

Find the Ore in the Data: Pricing AI-Assisted Mineral Discovery

A Decision Framework for Geological Data, Prospect Rankings and Staged Exploration Capital

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

Artificial intelligence can help assemble geological evidence and rank exploration targets. Its commercial value depends on whether the resulting decision is better than a credible conventional alternative. A prospectivity score, an anomaly and a trained model do not establish an orebody or a reportable mineral resource. This paper develops a transaction framework for deciding what to pay for data, a model, an exploration right and a verified technical result. It separates the rights to use source data from the performance of the model and the rights to explore a ranked target. The proposed valuation unit is the incremental, risk-adjusted value of a specified field programme relative to the same programme selected without the model.

The paper draws on public geological-data initiatives, research on exploration bias, the US Geological Survey's artificial-intelligence assessment work, and mineral-disclosure standards. It proposes a data-rights ledger, spatially independent validation, a calibration test, an evidence ladder and staged payment terms. These are author-developed decision tools; no published model or exploration company is represented as having achieved the paper's scenario outcomes. A fictional eight-target programme illustrates the arithmetic. With expressly author-assumed probabilities of a qualifying verification milestone of 10% under the conventional selection and 16% under model-assisted selection, a hypothetical USD 18 million conditional value per success, and USD 3 million of incremental data and model cost, the model-assisted choice adds USD 5.64 million of expected value. That result disappears if the probability advantage is not reproducible or the cost exceeds the corresponding value of information.

The paper closes with an independent-validation protocol, contract terms, sensitivities, diligence questions and a staged implementation route. Its central recommendation is to buy a testable improvement in exploration decisions, pay for the rights actually received, and defer resource or mine valuation until competent geological and economic evidence exists. The numerical example is educational and does not state observed discovery rates, transaction prices or expected returns for any real prospect.

JEL Classification: C53, G31, G32, L72, Q32

Keywords: mineral exploration, artificial intelligence, prospectivity, geological data, value of information, target ranking, exploration finance

1 Define the purchase decision

An investor considering an AI-assisted mineral-discovery business can be offered several different things. One seller may own historical survey files. Another may have cleaned and georeferenced those files. A software company may own a trained model that ranks ground for field investigation. A licence holder may control the right to test a ranked target. An explorer may have already drilled and reported technical results. Those assets do not confer the same claim on future cash flow. The first investment question is therefore what legal and economic interest the investor would actually acquire, at what price, and before which next experiment.

The US Geological Survey's Earth Mapping Resources Initiative collects geologic, geophysical, geochemical and other data to improve understanding of mineral systems [1]. A US, Canadian and Australian compilation offers national-scale earth-science layers for critical-mineral analysis and machine-learning work [2]. Publicly available inputs create opportunities for analysis. They also require the buyer to identify which part of a seller's product is proprietary and defensible. A map built from open data can be useful without giving its creator exclusive ownership of the mapped ground. The payment should reflect the specific service, intellectual property, access right or exploration title conveyed.

This paper asks whether a model changes the choice of targets within a fixed exploration budget. Its benchmark is a documented conventional process using available geological expertise and data. The model has value when the change in target selection raises the expected value of a defined verification milestone enough to cover incremental costs and risks. A higher score or more vivid map is not itself a monetisable discovery. The later steps of resource definition, metallurgy, permits, infrastructure and market access remain separate decisions. An investment committee should approve the next cash commitment against the evidence appropriate to that stage.

2 Specify the geological question before fitting a model

Mineral prospectivity is conditional on a deposit concept. A model trained to identify one kind of mineral system should not automatically be used for a different host rock, depth, climate, survey method or commodity. The geologist should state the mineral system, its diagnostic features, plausible spatial scale, depth range and rival explanations for a positive signal. Those assumptions determine which geophysical, geochemical, structural and remote-sensing layers are relevant. They also determine what a negative field test can teach. A generic probability of finding “ore” has little decision value without a defined event and an observation protocol.

The target event might be a verified mineralised intercept meeting a technical threshold, a deposit-scale continuity result, a resource estimate under an applicable reporting framework, or an agreed sale of exploration rights. These are different endpoints. The farther the endpoint is from the model's input data, the more other hazards stand between ranking and value. The USGS's three-part quantitative assessment approach distinguishes permissive tracts, potential numbers of deposits and distributions of deposit size or content [3]. It represents uncertainty explicitly. That methodology is a reminder to define the geological event before converting a prospectivity map into a financial probability.

For each target, the investment file should specify the licence polygon, mineral-system hypothesis, sampling plan, permitted access, decision threshold, programme cost and the result that would cause further funding. If the target lies outside an enforceable right, a favourable model output may have no direct transactional value to the current holder. If the target is inaccessible or protected, field verification may be impossible or delayed. A buyer should therefore value the actionable intersection of scientific signal and legal, physical and financial ability to test it. This narrows the object being priced to something an investor can actually execute.

3 Separate a score from a probability

A classifier can rank targets without producing calibrated probabilities. A score of 0.82 may mean only that one cell ranks above another under a particular algorithm. It does not automatically mean an 82% chance of a discovery. Calibration requires comparing predicted probabilities with observed frequencies for the same outcome definition and a validation set representative of the proposed use. The sample must also account for the fact that most untested ground has an unknown label. Calling every unexplored cell a negative example can distort both training and performance claims.

Historical deposits are also an imperfect training set. A USGS-published study describes how discovered resources reflect accessibility and previous exploration criteria as well as geology [4]. Models can therefore learn where people drilled, roads existed or old surveys were dense, instead of learning a transferable mineral-system relationship. The investment committee should see the spatial distribution of known deposits, barren tests and untested ground. It should ask whether the model's high-ranking areas coincide with established camps or survey coverage. That pattern can be commercially useful, yet it may not demonstrate discovery in a new geological province.

4 Inventory the actual data asset

Data value begins with provenance and permitted use. A geological-data room should catalogue survey type, acquisition date, spatial reference system, resolution, sensor calibration, sampling method, laboratory quality control, missing values, processing history, owner, licence terms and restrictions on transfer or model training. Geological maps, airborne geophysics, geochemistry, drilling logs, assays and field notes have different error structures. The useful unit is a reproducible layer with documented limitations, not a raw terabyte count. An old map can be valuable when it covers a decision-critical gap, while a much larger duplicate dataset can add little.

The USGS and partner agencies have released broad datasets intended to support critical-mineral assessment [1, 2]. Geoscience Australia describes precompetitive geoscience as a way to address exploration challenges and identify new mineral provinces [5]. These programmes illustrate a public-good baseline. A commercial provider can add value through quality control, integration, interpretation, local acquisition or exclusive contractual access. The buyer should avoid paying a proprietary-data premium for public layers merely repackaged without a durable transformation. It should also avoid assuming that public availability means unrestricted commercial use of every derived dataset or third-party layer.

The diligence file should identify each source's licence and any obligations to attribute, share derivatives, protect sensitive locations or obtain government consent. It should show whether the model can be reproduced if a

vendor contract ends. A data escrow can include raw files, metadata, preprocessing code, feature definitions, model version and evaluation results. Rights to reuse a trained model across other licences can differ from rights to disclose a target map or transfer a geological database in an acquisition. These differences affect the amount of value retained by the buyer after the transaction. Specialist counsel should review the rights under the relevant jurisdictions and contracts.

Table 1. What the buyer receives at each evidence stage

Asset or result	Minimum diligence	Principal value question	Major unresolved risk
Source data	Provenance, coverage, quality and licence	Does it change an actionable decision?	Missing or restricted rights
Trained model	Reproducible code and independent validation	Does it outperform a defined baseline?	Leakage or transfer failure
Exploration right	Title, access, work commitments and term	Can the target be tested and retained?	Licence or access failure
Verified technical result	Samples, QA, competent review	Does new evidence justify the next gate?	Geological discontinuity

Author-developed diligence framework; a listed value driver does not establish a market price.

5 Test how the data were transformed

Machine-learning performance can depend more on preprocessing than on the algorithm's brand. Old maps may need georeferencing. Geophysical grids may be resampled. Assays can have detection limits, unequal sampling density and inconsistent units. A spatial join can associate a known deposit with a nearby feature that was measured after the deposit was found. Such leakage can make retrospective prediction appear impressive. The buyer needs a processing lineage from raw file to model feature, with dates and code or a sufficiently detailed reproducible procedure.

The USGS's 2025 account of an AI competition for historical geological maps reported measurable but uneven performance across georeferencing and feature extraction tasks [6]. It concerned map-processing functions, not a demonstrated economic discovery rate. That distinction matters commercially. Better digitisation can reduce analyst time or improve the input catalogue, yet a seller should not convert an extraction benchmark into a claim that drill success has improved by the same percentage. The budget should separately price map digitisation, geological interpretation, model development, validation and field testing. Each has a different acceptance criterion.

The data audit should include an error sample that a buyer can inspect without access to the vendor's preferred examples. Select pages, survey lines or sample records at random and compare them with raw sources. Check coordinate systems, transformations, missing-data imputation and assay units. Compare the model's ranking with a simple expert screen and an elementary statistical baseline using the same information. If the advanced model wins only when the benchmark is deliberately weak, its incremental commercial value is overstated. A reproducible but modest improvement can still merit payment when exploration decisions are expensive and the underlying rights are actionable.


6 Validate across space, time and geology

Randomly dividing neighbouring map pixels into training and test sets can allow highly similar locations into both groups. The resulting accuracy may measure interpolation near known sites rather than prediction in a new basin. A more informative challenge holds out entire geological districts or licence blocks, keeps later observations out of earlier training sets, and reports results by mineral-system type. The exact split should match the investor's intended use. A model to rank targets within one well-studied camp needs a different test from a model marketed as transferable across countries.

The validation protocol should be registered before the vendor sees its outcomes. Record the model version, target list, ranking date, data cutoff and statistical measures. Keep a sealed set of ground observations or newly acquired data for a prospective test. A reviewer independent of the vendor should reproduce the ranking from the materials actually deliverable under the contract. The USGS CriticalMAAS programme covers data extraction, mineral-potential mapping and human-in-the-loop learning [7]. Its training-and-validation data release documents a specific map-analysis benchmark [11]. These public efforts do not validate a particular commercial vendor or promised return.

The committee should compare a full portfolio, not only the best success story. Report how many targets were ranked, selected, tested, positive, negative and still unresolved. State the cost and decision date for each. A target that cannot be drilled should not silently disappear from the denominator. Repeatability across another licence or later campaign matters because an acquisition or long-term software contract prices future use. Weak transferability can justify a lower upfront fee and a fresh gate before geographic expansion. The contract should define whether failure on that gate changes price, exclusivity or termination rights.

Figure 1. Author-developed evidence path from data to capital commitment



Only independently testable increments support a performance premium

A ranked target remains an exploration proposition until field evidence changes its status

Arrows indicate decisions and verification work, not observed discovery probabilities.

7 Build a credible conventional benchmark

The counterfactual should be an exploration team using the same available maps, surveys and licence information without the contested model. It should have the same budget, target definition and time window. Otherwise a model can seem valuable simply because its comparison group spent less, used poorer data or screened fewer sites. The committee should ask the conventional team to rank targets and explain its geological rationale before seeing the AI list. An independent adjudicator can then compare the selected sets, the overlap, and the expected costs of testing each. Overlap is important: the model cannot claim full credit for a target that competent geologists would have selected anyway.

The benchmark need not be a straw-man manual process. An incumbent explorer may already use statistical prospectivity maps, GIS, remote sensing and expert systems. The relevant economic test concerns the incremental improvement from the vendor's complete offer over the best accessible alternative at the decision date. Public survey data can increase the quality of both approaches. The committee should include the cost of maintaining the baseline; it should also include the new model's costs of integration, cloud use, specialist staff, retraining and data acquisition. A low annual software fee can conceal a much larger dependent field programme or data-preparation expense.

The benchmark should be written into a pilot contract. Define selected-target count, field-work standard, success milestone, evaluation period and the treatment of untestable sites. Give the vendor credit for faster decisions only if the time saving is measured and economically consequential. If a licence expires before a conventional analysis would finish, speed can have value; the expiry date and avoided loss must be evidenced. A vendor should also receive credit for a newly identified target outside the conventional shortlist where the buyer has a valid

exploration right and a feasible test plan. This keeps value attribution tied to decisions rather than marketing claims.

8 Translate a ranking into an investment choice

The investment unit is a portfolio of potential tests. Suppose an investor can fund eight drill or sampling targets this season. The model changes which eight are selected, and perhaps the order in which they are tested. Expected value depends on the chance of a precisely defined verification milestone, the value conditional on that milestone, the test cost, and the cash required for later stages. The probability must correspond to the programme being valued. A regional prediction that some deposit may exist cannot be applied as the chance that a particular licence will reach a saleable technical result within a specified term.

An author-developed decision formula is: incremental programme value equals the expected conditional value of model-selected targets minus the expected conditional value of baseline-selected targets, less additional data, software, validation and execution costs. If the two selections contain the same target, the corresponding value appears in both programmes and cancels. A fully developed valuation should also model correlations between targets, rights constraints, follow-on funding and the ability to abandon after a failed test. The simple case later in this paper holds those complications apart to make the incremental logic auditable. It does not estimate mine NPV from a model score.

The investment committee should present probability ranges rather than a single precise number where field observations are scarce. It should disclose the sample used for calibration and separate independent holdout results from retrospective training results. It should state the conditional-value basis: a signed purchase offer, a comparable licence transaction adjusted for stage, or a hypothetical teaching assumption. Those bases have different evidence strength. The model's sensitivity to commodity price and later project costs should be considered only when the chosen milestone genuinely depends on them. Otherwise a high price scenario can become a way to inflate an early-stage algorithm premium.

9 Do not capitalise an untested mineral resource

Public mineral-disclosure systems distinguish exploration information from resources and reserves. The CRIRSCO International Reporting Template describes the separate categories and the role of a Competent Person [8]. The US Securities and Exchange Commission's mining-property rules define the technical disclosure framework for covered registrants [9]. Canadian NI 43-101 requires qualified-person involvement and sets disclosure and technical-report obligations for issuers within its scope [10]. The applicable code depends on the issuer and jurisdiction. An AI-generated colour map or inferred anomaly is not a substitute for sampling, resource estimation or the modifying factors needed for a reserve.

A private investor can assess an exploration prospect without issuing a public resource statement. It still needs honest terminology. The investment memorandum should distinguish known assay results, historical records of uncertain quality, author-developed models, exploration targets and reportable resources. A prospective buyer should ask which independent technical professional reviewed the method and whether original data can be inspected. The company should not present predicted tonnes or contained metal as existing inventory merely because a machine-learning model draws a boundary around a favourable area. A financing term sheet tied to a future resource milestone should define the governing technical standard and who determines whether it has been met.

The value of a model can be meaningful before a resource exists because better information can prevent wasteful work or direct a fixed budget toward stronger hypotheses. That value is an option on learning, subject to proof. The subsequent value of a deposit belongs to rights, geology, development economics and market conditions. A vendor may bargain for a royalty, equity interest or success fee, but that is a commercial allocation of uncertain future value. It is not geological evidence. The buyer should price such claims with the same discipline as any other contingent transaction claim and include their effect on later financing.

10 Map ownership, access and transfer

An exploration company can lack the right to use a dataset outside the original licence or to transfer it with shares in an acquisition. A software licence can permit internal scoring while prohibiting disclosure to an investor, independent technical adviser or buyer. A government survey can be open while a proprietary drill database remains restricted. The buyer should construct a data-rights ledger that maps every layer and model output to owner, licence, term, geography, permitted use, sublicensing, confidentiality and survival after termination. The ledger should be reviewed alongside the mineral-title register.

The commercial package may be a subscription, an asset purchase, a joint venture, a funded farm-in or an option over ranked ground. Each package grants a different claim. Subscription value depends on continued access and service quality. An asset purchase depends on deliverability of data, code and rights. A farm-in depends on the cost and work needed to earn an interest in the ground. A model provider may seek rights over every target it identifies, which can encumber future exploration financing. The buyer should distinguish payment for a repeatable software capability from payment for a particular exclusive prospect.

Transfer tests should be concrete. Can the buyer reproduce a model output with delivered data and code? Can the buyer grant an independent reviewer controlled access? Can an acquirer receive the same rights? What happens if the vendor is sold, insolvent or stops supporting the product? Can the buyer export features and results in documented formats? The answers determine retention value. A contract promising “ownership of insights” may leave critical raw-data and model rights unresolved. Counsel, geologists and information-security specialists should jointly review the schedule before the investor pays an exclusivity premium.

11 Price the information in a bounded field programme

The simplest defensible price for a model is bounded by the value of the decision it improves. A useful exploration programme states its maximum test budget, expected timing, field standard and exit rule before a model is evaluated. A vendor can then demonstrate how its rankings change the selected set or sequence. The upper bound for a rational buyer is the incremental expected benefit after the model's acquisition, validation and operation costs, adjusted for uncertainty and bargaining position. The upper bound is not the total expected value of all model-selected ground, because the buyer may have found much of it through the baseline process or already owned the underlying right.

Value of information depends on the decision threshold. A better ranking can be worth little when the investor can cheaply test every target. It can be valuable when field work is costly and only a small subset can be tested before a licence, season or budget expires. Conversely, a model that confidently identifies ground the buyer cannot access has limited immediate value. The committee should report how the target ordering changes the use of a scarce constraint: drill metres, assay capacity, land access, crew time, legal title or capital. This specifies why an information product might earn an investment premium.

The buyer should also price the ability to learn from negative work. A properly recorded barren hole can revise a mineral-system hypothesis and improve later target ranking. That benefit is real only if results and model updates are retained under the contract. A royalty or success fee based on a later discovery can make a low upfront software price expensive after the model has been replaced or the target materially reworked by geologists. A staged price should therefore allocate payment among delivered data, demonstrated ranking performance and retained rights. It should specify whether the vendor shares the cost of false positives or only participates in upside.

12 Use a fictional case with explicit assumptions

Consider a fictional licence portfolio containing 40 candidate targets. The investor can test eight in one field season. A conventional geoscience team and a model-assisted team select eight under the same USD 1.5 million cost per target. The author assumes, solely for teaching, that each conventionally selected target has a 10% probability of meeting a defined independent verification milestone, while each model-selected target has a 16% probability. The author also assumes a USD 18 million gross value to the licence holder for each successful milestone, after all further work needed to make that milestone saleable has been allowed for in the assumed amount. These figures are not estimates drawn from field data or comparable transactions.

The conventional programme spends USD 12 million on eight tests. Its illustrative expected milestone value is eight multiplied by 10% multiplied by USD 18 million, or USD 14.4 million; expected value after test cost is

USD 2.4 million. The model-assisted programme has illustrative expected milestone value of eight multiplied by 16% multiplied by USD 18 million, or USD 23.04 million. After the same USD 12 million field cost and USD 3 million of additional data, model and independent-validation cost, its expected value is USD 8.04 million. The difference is USD 5.64 million. No mine or resource value is implied; the fictional outcome is an earlier verification milestone with a conditional transaction value.

This simple comparison assumes independent target outcomes, identical conditional milestone value, full ability to exploit a result, and no difference in timing or land access. It assumes the two methods' probabilities have been validated for their selected sets. If several targets share one geological system, the probability of multiple successes can be correlated and total value may be lower. If the licence term is short or the investor lacks rights to model-selected targets, the value can fall sharply. The calculation is useful because it exposes these assumptions. It is not a valuation of any named company, model or mineral property.

Table 2. Fictional eight-target programme under author assumptions

Measure	Conventional selection	Model-assisted selection
Candidate targets in portfolio	40	40
Targets tested	8	8
Assumed milestone probability per test	10%	16%
Assumed conditional value per milestone	18	18
Expected milestone value	14.40	23.04
Field cost	12.00	12.00
Additional data, model and validation cost	0.00	3.00
Expected value after stated costs	2.40	8.04


USD millions except probabilities and target counts; arithmetic is illustrative, with no empirical discovery-rate claim.

13 Stress the advantage before offering a price

In the fictional case, incremental value equals eight times the difference between the model-assisted and conventional probabilities, multiplied by USD 18 million, less the incremental model cost. At a USD 3 million cost and a conventional probability of 10%, the model-assisted programme needs an assumed probability above approximately 12.08% merely to produce a positive incremental expected value. This is a break-even threshold under the case's assumptions, not a performance claim. If independent validation supports only a 12% model probability, the incremental value is negative USD 0.12 million. At 16%, it is positive USD 5.64 million.

The committee should stress the conditional milestone value as well as the probability. The USD 18 million teaching input might be too high if later sampling, title defects or commodity-price conditions reduce what a buyer would pay. It might also be too low for a scarce, well-located and well-documented right. Without an observed sale or signed offer, the committee should label it as an assumption. It should show how much of the model's alleged advantage survives when the field programme costs more, validation fails, the result takes longer or an additional funding round dilutes the investor. A vendor performance metric belongs in the model only after its economic meaning has been established.

Sensitivity analysis should also consider the baseline probability. A strong conventional team may already choose most of the successful targets. If the vendor's chosen eight overlap six of the conventional eight, only two substitutions can produce incremental value; the total value of all eight model choices should not be assigned to the vendor. If the model's edge depends on proprietary data acquired after the evaluation date, the buyer must include the ongoing cost of that data. If target definitions change after results arrive, the backtest is invalid for the original claim. Price negotiation should use the lower end of a defensible value range until these matters have been tested.

Figure 2. Fictional incremental value across model probability and cost assumptions

All probabilities, values and costs are author assumptions; values are USD millions and do not describe observed exploration outcomes.

14 Compare a licence premium with a software fee

An investor might pay a model vendor a subscription and separately buy exploration rights. Another deal might grant the vendor equity or a discovery-contingent payment instead of a large fee. These structures should be compared on an equivalent basis: cash paid now, work obligations, dilution, future royalties, control of data and rights after termination. A small subscription can be attractive when transferability is uncertain. A significant upfront acquisition price requires stronger evidence that the buyer receives durable code, data rights, customer relationships or validated performance. A royalty can conserve early cash but reduce the value of a successful asset to later investors.

For a licence premium, the starting point is the ground interest and its obligations. A ranked target within the licence may justify more expenditure than a low-ranked one, but its value depends on title, exploration term, work requirements and competition for the right. The model vendor should not automatically capture that entire premium if the licence holder already owned the ground and the same target appeared in existing geological work. A transaction model should separate payments to the licence seller, the data owner, the software provider and the exploration operator. It should also show whether an offtaker, strategic investor or local partner acquires rights that restrict a later sale.

Milestones can reduce premature valuation. An initial fee may cover delivery and independent reproduction. A second payment can follow a specified out-of-sample ranking test. A later payment can follow a technically verified field result under the governing disclosure and contract definitions. The parties must negotiate negative and inconclusive outcomes as well as success: who owns the data, who funds another test, and when exclusivity ends. Such a structure does not guarantee fairness or geological success. It makes consideration better aligned with evidence and supplies a clearer basis for future financing discussions.

15 Examine market benchmarks without false equivalence

Mining data assets, exploration software, licence interests and drill-stage discoveries are often described together in sale materials. Their transaction prices are not automatically comparable. A quoted enterprise value for a listed explorer can reflect cash, several licences, a management team, a resource estimate, commodity exposure and market sentiment. A software valuation can reflect recurring revenue and broad customer use. A licence farm-in can price a programme of spending, control rights and earn-in milestones rather than a simple cash purchase. A benchmark is useful only after the transaction perimeter and stage have been adjusted.

The investor can build a market evidence book from published transaction documents, signed terms it may lawfully use, and technical-report dates. Each record should show stage, commodity, geography, rights area, work commitments, royalty or stream burdens, ownership earned, and the amount actually paid or committed. Missing consideration should remain missing. A reported “up to” earn-in amount should not be treated as cash

paid at signing. The committee can use the book to bound the assumed conditional value in a scenario, while stating where transactions differ. This approach is especially important when a model provider cites a spectacular deposit discovery as a direct price precedent for its algorithm.

16 Allocate geological and model risk between parties

A model vendor controls some risks: data processing, version control, reproducibility, documentation and whether reported tests were performed as agreed. The explorer controls other risks: access, permitting, sampling quality, drilling execution and reporting. Geological uncertainty remains even when both perform well. A contract should place each obligation with the party able to manage it. A vendor warranty can cover ownership of deliverables and non-manipulation of validation records. It cannot rationally warrant that a deposit exists merely because a model selected a site. An operator warranty can cover following an agreed field protocol, while a negative result remains a possible outcome.

The contract should specify the validation population, excluded targets, data cutoff, permitted retraining, model version and metric. It should specify who adjudicates whether a target is independently distinct, whether a result meets the milestone, and how a delayed or inconclusive test is treated. Escrow or reproducible delivery can preserve the buyer's ability to use paid work if the relationship ends. Security provisions should protect sensitive location and assay data. A party granting exclusive rights over identified targets should also specify geographic and temporal bounds, so stale rankings do not indefinitely block the licence holder's other work.

Risk allocation should be reflected in the capital structure. A buyer facing unverified transferability can offer a small initial payment and a funded validation campaign. A vendor confident in an independent edge can accept contingent consideration tied to that edge. An explorer with strong rights and a capable technical team can retain more upside while bringing in a funder for a defined programme. The allocation is a negotiating choice, not a formula dictated by this paper. The investment memorandum should show the same outcome tree for all offers so the parties can see the cost of rights, data, field work and contingent claims separately.

17 Include field execution and adverse evidence

The model's commercial test ends in physical work. A high-ranking prospect can be rejected after surface inspection, sampling, geophysical interpretation or drilling. The field programme should use documented sampling, chain of custody, laboratory QA and competent interpretation. The programme should capture negative results with the same care as positive results. This protects later model evaluation and reduces selective reporting. An untested target should be marked untested rather than counted as a model failure or success. A target blocked by access should be reported separately because it affects commercial usefulness even though it does not resolve geology.

The USGS's AI-supported assessment work distinguishes map processing, mineral-potential mapping and human-in-the-loop interpretation [7]. A practical investment process can follow the same separation. Geologists should examine the model's explanation and alternatives, then design field tests that discriminate between plausible geological causes. A magnetic anomaly, for example, may admit several interpretations. The work programme should state which observation would change the investment decision. A drilling campaign that merely confirms the model saw an anomaly has not necessarily advanced a mineral resource or conditional asset value.

Field outcomes should update the probability model and the capital plan. The update should use the result as it was observed, including assay limits and spatial uncertainty, rather than a post-hoc success label chosen to improve marketing. Costs of remobilisation, permits, rehabilitation and further testing should be recorded against the originally approved budget. If results reveal a different mineral system, the model may need redesign rather than minor recalibration. A committee should authorize that new hypothesis as a separate work stage. Disciplined adverse-evidence handling makes future financing and buyer diligence more credible.

18 Test concentration, transfer and portfolio effects

Eight high-ranked targets in one structural corridor may all depend on the same geological interpretation. An investor should map the geological and data correlation among targets. If a shared survey artifact or false deposit model drives the rankings, one negative observation can reduce the value of the whole group. Diversifying across mineral systems can reduce common-model exposure, although it may also reduce operational efficiency and require different expertise. The decision model should therefore report both target-level and portfolio-level uncertainty. Independence in the fictional calculation is a teaching simplification, not a characteristic to assume in a real field plan.

Geographic transfer can be tested by withholding an entire region and applying the frozen model there. Commodity transfer requires an even stronger geological rationale. The mineral-system features that matter for one deposit type may not generalise to another. The national-scale compilation cited earlier spans the United States, Canada and Australia [2], but data availability across countries does not itself establish a transferable predictive relationship. The validation record should show which features are measured comparably and which were imputed or excluded. A vendor seeking a global licence price should demonstrate performance for the jurisdictions and mineral systems included in that price.

The buyer should also examine operational concentration. If one specialist controls a model that cannot be reproduced or maintained, the apparent IP asset may have significant key-person risk. If a critical dataset can be withdrawn, the valuation should account for replacement cost and possible downtime. If nearly all model customers are exploration companies funded by short capital cycles, renewal revenue can be volatile. Those are business risks distinct from geological success. The transaction structure can separate software-business value from prospective mineral-right value and price each with its own evidence and downside case.

Table 3. Investment committee tests for an AI-assisted exploration transaction

Gate	Evidence to request	Decision if unresolved
Data and rights	Provenance, licences, transfer and access	Defer exclusive-data premium
Model validation	Frozen version, spatial holdout, baseline and calibration	Limit payment to bounded pilot
Actionable targets	Mineral title, access, work budget and field plan	Exclude inaccessible targets from value
Technical result	Raw observations, QA and competent review	Withhold resource-stage valuation
Further capital	Rights to stop, follow-on budget and financing route	Cap commitments at current gate

Author-developed controls; each test requires project-specific evidence.

19 Prepare an acquisition or financing diligence file

The technical room should include raw data, metadata, source licences, preprocessing logs, model code or a reproducible executable, training and validation sets, frozen rankings, field results and adverse cases. The geological room should include title documents, maps, work obligations, permits, sampling protocols, assays and any reports prepared under the applicable disclosure regime. The commercial room should include customer contracts, IP assignments, royalties, farm-ins, software subscriptions, contingent consideration and the fully diluted capitalization table. A simple cross-index should link every value-driving claim to its evidence and owner.

An independent reviewer should be able to reproduce the claimed target ranking and identify what would have changed under the conventional baseline. The reviewer should document data leakage, survey-coverage bias, changing label definitions and any decisions made after validation results were known. Counsel should test ownership and transfer. A qualified technical professional should review exploration claims. Financial advisers can then reconcile the rights and evidence into a stage-specific valuation. These roles complement one another; a financial model cannot validate geology, and a geological report cannot by itself establish IP rights or a transaction price.

The diligence report should show disputed and unverified items plainly. A public deposit database may contain historical results with unknown QA. A vendor's commercial slide can report a backtest without source code. A


licence register can be current while access negotiations remain open. Each item should have a status, responsible reviewer, required evidence, deadline and effect on payment or work authorisation. This avoids the common failure of converting an unanswered technical question into an optimistic spreadsheet input. It also gives a prospective lender or buyer a record of how the exploration thesis evolved as new work was completed.

20 Implement a staged decision and payment plan

The first gate is a paper review of rights and data provenance. Approve only the expenditure needed to establish whether the target universe is legally testable and technically coherent. The second gate is a reproducible model challenge against a pre-registered expert baseline with a spatial or temporal holdout matched to the proposed use. The third is a bounded field pilot that records all selected targets and negative outcomes. The fourth is a technical and commercial review of any verified result. A resource-definition or development decision follows only when the relevant evidence exists. Each gate should have an approved budget, named decision maker and stop condition.

The payment schedule can mirror those gates. The vendor receives a delivery payment for an auditable data and model package. A performance payment can follow a specified independent test. A field-result payment can follow a contractually defined milestone, with exclusions for results materially produced by different targets or methods. A mineral-right seller can receive a separate option fee and later exercise consideration. A funding investor should see the total encumbrance of royalties, contingent payments and required work before accepting the capital stack. The schedule should preserve enough cash for an unexpectedly positive result that needs follow-on verification.

Figure 3. Author-developed four-gate funding and evidence sequence



Proceed, repeat within a cap, or stop after each evidence gate

Budgets and consideration are set by transaction documents; the diagram provides a governance sequence rather than a timetable forecast.

21 Conclusion

AI-assisted mineral discovery should be valued as a specific improvement in exploration decisions. The buyer needs a defined mineral-system question, defensible data rights, a reproducible model, independent validation and legal access to the targets selected. A model score is a ranking unless it has been calibrated for a stated event and context. Exploration targets require physical testing and competent interpretation before later resource or mine economics can be considered. The price paid today should reflect the rights and evidence that actually exist today, together with any contingent exposure the buyer accepts.

The fictional eight-target case illustrates the mechanics. Under author-assumed probabilities of 10% and 16%, a conditional milestone value of USD 18 million, eight tests and USD 3 million of incremental model cost, the model-assisted choice adds USD 5.64 million in expected value. That number depends on a genuine out-of-sample probability advantage, comparable rights and field costs, and the ability to realise the milestone. It

is not an observed return or a price for a software company or mineral property. The sensitivity test shows how quickly value falls when the advantage narrows or costs rise.

A workable transaction can pay for delivered data and model rights, reserve performance payments for independently verified improvements, and fund field work through explicit stop-and-proceed gates. Diligence should preserve adverse data and track the counterfactual. Geologists, qualified technical reporters, counsel and finance teams each have a defined role. The investor's central discipline is to pay for a measurable change in a decision while keeping the uncertainty of the underlying geology visible in every subsequent capital commitment.

Frequently Asked Questions

Can a prospectivity score be used as the probability that a mine will be built?

No. A score may only rank targets. A probability requires calibration for a defined geological or commercial event and relevant validation data; mine development has additional technical, legal and economic conditions.

What should an investor buy first?

The first funded stage should establish usable data rights, mineral title and a reproducible test of whether the model changes target selection against a credible conventional baseline. The specific transaction depends on the rights already held.

How should public geological data affect a vendor price?

Public inputs can support valuable analysis, while the vendor's exclusivity claim requires separate evidence of proprietary data, transformation, model performance, service or rights. Diligence should identify each layer's owner and permitted use.

Why use a spatially independent validation set?

Nearby sites can share geology and survey characteristics. Holding out districts or licence blocks can better test the proposed geographic use than randomly splitting neighbouring map cells, provided the outcome definition and data cutoff are fixed in advance.

What does the paper's USD 5.64 million figure represent?

It is the incremental expected value in a fictional eight-target example under expressly author-assumed probabilities, conditional value and cost. It is neither an observed market result nor a proposed price for any real company or property.

Can an AI model create a mineral resource estimate?

A model can inform exploration and later technical work. A reportable mineral resource depends on the applicable disclosure framework, adequate geological evidence and competent professional judgement; a ranked target alone does not meet that standard.

Which deal structure contains early uncertainty?

A bounded pilot, option, subscription or farm-in can stage consideration against delivery, independent validation and field milestones. Counsel should check mineral title, IP, licences, transfer and contingent claims for the actual jurisdiction.

What should happen after a negative drill result?

Record and review the result, update the geological hypothesis and model, account for all costs and obligations, and make an explicit decision to stop, repeat within a cap or test a different target. Do not erase negative evidence from the performance record.

References

1. US Geological Survey. Earth Mapping Resources Initiative: About. Accessed 5 October 2026. <https://www.usgs.gov/special-topics/earth-mri/about>
2. US Geological Survey. National-Scale Geophysical, Geologic, and Mineral Resource Data and Grids for the United States, Canada, and Australia, version 1.1. Accessed 5 October 2026. <https://www.usgs.gov/data/national-scale-geophysical-geologic-and-mineral-resource-data-and-grids-united-states-canada>
3. US Geological Survey. Basic Concepts in Three-Part Quantitative Assessments of Undiscovered Mineral Resources. Accessed 5 October 2026. <https://www.usgs.gov/publications/basic-concepts-three-part-quantitative-assessments-undiscovered-mineral-resources>
4. Coolbaugh, M., Raines, G., and Zehner, R. (2007). Assessment of Exploration Bias in Data-Driven Predictive Models and the Estimation of Undiscovered Resources. US Geological Survey publication record. <https://pubs.usgs.gov/publication/70029904>
5. Geoscience Australia. Minerals: Precompetitive Geoscience and Exploration Information. Accessed 5 October 2026. <https://www.ga.gov.au/scientific-topics/minerals>
6. Goldman, M. et al. (2025). Extracting Data from Maps: Lessons Learned from the Artificial Intelligence for Critical Mineral Assessment Competition. US Geological Survey publication record. <https://www.usgs.gov/publications/extracting-data-maps-lessons-learned-artificial-intelligence-critical-mineral>
7. US Geological Survey. Collaborative Workshop Spotlights Machine Learning to Accelerate USGS Critical Mineral Assessments. Accessed 5 October 2026. <https://www.usgs.gov/news/featured-story/collaborative-workshop-spotlights-machine-learning-accelerate-usgs-critical>
8. Committee for Mineral Reserves International Reporting Standards. International Reporting Template. Accessed 5 October 2026. https://crirsc.com/wp-content/uploads/2024/04/crirsc_template_v2.pdf

9. US Securities and Exchange Commission. Modernization of Property Disclosures for Mining Registrants: Small Entity Compliance Guide. Accessed 5 October 2026. <https://www.sec.gov/resources-small-businesses/small-business-compliance-guides/modernization-property-disclosures-mining-registrants-small-entity-compliance-guide>
 10. British Columbia Securities Commission. National Instrument 43-101: Standards of Disclosure for Mineral Projects. Accessed 5 October 2026. <https://www.bsc.bc.ca/securities-law/law-and-policy/instruments-and-policies/4-distribution-requirements/current/43-101>
 11. US Geological Survey. Training and Validation Data from the AI for Critical Mineral Assessment Competition, version 2.0. Accessed 5 October 2026. <https://www.usgs.gov/data/training-and-validation-data-ai-critical-mineral-assessment-competition>
- Sources reviewed 5 October 2026. The programme, probabilities, values, costs, figures, tables and transaction example explicitly described as fictional or author-developed are educational illustrations. This paper provides no geological, engineering, legal, tax, accounting, regulatory or investment opinion on a particular property or transaction.

About the Author

Authored by Chennakeshav Adya Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.