


Geologists Martin Tuchscherer [left] and Patrice Rioux stand in the middle of a boulder field in Nunavik, Que., Canada. Predictions from AI models can identify possible minable ore deposits in a vast area.



AI Hunts for Hidden Minerals

Machine learning is uncovering hoards of vital EV battery metals By JOSH GOLDMAN & KURT HOUSE

22 SPECTRUM.IEEE.ORG JUNE 2023




One way to find potential minable ore deposits is to use a transmitter coil loop [below], which—similar to a metal detector you'd use on a beach—detects induced currents in conductive bodies deep underground.

In June 2022, six Boeing 737s—fully loaded with tents, food, satellite Internet equipment, drones, geophysical survey gear, drilling equipment, and a team of experienced geologists—flew to a remote airstrip in northern Quebec. The geologists were hunting for major deposits of the minerals needed to power a clean-energy future. Given the mix of cutting-edge scientific computing and old-school bravado, it was as though they were channeling Alan Turing and Indiana Jones simultaneously. • Our startup, KoBold Metals, acquired an 800-square-kilometer mineral claim in this region of Canada based in part on predictions from our artificial intelligence systems. According to the AI, there was good reason to believe we'd find valuable deposits of nickel and cobalt buried below the surface. Summer snowmelts in this near-arctic area created a brief window to bring in a small village's worth of equipment and personnel to test our predictions.



In order to survey large areas quickly, KoBold uses a helicopter [right] carrying a transmitter coil loop 35 meters in diameter that can detect conductive bodies, such as ore deposits, below the surface. Conductivity data from KoBold's helicopter surveys can be used to produce models of potential underground distributions [far right]. The blue represents nonconductive igneous rock, while the yellow, orange, and red circles indicate areas of conductivity, from least to most conductive.



We cofounded KoBold in 2018 with backing from Bill Gates's Breakthrough Energy Ventures and Silicon Valley venture capital firm Andreessen Horowitz. Our goal is to develop ways to discover major new deposits of vital metals needed for electric vehicle (EV) batteries—for which there is an enormous and growing need.

We're trying to transform mineral exploration from a manual, judgment-guided, trial-and-error process into a data-driven and scalable science. It's the mother of all needle-in-a-haystack problems: Find the significant minable deposits of cobalt, copper, lithium, and nickel resting anywhere from 100 to 2,000 meters deep in the Earth's surface.

PREVENTING THE MOST catastrophic impacts of climate change requires achieving net-zero greenhouse gas emissions by 2050, which includes, among many other things, replacing all fossil-fuel-powered light cars and trucks with electric vehicles. That, in turn, will require manufacturing billions of EV batteries. Even today's demand for the metals outstrips supply—as evidenced by nickel prices doubling and lithium prices quintupling over the last year. To realize a global transition to electric vehicles, we'll need to discover and mine an additional US \$15 trillion worth of cobalt, copper, lithium, and nickel by midcentury. (We're currently on target to mine about \$3.6 trillion worth of these metals by 2050).

World leaders are well aware of the need. In the United States, for example, President Biden invoked the Cold War-era Defense Production Act in March 2022 to use the presidential


powers it grants to encourage domestic production of the minerals required in EV batteries. The Inflation Reduction Act, signed into law in August 2022, included billions of dollars to subsidize the development and operation of metals mines, both in the United States and globally.

Investors are aware of the supply challenge as well. In February 2022, KoBold raised \$192.5 million in Series B financing, which has gone toward securing more than 50 exploration sites in Australia, Canada, Greenland, sub-Saharan Africa, and the United States. We plan to use AI to streamline the largely scattershot process of discovering new ore deposits. Once they're discovered, we plan to partner with mining companies for the actual mining operations and advise them on efficient extraction, again using our AI tools.

FOR THOUSANDS OF YEARS, humans have noticed the striking appearances of rocks containing useful minerals. For example, the iron sulfides that are the predominant mineral in nickel sulfide deposits produce distinct reddish rust when exposed to air and rainwater. Weathering turns copper sulfide into a variety of brightly colored minerals, including the brilliant green ones found in the Statue of Liberty's patina. These visual clues were, for thousands of years, one of the most reliable ways to distinguish useful minerals and metals from useless rock.

The mining industry's rate of successful exploration—meaning the number of big deposit discoveries found per dollar invested—has been declining for decades. At KoBold, we sometimes talk about "Eroom's law of mining." As its reversed name

ALL PHOTOS: KOBOLD METALS



suggests, it's like the opposite of Moore's law. In accordance with Eroom's law of mining, the number of ore deposits discovered per dollar of capital invested has decreased by a factor of 8 over the last 30 years. (The original Eroom's law refers to a similar trend in the cost of new pharmaceutical discoveries.)

Geologically speaking, the decline in new discoveries is largely because most of the easy-to-spot deposits, such as those on the surface, have been found. New discoveries will be deeper underground, concealed by layers of rock.

In fact, the vast majority of Earth's ore deposits are still waiting to be found. The chemical and physical processes that form these ores occur at temperatures and pressures that exist kilometers below the surface. That is, these ore deposits are not formed at the surface; tectonic processes bring only a small minority of them there long after they were formed. That small minority constitutes the bulk of the deposits being mined today. The mining industry has the equipment and the technology to mine ore deposits that lie deep underground—the problem is finding those deposits in the first place.

You might expect the mining industry to be investing heavily in exploration, as well as in R&D to improve its exploration methods. But it's not. Over the past several decades, large companies have relied less on their own exploration programs and more on acquiring discoveries made by other companies. Mining-company shareholders expect dividends, not innovation.

AT **KOBOLD**, we're treating exploration as an information problem—finding and analyzing multiple types of data in order to uncover what we're looking for. In particular, it's an information problem in which acquiring more of those data types comes at a high cost. Our solution is to combine AI systems with geoscience expertise to figure out what piece of information reduces our uncertainty the most.

There is a vast body of geoscience information already in the public domain, but it's dispersed and fragmented. Some of it comes from government-funded geological surveys, and some comes from surveys conducted by private companies that were required to make their findings public. This information is spread across millions of data sets, including geological maps showing types of rocks observed in different locations; geochemical measurements of the concentrations of dozens of elements in samples of rock, soil, drill cores, plants, and groundwater; geophysical measurements of the gravitational

field, magnetic field, natural and induced electric currents, seismic waves, and radiation from the decay of heavy-element nuclei in Earth's crust; satellite imagery—in both visual and infrared bands—measuring the spectral reflectance of minerals at the Earth's surface; and text reports describing field observations. The volume of data is, in a word, overwhelming.

What's more, these data sets range from state-of-the-art mass spectrometry measurements to hundred-year-old maps hand-painted on linen. Each data set is useful, and, combined in the right way, the full collection is potent—if you can make sense of it.

Our data system, called TerraShed, parses this information and brings it into a standard form to make it accessible and searchable by both humans and algorithms. Curating the data and putting it through quality control are just the first steps. We then use various algorithms to guide our decisions about what data to collect at each stage in the exploration process, from getting a sense of whether a particular deposit is worth mining all the way to construction of the mine itself.

TerraShed doesn't produce simplistic treasure maps: It doesn't spit out an "X marks the spot" in response to the data. Instead, we have hundreds of different proprietary modules that guide each decision in the exploration process.

Our exploration program in northern Quebec provides a good case study. We began by using machine learning to predict where we were most likely to find nickel in concentrations significant enough to be worth mining. We train our models using any available data on a region's underlying physics and geology, and supplement the results with expert insights from our geologists. In Quebec, the models pointed us to land less than 20 km from currently operating mines.

After we acquired the relevant land rights, our geologists worked out of a field camp on-site, making observations and taking measurements of rock outcrops. Across the more than 800 km² of our claims, the choice of which rocks to sample is practically limitless. Time and money, however, are not—and in the region we were working, there's less than a three-month window when the ground is free of snow.

So, the information challenge becomes: How do we decide which rocks to sample?

We built Machine Prospector, which comprises the machine-learning models, with historic data, such as information from previous discoveries elsewhere in the province. It helped us predict which rocks we should sample, given the



Geologist Dave Freedman [left] stands behind a row of core samples at KoBold's Cape Smith site in Nunavik, Que., Canada, taken from areas of interest to confirm AI predictions. Brownish flecks in the core sample above, KSC 22-07, reveal the presence of magmatic pyrrhotite.

limited time we had. Specifically, we were looking for spots where eons-long geologic processes would have formed nickel- and cobalt-rich magmatic sulfide deposits.

Predictions in hand, our field geologists fanned out. Some headed to the places that seemed most likely to yield these magmatic sulfides. Others went to locations where the predictions were the most uncertain. Collecting data from places with uncertain predictions improves the next generation of models more than just collecting data where the models are already confident.

When the field team returned to camp in the evening, they uploaded that day's data via satellite. Our data scientists, working all over the globe, then retrained the models based on the new data points. The resulting new predictions changed the map of potential sample sites over the whole region and guided the team's decisions on where to go next. By incorporating new field information in almost real time, our model's adaptive predictions effectively shortened the learning cycle from a season to a day.

Our models generated predictions with 80 percent lower false positive and false negative rates compared to conventional predictions from geological maps. Such maps are constructed by making observations of the rocks at a relatively small number of locations and then using a set of rules and principles to extend those observations to larger regions. That means the conventional predictions are largely inference—and worse, they result

in unquantified uncertainty. In other words, we don't know what we don't know about how accurate those maps are. By comparison, KoBold's predictive models do quantify uncertainty, which in turn guides our data collection, as the most uncertain rocks often represent the most valuable ones to sample.

THE RESULTS FROM one of our staked claims during that 2022 field season in northern Quebec are a perfect example of how our unique approach to exploration pays off.

Guided by the results from our AI systems, our field team found a large boulder field that geologist Lucie Mathieu identified as very anomalous, and not typical of the kind of igneous rock making up most of the region's boulders.

The boulder field originally piqued our interest after electromagnetic measurements we had taken indicated unusually high conductivity—consistent with the kinds of minerals we were seeking. The electromagnetic data was gathered by a helicopter towing a 30.5-meter-diameter transmitter coil loop for a daily time-domain electromagnetic survey. For these surveys, the transmitter pulses current through the loop at 7.5 hertz, which induces currents in conductive materials underground. When the transmitter pulse ends, the receiver coil detects the decay of those induced subsurface currents, enabling us to build a three-dimensional model of the subsurface rocks' con-

ductivity. The high electrical conductivity of the ore minerals we're seeking is just one of several things that we can use to distinguish ore from other rock.

Using helicopter and geophysical survey equipment is expensive, and in the north the windows of good weather are short and unpredictable. Where we send the bird, and how we manage the trade-off between aerial coverage and spatial resolution, are vital considerations.

We can use the collected data to build three-dimensional models of the probable locations of ore in the subsurface, which

To do better, we quantify the uncertainty in our predictions about the subsurface. Our machine-learning models are trained on many fewer parameters than traditional best-estimate models, and the parameters are directly related to the key exploration questions: How many conductive bodies are present? How deep are they? What is their orientation? Is their conductivity in the range that's consistent with high concentrations of ore minerals? The output of our models is the joint probability distribution of these parameters.

Ultimately, the most useful data to collect is that which reduces the uncertainty of finding an ore deposit that can be mined. Together with our collaborators at Stanford University's Mineral-X initiative, we have developed a novel way of quantifying how useful an incremental piece of data is. We published the framework, which we call "efficacy of information," in *Natural Resources Research* in March 2022, and we used it to design our drilling program for our northern Quebec exploration and for our other expeditions.


Over the course of the summer in Quebec, we drilled 10 exploration holes, each more than a kilometer away from the last. Each drilling location was determined by combining the results from our predictive models with the expert judgment of our geologists. In each instance, the collected data indicated we'd find conductive bodies in the right geologic setting—possible minable ore deposits, in other words—below the surface. Ultimately, we hit nickel-sulfide mineralization in 8 of the 10 drill holes, which equates to easily 10 times better than the industry average for similarly isolated drill holes.

We were also pleased with how accurate and specific the predictions were. For instance, at hole KSC-22-004, our data scientists predicted a conductive body to be located somewhere between

130 and 170 meters below the surface. Upon drilling, we encountered highly conductive rocks at 146 meters.

That particular discovery was made just days before the end of the field season. The data helped define the subsurface geology so that our team will start the next season—which begins soon—by making the most effective drill holes to establish the shape and size of that ore deposit.

Assuming that ore deposit and others we've begun to identify in the area turn out to be as promising as we hope, we'll be well on our way toward another mine for one of the crucial metals needed to electrify the planet. Collectively, the world needs at least 1,000 new mines to be developed by midcentury to provide enough critical metals to produce enough EVs and avoid the worst consequences of climate change. That's a tall order. But by applying new AI systems like KoBold's, we may just be able to dig up new opportunities fast enough. ■



Audry Afango [far left] and Morgan McNeill work with a receiver on a ground electromagnetic loop in the middle of an angular boulder field in Nunavik, Que., Canada. On the right, McNeill, a geophysics technician, uses a superconducting quantum interference device (SQUID) to conduct a ground electromagnetic survey to pick up faint magnetic fields from conductive bodies underground.

is a computationally difficult problem. Put simply, we have a limited set of measurements of the induced fields taken in a two-dimensional plane just above the surface, and from that we are trying to infer the properties (here, the conductivity) of a three-dimensional volume of the subsurface. There are an infinite number of subsurface rock configurations that are consistent with the surface data.

The conventional approach in the industry is to build a best-estimate model that tries to fit a huge number of parameters, which can easily exceed the number of data points. Anyone who has tried to solve a system of n equations for $2n$ unknowns knows that there is no unique solution to the problem. The traditional methods used in the industry to choose one of the many potential solutions can often incorporate assumptions that are inconsistent with geologic processes and are prone to confirmation bias.