



Data rich but knowledge poor

By Munkhtsolmon Munkhchuluun and Davide Elmo

The mining industry generates data from millions of metres of drill core, grade and mineralization data, alongside terabytes of surface monitoring data, convergence measurements, maintenance records and inspection sheets. The risk is not about running out of data but about running out of people who know what the data means and why they were collected in the first place.

How much knowledge truly exists at any given time on a mine site? Where is that knowledge stored? And, more importantly, are there mechanisms that lead to a loss of knowledge, and, if so, how much of that knowledge can be recovered? These are not rhetorical questions. Roughly half of the industry's senior engineers are expected to retire within the next decade, taking with them the knowledge unique to each employee. That form of knowledge grows within individuals and disappears with them once they retire or move on. It is a combination of experience and physical observations that are not stored on any server. A recent upturn in undergraduate enrolment cannot offset what decades of attrition have already cost in institutional knowledge: new graduates bring technical training, but that alone cannot recreate the link between the *why* and the *what*.

Consider what happens when a newly hired geotechnical engineer walks into an underground mine for the first time. They may have read the reports and technical memos that underlie the ground control management plan, but the decisions recorded there were made without their input, and all they have are the parameters used to characterize the rock mass and to design the support. The raw data behind those decisions may exist on a server somewhere, but the narrative that gives them meaning rests with the engineers who collected and interpreted it.

That narrative, built from experience, can only be transferred slowly, often in fragments, across the limited medium of documentation and handover notes. However, the chain of reasoning that held those observations together over years of working underground does not survive in the record. When engineers leave, the chain breaks, and what remains is a series of numbers and files that describe what was measured but not what it meant.

No amount of computational power can resolve this problem on its own. Mining, like every other industry, is not immune to the transformation driven by artificial intelligence. But the industry's response has been to invest in artificial intelligence built on the same fragmented data that already fail to capture what engineers actually know. A model trained on that data does not learn from the mine; it learns from incomplete human records about the mine, inheriting their gaps and replicating subjectivity at scale. For example, ground falls account for roughly 40 per cent of all underground mining fatalities. The knowledge that might prevent them is not in any computer system, and there is no language model that can tell us what is happening at Stope 14B,

what has happened there over the last three years, and how the ground support is performing relative to the expected failure mechanism because the knowledge was never captured in a form that the technology could use.

Every mine is a repository not just of what was measured and how, but of why, with data and narrative dispersed along its drifts, stopes and ramps.

Bridging that gap requires a different conception of mining knowledge. Every mine is a repository not just of what was measured and how, but of *why*, with data and narrative dispersed along its drifts, stopes and ramps. But without proper intervention, that repository may look more like a labyrinth: one in which an engineer resolves immediate local challenges while losing sight of the broader question of how the mine behaves as a whole. What is needed is a better thread that connects every decision to the ground from which it came, every observation to the context that made it meaningful, and every new engineer to the understanding their predecessors built and could not leave behind. The thread must have a structure to which exploration data, development scans, production records, monitoring readings and closure documentation all attach, indexed by when and where they were collected. The thread is not a map, but a continuous connection, maintained throughout the mine's life.

With such a thread in place, the repository takes a specific form. It becomes spatially and temporally anchored: every observation, measurement and decision tied to a location—a drift designation, a stope number, a coordinate on the surface—and to a stage of mine life, so that exploration data, development assessments, production anomalies and rehabilitation records can be read in sequence and in context. It must carry authorship: the name and role of the engineer who made the observation, and what they believed it meant at the time. And it must record operational consequence: when an observation led to a design change, a support revision or a production hold, that connection must be captured, so the next engineer can trace not just what was measured, but what it triggered and why.

The data exists; the sensors are in place. What is missing is the thread and the willingness to treat it as the primary asset rather than an afterthought to the numbers. **CIM**

Munkhtsolmon Munkhchuluun is a mining engineer and co-founder of Ayavex Solutions. Davide Elmo is a professor of rock engineering at the Norman B. Keevil Institute of Mining Engineering, University of British Columbia.

Send comments to editor@cim.org