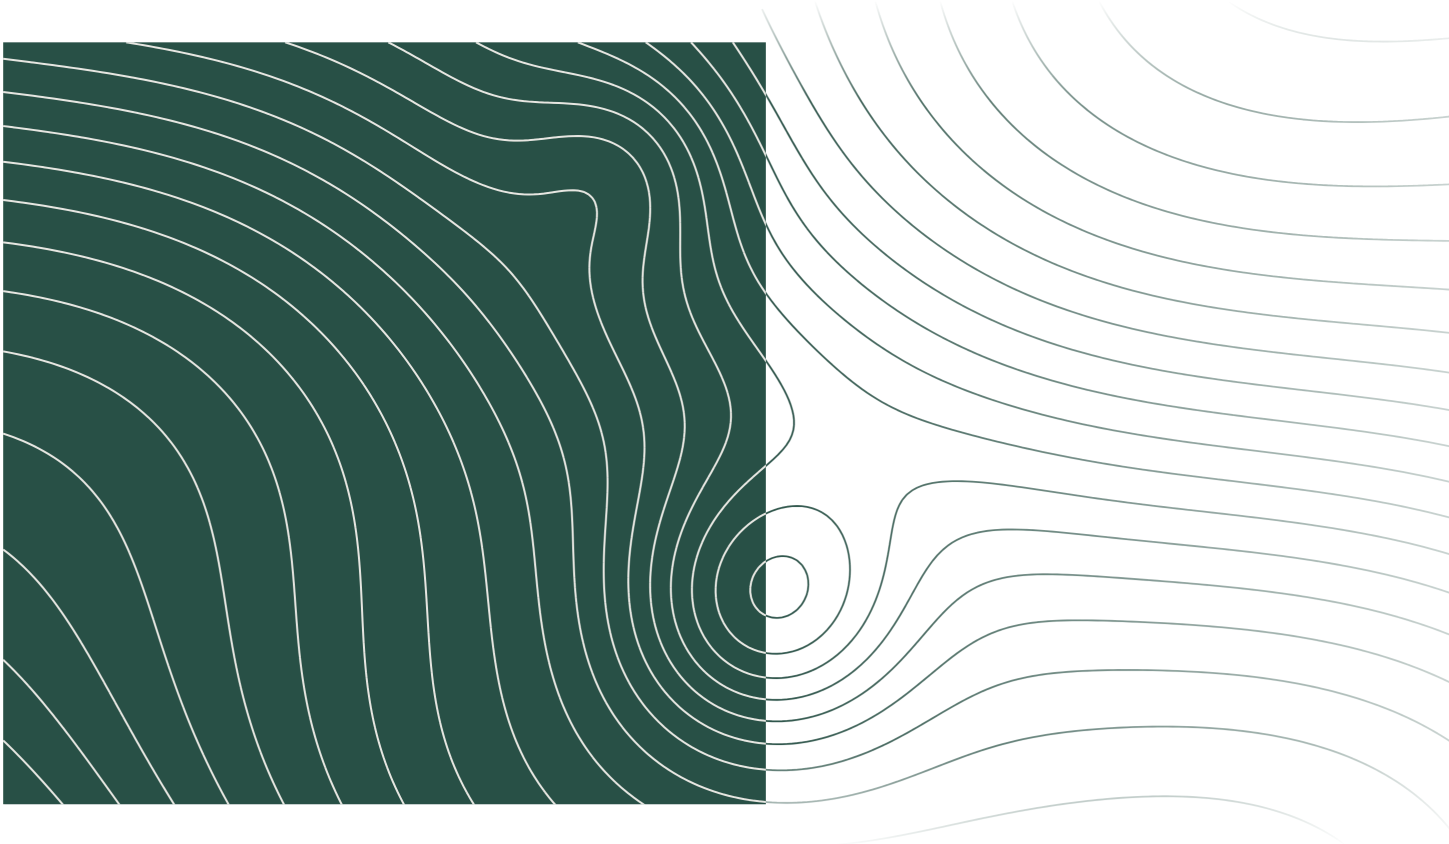

GIA · GABON INFRASTRUCTURE ASSETS

Analytical methods for territorial planning policy

A geospatial decision-support platform for infrastructure planning in Gabon, applied to access to healthcare, with the limits of its data stated before its results.

AUTHORS
Hodologic Research

PUBLISHED
21 September 2026

LANGUAGE
English

Analytical methods for territorial planning policy

Hodologic Research

21 September 2026

Abstract

Infrastructure planning in Gabon rests on data that is scattered, heterogeneous and unevenly reliable. We present GIA, a geospatial decision-support platform that brings open layers (population, buildings, networks, environment, services, works in progress, event feeds) onto a common base, and derives from it reproducible indicators for three kinds of decision: diagnosis, prioritisation and vigilance. We formalise its analytical modules — dasymetric allocation of the 2026 census population, accessibility over the road network, facility siting by maximal covering, network fragility through the bridges of the road graph, section load through betweenness centrality, alert rules — then apply them to access to healthcare.

The inventory used is the OpenStreetMap volunteered survey, 253 facilities, each of them a building surveyed on the ground or traced on imagery. **Its coverage is uneven, and we state that before any result:** 218 of its 253 points lie in one province out of nine, and one whole province carries none. On that basis, **90.7 % of the population lives within one hour of a centre** and 327,232 people lie beyond; half of them fall in six departments, five well-placed centres would cover 46.1 % of them, and 45.7 % of the population depends on a single road link.

The first location the model designates falls in the province the inventory leaves empty. We publish it while naming it as such, because a declared artefact remains usable and a silent one does not. That is the methodological result this work defends: a geographic indicator measures first the map that carries it, and a decision-support platform must make that measurement visible before the one it is asked for.

We draw three transferable contributions from it: *the chain of positional uncertainty* (§ 3.9), which describes how a property known only to the producer of an inventory propagates or is lost all the way to the screen; *a measure of road network load* by betweenness centrality (§ 3.10, § 4.5), where one per cent of the kilometres carries 39 % of the journeys; and *an external validation protocol for positions* (§ 5.3), which would turn adjectives into kilometres.

Keywords: spatial decision support, geographic information system, infrastructure, healthcare accessibility, maximal covering, positional uncertainty, data provenance, open data, Gabon.

1 Introduction

1.1 Context

Deciding where to build a road, a health centre or a power line, which sections to maintain first and which works to watch, requires cross-referencing information of very different kinds: where people live, how they travel, which services exist, which projects are under way and which events threaten them. In Gabon this information often exists, but it is scattered across the census, volunteered mapping, international databases, satellite feeds and the administrative tracking of projects, with uneven dates, precisions and coverage. The results of the 2026 general population and housing census (RGPL), published at province level, renew the demographic base without providing a fine-grained distribution.

Geographic indicators such as travel time to the nearest health centre are now produced at global scale [8]. Spatial decision support systems [2] answer this need by combining geographic data, analytical models and interfaces intended for decision-makers. Their value does not rest on model quality alone: an indicator that is exact but computed on incomplete data, or presented without its assumptions, can steer a decision towards an artefact of the map. The GIA platform was designed for Gabon with that constraint at its centre.

1.2 Questions

Q1. Integration. What common base of data and models makes it possible to answer different decision questions with the same building blocks?

Q2. Translation into decisions. Which indicators does the platform produce for diagnosis, prioritisation and vigilance, and what do they say about a concrete case?

Q3. Reliability. Under what conditions can these indicators ground a decision, and how does the platform make their limits visible?

1.3 Contributions

An architecture for decision support that separates sources, common base, analytical modules and alerts, and that attaches to every indicator its sources and its assumptions (§ 2).

Formalised generic modules: dasymetric allocation, accessibility, exact siting by bounded enumeration, fragility without double counting, section load through betweenness centrality (§ 3.10) and alert rules (§ 3). They depend on the service studied only through its inventory.

The chain of positional uncertainty (§ 3.9): how a property declared by the producer of an inventory propagates, or is lost, all the way to the screen.

A complete case study tying each indicator to a decision (§ 4), and publishing the artefact its own coverage produces instead of concealing it.

An external validation protocol for positions (§ 5.3) and a roadmap of features to develop (§ 6).

2 The GIA platform

2.1 Architecture

GIA is a web mapping application organised in four levels (figure 1). The sources bring together some fifty layers across seventeen themes (rail, road and power networks, infrastructure, hydrography, administrative divisions, environment, hazards and extraction, local services, forest, places, activity zones, works in progress), together with event feeds (fire detections, air traffic, press). The common base reduces them to four shared objects: a located population, a road graph, service inventories and a register of works. The modules compute indicators on that base. The alert rules cross those indicators with works and events. The interface presents the results as thematic dossiers, interactive tools and printable sheets.

Three names, and they do not overlap. *GIA* is the platform: the application, its layers and its indicators. *Hodologic* is the organisation that makes it available and runs it publicly. *Hodologic Research* is the entity that signs methodological work, including this document. Distinguishing them is not a formality: a platform, an operator and an author do not carry the same responsibilities, and what follows criticises the platform without the operator having to defend itself.

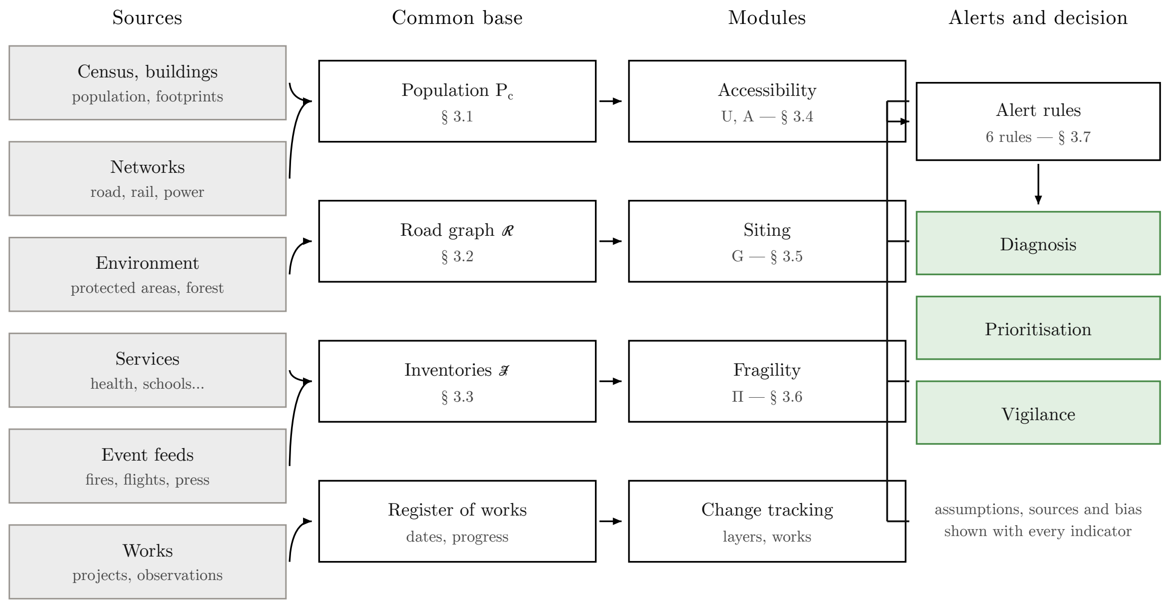

Figure 1. Platform architecture: sources, common base, analytical modules, alert rules and the kinds of decision served. Every indicator is displayed with its sources and its assumptions. *(Figure labels are in French in the original.)*

2.2 Features and kinds of decision

We classify features by the decision they serve: diagnosis (describing a situation and locating it), prioritisation (comparing courses of action) and vigilance (detecting what requires attention). Table 1 lists the features of the current version.

Table 1. Features of the current version, by kind of decision.

Kind	Feature	What the user gets
Diagnosis	Layer catalogue	52 layers in 17 themes, search by name, stack with visibility and opacity
	"Access to care" dossier	Population within 30, 60 and 120 min of a centre, by department, with a written reading
	Ranking–map link	Hovering a province in a ranking highlights it on the map
Prioritisation	Siting workshop	Gain from a centre placed at a point, plateau of equivalent sites, multi-centre plan
Vigilance	"Network fragility" dossier	Road links whose cutting would isolate inhabitants from every centre
	"Works" dossier	List of projects, deadlines and progress
	"Signals" dossier	Alerts from the rules and press signals
	Field observations	Notes taken at a work site, held offline until sent
Reporting	Sheets	Printable sheets by province

2.3 Design principles

Four principles govern the production of indicators. **(i) Provenance**: every indicator carries the source, the date and the licence of the data that produced it. **(ii) Declared assumptions**: the model's conventions are named, together with the direction of the bias they introduce. **(iii)**

Ranges: where a convention is not settled by the data, the indicator is given as an interval. **(iv) Prudent thresholds**: a threshold separating two statements is rounded towards the side that makes neither statement false, and whatever a threshold excludes is listed by name.

3 Methodology

Distances are geodesic (km) and times are expressed in hours unless stated otherwise. The modules in this section depend on the kind of service only through the inventory $\mathcal{F}$; we describe them for health facilities.

Table 2 collects the notation.

Table 2. Notation.

Symbol	Definition	Unit
c, x_c, a_c	Cell of the 0.025° × 0.025° grid, its centre, its area	—, —, km²
ϱ_c	Radius of the disc of equal area, $\sqrt{(a_c/\pi)}$	km
b_c	Number of buildings in the cell	—
P_p, P_d, P_c	Population of a province (enumerated), of a department, of a cell (estimated)	pers.
$\mathcal{R} = (V, E)$	Undirected road graph	—
d_H, δ	Haversine distance; shortest-path time in $\mathcal{R}$	km, h
v_e, v_w	Speed on section e; walking speed (5)	km/h
r_{max}	Snapping range to the network (25)	km
$\mathcal{F}$, $s(f)$	Inventory of facilities; snapping node of f	—
$T(n)$, τ_c	Time from node n to the nearest centre; access time of the cell	h, min
A_t, U	Population within t min; uncovered population	pers.
$\mathcal{U}$, G_k	Uncovered cells; gain of candidate cell k	—, pers.
$\mathcal{P}_\varepsilon$	Tolerance plateau ε	—
π_j, h_j	Population and facilities assigned to junction j	pers., —
$S_\pi(v)$, $S_h(v)$	Sums of π and h over the traversal subtree of v	pers., —
Π	Population behind a single road link	pers.

Definition 1 (coverage). A cell is covered if $\tau_c \le 60$; it is unreachable, and therefore uncovered, if no node connected to a facility lies within r_{max}.

Definition 2 (gain). The gain G_k is the population of the uncovered cells that would become covered if a facility were added in cell k.

Definition 3 (critical link). A section is critical if removing it disconnects the graph and if the detached part contains no facility and at least 25 inhabitants; it is maximal if its detached part is contained in that of no other critical link. A weak point groups the maximal links falling in the same tile of about one kilometre.

3.1 Population

Provincial census counts are allocated in proportion to buildings, following the principle of dasymetric allocation [4] (figure 2):

$$P_d = P_p \cdot B_d \,/\, \Sigma_{d' \in p} B_{d'} \;, \quad P_c = P_d \cdot b_c \,/\, B_d \;, \quad B_d = \Sigma_{c \in d}\, b_c \tag{1}$$

with $P_c = 0$ if $B_d = 0$. The grid holds 35,167 cells of about 7.7 km². The building layer combines OpenStreetMap footprints and Open Buildings detections [6]: overlapping footprints form a group, counted as many times as the finer of the two delineations (figure 3), giving 816,791 buildings (748,821 if each group counts as one).

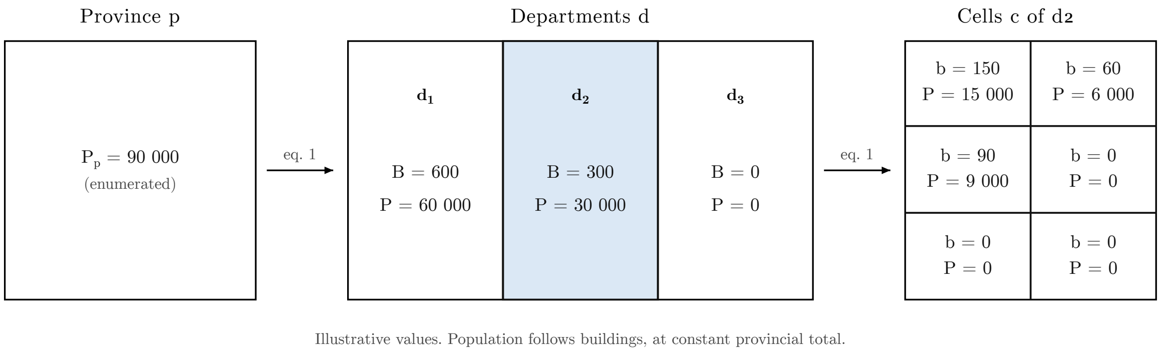

Figure 2. Dasymetric allocation (eq. 1): the provincial count is distributed among departments, then among cells in proportion to the number of buildings.

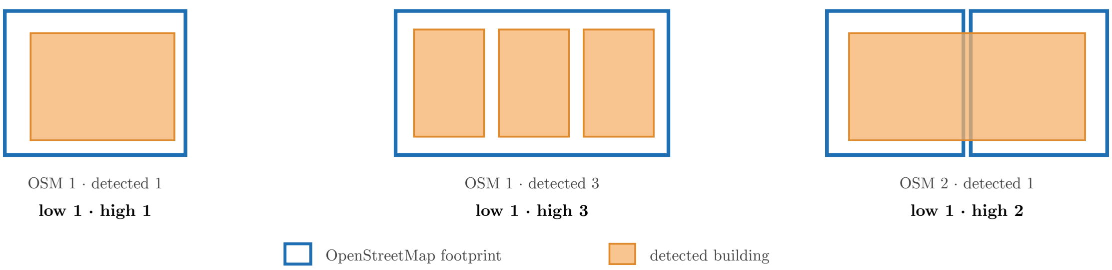

Figure 3. Counting buildings over three groups of overlapping footprints: a group counts as the larger of the two delineations (upper bound 6, lower bound 3).

3.2 Road network

The network is extracted from OpenStreetMap (886,354 nodes as of 12 September 2026). Coordinates are rounded to 10^{-5} degree so that coincident endpoints merge, and each segment becomes a section of duration $t_e = d_H(a,b) \,/\, v_e$, where d_H is the haversine distance and v_e depends on the road class (30 km/h by default).

3.3 Inventory of facilities

The service inventory comes from OpenStreetMap, queried through the Overpass API: **253 health facilities** — 173 dispensaries, 46 hospitals, 26 medical practices and 6 nursing posts. Every point in it is a building surveyed on the ground or traced on imagery by a contributor; the coordinate therefore designates the facility itself, and not the inhabited place that names it. This property cannot be derived from the coordinate: it follows from how the data was produced, and it is what section 3.9 formalises.

What this inventory does not cover, and it must be stated before any result. Volunteered mapping is thick where contributors are many and thin elsewhere: **218 of the 253 facilities are in the single province of Estuaire**, which holds 62 % of the population, and **Nyanga carries none at all**. The other eight provinces share 35 points. A low figure in a rural department is therefore a more fragile claim than a high one, and the gap pushes in a single direction: *real access is better than what is measured.* Every section that follows inherits this bound, and section 4 restates it at each result it draws.

The modules of this section depend on the kind of service only through the set $\mathcal{F}$. An inventory of schools, water points or markets substitutes for it without changing a line, provided its coverage is documented as this one's is.

3.4 Accessibility

Each facility is snapped at no cost to the nearest node within range r_{max}. A multi-source Dijkstra algorithm [3] gives the network time of every node, and the inhabitant walks to the node that minimises total time (figure 4):

$$T(n) = \min_{f\in\mathcal{F}} \delta(s(f), n), \quad \tau_c = 60 \cdot \min_{n \,:\, d_H(x_c,n) \le r_{max}} [\, T(n) + d_H(x_c,n) \,/\, v_w \,] \tag{2}$$

$$A_t = \Sigma_c\, P_c \cdot 1[\tau_c \le t], \quad U = \Sigma_c\, P_c \cdot 1[c \text{ uncovered}] \tag{3}$$

The search examines nodes in rings of 1, 3, 8, 15 and 25 km and stops as soon as the walk to the ring exceeds the best time so far, which makes the minimum exact. Free snapping is bracketed by an upper bound U^+, computed by initialising each source at the walking time $d_H(f, s(f)) \,/\, v_w$.

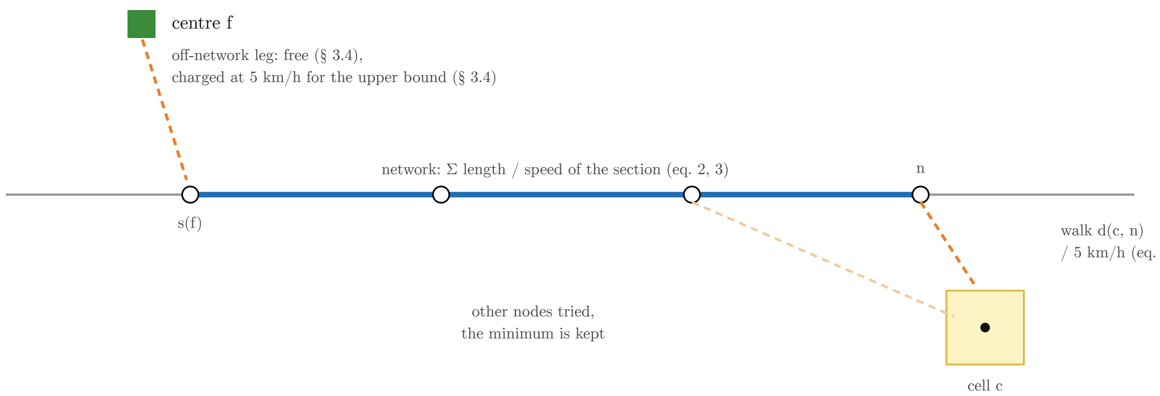

Figure 4. Components of a cell's access time (§ 3.4): snapping of the facility (free, except for the upper bound), network time, and the inhabitant's walk to the node that minimises total time.

3.5 Siting

Only uncovered cells contribute to the gain, through nodes lying within one hour's walk. For a candidate k snapped at n_k, with $D_k(n) = \delta(n_k, n)$ and $w_{c,n} = d_H(x_c,n) \,/\, v_w$:

$$G_k = \Sigma_{c\in\mathcal{U}}\, P_c \cdot 1[\, \min_{n \,:\, w_{c,n} \le 1} (D_k(n) + w_{c,n}) \le 1 \,] \tag{4}$$

The exploration from n_k stops at any node where $D_k(n) > 1$ or $D_k(n) \ge T(n)$.

Proposition 1. The cut $D_k(n) \ge T(n)$ does not change G_k.
Proof. For any node m reached beyond n, $T(m) \le T(n) + \delta(n,m) \le D_k(n) + \delta(n,m)$: the new facility improves the time neither of m nor of the cells that reach it. ∎

The computation takes about 4 ms per candidate, against 113 s for a full exploration, which allows an interactive response. What is presented is not a cell but a plateau $\mathcal{P}_\varepsilon = \{k : G_k \geq (1 - \varepsilon/100) \max G\}$ with $\varepsilon = 5$ (figure 5). For several centres, the greedy method adds at each step the cell of maximal gain and then recomputes the gains; since coverage is submodular, it reaches at least $1 - 1/e$ of the optimum of the maximal covering problem [1, 5].

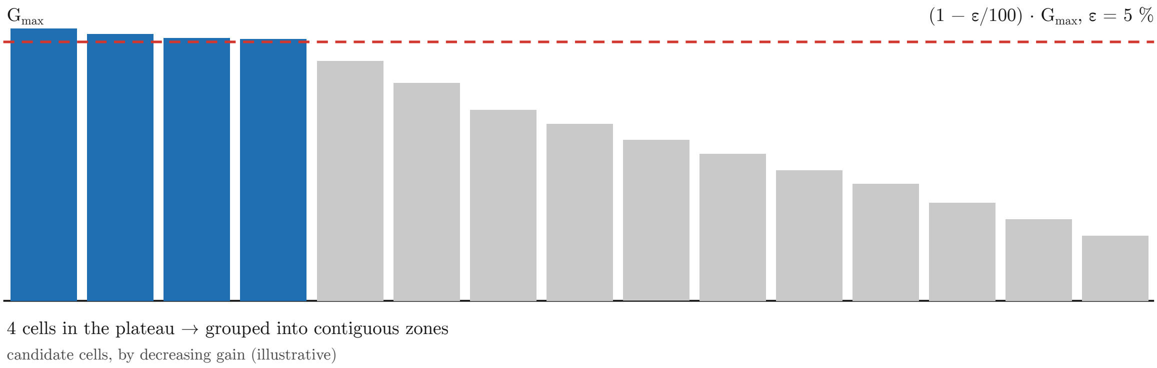

Figure 5. Tolerance plateau ε (illustrative values): every cell whose gain is within 5 % of the best is retained and grouped into contiguous zones.

3.6 Fragility

Chains of degree-2 nodes are contracted into links between junctions (60,340 junctions, 70,154 links), the population of each cell is distributed among the junctions of its equivalent disc, and each facility is assigned to its junction (figure 6). A depth-first traversal rooted at the junctions carrying a facility identifies the bridges [7]:

$$(u,v) \text{ critical} \iff \mathrm{low}(v) > \mathrm{disc}(u) \wedge S_h(v) = 0 \wedge S_\pi(v) \geq 25, \quad \Pi = \Sigma_{(u,v) \text{ maximal}} S_\pi(v) \quad (5)$$

Since the detached parts of maximal links are disjoint, Π counts each inhabitant at most once (figure 7).

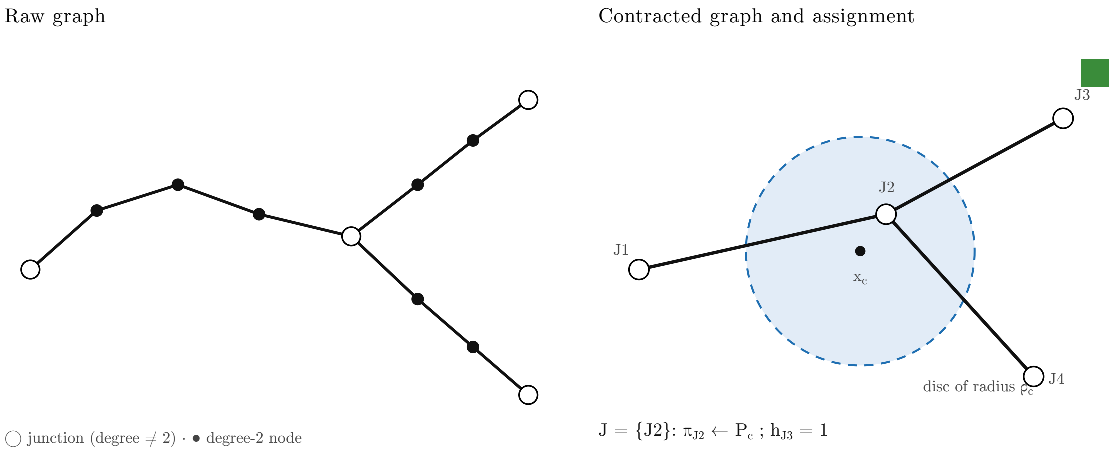

Figure 6. Contraction of degree-2 chains and assignment (§ 3.6): the population of a cell is distributed among the junctions of its equivalent disc, and each facility is assigned to its junction.

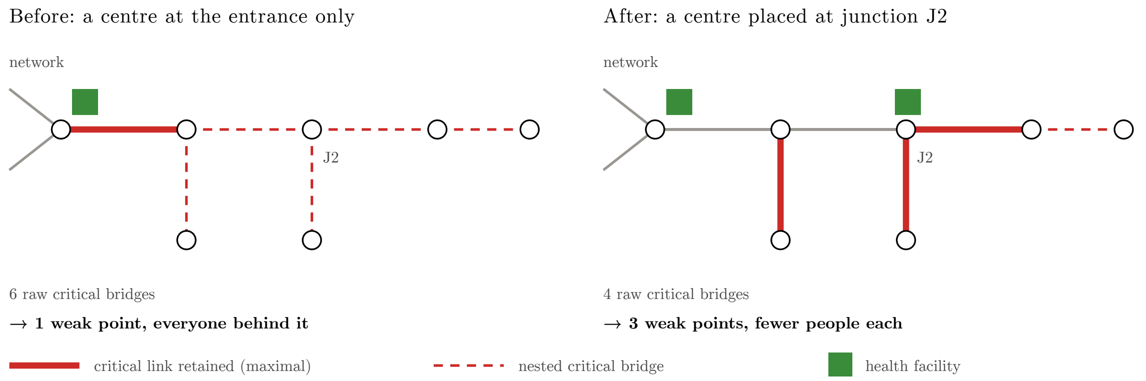

Figure 7. Effect of a facility added on a dead end (eq. 5). Critical links fall from 6 to 4, maximal links rise from 1 to 3, and the population behind each maximal link decreases.

3.7 Alert rules

Alerts are deterministic functions of the indicators, the works in progress and the event feeds. For a department d, a work site z, with n_7 the number of press signals over 7 days and s_7, s_{30} their negative shares over 7 and 30 days:

$$\mathrm{Underserved}(d) \iff 100 \cdot A_{60}(d)/P_d < 25 \wedge P_d \geq 2{,}000 \tag{6}$$

$$\mathrm{SensitiveWork}(z) \iff \exists\ (u,v) \text{ critical} : d_H(z,(u,v)) \leq 15 \wedge S_\pi(v) \geq 1{,}000 \tag{7}$$

$$\mathrm{NegativeSpike} \iff n_7 \geq 5 \wedge s_7 \geq 0.5 \wedge s_7 - s_{30} \geq 0.2 \tag{8}$$

To these are added the delayed work site, the work site located in a province under a negative spike, and the fire detected within 10 km of a work site with a confidence of at least 50. A rule whose observation window is empty is not evaluated: it is declared not evaluated, with the date at which collection stops. “Did not fire” and “was not evaluated” do not belong under the same word.

3.8 Assumptions and parameters

Table 3. Assumptions and direction of bias (+ overestimates, − underestimates, ? unknown, · no effect).

Assumption	Access	Gain	Fragility
Incomplete inventory of facilities	−	+	+
Position at the inhabited place, not at the building (§ 3.9)	?	?	?
Facilities snapped at no cost	+	+	·
Listed facilities assumed open and equivalent	+	?	−
Mapped network assumed complete	−	?	+
Constant speeds, without season or relief	+	?	·
Straight-line walking at constant speed	+	+	·
Population proportional to buildings	?	?	?
Components without a facility excluded	·	·	−
Single and total cut	·	·	−
Greedy plan	·	−	·

Parameters are declared within the platform: grid step 0.025°, default speed 30 km/h, walking 5 km/h, snapping range 25 km, matching radii 0.5 and 2 km, coverage threshold 60 min, plateau tolerance 5 %, minimum population of a critical link 25.

3.9 The chain of positional uncertainty

A coordinate does not say what it designates. The same point may hold at the building, at the inhabited place that names it, or at the centroid of an administrative unit: the difference is measured in kilometres, and it cannot be derived from the coordinate. Only the producer of the inventory knows it. It then propagates through the whole computation, changing nature at every step (figure 4), and can be lost at each of them.

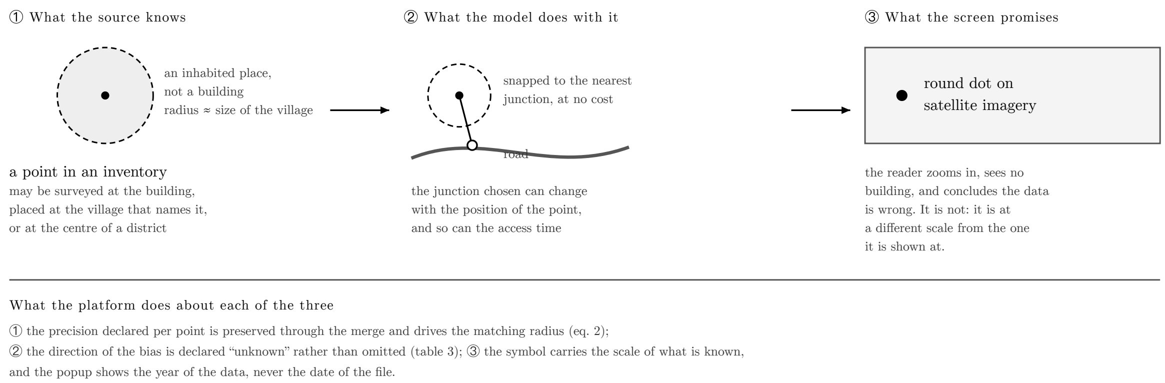

Figure 8. The chain of positional uncertainty, from source to screen. At every step the information "this position holds at the village" can be lost: when copied from one base to another, when snapped to the network, or when drawn. Once lost it never reappears downstream, but the screen goes on asserting.

Three requirements follow, and each corresponds to a defect we found in the platform itself. **(i) Precision is preserved**: the field that declares, for each point, whether the position holds at the building or at the inhabited place must survive the merge and be read by the screen, failing which the producer has documented a property that nobody uses. **(ii) The date is the date of the data**: the date a file was deposited is not the date of a survey, and taking the most recent of several sources' dates amounts to dating the whole by its most flattering half. **(iii) The symbol carries the scale**: a dot of fixed radius, laid over a satellite image, promises metric precision whatever the data behind it.

3.10 Through routes: the load on a section

Fragility (§ 3.6) answers "which cut would isolate inhabitants". It says nothing about *load*: a section may isolate nobody (because another route exists) and still carry the bulk of the country's journeys. The two questions are independent (figure 9), and the second belongs to a classical measure of graph theory, betweenness centrality [9].

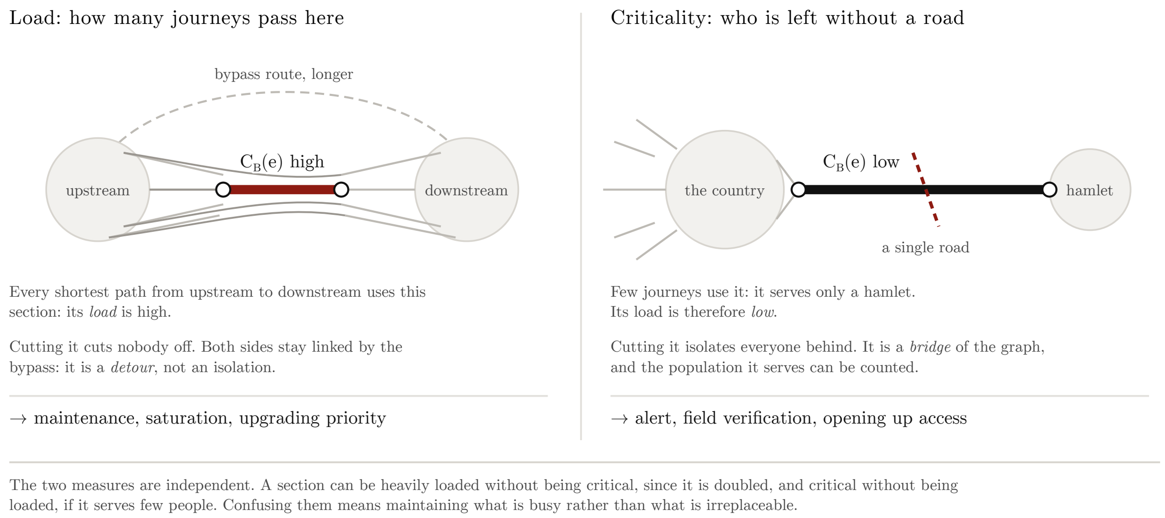

Figure 9. Load and criticality are two independent properties of the same section. On the left, an axis used by all shortest paths: cutting it imposes a detour, it isolates nobody. On the right, a single access road: few journeys use it, and cutting it isolates everything behind. Confusing the two would mean maintaining what is busy rather than what is irreplaceable.

For a section e of the road graph $\mathcal{R}$, writing $\sigma(s,t)$ for the number of shortest paths from s to t and $\sigma(s,t|e)$ for those passing through e:

$$C_B(e) = \Sigma_{s\neq t}\ \sigma(s,t|e)\ /\ \sigma(s,t) \qquad (9)$$

The published layer carries an **unnormalised, integer** variant of this quantity: a count of shortest paths crossing the section, over a set of origin–destination pairs. Across the 34,134 sections and 63,091 km of the network, the distribution is extremely skewed: **half the sections carry the value 1**, the median is 1, the top decile begins at 27 and the maximum reaches 7,111. The top percentile, 1,769 km, consists mostly of primary and secondary roads and expressways: *the measure recovers the road hierarchy without knowing it*, which is the first check one can ask of it.

The screen does not paint this value as it stands. A linear scale would crush 99 % of the network against zero; the layer is therefore coloured on **$\log_{10}$ capped at 4**, which spreads the four orders of magnitude over the ramp. It is a display choice, declared as such: it does not change the measure, it makes it legible.

Three reservations, two of which bear on what the figure allows one to say

(i) The layer is not recomputed by the current pipeline. The values come from a March 2023 extract and no script in the repository reproduces them; the other modules, by contrast, route on the September 2026 network. *An indicator one cannot redo is an indicator one cannot date*, and it ages without anything signalling it. This is the same family as the dating of inventories (§ 5.4), applied to an object we produce ourselves. **(ii) The weight of a section comes from its label, not from a measurement.** Speed is inferred from the class declared in OpenStreetMap: six values, from 30 to 100 km/h. A width surveyed by remote sensing exists and is not used: the median detected width is 11.9 m, and 4.5 % of segments exceed 20 m. *A thirty-metre road is not a nine-metre road, whatever its label says.* **(iii) High betweenness is not criticality.** A heavily loaded axis doubled by another route appears in none of our alerts; a section of low betweenness but sole access does. Confusing the two would mean maintaining what is busy rather than what is irreplaceable, and the converse is equally false. **The platform publishes both layers side by side and names the question each one answers.**

4 Case study: access to healthcare

Sections 4.1 to 4.4 are computed on the network and population of 17 September 2026, with the inventory described in § 3.3. **Every figure in this section therefore carries an inventory in its statement**; none is a measurement of the country.

4.1 Diagnosis

Nationally, **90.7 % of the population lives within one hour** of a facility, 85.6 % within thirty minutes and 97.1 % within two hours. The population beyond one hour is **327,232 people (9.3 %)**, 0.7 % of whom have no road access to a facility at all. It is concentrated (figure 10, table 4): six departments hold 46 % of it.

The first of them is not a field result, and the platform says so itself. Mougoutsi, in Nyanga, has 44,921 inhabitants beyond one hour, that is *all* of its population, because Nyanga carries no facility at all in the inventory. The published artefact records this absence in plain words (“No centre mapped in: Nyanga”), and that is the only reason the figure is legible: *a department at 100 % distant is a statement about the map, not about the country.*

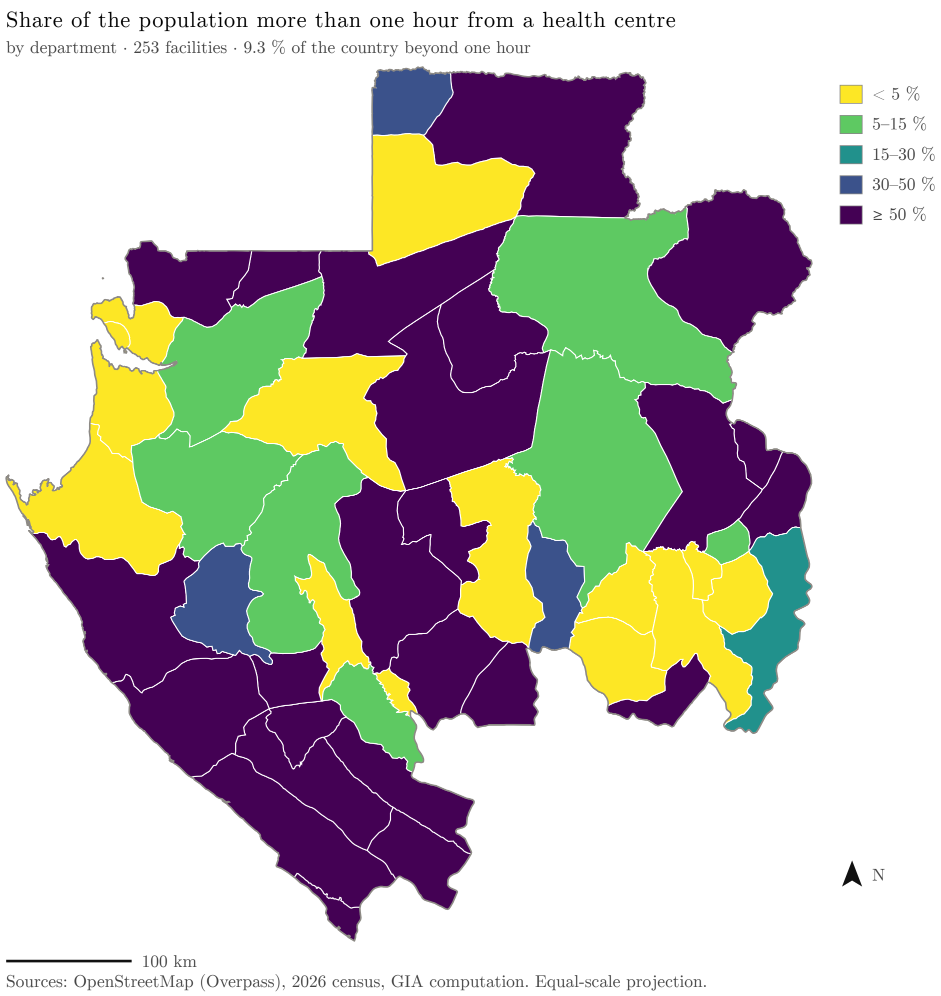

Figure 10. Share of the population more than one hour from a health centre, by department. The departments in dark purple are not without care: they are the ones the inventory does not cover, and the whole of Nyanga is among them (§ 3.3).

Table 4. Departments with the most inhabitants more than one hour from a facility (253 facilities, 17 September 2026). A "> 1 h" column at 100 % signals a province with no mapped facility at all, not a department without care.

Department	Province	Population	> 1 h	People
Mougoutsi	Nyanga	44,921	100.0 %	44,921
Ntem	Woleu-Ntem	61,124	44.4 %	27,139
Okano	Woleu-Ntem	21,262	98.6 %	20,964
Zadie	Ogooué-Ivindo	21,119	95.8 %	20,232
Sébé-Brikolo	Haut-Ogooué	21,987	90.9 %	19,986
Haut-Ntem	Woleu-Ntem	21,009	89.0 %	18,698

By province, Woleu-Ntem has the most distant inhabitants (74,393, 30.7 %), ahead of Nyanga (68,215, 100 %) and Ogooué-Ivindo (45,981, 49.2 %); Estuaire, which holds 62 % of the population, has only 1.1 %. **This ranking measures two things at once, and the figure does not separate them:** the real distance of inhabitants, and the uneven thickness of volunteered mapping, which puts 218 of its 253 points in Estuaire alone. A rural department at 90 % distant says "we have mapped nothing here" as loudly as it says "there is nothing here". The diagnosis nevertheless remains usable for what it is: a *checklist*, ordered by the number of inhabitants concerned, of what must be looked at on the ground.

4.2 Prioritisation

A first well-placed centre would cover **52,221 people**, that is 16.0 % of the population beyond one hour; five centres would cover **150,811, or 46.1 %** (table 5). The 5 % tolerance plateau of the first step holds 124 cells in four distinct zones, about 950 km^2: the model designates *regions*, not a point. Returns fall, but slowly — 52,221, 35,338, 22,142, 21,628 then 19,483 people — because distance here bears on broad rural areas rather than on a few residual pockets. Overlap between centres is nil at all five steps, so the plan can be undertaken in any order without loss.

And the first step falls in Nyanga. Basse-Banio, near Milolo, is designated because the province carries no facility in the inventory. *This is not a siting recommendation, it is the designation of an area to verify*, and that is how the platform presents it. The four following steps — Ntem, Okano, Sébé-Brikolo, Zadie — bear on rural departments whose remoteness is contradicted by no other source, and they are directly usable.

Table 5. Greedy five-centre plan (253 facilities, network and population of 17 September 2026). Shares are given against U = 327,232 people. The first step is to be read with § 3.3: Nyanga carries no facility in this inventory.

Step	Department and location	Gain	Cumulative	Cumulative share
1	Basse-Banio, Nyanga, Milolo at 1.7 km	52,221	52,221	16.0 %
2	Ntem, Woleu-Ntem, Méloh at 5.6 km	35,338	87,559	26.8 %
3	Okano, Woleu-Ntem, Essong-Ville at 3 km	22,142	109,700	33.5 %
4	Sébé-Brikolo, Haut-Ogooué, Makatamangoye II at 4.3 km	21,628	131,328	40.1 %
5	Zadie, Ogooué-Ivindo, Mékambo at 1.7 km	19,483	150,811	46.1 %

4.3 Vigilance

The population behind a single road link reaches **Π = 1,606,766 people (45.7 %)**, spread over 1,375 weak points arising from 7,075 critical links (figure 10). The first point isolates **27,542 people** on a secondary road in Estuaire, the second 26,640 on an expressway in the same area. In an urban setting, a street that appears irreplaceable often signals a connection missing from the map, which calls for verification before any maintenance decision.

The number of weak points and the population behind them do not move together, which is why the platform publishes the second and not the first: a facility added on a dead end cuts the detached part into independent branches, so the points multiply while the isolated population falls. A dashboard counting points would read a deterioration where the situation improves — the defect described in § 5.1.

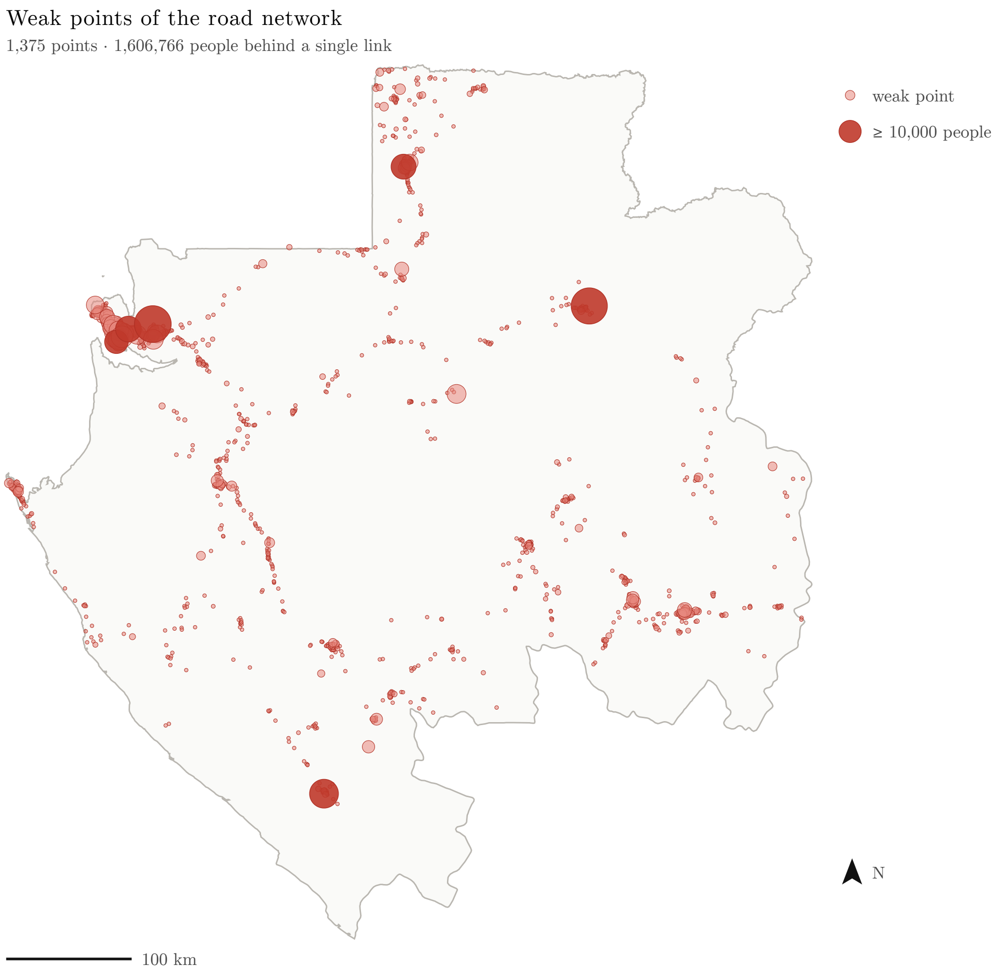

Sources: OpenStreetMap (September 2026), GIA computation. Each circle's area is proportional to the isolated population.

Figure 11. Weak points of the road network. Each circle's area is proportional to the population that would lose road access to every facility if the link were cut.

4.4 From indicators to decisions

Table 6. Decision questions, indicators and case-study results.

Kind	Decision-maker's question	Indicator	Result	Possible action
Diagnosis	Which areas are far from care?	U, $A_{60}(d)$	Mougoutsi: 44,921 pers. (100 %)	Check the map first, then the ground
Prioritisation	Where would a centre serve most?	G_k, $\mathcal{P}_\varepsilon$	Basse-Banio: 52,221 pers.	Feasibility study over the plateau
Prioritisation	How many centres for what effect?	greedy plan	5 centres: 46.1 % of the uncovered	Size a programme
Vigilance	Which roads to protect?	Π, weak points	27,542 pers. behind one section	Verification, priority maintenance
Vigilance	Which works to follow?	Rules (7) and delay	Work site near a critical link	Close monitoring

4.5 Through routes: where the load concentrates

The computation of § 3.10, applied to the 34,134 sections and 63,091 km of the network, yields a concentration the decision-maker can use directly (table 7, figure 12). **One per cent of the kilometres carries 39 % of the load, five per cent carries 78 %**, and half the network carries 98 %: in other words, half the kilometres serve almost no through journey at all.

Table 7. Concentration of load: share of the most heavily loaded kilometres and share of total load they carry. Total load here is the *sum of passages* over sections, unweighted by their length; weighted by kilometres the concentration is sharper still (5 % of the kilometres then carry 83 % of the load).

Share of network (km)	Kilometres	Share of load carried
1 %	1,769	39 %
5 %	3,155	78 %
10 %	6,309	88 %
25 %	15,773	96 %
50 %	31,546	98 %

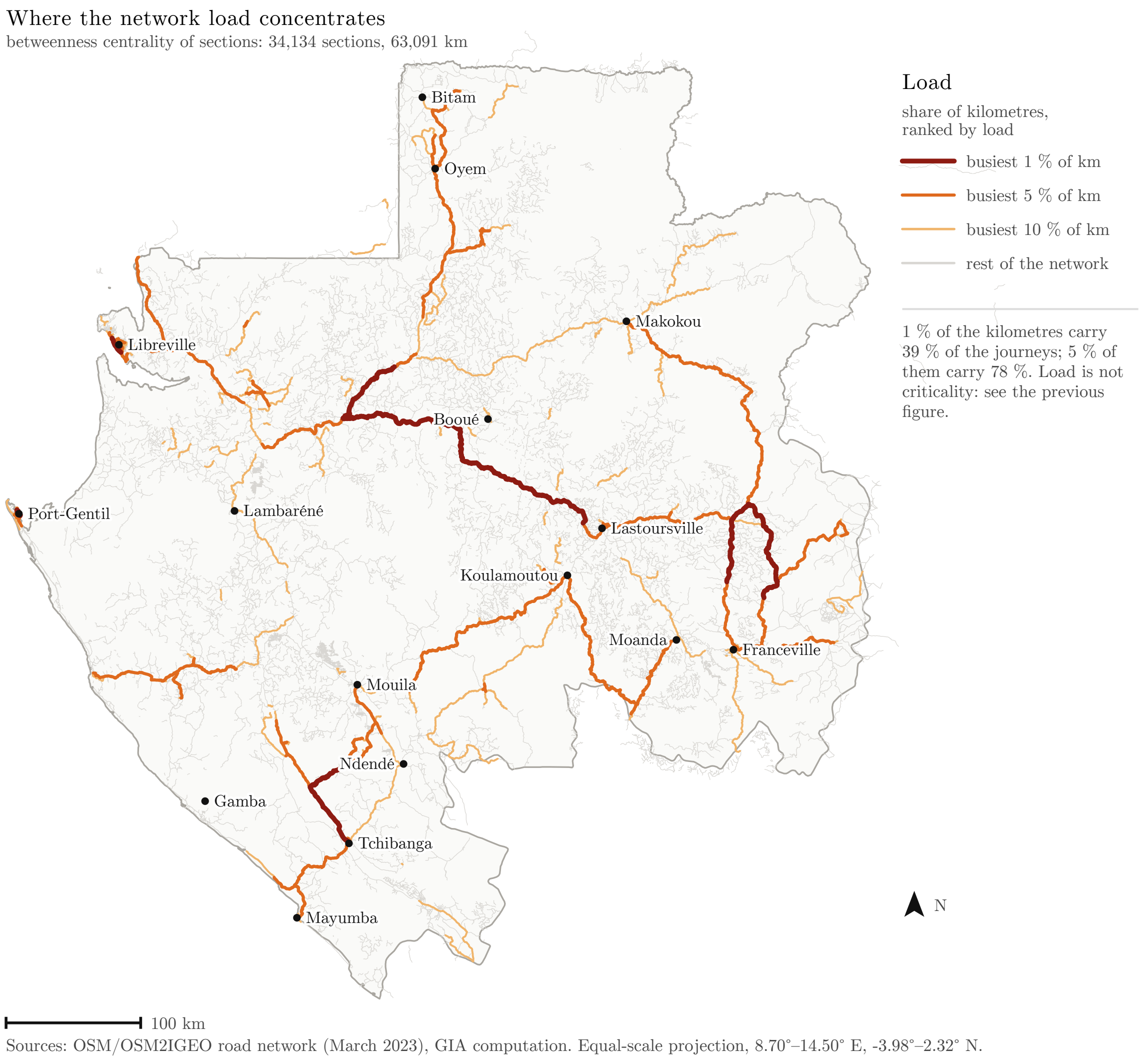

Figure 12. The through routes of the Gabonese network. In dark red, the most heavily loaded percentile, 1,769 km carrying 39 % of the load; in orange the most heavily loaded 5 %, in amber the 10 %, in grey the rest of the network. The main corridor links Estuaire to Haut-Ogooué through Booué and Lastoursville; it is doubled by an axis towards the south-west, from Mouila to Tchibanga. The background is the *land* outline of the country, obtained by cancelling the shared boundaries of the 48 departments: the provincial relations of OpenStreetMap carry territorial waters, and gave the country a straight coastline that does not exist. The projection preserves proportions: one scale on both axes.

The measure recovers the road hierarchy while knowing nothing of it: motorways and expressways carry median loads of 544 and 224, primary roads 162, while the 13,018 unsurfaced tracks, which nonetheless make up 42,776 km, two thirds of the network, have a median of 2 (table 8). **This is the first check one can ask of a topological indicator: to recover, without the information, what the information already says.**

Table 8. Load by road class. The median clearly separates the structuring classes from the rest, although the computation ignores class.

Declared class	Sections	Kilometres	Median load	Maximum
Motorway	12	20	544	700
Expressway	309	1,112	224	6,555
Primary road	441	2,626	162	7,111
Secondary road	465	1,668	129	3,481
Tertiary road	680	1,745	29	3,151
Unclassified road	2,488	8,846	4	4,122
Unsurfaced track	13,018	42,776	2	777
Residential road	11,191	3,110	1	4,396

What these two tables allow, and what they do not. They allow a maintenance programme to be sized: 1,769 km, less than a twentieth of the network, concentrate the largest share of through journeys, and a euro of maintenance there touches more journeys than elsewhere. They do not allow any conclusion about isolation: the "maximum" column shows that modest classes sometimes carry very high loads, because they are the only ones linking two parts of the country, and it is fragility (§ 4.3), not load, that says who would be left without a road. *The two indicators are read together, never one in place of the other.*

5 Reliability of the decision support

5.1 Counting in people rather than in objects

The number of maximal critical links rises by 7.7 % while fragility falls. At constant scope, critical links decrease (14,550 → 14,072) but maximal links increase (6,999 → 7,417): a facility added on a dead end cuts the detached part into independent branches (figure 7), and widening the scope to newly served components adds 118 links. A dashboard displaying the number of links as a measure of fragility would signal a deterioration where the situation improves. The platform therefore presents fragility in population.

5.2 Thresholds and tolerances

Rule (6) flags only Offoué-Onoye, although Bayi-Brikolo has more than twice as many distant inhabitants: with 32 % of its population within one hour, it stays above the 25 % threshold. A threshold expressed as a proportion favours small, very isolated territories; a threshold expressed as a headcount would favour populous ones. That choice belongs to the decision-maker, and the platform must make it explicit by listing what each threshold excludes. Likewise, the tolerance plateau of the plan's first step holds **124 cells across four distinct zones**, about 950 km^2: presenting a single site would claim a precision that neither the grid nor the positions of the facilities can carry, hence the presentation as a plateau.

5.3 External validation protocol for positions

Adjectives ("at the building", "at the village") cannot be subtracted from a figure; kilometres can. We propose the smallest protocol that turns the former into the latter.

Sample. Thirty points per geocoding method, that is 150, drawn at random from the whole method and not from among the suspect points: a sample drawn on the very criterion one wants to measure measures the criterion, not the population.

Prior publication. The seed of the draw and the list of the 150 points are published before verification, so that no point can be set aside afterwards in the light of its result.

Instrument. High-resolution imagery, from a provider that is none of the inventory's sources: verifying one source by another source of the same inventory is circular.

Output. For each method: median positional error, upper decile, maximum; and the number of facilities not found, which is not waste but a result. A facility absent from recent imagery may have closed, information we have today no means of obtaining.

This protocol would allow three things that are impossible today: publishing "n centres whose position is accurate to within so many kilometres" instead of an adjective; publishing the access indicator with an error bar obtained by displacing each centre by the error of its class; and checking the classification itself, which for now rests on the authors' description and on a distribution of distances, not on a measurement.

5.4 Principles that follow

(i) A published indicator is dated by its data, not by the file that carries them, and any comparison over time separates the evolution of the map from that of the ground. **(ii)** A property that only the producer knows is declared alongside the data and must survive to the screen: the precision of a position, the refresh rate of a layer, the window of a rule. **(iii)** Count-based indicators are subordinate to population-based indicators. **(iv)** Decision thresholds are explicit, and what they exclude is visible. **(v)** An absence of measurement is distinguished from a null measurement: "not evaluated", "not distinguished" and "zero" are three different states, and confusing them makes a silence read as a success.

6 Features to develop

6.1 Features of the previous version to restore

The previous version of the platform offered features that have not yet been restored in the current one. Table 12 ties them to the decision they serve. Four of them are under development.

Table 9. Features of the previous version to restore.

Feature	Decision served	Status
Population and buildings within a drawn area	Local diagnosis, sizing a project	in progress
Place search	Access to every diagnosis	in progress
Acknowledging and re-raising alerts	Vigilance	in progress
Shareable link to a view	Consultation, arbitration	in progress
Isochrones from a point	Diagnosis of a site	to restore
Comparison of two layers by swipe	Diagnosis, before/after	to restore
Layer of map changes	Reliability (what changed in the sources)	to restore
Entering and editing works, discussion	Project monitoring	to restore
Active fires, press tone, trends	Vigilance	to restore
Most critical roads by centrality, density by region	Maintenance prioritisation	to restore
Live map queries (mines, flood zones, services)	Hazard diagnosis	to restore

Restoring is not simply copying. The previous version estimated the population of a drawn area in proportion to the department's surface; the current version will use the per-cell population allocated by buildings (equation 1), and will report the share of the area covered by partially included cells.

6.2 New features

The reliability analysis suggests six developments. **(i) History of results**: keeping every computation with the version of the data that produced it. **(ii) Propagation of positional uncertainty**: drawing the symbol at the scale of what the source knows, and recomputing the indicator under displacement of the doubtful positions once the protocol of § 5.6 has been run. **(iii) Extension to other services**: since the modules depend on the inventory only through $\mathscr{I}$, access to schools, water points or markets is computed with the same tools, provided the completeness of those inventories is documented. **(iv) Scenarios**: simulating the cutting of a section, the closing or opening of a facility. **(v) Alert subscription by territory. (vi) Exportable decision sheets**, carrying a result together with its assumptions, its sources and its range.

7 Conclusion

Q1. Integration. A base of four shared objects (located population, road graph, service inventories, register of works) is enough to feed modules for diagnosis, prioritisation and vigilance, and these modules depend on a kind of service only through its inventory.

Q2. Translation into decisions. On access to healthcare, the platform locates the 327,232 people more than one hour from a centre (9.3 %), orders departments by the number of inhabitants concerned, quantifies what a five-centre programme would bring — 150,811 people covered, 46.1 % of the uncovered — and identifies the roads on which 45.7 % of the population depends.

Q3. Reliability. These indicators are worth only what the coverage of the inventory that produces them is worth, and that coverage is uneven: 218 of the 253 facilities are in one province, and one whole province carries none. **The first location the model designates falls precisely in that province**, and it is an artefact of the map, not a field result. The platform answers with provenance, assumptions declared together with the direction of their bias, ranges, plateaus rather than points, measurement in people rather than in objects, the preservation of positional precision all the way to the screen, and the recording in plain words of what the inventory does not cover.

Limitations. The main result depends on an inventory that covers rural Gabon poorly, and none of its figures should be read as a measure of access to healthcare: it is a measure of what the map carries. Positions have been verified neither in the field nor against independent imagery: the protocol of § 5.6 has not yet been run, and until it is, positional uncertainty remains declared but unquantified. The road network is assumed complete and travelled at constant speed. Listed facilities are assumed open and equivalent, whereas the national health development plan indicates that several dispensaries are not functional. Finally, the platform's usefulness to decision-makers has not yet been evaluated; that will be the subject of later work, alongside the developments described in section 6.

References

1. R. Church and C. ReVelle. The maximal covering location problem. *Papers of the Regional Science Association*, 32(1): 101–118, 1974.
2. P. J. Densham. Spatial decision support systems. In D. J. Maguire, M. F. Goodchild and D. W. Rhind (eds.), *Geographical Information Systems: Principles and Applications*, vol. 1, pp. 403–412. Longman, 1991.
3. E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1: 269–271, 1959.
4. J. Mennis. Generating surface models of population using dasymetric mapping. *The Professional Geographer*, 55(1): 31–42, 2003.
5. G. L. Nemhauser, L. A. Wolsey and M. L. Fisher. An analysis of approximations for maximizing submodular set functions — I. *Mathematical Programming*, 14: 265–294, 1978.
6. W. Sirko, S. Kashubin, M. Ritter *et al.* Continental-scale building detection from high resolution satellite imagery. arXiv:2107.12283, 2021.
7. R. E. Tarjan. A note on finding the bridges of a graph. *Information Processing Letters*, 2(6): 160–161, 1974.
8. D. J. Weiss, A. Nelson, C. A. Vargas-Ruiz *et al.* Global maps of travel time to healthcare facilities. *Nature Medicine*, 26: 1835–1838, 2020.
9. L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1): 35–41, 1977.