`), and the route is recovered via Dijkstra's algorithm after snapping cluster and facility points to the nearest graph node. For the 535 clusters this produces 1,605 Google routes and 1,605 Weiss LCP routes; OSM contributes 148 routes restricted to the 53 Luapula clusters.

Feature extraction is performed along the Google and Weiss routes to the rank-1 matched facility. From route geometry we compute route length, mean absolute turn angle, and turn-density comparisons between the Google and Weiss paths. From CHIRPS monthly precipitation rasters we sample along-route rainfall and summarize it using maximum, 90th-percentile, minimum, and high-rainfall-share statistics, with separate summaries for the Google and Weiss representations. These route-level summaries are combined with cluster-level DHS spatial attributes (coordinates, rural/urban status, and Weiss accessibility) and seasonal information derived from interview month. We also tested richer route representations, including 5 km route buffers and segment-weighted exposure summaries, but these did not improve held-out performance enough to justify the added complexity. The result is a route-feature table with one row per modeled DHS cluster. Although the dashboard can visualize all three candidate facilities and their associated routes, the final predictive model uses only the rank-1 facility for feature extraction and prediction, which keeps the feature set interpretable and avoids introducing an additional aggregation assumption over multiple facilities.

Predictive modeling framework. The predictive task uses the 2018–19 DHS wave and treats V467D as a binary respondent-level outcome (1 = the respondent reports distance to a health facility as a "big problem," 0 = otherwise). Route and environmental predictors are measured at the DHS-cluster level, so every respondent in the same cluster receives the same route-derived feature vector. In that sense, the model is trained on respondent-level labels but functions as a cluster-risk model.

The final feature set is intentionally compact. Categorical predictors are `urban/rural` and a wet/dry `season` label derived from interview month (V006). Numeric predictors are cluster latitude, matched-facility latitude and longitude, the cluster's Weiss accessibility value, Google-route distance, a route-duration-to-Weiss ratio, Weiss route length, Google-route mean absolute turn angle, a turn-density gap between Google and Weiss routes, CHIRPS rainfall summaries sampled along both Google and Weiss routes, a cyclical month term, and a wet-season rainfall interaction. This reduced set was chosen after testing richer route-buffer and segment-weighted summaries that did not improve held-out generalization.

All preprocessing is implemented in a single scikit-learn pipeline. Numeric variables are median-imputed and standardized; categorical variables are imputed with the most frequent level and one-hot encoded. We compare four candidate models: logistic regression, Random Forest, ExtraTrees, and a weighted ensemble (`WeightedBlend`) combining logistic regression, Random Forest, and ExtraTrees. Ensemble weights are tuned by inner grouped cross-validation over a coarse simplex grid, using cluster-level R^2 as the primary objective and MAE as a tiebreaker.

We evaluate the models under two grouped validation schemes. The first is cluster-grouped 5-fold cross-validation, which measures within-dataset generalization while preventing the same DHS cluster from appearing in both train and test folds. The second, and primary, evaluation is leave-one-province-out validation, a stricter spatial generalization test in which the model is trained on nine provinces and tested on an unseen tenth. Province is used only to define holdout folds and is not included as a predictor in the final model.

For each held-out respondent, the model outputs a probability that V467D = 1. These probabilities are then averaged within each DHS cluster to produce the final cluster-level prediction:

$$\widehat{\mathrm{V467D}}_c = \frac{1}{n_c} \sum_{i \in c} \hat{p}_i,$$

where $\hat{p}_i$ is the predicted respondent-level probability and n_c is the number of respondents in cluster c. This cluster-level mean prediction, bounded in $[0, 1]$, is the seasonal risk score displayed in the dashboard. Performance is summarized using cluster-level R^2 and MAE, respondent-level AUC, and cluster-level Spearman rank correlation.

On the final 535-cluster dataset, the selected WeightedBlend model achieves $R^2 = 0.357$ and AUC $= 0.730$ under cluster-grouped 5-fold cross-validation, and $R^2 = 0.293$, MAE $= 0.191$, AUC $= 0.695$, and Spearman $\rho = 0.543$ under province holdout. After evaluation, the fitted pipeline is used to generate dry- and wet-season cluster predictions for the dashboard by updating the seasonal and rainfall-related inputs while holding all other cluster attributes fixed.

Model selection emphasizes interpretability alongside predictive performance, since the end users of the tool are planners rather than machine-learning specialists. All analysis code, including the reproducible modeling notebook, is maintained in the project repository.

1. DH Peters, et al., Poverty and access to health care in developing countries. *Annals New York Acad. Sci.* **1136**, 161–171 (2008).
2. DJ Weiss, et al., Global maps of travel time to healthcare facilities. *Nat. Medicine* **26**, 1835–1838 (2020).
3. DJ Weiss, et al., A global map of travel time to cities to assess inequalities in accessibility in 2015. *Nature* **553**, 333–336 (2018).
4. EJ Mroz, et al., Impacts of seasonal flooding on geographical access to maternal healthcare in the Barotse Floodplain, Zambia. *Int. J. Heal. Geogr.* **22** (2023).
5. GT Sehi, CA Houngbedji, DM Parker, PM Macharia, Geographic accessibility to public healthcare facilities and spatial clustering during the wet and dry seasons in Côte d'Ivoire. *medRxiv* (2023) Preprint.
6. Zambia Development Agency, Health sector profile — Zambia, (Zambia Development Agency), Technical report (2013).
7. Google, Routes API documentation (https://developers.google.com/maps/documentation/routes) (2024) Accessed 1 May 2026.