Can GRAND and TAMBO share a mountain?

Mauricio Bustamante

Niels Bohr Institute, University of Copenhagen

TAMBO Collaboration Meeting
August 18, 2026

UNIVERSITY OF COPENHAGEN

VILLUM FONDEN

Why put two detectors on the same mountain

TAMBO: the Tau Air-Shower Mountain-Based Observatory

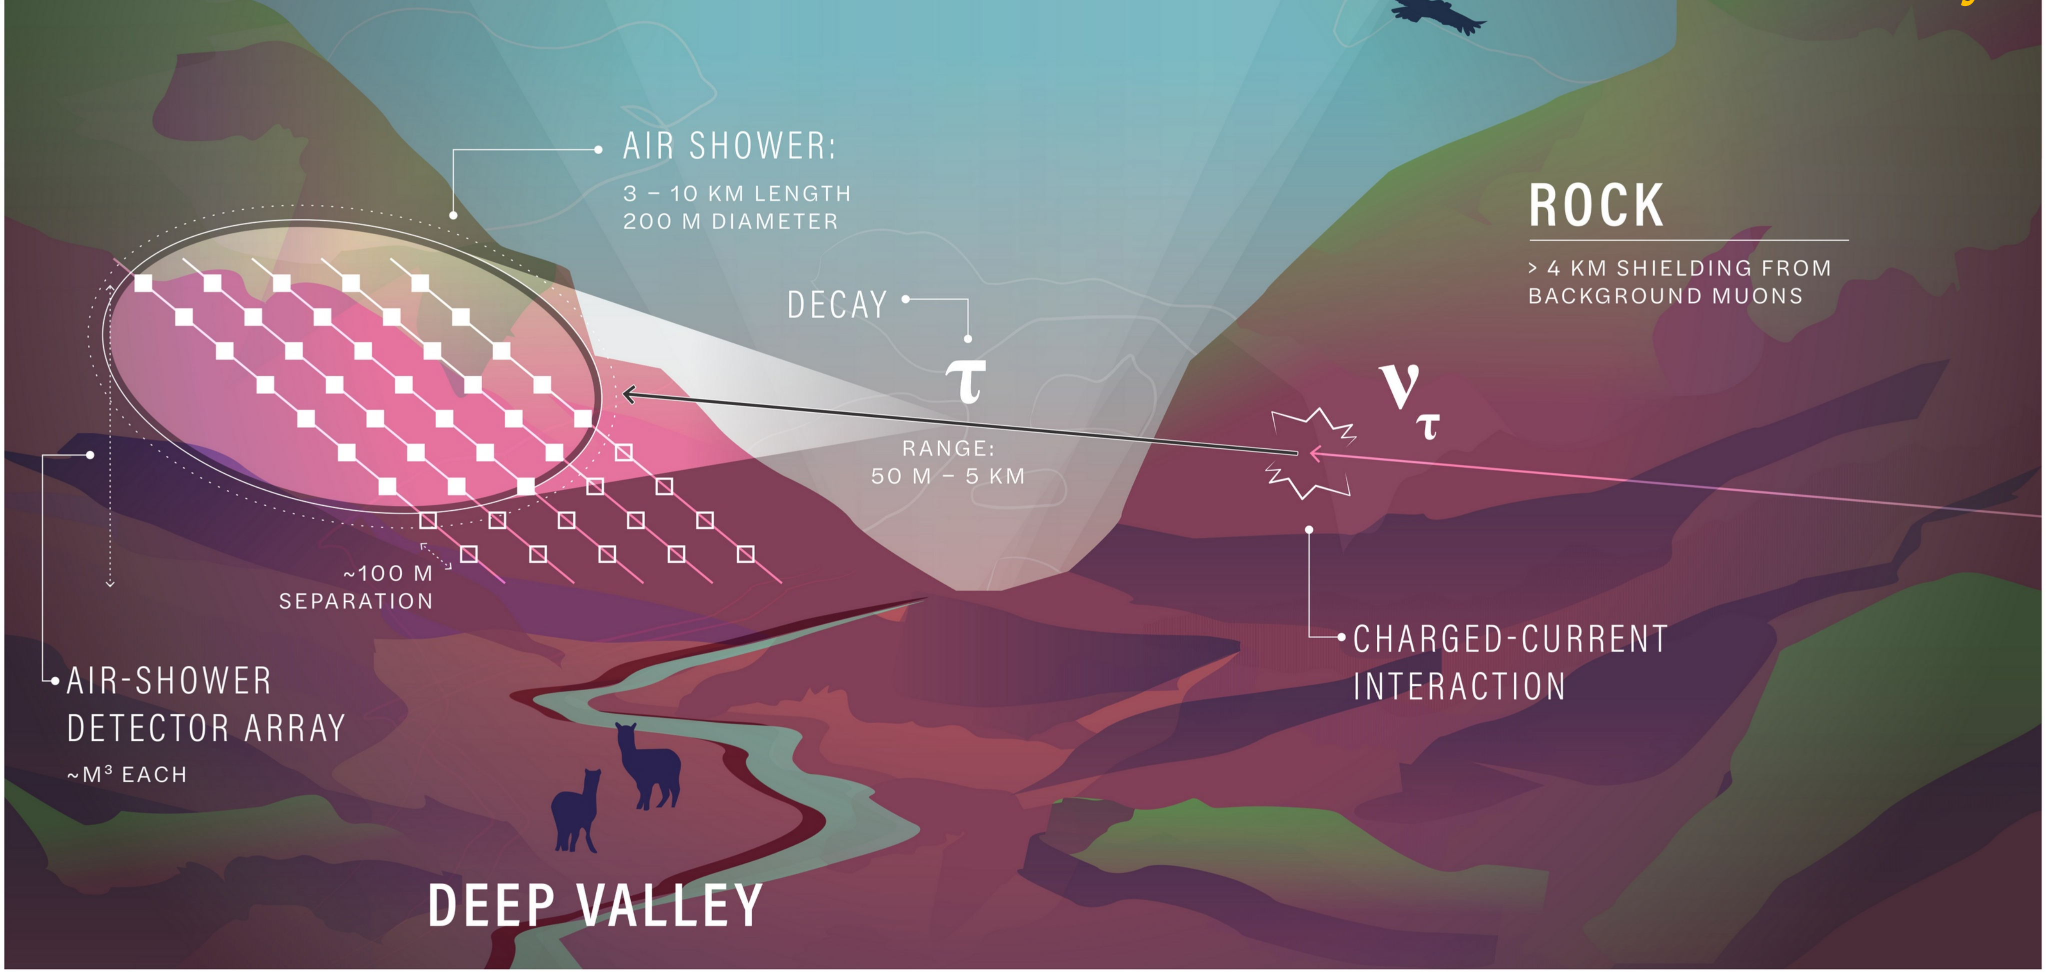

TAU AIR-SHOWER MOUNTAIN-BASED OBSERVATORY (TAMBO) · COLCA VALLEY, PERU

GRAND: Giant Radio Array for Neutrino Detection

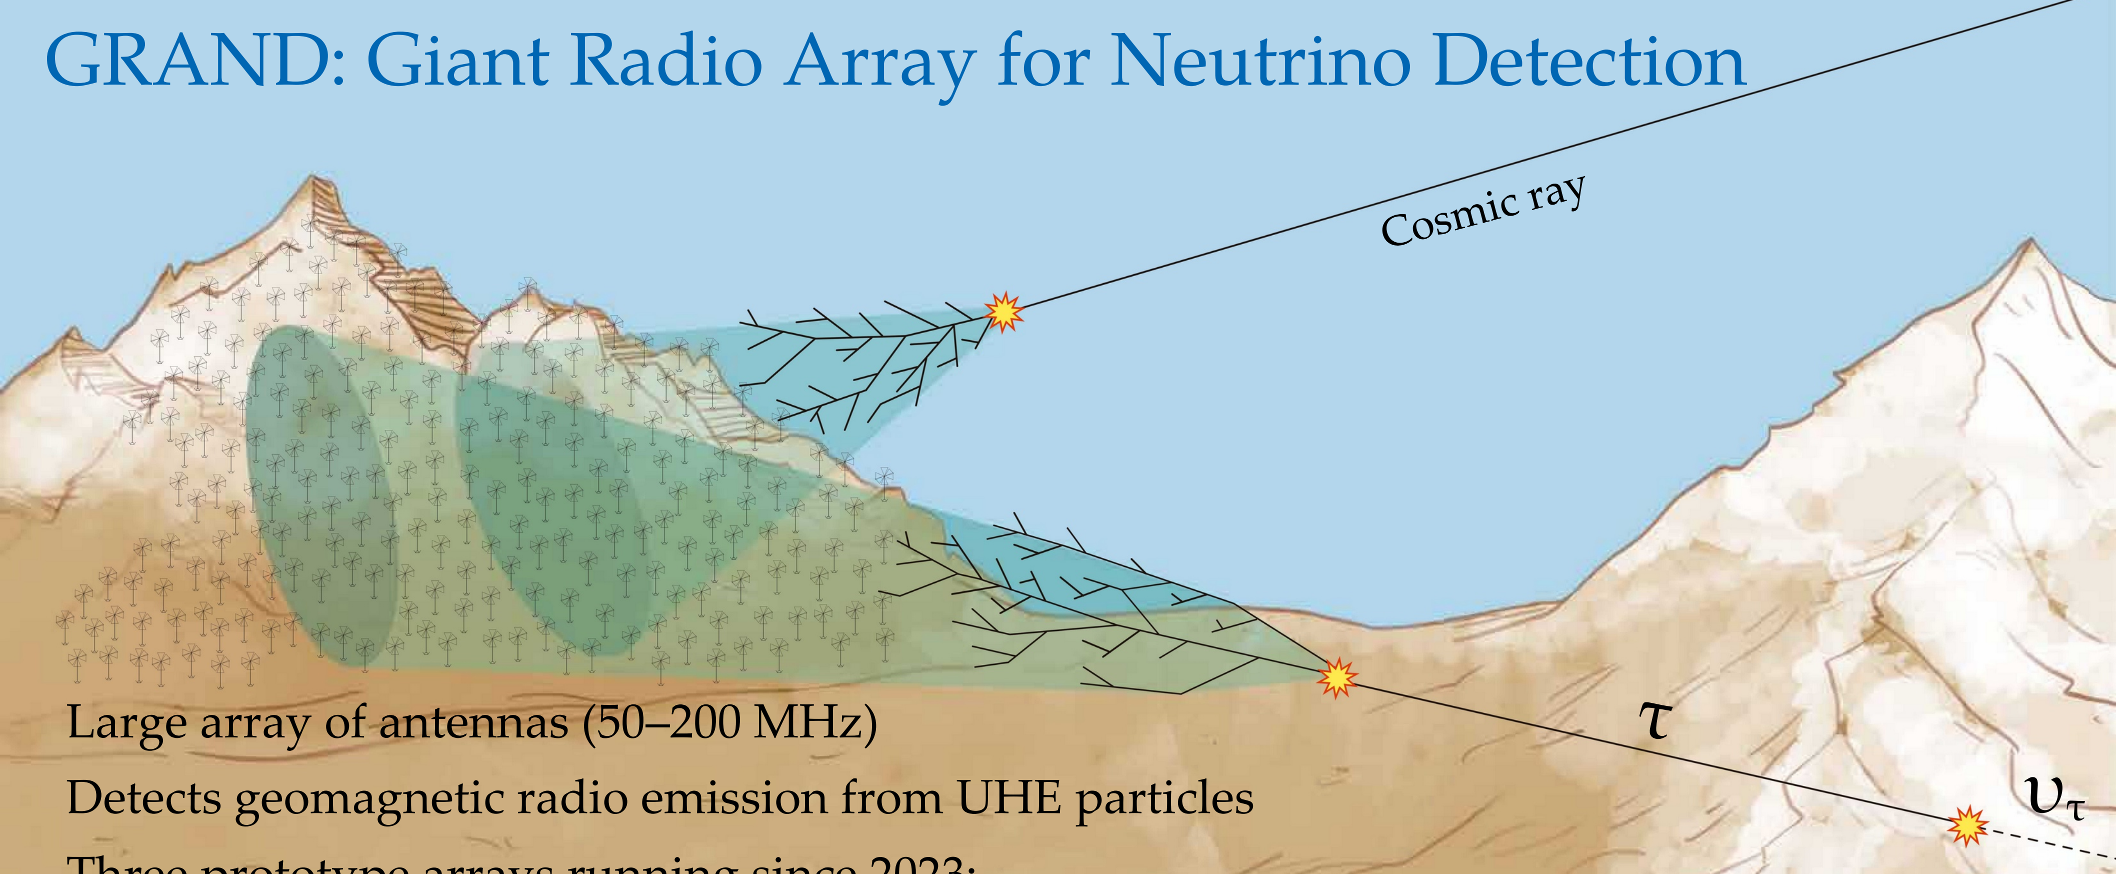

Large array of antennas (50–200 MHz)

Detects geomagnetic radio emission from UHE particles

Three prototype arrays running since 2023:

GRANDProto300	•	GRAND@Auger	•	GRAND@Nançay
65 antennas, China		10 antennas, Argentina		5 antennas, France

First cosmic-ray candidates shown at ICRC 2025

Funded: extension to 300 antennas

GRAND: Giant Radio Array for Neutrino Detection

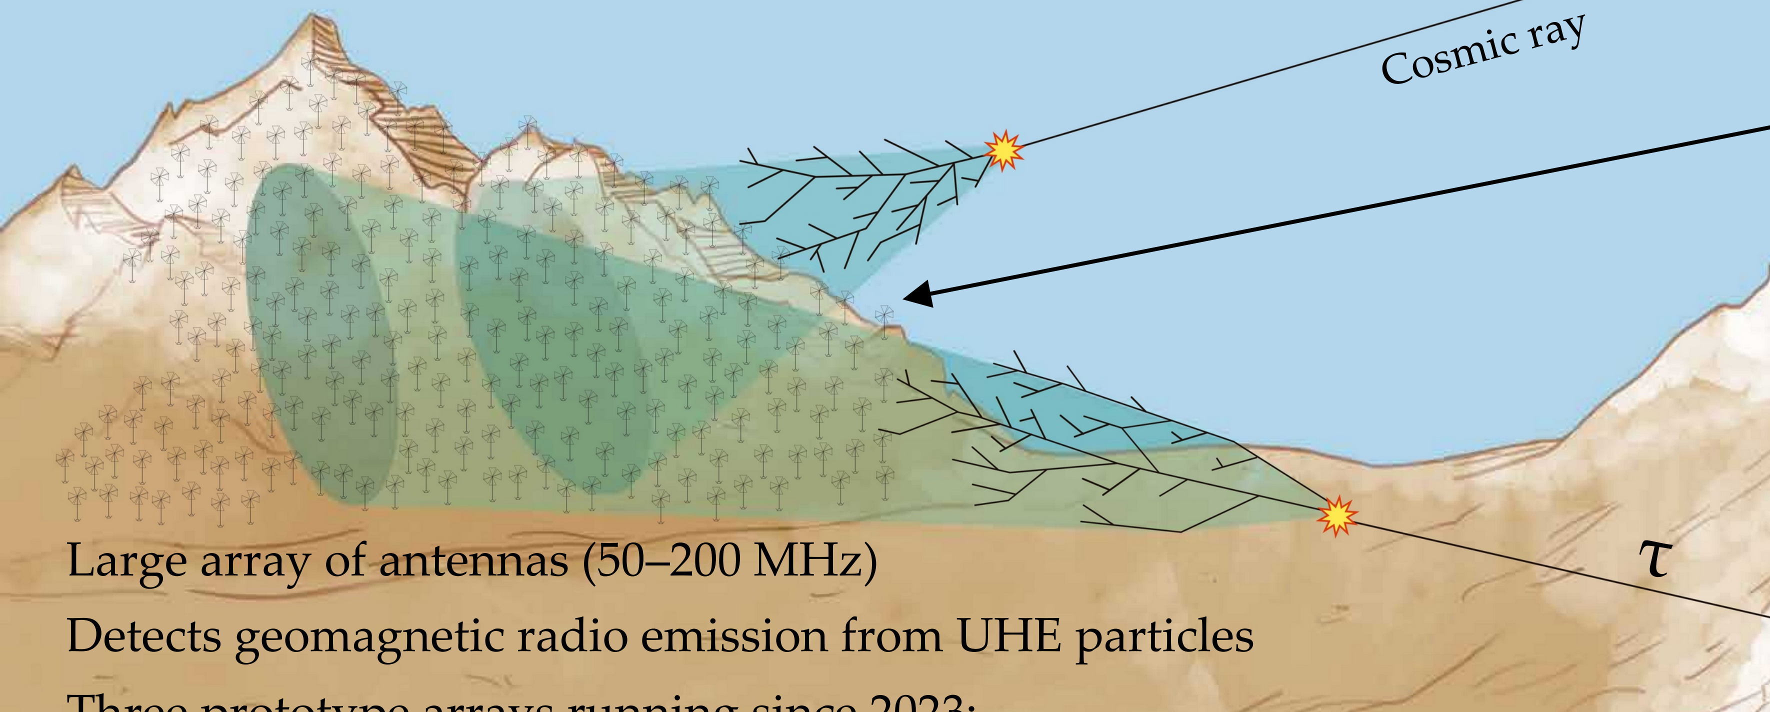

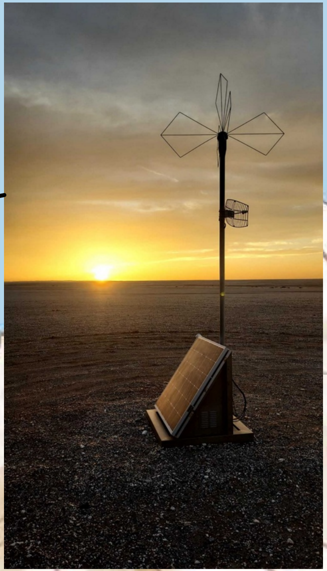

Large array of antennas (50–200 MHz)

Detects geomagnetic radio emission from UHE particles

Three prototype arrays running since 2023:

GRANDProto300	GRAND@Auger	GRAND@Nançay
65 antennas, China	10 antennas, Argentina	5 antennas, France

First cosmic-ray candidates shown at ICRC 2025

Funded: extension to 300 antennas

The messenger: Earth-skimming tau neutrinos

A ν_τ crosses the Earth, interacts in rock, and makes a τ.

If the τ escapes to the surface, it decays in air and starts an extensive air shower.

The Earth is the target. The atmosphere is the calorimeter. A mountain is the interface.

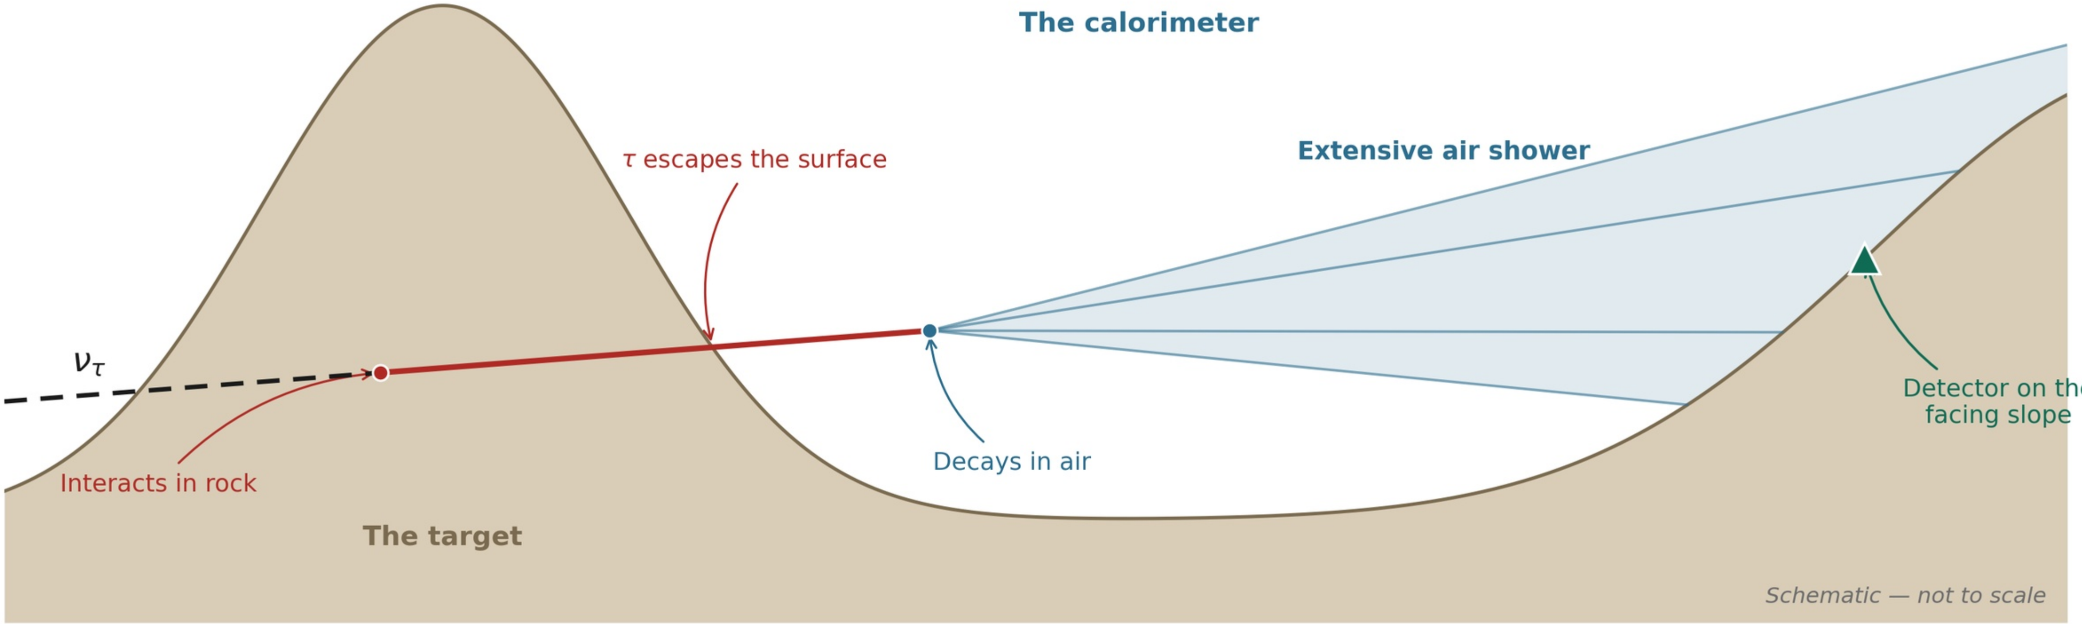

Same shower, two readings

	GRAND	TAMBO
Channel	Radio from the shower	The particles themselves
Sees of the shower	Only the electromagnetic part	All of it
Energy reach	$\gtrsim$100 PeV	~3 PeV – 1 EeV

Radio reflects only the electromagnetic component of the shower; the particle detector samples all of it. Two independent channels, not one signal read twice.

Energy ranges **overlap but are offset**; together, they measure more than either alone.

GRAND: arXiv:1810.09994 — TAMBO: Nature Astronomy (2026), doi:10.1038/s41550-026-02916-4

Why co-locate at all?

1. Cross-calibration

An established technique beside a new one: same sky, same energy decade, two channels that sample different parts of the same shower.

2. A joint spectrum wider than either alone

PeV to EeV, from one site.

3. Shared access, power, data, permits, community relations

At 4,000 m in the Andes these are not a footnote in the budget.

4. Hybrid events

If the arrival directions and energies genuinely coincide.

Why Peru: each experiment's requirements

	GRAND needs	TAMBO needs
Ground	Gentle slope, 3–25°	A canyon wall, 20–60°
Spacing	Antennas every 1 km	Detection units every 150 m
A target	Massifs 10–40 km away	A far wall 2–5 km away
A view	±3° about the horizon	±20°

The Andes supply both: high plateau + canyons among the deepest on Earth, within tens of km of each other. Colca Canyon is TAMBO's chosen site: high, dry, sparsely populated.

The two lists name the same resource only once: the ground itself. Hold that thought.

What competes, and what does not

Standing. A patch of ground has **one slope**. GRAND builds on 3–25°, TAMBO on 20–60°. The only ground both can build on lies in **20–25°**: a five-degree sliver.

Looking. GRAND watches massifs **10–40 km** out, within 3° of the horizon. TAMBO watches a canyon wall **2–5 km** out, within 20°. Different targets, different ranges, different angles. Nothing to compete for.

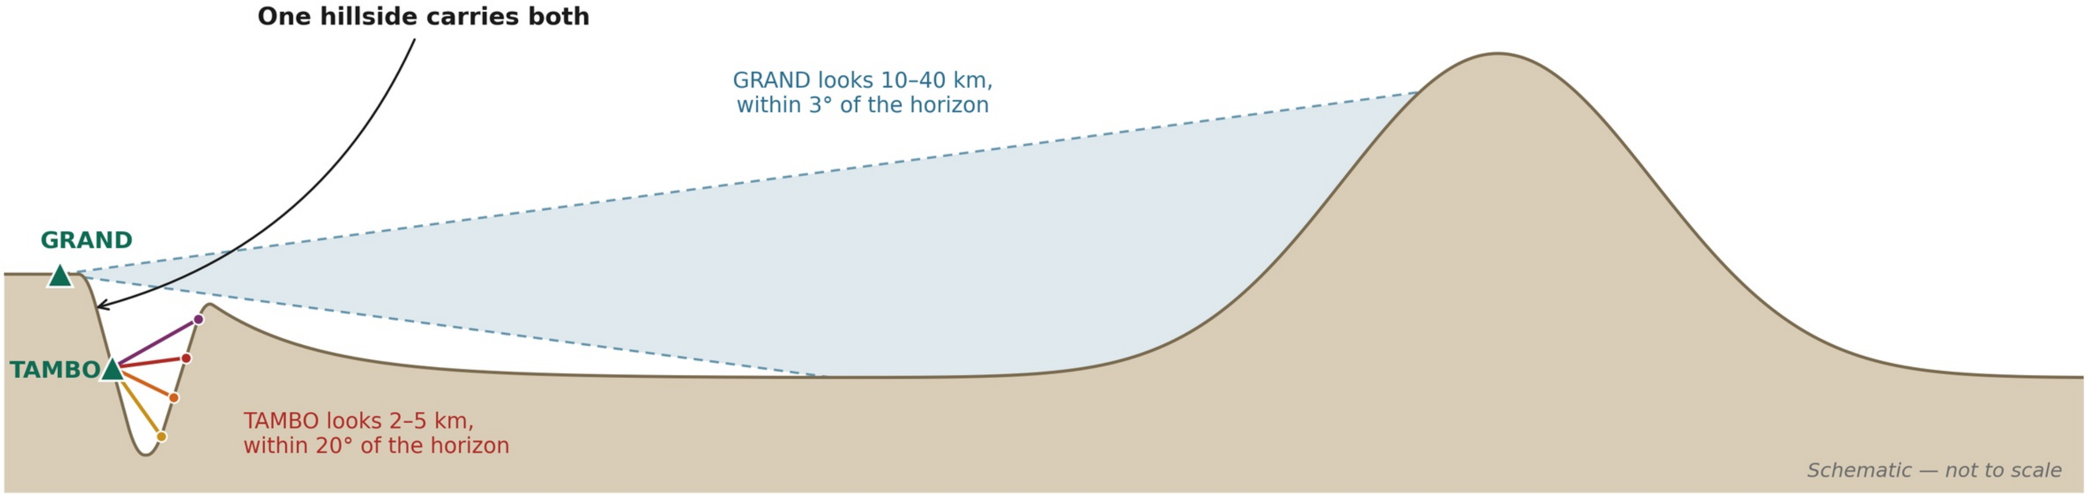

Systematically searching for array sites

What do we want answered?

> From this patch of ground, is there a **target surface** at the **right range**, in the **right direction**, at the **right relative orientation**, with the **right matter behind it**?

One engine serves both experiments: **they differ in their numbers, not their structure.**

(Adding an experiment is a configuration file, not a code redo.)

The data: digital elevation model (DEM)

A DEM is a raster: a grid of pixels, each containing the elevation of the ground it covers.

We use SRTMGL1 DEM files (1 arc-second, ≈ 30 m per pixel) from OpenTopography (70–130 Mpx for each of Peru's department we explore).

(a) The Colca DEM — 3,061 × 1,981 pixels (6.1 Mpx)

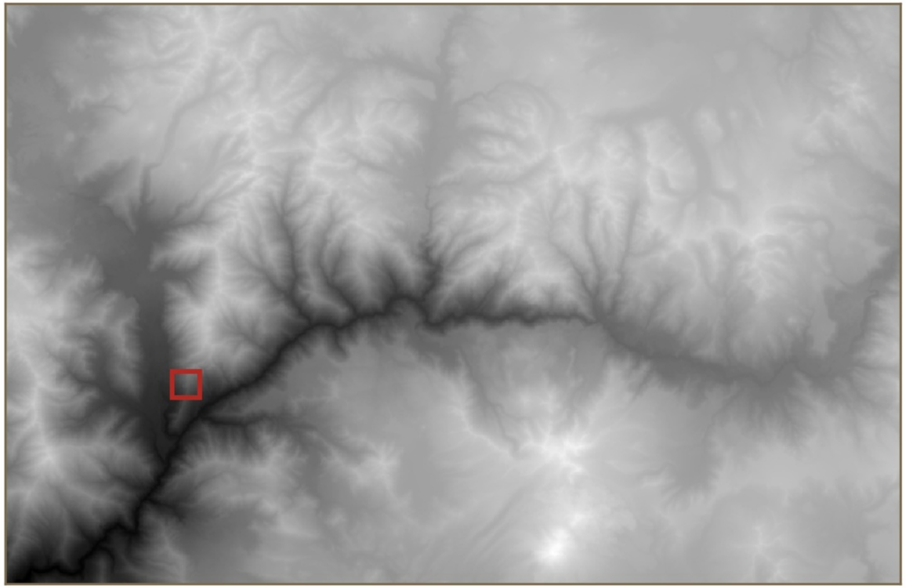

(b) One pixel = 30 m of ground = one elevation (m)

2786	2810	2844	2868	2896	2918	2951	2973	2994
2782	2807	2849	2881	2916	2935	2949	2970	3009
2779	2812	2820	2877	2922	2940	2960	2980	3010
2777	2811	2820	2867	2894	2935	2955	2986	3014
2770	2804	2817	2846	2884	2917	2952	2987	3008
2748	2786	2817	2835	2855	2892	2926	2970	2994

Elevation: SRTMGL1 via OpenTopography.

Oroscope: a new tool to search for astroparticle array sites

A terrain search for tau-neutrino detectors: an elevation model in; out come the ground that qualifies, ranked sites, and the assumptions that produced them.

Install: `pip install oroscope` – then simply `import oroscope` in your code

Code: github.com/mbustama/oroscope

Docs + API ref.: mbustama.github.io/oroscope — 13 tutorial Jupyter notebooks

```
# Fetch the elevation model (a free OpenTopography key)
$ oroscope-fetch-dem --region arequipa --open_topography_api_key KEY

# Cut the Colca window out of it
$ oroscope-crop input/dem/arequipa_SRTMGL1.tif input/dem/colca.tif \
     --north -15.30 --south -15.85 --west -72.40 --east -71.55

# Run the search – the experiment is the configuration file
$ oroscope --config_path config/tambo_colca_config.json
```

NumPy- and Numba-optimized and parallelized, with memory safeguards to prevent blow-ups:
a department-sized scan at 30-m resolution runs in **26.8 minutes**; all of Peru at 90-m, in **4 minutes**.

Striding, closing, pruning

Striding: Before the expensive scan, we keep one surviving candidate pixel in N (here 1 in 5).
A memory-cost control, not a physics criterion

The price is paid downstream: **closing** must reconnect what striding cut apart, and a fragment that stays too small fails the minimum-array-size cut at **pruning**.

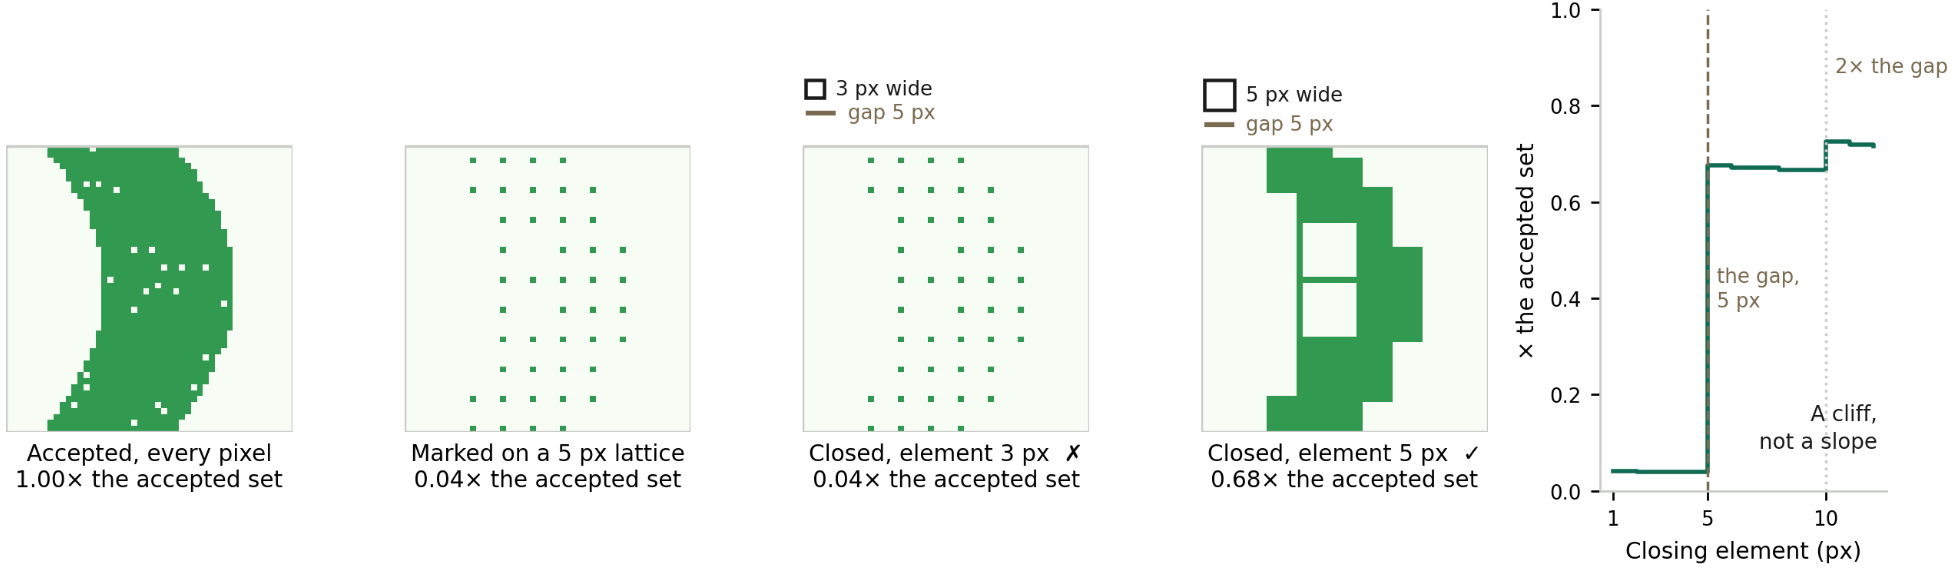

Scoring: ten components, one product

Each score component grades one criterion in [0, 1]:

clearance: The radio path stays clear — Sight-line clearance to the target, in Fresnel radii; full marks at one clear zone.

decay: The τ decays inside the gap — $1 - \exp(-\text{gap}/\text{decay length})$, folded over a power-law spectrum.

depth: Rock enough behind the exit — A band on column depth: enough rock to produce the τ, little enough that it escapes.

distance: Exit in the distance window — 1 inside the experiment's window, falling to zero outside it.

footprint: Spacing samples the footprint — The radio footprint (cone angle × distance, narrower at altitude) vs. the detector spacing.

geomagnetic: Emission in local field — Mean angle between accepted directions and geomagnetic field; 0 along it, so N–S showers fail.

muon_shielding: Rock enough to block muons — pass or fail: the overburden must stop atmospheric muons (TAMBO has > 4 km).

nu_survival: Survives the Earth chord — exp(−chord/interaction length): the ν must reach the exit unabsorbed.

shower: Matures before arrival — Air along the path, in g cm^{-2}: the shower must reach its maximum before it reaches the detector.

solid_angle: How much sky is accepted — Accepted solid angle as a share of what the fan could accept; saturating, half marks near 8%.

Multiplied together: a site must be good at everything; a single zero rejects it outright

Scoring: ten components, one product

We keep scores only > min_score:
0 (GRAND), 0.35 (TAMBO)

(Not systematically explored; other rankings available in code)

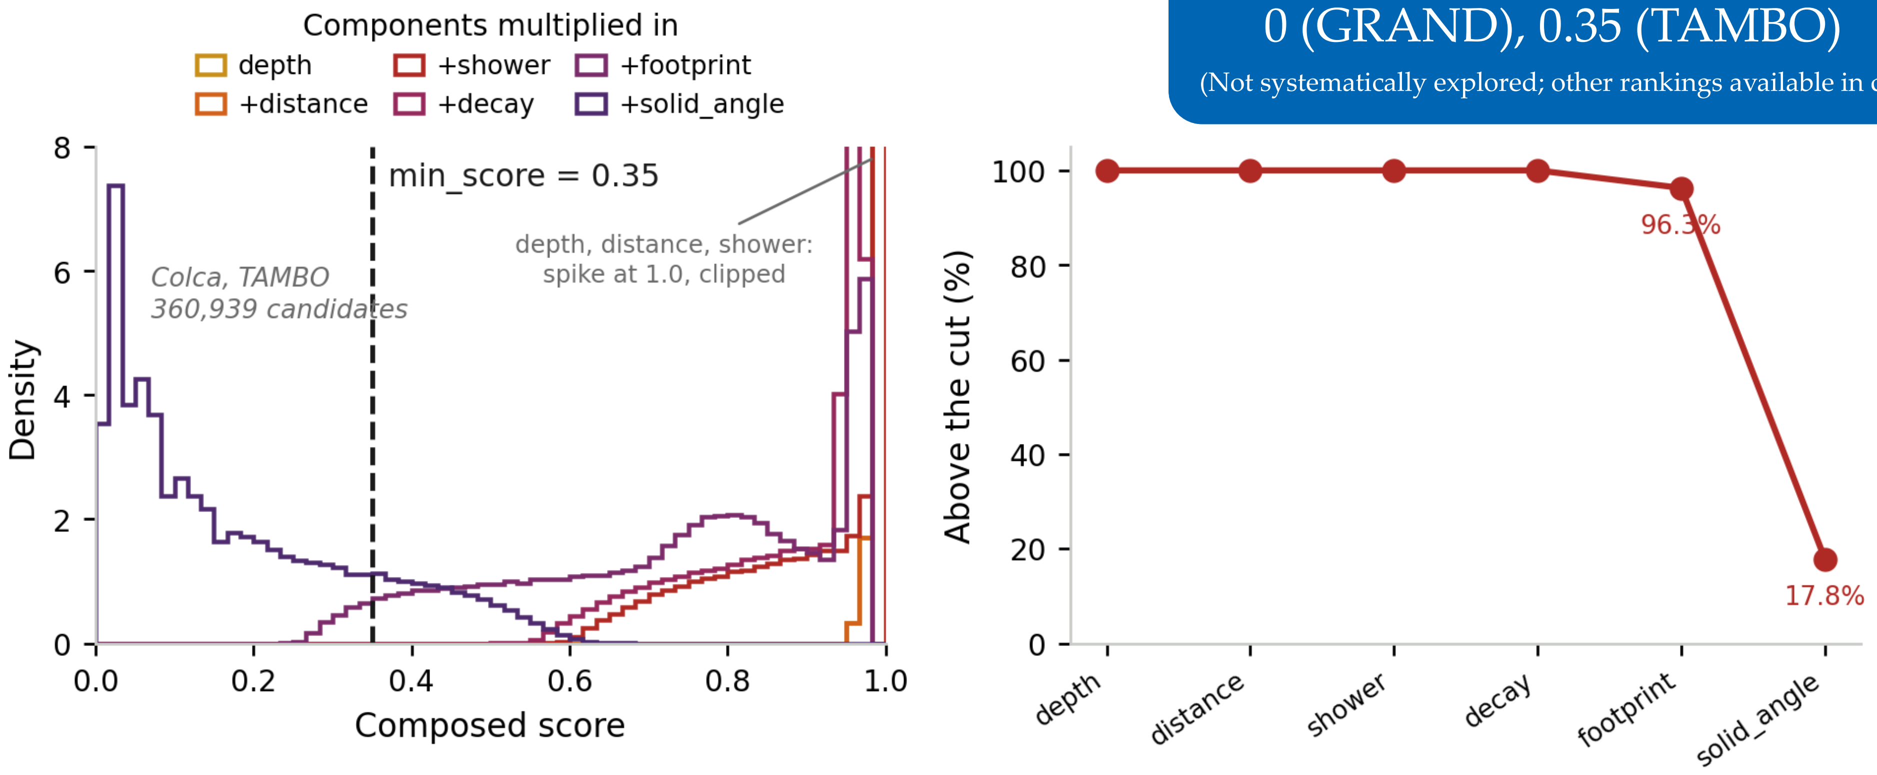

Multiplied together: a site must be good at everything; a single zero rejects it outright

Arrival scan: when does a direction count?

A direction is **accepted** when the first terrain the ray meets lands inside two windows — both drawn here — and, for TAMBO, strikes ground steep enough:

Distance window — how far away a target may sit: GRAND **10–40 km**, TAMBO **2–5 km**.

Elevation window — how far above or below the horizon: GRAND **±3°**, TAMBO **±20°**.

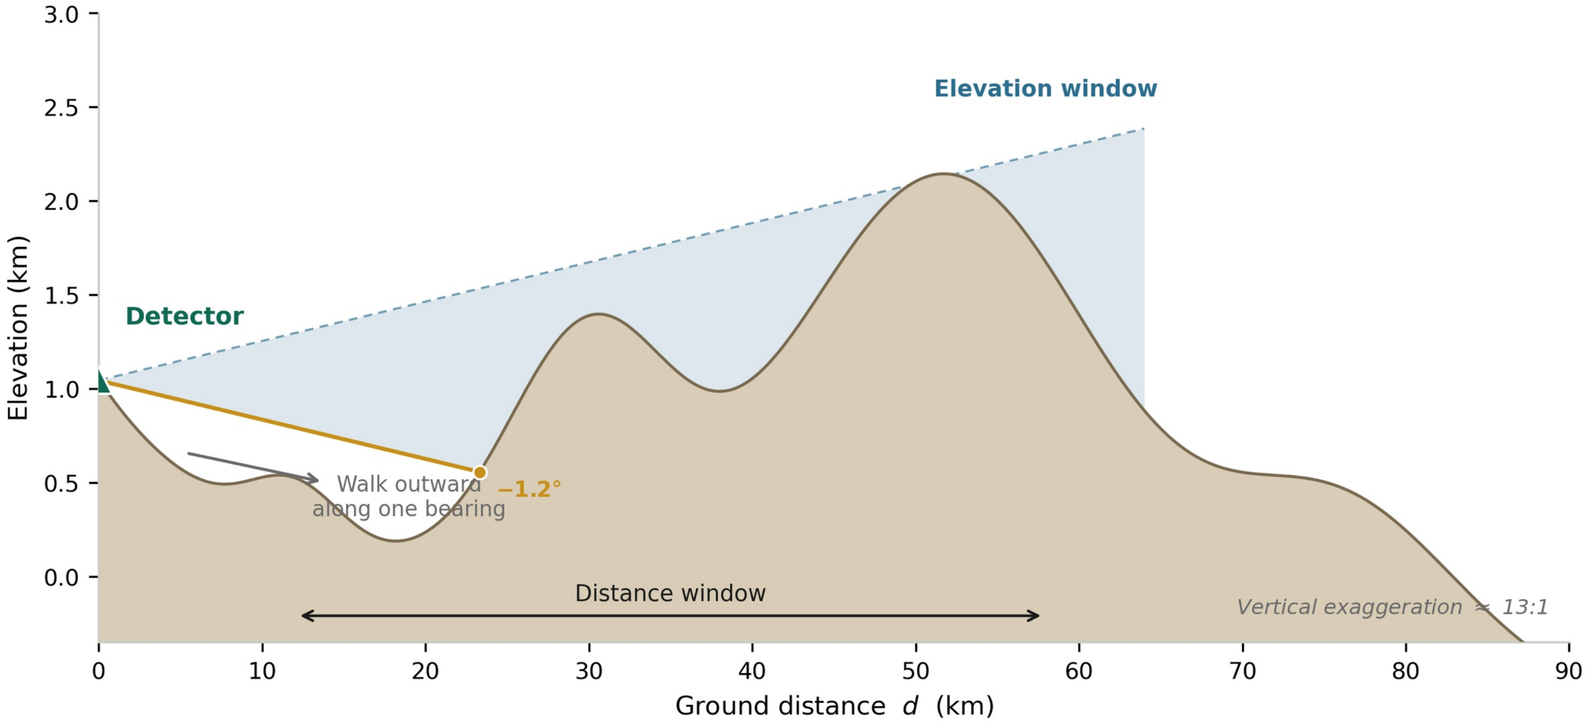

Arrival scan: when does a direction count?

A direction is **accepted** when the first terrain the ray meets lands inside two windows — both drawn here — and, for TAMBO, strikes ground steep enough:

Distance window — how far away a target may sit: GRAND **10–40 km**, TAMBO **2–5 km**.

Elevation window — how far above or below the horizon: GRAND **±3°**, TAMBO **±20°**.

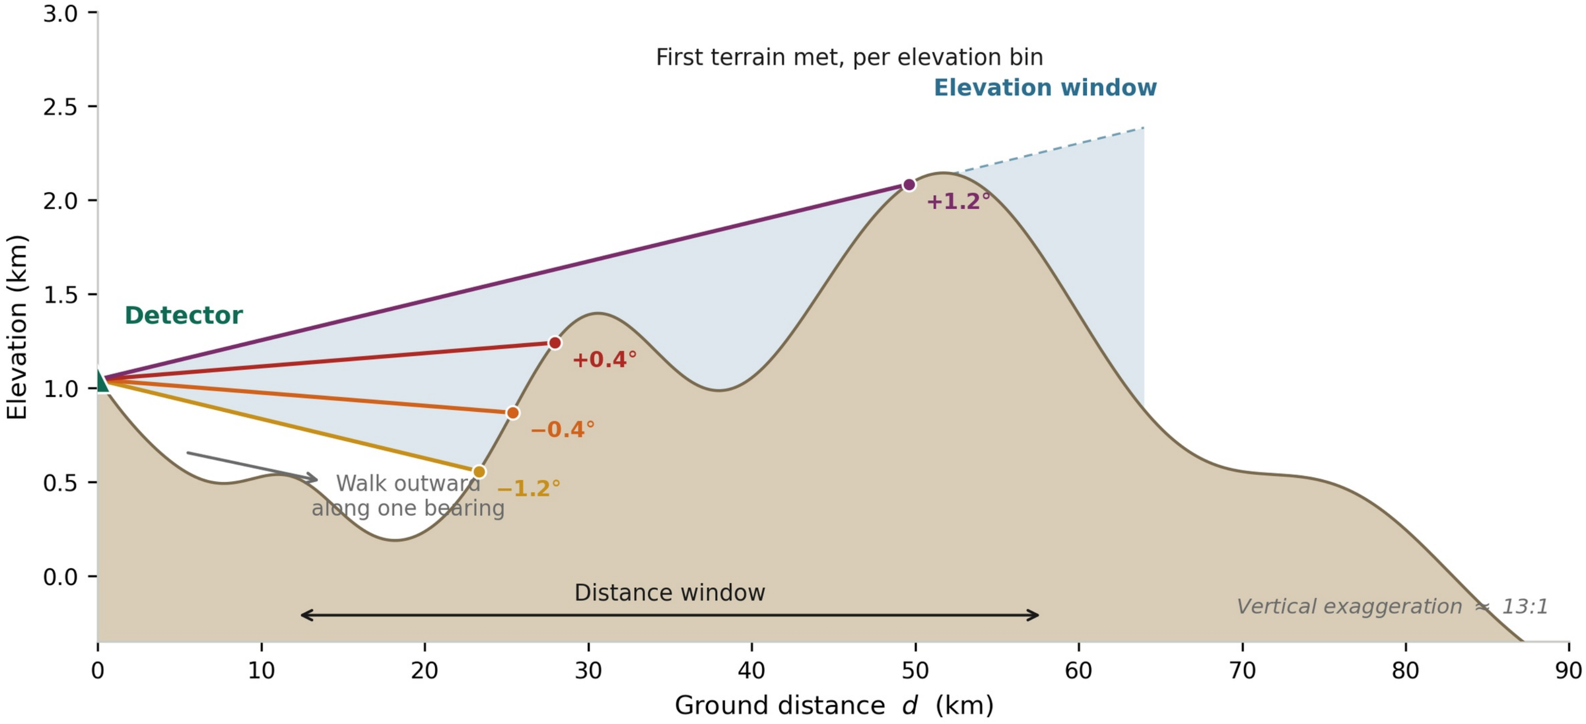

The Oroscope funnel

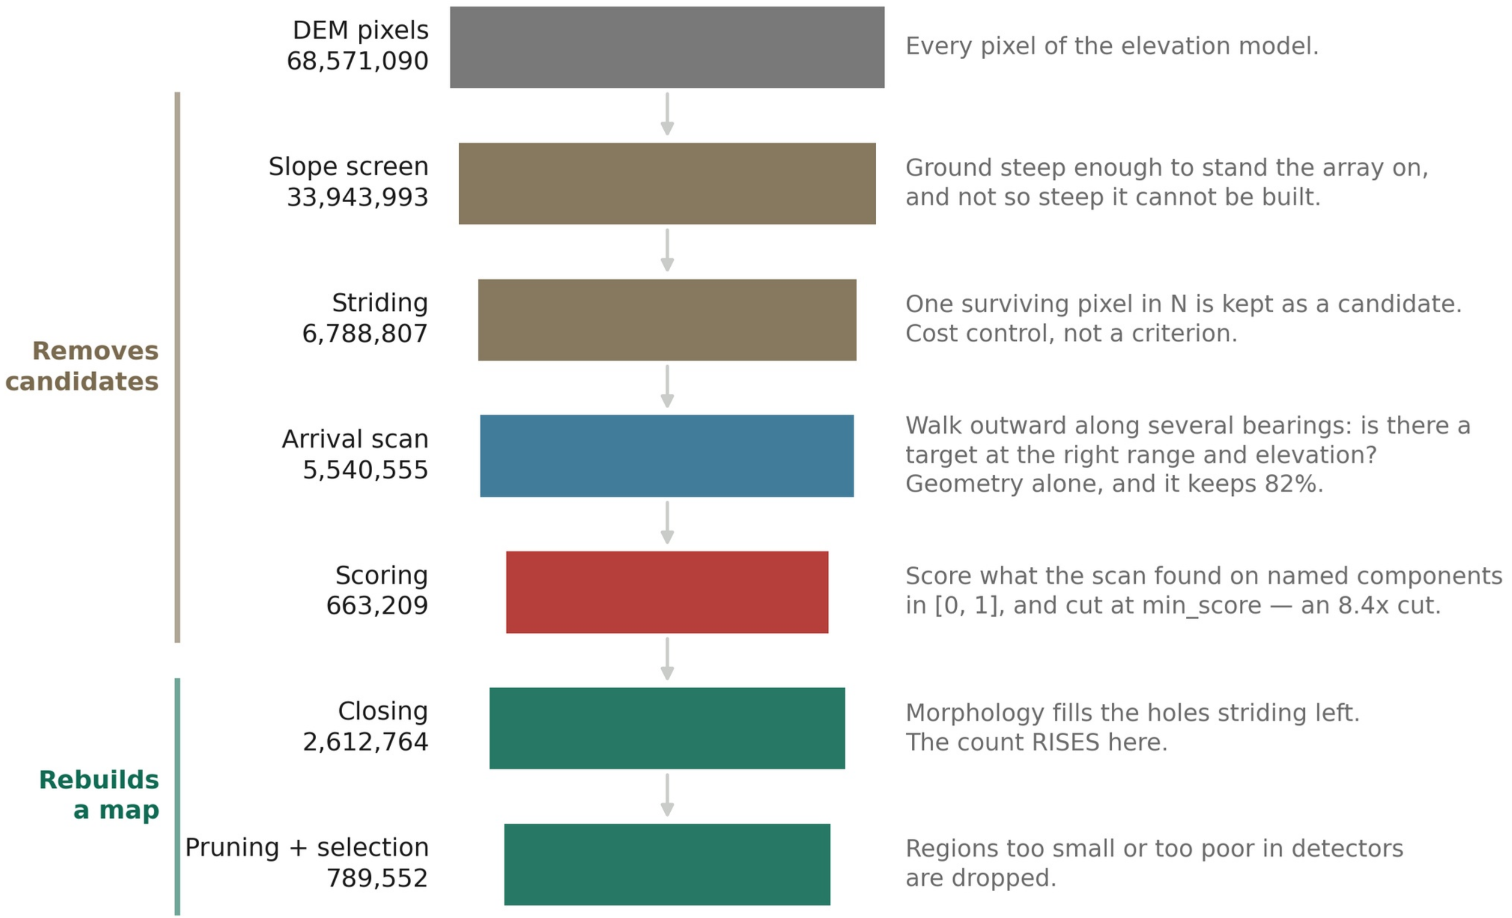

TAMBO over the full Ancash DEM. Bar widths are logarithmic in the survivor count.

What we found

Arequipa, in full – placing **GRAND**

The blue boundary encloses **87,589 km²** — plateau and gentle ground, most of the department.

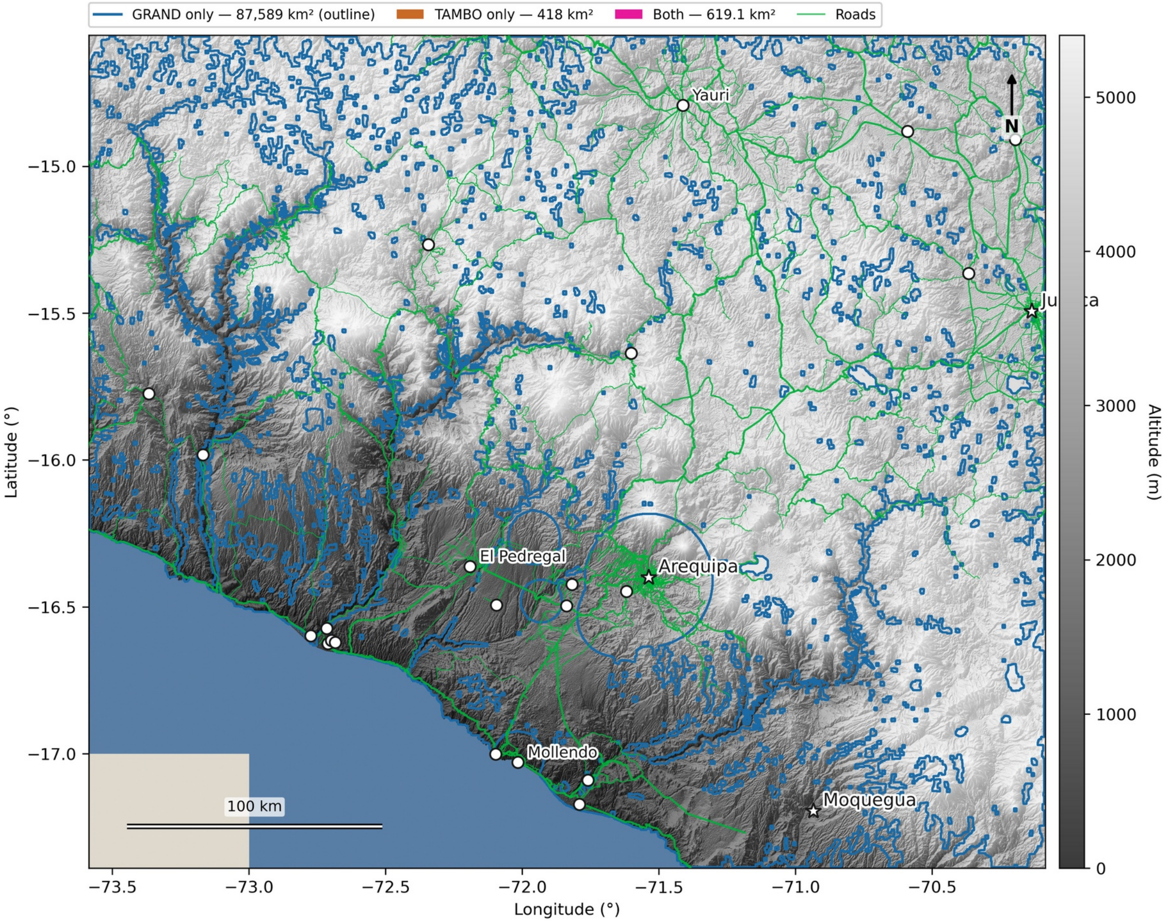

Roads and place names © OpenStreetMap contributors (ODbL). Elevation: SRTMGL1 via OpenTopography.

Arequipa, in full – placing **TAMBO**

418 km², strung along the canyons — TAMBO is the scarce partner from the start.

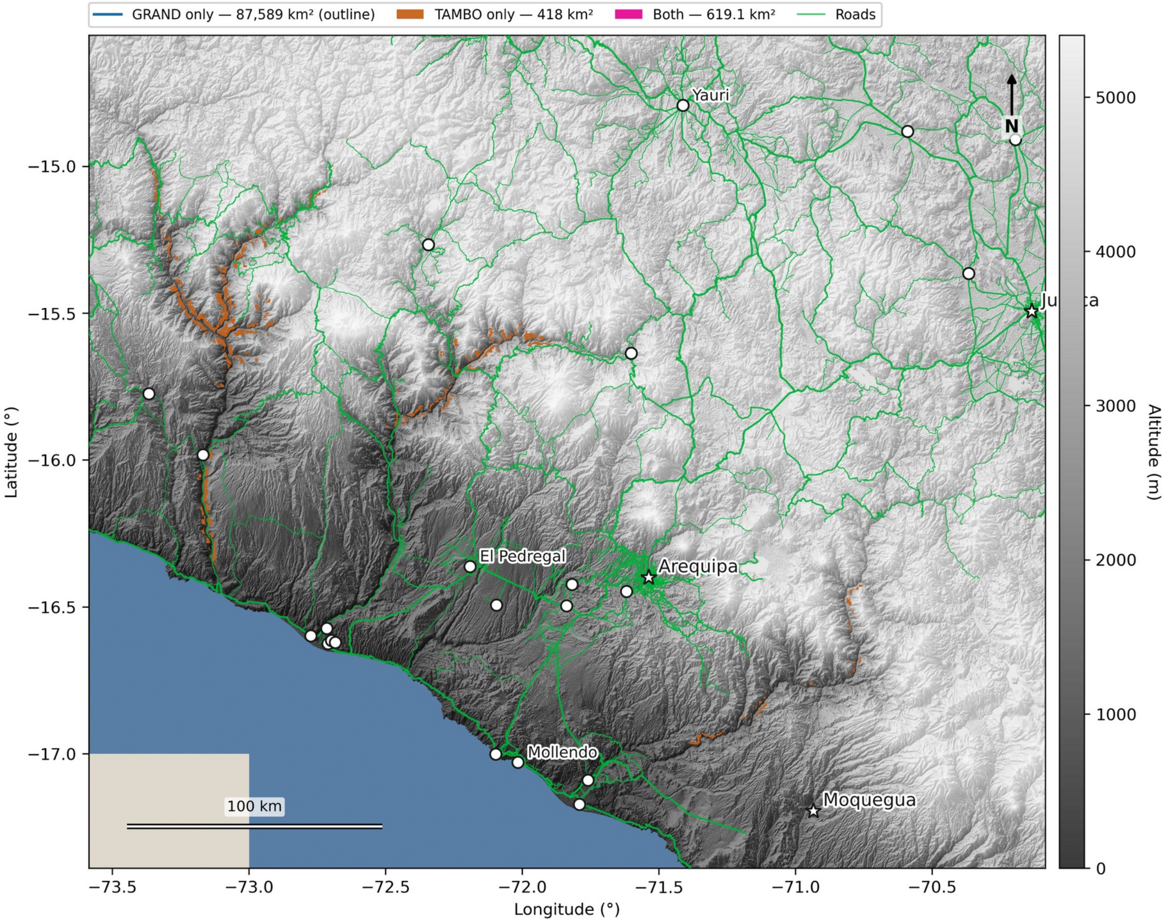

Roads and place names © OpenStreetMap contributors (ODbL). Elevation: SRTMGL1 via OpenTopography.

Arequipa, in full – placing **GRAND + TAMBO**

619 km² shared — **59.7%** of everything TAMBO-viable (619 of 1,037 km²) is also GRAND-viable.

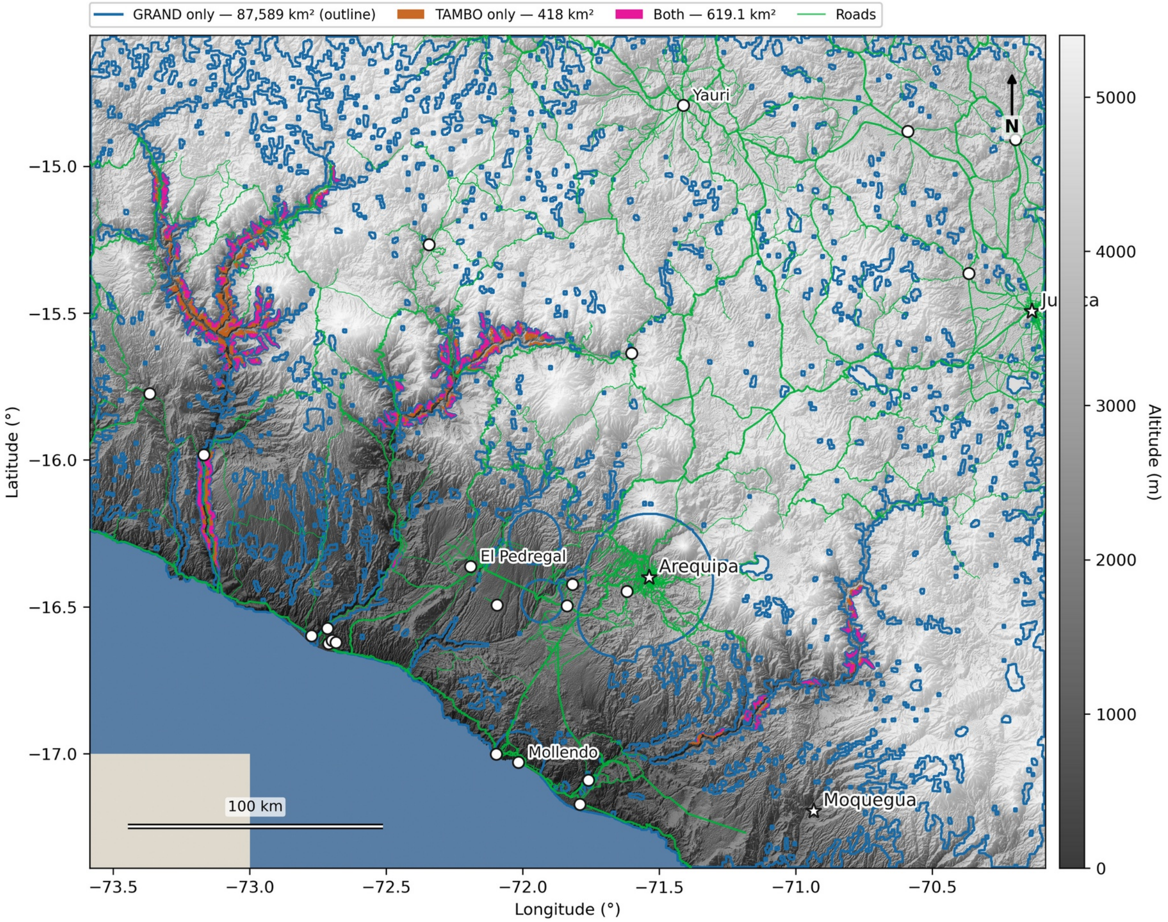

Roads and place names © OpenStreetMap contributors (ODbL). Elevation: SRTMGL1 via OpenTopography.

Every run explains itself

Oroscope output includes: a map, **GeoTIFF mask**, **KML** (for Google Earth), **results JSON**

Plus, plain-language account of the results (**explanation.txt**): what was found, which stage bound it, *etc.*

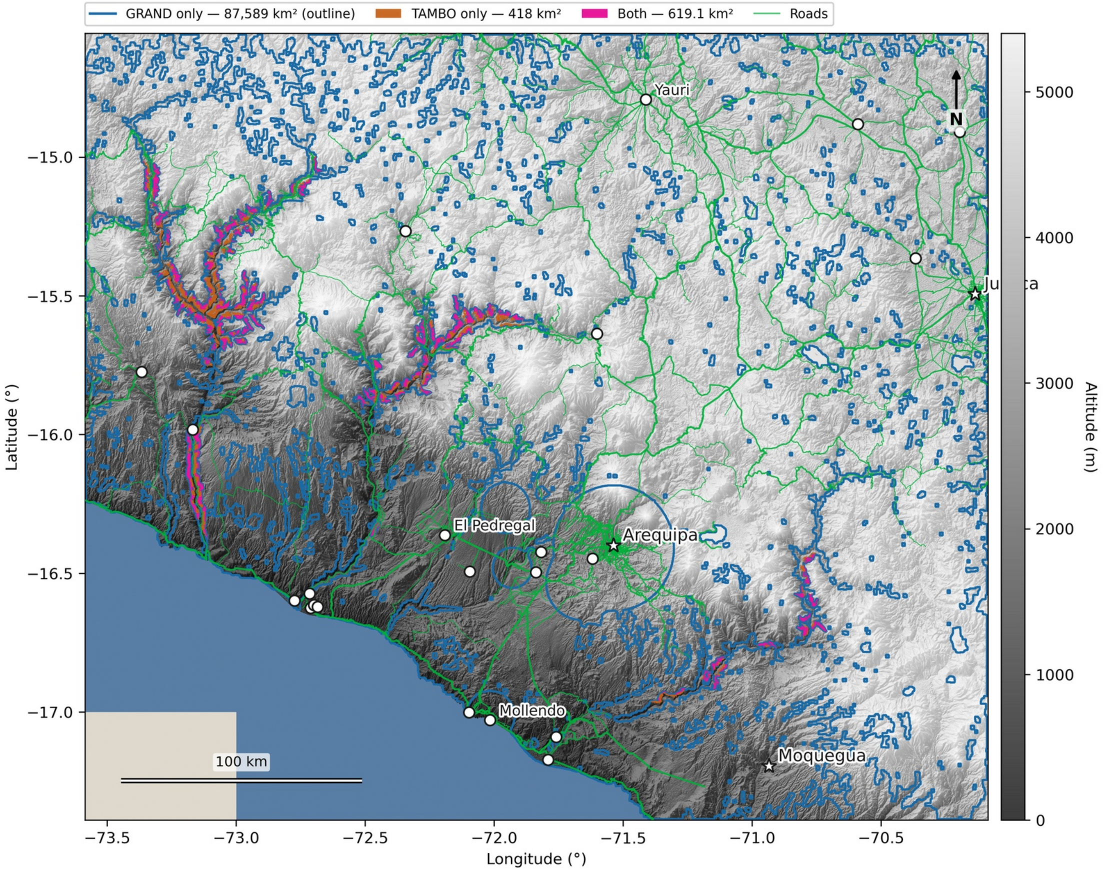


```
================================================================================
 WHERE THESE EXPERIMENTS CAN SHARE GROUND, AND WHY
================================================================================
WHAT EACH BRINGS
----------------
  experiment |     area km² |   sites |  detectors |  in joint
  ------------------------------------------------------------
  GRAND      |     88,208.2 |       1 |    101,584 |      0.7%
  TAMBO      |      1,036.9 |      85 |     49,271 |     59.7%
  Ground satisfying GRAND and TAMBO at once: 619.1 km². Ground useful to any
  of them: 88,626.0 km². The first is the number that matters for one site,
  one road and one power feed serving two experiments; the second is what the
  programme as a whole could use.

WHAT DECIDES IT
---------------
  A pixel has one slope, one altitude and one aspect, and both experiments
  must accept those same values to stand on it. So co-location is settled by
  whichever of those bands they share least of — before any arrival geometry
  is considered.
  band                 |            GRAND |            TAMBO | shared
  -----------------------------------------------------------------------
  deployable slope     |            3–25° |           20–60° | 20–25°  (23%)
  “deployable slope” is what limits the sharing: they overlap only over
  20–25°, which is 23% of the narrower of the two bands. Co-location is
  decided there, not by anything about the physics of either experiment.
  What they ask of the *view* differs, and that is no obstacle: two
  experiments can look out from the same hillside at different ranges and
  different elevations without conflict.
  target distance     | GRAND: 10–40 km, TAMBO: 2–5 km
  arrival elevation   | GRAND: -3–3°, TAMBO: -20–20°

PAIRWISE
--------
  GRAND & TAMBO: 619.1 km² shared, Jaccard 0.0070.
        0.7% of GRAND's own ground
       59.7% of TAMBO's own ground

HOW TO READ THIS
----------------
  Every caveat on the individual runs applies here and compounds: these are
  the same masks, so morphological closing has already moved each area, and
  the joint of two inflated masks is inflated twice over. The overlay is
  exact — the masks are pixel-aligned and the alignment is checked, not
  assumed — but what it overlays is only as good as each run.
  And co-location is a question about ground, not about detectors. Two
  experiments sharing a hillside still need their own arrays, their own
  spacing and their own trigger; what they share is the site, the access and
  the power.
```

Colca (Arequipa) zoom – placing **GRAND**

4,446 km² — most of the frame, plateau and gentle ground.

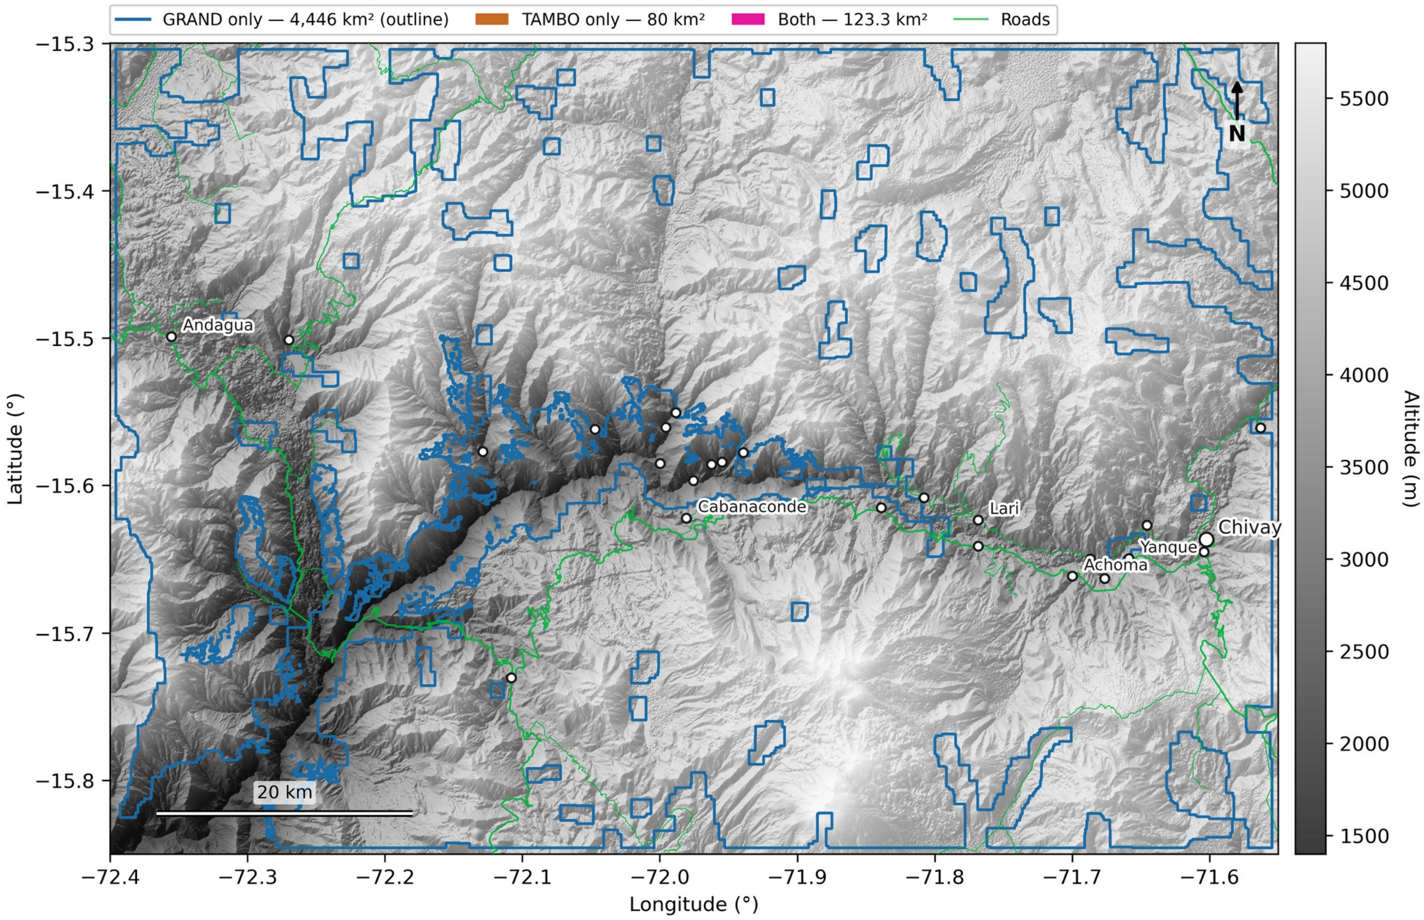

Colca (Arequipa) zoom – placing **TAMBO**

80 km² — a scatter of patches threading the canyon system.

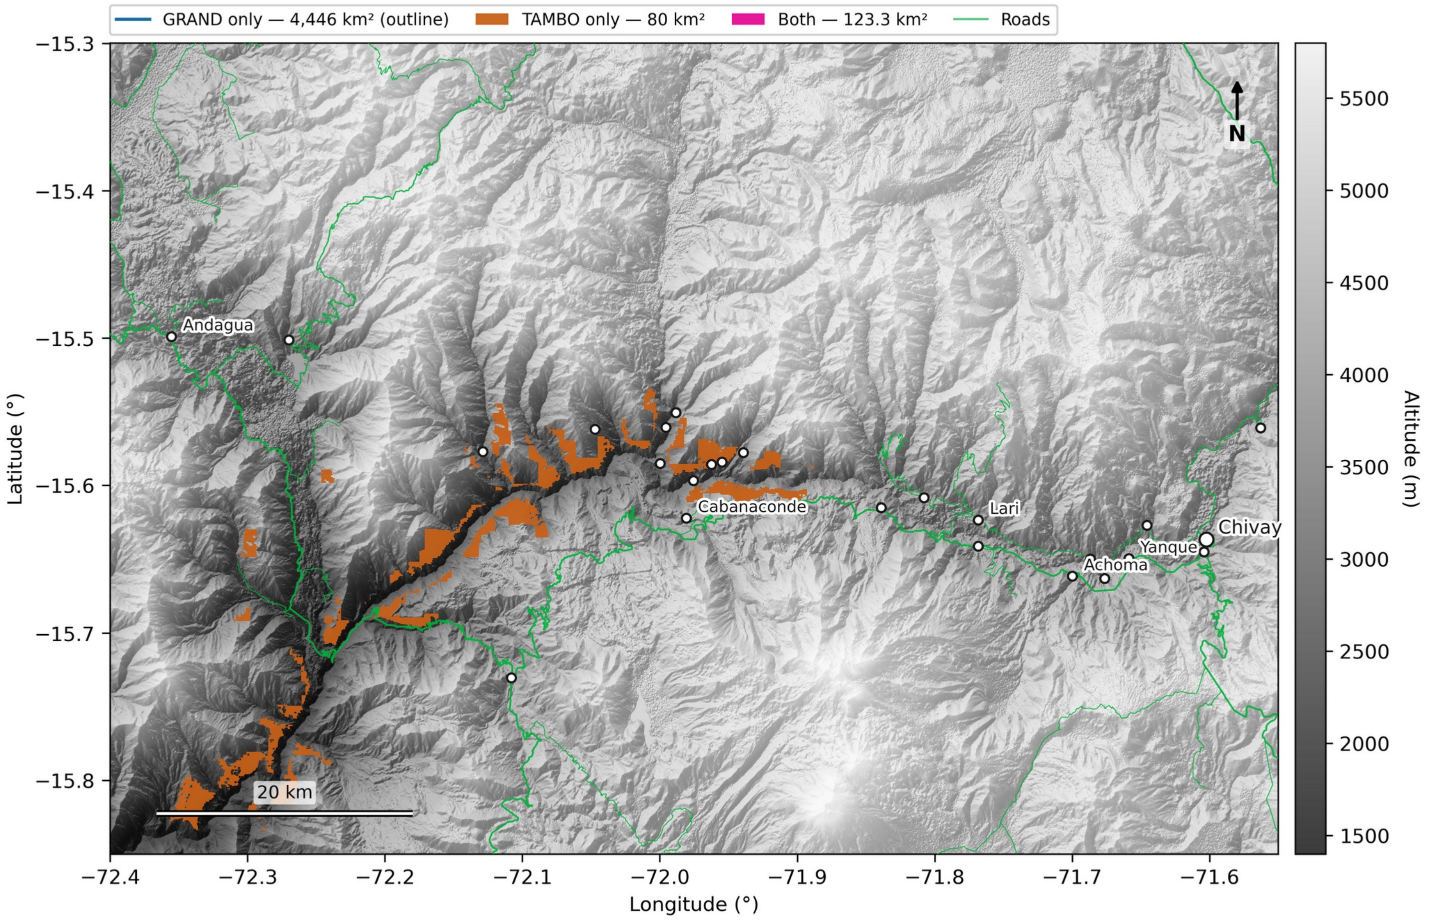

Roads and place names © OpenStreetMap contributors (ODbL). Elevation: SRTMGL1 via OpenTopography.

Colca (Arequipa) zoom – placing **GRAND + TAMBO**

123.3 km² shared — **60.7%** of TAMBO's 203 km², tracing the canyon rims — where the roads also run.

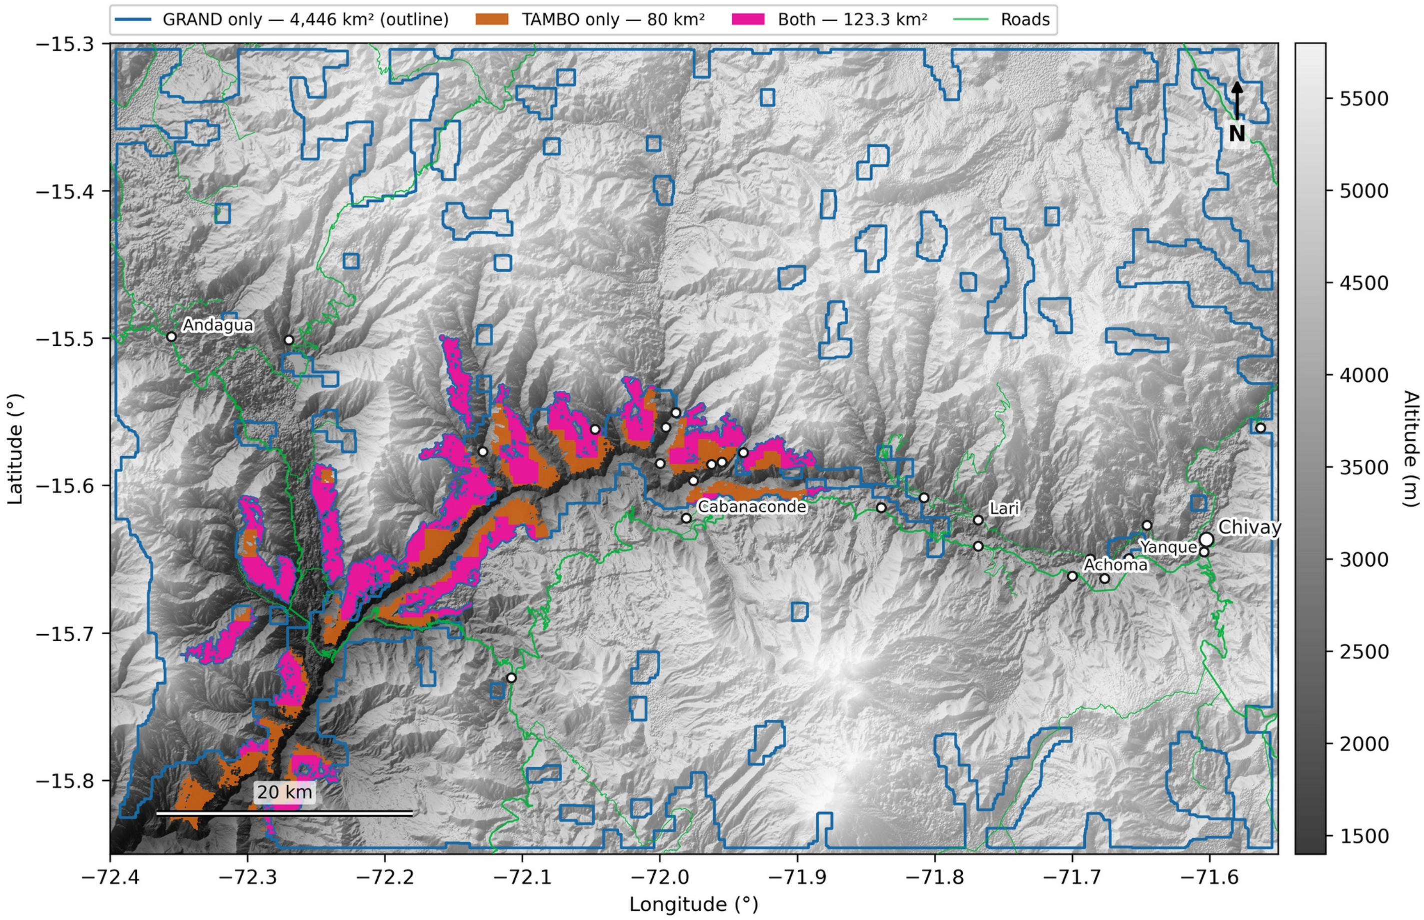

Roads and place names © OpenStreetMap contributors (ODbL). Elevation: SRTMGL1 via OpenTopography.

Áncash, in full – placing **GRAND + TAMBO**

411.1 km² shared — **55.6%** of TAMBO's 740 km².

Subsampled, so a lower bound.

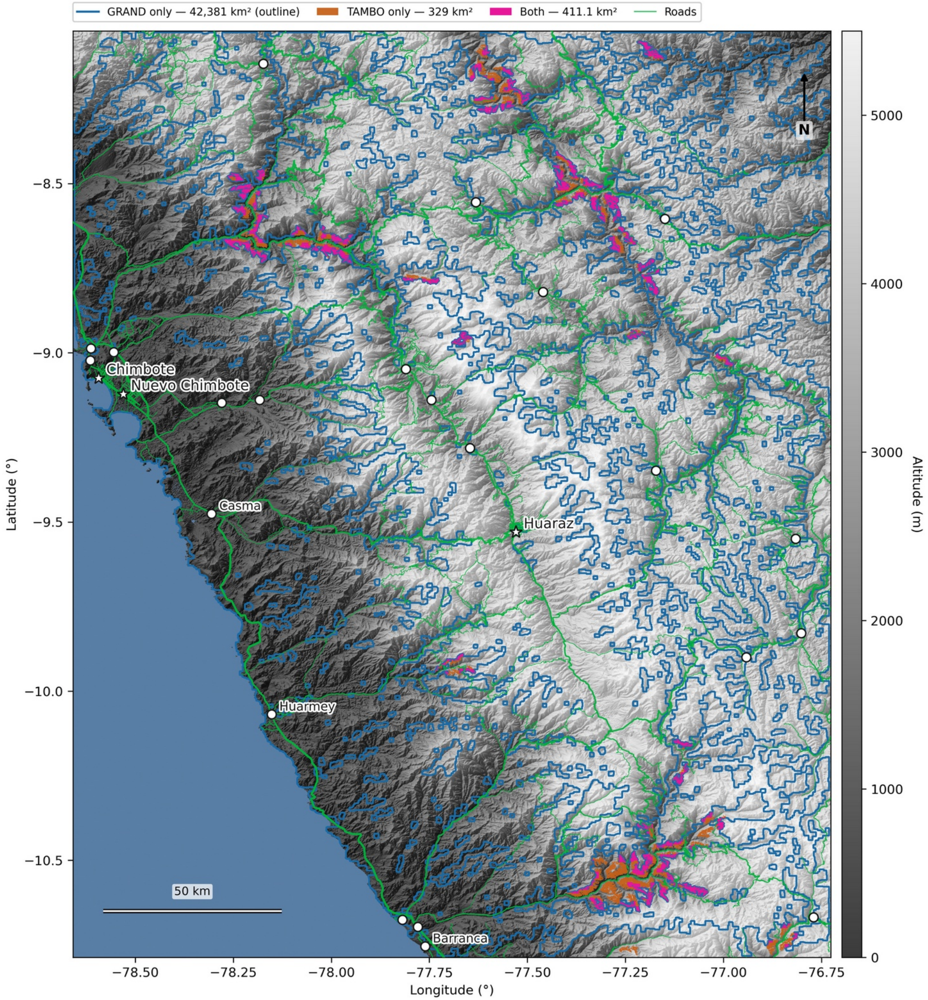

Callejón de Huaylas (Áncash) – placing **GRAND + TAMBO**

No subsampling

228.7 km² of TAMBO's 291.3 km² is also GRAND-viable: **78.5%** (highest measured).

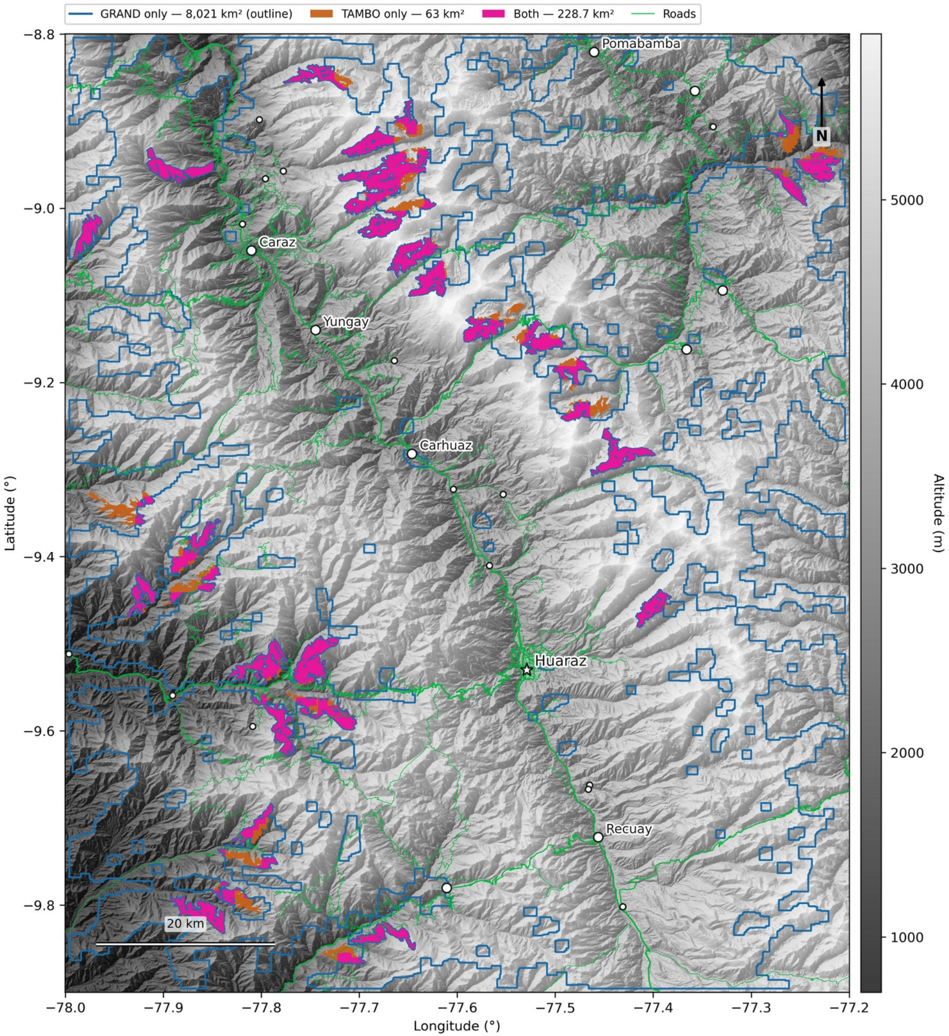

Lima, in full – placing **GRAND + TAMBO**

509.8 km² shared — **55.7%** of TAMBO's 915 km².

Subsampled, so a lower bound.

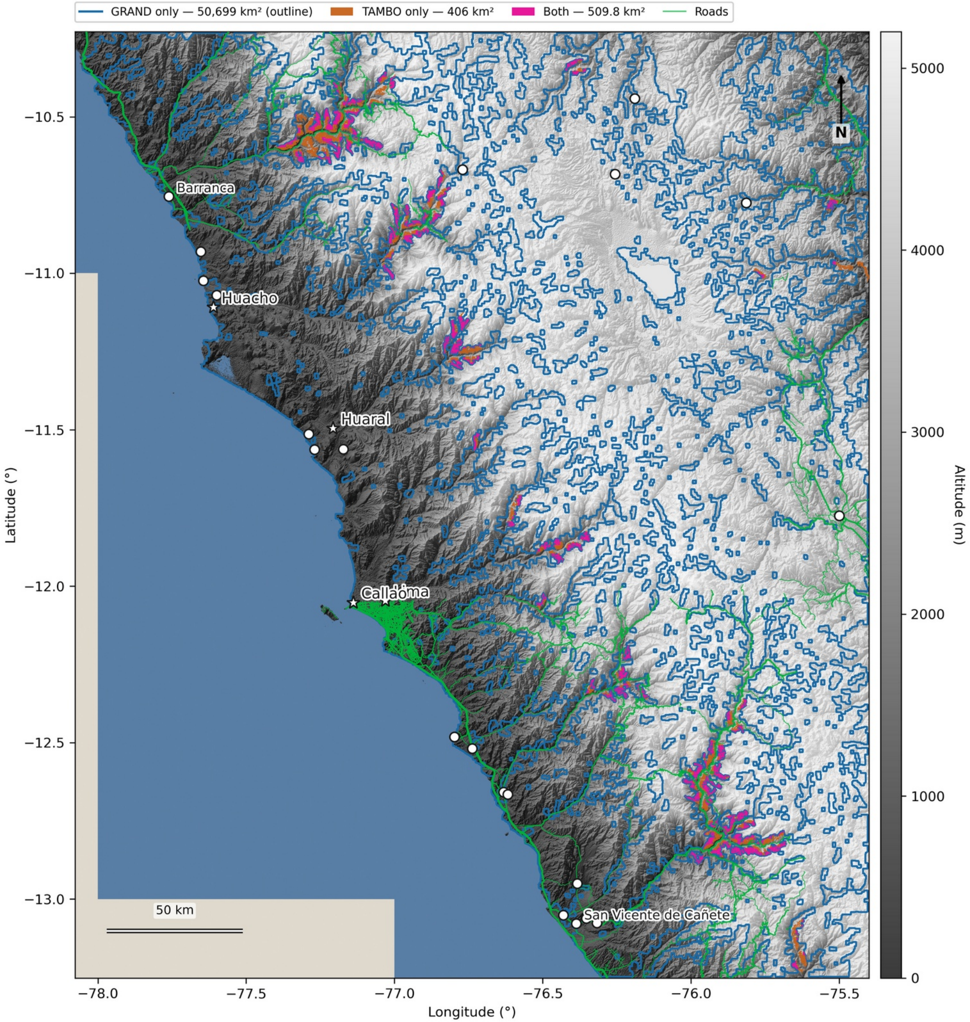

Roads and place names © OpenStreetMap contributors (ODbL). Elevation: SRTMGL1 via OpenTopography.

Cajatambo (Lima) – placing **GRAND + TAMBO**

509.8 km² shared — **55.7%** of TAMBO's 915 km².

Subsampled, so a lower bound.

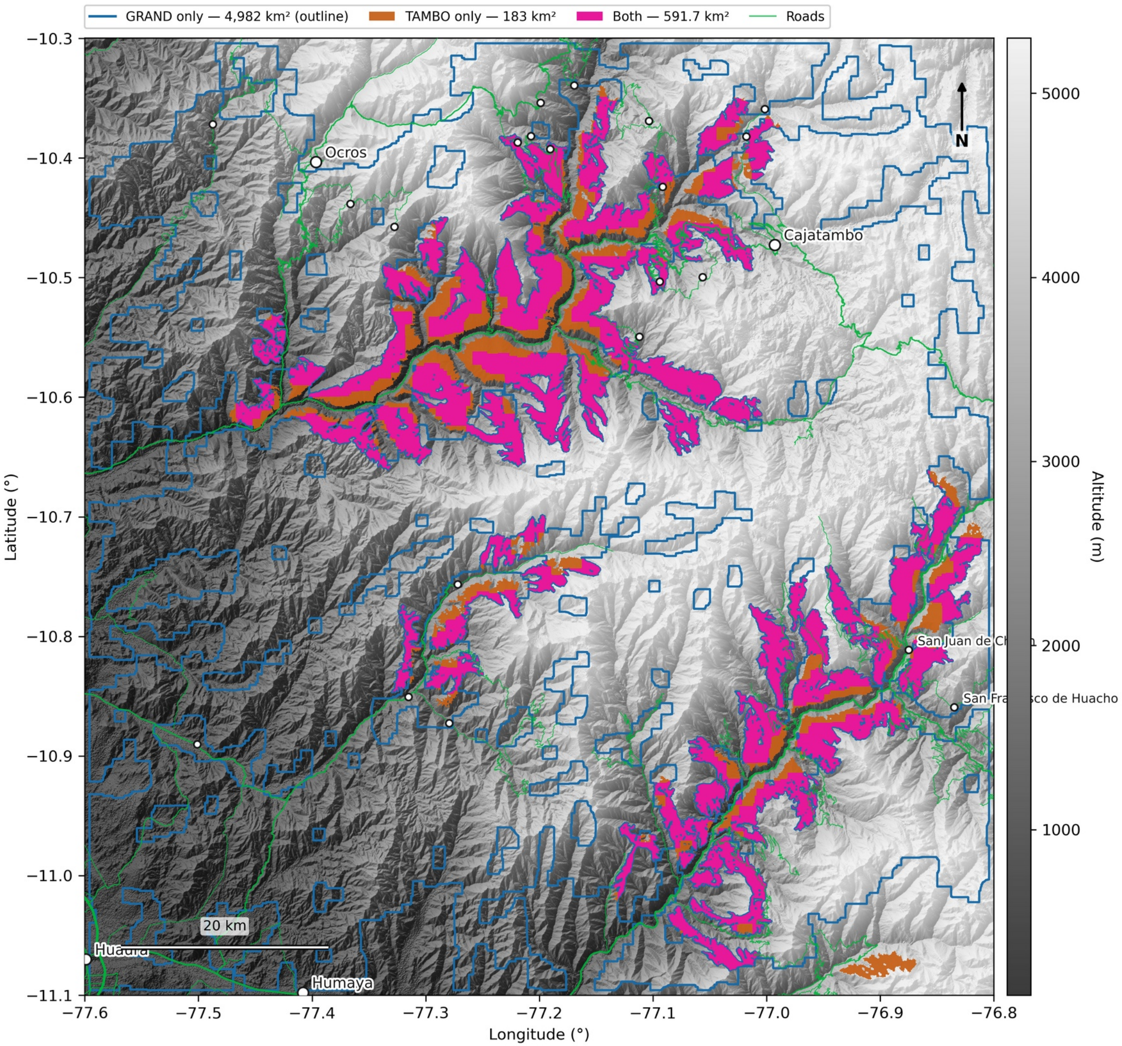

All of Peru

Only on coarse 90-m grid (SRTMGL3, 339 Mpx):

400,000–600,000 km² of GRAND-viable ground.

The slope screen is the binding constraint.

There is no national TAMBO or joint answer yet. It needs the 30-m grid over the whole country, plus tiling; that has not been run.

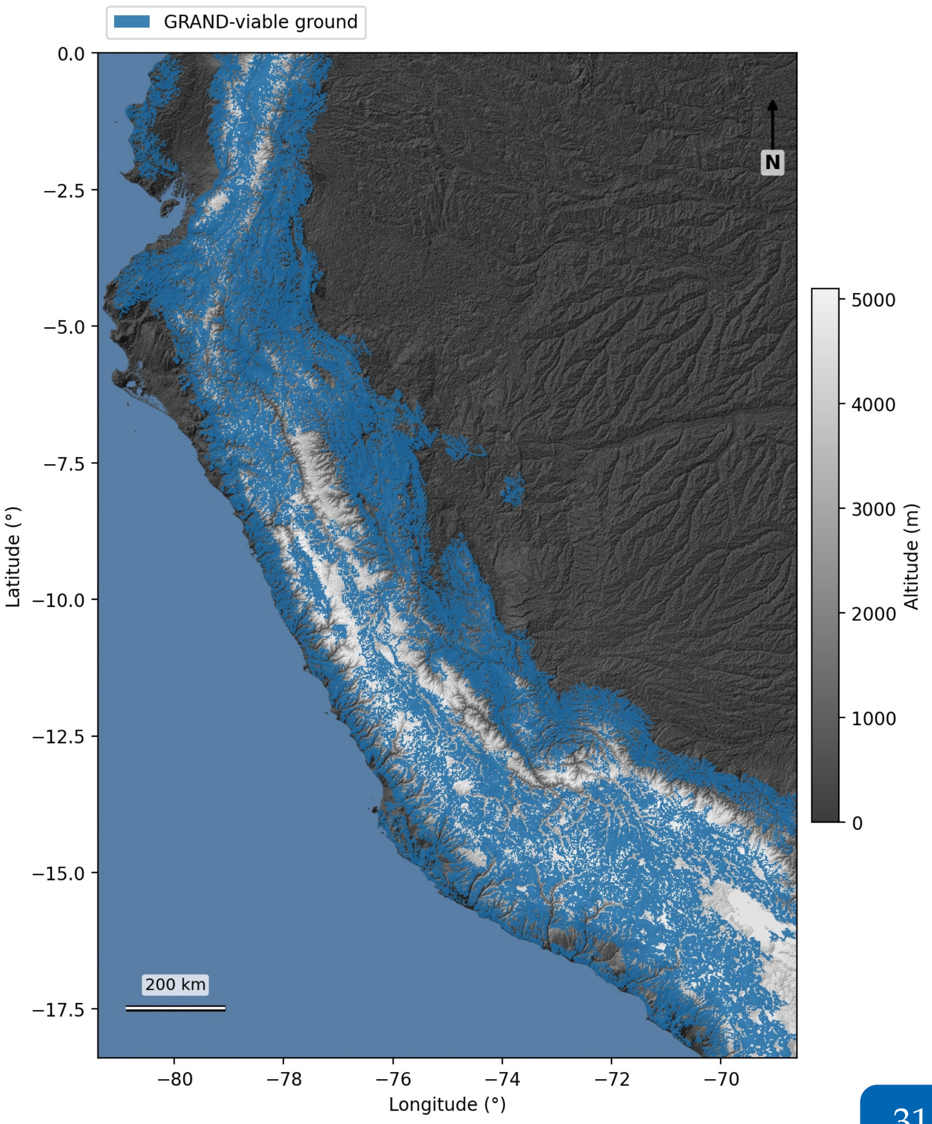

Elevation: SRTMGL3 via OpenTopography. Mask: the stored survey raster, redrawn.

Shared area in Arequipa

The **619 km²** of shared GRAND + TAMBO ground, divided by each experiment's own:

0.7%

619 / 88,208 km²

of GRAND's viable ground

59.7%

619 / 1,037 km²

of TAMBO's viable ground

Co-location costs GRAND almost nothing, and consumes most of what TAMBO has.

The shared ground is TAMBO-limited in every region measured — so the site choice is TAMBO's to lead, and GRAND's to join.

TAMBO takes the canyon, GRAND stands a ways off

The shared ground is canyon wall: median slope 35° at Colca

At 150-m spacing, it holds 6,322 TAMBO units (enough for full TAMBO)

Ground gentle enough for GRAND is 10 km² of slivers there, about 8 antennas' worth.

GRAND was never meant to stand on it

What couples the arrays is a shared line of sight to the same massif — its lattice lies a median 1.0 km from the shared ground (83% of it within 3 km)

From each region's best TAMBO site: 100 antennas fit within ~10 km, 1,000 within 20–30 km, 5,000 within 60 km.

Six regions explored

Region	GRAND km²	TAMBO km²	Joint km²	Share of TAMBO
Huaylas	8,249.5	291.3	228.7	78.5%
Cajatambo	5,573.8	774.5	591.7	76.4%
Colca	4,569.4	203.0	123.3	60.7%
Ancash	42,791.9	740.0	411.1	55.6%
Lima	51,209.0	915.4	509.8	55.7%
Arequipa	88,208.2	1,036.9	619.1	59.7%

Share of TAMBO = joint area ÷ TAMBO's own viable area — the fraction of TAMBO-viable ground that GRAND could also stand on.

Shaded rows: every pixel was tested. The rest were subsampled to fit in memory — read those as lower bounds.

Roughly **three quarters to four fifths** of TAMBO-viable ground is also GRAND-viable.

How many detectors is that?

Spacing is the input (GRAND 1 km, TAMBO 150 m). **Capacity is computed**: a hexagonal lattice is packed into each region's mask, position by position — not area ÷ spacing^2.

Region	GRAND antennas	TAMBO units
Huaylas	9,559	14,925
Cajatambo	6,457	39,658
Colca	5,315	10,437
Ancash	49,059	34,275
Lima	58,669	42,549
Arequipa	101,584	49,271

TAMBO's design target is ~5,000 units — Colca alone holds twice that. Every region measured clears it. **GRAND's 10,000-antenna design** fits in any of the three departments; no single crop holds it. Shaded rows unbiased, the rest subsampled — the same lower-bound caveat as the areas.

What this does not tell you

What is missing, and what's next

Nothing has been checked against an external simulation. Everything here is internally consistent and externally unvalidated. The cheapest test is the Earth-absorption prediction — an arrival-window edge of **−4.4° at 100 PeV rising to −0.9° at 10 EeV** — and it has not been done.

Detector acceptance A(E) is modeled poorly. Two published *integral* curves are digitized, and an array-size correction cannot correct for the site.

Neutral-current regeneration is not modeled, so Earth-chord suppression is somewhat overstated.

Magnetic declination falls back to the Arequipa value rather than being per-site.

A **national TAMBO answer** needs the 30 m grid over the whole country, and tiling — the Peru map you saw is GRAND-only.

Price tag

1 GRAND DU ~ 3.5k € ~ 1 TAMBO unit

What is the minimum number of GRAND DUs needed?

About 10 (for five-DU coincidence triggers), so ~35k €

Take-home

- Two detectors that want opposite hillsides do not want opposite **views**.
- Co-location limited by **slope**, not by sight: bands overlap only over 20–25°.
- Even so, **three quarters to four fifths** of TAMBO-viable ground is also GRAND-viable.
- It costs GRAND **0.7%** of its ground and gives TAMBO a partner within **10 km** of its best sites.

The terrain does not forbid it. The question worth asking next is not **whether** they can share — but **which massif**.

Thank you!

Backup slides

HERON: Hybrid Elevated Radio Observatory for Neutrinos

BEACON phased-array stations (low-SNR detection) + GRAND antennas (self-triggering)

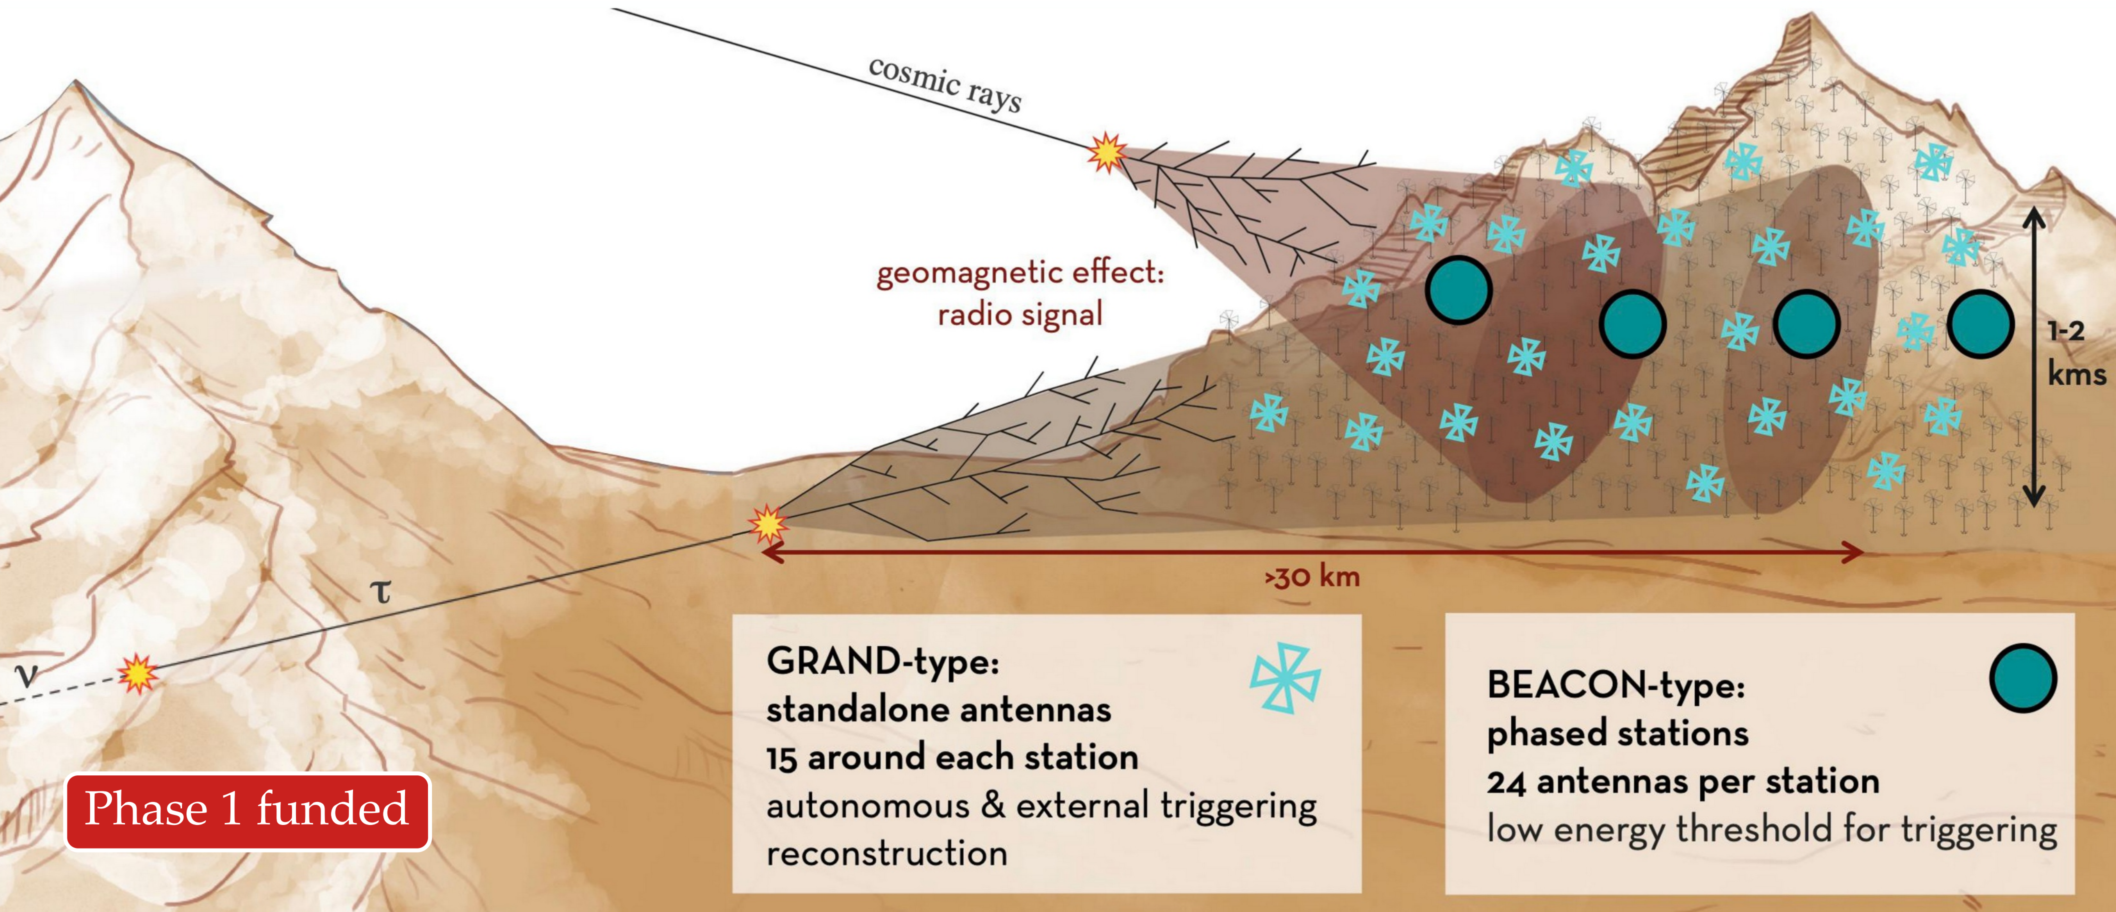

HERON: Hybrid Elevated Radio Observatory for Neutrinos

BEACON phased-array stations (low-SNR detection) + GRAND antennas (self-triggering)

Narrow field of view, high sensitivity:
Phase 1: 24 BEACON stations + 360 GRAND antennas [≈ GRAND10k sens. for point sources]
Phase 2: 200 BEACON stations + 3k GRAND antennas [≈ GRAND, IC-Gen2 sens. for diffuse flux]

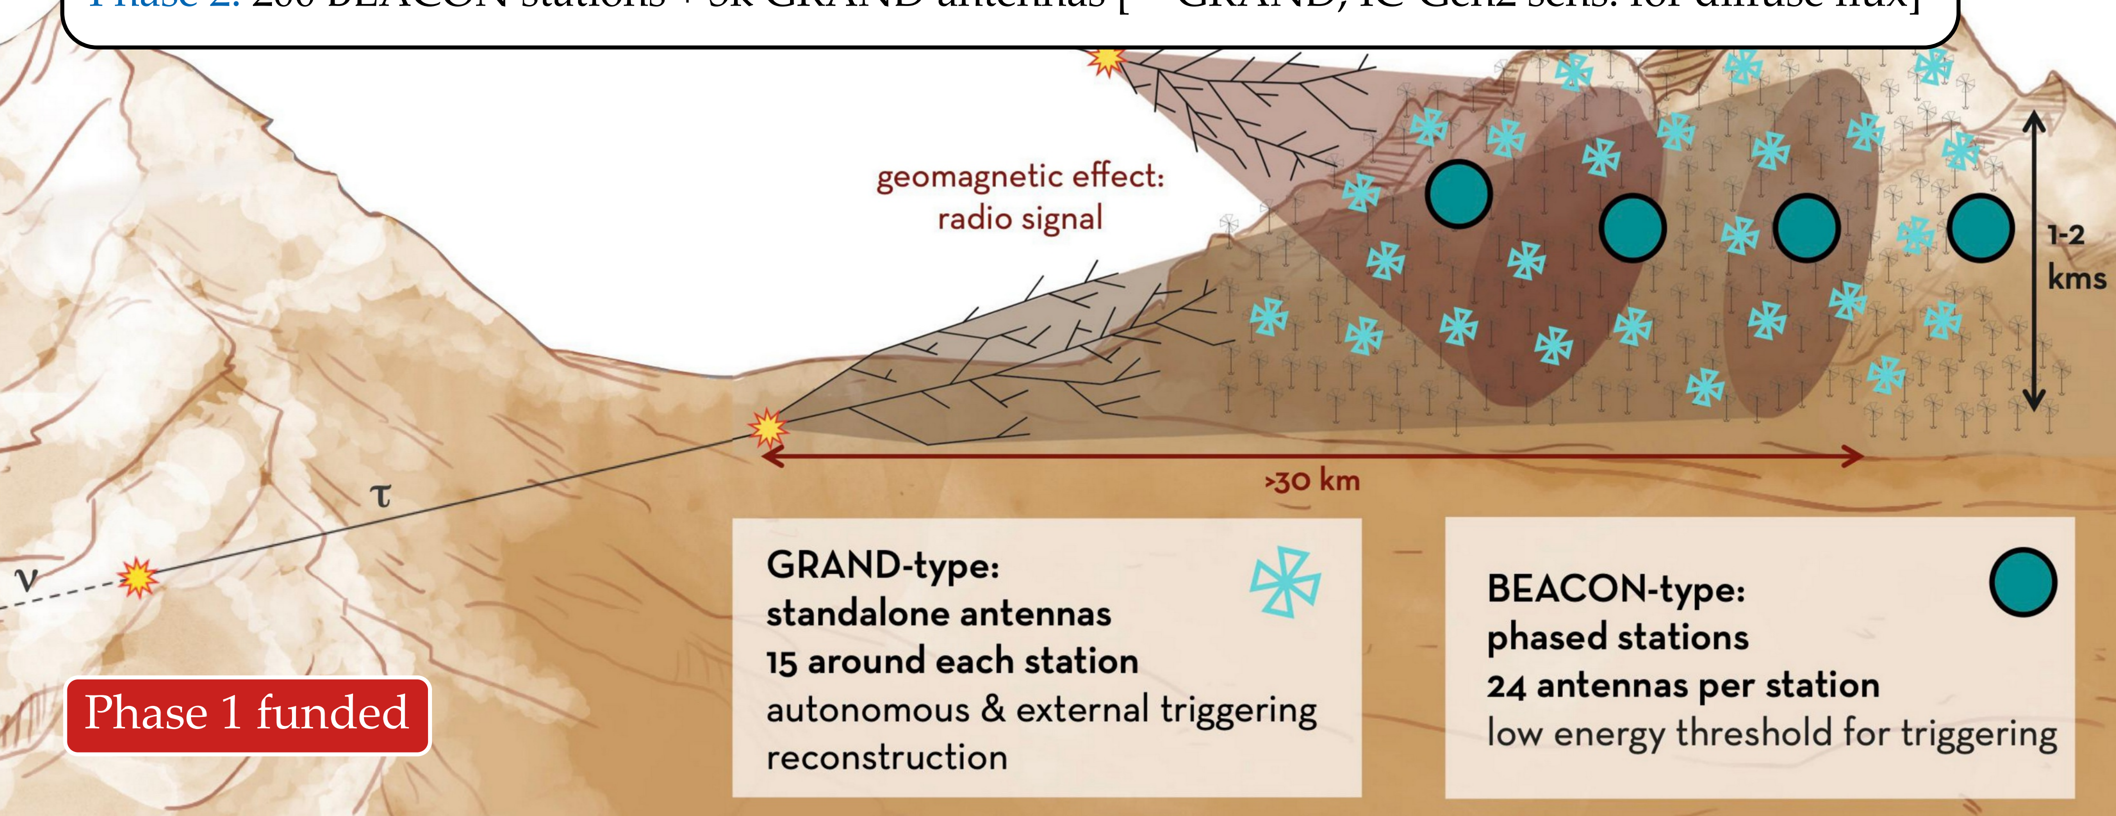

Phase 1 funded

Partner-array distances, per region

How far from each region's best TAMBO site a GRAND array of a given size fits.

From each region's best TAMBO site	100 antennas	1,000	5,000
Arequipa	10 km	20 km	60 km
Ancash	10 km	30 km	60 km
Lima	10 km	30 km	60 km

The limit is only that 1,000 antennas at 1 km spacing need 866 km^2, and steep country does not offer that within a few kilometres. These are capacity counts from packing a lattice, not optimised layouts.

Four caveats

1 **Reported area is an upper bound.** Closing — the stage that fills the holes striding leaves — inflates the mask before area is measured: **2.35×** at Colca, against a stride-1 control.

2 **A subsampled TAMBO area is a lower bound** — by **1.51×** at Colca, **23.0×** on the Callejón de Huaylas. Striding does not bias which pixels pass; the loss is in rebuilding the map, where a fragmented mask fails the minimum-array-size cut at pruning.

3 **min_score is the dominant assumption — and it is TAMBO's alone.** GRAND's configurations cut nothing. The shipped 0.35 sits at the **17.8th percentile** of a product that piles up near zero, and a sweep across the cut finds **no knee anywhere** — every TAMBO area here moves if you move it.

4 **solid_angle is the weakest component at every selected TAMBO site in every region** — 16/16 Colca, 62/62 Ancash, 84/84 Lima, 85/85 Arequipa. Under a product, the weakest term bounds the site — so the TAMBO result rests on that one component.