

CASE STUDY

GOLD · NEVADA, USA

NOUS ranked a region of Nevada it had never been shown, and the gold was where the model said it would be.

PROJECT

Nevada is the best-documented gold province in the United States, which makes it the one place a targeting model can be graded honestly. Pluton ran NOUS across three areas of interest totalling roughly 42,000 square miles, about 40% of the state. Southern Walker Lane and the Northern Nevada Rift were used for training. Northwestern Nevada was set aside before any modelling began.

SUMMARY

NOUS assimilates terrain, spectral, structural and geochemical evidence into a single ranked prospectivity surface. Trained on the two known regions, it scored every 30-metre cell of the withheld region once, with no opportunity to tune against the result. In target districts, meaning those with epithermal gold present at surface, over 80% of known gold fell inside the top 5% of ranked ground. Separation on the withheld region reached 0.76 ROC-AUC against 0.80 on the training regions.

CHALLENGE

Almost every mineral targeting model is reported on ground it has already seen. If a model is fitted to a mapped district, it will find that district again. That exercise says nothing about where to drill next.

Splitting the data at random does not fix the problem. A mining district is drawn around known workings, so two records inside one share a geological system. If the records are split at random, the model sees part of a system in training and is scored on the remainder. The result is a number the ground will not honour.

The only honest test is to withhold real ground, model without it, and score once.

42,000 mi²

AREA MODELLED · THREE
AREAS OF INTEREST

80%+

PER-DISTRICT CAPTURE IN
THE TOP 5%

0.76

ROC-AUC ON THE WITHHELD
REGION

30 m

SCORING GRID RESOLUTION

The relevant baseline is the anomaly score an experienced geologist builds by hand over the classic pathfinder elements. On the withheld records the strongest single pathfinder, soil antimony, separates deposits at **0.604**. No single input does better. The model reaches **0.756** on those same records.

WITHHELD-REGION EVALUATION · NOUS NEVADA V1

Approach

NOUS reads ground as a stack of evidence. Each family of evidence enters as its own layer. The layers are an elevation model and its terrain derivatives, satellite band ratios carrying the alteration signature, fault architecture expressed as proximity and density, and assays kriged between sample points. A Bayesian hierarchical model combines the families into a single posterior surface, so every cell carries a mean and a standard deviation rather than a bare score.

Cross-validation folds were cut along mining-district boundaries rather than at random, so no fold could be trained on part of a geological system and scored on the rest. Isolated occurrences outside any district were each treated as their own fold group. Only the spatial-split result is reported.

Effects that absorb drainage, satellite acquisition and district structure are switched off when scoring new ground. The rock-type effect stays on, because a mapped geological unit is a property of the ground itself. Rankings are therefore directly comparable between cells, and absolute probabilities are quoted only against a stated base rate.

The top of the stack is the deliverable. It holds ranked drill targets, each with the evidence behind its rank, and a value-of-information score. That score states which next measurement, whether a drone survey, ground sampling or a hole, would remove the most uncertainty for the least spend.




Figure 1. The figure shows a 31 × 31 km window over the Ten Mile District, Humboldt County, at 30 m resolution. Reading from the bottom, the layers are terrain, the spectral alteration signature, fault architecture and kriged geochemistry, followed by every family combined into a posterior mean. Ranked drill targets sit at the top.

TEST DESIGN

TRAIN	Southern Walker Lane
TRAIN	Northern Nevada Rift
WITHHELD	Northwestern Nevada, scored once at the end

Ranked drill targets

Hotspots are listed in order and priced against a drilling budget, with the geological reasoning behind every entry.

Hotspot briefs

Each target has its own document covering the evidence behind its rank and how well it fits the deposit class.

VOI-scored next step

Each candidate follow-up is scored by how much uncertainty that measurement would remove. The next step is recommended rather than guessed.

Results

WITHIN A TARGET DISTRICT

In districts with epithermal gold present at surface, **over 80% of known gold falls in the top 5% of ranked ground**. It shows how sharply the ranking separates the best ground inside a district where the target deposit class is known to occur.

ACROSS THE WHOLE WITHHELD REGION

Capture is lower when it is measured over the entire withheld area, including ground that carries no epithermal system at surface. The comparison that matters here is against chance.

SLICE OF RANKED AREA	BAYESIAN MODEL	TREE MODEL	RANDOM	CONCENTRATION
Top 1%	n/a	n/a	1%	Bayesian model leads
Top 5%	18%	16%	5%	3.6× chance
Top 10%	29%	27%	10%	2.9× chance
Top 20%	47%	47%	20%	2.4× chance

These figures were measured on Northwestern Nevada, which was withheld before modelling and scored once. The margin between the cross-validated and withheld scores is 0.047 for the Bayesian model and 0.065 for the tree model. The tree model scores marginally higher on whole-ranking ROC-AUC, but the Bayesian model leads at the top 1%, 5% and 10%, which are the slices a drilling budget actually consumes. The full method and test design are set out in the *Technical Overview, NOUS Nevada v1*, which is available on request.

Conclusion

The ranking survives new ground. The withheld region was scored once, at the end, with no opportunity to tune against it. Performance held to within 0.04 ROC-AUC of the training regions, which is the margin a model shows when it has learned geology rather than memorised a map.

The output is ranked ground, not a probability. Every cell carries a posterior mean and a standard deviation. The deliverable is an ordered list of ground worth walking, with the uncertainty attached to each entry, and a stated next measurement for the targets where uncertainty is worth buying down.

The model has a stated envelope. The model is calibrated to epithermal gold at surface. Districts whose surface signature is porphyry molybdenum-copper, or is dominated by antimony rather than gold, sit outside that envelope and are not captured. The reach of the current design is reported, not implied.

Bring us ground you already know.

The fastest way to judge a targeting model is to give it a question you can already mark. Send us a licence where you know what was found, hold the answer back, and we will rank it blind and hand you the result against your own record. Bring whatever data you already hold, such as magnetics, geochemistry or mapping, and it will go into the stack.

PLUTON.EARTH