Earth Dynamics

Multi-Modal Geologic Intelligence

3D Inversion and Map Synthesis via Generative Foundation Models

Michal Holtzman Gazit

IMVC 2026 · 3D Vision Session

The Problem: A Fragmented Earth

Earth Dynamics

Sparse surveys

Patchy geophysical coverage with large gaps between surveyed zones.

Coarse legacy maps

Regional gravity and magnetics only at km-scale resolution.

Unstructured reports

Decades of expert observations locked in legacy text archives.

Three vision problems in disguise:
3D completion · conditional generation · super-resolution

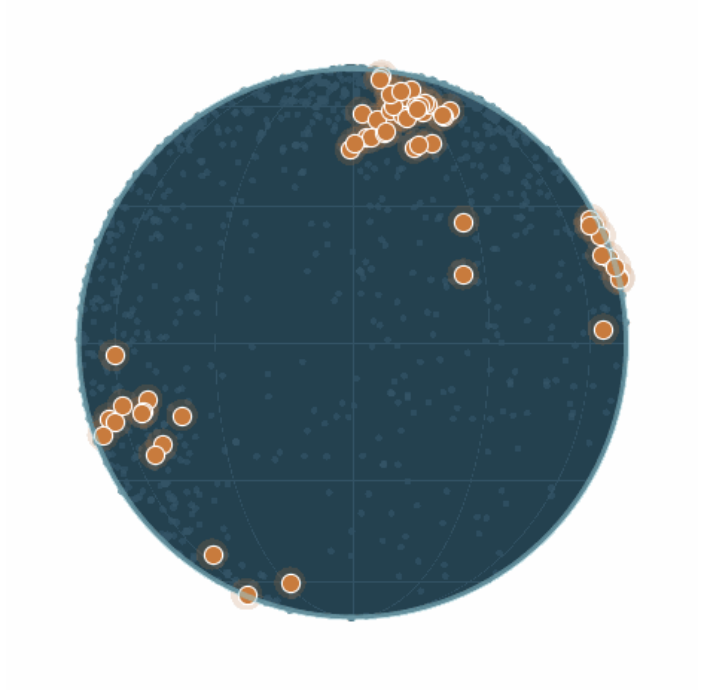

Sparse survey coverage — vast “white spaces” between data

The 3D Borehole Inversion Challenge

- **Boreholes are sparse, irregular 1D samples of a 3D world**
- Observed voxels: ~0.001% of the target volume
- Massively ill-posed and non-unique — infinitely many volumes fit the data
- Goal: reconstruct the full 3D geological model, with uncertainty

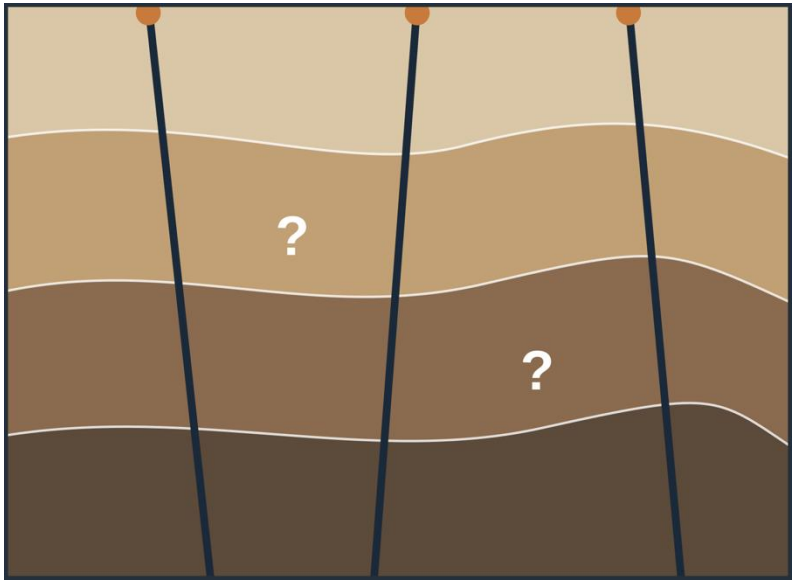

A handful of 1D drill lines through an unknown 3D volume

Simulation-Driven Generative Inversion

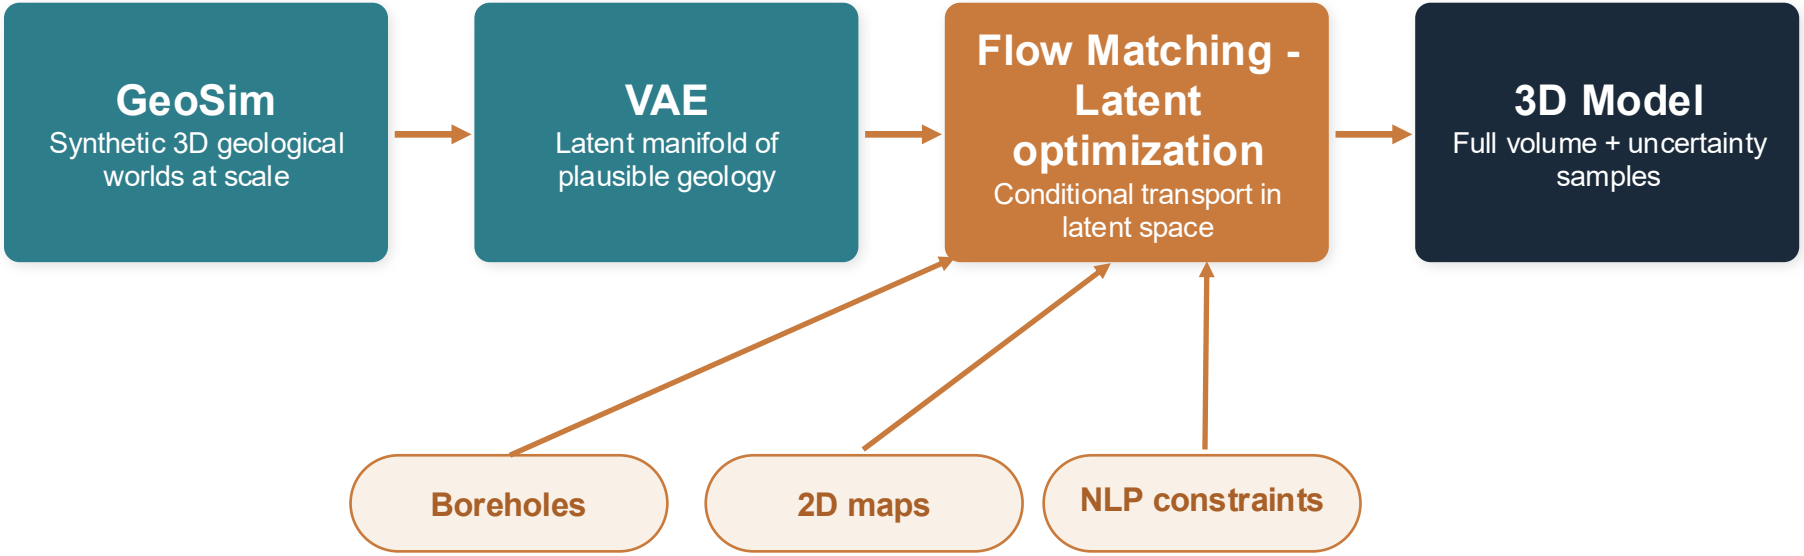

GeoSim: A 3D Time Machine for the Subsurface

Earth Dynamics

We can’t see through miles of rock.

GeoSim uses generative modeling to build high-fidelity 3D simulations of what the ground could look like — not random drawings, but geologically and physically plausible worlds.

- **Plausible, not arbitrary** — Honors stratigraphy, folding and faulting.
- **The space between holes** — Visualizes structure where no data exists.
- **Unlimited training worlds** — Endless paired data for the generative models.

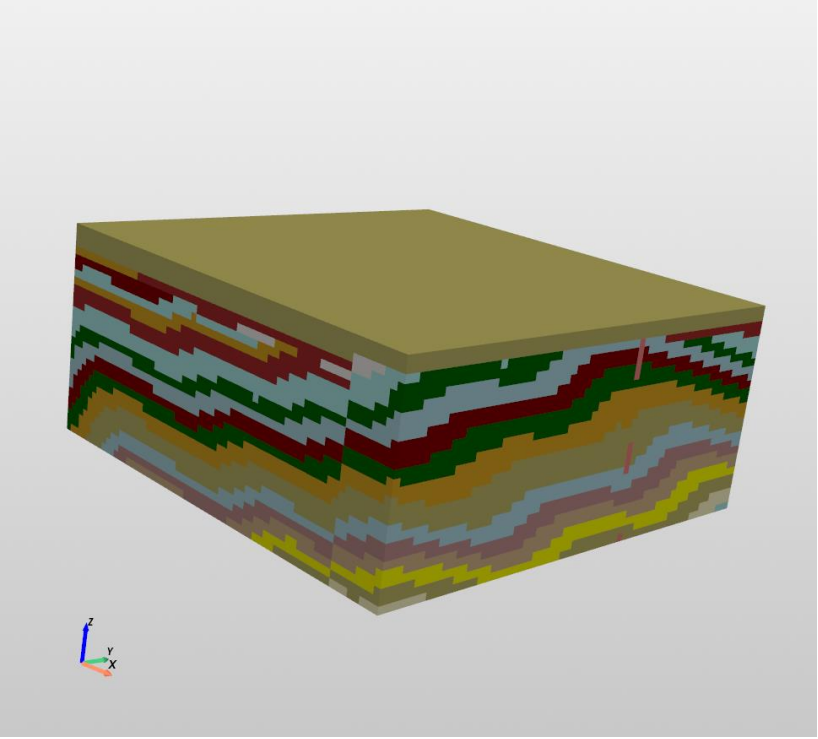

Example GeoSim volume — folded, faulted strata generated by the engine

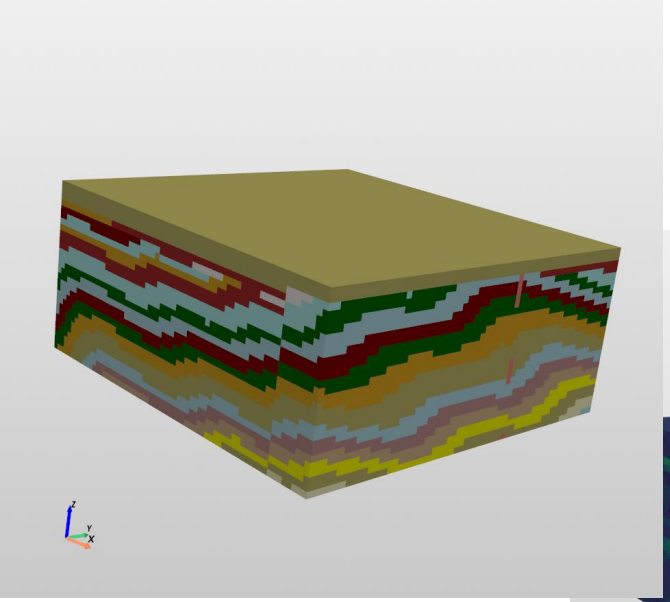

Lithology

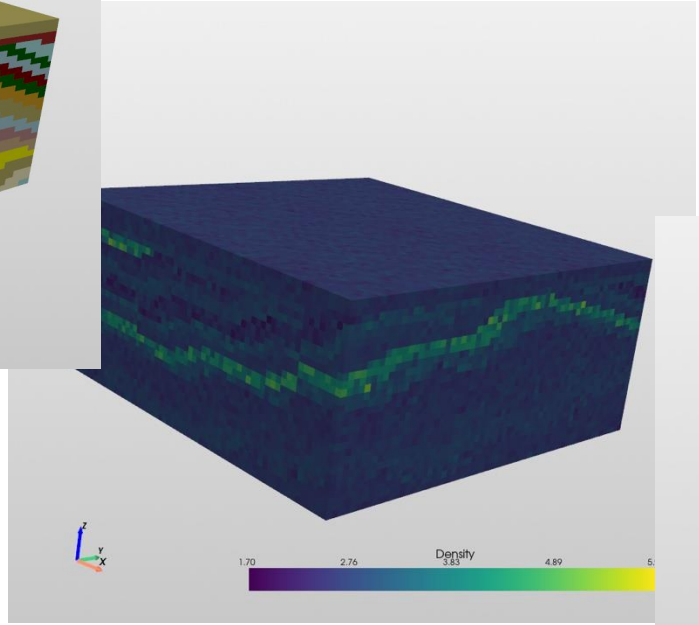

Mass density

Magnetic susceptiblity

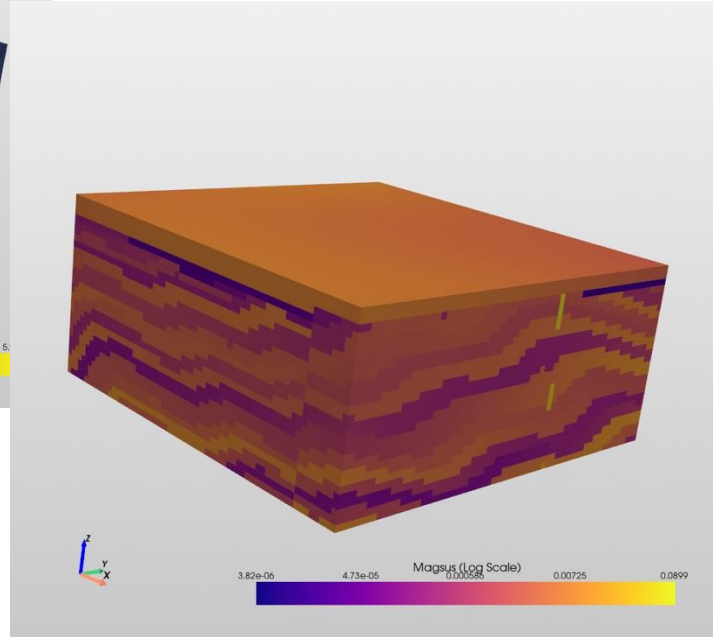

GeoSim: Synthetic Ground Truth at Scale

Earth Dynamics

Labeled 3D ground truth barely exists in the real world. GeoSim generates vast libraries of geologically plausible 3D environments — one simulation engine feeding both capabilities of the platform.

Unlimited virtual boreholes

Drill anywhere in simulation → perfect supervised (boreholes, volume) pairs for 3D inversion.

Paired low-resolution / high-resolution maps

Forward-model gravity and magnetics at multiple resolutions → training data for map super-resolution.

Physically consistent modalities

Seismic, gravity, magnetics and surface geology generated from the same world — the model learns true cross-modal structure.

VAE + Flow Matching in Latent Space

Earth Dynamics

VAE — the manifold of plausible geology

- Compresses 3D geological volumes into a structured latent space
- Generation happens on the manifold — not in voxel space
- Implausible reconstructions excluded by construction → tames non-uniqueness

Conditional Flow Matching — inference

- Learns a velocity field transporting noise → latent codes, conditioned on observations
- Deterministic and stable trajectories
- Far fewer function evaluations than diffusion → fast 3D reconstruction

Results: 3D Inversion from Boreholes

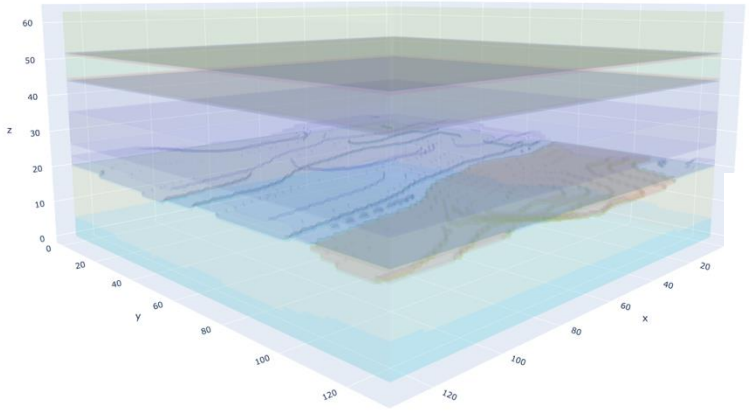

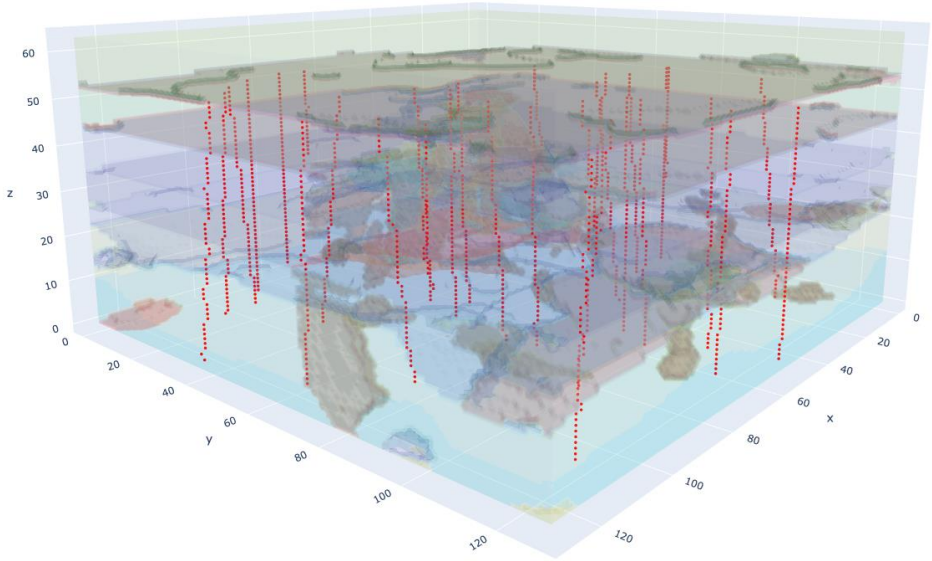

Results: 3D Inversion from Boreholes

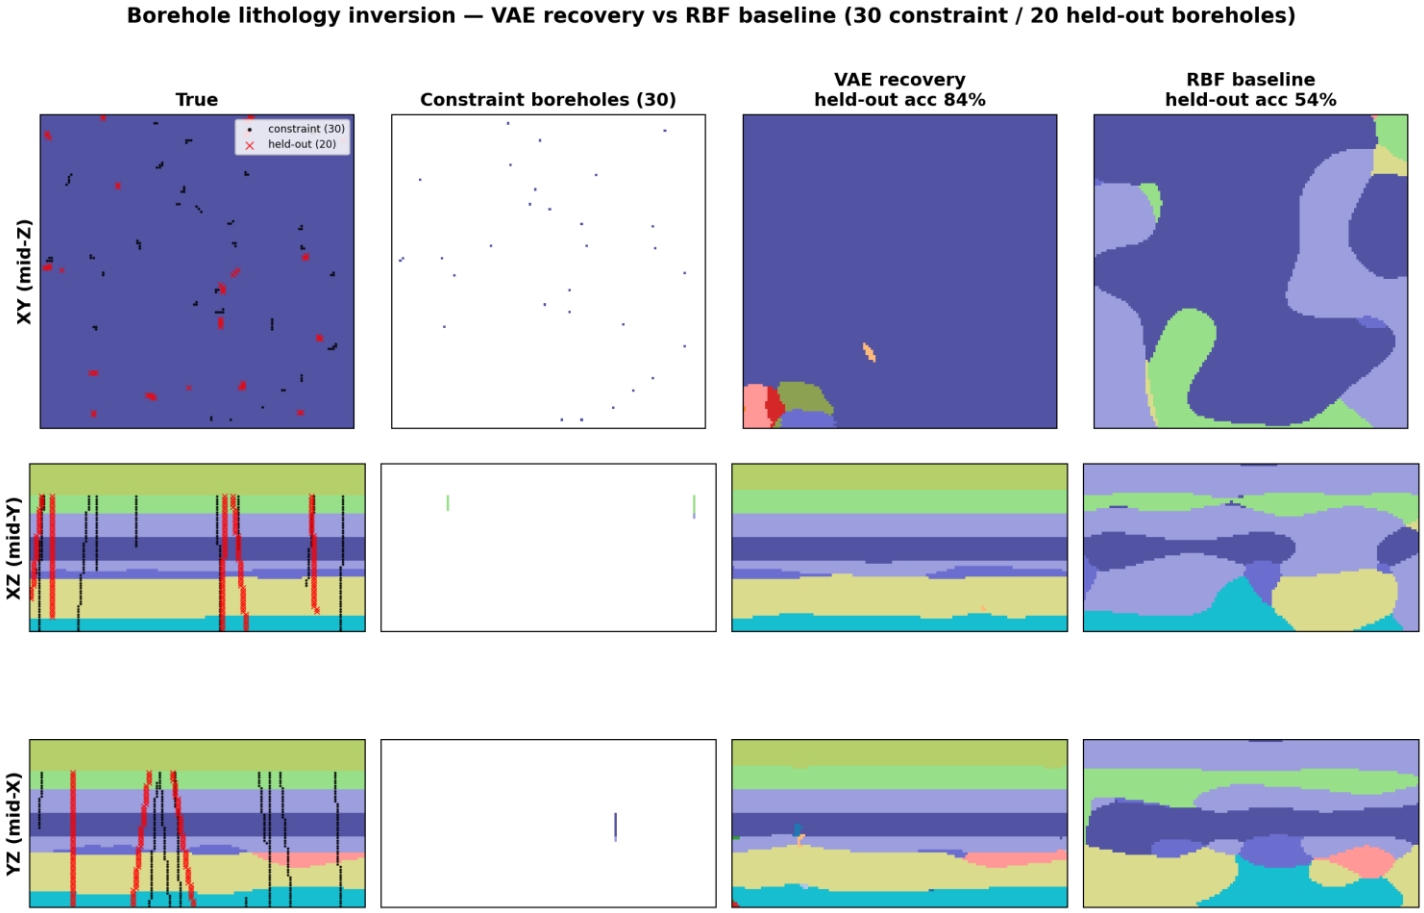

Results: 3D Inversion from Boreholes

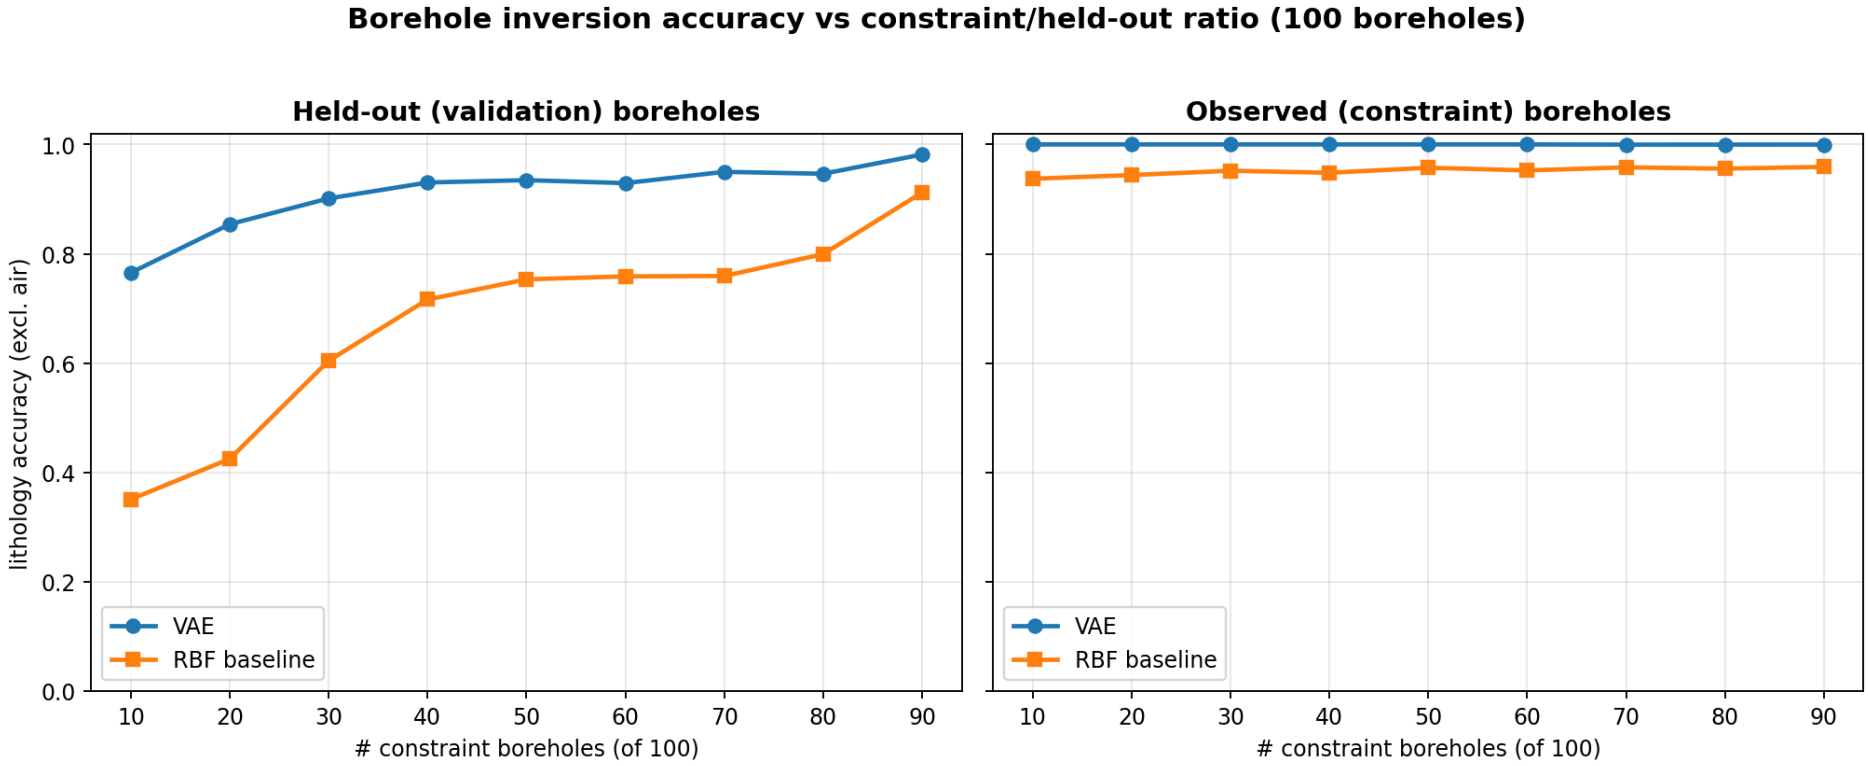

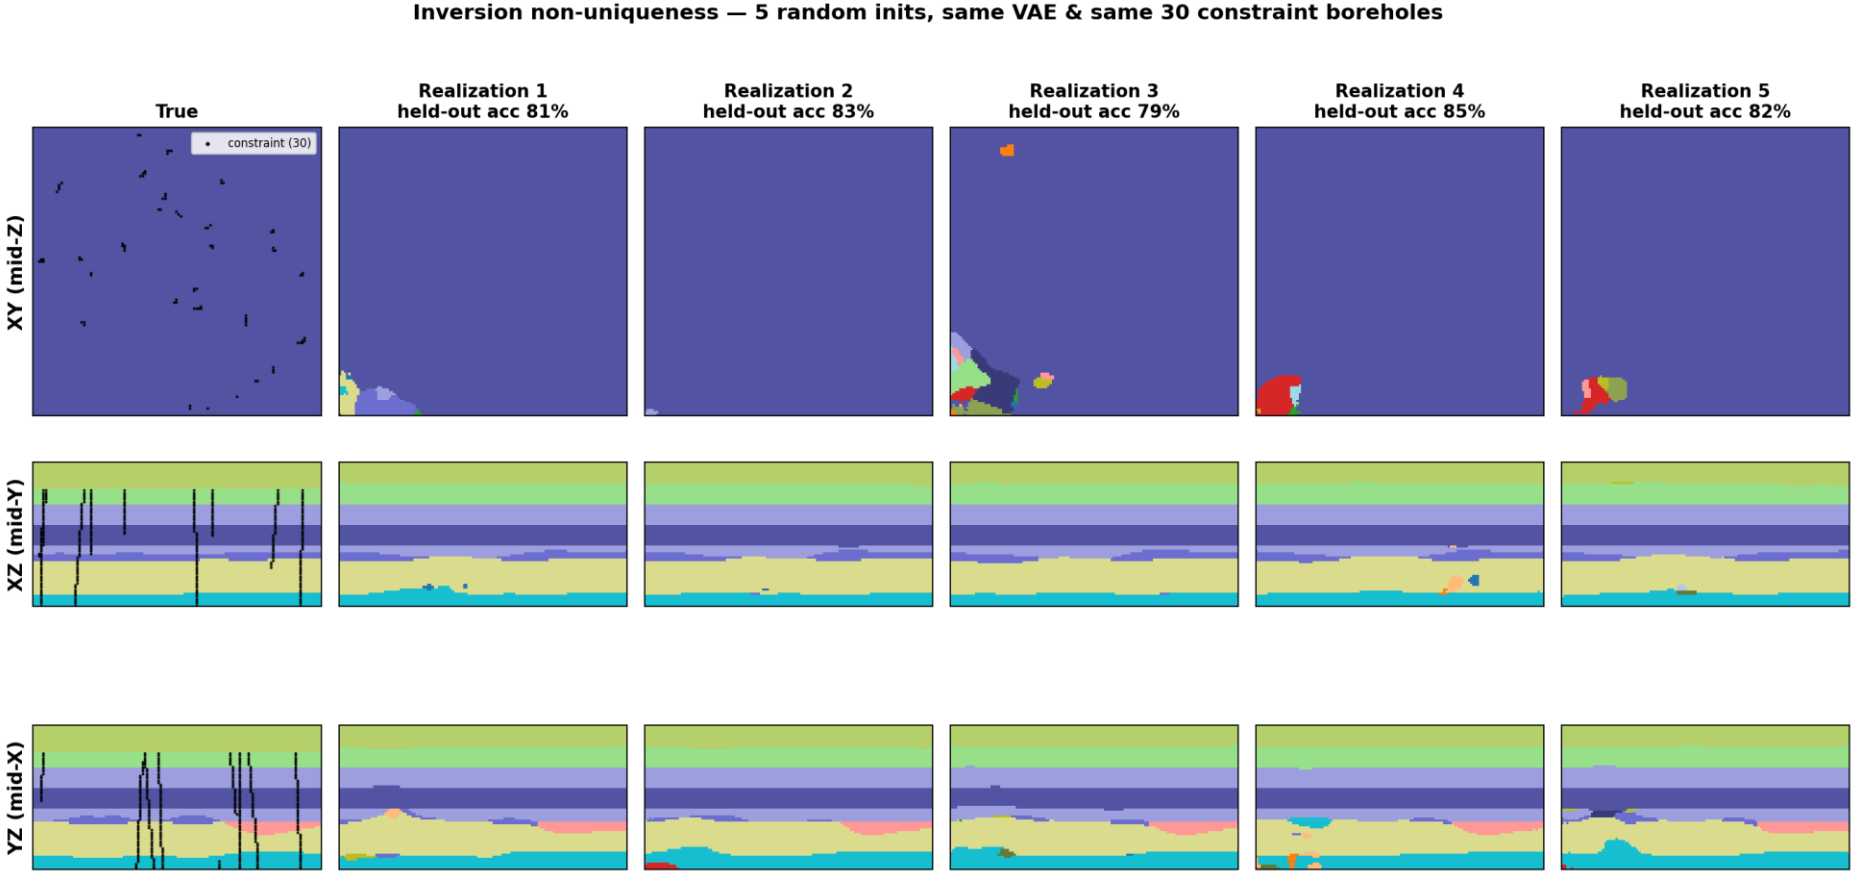
Inversion non-uniqueness — 5 random inits, same VAE & same 30 constraint boreholes
True
constraint (30)
Realization 1
held-out acc 81%
Realization 2
held-out acc 83%
Realization 3
held-out acc 79%
Realization 4
held-out acc 85%
Realization 5
held-out acc 82%
XY (mid-Z)
XZ (mid-Y)
YZ (mid-X)

New: Super-Resolution for Gravity Maps

Earth Dynamics

The same gaps exist in 2D — regional gravity is coarse. But two complementary high-resolution sources are available:

Auxiliary physical maps

Magnetics, topography and other correlated fields — dense grids covering the full area.

Scattered gravity stations

Ground measurements — sparse and irregular, but exact where they exist.

Reference-guided super-resolution + hard point constraints

Think pansharpening: a sharp correlated signal guides a blurry one — plus exact anchors at station locations.

Architecture: UNet + PointNet

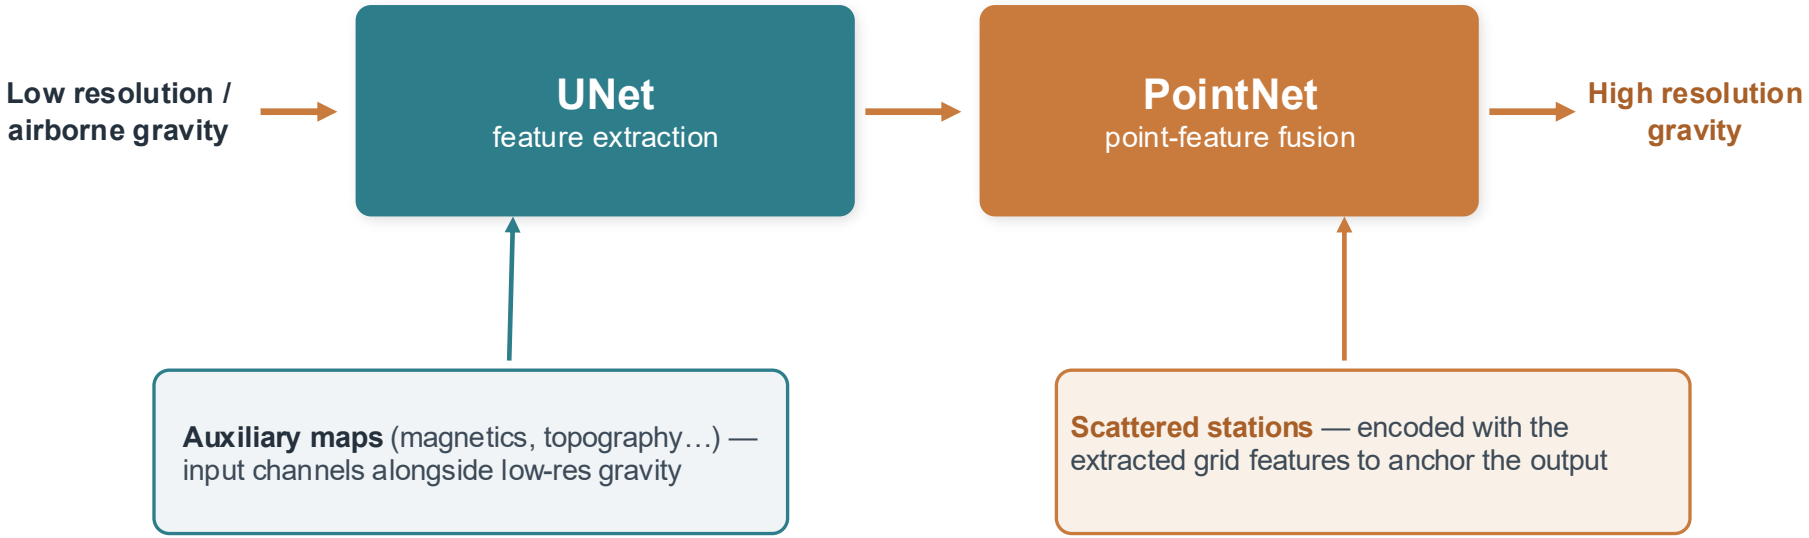

*Each data type gets the architecture suited to its geometry — **grids → convolutions, points → permutation-invariant set encoding.***

Results: Super-Resolved Gravity - 1st example

Earth Dynamics

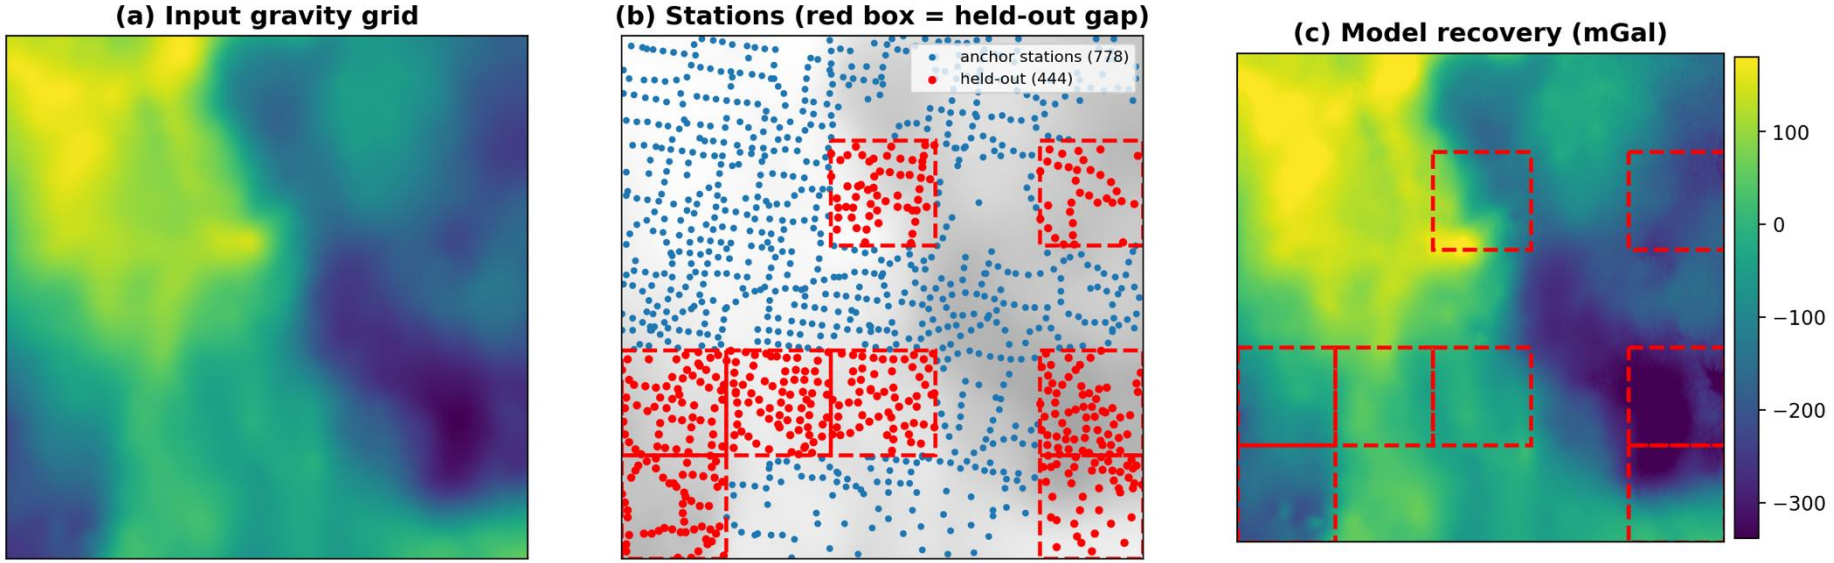

Model beats RBF by 62% MAE — held-out regions across val set (129 patches, 32,458 stations)

	MAE (mGal)	RMSE (mGal)	R^2
Our model	15.26	23.13	0.936
RBF interpolation	40.12	57.82	0.551

Results: Zoom in

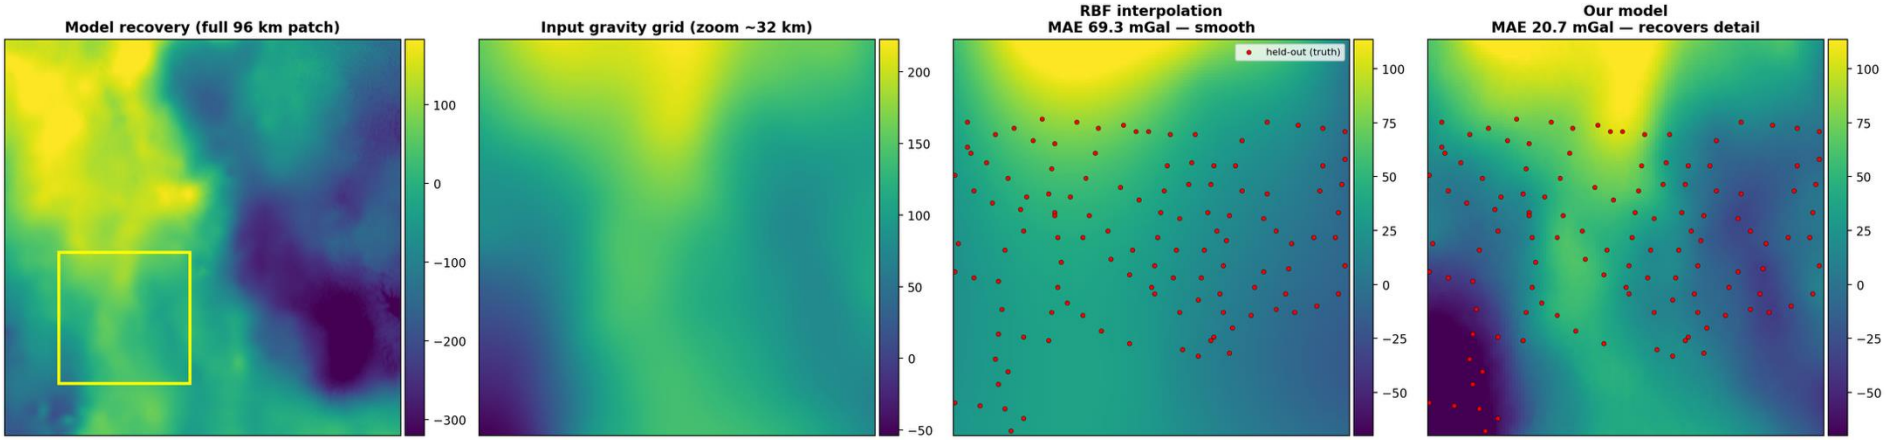

Results: Super-Resolved Gravity – 2nd example

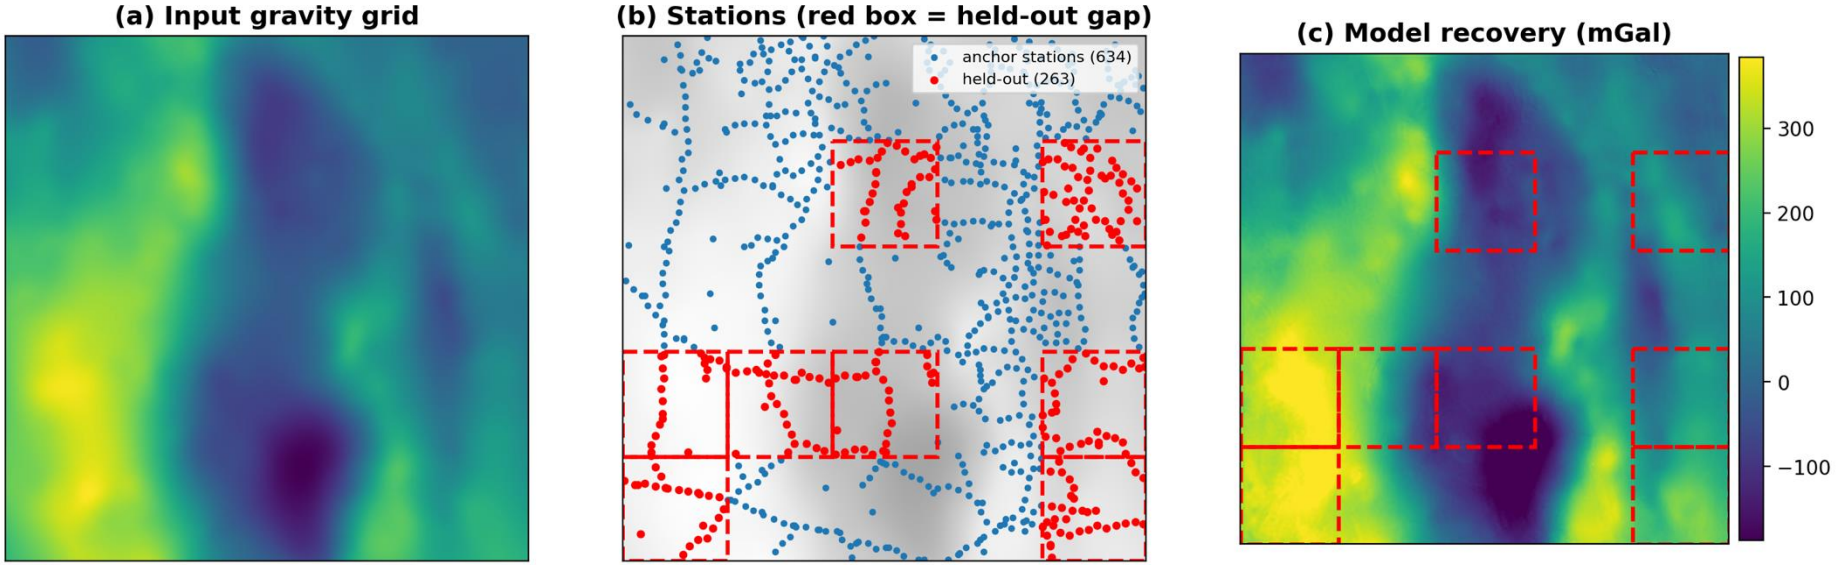

Model beats RBF by 62% MAE — held-out regions across val set (129 patches, 32,458 stations)

	MAE (mGal)	RMSE (mGal)	R^2
Our model	15.26	23.13	0.936
RBF interpolation	40.12	57.82	0.551

Results: Zoom in

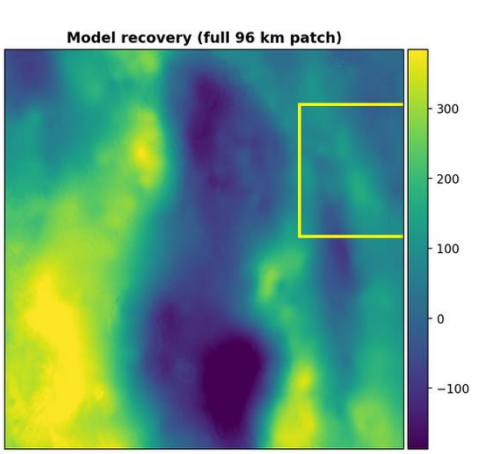

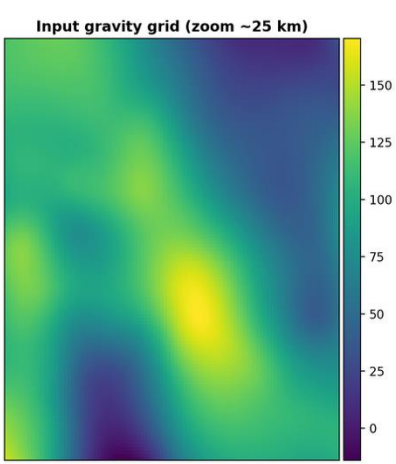

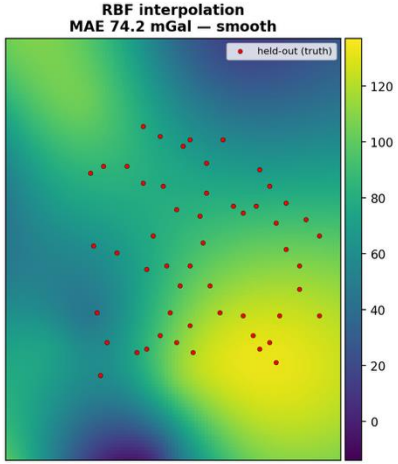

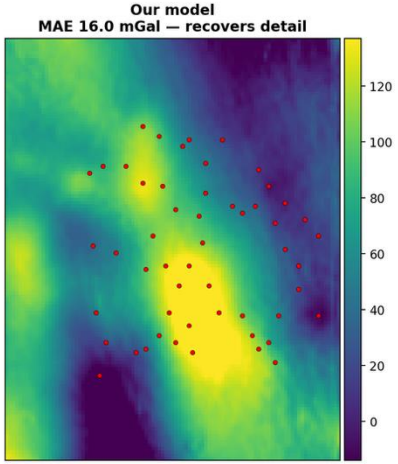

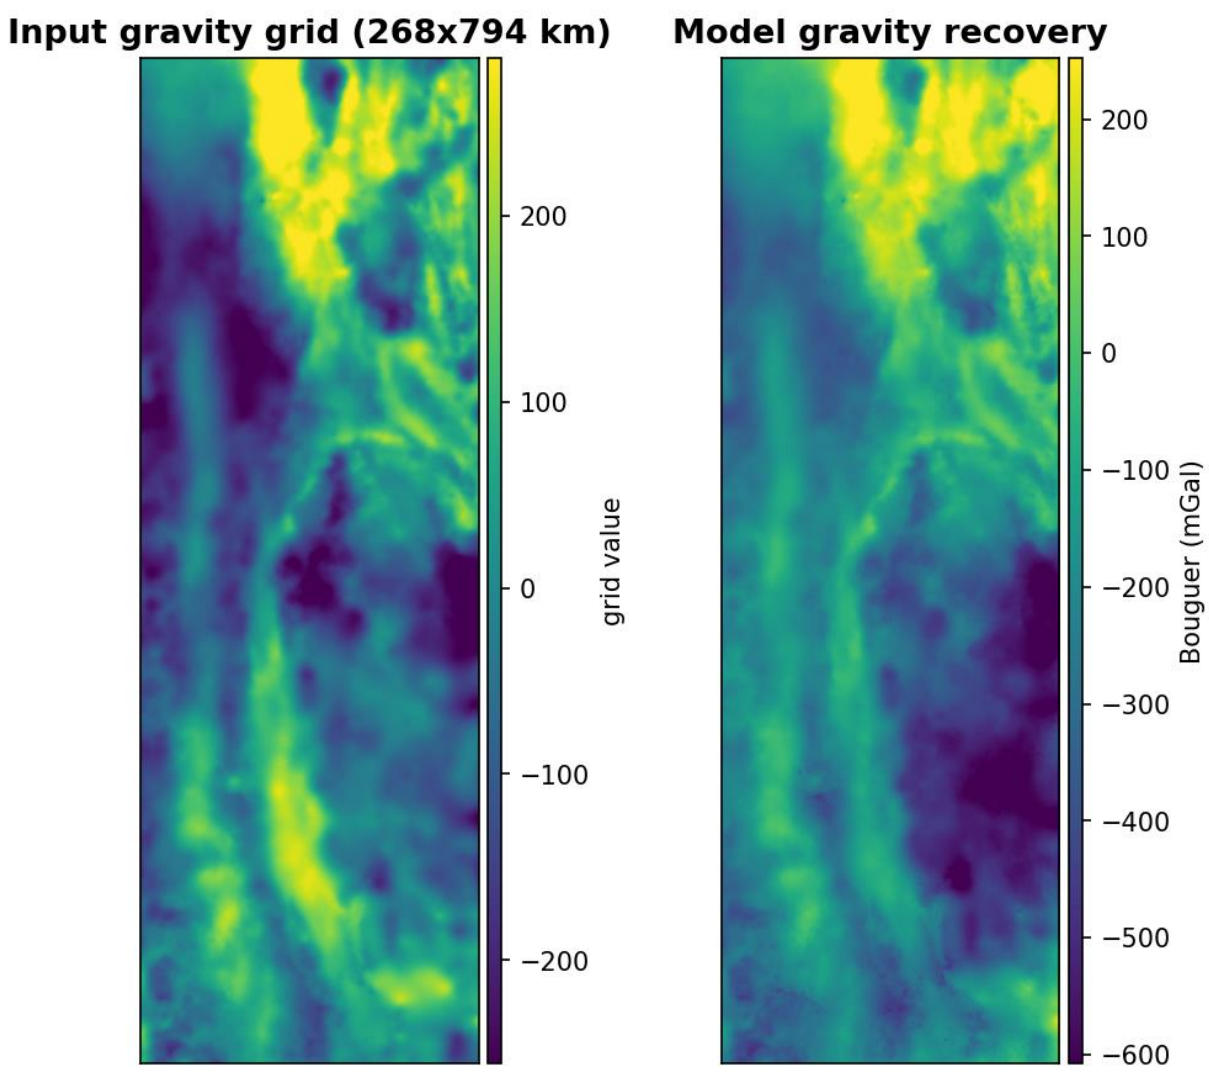
Input gravity grid (268x794 km)
Model gravity recovery
200
100
0
−100
−200
grid value
200
100
0
−100
−200
−300
−400
−500
−600
Bouguer (mGal)

Why It Matters for Exploration

Earth Dynamics

Holistic regional view

Gaps filled and maps extended → one continuous, high-resolution picture of the subsurface instead of fragmented patches.

Non-uniqueness, managed

GeoSim priors keep completions geologically realistic; generative sampling exposes the remaining uncertainty honestly.

Operational efficiency

Fewer redundant field surveys and drill holes; faster identification of high-value exploration targets.

From expensive, redundant surveying to AI-guided discovery.

Earth Dynamics

Subsurface understanding as conditional generation

in 3D from boreholes · in 2D from maps and points · all trained in simulation

GeoSim synthetic worlds

VAE + flow matching

Cascaded UNet + PointNet

Thank you · Questions?

Michal Holtzman Gazit · Earth Dynamics