McKinsey & Company

Impact Story

Inside a mining company's AI transformation

February 2020

Inside a mining company's AI transformation

How copper-mining giant Freeport-McMoRan unlocked next-level performance with help from McKinsey data scientists and agile coaches.

Authors
Red Conger
Harry Robinson
Richard Sellschop

The mood was apprehensive as data scientists, metallurgists, and engineers from Freeport-McMoRan filed into the control room of a copper-ore concentrating mill in Bagdad, Arizona, on the morning of October 19, 2018. They had come to learn what would happen when they cranked the big mill up to a work rate that had never been tried.

The possibility of causing problems at the mill weighed on everyone's mind. The team members had initially resisted the idea of running the mill faster. They wanted to keep the stockpile of ore that feeds the mill from dropping below the minimum size they had long maintained. Their concern was that a too-small stockpile would hamper the mill's performance.

Whether the minimum stockpile size actually helped the mill run better was another matter. No one really knew for sure. Nor could the mill's managers and staff say what would happen if the stockpile shrank to less than the traditional minimum.

As the weeks went by, the copper-ore concentrating mill sustained a faster pace with no loss of efficiency. The data model had been right: the mill could handle more ore than its operators thought.

What they did know is that a custom-built artificial-intelligence (AI) model, loaded with three years' worth of operating data from the mill and programmed to look for operational tweaks that would boost output, kept saying copper production would rise if the mill were fed with more ore per minute.

To the mill operators, that notion sounded logical enough—except that it didn't account for the minimum stockpile size they had in mind. But the model didn't know, or care, about minimum stockpile size or any of the mill operators' other ideas about how the mill ought to be run.

With permission from company executives, the crew members at the Bagdad site decided to turn up the pace of the mill as the model had suggested. They also prepared to ramp up mining and crushing activities so the stockpile of ore wouldn't run out.

At ten o'clock in the morning, a technician clicked a control on his computer screen to speed up the system of conveyor belts carrying chunks of ore from the crusher to the stockpile and from the stockpile to the mill.

Everyone in the room kept watch on the 13 oversize monitors in the control room, which were lit up with readings from hundreds of performance sensors placed around the mill. The quantity of ore grinding through the mill rose. No warnings went up.

Twelve hours passed. The mill held steady. Even when its stockpile of ore dipped below the usual minimum, the accelerated delivery of ore from the crusher and the mine allowed the mill to keep going. As the weeks went by, the mill sustained the faster pace with no loss of efficiency. The data model had been right: the mill could handle more ore than its operators thought.

"That was the breakthrough we'd been looking for," Justin Cross, the Bagdad site's general manager, told us. "Once we started to run the mill at full speed, we knew we could get results from more of the recommendations that the model was making."

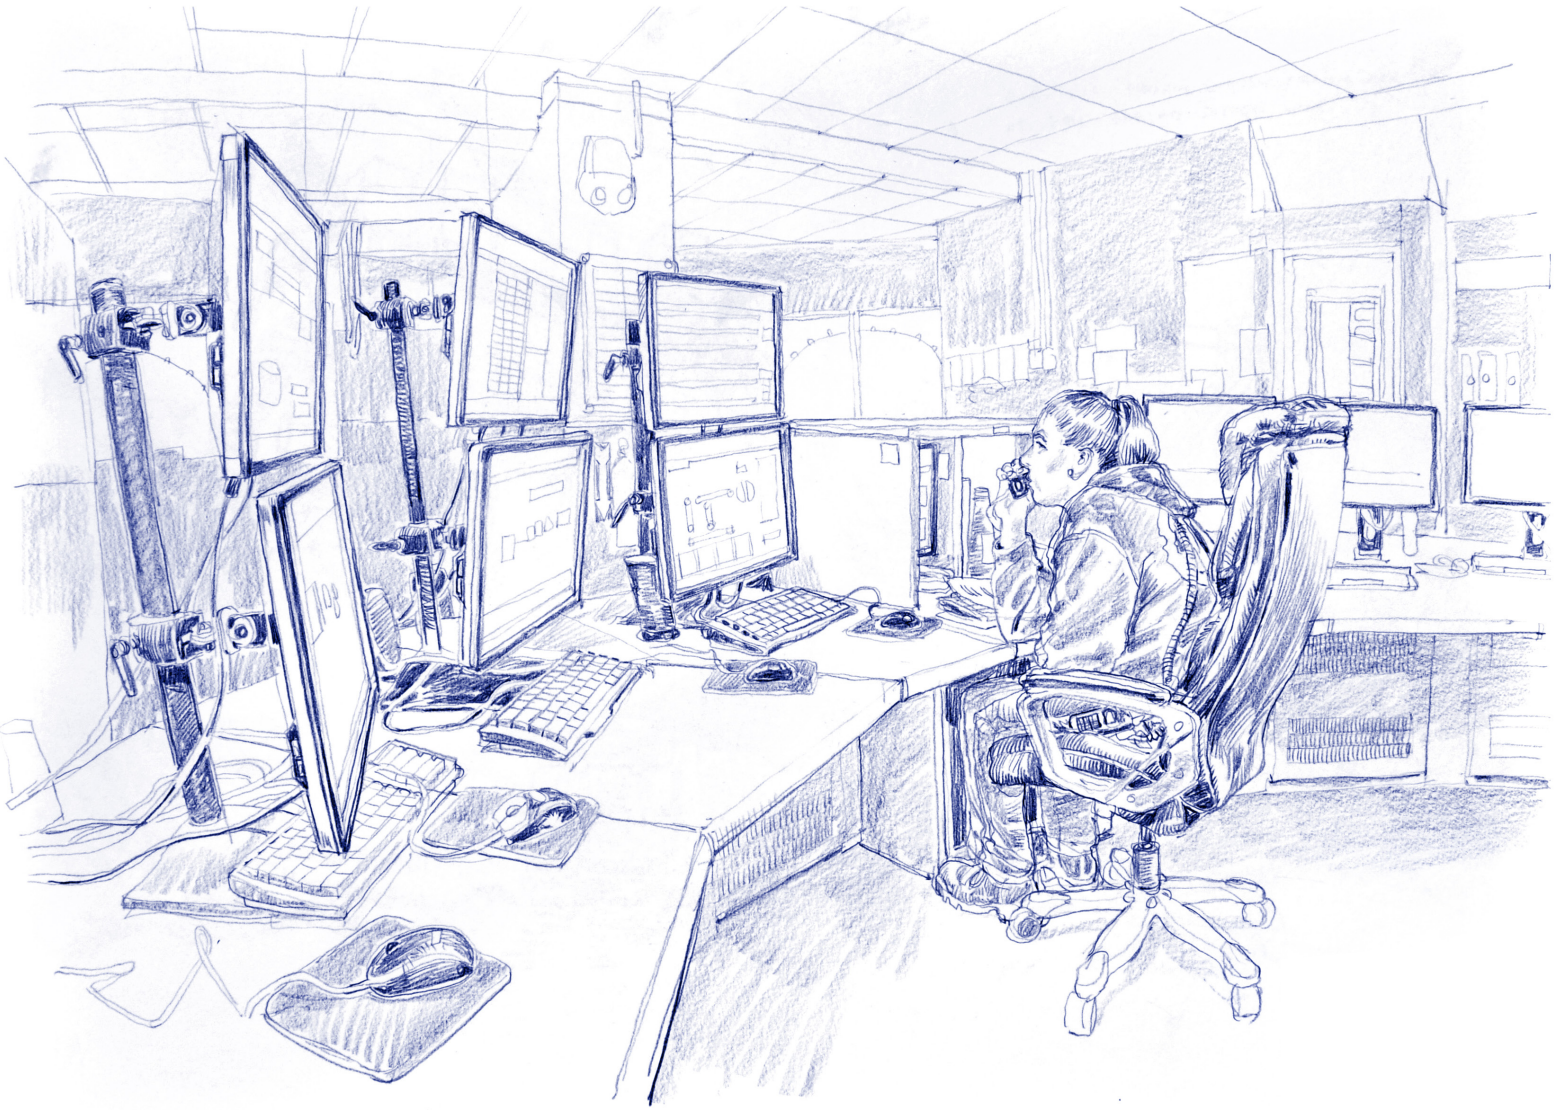

Grinding-mill operator Megan Alford monitors the bank of screens displaying details of the copper ore churning in the main mill just beyond the glass.

The ‘age of the operator’

The story of how Freeport-McMoRan learned to rely on an AI model as much as the intuition of veteran mining engineers and metallurgists might not raise eyebrows outside the tech industry.

For mining companies, though, it illustrates a quiet but profound shift into an era we think of as the “age of the operator,” when the best-run businesses wring profits out of low-grade ore that miners would have waved off as waste just ten years ago.

One mine where Freeport-McMoRan had been processing declining ore grades is Bagdad, a sprawling Arizona complex where prospectors staked their first claims in 1882. Bagdad’s reserves of higher-grade ore have been depleted for some time, but Freeport-McMoRan has sustained the mine’s production of copper by making various process improvements.

By the end of 2017, executives believed that Bagdad had gotten as efficient as it could get with its existing equipment, so they reasoned that adding capacity would be the surest way to get even more copper out of the site. Early in 2018, they started planning a $200 million capital expansion of Bagdad’s ore-concentrating mill that would lift production by 20 percent.

Copper prices were high at the time. The investment looked certain to pay off.

Then copper prices dropped from a five-year peak of around $3.30 per pound in early June to $2.75 or so a month later. All of a sudden, investing $200 million to expand Bagdad no longer seemed practical.

Instead, Freeport-McMoRan’s leaders resolved to find new process changes that would increase Bagdad’s copper output without a massive injection of capital.

Discovering improvements at an efficient mine wouldn’t be easy. But Freeport-McMoRan had plenty of high-quality information to study. Around ten years before, Bert Odinet, Freeport-McMoRan’s chief information officer, coordinated an effort to standardize the way that each site measures and reports its performance and to build a central data warehouse for storing those measurements.

Several years later, maintenance teams lobbied for the installation of additional network equipment and performance sensors on the company’s trucks, power shovels, and stationary machines. The teams would manually download data from those sensors to the data warehouse so they could further sharpen their maintenance practices and improve the functioning of equipment.

Freeport-McMoRan’s leaders resolved to find new process changes that would increase the site’s copper output without a massive injection of capital.

When wireless mesh networks became cost effective and reliable, Freeport-McMoRan installed them at all its sites. Now the company could capture and correlate second-by-second performance readings in the data warehouse, all in real time.

"We learned things we'd have never predicted," Odinet said. "That project taught us to be more receptive to what the data was telling us. And it gave us the confidence to try more complicated analyses."

With advanced analytics and AI techniques, Freeport-McMoRan could scan the vast quantity of data it collected, identify even more operational changes that might raise performance, and put them to the test in the field.

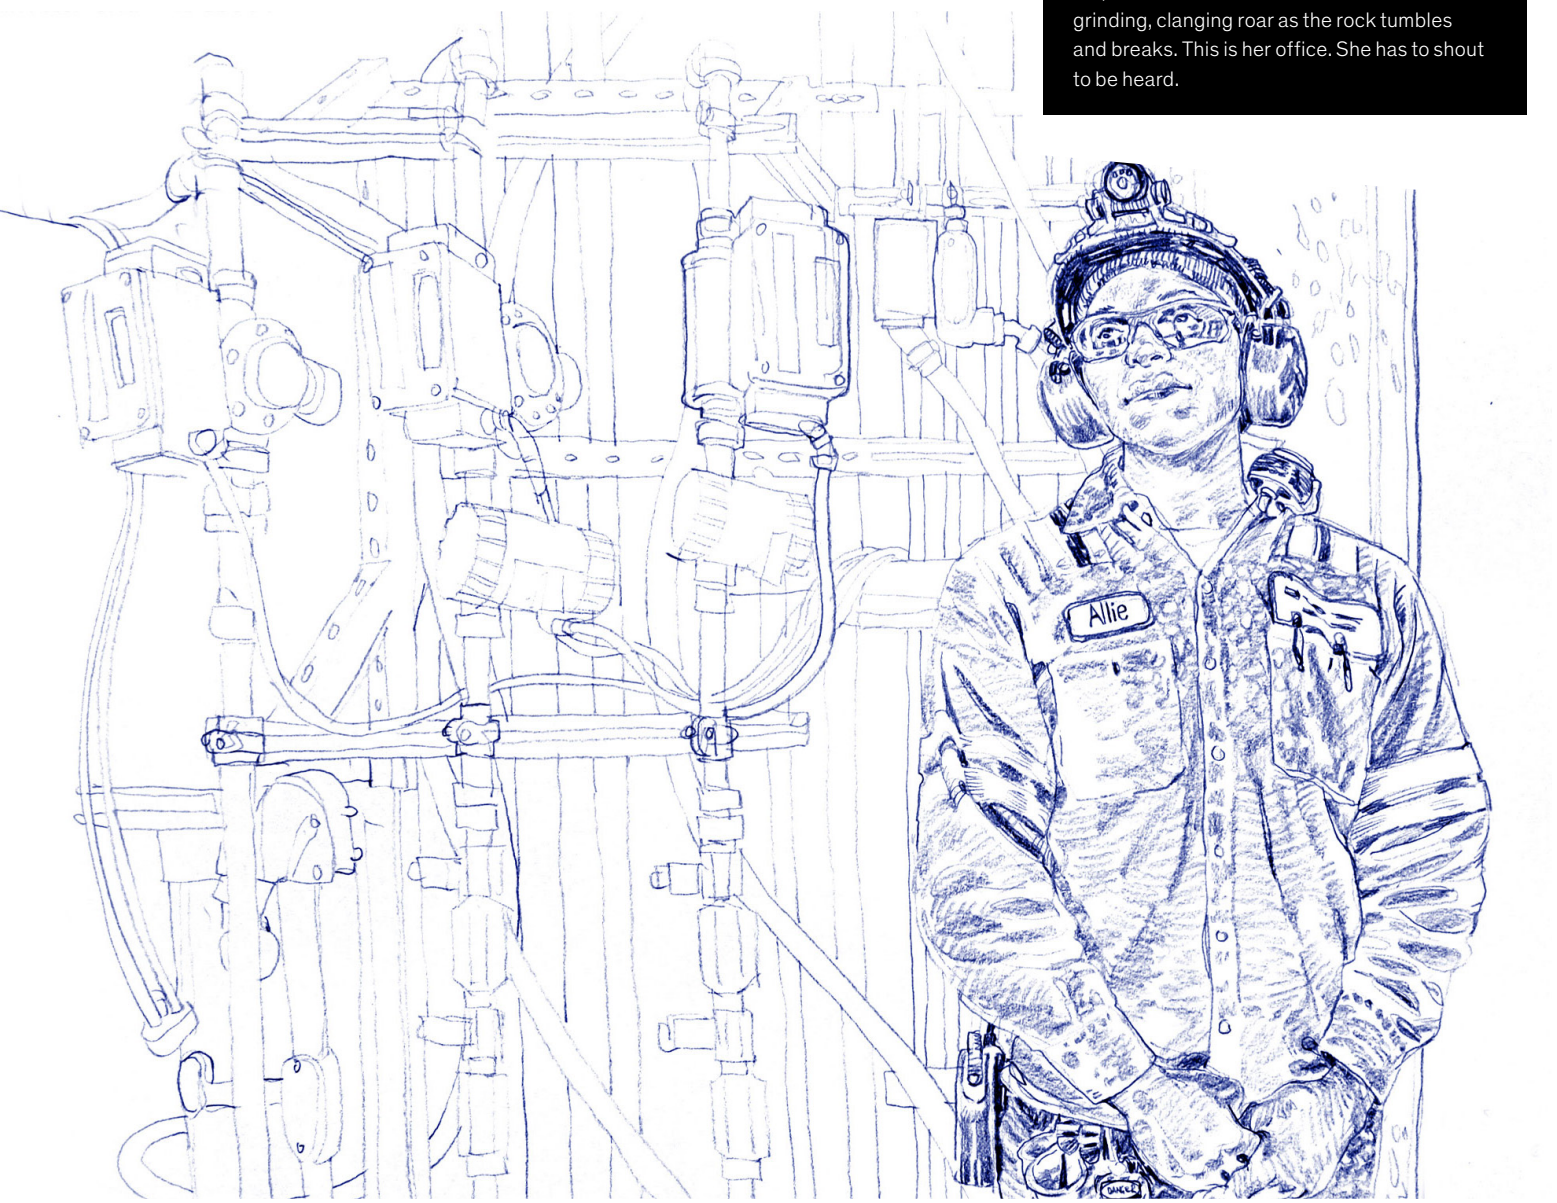

Allison (Allie) Naltazan is in training to become a mill operator on the ore-grinding line. Around her, the ore chutes and barrels create a constant grinding, clanging roar as the rock tumbles and breaks. This is her office. She has to shout to be heard.

Bagdad looked like a good proving ground for this method. The site is staffed with creative, open-minded engineers, metallurgists, and equipment operators who stood out at Freeport-McMoRan for their willingness to try new things. Their earlier efforts to enhance Bagdad's metallurgical processes, for example, had resulted in higher copper recovery.

And Cross, Bagdad's general manager, was a natural tinkerer who liked to spend his free time outfitting his pickup truck for off-road use and milling mesquite logs into lumber for a house he planned to build himself. Since joining Freeport-McMoRan in 2006, Cross had led a series of projects to streamline operations at the company's mines.

It also helped that Bagdad's operations were stable. With few equipment problems or process hiccups to straighten out, workers would have time to help increase copper production.

Executives felt that promising opportunities to boost Bagdad's production could be found in one part of the operation: the ore-concentrating mill, a noisy facility where huge milling machines and flotation cells bubbling with chemical solutions turn grapefruit-size rocks containing around 0.4 percent copper into a fine-ground mix of 25 percent copper and 75 percent rock.

The mill's technicians ran the facility strictly by the book, a set of operating instructions that Freeport-McMoRan engineers in Phoenix had developed. There had to be ways of building on those instructions that would better the mill's performance.

After talking things over with engineers and operations specialists from Bagdad and from headquarters, Freeport-McMoRan CEO Richard Adkerson and CFO Kathleen Quirk decided to let the crew at Bagdad work with McKinsey on a new kind of mining project: data mining in an agile way.

Learning agile

The project called for more than sophisticated data science. It also called for a new approach to solving tricky operational problems.

"Usually when we run operational projects, we overengineer them. We test every conceivable scenario, build in safeguards, and do everything we can to ensure that a process change will result in an improvement before we make it," Cross observed. "It's a dependable way to get good results. But it takes a huge amount of time, effort, and capital investment."

McKinsey's consultants reckoned that the crew at Bagdad might get better results, more quickly, by carrying out the analytics project differently. They introduced the idea of working under agile principles, which emphasize quick development of functional solutions that teams then improve, little by little, according to feedback from users.

The general manager assembled a team of people representing every division of the mill, along with other parts of the organization they would need to work with, such as the mine and Freeport-McMoRan's central data-science group.

Another essential feature of agile methods is face-to-face collaboration within well-rounded teams. Cross assembled a team of people representing every division of the mill, along with other parts of the organization they would need to work with, such as the mine and Freeport-McMoRan's central data-science group.

The composition of the team allowed it to tap the expertise and account for the interests of each division of Bagdad that its work might affect. It also enabled the team to contend better with challenges that involved different divisions and couldn't be solved by one division alone.

The team's agile approach was to work in "sprints"—two-week bouts of activity in which the team conceived a data-modeling function or operational change, tested it, and learned what would make it better.

As improvements came to light, the team would add them to a backlog. Then it would plow through the items on the backlog in subsequent sprints, starting with the easiest, most beneficial tasks.

For Bagdad's crew members, this agile style of working wasn't just different from business as usual. It represented a radical departure from the way they'd been doing things.

"It took us a while to get comfortable with agile," Cross said. "We had to let go of a lot of old habits."

McKinsey brought in agile coaches to help. The coaches explained the rudiments of agile—building a backlog, deciding what to accomplish in each sprint, holding morning meetings to agree on the work the crew would perform each day and to note any difficulties that might slow it down—but were mainly there to join the team's activities and teach its members to work together in agile ways.

Shannon Lijek, a McKinsey partner who specializes in helping organizations apply agile methods, was one of the coaches who came to help the Bagdad team get the hang of agile.

“We’ve found that the best way to learn agile is to jump right in,” said McKinsey partner Shannon Lijek.

“Agile can be tricky to adopt at first because it isn’t a process you can memorize. It’s a set of principles for minimizing wasted effort and getting more work done. And we’ve found that the best way to learn agile is to jump right in,” Lijek said.

One way that Bagdad’s agile team cut out needless effort was by introducing solutions as soon as it had built “minimum viable products,” or MVPs, that were good enough to use, rather than laboring to perfect those products first.

“If we’d built the model ourselves, we’d have tried to get it 100 percent right before doing anything with it,” Cross told us.

“Shannon and the McKinsey coaches encouraged us to work with solutions that weren’t finished. They’d say, ‘You can get 60 percent of the improvement with an MVP, and that’s a lot. So just start using it. Then you can worry about making it better.’”

From predicting to optimizing

Once the team Cross formed came together, it began investigating the possibility of improving the mill’s performance. The idea was to spend a month examining data from the mill for patterns that revealed potential improvements. If those improvements looked promising enough, the team would pursue them.

Beginning in late June, the Bagdad team and data scientists from McKinsey built a machine-learning model to check whether the mill truly ran as efficiently as people believed. The model, a type of extreme gradient-boosting model, consisted of an ensemble of thousands of decision trees that had been engineered to include a great deal of metallurgical knowledge.

The staff at Bagdad and Freeport-McMoRan’s central operations group believed all the ore entering the mill was of the same type. Consequently, they had defined a single “recipe” of lower and upper parameters for the mill’s 42 control settings: the mix of differently sized ore chunks being fed into the mill, the pH level in the flotation cells, and so on.

But when the agile team at Bagdad ran the data from the mill’s performance sensors through its model, the members of the team learned something new. From the mill’s perspective, the mine was actually producing seven distinct types of ore.

What’s more, the mill’s standard recipe for control settings didn’t match the properties of all those ore types. Ore containing more iron pyrite, for example, would yield more copper if the pH level in the flotation cells were set higher than the recipe prescribed.

"Thinking about ore clusters in terms of data from the mill's instruments, rather than classifications from traditional geology, was a major mind-set shift—and it opened up many new possibilities for improving performance," said Sean Buckley, a McKinsey partner who led the analytics work.

All told, the team's analysis suggested that adjusting the mill's controls to suit each of the seven ore types could increase copper production by 10 percent or more.

That prospect convinced Freeport-McMoRan's leaders to let the agile team at Bagdad build an AI model that would look at the ore coming into the mill and suggest control settings to heighten production of copper from that ore.

To determine just how much copper Bagdad could yield, staff decided to establish a new mandate—maximizing copper production at a reasonable cost, with little new capital investment.

Mill operator II Floyd (George) Mocaby and senior metallurgist Rahul Singh review the alarms on the gantry system.

Team members wrote algorithms to discern the connections among the ore type, the operational readings from the plant's sensors, the amount of ore running through the mill, and the amount of copper recovered. Next, they developed more algorithms to predict the plant's performance based on measurements from the sensors.

After several weeks of development sprints, the team had raised the accuracy of the model's performance predictions to 96 percent—high enough to know that the model was properly interpreting the data streaming in from the mill's sensors and relating it to the mill's control settings.

The team then turned its attention from predicting performance to improving it. Staff began by asking a simple question that no one had asked in some time: What measure of performance do we want to optimize?

For years, the team at Bagdad had oriented its decisions and activities toward particular targets for copper production and operating cost. That approach made a certain kind of sense. It meant that Bagdad consistently generated profits.

Now, to determine just how much copper Bagdad could yield, the team decided to establish a new mandate—*maximizing* copper production at a reasonable cost, with little new capital investment.

Sidebar

Introducing AI and agile to mining operations: Lessons from Bagdad

by Red Conger

Freeport-McMoRan's effort to increase copper production at Bagdad taught us a good deal about how to use agile methods and AI tools at our sites, where it can be difficult to alter accepted routines. Here are a few things we're keeping in mind as we expand the use of agile and AI to more of Freeport-McMoRan's operations:

- Don't wait for the "perfect" product or solution to begin using it. Once it's working well enough, implement it right away. Immediate action brings immediate results.
- Be willing to reconsider and discard long-standing assumptions and processes if you find better ways to do things. That means validating your new ideas through data analysis and fieldwork.
- Empower frontline teams to take risks. That's how testing and learning happens. Set clear boundaries on what teams can try. Make it clear they won't be blamed if their experiments come up short or incur extra costs.
- Use data science to catalyze decision making. Human judgment and intuition are hard to replace, but people can make better decisions when they're informed by analytical findings.
- Once you create value with agile and AI, spread the word about what you did and how you did it. Showcasing success will attract interest in these capabilities and motivate colleagues to adopt them.

Red Conger is president and COO, Americas, of Freeport-McMoRan.

Maximizing production could lessen performance in other areas. Nonetheless, executives agreed that if Bagdad could increase production as the model predicted, the short-term cost would be worth it.

Cross and Cory Stevens, Freeport-McMoRan's vice president of operational improvement, knew that maximizing production could lessen Bagdad's performance in other areas. The mill's recovery rate—the percentage of copper extracted from the ore—might drop. Or the whole operation could come to a halt for hours.

Stevens went to other executives to explain that Bagdad's experiment could be costly. The performance numbers they'd see for the next few months might be dismal, he warned.

Nonetheless, the executives agreed with Stevens that if Bagdad could achieve the 10 percent increase in production that the model predicted, the short-term cost would be worth it. They gave him the go-ahead to try maximizing production.

With that approval, Cross granted Bagdad's staff the latitude to make operational changes that deviated from standard procedures and could cause the mill to miss its performance targets. Worker safety and equipment integrity were the only areas where no compromises or experiments would be allowed. Any other changes were fair game.

A big breakthrough

Over a series of iterations during the next month or so, the team conceived, tested, and refined algorithms that would look at sensor-generated data and recommend control settings to maximize copper output. The new algorithms, known as genetic algorithms, used the principles of natural selection to "evolve" settings that would produce the most copper, given a particular type of ore.

Most challenging were the model's recommendations to depart from the operational recipe that the staff at Bagdad had been following for years. The agile team spent a lot of time debating what to do with those.

Metallurgist I Alaina Mallard, senior metallurgist Rahul Singh, and mill operator II Floyd (George) Mocaby discuss the expert system inside the ore crusher's control room.

By early September, the team had expanded the prediction model into an MVP of an optimization model, dubbed TROI, that was capable of issuing recommendations every 12 hours, once for each of the mill's two daily shifts.

When each new set of recommendations came out, the engineers, equipment operators, and metallurgists on the team would huddle and decide what do to with them.

TROI was a work in progress, so its earliest recommendations weren't entirely reliable. At every shift, metallurgists from Freeport-McMoRan and from McKinsey would study the model's recommendations and question whether they were credible. Then the metallurgists would take note of the problem recommendations so the agile team could look into them.

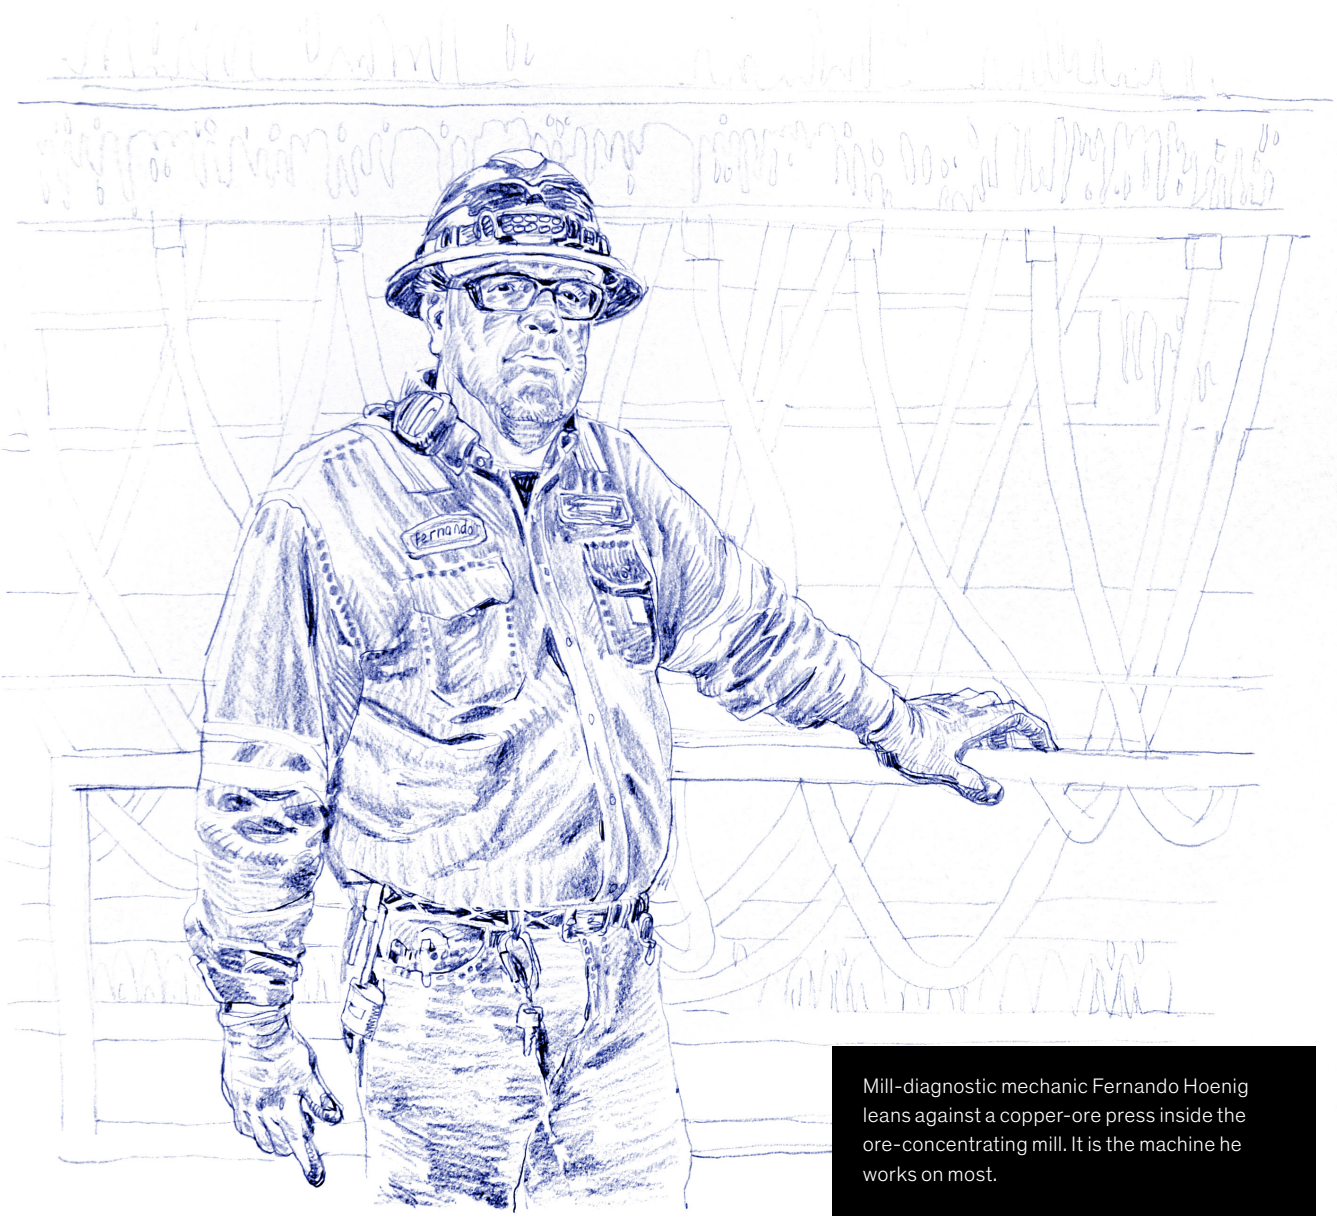

Mill-diagnostic mechanic Fernando Hoenig leans against a copper-ore press inside the ore-concentrating mill. It is the machine he works on most.

Some recommendations led the team members to discover flaws in TROI's logic, which they added to their backlog and corrected in subsequent development sprints. Others indicated that the underlying performance data were faulty and prompted the team to look for fixes.

"TROI helps us to improve the quality of our instrumentation and highlights sensors that need attention," said Frank Ochoa, one of Bagdad's process-control and instrumentation engineers.

Most challenging were the recommendations to depart from the operational recipe the staff had been following for years. The agile team spent a lot of time debating what to do with those.

Gradually, as the team fine-tuned TROI, its recommendations became more plausible, and the staff at Bagdad began following them. Yet many of those recommendations resulted in slim performance gains, if any.

Mid-October arrived. The team was nowhere close to the 10 percent production boost it thought possible.

Cross and Stevens decided it was time to act on a weighty recommendation that no one was especially eager to try: speeding the flow of ore from the mine and the crusher to the mill. Cross asked the mine operators to redline their activities—and reassured them that they wouldn't be blamed for spending more money or triggering operational breakdowns.

The mine operators ramped up blasting, even though they had to use more explosives. They queued up trucks to carry rocks to the crushing plant, in violation of a long-standing directive to keep trucks from standing idle. They choke-fed the giant crusher with run-of-mine, or unprocessed, ore to find out how much it could handle.

Finally, on October 19, the team pushed up the mill's processing rate. Right away, copper production jumped 5 percent. TROI had helped the team unlock a record level of performance.

Small gains add up

Having achieved a major performance gain, Bagdad's agile team turned to enhancing the model's ability to recommend mill-control settings that would increase copper production.

TROI could already identify which type of ore was running through the mill at any moment. In the next round of sprints, the team added functions to account for other incoming operational data.

Each time TROI recommended a set of control settings, the metallurgists at the plant would consider the recommendations, choose which ones to accept, and pass them to shift supervisors and operators, who would adjust the mill's controls accordingly.

Letting the metallurgists decide which recommendations to follow helped the agile team learn more quickly. Occasionally, the metallurgists applied settings that looked questionable just to find out whether they would work. And when the metallurgists rejected the recommended settings, they typed notes into the model to explain their decisions.

"TROI doesn't always give fully accurate recommendations, but it provides a new perspective on how to manage the plant and challenges our assumptions," said Lulu Raymond, a senior metallurgist at Bagdad.

As soon as the team pushed up the mill's processing rate, copper production jumped 5 percent. The model had helped the team unlock a record level of performance.

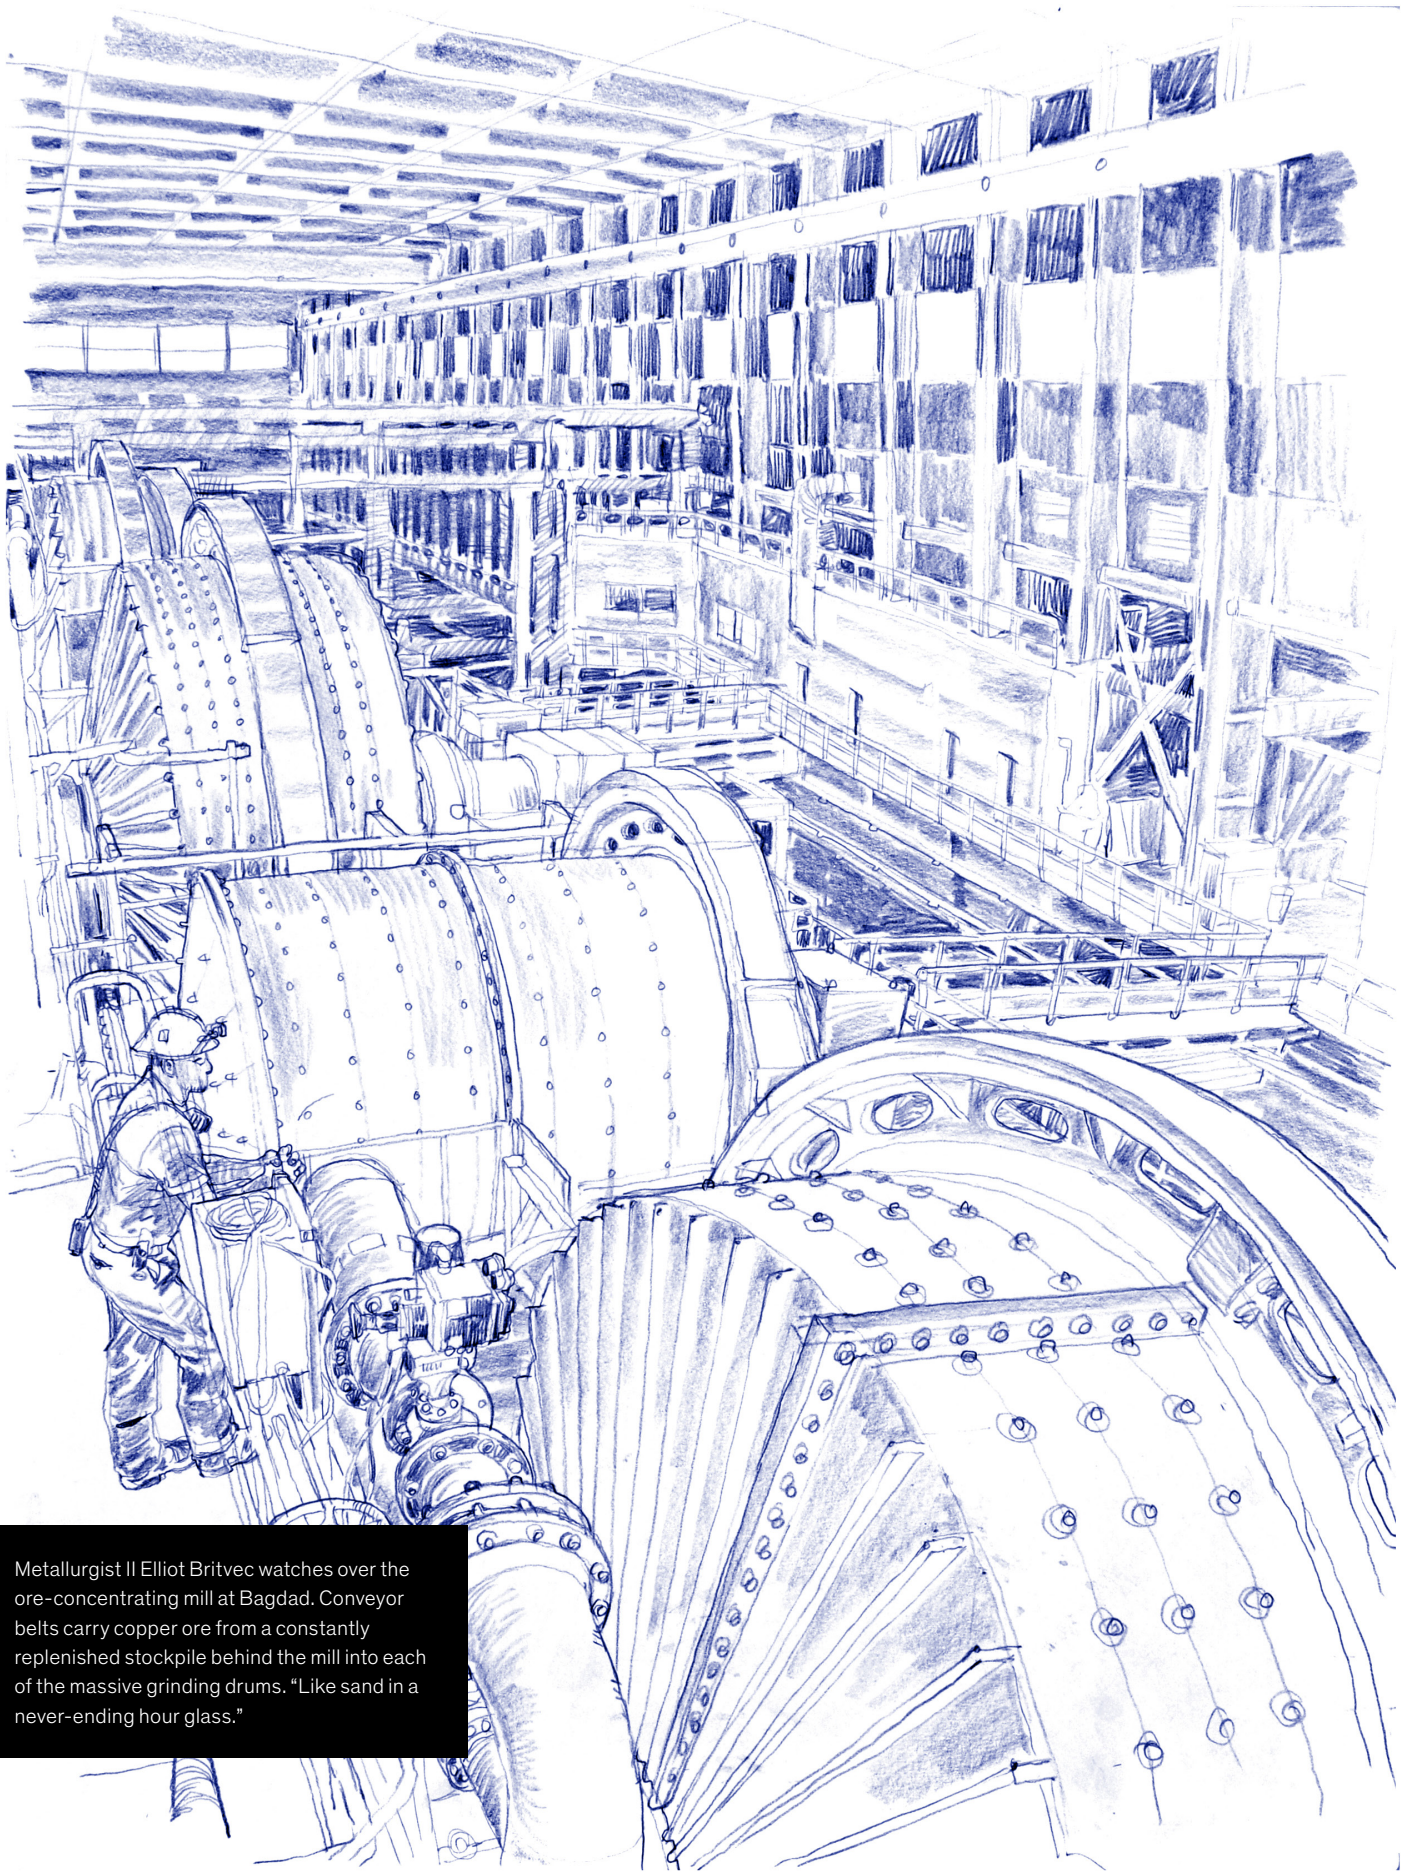

Metallurgist II Elliot Britvec watches over the ore-concentrating mill at Bagdad. Conveyor belts carry copper ore from a constantly replenished stockpile behind the mill into each of the massive grinding drums. “Like sand in a never-ending hour glass.”

The agile team reviewed the data model's recordings and the metallurgists' notes every day and kept working through a backlog of upgrades. Within several weeks, the team had refined the model to the point that metallurgists were accepting more than 80 percent of its recommendations.

All the while, sensors gauged the mill's performance. The model's machine-learning algorithms recorded which settings improved performance and which ones didn't, and whether the recommendations were helping.

The agile team reviewed the model's recordings and the metallurgists' notes every day, added items to the backlog of upgrades it planned to make, and kept working on those upgrades. By early December, the team had refined TROI to the point that metallurgists were accepting more than 80 percent of its recommendations.

It wasn't long before the metallurgists and mill operators began trying to outsmart TROI. They would monitor the type of ore passing into the mill, anticipate the control settings that the model might suggest, and apply those settings before the model made its twice-daily recommendations (later increased to every three hours). This became a kind of competition: Who can run the mill better than TROI would?

Most important, the mill's production increased substantially. In the fourth quarter of 2018, Bagdad's throughput exceeded 85,000 tons of ore per day—10 percent more than the previous quarter—while its copper-recovery rate rose by one percentage point and its operations became more stable (exhibit). The following quarter, copper production at Bagdad went up yet again.

Those gains should lift Bagdad's copper production by 20 million pounds per year, an increase that has allowed Freeport-McMoRan to avoid most of the $200 million capital expansion of the Bagdad concentrator complex.

Exhibit

By following recommendations from the artificial-intelligence model, a copper mill's operators increased throughput by 10 percent.

Average Bagdad mine dry ore milled daily, thousand tons

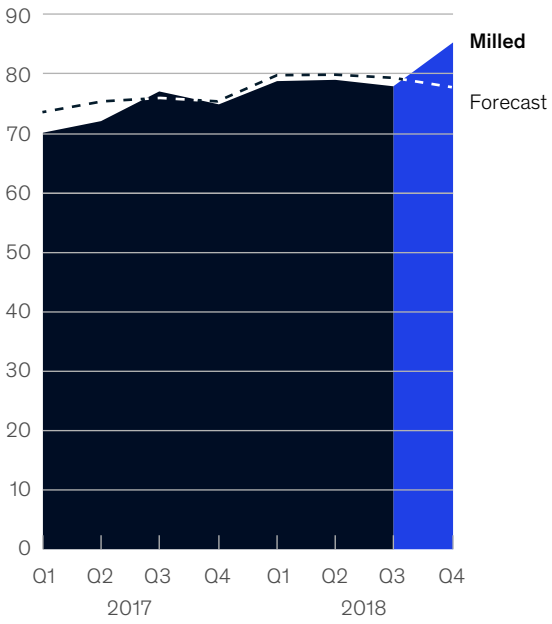

Source: Freeport-McMoRan

We think this is just the beginning for Freeport-McMoRan.

Having learned to maintain TROI during the project, the company's metallurgists and data scientists now run the model themselves, without ongoing support from McKinsey. They study daily and weekly reports that compare the mill's performance with TROI's predictions, and they continue enhancing the model's ability to make recommendations.

Freeport-McMoRan executives have also sponsored the creation of a second agile team at Bagdad to test and make process improvements at the mine. This team, too, is working without help from McKinsey, using the agile methods that it learned on the mill project.

At another one of Freeport-McMoRan's Arizona copper mines, Morenci, managers have kicked off an agile and analytics effort like Bagdad's. And the company will soon launch its most ambitious program of this kind at Cerro Verde, a copper mine in Peru with five times the capacity of Bagdad.

The age of the operator is here, and Freeport-McMoRan is adapting to it with agile methods and AI tools.

Red Conger is president and COO, Americas, of Freeport-McMoRan. **Harry Robinson** is a senior partner in McKinsey's Southern California office, and **Richard Sellschop** is a partner in the Stamford office.

Illustrations and photography by Richard Johnson.
Designed by Global Editorial Services