

Perspective

The future of AI in critical mineral exploration

Jef Caers^{1,*}

¹ Mineral-X, Department of Earth & Planetary Sciences, Stanford University, Stanford, CA 94305-2115, USA

* Correspondence: jcaers@stanford.edu

Received: November 21, 2025

Revised: December 17, 2025

Accepted: December 22, 2025

Published: December 23, 2025

Cited as:

Caers, J. The future of AI in critical mineral exploration. *Sustainable Earth Resources Communications*, 2025, 1(2): 69-82.

<https://doi.org/10.46690/serc.2025.02.05>

Abstract: The energy transition through increased electrification has put the world's attention on critical mineral exploration. Despite the promise of a growing demand, the global exploration industry is a money-losing enterprise. Even with increased investments, a decrease in new discoveries has taken place over the last two decades. In the paper, I propose a solution to this problem where AI is implemented as the enabler of a rigorous scientific method for mineral exploration that aims to reduce cognitive bias & false positives, enhances the role of domain experts, and drive down the cost of exploration. The current organization of exploration activities involving many fields of science (geology, geochemistry, geophysics) is no longer effective in discovering deposits under cover. In particular, the current approach fails to adequately quantify uncertainty, leading to sub-optimal decision-making and \$ spent on drilling that often result in false positives. Instead, I propose a new scientific method that is based on a philosophical approach founded on the principles of Bayesianism and falsification. In this approach, data acquisition is, in the first place, seen as a means to falsify human-generated hypotheses. The decision of what data to acquire next is quantified with verifiable metrics and based on rational decision-making. A practical protocol is provided that can be used as a template in any exploration campaign. However, in order to make this protocol practical, various forms of artificial intelligence are needed. I will argue that the most important forms are 1) novel unsupervised learning methods that collaborate with domain experts to better understand data and generate multiple competing geological hypotheses, and 2) human-in-the-loop AI algorithms that can optimally plan various geological, geophysical, geochemical, and drilling data acquisition, where uncertainty reduction of geological hypotheses precedes the uncertainty reduction on grade and tonnage. The approach will be illustrated using several ongoing exploration cases.

Keywords: Critical minerals; mineral exploration; artificial intelligence; sustainability; education

1. Introduction

The purpose of this paper is to provide a forward-looking essay on the use of AI in critical mineral exploration. This is a personal perspective. The writing is subjective and contains personal opinions and observations, but these are sourced from three decades of unique experience in the use of data science, decision science, and AI in the context of Earth resources exploration and development (Caers et al., 2026; Caers and Journel, 1998; Caers and Ma, 2002). In the oil & gas industry, it has long been recognized that imposing geological realism is essential to accurately predict fluid flow in porous media.

Machine learning and deep learning methods are now standard in generating reservoir models whose spatial variability is trained on realistic examples such as those obtained from outcrop data or from physical models of deposition (He et al., 2022; Kani and Elsheikh, 2018; Laloy et al., 2018, 2017; Lewis and Vigh, 2017; Liu et al., 2021; Nasir and Durlofsky, 2023, to list a few). Secondly, various diverse sources of information, such as from wells, seismic, and production data, need to be integrated into an ensemble of models (Ringrose and Bentley, 2015; Scheidt et al., 2018) and allow for effective decision making on field development scenarios (e.g., location of injection & producing wells). A key issue that is addressed

in reservoir modeling is that the different scale of information of these various sources needs to be carefully accounted for; for example, the volume of a core in a bore-hole differs significantly from the scale of investigation of an seismic amplitude signal at thousands of feet in the subsurface. This scale issue will be covered at length in this paper because it equally applies to mineral exploration.

Today, one can hardly imagine oil & gas reservoir exploration and engineering not using various forms of AI, whether it is machine learning, reinforcement learning or AI-based decision making. While oil/gas development based on AI, uncertainty quantification, and Bayesian philosophy (Caers, 2018) accelerated in the 21st century, its use in exploration, mining, and processing remains sporadic and anecdotal. Even though sophisticated methodologies have been developed (Rossi and Deutsch, 2014), they are not broadly applied and JORC reports still produce inverse distance weighting for resource estimation (JORC, 2024). At the same time, the ROI (return on investment) in exploration over the last few decades has become largely negative (Schodde, 2017) (except for certain commodities such as gold) and remain well below the oil & gas industry. While oil/gas companies wouldn't dream of building a large offshore platform without quantifying geological uncertainty, and thereby risk, most mines today are still developed using deterministic models. In exploration, deterministic modeling and inversion of multi-physics data is still the norm. Knowledge about the advances that happened regarding quantitative risk assessment in the oil/gas industry appears absent in the minerals industry.

However, things are changing rapidly: over the past decade, the global mining industry is now finding itself at the nucleus of the energy transition, with the demand for critical minerals likely increasing, as well as the corresponding geopolitical, social & environmental aspects. This is now generating the formation of new types of companies, often backed by venture capital, who are seeking to disrupt the mining industry, as they already have in healthcare, renewable energy and climate technology. A new wave of digital technology is emerging that may change the current economic, educational, social & industrial systems as we know them through the use of large language models. Whether this actually will have the impact it promises still remains to be seen, but the fear of missing out (FOMO) is ubiquitous. I believe that for the mining industry, no fear is warranted, just opportunity in the use of quantitative methods such as artificial intelligence. This paper looks at where those opportunities lie and how AI may contribute.

This text does not focus on listing possible applications of AI to exploration, nor does it write an exhaustive review paper, as these already exist (Lindi et al., 2024; Yang et al., 2024). Instead, I will view AI through the lens of better decision-making founded on a new scientific method. In part I of this paper, I will lay down the key decision questions that need to be made in exploration and how developing a scientific method to exploration will make the endeavor more efficient. Indeed, today, in mineral exploration, no agreed-upon scientific method exists. What exists is a set of tasks executed by domain experts, each with their own scientific method agreed upon as a

paradigm within their own fields. I believe that a broader scientific method for the mineral exploration endeavor is largely lacking, and I plan to argue that this is one of the most important issues in making exploration more efficient, meaning less costs to make a discovery. In Part II, I will then review the current most common applications of AI for exploration and point out common errors. Then, in Part III, I will conjecture what the most likely evolution of AI is in exploration, again through the lens of decision-making.

2. Part I: A new scientific method

2.1. What decisions need to be made in exploration?

The ultimate goal of exploration is to make a new discovery. Discoveries can be made near known deposits (brownfield exploration) or in an unexplored area (greenfield exploration). In the next decade or so, we are likely to see more brownfield discoveries than greenfield ones. Exploration can be seen as an exercise in reducing the scale of investigation by reducing uncertainty in knowledge of the existence of economic deposits. Exploration often starts at the continental scale, looking for example at large igneous provinces that may have the favorable tectonic setting where magma (source) can rise to the surface crust (pathway), then get trapped by physical means and chemical precipitation in large and shallow enough concentrations to be economical. At the larger scales, geologists are guided by geological mapping and airborne geophysical data. Anomalies detected on airborne multi-physics surveys are identified (possibly by comparison with those of known deposits) to reduce the scale of exploration on the ground by more detailed geological field work, ground geophysical surveys, and geochemical assaying of soils. The next steps are by far the costliest: drilling to confirm or falsify hypothesis of a promising geometries/lithologies (e.g., intrusions) and mineralization. In an ideal setting, such decisions are made rationally using decision science (Howard, 1968), uncertainty quantification (Scheidt et al., 2018) and understanding the value of information (Eidsvik et al., 2015) to minimize cost by optimally quantifying uncertainty relative to some economic threshold or other property of interest. Various data sources create an understanding of how the mineral system was developed. Data is used to make predictions for regions/zones/location of interest. Data may not just be acquired to directly detect mineralization but to test geological hypotheses or reduce uncertainty on orebody geometry. Not only are various decisions made, but the objective (reward in AI) of the decision changes during the exploration project.

By decisions, we mean the irrevocable allocations of resources (\$ amounts), not just mental constructs. Obviously, in exploration, one of the main decisions is: what data should we acquire next, or should we walk away from the current project. In Part III, I will argue that such rational decision making is largely absent for reasons that have nothing to do with science, instead, with the way the exploration industry is currently organized.

To illustrate the need for a scientific method, I will cover one particular exploration activity: exploration under cover.

2.2. A typical approach to exploration under cover

The world has seen a decline in new (greenfield) discoveries over the last decade. Not because of a lack of investments, in fact, they went up in \$ values (Schodde, 2017), but because many outcropping deposits have been found. Exploration now needs to go under cover. This means that indirect observations, such as geophysical & geochemical anomalies, become increasingly important. A direct observation consists of an observation of mineralization at the surface during geological field work, or during drilling. Here, we discuss a hypothetical case where drilling has not been done, and hence, in a promising case, needs to be planned. Here would be a common approach today

- Deterministic inversion of gravity, magnetics, and EM into 3D subsurface models of density, magnetic susceptibility, and conductivity.
- Interpretation of the 3D property variation into features/zones of interest (e.g. and intrusion, an alteration, a lithology, a fault).
- Possible follow-up with additional surveys now on the surface of the Earth, then possible re-interpretation of features of interest.
- An expert-guided decision on drilling, informed by the deterministically interpreted anomaly, often with the goal to maximize intersection.

While this approach appears logical and certainly has led to successes, my observation in working with companies is that in many cases nothing of interest (and rarely reported to the public) is intersected when drilling. In other words, a false positive is observed and worse: the reason why is not understood or evaluated, one simply moves on to the next anomaly. Why is this the case?

The largest problem is determinism. A single geophysical inversion is given a single geological interpretation. Uncertainty at every step of this process is ignored.

In the next section, we will take a very different approach, where the following changes are made

- Deterministic inversion becomes stochastic inversion.
- Interpretation includes a falsification step to eliminate wrong modeling assumptions because they may induce false positives.
- Geological realism is included right from the start, not after inversion. We invert for geology, without vague (non-geological) regularization terms that are typically used in deterministic inversion.
- We propose drilling plans as a collaboration between experts and metrics of uncertainty quantification.
- We plan drilling with a given objective in mind (maximal intersection, volume of the intrusion) and do such planning based on quantitative metrics of optimality for the proposed drilling plan.

These changes are not just random choices, they are motivated by a purposeful way of thinking about data, models, objectives, and decisions that emanate from a long history of scientific approaches. In the next section, we will emphasize this through a philosophy of science reasoning, in particular focus on two major philosophies: Falsification & Bayesianism.

2.3. A scientific method for mineral exploration

2.3.1. The scientific method

Humans do science, but also study how science is done. This raises philosophical concepts that question how science should be done. What is a good scientific method? The practice of science has evolved over human history, starting from empiricism to questions that involve the relationship between models and data. It also considers how humans organize to do science, because various types of social interactions define what standards and norms are, even if they may later be proven incorrect. Insiders often don't see faults or epistemic errors, because the community has accepted them as standards. Scientists may be overly optimistic about what they know (Bond et al., 2007). Thomas Kuhn, in particular, studied these organizations and called them "paradigms". He sees paradigms as consisting of certain (theoretical) assumptions, laws, methodologies, and applications adapted by members of a scientific community (e.g., evolution, plate tectonics, genetics, theory of general relativity). For example, probabilistic methods can be seen as such a paradigm: they rely on axioms of probability and the definition of a conditional probability, maximum entropy, the principle of indifference, Monte Carlo simulations, and so on. Researchers within this paradigm do not question the fundamentals of this paradigm. Activities within the paradigm are then puzzle-solving activities (e.g., interpreting faults on a seismic survey) governed by the rules of the paradigm. But often unresolved issues fester under the surface. People who question paradigms are often considered crazy, ignored, and ostracized from the community. Many such examples in human history. Some end up winning Nobel prizes (a recent example is Katalin Karikó, Nobel Prize in Medicine, 2023, who was shunned, ignored, belittled, and "not of faculty quality" at the University of Pennsylvania even when her groundbreaking Nobel Prize work was already established).

When it comes to mineral exploration, there does not seem to be any clearly defined scientific method. What does exist are individual fields of science, such as economic geology, geochemistry, and geophysics, each with their own organization of puzzle-solving activities. While each component is essential, the question begs if the absence of an overarching scientific method in exploration leads to sub-optimal outcomes.

2.3.2. The Bayesian philosophy

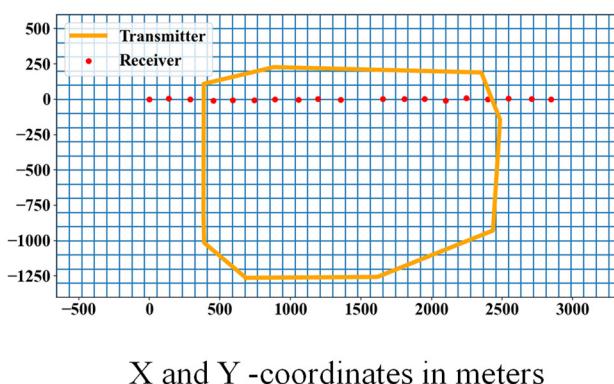
While Bayes' rule has been around for more than 250 years, the Bayesian approach is only a recent trend in defining a new scientific method. What sets Bayesian probability from "regular" probability (in the notion of Pascal) is the use of prior information. Bayes' came up with a problem of gambling where the odds of winning itself were uncertain (certainly true in exploration). There

was no true p , the probability of winning. Previous scientists, such as Pascal, had tried to solve such problems using a classical probability approach, such as maximum likelihood estimation of p , which relied on defining how likely a dataset would be, if we were to know the truth. However, Pascal's answer to such a problem was shown to be biased by Bayes'. Even though we do not know a true p , we can conjecture on the plausibility of possible values. This idea is captured quantitatively in the notion of a prior probability model on the parameter p of a gamble. For an unknown chance of success, we may then assume it is between 0 and 1, and we could assume a uniform distribution. Why uniform? That's the subjectivity of the prior, but is better than assuming one should simply estimate a single p . Accepting the notion of prior uncertainty is counter-intuitive, yet essential to quantify uncertainty and thereby making optimal decision. Tversky & Kahneman demonstrated through experiments that humans ignore Bayes' and they ended up winning the Nobel Prize for economics in 1973. Yet, my own experience in working with those who do mineral exploration is the avoidance of a prior exactly because it is considered too subjective. They reason that assuming there is only a single true p is considered not subjective! This serious

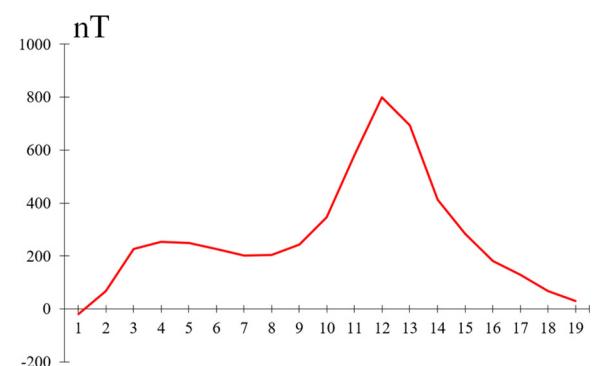
contradiction has hampered the rigorous development of a scientific method for mineral exploration.

Bayes' rule essentially comes down to stating that in order to make predictions with models and data, we need to know two elements: 1) how likely the new data is assuming a given truth, and 2) how likely some truth is before acquiring the new data. Traditional statistical approaches, such as deterministic inversion only use the first (likelihood).

Let's consider the problem described in Fig. 1, where we have a time-domain EM survey and would like to determine electrical properties in the subsurface. The transmitter is a large current loop transmitting an EM signal into the Earth, and 19 receivers record the vertical component of the secondary magnetic field, produced by a conductivity anomaly underground, over 12 time channels along a survey line. Only one time channel is shown with a clear anomaly at stations 11-14. The surface topography is mostly flat, and the subsurface is discretized into $45 \times 20 \times 45$ cells in X, Y, and Z directions, extending 4000m, 2000m, and 4000m in these directions. The depth of discretization is determined by the diffusion distance corresponding to the last time channel.



X and Y -coordinates in meters



19 stations (red dots in left figure)

Fig. 1. (left) layout of the surface EM survey, with one loop and 19 stations (receivers), (right) one of 16 time channels. The inversion uses all time channels. (courtesy of Zhuo Liu, Stanford University)

The traditional approach is to perform a deterministic inversion, then use that to make a deterministic interpretation, which is shown in Fig. 2. Then an expert geologist will interpret the deterministic inversion in terms of shapes, such as intrusions relevant to exploration. Interpretations of the geometry are done in software that focuses on a single deterministic model.

With this approach, geophysical inversion happens before geological interpretation. Again, one may argue that the regularization terms help impose geological knowledge, but this is too vague and very limited. In a way

the current approach causes a double dip of the same geological knowledge: first, we have to impose it in the regularization term, then we impose it once more in the interpretation. It becomes a self-fulfilling prophecy and fraud with circular reasoning. In Bayes' both come together and are solved as a single problem. Bayes' requires one to state plausible solutions (geometries) before doing any inversion.

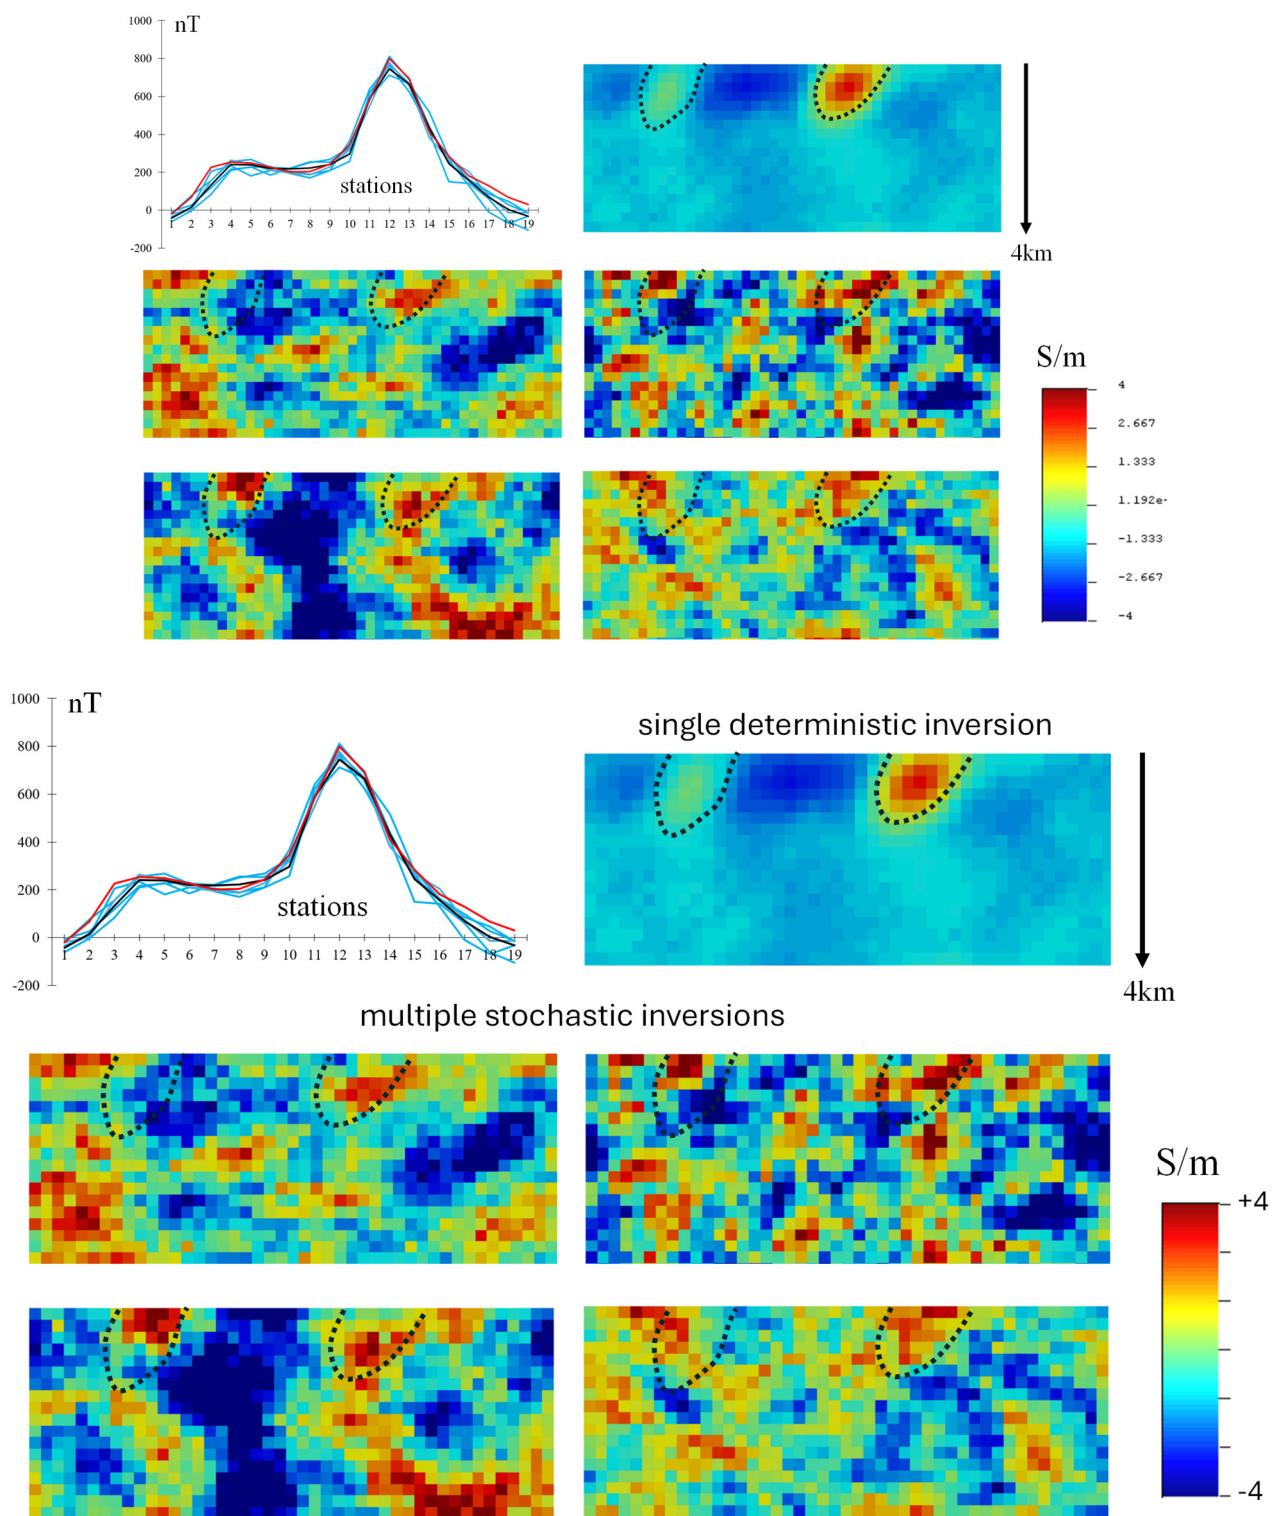


Fig. 2. deterministic vs stochastic information. (courtesy of Zhuo Liu, Stanford University). In deterministic inversion, one first inverts the data, then makes a geological interpretation (dotted line). (bottom) are shown four stochastic inversion all matching the same data (blue lines in top left figure) compared to deterministic inversion (single red line). On the four realizations the same deterministic interpretation is overlayed to show that the clearly shaped intrusion may be very different in reality.

Bayesian inversion offers relatively simple solutions to the problem, which can be coded up surprisingly easily and are general as well. Let's revisit the case in Fig. 2 and discuss how a prior probability model can be designed. We start with the following question: what is the minimum amount of information we have about the subsurface

before acquiring any data? First, we have previously observed that petrophysical properties have certain ranges, and if we know what lithologies we are dealing with, we could possibly even get more specific about these ranges. Secondly, we know the spatial variation of these properties is not random, nor constant over very large areas.

Indeed, we know that the formation of deposits has very specific variations of these properties. Conceptually, the minimum information is constraints on histograms and the presence of spatial correlation. Mathematically, we can translate the latter into a spatial covariance model (in geostatistics: variogram), but we are uncertain about the parameters of this variogram, such as the mean, ranges, and anisotropy, in fact, they could be pretty much anything (except random!). What we have established is a hierarchical way of thinking as follows:

- Conceptual idea → high-level mathematical assumptions → detailed mathematical assumptions

which for this case becomes:

- The idea of spatial correlation → spatial covariance → parameters of the spatial covariance

Fig. 2 now shows that one can generate multiple stochastic inversions, all matching the data equally well as the deterministic inversion, but the spatial variability (variogram) is different for each inversion. Overlaying the deterministically interpreted intrusion, we notice that intrusions may be present in many different shapes, sizes, and orientations, all very different from the single interpretation.

This type of approach can be used in any setting: geologists describe conceptual models (the diagrams found in papers/reports), mathematicians turn these into constraints on the prior probability model, geophysicists use this prior to perform Bayesian inversion.

Fig. 3 shows an actual case study where this is applied (Wei et al., 2025). The case concerns a well-studied intrusion (O’brian, 2018), the Crystal Lake Gabbro (CLG), a mineralized mafic intrusion located along the western shore of Lake Superior and about 7 km north of the USA-Canada. Decision need to be made on where and how to drill boreholes to further delineate the intrusion. Decision-making under uncertainty requires quantifying the uncertainty of the location and shape of the intrusion. However, uncertainty exists in the style of the intrusion, which can be compared to known classes of intrusions (Barnes et al., 2016). One hypothesis is that of Nebo-Babel type magmatic Ni-Cu sulfide mineralization; another possibility offered is the Eagle-Kalatongke type mineralization, which presents a funnel shape. These shapes can be represented using simplified mathematical models, of which then many realizations can be generated. The set of realizations constitutes a quantification of uncertainty on the shape. These realizations are a direct input to AI-based decision making as will be discussed in Part III.

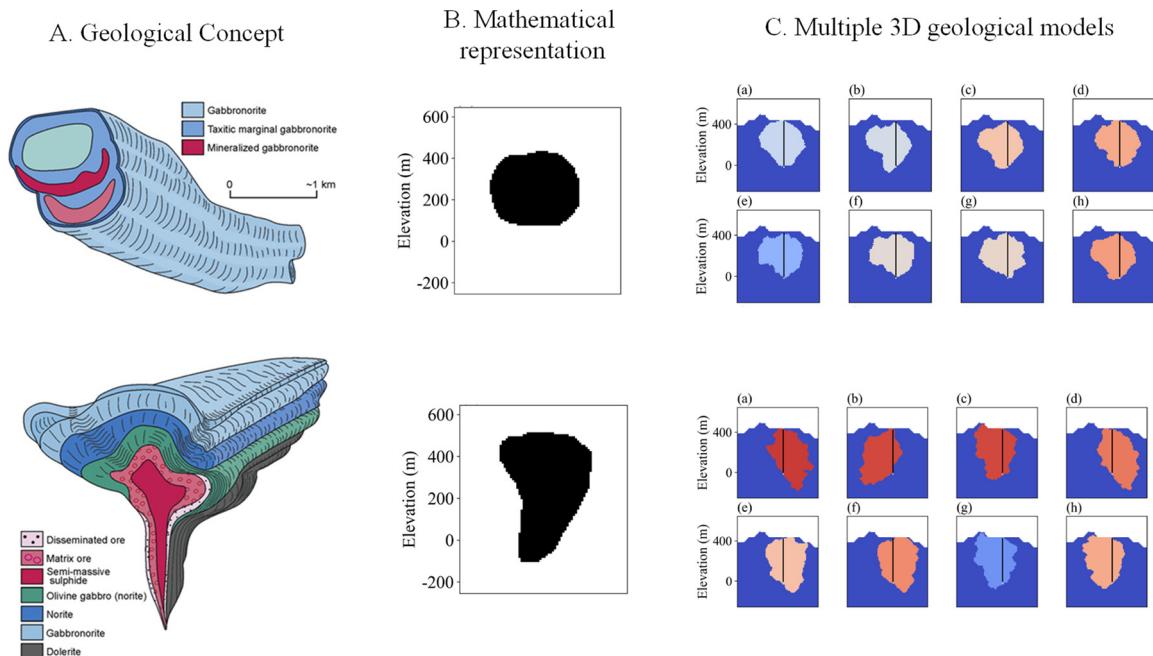


Fig. 3. several geological concepts are hypothesized on the geometry of magmatic intrusion containing massive sulfides based on studies of intrusions worldwide, here we show two such concepts. When applied to the actual CLG case, these concepts have to be turned into mathematical models, which are often simplifications of the actual intrusion. Using this model, multiple realizations can be generated (possibly constrained to existing borehole data), explicitly representing uncertainty in the intrusion shape. Models are in 3D, but only a vertical cross-section is shown (Wei et al., 2025).

There is a snag in this framework: how do we know that the hypothesized conceptual idea, and resulting mathematical model, is not conflicting with the geophysical data or any other data? After all, the initial conceptual idea may be completely incorrect, or the mathematical model oversimplified. This type of problem is an example of a wider problem of science: how do we prove that the assumptions we make are correct? It is the Austrian

philosopher Karl Popper who spent a lifetime analyzing philosophical questions like these.

2.3.3. Karl Popper: Falsification

Popper thought that science should not involve induction (theories derived from observations). Instead, theories are seen as speculative or tentative, as created by the human intellect, usually to overcome limitations of previous theories. Once stated, such theories need to be

tested rigorously with observations. Theories that are inconsistent with such observation should be rejected (falsified). The theories that survive are the best current theories.

In terms of UQ, one can then see models not as true representations of actual reality but as hypotheses. One has as many hypotheses as models. Such a hypothesis can be constrained by previous knowledge, but real field data should not be used only to match models (confirmation with data) but to falsify a model hypothesis (reject, the model does not conform with data).

Popper offers a way to address (but not fully solve) the subjectivity of Bayesian reasoning, which is a form of inductive reasoning, as opposed to the deductive form of falsification. If the prior distribution cannot predict the current data (up to data error), then the Bayesian prior is falsified. It basically means that no matter how many models one samples from the prior, each model conflict the data. While Popper's falsification may render the prior less subjective (than before testing), he does not offer a way to revise the prior that will still be up to human subjectivity.

2.3.4. Moving forward with a new scientific method

A scientific method for mineral exploration can be built based on Bayes' and Popper, let's term it a Popper-Bayes approach. In such a scientific method, a certain paradigm/protocol can be followed:

- Domain scientists generate conceptual model hypotheses by interpreting existing data. These are high-level hypotheses, conveyed descriptively
- Experts in numerical modeling turn such high-level hypotheses into a numerical representation that at the highest level has many quantitative hypotheses translated from the conceptual hypothesis but also implement details through model parametrization.
- Model hypotheses are assigned discrete probabilities, while detailed variables require probability distributions.
- This prior distribution model needs to be tested with data. This involves numerical modeling (also termed forward modeling) of the data, done by domain experts of these data sets. A falsification test is a statistical test that requires statistical metrics and p-values.
- Domain experts in Monte Carlo simulation will provide a Bayesian inversion of the model hypothesis and variables given the data.
- Make decisions based on the Bayesian uncertainty quantification.

To achieve this, all domain experts need to be in the same room. In Part III, we will discuss the various opportunities AI brings to make this paradigm possible. Before that, I will review some of the most common application of AI in mineral exploration.

3. Part II: Current applications of AI

3.1 What is AI?

To discuss the various forms of AI (Kochenderfer, 2015; Russell and Norvig, 2016; Sutton and Barto, 1998), I will use a specific example instead of formal definitions. Consider a self-driving car. Consider such a car given the task of getting from A to B. The goal is to arrive safely & lawfully, within a reasonable time, at its destination. To do so, it needs to take various actions of steering, accelerating, decelerating and steering. Sensors classify the current state of the road, and predictions are made for future actions of itself and other road users or objects. Let's identify the major components

- *Rewards*: the goal: safe, lawful, fast
- *Actions*: steering, braking
- *State*: the current situation on the road with markers, other cars & people, and their direction and velocity
- *Data*: various sensors of the car are used to predict the current state and the next.
- *Belief*: sensors provide imperfect measurements, so the true state is not fully observed, only a belief of the true state.

Solutions to these problems start by formulating the problem into mathematical equations. These equations are essential, and it would be ineffective to right away start solving the problem by collating a set of algorithms. Without going into the details of these equations, we can identify the following mathematical components

- From the current state and actions, a prediction needs to be made on what state we find ourselves in at the next time step. However, actions may not be executed perfectly, so a transition probability model is defined.
- Sensors allows to measure the current state, but sensors are not perfect. This uncertainty (belief in AI) depends on what state we are in. A prediction model, termed a belief model, allows updating the current state with new data. Bayes' rule is often used.
- Routes need to be planned and executed, but updates on such plans need to be made. A planning of a route beforehand is called offline planning, while replanning while executing is termed online planning.

Notice that two fundamental uncertainties exist: state transition uncertainty and model (state) uncertainty (belief). Cases with transition uncertainty are often solved using reinforcement learning, while those with model uncertainty are formulated as partially observable Markov decision processes. Partially observed refers to an inexact observation of the true world, while Markov refers to the sequential planning aspect.

Once we have a formulation, we can apply algorithms and tools to solve them:

- Predictions can be modeled using machine learning methods, such as neural networks or generative (AI) algorithms.
- Offline and online sequential planning under uncertainty can be solved using Monte Carlo Tree search algorithms

Having identified these two major components of AI (predicting and deciding), I will review how they are currently applied in exploration.

3.1.1. AI that makes predictions

Making predictions requires datasets. The most important challenge in mineral exploration is the creation and curation of datasets, not the type of machine learning method that is used. Several challenges exist in creating such datasets that can be multi-variate, 3D spatial, compositional, and extremal. Rarely does one encounter such a challenging prediction/learning problem in other domains. I will consider three crucial and often overlooked or unaddressed components of data curation:

- Not all data is on the same scale
- Not all data should be pooled into a single data set
- Learning in area A may not apply to area B

Common errors:

- Upscaling ground truth data (e.g. a core sample) to the geophysical scale.
- Deterministic interpolation of geophysical and geochemical datasets.
- Merging data from different zones cannot be explained geologically.
- Applying machine learning models from Area A to Area B without testing the validity of doing so.

All these common errors lead to an increase in false positives (see Fig. 4, for an example), and gives AI a bad name in the industry, while the actual problem is not AI, it is people curating the data.

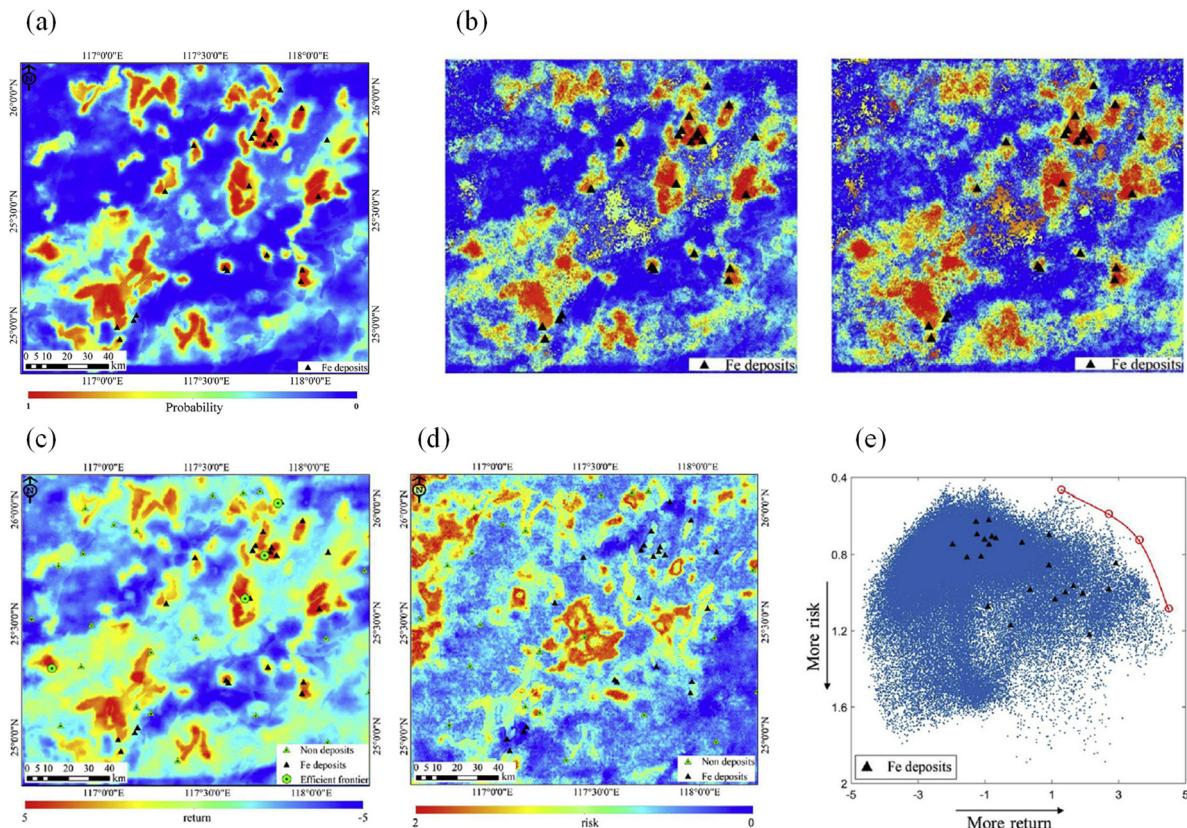


Fig. 4. case study of mineral prospectivity mapping comparing deterministic vs stochastic interpolation of input data (see (Wang et al., 2020)) for details. In deterministic interpolation (a), all data (geophysical, geological, geochemical) are interpolated deterministically, producing input that has less variance than the data itself. In stochastic interpolation (called simulation), the scale (i.e., variance) of the data is accounted for. As a result, one does not simply obtain one possible MPM, but many possible ((b) shows two realizations) due to the propagation of interpolation uncertainty. When comparing the two approaches with known deposits, the recall of the deterministic prediction equals 0.4, while the recall of the stochastic approach is 0.7. Additionally, in a stochastic approach, one can assess a return (c) vs a risk (d). Here, the return is the mean of the log odds ratio, while the risk is its variance. Locations where risk and return are optimally traded off are good candidates for further exploration (e).

A significant error in the current paradigm is something that has long been known in mining geostatistics (Journel and Huijbregts, 1978): data at various scales may have significant differences in variance. An example is that of comparing interpolated geophysical data with outcrop sampling or core measurement, see Fig. 5. Any deterministic interpolation will, with certainty, reduce the variance: the interpolated geophysics has a

smaller variance than the actual data. Second, label data from drilling at cm scale is often upscaled (or worse, moved) to enable a comparison with the geophysical data. The result is an increased correlation between label and predictor that is induced (not real) because of these operations. As a result, current mineral potential or prospective mapping have a significant number of false positives.

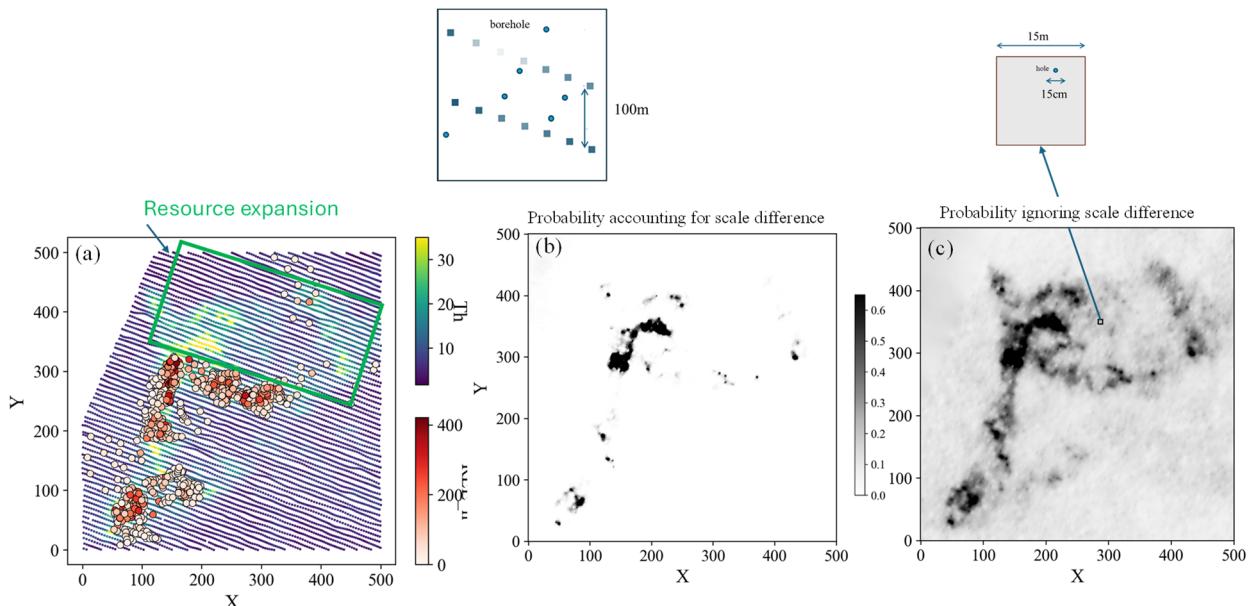


Fig. 5. Example of the impact of ignoring the difference in scale when performing machine learning with geophysical data and borehole data (size = ~15cm, see (c)). Shown is a case of rare earth element resource expansion (Caers et al., 2025), where an area of high drill density (in (a)) is used to expand the resource into an area with almost no drilling. The geophysical data here are radiometric data (Th and U) along flightlines, which can be correlated with the REE borehole concentration. The line spacing (a) is about 100m, and the grid onto which the radiometric data is interpolated is about 15m (b and c). The interpolation is needed in order to compare the radiometric data with the REE concentration. In deterministic interpolation, this scale difference is largely ignored, and the assessed correlation coefficient is artificially high because of the smoothing of the interpolators (Wang et al., 2020). In stochastic simulation, one accounts for the variability of the data and addresses the uncertainty of the interpolation. As a result, the predictions are much more focused (compare (b) with (c)), while the traditional approach of ignoring scales leads to overly optimistic predictions (and possibly false positives as a consequence).

Instead, mineral exploration can be seen as a reduction of uncertainty over multiple scale. Mineral prospectivity maps are informative at the large scale but for the smaller scale, additional data through field work is required to narrow down to the prospect scale. Fig. 6 shows an example in Ni-Cu-Co exploration in Canada, where exploration went from the 1000km scale to the 30cm scale. The role of machine learning here can be essential, in the sense that any new information can be used to directly update mineral potential maps and thereby guide geological field work in real time.

3.1.2. AI that makes plans

While machine learning and generative techniques are popular and ubiquitous, the applications AI in decision-making is less common. Indeed, decision-making requires the irrevocable allocation of resources, which is not with the realm of academic research. To date, the only

significant and impactful application of decision-making AI in mineral exploration is the characterization of the high-grade portion of the Mingomba deposit in Zambia (Dempsey, 2024). In this case, historical drilling was re-analyzed using more recent data analysis tools, and new hypotheses were generated that include a higher-grade portion than was previously considered. AI (formulated as a POMDP, Partially Observable Decision Processes) was used to sequentially plan drilling. The initial goal of such drilling included the option for the AI to reduce uncertainty on geological hypotheses on the spatial extent of the high-grade mineralization. This method was shown to be much more efficient than the traditional grid-based drilling (Mern and Caers, 2023). The efficiency comes from focusing on falsifying geological hypotheses rather than merely attempting to hit the mineralization.

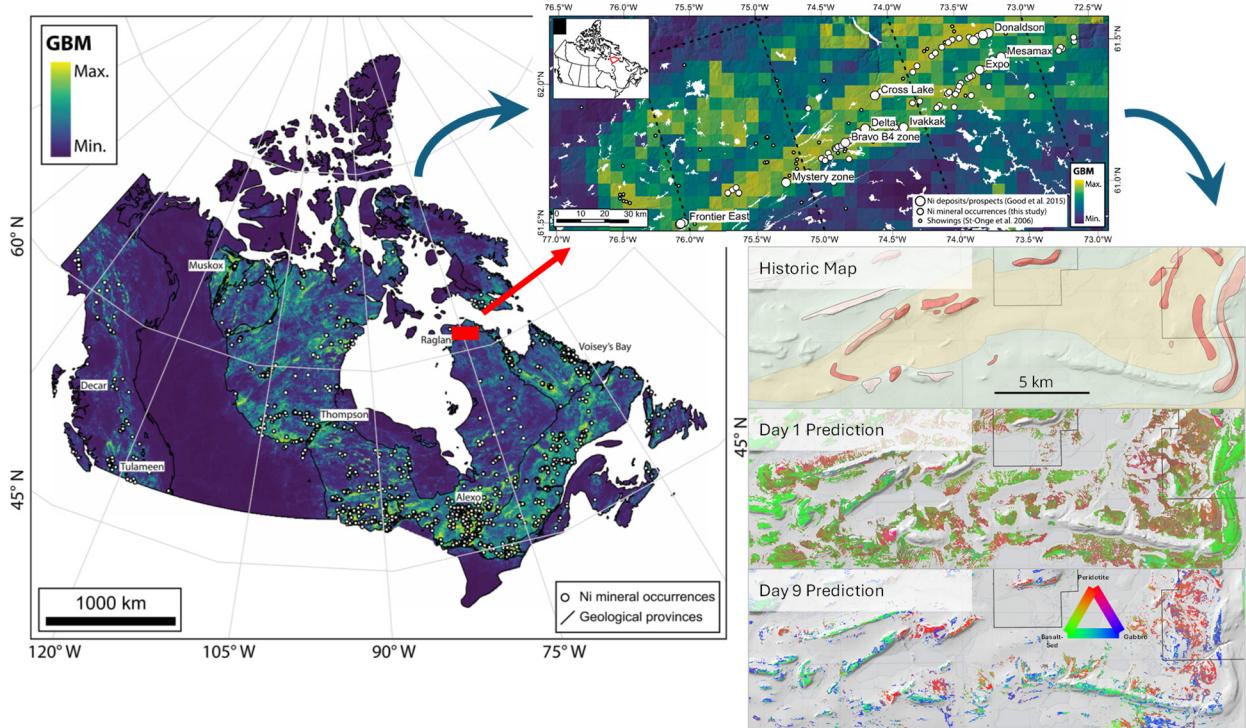


Fig. 6. (left) a map showing the potential for Ni-Cu-Co mineralization, the yellow color indicates high potential (from Lawley et al., 2021). (top) a zoom into the Cape Smith area, which has high potential. The resolution of the map (the size of the pixel) is 5×5 km. (right) AI-assisted geological field work, the resolution of this map is 30 cm (multispectral imaging data), allowing geologists to identify efficiently locations of mineralized rocks (red color). Several of the red areas in the day 9 prediction were later confirmed to have Ni-Cu-Co mineralization.

4. Part III: The future of AI in mineral exploration

In Part III, I will focus on how AI can help accelerate the discovery and characterization of critical mineral deposits. I will focus first on data analysis and prediction, then on AI-based decision making.

4.1. Human-in-the-loop data science & ML

The initial deployment of AI was met, in some corners, with a level of defensiveness. Indeed, the perception was that AI would replace or somehow reduce the importance of domain experts. Nobody likes to see their jobs targeted, and this is indeed a common fear today. The perception was that AI would make a new discovery. This will not happen in the way of an AI putting a big red cross on a map where robots will go and drill, make a discovery. The way I view AI in this context is to enhance the domain expert, by removing tasks that the expert no longer needs to do or is no good at, and focus on the creativity of their process-based thinking (geology and geochemistry). Let's start with what humans are good at:

- Generate conceptual hypotheses on how mineralization took place from large-scale to deposit scale.
- Interpret relationships in time and space of various geological units.
- Discover patterns in low-dimensional datasets.
- Build complex deterministic numerical models, ensuring geological plausibility.

and what they are not good at

- Read 1000s of documents and create summaries and query systems in a decision-relevant time.
- Discover patterns in high-dimensional multivariate and/or spatial datasets.
- Avoiding cognitive bias.
- Realistic uncertainty quantification.
- Make rational decisions that involve uncertainty.

The most significant potential is an AI-system that interacts with humans and applies novel data science on high-dimensional space-time data to generate plausible hypotheses of mineralization, its extent and grade, and that is decision-ready (e.g. for further data acquisition decisions). And these hypotheses should be falsifiable, meaning that next to stating the hypotheses, one should state what data or data outcome will falsify the hypothesis. The latter is critical in avoiding cognitive bias. The decision-ready is essential; deterministic mapping does not provide such readiness.

Foundational to this system are data-science tools. As mentioned above, the most important problem is not the choice of the machine learning algorithm, but the proper curating & thorough analysis of data to generate hypotheses. Humans tend to focus on what they know are indicative relationships of mineralization. For example, in the exploration of massive Ni sulfide deposits geologists use the scatterplot between MgO and Ni to classify the

potential for a massive sulfide mineralization (Barnes et al., 2016). This bivariate plot, possibly supported by a few more, can be enhanced by data science and AI methods that may discover additional high-dimensional relationships (four or more) that were previously not discovered. As such, development and use of advanced data science and AI-based or enabled modeling & analysis tools will lead the way in aiding domain experts, for example

- Detection of anomalies on high-dimensional geological & geochemical data.
- High-dimensional pattern analysis of geophysical data to make comparisons with known deposits.
- Non-linear versions of basic data science tools such as PCA, MDS, and FA.
- Spatial clustering techniques that can detect complex spatial relationships.
- Advanced geostatistical methods can create complex surfaces constrained to various datasets.
- Data science tool that focuses on falsification, instead of data fitting.
- Multi-physics stochastic inversion from borehole to country-scale.
- Surrogate modeling to enable multi-physics inversion.
- Generative AI to create realistic 3D variation of structure, lithology & mineralogy.

This also means that a collaboration between a data scientist and a domain expert is where the magic will happen. We will comment on this topic later.

4.2. AI to support the Popper-Bayes protocol for optimal decision making

While data science and ML aid in generating hypotheses of locations of interest, decisions need to be made on what to do next, mainly, what type of data is now needed, typically

- Dense airborne geophysics
- Surface geophysics & geochemistry
- Drilling

Rational decision-making based on the approach we discussed above means

- Generating multiple high-level geological hypotheses at the deposit scale.
- Generating multiple models of the subsurface constrained to current data, under each geological hypothesis. Generative techniques from the oil & gas industry are readily available (see introduction).
- Determining the goal of the drilling, i.e., the reward, or multiple rewards, such as better knowing

- Depth of mineralization
- Extent of mineralization
- Maximizing intersection with mineralization

Once multiple hypotheses are generated, metrics of optimality can be used to determine what data should be acquired next. Value of information is a common metric in decision science that can be used to rank data acquisition campaigns based on rigorous economics. In the absence of economics, a dollar-free metric, the efficacy of information (Caers et al., 2022) has been introduced that quantifies how much future data will reduce uncertainty on average on some quantity of interest. Mathematically, this metric compares what we know about the quantity x without the data (the prior probability distribution $f(x)$) and what we would know when we have the data (the posterior distribution $f(x|d)$ for various possible outcomes of the data d and for a given survey design y), the equation goes as follows:

$$EOI(y) = \int_d \int_x f(x|d(y)) \log \left(\frac{f(x|d(y))}{f(x)} \right) dddx$$

This metric can be calculated for various data acquisition proposals y using Monte Carlo simulation (Caers et al., 2026).

When performing sequences of measurements, such as planning a sequence of drilling, one can use AI frameworks, such as Partially Observable Markov Decision processes (POMDP) solved using Monte Carlo tree search, to determine the optimal sequential drilling plan (Mern and Caers, 2023). The properties of interest, however, will change, namely the following hierarchy can be defined. Drill for hypotheses, then volume, then grade. Indeed, once drilling has falsified all hypotheses but one, focus should be on the volume of the mineral system, then the grade, or, if desired, combine in a grade-tonnage curve.

4.3. Implementation in industry

Despite the promise of a comprehensive scientific method to make exploration more effective and thereby less costly, several non-scientific barriers exist in the exploration enterprise. These are 1) ingrained practices & company organization 2) the exploration industry organization 3) no comprehensive software.

The mining industry remains a conservative industry where innovation is met with resistance. Most large companies focus on making a profit from large mining operations and their investment in exploration remains relatively low. Large companies are also hierarchically organized in disciplines, rather than in problem-solving. I know of at least one large company that has ~1% of experts in data science and ~99% in the geosciences.

Companies are organized such that experts in each field handle the responsibilities best suited to their training and experience. The emphasis is often on figuring out details, without knowing whether they in fact affect the ultimate decision. A successful implementation of AI requires a re-organization of large companies around decision-making problems as opposed to task-executing mind-set. This is not simply a matter of reorganization but also a matter of education. Geologists don't need to

become data scientists, but there is a great need to learn each other's language, a matter I will revisit later in the education section.

Socialization still happens around a single model, performed with software that is complicit to allow biased & deterministic interpretation. Over time, these models are attributed to a sense of truth. As people move on, the reason for the actual interpretation is lost, and the particular geometry created is now considered data, instead of a model. Uncertainty in a company is often considered a failure. Not knowing is equated with "lack of knowledge" of a person. If only the person spent more time on the problem and worked harder, uncertainty would be reduced.

Proper software for a rational decision-making approach is lacking. The goal of a software company is to sell software that people want to buy, and those who buy it like software that does what they want it to do, according to the leading paradigm. Rational decision making based on rigorous uncertainty quantification often requires high-performance computing, such as in the implementation of Monte Carlo sampling methods, and comes at a \$ cost. Small exploration companies are not going to spend tens of thousands of dollars on a cloud computing system, nor will they hire people with the expertise to make this happen.

Finally, most exploration is done using junior mining companies, i.e., small companies owning one or two prospects, employing a few geoscientists. Funding for such an operation is often based on showing some form of success (hitting mineralization with one borehole), and the pursuit of additional drilling dollars from investors trumps a rigorous scientific approach to decision-making. Falsifying a geological hypothesis requires long-term thinking in terms of data acquisition, while junior mining companies focus on the short-term only.

Major mining companies have made venture capital available that brings new talent and interest to the mining industry. The emphasis lies in bringing technological innovation in either sensing, computing, or predicting, with often a focus on AI. These are much-needed initiatives. Besides technological innovation, a systemic change needs to happen within which these technologies can thrive. This means that resource assessment protocols need to become open to new technologies, even encourage them. The JORC reporting standard is still dominating how exploration is done, because these are used by investors. A proposal would be to extend the current framework, setting standards for uncertainty questions that are steeped in the same practices as in the oil & gas industry. Such standards may be accompanied by reporting on computing efforts for verification and validation of results, in addition to the mandate of a qualified person. After all, the more information and effort is put into an assessment, and supported by those expert with experience, the more likely the resource assessment will be rigorous, not just in the estimate but in the quantification of uncertainty of such estimates. To date, most uncertainty quantification in the mineral resource industry do not rely on Bayesian methods, while it is much in the interest of someone owning property to be Bayesian in order to know their risk exposure.

It is also clear that new financial vehicles are needed. A one-borehole-at-a-time investment is like betting on a single stock each time. However, the common way to invest in stocks is to have a portfolio where risk and return are traded off. The industry should therefore move to investment vehicles that bet on a portfolio of prospects, by understanding the trade-off between risk (uncertainty) and return (estimated present value).

4.4. Sustainable decision-making

Mining has a clear environmental impact and affects the communities that live near these mines. This means that not every economic deposit should be mined if the effect outweighs the impact. Nevertheless, one can easily argue that this impact is inversely proportional to the grade and size of the deposit. Low-grade copper mining in the US, now mined at less than 0.5%, requires large open-pit mining, while a 6% grade mine in Zambia can be mined using underground mining techniques. Thus, exploration for high-grade deposits will lead to more sustainable mining. Unfortunately, sustainability (social & environmental) is rarely a metric in exploration, and the junior mining company model is not helping either. Sustainability is often a secondary consideration after resources are turned into economic reserves. This may well turn out costly to the company, as construction is delayed due to the ensuing lawsuits. Sustainability must therefore be grounded in jurisdiction-specific social and environmental contexts and regulations in federal, state, and local levels. Decision-making around exploration should include these additional metrics and preferably should be done on a portfolio of prospects, rather than looking at an individual one. In such a portfolio, the risk-return profile of a single deposit now includes all factors and can be traded-off with other prospects.

4.5. Knowledge retention and decimation

Having been in Academia for more than 30 years, arguably, its largest transformation started in the last few years with the arrival of LLMs. The professor needed to rush to rewrite their syllabus, the way homework and exams were done. I believe this transformation is long overdue. Problems facing humanity, such as climate change and the energy transition, are far exceeding the breadth of knowledge of a single scientific discipline, yet we still don't offer a degree in Mineral Exploration. Secondly, information is now produced instantaneously, and its intelligence is still improving. Going to the library to copy from a journal is transformed into an algorithm that can search and summarize 1000s of documents, almost in real time.

At the same time, interest in degrees such as economic geology and mining engineering keeps decreasing in the Western world. Now more than ever it is urgent to retain the century-old knowledge about the formation of ore deposits, some stored in dusty geological reports in forgotten closets around the world. Generation of high-quality databases such in Canada and Australia is much needed all over the world. In Africa, only one country has opened such a database (<https://gsd.gov.zm> (accessed on December 2025)). Making such information accessible

and open on web platform is essential to attract investment.

5. Summary

In this personal essay, I advocate for the scientific and industrial community to come together around a new scientific method for mineral exploration. As exploration is moving “under cover” the existing paradigm of starting with surface outcrops assisted with deterministic geophysical inversion is no longer effective, as indicated by the declining discovery rate. Instead, a decision-focused approach that explicitly assesses risk, through uncertainty quantification, will reduce the number of false positive drilling outcomes. This uncertainty quantification will need to acknowledge the fundamental lack of understanding we still have about the nature of orebodies, as modeled through a quantification of epistemic uncertainty. Additionally, proper treatment of data at their respective scales, though proper geostatistical methods today are lacking and much needed, to reduce false positives.

This change will not happen overnight, as the industry and parts of academia remain set on socializing decisions on a single conceptual and 3D model. It will require reorganization of the industrial apparatus, new funding models, and a new form of education that focuses on problem solving in addition to the useful domain expertise focus.

Funding

This research was funded by members of the Stanford Mineral-X Industrial Affiliate Program.

Acknowledgments

I’d like to acknowledge members of the Mineral-X team, in particular, Zhuo Liu & Jonas Kloekner for providing material for this paper.

Conflict of interest

The authors declare no conflict of interest.

Use of AI and AI-assisted Technologies

No AI tools were utilized for this paper.

Open Access

This article is distributed under the terms and conditions of the Creative Commons Attribution (CC BY-NC-ND) license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

References

Barnes, S. J., Cruden, A. R., Arndt, N., et al. The mineral system approach applied to magmatic Ni–Cu–PGE sulphide deposits. *Ore Geology Reviews*, 2016, 76: 296–316.

Bond, C. E., Gibbs, A. D., Shipton, Z. K., et al. What do you think this is? “Conceptual uncertainty” in geoscience interpretation. *GSA Today*, 2007, 17(11): 4–11.

Caers, J. Bayesianism in the Geosciences, in *Handbook of Mathematical Geosciences*, edited by B. S. Daya Sagar, Q. Cheng and F. Agterberg, Springer Nature, pp. 527–566, 2018.

Caers, J., Journel, A. G. Stochastic Reservoir Simulation Using Neural Networks Trained on Outcrop Data. Presented at the SPE Annual Technical Conference and Exhibition, New Orleans, Louisiana, USA, 27–30 September, 1998.

Caers, J., Ma, X. Modeling Conditional Distributions of Facies from Seismic Using Neural Nets. *Mathematical Geology*, 2002, 34: 143–167.

Caers, J., Scheidt, C., Li, L. Decision-Making for Earth Resource Development: Quantifying Uncertainty, American Geophysical Union, 2026.

Caers, J., Scheidt, C., Yin, Z., et al. Efficacy of Information in Mineral Exploration Drilling. *Natural Resources Research*, 2022, 31: 1157–1173.

Caers, J., Yin, D. Z., Asadi, A., et al. Resource expansion with uncertainty quantification of regolith-hosted REE deposits using radiometric data. *Earth ArXiv* 2025, 10.31223/X5D18H. Available online: <https://eartharxiv.org/repository/view/11194/> (accessed on 17 December 2025).

Dempsey, H. Bill Gates-backed mining company discovers vast Zambian copper deposit. *Financial Times*, 5 February, 2024.

Eidsvik, J., Mukerji, T., Bhattacharjya, D. *Value of Information in the Earth Sciences: Integrating Spatial Modeling and Decision Analysis*. Cambridge University Press, 2015.

He, J., Tang, M., Hu, C., et al. Deep Reinforcement Learning for Generalizable Field Development Optimization. *SPE Journal*, 2022, 27: 226–245.

Howard, R. A. The Foundations of Decision Analysis. *IEEE Transactions on Systems Science and Cybernetics*, 1968, 4: 211–219.

JORC. *JORC Code Update*. 2024.

Journel, A. G., Huijbregts, C. J. *Mining Geostatistics*. The Blackburn press, 1978.

Kani, J. N., Elsheikh, A. H. Reduced-Order Modeling of Subsurface Multi-phase Flow Models Using Deep Residual Recurrent Neural Networks. *Transport in Porous Media*, 2018, 126: 713–741.

Kochenderfer, M. J. *Decision Making Under Uncertainty: Theory and Application*. MIT Press, 2015.

Laloy, E., Hérault, R., Jacques, D. et al. Training - Image Based Geostatistical Inversion Using a Spatial Generative Adversarial Neural Network. *Water Resources Research*, 2018, 54: 381–406.

Laloy, E., Hérault, R., Lee, J., et al. Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network. *Advances in Water Resources*, 2017, 110: 387–405.

Lawley, C. J. M., Tschorhart, V., Smith, J. W., et al. Prospective modelling of Canadian magmatic Ni (\pm Cu \pm Co \pm PGE) sulphide mineral systems. *Ore Geology Reviews*, 2021, 132: 103985.

Lewis, W., Vigh, D. Deep learning prior models from seismic images for full-waveform inversion. Presented at

the SEG Technical Program Expanded Abstracts 2017, Society of Exploration Geophysicists, Houston, Texas, USA, 24-29 September 2017.

Lindi, O. T., Aladejare, A. E., Ozoji, T. M., et al. Uncertainty Quantification in Mineral Resource Estimation. *Natural Resources Research*, 2024, 33: 2503–2526.

Liu, B., Yang, S., Ren, Y., et al. Deep-learning seismic full-waveform inversion for realistic structural models. *Geophysics*, 2021, 86: R31–R44.

Mern, J., Caers, J. The Intelligent Prospector v1.0: geoscientific model development and prediction by sequential data acquisition planning with application to mineral exploration. *Geoscientific Model Development*, 2023, 16: 289–313.

Nasir, Y., Durlofsky, L. J.. Deep reinforcement learning for optimal well control in subsurface systems with uncertain geology. *Journal of Computational Physics*, 2023, 477: 111945.

O'brian, S. Petrology of the Crystal Lake Gabbro and the Mount Mollie Dyke, Midcontinent Rift, Northwest Ontario. Thunder Bay, Lakehead University, 2018.

Ringrose, P., Bentley, M. *Reservoir Model Design*. Springer Nature, 2015.

Rossi, M. E., Deutsch, C. V. *Mineral Resource Estimation*. Springer Nature, 2014.

Russell, S. J., Norvig, P. *Artificial intelligence: A modern approach*. Pearson, 2016

Scheidt, C., Li, L., Caers, J. *Quantifying Uncertainty in Subsurface Systems, Geophysical Mono-graph Series*. John Wiley & Sons, 2018

Schodde, R. Recent trends and outlook for global exploration. Presented at the International Convention, Trade Show & Investors Exchange, Toronto, Canada, 9 February, 2017

Sutton, R. S., Barto, A. G., *Reinforcement Learning: An Introduction*. MIT Press, 1998.

Wang, Z., Yin, Z., Caers, J., et al. A Monte Carlo-based framework for risk-return analysis in mineral prospectivity mapping. *Geoscience Frontiers*, 2020, 11: 2297–2308.

Wei, X., Yin, Z., Bonner, W., et al. Knowledge-driven stochastic modeling of geological geometry features conditioned on drillholes and outcrop contacts. *Computers & Geosciences*, 2025, 196: 105779.

Yang, F., Zuo, R., Kreuzer, O. P. Artificial intelligence for mineral exploration: A review and perspectives on future directions from data science. *Earth-Science Reviews*, 2024, 258: 104941.