Int. j. adv. multidisc. res. stud. 2026; 6(5):917-925

Received: 06-08-2026
Accepted: 16-09-2026

International Journal of Advanced Multidisciplinary Research and Studies

ISSN: 2583-049X

Spatiotemporal Optimization of Mobile Health-Clinic Deployment Under Seasonal Flooding

[1] **Philipus Nghifikepunye Nangolo,** [2] **Kaarina Nakale**
[1] Department of Mathematics and Economics, International University of Management, Nkurenkuru Campus, Namibia
[2] Department of Information and Communications Technology, International University of Management, Nkurenkuru Campus, Namibia

Corresponding Author: **Philipus Nghifikepunye Nangolo**

Abstract
Seasonal flooding can fragment rural road networks precisely when dispersed communities require dependable primary care. This study develops a two-stage stochastic robust location-routing model for stationing and routing mobile health clinics under uncertain road closures, travel times and service demand. Candidate bases are selected before flood conditions are known; community assignments and vehicle routes adapt after a scenario is observed. Scenario-specific shortest paths are computed on a flood-modified road graph, while a population-weighted spatial accessibility index, unmet-demand penalties, an equity term and conditional value-at-risk control the efficiency–resilience trade-off. A reproducible synthetic case representing a linear rural settlement system contains 24 communities, six candidate bases, 76 road links, three clinics and 12 seasonal flood scenarios. The robust policy selected bases B2, B4 and B5. Relative to stationing optimized for dry conditions, it increased the probability-weighted objective by 3.3% but reduced the worst-scenario objective by 9.0% and the mean of the three most adverse outcomes by 2.8%. All 900 expected weekly consultations remained allocable in the tested scenarios, although the population within 45 minutes of a selected base fell to 89.1% during central, prolonged and bridge-failure scenarios. The framework converts flood forecasts and road-status observations into stationing, routing and accessibility outputs that health managers can update before and during extreme weather. Results are illustrative rather than estimates of any named district and require field calibration before operational use.

Keywords: Mobile Health Clinics, Location-Routing, Seasonal Flooding, Robust Optimization, Spatial Accessibility, Rural Health

Introduction

Geographical access is a basic condition for effective primary health care. In sparsely populated rural regions, a clinic may be administratively present while remaining practically inaccessible because travel is long, transport is costly, or roads become impassable during the rainy season. Network-based analysis is therefore more informative than straight-line distance because it represents the routes that patients and service vehicles can actually use [1-3]. The problem becomes more severe when flooding affects bridges, culverts and unpaved feeder roads. Climate-resilient health-system guidance accordingly treats transport continuity and service access as part of preparedness [22, 23]. A road network that is connected under ordinary conditions can separate into weakly connected components during an extreme event, and a station that appears central in a dry-season map can become operationally peripheral.

Mobile clinics partly overcome the fixed nature of rural health infrastructure. They can carry staff, medicines, diagnostic equipment, vaccination supplies and referral coordination to communities on a rotating schedule. Their flexibility does not eliminate the planning problem. Managers must decide where vehicles should be stationed before a flood develops, which communities each unit should serve, in what order those communities should be visited, and how routes should change when a road is closed. These decisions couple strategic facility location with operational vehicle routing. The resulting location-routing problem is substantially harder than either problem considered separately [4-7].

Classical location models select facilities to minimize weighted distance or maximize covered demand [4, 5, 25]. Health-care applications have extended these ideas through capacity, service hierarchy and preventive-care objectives [8]. Spatial-accessibility research has also developed gravity-type and floating-catchment measures that distinguish nominal proximity from effective access [2, 3]. However, a static accessibility surface assumes that the travel network and service capacity remain unchanged. That assumption is unsuitable when seasonal floods alter edge availability and travel speed from one week to another.

Vehicle-routing models determine tours that satisfy vehicle capacity, travel-time or time-window restrictions [6, 7, 21]. Stochastic and robust extensions protect routes against travel-time or demand uncertainty [9-12, 20]. Humanitarian logistics models similarly consider pre-positioning, network disruption and equitable distribution [13-17]. Yet three separations remain common. First, location and routing are frequently optimized in different planning stages. Second, flood exposure is often represented only by an average travel-time multiplier rather than by explicit edge closures and scenario-specific shortest paths. Third, performance is commonly summarized by total distance or cost, without a population-weighted measure of spatial access or an explicit penalty for the upper tail of adverse scenarios.

This study addresses those gaps through a spatiotemporal decision-support framework. The principal objective is to determine robust staging sites for a limited fleet of mobile health clinics and adaptive weekly routes that remain useful when floods disrupt road access. The specific objectives are to: (i) represent the rural transport system as a time- and scenario-dependent graph; (ii) integrate candidate-base selection, community assignment and routing in a two-stage stochastic robust model; (iii) embed population-weighted accessibility and equity within the objective; (iv) solve the resulting model by shortest-path preprocessing and scenario decomposition; and (v) demonstrate how the framework reports operational trade-offs through a transparent synthetic case.

The contribution is methodological and practical. Methodologically, the model combines network disruption, robust location-routing, a gravity-type accessibility index and conditional value-at-risk in one formulation. Practically, it produces a compact set of outputs, selected bases, feasible routes, communities at risk of delayed service and scenario-specific performance, that can be linked to a geographic information system. The case study does not claim to measure health access in a named locality. It is a reproducible stress test showing what data are required, how uncertainty enters the model, and how a health authority could update the plan when field data become available.

Materials and Methods

Study design and decision setting

The study is an applied mathematical modelling investigation with a synthetic computational experiment. No patients, clinical records or personal data were used. A planning horizon is divided into weekly periods indexed by $t \in T$. Before the rainy season, the authority chooses p candidate bases at which mobile clinics will be staged. After a flood scenario $\omega \in \Omega$ is observed, it updates road travel times, assigns communities to available clinics and constructs routes for that week. This timing leads naturally to a two-stage model: stationing variables are here-and-now decisions, while assignment, service and routing variables are recourse decisions.

Let I be the set of rural communities, J the set of candidate clinic bases, K the fleet of mobile clinics, and $V = I \cup J \cup R$ the vertices of the road network, where R contains junctions or bridge nodes. The undirected baseline road graph is $G = (V, E)$. Each edge $e \in E$ has length ℓ_e, normal speed v_e, road-class indicator r_e, and flood-exposure score $h_e \in [0,1]$. Community i has population P_i, expected service demand d_{it}^{ω}, and vulnerability weight $q_i \geq 1$. The scenario probability is π_ω, with $\sum_{\omega\in\Omega} \pi_\omega = 1$.

Flood-dependent road network

Flooding changes both edge availability and traversal time. Let $a_{et}^{\omega} \in \{0,1\}$ equal one when edge e is open in period t under scenario ω. The scenario travel time is:

$$\tau_{et}^{\omega} = a_{et}^{\omega} \frac{60\ell_e}{v_e}[1 + \alpha(h_e f_t^{\omega})^{\gamma} + \delta r_e f_t^{\omega}] + (1 - a_{et}^{\omega})M \tag{1}$$

Where $f_t^{\omega} \in [0,1]$ is flood severity, $\alpha > 0$ controls nonlinear delay, $\gamma > 1$ makes delay increase rapidly near severe conditions, δ represents the additional sensitivity of unpaved roads, and M is a prohibitive constant. Edge closure is generated by a threshold rule,

$$a_{et}^{\omega} = \mathbf{1}\{h_e f_t^{\omega} + \zeta_{et}^{\omega} \leq \kappa_e\} \tag{2}$$

Where ζ_{et}^{ω} captures local effects such as overtopping, a damaged culvert or a bridge inspection, and κ_e is the edge resistance threshold. In a field implementation, a_{et}^{ω} can instead be read directly from remote-sensing classifications, road-agency reports or community observations.

For each open scenario graph $G_t^{\omega} = (V, E_t^{\omega})$, the minimum travel time from node u to node v is:

$$D_{uvt}^{\omega} = \min_{\mathcal{P}\in\mathcal{P}_{uv}(G_t^{\omega})} \sum_{e\in\mathcal{P}} \tau_{et}^{\omega} \tag{3}$$

Non-negative edge times permit Dijkstra's label-setting algorithm [18]. If no path exists, $D_{uvt}^{\omega} = M$. This explicit shortest-path layer is important: adding a uniform flood multiplier cannot represent a closed bridge that forces a long detour or disconnects a community.

Spatial accessibility and equity

Accessibility is measured from the deployed service rather than from the nearest permanent facility. Let C_k be the weekly consultation capacity of clinic k, and let x_{ijkt}^{ω} equal one if community i is assigned to clinic k staged at base j. A distance-decay function $g(D) = \exp(-D/\beta)$ converts travel time into an access contribution, where β is a decay scale in minutes. The population-weighted accessibility index for scenario ω is:

$$A_t^{\omega} = \frac{\sum_{i\in I} P_i\, q_i \sum_{j\in J} \sum_{k\in K} x_{ijkt}^{\omega}\, g\left(D_{jit}^{\omega}\right)}{\sum_{i\in I} P_i\, q_i} \tag{4}$$

The index lies in $[0,1]$: larger values indicate that more vulnerable population is served from bases with shorter travel times. A threshold coverage indicator supplements the continuous measure,

$$H_t^{\omega}(L) = \frac{\sum_{i\in I} P_i\, \mathbf{1}\{\min_{j:z_j=1} D_{jit}^{\omega} \le L\}}{\sum_{i\in I} P_i} \tag{5}$$

Where L is a policy-relevant response threshold. Equity is represented by the weighted upper quantile of base-to-community travel time. In the computational experiment, the 90th percentile was used so that a small number of remote settlements could not be hidden by a low average.

Proposition 1: *If edge travel times weakly increase and no additional edge opens, the accessibility index in (4) cannot increase for a fixed stationing and assignment plan.*

Proof: *Under the stated change, every feasible path remains unchanged or becomes slower, and some paths may disappear. Hence each shortest-path time D_{jit}^{ω} weakly increases. Because $g(D) = \exp(-D/\beta)$ is decreasing, every summand in the numerator of (4) weakly decreases while the denominator is constant. Therefore A_t^{ω} cannot increase. This monotonicity provides a useful data-quality check: a more severe road state should not produce a higher fixed-plan access score unless another input has changed.*

Two-stage stochastic robust location-routing model

The first-stage binary variable z_j equals one if base j is selected. The second-stage variables are x_{ijkt}^{ω} for assignment, y_{uvkt}^{ω} for travel of clinic k from service node u to v, s_{ikt}^{ω} for consultations delivered, and u_{it}^{ω} for unmet demand. Let F_j denote stationing cost, c_{uvt}^{ω} route cost, B_k the maximum weekly route duration, and b_i the on-site service time per consultation. A generalized scenario loss is:

$$L_t^{\omega} = \sum_{k\in K} \sum_{u,v\in I\cup J} c_{uvt}^{\omega}\, y_{uvkt}^{\omega} + \lambda_u \sum_{i\in I} q_i\, u_{it}^{\omega} + \lambda_e Q_{0.90}^{\omega} - \lambda_a A_t^{\omega} \tag{6}$$

Where $Q_{0.90}^{\omega}$ is the 90th-percentile access time and the λ parameters convert service, equity and access into comparable planning units. The combined objective is:

$$\min \sum_{j\in J} F_j\, z_j + (1-\rho-\eta) \sum_{\omega\in\Omega} \pi_{\omega}\, L^{\omega} + \rho \max_{\omega\in\Omega} L^{\omega} + \eta \mathrm{CVaR}_{\alpha}(L) \tag{7}$$

Where $L^{\omega} = \sum_t L_t^{\omega}$, ρ controls worst-case protection, η controls upper-tail protection, and α is the CVaR confidence level. The standard linear representation is:

$$\mathrm{CVaR}_{\alpha}(L) = \min_{\xi, \nu_{\omega}\ge 0} \left\{ \xi + \frac{1}{1-\alpha} \sum_{\omega\in\Omega} \pi_{\omega}\, \nu_{\omega} : \nu_{\omega} \ge L^{\omega} - \xi \right\} \tag{8}$$

The main constraints are:

$$\sum_{j\in J} z_j = p \tag{9}$$

$$\sum_{j\in J} \sum_{k\in K} x_{ijkt}^{\omega} \le 1, \quad x_{ijkt}^{\omega} \le z_j, \qquad i\in I,\ t\in T,\ \omega\in\Omega \tag{10}$$

$$s_{ikt}^{\omega} + u_{it}^{\omega} = d_{it}^{\omega}, \quad 0 \le s_{ikt}^{\omega} \le d_{it}^{\omega} \sum_{j\in J} x_{ijkt}^{\omega} \tag{11}$$

$$\sum_{i\in I} s_{ikt}^{\omega} \le C_k, \qquad k\in K,\ t\in T,\ \omega\in\Omega \tag{12}$$

$$\sum_{v} y_{ivkt}^{\omega} = \sum_{v} y_{vikt}^{\omega} = \sum_{j} x_{ijkt}^{\omega}, \qquad i\in I,\ k\in K \tag{13}$$

$$\sum_{u,v} D_{uvt}^{\omega}\, y_{uvkt}^{\omega} + \sum_{i\in I} b_i\, s_{ikt}^{\omega} \le B_k \tag{14}$$

$$y_{uvkt}^{\omega} = 0 \quad \text{if} \quad D_{uvt}^{\omega} = M \tag{15}$$

Constraint (9) limits the number of selected bases. Constraints (10) and (11) link assignment to stationing and decompose demand into served and unmet portions. Constraint (12) imposes clinic capacity. Equations (13) are route-flow conservation conditions; standard subtour-elimination constraints are added in the implementation. Constraint (14) limits driving plus clinical service time, and (15) prohibits movement between disconnected nodes.

The finite-scenario model is robust in the following operational sense.

Proposition 2: *If (z, x, y, s, u) satisfies constraints (9)–(15) for every $\omega \in \Omega$, then the stationing plan z admits a feasible recourse plan for every flood realization represented in the scenario set.*

Proof: *The first-stage vector z is common to all scenarios. For each ω, the associated variables $(x^{\omega}, y^{\omega}, s^{\omega}, u^{\omega})$ satisfy assignment, capacity, flow, duration and network-availability constraints on G_t^{ω}. Selecting the corresponding recourse vector after observing ω therefore produces a feasible operational plan. The statement does not extend to unrepresented events; this is why scenario coverage and out-of-sample stress testing are essential.*

Scenario Construction

Scenarios should combine seasonal forecasts with operational road information. The proposed workflow has five steps: (i) create a baseline road graph from authoritative road centre-lines; (ii) intersect road segments with flood-depth, water-persistence and drainage layers; (iii) estimate closure and speed-reduction functions by road class; (iv) cluster historical or simulated flood maps into a manageable scenario set; and (v) assign probabilities from empirical frequency or ensemble weights. The scenario set must include spatially distinct events. Two floods with equal regional rainfall can have different routing consequences when one affects a western feeder corridor and the other affects a central bridge.

For the computational demonstration, 12 scenarios were constructed: dry season, early rains, two moderate floods, spatially concentrated severe floods in western, central and eastern subregions, prolonged flooding, a bridge failure, two recovery stages and a late-season local flood. Probabilities sum to one. The case uses no observed flood frequencies; weights are modelling inputs chosen to exercise the framework.

Solution algorithm

The model is solved through shortest-path preprocessing and scenario decomposition. For each scenario, Dijkstra's algorithm generates the matrix D^{ω}. The master problem selects bases and contains an epigraph for worst-case loss and the CVaR variables. Scenario subproblems then construct capacity-feasible assignments and routes. Violated cost or feasibility information is returned to the master as a cut, following the logic of column-and-constraint generation for two-stage robust optimization [19].

For the small demonstrator, all $\binom{6}{3} = 20$ stationing combinations were enumerated. Communities were ordered from most difficult to least difficult according to minimum scenario travel time and vulnerability; each was assigned to the nearest clinic with remaining capacity. A nearest-neighbour tour on the scenario shortest-path metric was then improved by pairwise exchanges. Enumeration makes the reported stationing choice auditable. Larger applications should use mixed-integer programming, decomposition and route-generation techniques [6, 7, 19].

The algorithm terminates when the relative robust gap satisfies:

$$\frac{UB - LB}{\max\{1, |UB|\}} \le \varepsilon \tag{16}$$

Where UB and LB are valid upper and lower bounds and ε is the chosen tolerance. Operationally, the same software can be run in rolling-horizon form: base decisions are retained for the season, but travel matrices, assignments and routes are recomputed whenever a road-status update arrives.

Synthetic case and parameterization

The test system represents a 73 km east–west rural corridor with settlements on both sides of a flood-prone central band. It contains 24 communities, six candidate bases and 76 road links. The communities have a combined synthetic population of 24,147 and expected weekly demand of 900 consultations. Three identical mobile clinics are available; each has capacity for 510 consultations per week, which leaves reserve capacity for spatial imbalance and route disruption. Baseline speeds are 62 km/h on paved links and 38 km/h on unpaved links. The decay parameter in (4) is $\beta = 75$ minutes. The objective weights used

in the experiment are $\rho = 0.35$ for worst-case protection and $\eta = 0.20$ for the upper tail, with the remaining weight on probability-weighted performance.
Table 1 summarizes the principal inputs. All quantities are synthetic and are supplied to make the experiment reproducible, not to describe a real district.

Table 1: Principal inputs for the synthetic case study

Component	Value	Interpretation
Communities	24	Demand nodes
Candidate bases	6	Safe staging candidates B1–B6
Mobile clinics	3	One vehicle per selected base
Road links	76	Paved and unpaved undirected links
Total population	24,147	Synthetic residents
Weekly demand	900	Expected consultations
Clinic capacity	510	Consultations per clinic per week
Paved-road speed	62 km/h	Dry-condition speed
Unpaved-road speed	38 km/h	Dry-condition speed
Access decay	75 min	Parameter β in (4)
Coverage threshold	45 min	Parameter L in (5)
Scenarios	12	Seasonal and disruption states
Risk weights	$\rho = 0.35$, $\eta = 0.20$	Worst-case and CVaR weights

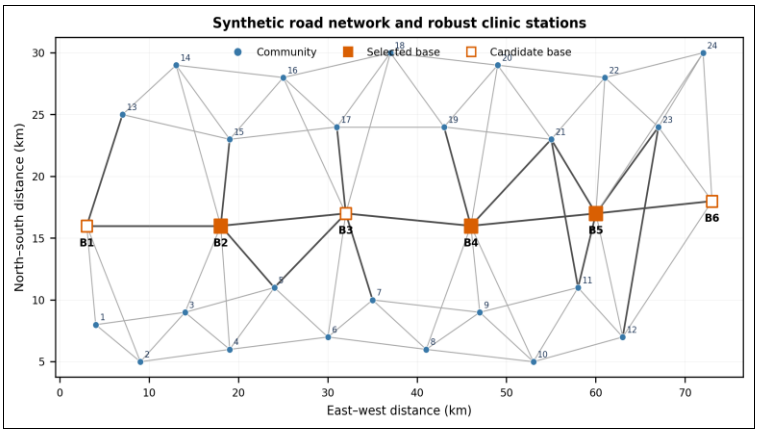

Fig 1: Synthetic road network and robust stationing decision. Dark links represent the paved corridor and pale links represent feeder roads. Coordinates are illustrative planar distances

Outcome measures and validation

Five outcomes were recorded: generalized objective value, total weekly route minutes, accessibility index A^{ω}, population coverage $H^{\omega}(45)$, and the 90th percentile of base-to-community travel time. The robust policy was compared with a dry-season policy obtained by optimizing only the first scenario. Internal validation comprised four checks. First, scenario probabilities were verified to sum to one. Second, all routes were restricted to open links. Third, served demand did not exceed clinic capacity. Fourth, accessibility obeyed the monotonicity in Proposition 1 when the network was degraded without compensating changes.
Because the data are synthetic, conventional inferential statistics would be misleading. The appropriate evidence is deterministic reproducibility under fixed inputs, comparison across scenario policies and sensitivity to adverse network states. No confidence interval is reported as though the synthetic nodes were a random sample from a real population.

Results and Discussion

Stationing and route adaptation

The robust objective selected B2, B4 and B5, spreading the three clinics across the west-central, central-east and eastern portions of the corridor (Fig 1). Dry-season optimization selected B2, B3 and B5. The difference is operationally meaningful: the robust solution moves the middle clinic from B3 to B4. B3 is efficient under ordinary travel conditions, but several central disruptions make routes originating there more sensitive to delays. B4 gives up some average efficiency while providing alternative approach paths to communities on both sides of the central flood band.
This result illustrates why a map-centre argument is insufficient. The geometric centre of demand need not be the network-resilient centre. A base can be close in Euclidean distance yet poorly placed relative to bridges, drainage bottlenecks and route redundancy. In the model, robust stationing values topological alternatives as well as short distance. This agrees with reliable location-routing research, where backup connectivity and disruption-aware assignments change the preferred facility pattern [12, 17].
The scenario solution retains the same seasonal bases but adapts community assignments and tour order after each road state is known. Such recourse is important. Requiring one fixed tour for every scenario would protect against disruption by excessive conservatism and would waste capacity in normal weeks. Conversely, reselecting bases after every event may be impractical

because staff accommodation, medicine storage, security and vehicle servicing must be arranged in advance. The two-stage structure therefore corresponds to the actual difference between slow strategic decisions and fast operational decisions.

Efficiency–resilience trade-off

Table 2 compares the two stationing policies. The dry-season policy has the lower probability-weighted objective, 645.0 generalized units, compared with 666.2 for the robust plan. The robust premium is therefore 3.3%. In return, the maximum scenario loss falls from 835.9 to 761.0, a reduction of 9.0%. The mean of the three largest scenario losses, used as an empirical upper-tail measure, falls from 774.5 to 752.7, or 2.8%.

Table 2: Comparison of dry-season and robust stationing policies

Measure	Dry-season stationing B2–B3–B5	Robust stationing B2–B4–B5	Relative effect of robust policy
Expected objective	645.0	666.2	+3.3%
Worst-case objective	835.9	761.0	−9.0%
Mean of three worst outcomes	774.5	752.7	−2.8%
Expected route time	493.1 min	512.8 min	+4.0%
Expected accessibility index	0.737	0.736	−0.1%
Expected population within 45 min	98.3%	96.7%	−1.6 percentage points
Expected consultations allocable	900	900	No change

The robust policy is not uniformly superior; that would contradict the trade-off in (7). It incurs longer expected routes and slightly lower average 45-minute coverage. Its benefit appears in the adverse tail. A decision-maker should therefore interpret robustness as insurance. If flood disruption has low operational or health consequences, the 3.3% expected-cost premium may not be justified. If missed maternal, vaccination or chronic-care visits have high consequences, the 9.0% reduction in the worst composite loss can be more important than a small average penalty.

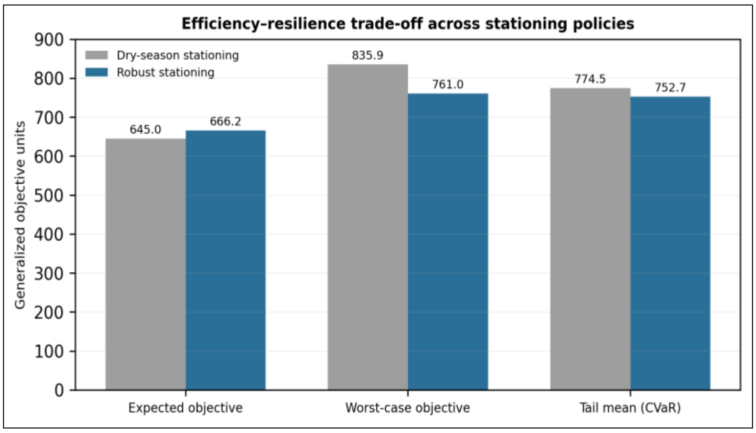

Fig 2: Efficiency–resilience trade-off across stationing policies. The robust plan sacrifices average efficiency but protects the most adverse scenarios

The expected accessibility indices of the two policies are almost identical, although their risk profiles differ. This is a useful warning against selecting a plan from one aggregate indicator. An expected score can conceal a severe week in which a subset of settlements experiences sharp deterioration. The model therefore reports expected value, worst case, tail mean, threshold coverage and a high travel-time quantile together. These measures answer different managerial questions.

Scenario-specific performance

Table 3 shows the robust policy under all 12 scenarios. Route time rises from 398.8 minutes in the dry season to 571.5 minutes in the bridge-failure scenario. The accessibility index declines from 0.796 to 0.688. The 90th-percentile base-to-community travel time remains below 45 minutes in every scenario, but population coverage within 45 minutes falls to 89.1% in the central, prolonged and bridge-failure scenarios. These results show that a percentile and a population share need not move identically: the quantile is community-based, whereas coverage is population-weighted.

Table 3: Performance of the robust policy by flood scenario

Scenario	Probability	Route time (min)	Accessibility index	Population within 45 min	90th-percentile time (min)
Dry season	0.08	398.8	0.796	100.0%	24.0
Early rains	0.08	447.4	0.781	100.0%	26.1
Moderate flood A	0.10	482.9	0.760	100.0%	29.4
Moderate flood B	0.10	501.4	0.750	100.0%	31.4
Severe west	0.10	551.3	0.695	91.8%	43.0
Severe central	0.10	567.1	0.698	89.1%	41.1
Severe east	0.10	568.9	0.716	100.0%	36.9
Prolonged flood	0.08	559.6	0.695	89.1%	40.8
Bridge failure	0.05	571.5	0.688	89.1%	42.1
Recovery week 1	0.07	509.7	0.745	100.0%	32.3
Recovery week 2	0.07	463.2	0.772	100.0%	27.5
Late-season local flood	0.07	521.0	0.739	100.0%	33.5

All 900 expected consultations were allocable in the experiment because total fleet capacity exceeded demand and no community became fully disconnected from every selected base. This result should not be interpreted as proof that flooding causes no unmet care. It indicates that, under the chosen synthetic scenarios and capacity assumptions, routing rather than clinical capacity is the active stressor. More severe closures, reduced staffing, medicine stock-outs or longer consultation times would activate the unmet-demand variables in (11).

The scenario table also identifies when contingency measures should be triggered. When $H^{\omega}(45)$ falls below a policy threshold, the authority can pre-position a boat, authorize an overnight clinic stop, transfer a community to a neighbouring district's mobile unit, or establish a temporary foot-access collection point. Thus the optimization does not replace local emergency judgment. It indicates where that judgment is most urgently required.

Mathematical interpretation

The change from B3 to B4 can be interpreted as a reduction in exposure to a network cut. Let C^{ω} be a minimal set of edges whose closure separates a selected base from a group of demand nodes. A dry-season solution mainly minimizes path length. The robust objective also responds to the probability and consequence of small cut sets. When B3 depends heavily on a central crossing, the marginal reduction in normal travel time can be outweighed by the tail loss produced when that crossing is degraded. B4 is selected because its incident routes provide better access to alternative arcs.

The CVaR term further distinguishes frequent moderate floods from rare severe disruptions. The expected term alone weights a 5% bridge failure lightly. The maximum term can overreact to one extreme scenario. CVaR provides an intermediate measure by averaging losses above a quantile. The combined objective in (7) is deliberately plural: expected performance supports routine service, the maximum limits catastrophic exposure, and CVaR stabilizes the adverse tail. The relative weights are policy choices, not universal constants.

There is also a useful relationship between network robustness and health equity. If vulnerability weights q_i are increased for communities with high maternal risk, chronic-disease burden, poverty or limited alternative transport, then unmet demand and accessibility loss at those communities become more expensive. The optimizer may respond by selecting a less central base or by preserving route capacity for a small high-priority settlement. Equity is therefore operationalized through explicit weights and quantiles rather than asserted after a cost-minimizing solution has been chosen.

Comparison with existing approaches

Maximum-coverage models are valuable when a single distance standard defines acceptable access [4]. Their binary coverage structure is less informative when flood delays vary continuously and the same community moves from 25 to 70 minutes across scenarios. The exponential decay in (4) retains that information. Floating-catchment approaches account for the balance between provider capacity and population demand [2, 3], but they generally evaluate a given supply configuration rather than jointly constructing mobile-clinic tours. The present model embeds a related accessibility principle inside the decision problem.

Location-routing formulations integrate facility and vehicle decisions [6, 7], while robust vehicle-routing models address uncertain travel time [9-12]. The present framework adds explicit flood-state graphs and a population-weighted access term. Humanitarian network models provide close methodological precedents because they combine uncertain demand, disrupted transport and equity [13-17]. Mobile primary care differs in two respects: service repeats across a season, and on-site consultation time consumes vehicle-day capacity. The model therefore includes a rolling horizon and clinical service time in (14).

Decision-support implementation

A deployable system can be organized in four layers. The data layer stores communities, population, health demand, road attributes, bridges, flood depth and candidate safe sites. The analytics layer updates G_t^{ω}, computes shortest paths and solves the location-routing model. The presentation layer displays selected bases, ordered routes, inaccessible links, expected arrival times and communities below an access threshold. The governance layer records data timestamps, overrides, reasons for deviations and post-event outcomes.

The minimum operational dataset is modest but must be maintained. Community coordinates and approximate population can be updated less frequently, and rural-access indicators provide a useful starting point for diagnosing all-season connectivity [24]. Road status, bridge passability, vehicle availability, staff availability and medicine capacity require near-real-time confirmation

during floods. Uncertain or conflicting road reports should be represented as alternative scenarios rather than silently converted into one assumed status. A route issued by the system must also identify the date and time of the underlying road information.
Implementation should proceed as decision support, not automatic dispatch. Local drivers may know that a mapped road is unsafe despite appearing open, or that a technically passable link is unsuitable for a heavy clinic vehicle. Health teams may also face priorities not encoded in expected consultation demand. The software should therefore allow a dispatcher to lock or prohibit an edge, force service to a community, re-run the optimization and record the override. This human-in-the-loop design strengthens accountability.

Limitations

The principal limitation is the synthetic evidence base. The reported numerical improvements demonstrate model behaviour and do not estimate benefits for Namibia or any other country. Before use, road speeds, closure functions, demand, service times, vehicle capabilities and vulnerability weights must be calibrated with local data. Scenario probabilities may be unstable under climate change, so historical frequency alone may understate future extremes.
The model simplifies health operations. It treats expected consultations as divisible service demand, assumes that each clinic carries a compatible service package and does not model staff specialities, cold-chain limits, medicine compatibility, referral transport or patient scheduling. The routing heuristic is transparent but does not guarantee a globally optimal tour for a large instance. Communication failures and delays in road-status reporting are also outside the present state model.
Accessibility is necessary but not sufficient for utilization. Cultural acceptability, information, clinic opening time, affordability and perceived quality affect whether people use a service. The index in (4) should therefore be combined with community consultation and utilization evidence. Finally, a plan feasible for enumerated scenarios is not protected against every conceivable event. Stress tests should add compound cases such as simultaneous bridge failure, vehicle breakdown and demand surge.

Conclusion

This study formulated mobile-clinic deployment under seasonal flooding as a two-stage stochastic robust location-routing problem on a changing road graph. The framework selects seasonal staging bases before uncertainty is resolved and adapts assignments and routes after a flood scenario is observed. Shortest-path preprocessing captures closures and detours; capacity and route constraints preserve operational feasibility; and accessibility, equity, worst-case loss and CVaR prevent average travel cost from becoming the only planning criterion.
In the synthetic experiment, robust stationing at B2, B4 and B5 reduced the worst-scenario objective by 9.0% and the upper-tail mean by 2.8% relative to dry-season stationing, at a 3.3% expected-cost premium. The result is not a field-effect estimate. It shows how modest changes in stationing can protect adverse outcomes without materially changing average accessibility. The most useful output is therefore not a single route but a contingent plan that links observable road states to revised routes and clearly identifies communities whose access deteriorates.
Future work should calibrate the model with road-condition histories, flood-depth rasters, community demand and actual mobile-clinic operating data. Methodological extensions should include distributionally robust scenario probabilities, multi-period medicine inventories, heterogeneous clinical teams, vehicle breakdown, boat–road intermodal service and endogenous temporary crossing repair. Prospective evaluation should compare model-supported deployment with existing practice using missed visits, travel time, cost, equity and patient outcomes.

Recommendations

Health authorities should begin with a pilot in one flood-prone service area. The pilot should geocode communities, candidate safe bases, bridges and health facilities; classify each road by surface, seasonal passability and vehicle suitability; and agree on a simple road-status reporting protocol. At least one normal, moderate and severe disruption scenario should be reviewed with drivers, nurses, emergency managers and community representatives.
The selected seasonal plan should include predefined triggers. Examples are a bridge-closure alert, a forecast flood-depth threshold, loss of communication with a road monitor, or 45-minute population coverage below 90%. Each trigger should map to an action such as route recomputation, overnight staging, deployment of an alternative transport mode or coordination with an adjacent district.
Managers should retain several indicators on the dashboard: expected route time, worst-case route time, unmet demand, population within the travel threshold, accessibility index and the high travel-time quantile. A low average alone should not be accepted as evidence of equitable access. After every flood season, planned and actual routes, failed links, missed visits and overrides should be reviewed to recalibrate the model.

Acknowledgement

The author acknowledges colleagues in applied mathematics, public health, transport planning and geographic information science whose interdisciplinary discussions motivated the decision-support framework. No external financial support was received for this study. The author declares no competing financial or non-financial interests. No human participants, patient records or identifiable personal data were used. All numerical results are based on synthetic scenarios, and the fixed computational design is described in the manuscript to support reproduction.

References

1. Tansley G, Schuurman N, Amram O, Yanchar NL. Spatial access to emergency services in low- and middle-income countries: A GIS analysis. PLoS One. 2015; 10:e0141113.

2. Luo W, Wang F. Measures of spatial accessibility to health care in a GIS environment: Synthesis and a case study in the Chicago region. Environment and Planning B: Planning and Design. 2003; 30:865-884.
3. Delamater PL. Spatial accessibility in suboptimally configured health care systems: A modified two-step floating catchment area metric. Health and Place. 2013; 24:30-43.
4. Church R, ReVelle C. The maximal covering location problem. Papers of the Regional Science Association. 1974; 32:101-118.
5. Daskin MS. Network and Discrete Location: Models, Algorithms and Applications. New York: Wiley, 1995.
6. Laporte G. Location-routing problems. In: Golden BL, Assad AA, editors. Vehicle Routing: Methods and Studies. Amsterdam: North-Holland, 1988, 163-197.
7. Nagy G, Salhi S. Location-routing: Issues, models and methods. European Journal of Operational Research. 2007; 177:649-672.
8. Gu W, Wang X, McGregor SE. Optimization of preventive health care facility locations. International Journal of Health Geographics. 2010; 9:17.
9. Gendreau M, Laporte G, Séguin R. Stochastic vehicle routing. European Journal of Operational Research. 1996; 88:3-12.
10. Ben-Tal A, El Ghaoui L, Nemirovski A. Robust Optimization. Princeton: Princeton University Press, 2009.
11. Bertsimas D, Brown DB, Caramanis C. Theory and applications of robust optimization. SIAM Review. 2011; 53:464-501.
12. Eufinger L, Kurtz J, Buchheim C, Clausen U. A robust approach to the capacitated vehicle routing problem with uncertain costs. INFORMS Journal on Optimization. 2020; 2:79-95.
13. Campbell AM, Vandenbussche D, Hermann W. Routing for relief efforts. Transportation Science. 2008; 42:127-145.
14. Mete HO, Zabinsky ZB. Stochastic optimization of medical supply location and distribution in disaster management. International Journal of Production Economics. 2010; 126:76-84.
15. Rawls CG, Turnquist MA. Pre-positioning of emergency supplies for disaster response. Transportation Research Part B: Methodological. 2010; 44:521-534.
16. Bozorgi-Amiri A, Jabalameli MS, Al-e-Hashem SMJM. A multi-objective robust stochastic programming model for disaster relief logistics under uncertainty. OR Spectrum. 2013; 35:905-933.
17. Xie W, Ouyang Y, Wong SC. Reliable location-routing design under probabilistic facility disruptions. Transportation Science. 2016; 50:1128-1138.
18. Dijkstra EW. A note on two problems in connexion with graphs. Numerische Mathematik. 1959; 1:269-271.
19. Zeng B, Zhao L. Solving two-stage robust optimization problems using a column-and-constraint generation method. Operations Research Letters. 2013; 41:457-461.
20. Snyder LV. Facility location under uncertainty: A review. IIE Transactions. 2006; 38:547-564.
21. Toth P, Vigo D, editors. Vehicle Routing: Problems, Methods, and Applications. 2nd ed. Philadelphia: Society for Industrial and Applied Mathematics, 2014.
22. World Health Organization. Operational Framework for Building Climate Resilient Health Systems. Geneva: World Health Organization, 2015.
23. World Health Organization. Operational Framework for Building Climate Resilient and Low Carbon Health Systems. Geneva: World Health Organization, 2023.
24. World Bank. Rural Access Index: A Key Development Indicator. Washington, DC: World Bank, 2016.
25. Hakimi SL. Optimum locations of switching centers and the absolute centers and medians of a graph. Operations Research. 1964; 12:450-459.