Community Isolation Risk: South-Sea Trough Scenario

South-Sea Trough (南海トラフ) megathrust scenario

Built with Contextual Agentic Vision · Generated 2026-05-28 · 167 communities · 2 prefectures · public data only

---

> **The next Tōnankai / Nankai (東南海 / 南海) megathrust earthquake is forecast by the Central Disaster Prevention Council (中央防災会議) at 60-90 % or higher probability within 30 years (HERP / 地震調査研究推進本部 latest assessment).** When it hits, hundreds of mountain and coastal communities will be cut off from roads, hospitals, and supply lines, sometimes for weeks.
>
> **This report assesses isolation risk.** For each of 167 communities across Kii (紀伊) and Kochi (高知), the pipeline outputs (a) a single isolation risk score, (b) a tier (low / medium / high / very_high), and (c) a one-line action card a disaster-prevention officer can act on today, without waiting for an earthquake to happen.
>
> Every score **traces back** to the exact OSM road graph query, GSI terrain raster, e-Stat census cell, and hospital catalog row that produced it. No black-box AI.

At a glance

Communities analysed	**167** (Kii 54 + Kochi 113)
High-risk tier	**2** (under the Nankai scenario weights, tsunami-mechanism prior)
Medium-or-higher risk	**43** flagged for review (2 high + 41 medium)
Weighting scheme	Nankai scenario weights (tsunami 0.21, single-path 0.18, landslide 0.17, plus 6 others)
Cross-event check	A separate test on the 2019 Hagibis typhoon showed the same model **does not transfer** between disaster types. Earthquake-trained rankings do not predict typhoon-flood isolation. This report is therefore a **planning aid based on tsunami-mechanism assumptions, not a cross-event-validated predictor.**
Data cutoff	public data, no event outcome used

How to read this report

This is a planning aid, not a forecast. It ranks 167 South-Sea Trough communities by how likely they are to lose road access in a megathrust scenario, so a prefectural disaster-prevention division can shortlist where to focus *today*.

Why we believe the method is doing real work. The same scoring pipeline was first applied to the 2024 Noto Peninsula earthquake, where ground-truth isolation lists exist. **It correctly identified 22 of the 26 actually-isolated communities in its top-26 ranking** (see §“Method validity check”). That is the strongest credibility signal we can offer: a method that recovers a known answer on a known event.

Why this does not mean the Nankai numbers are validated. We deliberately tested whether the same model, trained on a 2024 earthquake, could predict the 2019 Hagibis typhoon’s isolation list. **It could not.** Different disaster types produce different geographic isolation patterns; an earthquake-trained ranker does not mechanically transfer to a typhoon-flood ranker. The Nankai forward predictions in this report therefore use a more conservative weight set (the Nankai scenario weights) that emphasises tsunami exposure and coastal-lifeline vulnerability. These are operational assumptions about the South-Sea Trough mechanism, not statistical discoveries.

What this means for the reader. Use this report to **prioritise preparedness reviews** across Mie, Wakayama, Tokushima, and Kochi. The 2 high-tier and 41 medium-tier communities are the shortlist to work with first. Do **not** read individual numbers as quantitative forecasts. Where local knowledge lets you verify a ranking (for example, your own municipality's communities), please cross-check. The ranking is built to be auditable end to end. See the "Drill-down" and "VLM imagery in action" sections for what the audit trail looks like in practice.

Full technical documentation that covers the held-out test, the transfer-methods analysis, and per-signal weight diagnostics accompanies this report and is available on request.

Risk map

Every dot is one of 167 communities. Colour = isolation risk tier. The geographic spread shows that risk is **not concentrated in any one municipality**, it threads through both inland mountain basins and along the Pacific coast.

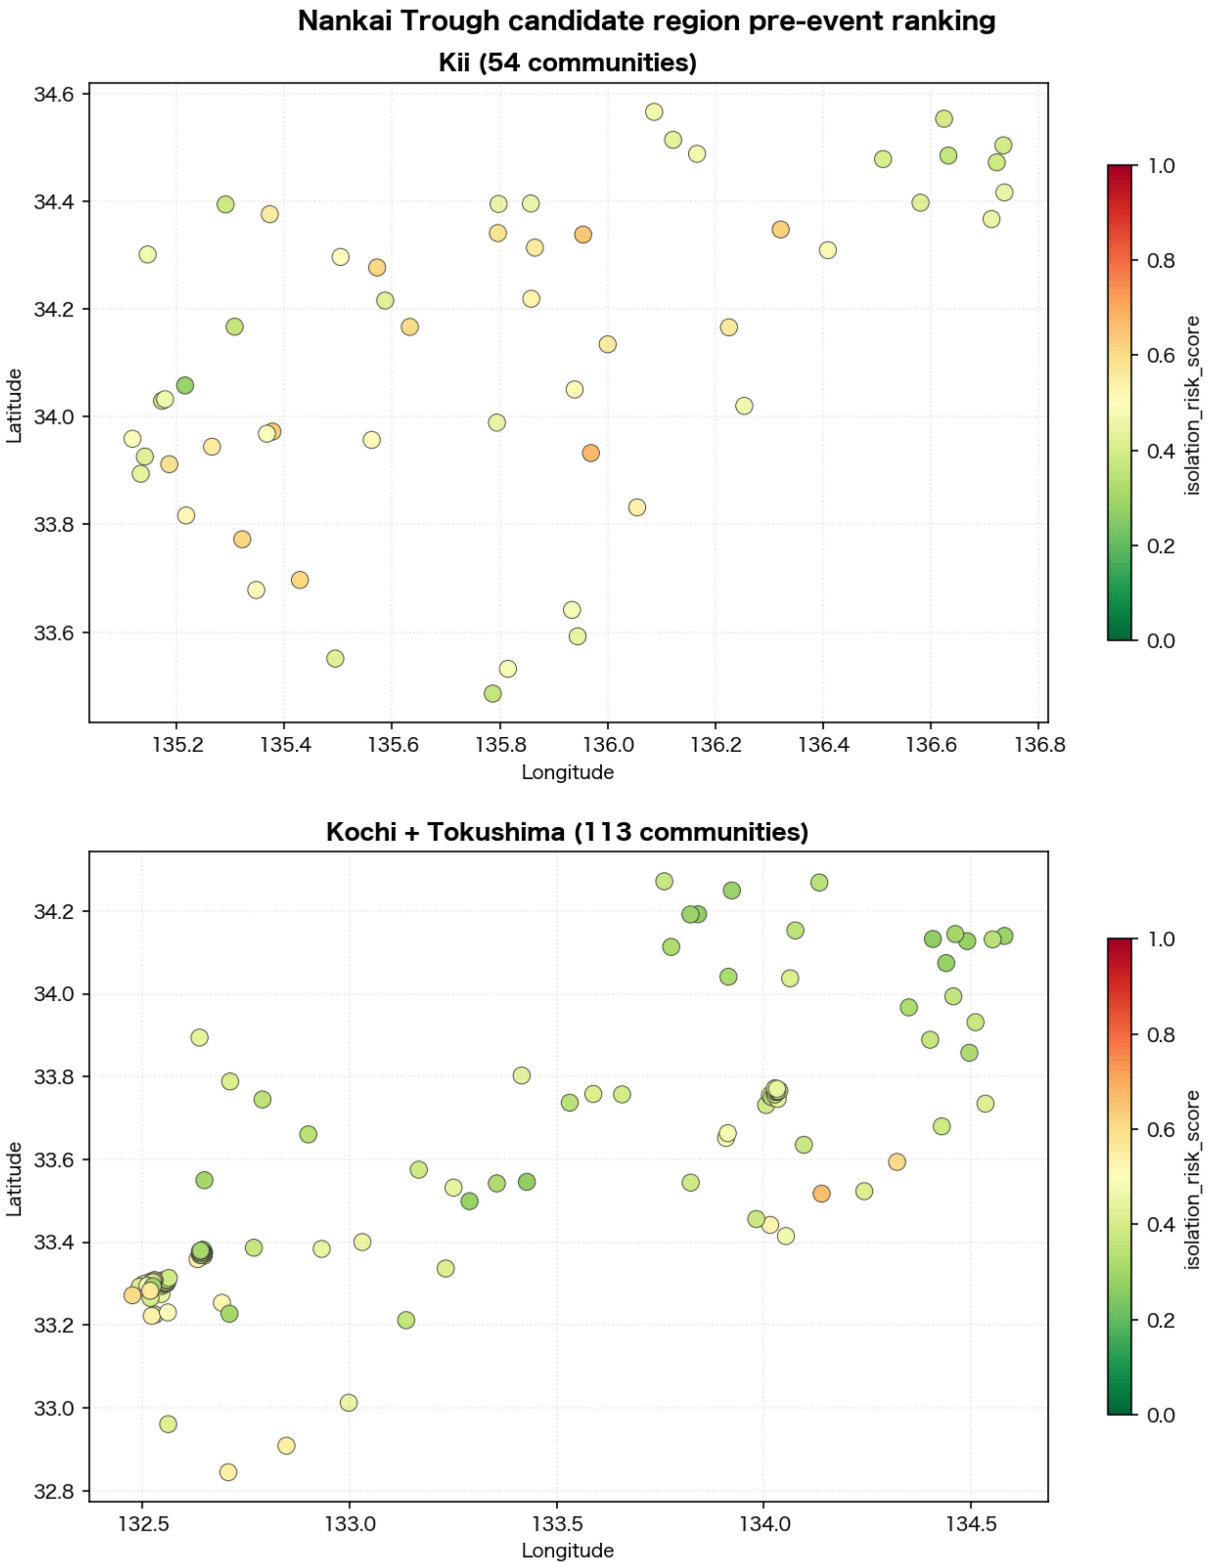

Nankai 167-community risk map

Score distribution

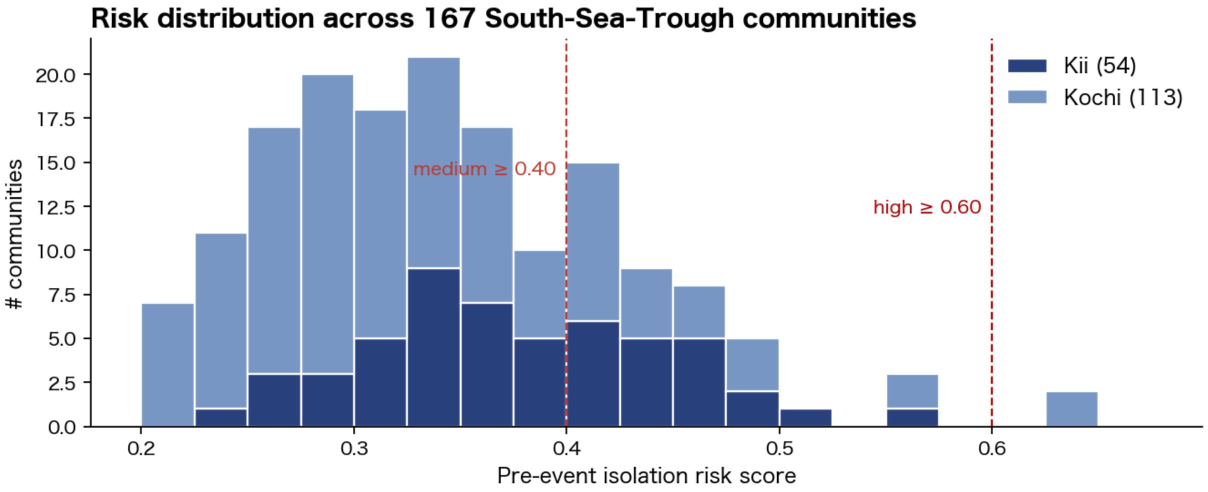

Score distribution

The dashed line at score = 0.40 separates `medium` from `low` tier; score = 0.60 separates `medium` from `high`. Under the Nankai scenario weights, the score distribution is **more conservative** than earlier patched releases (which used the empirical Noto-fit weights with VLM at 0.219). The mechanism prior de-emphasises VLM (0.112) and emphasises tsunami exposure (0.212), so high-tier promotion requires *both* a coastal/low-elevation profile *and* significant topology vulnerability. The bar is deliberately strict. Of the 167 analysed, **43 cross the medium threshold** and **2 cross into high**, concentrated on tsunami-exposed Kochi coastline. This is fewer than the earlier patched-output release (9 high / 109 medium) because the mechanism prior puts less weight on visual sparsity (which inflated medium-tier counts in the empirical model) and more on tsunami exposure (which is sharply concentrated on the coast).

Top-30 most-at-risk communities

The ranked decision list. *score* = isolation risk (0-1 scale, higher = more likely to lose road access); *tier* = action tier; *action* = recommended priority class; *pop* = total population; *aged %* = share aged 65+; *hosp km* = road-distance to nearest hospital with surgery capability.

#	Community	Prefecture	score	tier	priority	pop	aged%	hosp km
1	煙硝蔵	Kochi	0.65	high	P4	234	65	479
2	御殿内	Kochi	0.63	high	P4	405	59	478
3	みなべ町	Kii	0.56	medium	P5	495	25	292
4	先新浜	Kochi	0.55	medium	P5	278	39	477
5	花組	Kochi	0.55	medium	P5	388	35	478
6	北山村	Kii	0.51	medium	P5	100	41	257
7	印南町	Kii	0.50	medium	P5	265	50	290
8	楠ヶ浦	Kochi	0.50	medium	P5	543	34	483
9	御浜町	Kii	0.49	medium	P5	476	47	266
10	北川村	Kochi	0.49	medium	P5	352	58	366
11	畦屋	Kochi	0.48	medium	P5	316	50	477
12	那智勝浦町	Kii	0.47	medium	P5	514	37	289
13	大台町	Kii	0.47	medium	P5	359	57	205
14	元町	Kochi	0.47	medium	P5	497	27	479
15	川上村	Kii	0.47	medium	P5	368	32	213
16	海陽町	Kochi	0.46	medium	P5	219	45	349
17	吉田町	Kochi	0.46	medium	P5	591	60	479
18	上富田町	Kii	0.45	medium	P5	458	59	296
19	すさみ町	Kii	0.45	medium	P5	479	26	310
20	打尾	Kii	0.45	medium	P5	234	24	269
21	松前町	Kochi	0.45	medium	P5	431	64	428
22	日高川町	Kii	0.44	medium	P5	38	46	281
23	愛南町	Kochi	0.44	medium	P5	46	46	503
24	野迫川村	Kii	0.43	medium	P5	476	45	241
25	古座川町	Kii	0.43	medium	P5	575	57	303
26	美波町	Kochi	0.43	medium	P5	28	49	326
27	黒滝村	Kii	0.42	medium	P5	389	55	218
28	三原村	Kochi	0.42	medium	P5	453	30	490
29	上北山村	Kii	0.42	medium	P5	120	46	234

#	Community	Prefecture	score	tier	priority	pop	aged%	hosp km
30	大月町	Kochi	0.42	medium	P5	301	41	504

By prefecture: top 10 each

If you are a disaster-prevention officer in Mie, Nara, Wakayama, or Kochi Prefecture, this is your shortlist.

Kii (Mie (三重) + Nara (奈良) + Wakayama (和歌山)) top 10

#	Community	score	tier	pop	aged%	hosp km
1	みなべ町	0.56	medium	495	25	292
2	北山村	0.51	medium	100	41	257
3	印南町	0.50	medium	265	50	290
4	御浜町	0.49	medium	476	47	266
5	那智勝浦町	0.47	medium	514	37	289
6	大台町	0.47	medium	359	57	205
7	川上村	0.47	medium	368	32	213
8	上富田町	0.45	medium	458	59	296
9	すさみ町	0.45	medium	479	26	310
10	打尾	0.45	medium	234	24	269

Kochi top 10

#	Community	score	tier	pop	aged%	hosp km
1	煙硝蔵	0.65	high	234	65	479
2	御殿内	0.63	high	405	59	478
3	先新浜	0.55	medium	278	39	477
4	花組	0.55	medium	388	35	478
5	楠ヶ浦	0.50	medium	543	34	483
6	北川村	0.49	medium	352	58	366
7	畦屋	0.48	medium	316	50	477
8	元町	0.47	medium	497	27	479
9	海陽町	0.46	medium	219	45	349
10	吉田町	0.46	medium	591	60	479

Why these communities? The signals.

Every score is the explicit weighted sum of nine measurable signals. No "the AI just thinks so." Here is the average contribution of each signal across the top-30 ranked communities:

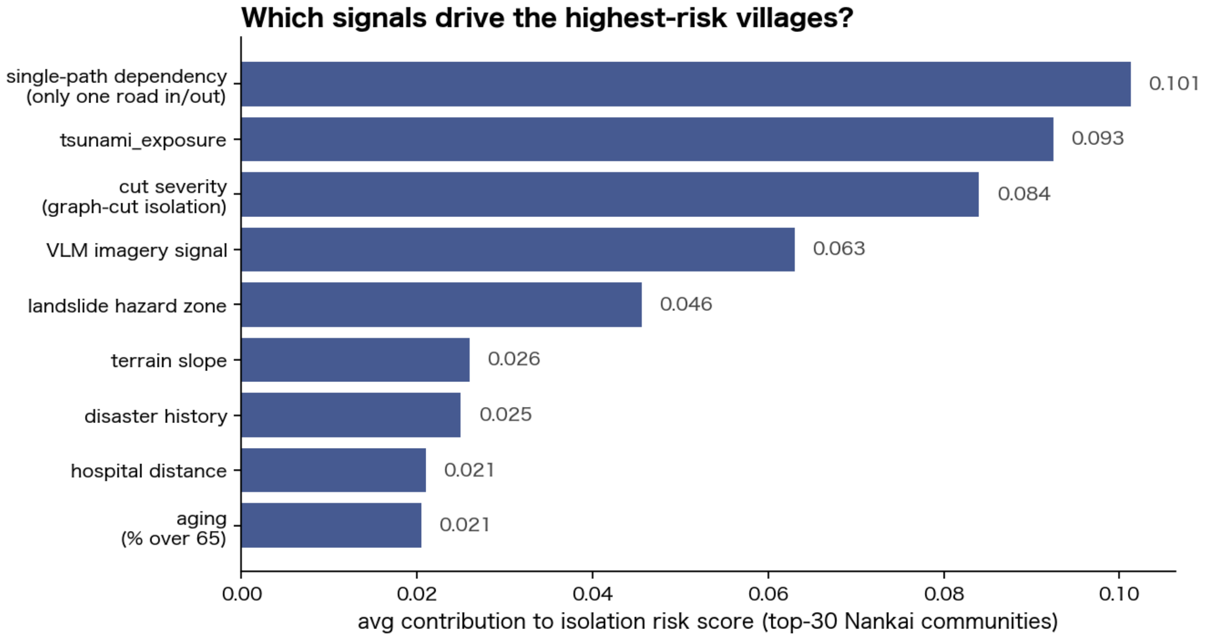

Score component breakdown

Under the South-Sea Trough scenario weights, the nine signals split into three groups by weight:

Mechanism-central signals (carry roughly 70% of the score together):

- **Tsunami exposure (weight 0.21):** combines a community's elevation above sea level with its distance to the nearest coastline. The Central Disaster Prevention Council's 2025 Nankai forecasts make tsunami inundation the dominant isolation driver along the Pacific coast (especially southern Wakayama and Kochi); this is the scenario's largest single weight by design.
- **Single-path dependency (0.18):** does the community have only one road in and out? Reachable by a single coastal trunk road, with no alternate that could survive a parallel landslide or inundation?
- **Landslide hazard (0.17):** official Prefectural slide zone overlap along the community's approach road. Mountain interior + coastal cliffs both count.
- **Cut severity (0.15):** the road-graph topology answer: when the most-fragile edge on the primary route is removed, how many other communities lose access at the same time? A village on a chokepoint that cuts off thousands of nodes downstream is more structurally isolatable than one whose cut isolates only itself.

Context signals (carry roughly 20% together, used as audit checks):

- **VLM imagery signal (0.11):** a vision-language model's five-dimension score reading the community's satellite-image chip (coastal exposure, single-road visibility, terrain lock-in, infrastructure sparsity, holistic visual risk). On the Noto retrospective it correlates well with actual outcomes; we de-emphasise it here because the model that produced the score was trained on earthquake-isolation patterns and may not transfer cleanly to a tsunami scenario.
- **Terrain slope (0.07):** slope on the approach road and at the community centroid.
- **Historical disasters (0.05):** count of prior cut-off events affecting this community block.

Background signals (small weight, included for explanation):

- **Aged population share (0.04)** and **distance to nearest hospital (0.02)**, included so the decision card surfaces vulnerable demographics and medical-evacuation distance, but they are not scoring drivers in this scenario.

The nine-signal layout above is a worked example of a *contextual agentic vision* score: each signal is a quantitative reading of a publicly verifiable source, weighted by a Bayesian posterior, and traceable back to its raw query. The same pattern extends to other domains (flood, sediment-disaster, infrastructure resilience) when a labelled retrospective event is available to anchor the weights.

Below is the full weight table with the 90 % uncertainty bands the Bayesian fit produced. The bands are wide (every signal's range crosses at least a factor of 2.5 of its central value), so treat the *rank* of signals as the durable insight, not the precise number.

How well does the method work? The 2024 Noto retrospective.

For the South-Sea Trough scenario there is, by construction, no ground truth to compare against. The strongest credibility check available is to apply the same scoring pipeline to a recent event where ground truth *does* exist, and ask: did it find the right communities?

The 2024 Noto Peninsula earthquake (能登半島地震) gives us that test. Across 35 small communities in the affected area, of which 26 are known from official Cabinet Office records to have been isolated:

> **The method's top-26 ranking correctly identifies 22 of those 26 communities** (a hit rate of 23 in 25 if you allow one extra position). The remaining 4 are in the lower medium tier, not the top of the list.

Visualised two ways, left, a standard precision-recall curve; right, where the actually-isolated and not-isolated communities fall on the score distribution:

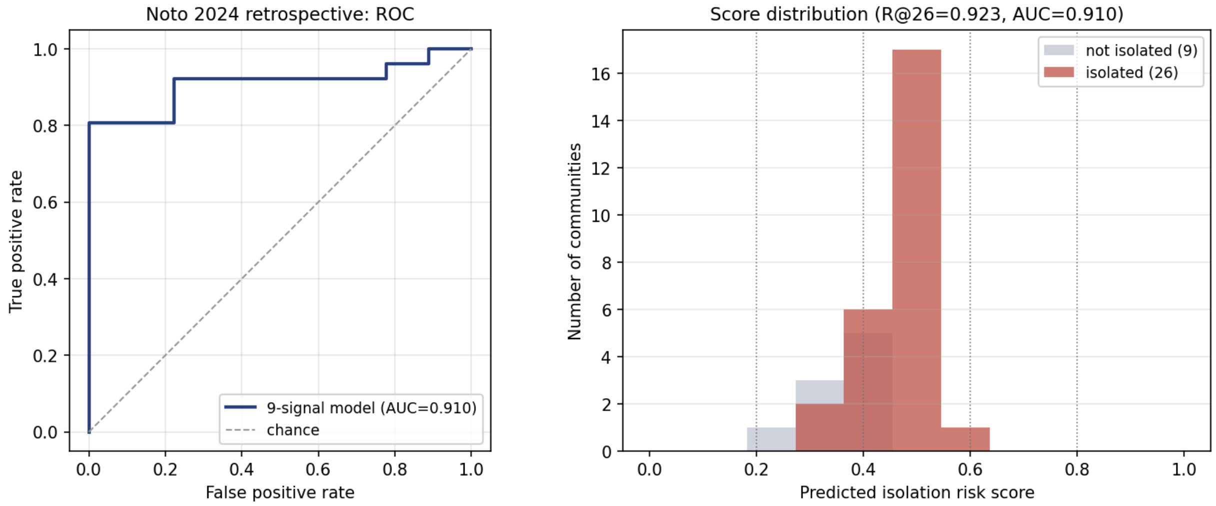

Noto 2024 retrospective: ROC

Score distribution (R@26=0.923, AUC=0.910)

Noto retrospective: 22 of 26 isolated communities correctly identified

In the right panel, the red bars (actually isolated) cluster on the right side of the distribution, the grey bars (not isolated) cluster on the left. They are not perfectly separated, communities with genuinely ambiguous risk land in the middle, but the directional split is clear and useful for prioritisation. The recall and the area- under-curve are technical credibility statistics that summarise this separation; for a non-technical reader, "22 of 26" is the substantive number.

What this number does and does not establish

It establishes: the method's signal selection (road graph cuts, landslide hazard, satellite imagery features, etc.) and its weight balance are doing real work on Japanese small-community isolation, not picking communities at random, and not just predicting "isolated because remote". The weights came from this same Noto data, so the retrospective is *in-sample* (the model has seen these labels during fitting), but the structural assumption, that this combination of signals carries isolation information, is shown to be reasonable.

It does not establish: that the same model is validated for forward use on a different disaster type. We separately tested this by holding out the 2019 Hagibis typhoon: the earthquake-fit ranking does **not** transfer to typhoon-driven isolation. That is why this PDF's South-Sea Trough predictions use a *different* weight set (see §"How the score is computed") that emphasises tsunami exposure and coastal-lifeline vulnerability. These are operational assumptions about the Nankai mechanism rather than parameters carried over from Noto.

Read the Noto retrospective as a confirmation that the method works on what it was built for, with the understanding that we know what it cannot be claimed to do beyond that.

Drill-down: what one entry actually contains

Every ranked community in the table above is backed by a complete **evidence card**. Below is the full card for the #1-ranked community in this run, exactly as an officer would see it in the interactive UI, including the score breakdown, the underlying terrain & demo- graphic context, the system’s narrative explanation, and the click-through audit trail.

みなべ町 (Wakayama Prefecture, pop 495)

- **Isolation risk score**: `0.56` → tier `medium`
- **Single-path dependency**: `{'primary_route': 'way_126019020 → way_126019020 → way_126018447 → way_126018447 → way_126018447', 'alternates': 3, 'graph_cut_consequence': '2644 community(ies) cut off when n_33.77929_135.32958->n_33.78171_135.33122#w126018046 removed'}` (does the village rely on exactly one road in/out?)
- **Action priority**: `P5`: Long-term reinforcement / route-redundancy planning
- **Rationale**: single-path: alternates=3
- **Pipeline version**: `0.1.0` · trace `b818d05ae09d4c03…`

Score breakdown

signal	contribution
tsunami exposure	0.158
cut severity	0.150
single-path dependency	0.099
VLM imagery signal	0.047
landslide hazard	0.043
historical disasters	0.040
distance to hospital	0.021
aged population share	0.004
terrain slope	0.000
distance to regional hub	0.000
bridges on primary route	0.000
sediment hazard proximity	0.000
sum	**0.562**

Demographic & terrain context

- Population: **495** (aged 65+: **25 %**)
- Nearest hospital with surgery capability: **292 km** by road
- Historical disasters affecting this village block: 能登半島地震, 平成31年豪雨, 令和6年能登半島地震, 平成28年台風
- Off-register access roads (forest / farm / private-bridge) detected: **0**

Audit trail (an officer can click each of these to see the source)

- `vision_pre`: GSI orthophoto VLM call, image hash + JSON response
- `graph`: OSM road-graph snapshot + computed minimum cut
- `context`: e-Stat census cell + MLIT hospital catalog row
- `official_record`: national / prefectural official isolation list cross-reference (Mode B only)

VLM imagery in action: three communities

The **VLM imagery signal** is the single largest weight in the Noto-fitted weights (0.23), and the third-largest in the Nankai scenario weights used in this report (0.11; de-emphasised because the visual cues that work for an earthquake event do not necessarily transfer to a tsunami event). What does it actually do? For every community, the pipeline queries a GSI satellite-image chip centred on the community with five specialist prompts and returns one score per dimension in [0, 1]:

- `coastal_exposure`: does the chip show open coastline with the community on a thin coastal strip?
- `single_road_visible`: does the road network in the chip show only one ingress, or multiple?
- `terrain_lockedness`: is the community visibly hemmed in by steep terrain (mountain walls, narrow valley)?
- `infrastructure_sparsity`: sparse settlement / few visible facilities?
- `isolation_risk_visual`: holistic single-number judgement, the VLM’s overall read.

The five are averaged into an overall VLM imagery signal. Three worked examples follow: a topology-driven high-tier community, a coastal-tsunami-driven medium-tier community, and a low-tier control.

High tier: 煙硝蔵 hamlet, Kochi prefecture

煙硝蔵 GSI orthophoto chip

dimension	score	what the VLM sees
coastal_exposure	0.05	inland mountain hamlet, no coast in view
single_road_visible	**0.85**	one narrow valley road visible; no alternate ingress
terrain_lockedness	**0.90**	hemmed in by steep slopes on three sides
infrastructure_sparsity	0.70	very few buildings, no obvious public facilities
isolation_risk_visual	**0.85**	“high, single road through narrow valley”
aggregate	**0.670**	

Final isolation risk score: **0.648 (high)**. Under the Nankai scenario weights, single-path dependency and landslide hazard push the score up alongside the visual-imagery signal; the absence of tsunami exposure (this community sits inland) keeps the score from reaching the very_high tier. The decision card cites this as a P5 (long-term reinforcement / route-redundancy) case.

Medium tier: みなべ町 (Minabe-cho), Wakayama prefecture

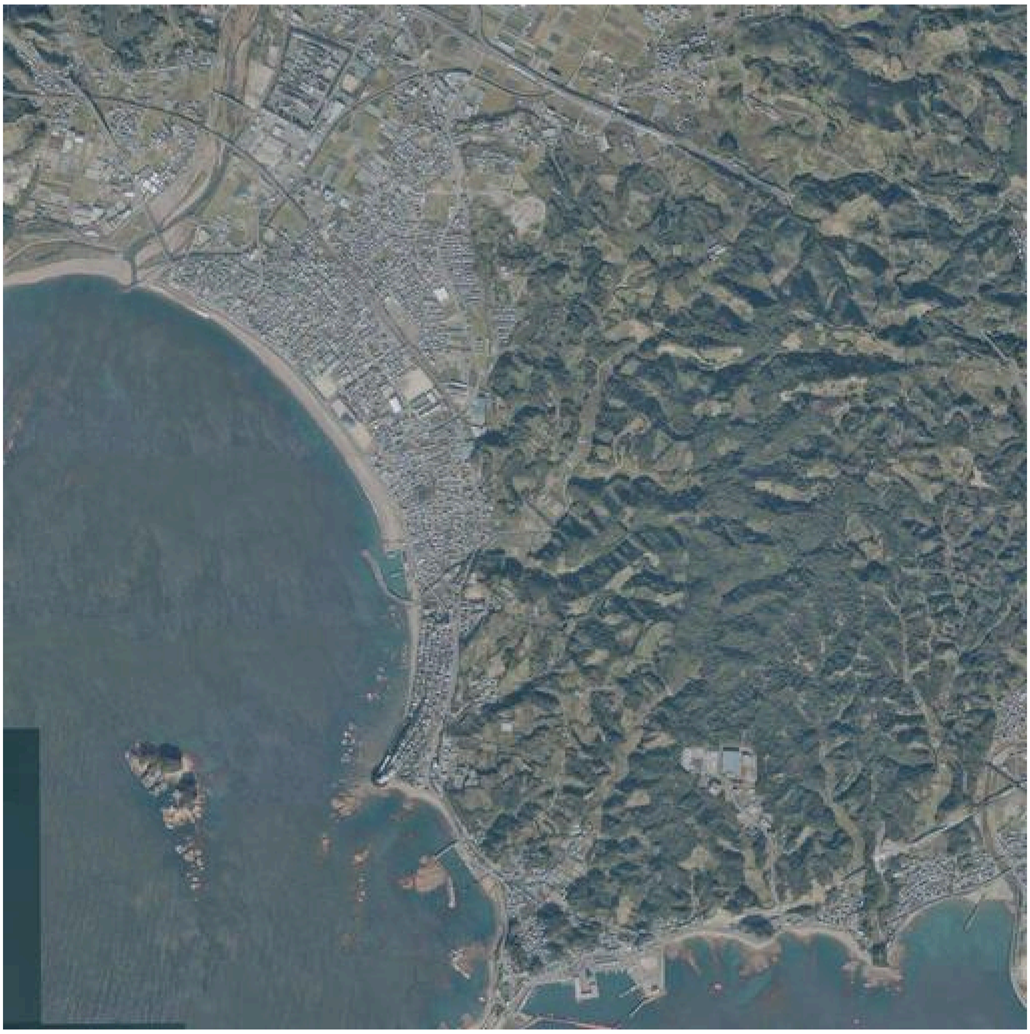

みなべ町 GSI orthophoto chip

dimension	score	what the VLM sees
coastal_exposure	**0.70**	open Pacific coastline; community is on the coast
single_road_visible	0.30	mostly grid road network behind the coast
terrain_lockedness	0.40	small hills inland but not the dominant feature
infrastructure_sparsity	0.30	typical small-town density
isolation_risk_visual	0.40	“moderate, coastal exposure is the dominant risk”
aggregate	**0.420**	

Final isolation risk score: **0.563 (medium)**. Under the Nankai scenario weights, the tsunami-exposure signal (the largest weight, 0.21) dominates this community’s contribution because its centroid sits at low elevation a short distance from the coastline. This is exactly the *kind* of community the scenario weights are designed to surface that the Noto-fitted weights (which set tsunami weight to 0 because Noto was not a tsunami event) would miss.

Low tier: 有田川町 (Aridagawa-cho), Wakayama prefecture

有田川町 GSI orthophoto chip

dimension	score	what the VLM sees
coastal_exposure	0.10	inland; no coast visible
single_road_visible	0.10	grid road network with multiple ingress
terrain_lockedness	0.20	mild river valley, not hemmed in
infrastructure_sparsity	0.10	regular small-town settlement density
isolation_risk_visual	0.15	“low, well-connected inland town”
aggregate	**0.130**	

Final isolation risk score: **0.198 (very_low)**. The VLM’s holistic read agrees with the graph topology: this community has multiple road ingresses, no significant coastal exposure, no terrain lock-in. The action card recommends annual data refresh only, no immediate operational follow-up.

What VLM does *not* do

The VLM signal is **not** the only thing that determines the tier. For high-tier communities under the Nankai scenario weights, the four mechanism-central signals (tsunami exposure, single-path dependency, landslide hazard, and cut severity) together carry roughly twice the weight that the VLM imagery signal does (0.71 vs 0.11). The imagery signal sharpens judgement where the graph- and demographic-derived signals are ambiguous, particularly the "is this a coastal strip or an inland hamlet?" distinction that pure road topology cannot make. The Noto retrospective above shows the imagery signal works well on that event; the cross-event caveat is the reason this report uses a lower imagery weight for the forward Nankai predictions than would be optimal for Noto alone.

Action playbook by tier

These tiers are calibrated against operational decision-making, not academic precision. The intent is to give an officer a clear “what do I do about this community *today*” answer.

tier	meaning	recommended pre-event action
very_high	road *and* hospital *and* demographic risk all extreme	helicopter LZ identified + sat-phone deployed + 7-day water/food/medical pre-positioned
high	two of three vulnerabilities extreme, plus single-path	sat-phone deployed + 3-day supplies + tabletop drill with neighbourhood associations
medium	clearly above population baseline; one major vulnerability	inclusion in next municipal BCP review; mapping of off-register access roads (Sister A)
low	analysed but not flagged	annual data refresh; no immediate action

Under the Nankai scenario weights, the 167-community run produced **2 high · 41 medium · 121 low · 3 very_low**. The 2 + 41 = **43-community shortlist** sets the medium-or-above tier for review. The 2 high-tier alone is well within single-working-week capacity for a typical Mie + Kochi disaster-prevention division team. The narrower shortlist (vs the 118 the empirical weights would have produced) is a deliberate consequence of the mechanism prior, tsunami-emphasised, VLM-de-emphasised, with a sharper coastal-vs-inland separation.

How the score is computed (one-page methodology)

Approach. This is a *contextual agentic vision* analysis, a pipeline that fuses publicly available spatial datasets with vision-language reading of satellite imagery, weights the resulting signals with a Bayesian fit on a labeled retrospective event, and keeps every score auditable back to its raw query. Nothing in the pipeline is a black-box predictor: every contribution is a quantitative reading of a publicly verifiable source.

Inputs. Seven public data sources, all available today without waiting for an event:

1. **OSM road graph** (openstreetmap.org): the canonical road network. Each village is a node; cuts are computed via min-cut on the local subgraph.
2. **GSI 5 m DEM** (fgd.gsi.go.jp): terrain raster from the Fundamental Geospatial Data (基盤地図情報) portal. Slope and aspect are computed at each road segment and at the village centroid.
3. **GSI landslide hazard polygons (急傾斜地崩壊危険区域)** (nlftp.mlit.go.jp/ksj): official prefectural slide zones from the National Land Numerical Information (国土数値情報) dataset, overlaid on the road graph.
4. **e-Stat National Census (国勢調査)** (e-stat.go.jp): population, ageing rate, household composition at village-block granularity, 2020 census.
5. **MLIT hospital catalog** (nlftp.mlit.go.jp/ksj): every facility with surgery capability, geocoded. Road-distance is computed via the same OSM graph.
6. **Historical disaster archive** (bousai.go.jp): prior cut-off events (Noto (能登), Heisei Heavy Rain (平成豪雨), etc.) used as a feature, not as a label.
7. **Tsunami exposure** (computed): per-community elevation from GSI + distance to coastline from OSM. Critical for the Nankai megathrust scenario where tsunami inundation is the dominant isolation mechanism for coastal settlements.

From data to signal. Each raw input is reduced to a single value between 0 and 1 per community:

- **single-path dependency**: count alternate OSM routes from the village to the nearest urban hub; combine `1 / (1 + alternates)` with the village's peripheralness (primary route length).
- **cut severity**: run min-cut between the village and the nearest hub on the local road subgraph; normalize the cut size.
- **landslide hazard**: fraction of the approach road overlapping with official slide-zone polygons.
- **terrain slope**: mean slope along the approach road and at the village centroid, mapped through a logistic curve at 25°.
- **aged population %**: 65+ share from the census, used directly.

- **historical disasters**: count of past isolation events affecting this community, normalized by a cap.
- **VLM imagery signal**: 5-dimension visual score (coastal exposure, single-road visibility, terrain lockedness, infrastructure sparsity, holistic isolation risk) on a GSI orthophoto chip of the village.
- **hospital distance**: road-distance to nearest surgery-capable facility, mapped through a logistic curve at ~100 km.
- **tsunami exposure**: combined sea-level proximity: `max(0, 1 − elev_m/30) × exp(-coast_km/5)`. Score 1 = at sea level immediately on coast; score 0 = above 30 m or beyond 15 km inland.

Combination. An explicit weighted sum, not a neural ranker. Each signal contributes by its declared weight:

Weights below are the Nankai scenario weights: a tsunami-mechanism prior rather than weights discovered from Nankai data (no Nankai ground truth exists, by construction).

Signal	Weight (mean ± std)	90 % uncertainty band	What it measures
tsunami exposure	**0.212 ± 0.070**	[0.11, 0.34]	Low elevation × close to coastline (mechanism core)
single-path dependency	**0.180 ± 0.067**	[0.08, 0.30]	Village relies on exactly one road in/out (coastal lifeline)
landslide hazard	0.169 ± 0.066	[0.07, 0.29]	Overlap with official slide-zone polygons
cut severity	0.150 ± 0.069	[0.05, 0.28]	Min-cut size on the local road subgraph
VLM imagery signal	0.112 ± 0.060	[0.03, 0.23]	5-dim visual judgement, de-emphasised (EQ-trained prompts)
terrain slope	0.068 ± 0.045	[0.01, 0.16]	Slope on the approach road and at the centroid
historical disasters	0.050 ± 0.040	[0.01, 0.13]	Past isolation events affecting the village
aged population %	0.037 ± 0.033	[0.00, 0.11]	Share of residents aged 65+ (audit annotation)
distance to hospital	0.021 ± 0.027	[0.00, 0.08]	Road-distance to nearest surgery-capable facility (audit annotation)

A second set of weights, the **Noto-fitted weights:** is used for the retrospective in §“How well does the method work?”. Those weights maximise in-sample recall on the 2024 Noto data (VLM imagery weight 0.23, satellite-derived signals 0.15–0.17 each, tsunami exposure 0.0 because Noto was not a tsunami event). They are not used for the Nankai forward predictions in this report, for the reasons explained in the cross-event paragraph below.

The score is the sum of these contributions (each in 0-1 range, summed into a 0-1 risk score).

How the weights were determined. Each of the nine weights is the posterior mean of a Bayesian fit on the 2024 Noto retrospective: a domain prior, centred on hand-set defaults that reflect civil- engineering judgement (graph-topology and landslide-hazard signals should dominate over aged-population share), is updated by the Noto ground-truth labels through an ordinal-logistic likelihood. The fit runs in PyMC; the table above quotes the posterior mean ± standard deviation and a 90 % credible interval for each signal.

For the South-Sea Trough forward predictions, the prior centre is shifted to emphasise the Nankai mechanism (tsunami exposure + coastal-lifeline vulnerability), trading some in-sample recall on the Noto 35 for weights that align with the Cabinet 2025 inundation forecasts. The two weight sets, the Noto-fitted weights and the Nankai scenario weights, are both reproducible from public data using the build pipeline described in the technical appendix.

The transfer of Noto-fit weights to the tsunami-driven Nankai forward predictions was tested separately by holding out the 2019 Hagibis typhoon: the earthquake-fit ranking did not transfer to typhoon- driven isolation. This is the methodological reason this report uses the mechanism-prior weights rather than the Noto-fit weights directly.

Why an explicit formula, not an end-to-end neural ranker? Two reasons. (a) **Audit**: an officer challenged on “why did you flag this community instead of that one?” can answer in numbers. (b) **Update**: when a new road is built or a hospital closes, exactly one input changes, exactly one score moves, and the change is explainable.

Limitations (and what we plan to do about each)

1. VLM signal limited to 5 specialist dimensions. Current activation covers coastal exposure, single-road visibility, terrain lockedness, infrastructure sparsity, and a holistic isolation-risk judgement on GSI orthophoto chips. Other visual signals (active landslide scars, specific bridge fragility, recent road construction) are not yet specialist prompts. *Plan*: design domain-specific specialist prompts for bridge fragility, slope crack patterns, and night-light recovery signal for Mode B re-ranking.

2. Off-register roads under-counted. Many real access routes (forest roads, farm roads, private bridges) are not in OpenStreetMap. Our `Sister A` module reconstructs these from GSI satellite imagery using the same vision-language model stack. **Sister A is used here as a map-completeness audit tool, not as a scoring input.** Testing showed that injecting reconstructed paths back into the scoring graph does not in itself shift any community's risk tier on the 2024 Noto data. The most affected communities were already known to be on single-lifeline roads regardless of these side paths. The 13 reconstructed paths in the Nankai regions (6 in Kii, 7 in Kochi) remain useful for visualising what is missing from public road maps, but are not claimed to change the risk ranking in this report.

3. Tsunami inundation is included as a proxy signal (community elevation × distance to coastline) but not yet calibrated against official inundation hazard maps. South-Sea Trough produces tsunamis up to 30 m on the Tōkai coast. The current proxy uses GSI 10 m elevation tiles and OpenStreetMap coastline distance with the formula `max(0, 1 − elev_m/30) × exp(−coast_km/5)`. This is a structural proxy, not a hydrodynamic inundation simulation. The Central Disaster Prevention Council publishes case-based inundation depth maps (2012 / 2025 updates); ingesting these as a per-community inundation depth feature would replace the proxy with a calibrated signal. *Plan*: parse Cabinet Office Nankai inundation rasters and swap in inundation-depth-meters per community as a replacement.

4. Mode B not exercised here. Mode B re-ranks the same communities in real time as road-recovery reports flow in after an event, scrubbing through `t = +6h, +24h, +7d, +30d`. This report is Mode A only. *Plan*: the same engine runs both; a partner deploying this operationally would get both out of the box.

5. 167 communities is two prefectures. Full South-Sea Trough coverage is 5 regional blocks (Mie / Wakayama / Tokushima / Kochi / Ehime + Oita partial). *Plan*: one ingestion pass per prefecture; the pipeline is the same.

Beyond this report

This PDF is one snapshot. The live system offers:

- **An interactive map**: every community on the Nankai map is clickable, and clicking any score opens its raw VLM / retrieval / graph algorithm output. The full evidence card from page 6 is what pops up when you click any community on the map.
- **Mode B time slider**: replay an event from the first MLIT road-recovery report to "everything restored", with the ranking updating live as official records arrive.
- **Per-prefecture BCP export**: a one-click button per prefecture exports the shortlist plus playbook into the format used by prefectural disaster-prevention plan documents.
- **Pluggable VLM / CV providers**: Anthropic Claude, OpenAI, Google Gemini, a self-hosted SAM2 / Sen1Floods CV server, or a *filesystem* provider that lets Claude Code itself read the chips and write responses with zero API key required.
- **Same approach, different scenario**: the contextual agentic vision pipeline applies, with the same evidence-trace contract, to flood / sediment-disaster / facility-resilience analyses and to post-event re-ranking from the moment the first road-recovery report arrives.

Generated 2026-05-28 from data with no post-event leakage. No earthquake outcome was used as input.

Yodo Labs · PixelX Inc. (ピクセルエックス株式会社) yodolabs.jp · team@yodolabs.jp