FOUNDATIONS OF **AI**

A FRAMEWORK FOR AI IN MINING

GMG

About GMG

The Global Mining Guidelines Group (GMG) is a network of representatives from mining companies, original equipment manufacturers (OEMs), original technology manufacturers (OTMs), research organizations, academia, regulatory agencies, consultancies, and industry associations who collaborate to tackle the challenges facing our industry. We aim to accelerate the improvement of mining performance, safety and sustainability by creating guidelines and white papers that address common industry challenges, facilitating collaboration and expanding the industry's knowledge base, and hosting and supporting events that bring mining stakeholders together along with external industries to address the industry's challenges, successes and innovations.

Layout by Clò Communications Inc. (www.clocommunications.com).

About the Project

This white paper was developed collaboratively by participants in the GMG Foundations for AI in Mining sub-committee of the AI in Mining Working Group to articulate a clear and unified understanding and framework of what AI is and how it can benefit the mining industry. Content development was done through online collaboration and over a series of conference calls and workshops.

Project Leader

Rob Johnston, Project Manager, CITIC Pacific Mining

Primary Contributors

Salil Ahuja, Technical Sales Leader, IBM

Michelle Levesque, Engineering Technical Lead, Natural Resources Canada / Government of Canada

Craig Lindley, Senior Principal Research Scientist, CSIRO

Todd McDonald, Senior Programmer / Analyst, Natural Resources Canada / Government of Canada

Mark O'Brien, Manager – Digital Transformation, CITIC Pacific Mining

Javier Orellana, Decision Automation Strategy and Innovation, BHP

Jess Robertson, Data Science Lead, Unearthed Solutions

Charlotte Sennersten, Research Scientist, CSIRO

Amanda Truscott, Co-Founder and CEO, Rithmik Solutions

Kevin Urbanski, Co-Founder and CTO, Rithmik Solutions

Farzi Yusufali, Co-Founder, Stratum AI

Organizations Involved in the Project

Agile GeoScience · Alex Atkins & Associates · Ambuja Cements · Andrés Bello National University · Antofagasta Minerals · Augmented Analytics · Barrick Gold · Bechtel · BHP · Boliden · C&O Servicios profesionales · Canary Systems · Canvass Analytics · Caterpillar · CITIC Pacific Mining · COREM · CSIRO · Curtin University · Datinfact · Deloitte · Department of Mines, Industry Regulation and Safety (DMIRS) · Digital Mine from GE Transportation, a Wabtec Company · École Polytechnique · Emergency Response Consulting · Enaex · ERAMET-SLN · ESG Solutions · Exergia · Flanders Electric · Flinders University · Fourdplan · Fundación Chile · Geological Survey of South Australia · Geomodelr · Global Insertion · Global Research · Hatch · Haultrax · Hitachi · Hudbay Minerals · IBM · Infosys · ISCTE - Lisbon University Institute · JVA · Larrain & Asociados · Les Entreprises Bektex · Liebherr · LKAB · Luleå University of Technology · Ma'aden · McGill University · Minetec · National Research Council Canada · Natural Resources Canada / Government of Canada · Newcrest · Newmont Goldcorp · Newtrax · Nutrien · Odgers Berndtson · Open World · Optika Solutions · Option · Organum Technology · OZ Minerals · Paradyn Systems · PETRA Data Science · Phoenix Contact · Resolution Systems · Rio Tinto · Rithmik Solutions · Rock Tech Centre · Rockwell Automation · Roy Hill · SafeAI · Santoencanto · Schneider Electric · Seed AI · Seequent · Seyo · Siemens · Sinese · Spectris Advance Mining · SRK Consulting · Stantec · Stratalis Group Consulting · Stratum AI · Symbiotic Innovations · Symboticware · Teck · TransformAction · TriStar Gold · Unearthed Solutions · University of Adelaide · University of South Africa · University of South Australia · Vale · Vedanta Resources · Wood · Yamana Gold · Zettascale Data Consulting

EXECUTIVE SUMMARY

The mining industry is increasingly using artificial intelligence (AI) as a tool to optimize processes, enhance decision-making, derive value from data, and improve safety. In this white paper, we provide a foundation for mining companies that are planning and implementing AI solutions.

Understanding what AI is and is not is an essential part of this foundation. AI, defined as the collection of techniques that allow for task automation by machines, fits into a bigger picture that includes machine learning, deep learning, data science, and big data. Although AI aims to mimic human intelligence, we emphasize that people are still central to any AI system and that human knowledge is required to interpret both the inputs and outputs of this technology. Nevertheless, implementing AI will bring about change that needs to be carefully and transparently managed.

There are several success factors associated with implementing AI. We identify some key success factors as being a coherent technology strategy, good data management, effective communication, clear expectations, internal support, agility and adaptability, consideration for the needs of end users, and long-term plans for scalability and repeatability.

The transition to an AI-enabled mine will look different for every organization, so we propose the following levels of maturity to help guide this process:

Level 1	**BASIC ▶ Exploration of what AI is, what it does, and what the benefits are.**
Level 2	**FOUNDATION ▶ Laying the foundations of the AI strategy, often through experimentation and investigation.**
Level 3	**INTEGRATED ▶ Integrating AI into business operations.**
Level 4	**DECISION SUPPORTED ▶ Leveraging analytical tools to provide centralized decision-making capabilities.**
Level 5	**AUTOMATED ▶ AI is at the core of the organization and most systems and processes are either fully automated or require minimal manual intervention.**

Introducing AI at any of the levels outlined above requires detailed planning and well-structured implementation. Choosing the right AI project involves brainstorming the possibilities, assessing and validating the potential projects from business and technical perspectives, and establishing a workflow for execution. To apply AI to a business problem, the project should be structured through well-defined processes. We suggest a general workflow that involves mapping and assessing the problem, analyzing the data, considering the data organization and collection systems, implementing an AI pilot project, and validating the findings.

While AI has many potential benefits, there are also many challenges with integrating it into an existing mining operation. These can be overcome through a robust foundation of planning, research, and assessment, and by establishing well-defined infrastructures and platforms, clear communication practices, and effective change management.

CONTENTS

INTRODUCTION

The mining industry is increasingly using AI-based innovation as a tool to optimize processes, enhance decision-making, derive value from data, and improve safety. That said, levels of maturity vary throughout the industry and there is still confusion about what AI is and how it can be applied to mining. As a result, mining operations still face many challenges with implementing AI, such as establishing a data infrastructure. Many mining stakeholders also have concerns about how AI will affect the workforce; they also worry about the risk of committing to a multi-year project and failing at it.

AI Adoption in Mining

On technology adoption by industry, the World Economic Forum (2018, Table 5) Future of Jobs Survey found the following among their mining and metals industry respondents:

Machine learning	69%
User and entity big data analytics	62%
Augmented and virtual reality	62%
Autonomous transport	50%

This white paper covers the basics needed to cut through the hype, address concerns, and clarify what methods are useful. It provides a realistic approach to moving forward with AI projects; it also acknowledges the importance of implementing an iterative process that focuses on quick wins and learning from failures and challenges in order to build the foundation needed for long-term success. This tool aims to help mining companies build the necessary foundation for planning for and implementing AI solutions within their businesses. To this end, it:

- Identifies background information on AI and its relevance to mining (Section 1)
- Describes a maturity model that will enable companies to plan their AI strategies (Section 2)
- Explores the high-level steps for implementing AI (Section 3)
- Describes some ways in which AI is being applied in mining (Section 4)

The primary audience is the members of mining companies who are tasked with introducing or further enabling AI in their organizations. The secondary audience is the ecosystem around mining companies that will help them implement the relevant technologies, culture, safety, and regulatory frameworks that a successful AI strategy requires. These audiences are familiar with AI, so this document does not aim to teach AI techniques; however, the Getting Started section (Section 1.1) identifies some useful resources that will help those who are new to the field learn the basics.

1

BACKGROUND

Before implementing AI technologies and transitioning to an AI-enabled mine, it is important to have a solid understanding of what AI is, what it is not, and some key success factors.

1.1 GETTING STARTED IN AI

The hype that surrounds AI often blurs the line between reality and myth. Understanding the basics of the technology is one way to cut through the hype. Much of the theory underpinning AI, machine learning, and data science is not new; these disciplines owe much to existing advanced mathematical and statistical techniques. AI leverages large quantities of data and computing power to apply statistics and mathematical modelling at an advanced level.

Recent technological advances and broad funding to generate further advances have contributed to the wider implementation of AI. For example, the advent of "elastic" cloud-based computing and storage and the appropriation of graphical processing units (GPUs) for building and training complex models have further revolutionized this field. Researchers and practitioners are now making new advances and discoveries at an incredible rate.

Some useful resources for those new to the field and experienced practitioners alike include:

Starting point: Andrew Ng's "AI for Everyone" course offered on Coursera (2019a). Ng has both academic and industry experience and is also the instructor of the Coursera course on "Machine Learning" (2019b) and the author of the AI Transformation Playbook (Ng, 2018).

Primer. Pedro Domingos's *The Master Algorithm* (2015) presents five main machine learning techniques in an easy-to-follow primer. The book treats the subject with respect while keeping it accessible.

A philosophical look: Nick Bostrom's *Superintelligence* (2014) investigates, at a high level, some dangers of AI systems.

1.2 THE BIG PICTURE

The terminology associated with AI is often a challenge because several frequently used terms overlap and there is little agreement on their exact definitions. This section defines these key terms at a high level. Figure 1 presents a way of visualizing how AI-related concepts intersect.

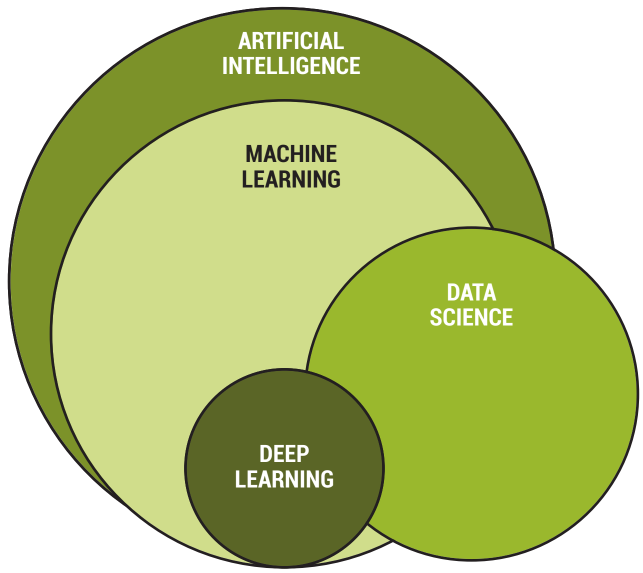

Figure 1. A Visualization of how AI-Related Concepts Intersect

Artificial intelligence: In industry applications, AI refers to a collection of techniques that allow for task automation by machines. These tasks are ones that humans would typically perform, and their automation implies that machines mimic certain aspects of human intelligence. A second definition of AI, artificial general intelligence, refers to a theoretical machine that has general human cognitive abilities.

Machine learning: A subfield of AI that focuses on machines that take data related to a specific task and learn from that data in order to build a model. This process enables the generation of an intelligent response to new data presented to the system through the use of algorithms and statistics models that identify patterns and make inferences.

Deep learning: A technique within machine learning that consists of algorithms that can take vast quantities of data and recognize patterns. For example, feeding a deep learning system thousands of pictures of different types of materials in order to build a material recognition system that could be mounted on an excavator. Goodfellow et al. (2016) describe deep learning as "representations that are expressed in terms of other, simpler representations (p. 5)." It is this concept of extending to build new capability that makes deep learning so exciting.

Output of AI, machine learning, and deep learning projects:
software that automates a certain task.

Data science: Analyzing data and extracting insights from them. Data science can be considered a separate activity to AI and machine learning, albeit one that overlaps them. AI can be done without the use of data science, but machine learning should not be conducted as such because algorithms should produce business intelligence and insights. Data science also provides tools that can be used to detect biases in the data used to train algorithms.

Output of a data science project: Likely a presentation to management and a list of recommendations. For example, analyzing rock fragmentation after a blast and recommending different blast criteria.

Big data: A term that recognizes the challenges inherent in analyzing large quantities of data. There are numerous definitions for big data but they all revolve around size. Big data can be defined by using the "Four V's": **volume**, **velocity**, **variety** and **veracity**.

1. The sheer **volume** of data that countries, organizations, and even individuals are producing is staggering. About 2.3 trillion gigabytes of data are created daily (IBM, n. d.). Some problems are impossible to solve on a single machine and therefore require clusters of computers and more complicated techniques to manage them.
2. The speed, or **velocity**, of data being generated creates challenges in how to effectively process vast quantities of real-time data. For example, a process plant may have tens of thousands of sensors that potentially send out many values per second.
3. The **variety** of data we now collect poses unique problems. Not only are organizations beginning to store structured data—for example, engine temperatures or alarm statuses—but they are also storing unstructured data, such as video documents.
4. Ultimately, if the data we're using is unreliable then there can be a high cost and so the **veracity** of the data is becoming increasingly important. For example, poor data quality costs the US government $3.1 trillion dollars a year (IBM, n.d.).

Useful Resource

IBM has created a useful infographic that outlines the **Four Vs of Big Data** and the statistics associated with them (IBM, n.d)

1.3 THE ROLE OF PEOPLE IN AI

Humans are central to any AI system because human knowledge is required to interpret the inputs and outputs. Inputs for AI models are built through carefully analyzing the external world, which can only be achieved with human involvement. If experts in a particular project's domain are not consulted, then valuable knowledge could be missed, and if too few people are consulted, then this could increase the risk of introducing bias.

For example, to build a predictive maintenance algorithm, reliability engineers will be consulted to help determine which data are important for identifying certain categories of faults. Maintenance work orders may also need to be inspected and interpreted to determine which breakdown events can be attributed to which maintenance activities. This human knowledge-gathering exercise is necessary for making sense of the complexities associated with all of the processes and systems before generating AI models for them.

Once a model has been built, it will require humans to interact with the output and interpret the results. For example, if the AI model produces a graph that shows the remaining useful life of a component, then an engineer would need to look at the output and decide if it means a part needs to be replaced.

When assessing the system output, it is also important to consider human–computer interaction concerns. This subfield of computer science relates to the effective interactions between a computer system and humans. The ability to develop a powerful interface or visualization may influence project success.

Implementing AI will also bring about many changes for people, so there should be strong organizational change management to ease the transition (see Section 3.5). Job loss, for example, is a key concern among the workforce. It is essential to be transparent about how some jobs may be replaced by AI and how many others will be created. Clear communication on the benefits is also crucial; most current mining-related AI projects provide tools to make people's jobs easier, safer, and more pleasant.

People That Might Need to be Involved in an AI Project

- Cloud specialists (for cloud solutions)
- Server teams (if the servers are on the premises)
- IT teams
- Support teams
- Data scientists
- Developers
- Domain experts
- Business stakeholders

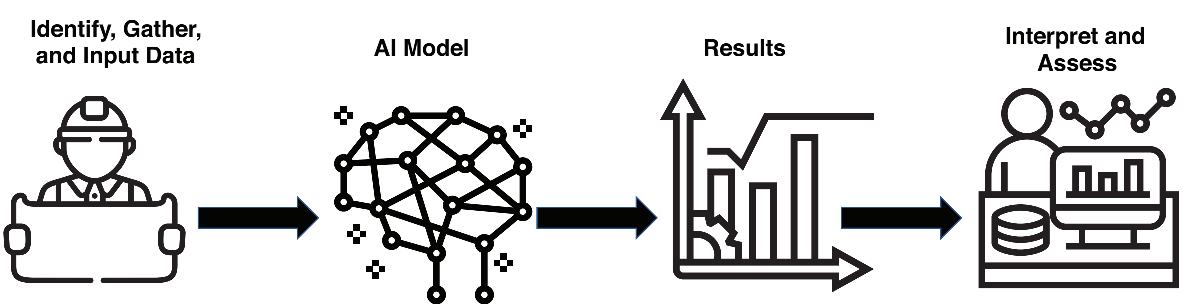

Figure 2. The Role of People in AI Input and Output Interfaces

1.4 SUCCESS FACTORS

Successful AI projects share many traits:

1. Successful AI projects have a coherent technology strategy.
Successful AI projects are undertaken to accomplish a goal for which using AI is the most efficient and effective path. When it comes to achieving long-term goals, AI needs to be part of the organization's larger strategy for evaluating and implementing new technologies.

2. Successful AI projects measure and monitor the quality of their data.
Good data are the foundation of AI because poor data quality reduces the quality of project output. AI can also expose deficiencies in the data as the project progresses, so measuring and monitoring data and being prepared with plans for filling gaps is key to success. It is also essential to measure and monitor data quality to ensure that the data are as complete as possible. Domain experts should be heavily involved in the entire process because they help ensure that the right data are collected, that the analyses generated are useful, and that the root causes of any problems are addressed.

3. Successful AI projects regularly evaluate the quality of communications among stakeholders.
A successful AI project depends upon clear communication. From the very beginning, those affected by the change should be engaged in discussions surrounding it, and they should feel confident that their opinions and concerns are being heard, understood, and considered. Unclear communications throughout the process can result in significant rework and make it less likely that those affected by the technology will understand its value. Everyone involved in the project should have a clear stake in it, and everyone with a stake in the project should be involved in some way.

4. Successful AI projects have a return on investment agreement.
Before a project begins, all of the participants should have a clear understanding of its intended outcomes and the expected return on investment. This understanding should be captured in writing, although it does not always need to be a binding legal document if the project scope is relatively small. If the expected project outcome delivers non-financial value, then the expected value should still be recorded and stored where stakeholders can refer to it as needed. If the expectations for a project and its stakeholders are not clear at the outset, then the project will be vulnerable to miscommunications, budget overruns, and missed deadlines.

5. Successful AI projects have at least one internal champion.
This internal champion is responsible for ensuring that all relevant stakeholders are updated regularly and those developing and implementing the technology have the tools and resources they need to succeed. The project's champion should be a good leader with the ability to inspire others and promote collaboration.

6. Successful AI projects are agile.

Things can change quickly in the mining industry, especially with increased digitalization, so bugetary and organizational changes often affect huge, pre-planned projects that have long timelines. It is also difficult to speculate on the limitations and challenges to be encountered with AI projects because they are complex and involve doing things that have not been done before. Agile project management is an approach that involves executing smaller parts of a project in rapid cycles while continuing toward a longer-term vision. This approach is a way to deliver the highest-value aspects of a project as quickly as possible and to ensure that at least some parts of the project do in fact come to fruition. Agile principles include:

- Prioritizing people over processes and tools
- Producing working prototypes over excessive documentation
- Behaving responsively and collaboratively

7. Successful AI projects focus on solving problems for the end users.

Successful projects focus on solving specific problems and easing the transition for the end user. Whether target users are technicians taking notes on equipment or engineers making operational decisions, dashboards, key performance indicators (KPIs) and visualizations should be easy to use and understand. As a result, their use is more likely to become habitual, making daily decision-making more impactful. Solutions that are habitually easy to use can also encourage users to be more open to adopting other new technologies and to recognize new applications for existing technologies.

8. Successful AI projects include long-term plans.

Even when starting small, having a repeatable and scalable long-term plan can lead to success in implementing the new technology company-wide and ingraining it in the organization's policies, procedures, tools, and habits. These plans should include provisions for any of the following: support, maintenance, change management, retraining and tuning algorithms, scaling, installation and setup, and acquiring or implementing new hardware.

2

THE TRANSITION TO THE AI-ENABLED MINE

The transition to an AI-enabled mine will look different for every organization. An AI maturity model allows each organization to put its AI transition into perspective and to visualize the transformation process it has in front of them.

This section outlines five levels of AI maturity.

2.1 MATURITY MODEL

Figure 3 outlines an AI Maturity Model for the mining industry.

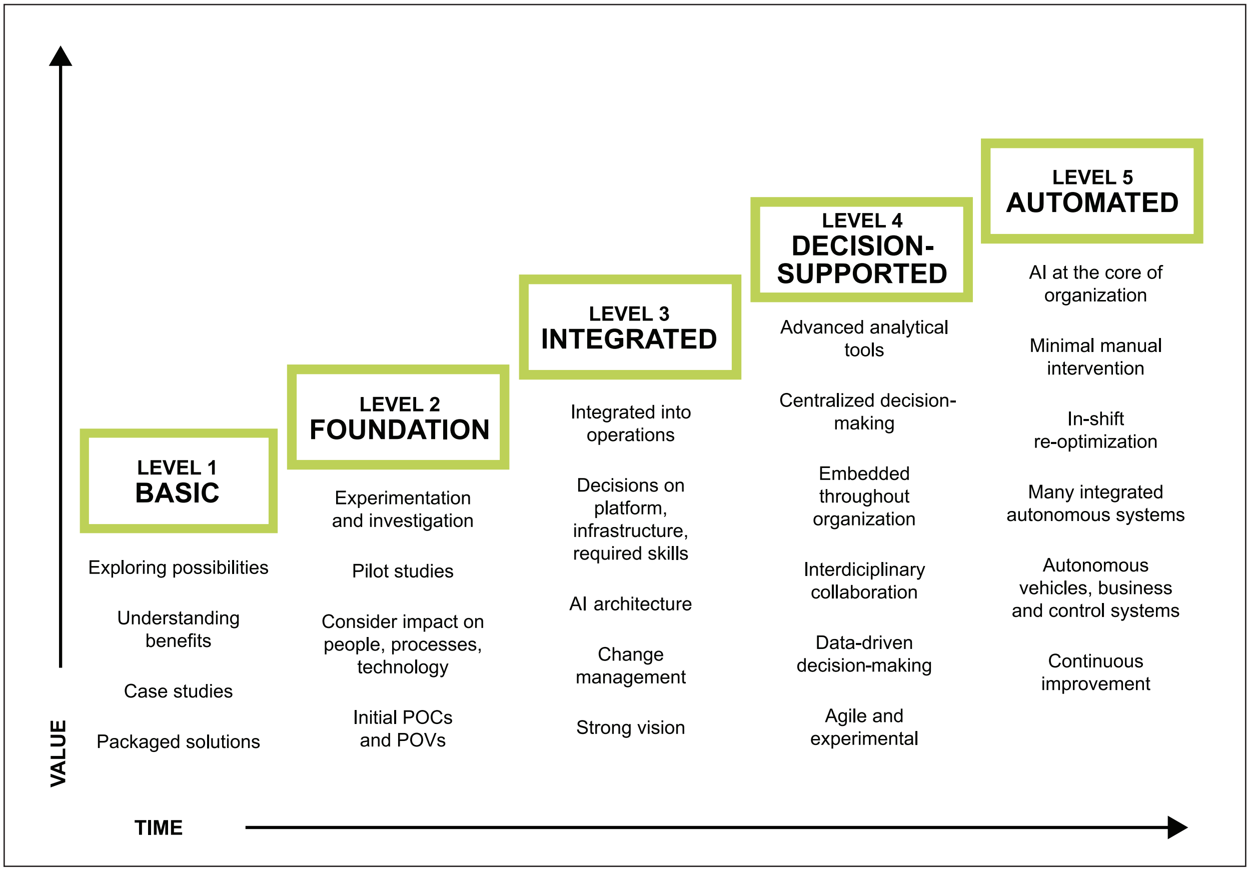

Figure 3. AI in Mining Maturity Model
Abbreviations: proof of concept (POC), proof of value (POV)

There have been several big data maturity frameworks proposed. Many follow a framework that resembles the Capability Maturity Model (Humphrey, 1989). These include:

- TDWI Big Data Maturity Model (Halper & Krishnan, 2013, p. 16)
- IBM's Big Data & Analytics Maturity Model (Nott, 2015)
- EMC's Big Data Business Model (Schmarzo, 2016)
- Hortonworks Big Data Maturity Model (Dhanuka, 2016)

Other types of maturity models have also emerged. For example, "The Roadmap to Machine Learning Maturity," developed by Algorithmia (2019), describes maturity in terms of four factors: data, training, deployment, and management and notes that each can mature at a different rate.

2.2 LEVEL 1: BASIC

A Level 1 organization is exploring what AI is, what it does, and how it can benefit them. These organizations can use documents like this one to understand how AI is relevant to the mining industry and to identify the resources required for gaining an introductory understanding of AI (see Section 1.1 and the section titled External Resources in Appendix B).

Organizations at this level are already using AI and machine learning tools, though most individuals may not be aware of it. Typically, AI at this level comes in packaged solutions that act as black boxes. For example:

- Facial feature recognition in driver safety systems. These systems use computer vision and machine learning to recognize potential distractions or micro-sleep events and alert the driver.
- The use of cameras to determine particle size in digger and mill operations, which is another computer vision application.
- Automated planning and scheduling, or AI planning, which is often already in use in rail and shipping scheduling software.
- Autonomous vehicles, which bring together numerous AI technologies that are mostly hidden from operational users.

In smaller organizations, this packaged use of AI fits the operational needs perfectly and there is no need to invest further. However, the impact of AI is limited at this level; embedding AI into the organization's core is the only way to achieve the technology's greatest potential value.

Does AI align with the organization's digital transformation strategy?

Implementing AI can be considered as part of an organization's digital transformation strategy and should only be implemented if it aligns with this strategy and the organization is at the right stage of its journey. Successful digital transformation requires that maturity is achieved through digitization and digitalization processes, as defined in Figure 4.

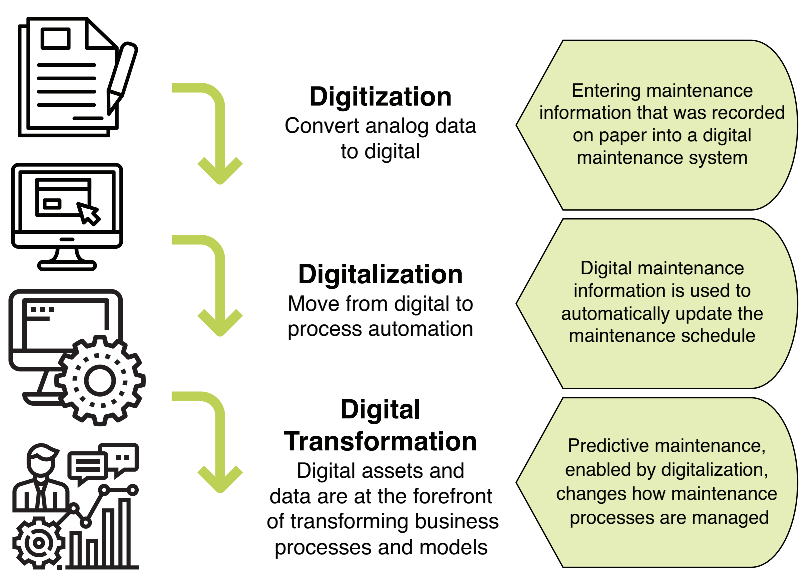

Figure 4. Digitization, Digitalization, and Digital Transformation

Icons made by "Vectors Market", "Icongeek26", "prettycons", "Becris" from www.flaticon.com

2.3 LEVEL 2: FOUNDATION

A Level 2 organization is starting to lay the foundations of their AI strategy. At this level, there is a degree of experimentation and investigation into what the benefits of the technology are and what the organization requires in order to realize them. At this stage, organizations will conduct pilot studies to gain experience. These initial investigations into AI-related projects are often limited to one or two business groups rather than the entire organization.

Distinguishing between data science and machine learning is also important at this stage because data science projects are typically easier starting points. Section 3.1 on Choosing an AI Project covers the process of identifying machine learning and data science projects.

In any business transformation program, the three pillars of people, process and technology need to be addressed. At the foundation level, the following should be considered:

- **People:** Do we have the people in the organization who understand AI and data science? What are the challenges in getting the workforce to accept these changes? Do we build a team by using internal or external resources (see Section 3.3)?
- **Process:** How will this affect the way we are working? Are AI technologies useful tools for improving processes? Are the data capture processes complete, accurate, and consistent? Are the necessary process changes in place that will enable the organization to accurately capture this information?
- **Technology:** How does AI fit into the organization's existing technology stack? Will an on-premise or cloud-based solution work better (see discussion on common platforms in Section 3.5)? How can the organization keep up with the rapid rate of technological change in AI, machine learning, and data science?

Proofs of concept or proofs of value are completed at this level in order to prove that the technology is beneficial and valuable. This stage can often be viewed as presenting an opportunity to play with the technology, but it should instead be seen as the time to prove to the organization that implementing AI can help it achieve significant business benefits.

What's the difference between a proof of concept and a proof of value?

Often a proof of concept focuses on proving that a technology will work. This could be trialling a particular piece of machinery or type of sensor. Success is often measured solely upon whether the trial worked.

A proof of value determines whether there is business value to a project. For example, a new type of drill bit might work successfully on the drill rig (proof of concept) and the increased durability has a potential cost saving of 10% (proof of value).

2.4 LEVEL 3: INTEGRATED

A Level 3 organization is beginning to mature, and AI is becoming integrated into the business operations. The organization will need to make serious decisions about building a solid platform from an infrastructure, skills, and process point of view.

An AI architecture is important at this stage; it defines both the infrastructure and the tools that will enable the quick and easy implementation of AI projects across the organization. Collating and identifying the data needed for a project is a key challenge at this stage, and a solid AI infrastructure defines a common way to store, retrieve, and label data.

Investing in people and culture is also important during this stage. Investing in the right experts, whether internal or external, is essential for success (see Section 3.3). Change management and training are also critical and there should be a strong management-led vision or message that clearly defines what is to be achieved and what the benefits and risks are (see Section 3.5).

2.5 LEVEL 4: DECISION SUPPORTED

Level 4 organizations leverage advanced analytical tools in order to provide centralized decision-making capabilities and further improve the operation. At this level, AI is embedded throughout the organization, and it is likely already using a variety of AI technologies. The new technology will significantly affect many business processes and it will be important to ensure that good governance is in place. The following three changes* need to occur in order to embed AI into the organization and to achieve this level of maturity:

1. **"From siloed work to interdisciplinary collaboration."** Interdisciplinary collaboration not only means that domain experts from different areas of operations are working with data scientists and analysts but also that there is collaboration across all areas of operations. For example, predictive maintenance strategies can be rolled out across mine, processing plant, rail, and port operations and the benefits can be fully realized with full collaboration between these entities.
2. **"From experience-based, leader-driven decision making to data-driven decision making at the front line."** In building AI-based systems, operator knowledge is embedded into the system so that when the system makes a recommendation, it should be trusted and not require management approval.
3. **"From rigid and risk-averse to agile, experimental, and adaptable."** The move to agile practices is being widely adopted in the mining industry. To fully embrace these practices, mindsets need to change throughout the entire organization (See Section 1.4).

*Key changes are quoted from "Building the AI-Powered Organization" by Tim Fountaine, Brian McCarthy, and Tamim Saleh. Harvard Business Review, July-August 2019

2.6 LEVEL 5: AUTOMATED

AI is at the core of the Level 5 organization and most systems and processes are either fully automated or require minimal manual intervention. Achieving this level of maturity is a multi-year journey: most major AI implementations are completed within three years but they can take up to five years (Fountaine et al., 2019). Once this level is achieved, AI is part of the core fabric of the organization and is no longer considered to be a separate technology. For example, autonomous vehicles are a use of AI technology that has been introduced into mining and is becoming mainstream to the extent that they are no longer thought of as AI systems.

The Level 5 mining operation has many AI-enabled systems to which it has aligned and optimized its workforce and business processes. These systems are not static; however, the data that power these processes continually produce new insights that enable continuous improvement and business adaptation.

Rio Tinto - Autonomous Haulage Trucks, West Angelas minesite

3

IMPLEMENTING AI

Introducing AI into an organization involves detailed planning and well-structured implementation.

3.1 CHOOSING AN AI PROJECT

Choosing the right AI project involves brainstorming possibilities and assessing the potential projects from business and technical perspectives, and establishing a workflow for execution (see also Ng, 2019b).

Brainstorming

In the brainstorming phase, getting input from a mix of technology and domain expert participants is important. First, come up with tasks that fit within the following categories:

1. **Assistive intelligence:** AI executes tasks that would be done by humans; decision-making is in the user's hands.
2. **Augmented intelligence:** AI and humans learn from each other; users glean insights from data generated.
3. **Autonomous intelligence:** The AI system is capable of acting on insights generated; decision-making capabilities are with the machine.

Next, identify the general scope of the projects. Ideas can be classified as short-term (< 1 year), medium-term (1-5 years), or long-term (> 5 years) projects based on the time it would take to achieve value. Table 1 provides an example of a scoping matrix.

Table 1. Example Project Scoping Matrix

	SHORT-TERM PROJECT	MEDIUM-TERM PROJECT	LONG-TERM PROJECT
ASSISTIVE	Sensors measuring grain size of rock	Drones taking spectral images of open pit	Hyperspectral images that indicate high-grade deposit
AUGMENTED	Sensors reporting on delays in haul trucks	Alert system for measuring equipment defects	Maintenance protocol for easily degraded equipment
AUTONOMOUS	A belt that sorts rock based on measured size and mineralogy	Routing system that plans haul truck routes	Automated mine plan updates based on dynamic data input

Business Metrics

Once potential AI projects are identified, determine and prioritize suboptimal processes and how each affects the bottom line. Consider the following metrics:

1. Time, personnel, and financial cost of addressing the target process
2. Number of non-AI improvements to be made to the target process
3. Projected value gained from optimizing the target process
4. Data available or the cost of procuring the data needed to assess and improve the target process
5. Number of dependencies between the target process and the non-target processes
6. Number of stakeholders with conflicting interests involved in target process
7. Time it takes to assess the improvement of the target process

Finding similar problems that have been solved inside or outside the mining industry can be useful for further narrowing down the options. Open access archives such as Arxiv (https://arxiv.org/) hold a wealth of published and readily available academic papers.

Note

The presence of AI does not immediately indicate higher efficiency and high return on investment; all other improvements that can be done before committing to an AI project should also be considered.

Technical Assessment

Internal and external validation is necessary to narrow down potential projects. Assess the following technical factors:

1. Knowledge of process inefficiencies and sources
2. All datasets that can measure or indicate inefficiency in the target process (i.e., measure the quantity and quality of the available data.)
3. Amount of time and personnel available to address the target process
4. Knowledge of all unknowns regarding the process and the associated data (usually via audit)
5. Standards and ethics compliance before and after improving the target process
6. Knowledge of short- and long-term financial and profitability impacts
7. Technical expertise available internally and/or the need for external intervention to tackle the project
8. Transferability of process tasks to other tasks along the operations pipeline

Acceptance Criteria

The next step is to determine the key performance improvement metrics for the target process. The acceptance criteria for the successful implementation of AI and machine learning projects generally revolve around the accuracy of the output and the projected return on investment, which are typically expressed in confidence levels. For example, a hypothetical material recognition project might define a high confidence level as an output accuracy of 95% in differentiating between materials, thus reducing the time it takes to do human classifications by 20 hours a week.

It is also prudent to know the optimal limits of the problem even after implementing an AI solution because:

- The hardware may have physical limitations.
- The quantity and quality of usable, clean data may be limited.
- The processes may depend on other inefficient inputs coming from other processes.

To address these limitations, a common trade-off is to choose between creating an extremely specific system at near 100% accuracy or creating a generalized and transferable system that only produces accurate outputs 85% of the time.

3.2 STRUCTURING AN AI PROJECT

Once a project has been chosen, it is important to define what success looks like and what is needed to achieve it. This process will help establish benchmarking and success criteria for the project. Workflows from mapping the problem to validating the solution are outlined in this section. Most AI and machine learning projects use an agile workflow where these tasks are done almost concurrently. Using this methodology, the tasks are further broken down to deliver small incremental improvements in one- to two-week "sprints" (Atlassian, 2019).

1. Map and assess the problem

As an extension of the technical validation used to choose the project, a small team will be needed to scope the following more thoroughly: the minimum performance requirements, business value, analytic methods, and the requirements for the data and domain expertise. Ideally, a combination of on-site expertise, financial forecasters, and AI specialists would conduct this step together. An approach for defining the problem and success criteria is outlined in Figure 5.

Resources on specific software processes

- An overview of software development processes: "The Software Development Process" (n.d.).
- Cross-industry standard process for data mining (CRISP-DM): Chapman et al. (2009)
- Microsoft's Team Data Science Process (TDSP): Microsoft (2017)

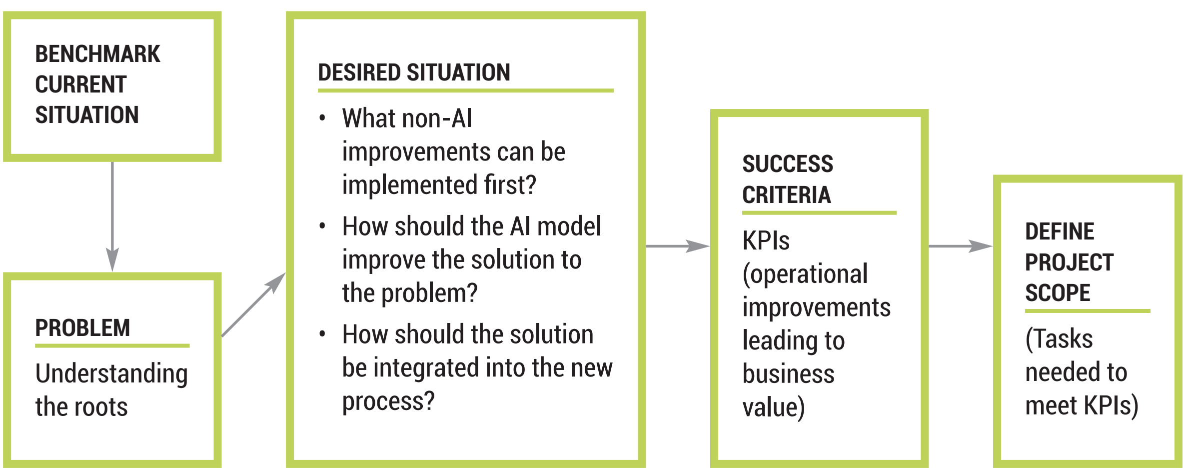

Figure 5. Problem and Success Criteria Definition Process

2. Analyze the data

The data needed to apply an AI-based improvement must be analyzed for their organization and relevance to the project and to assess whether additional data are needed. Figure 6 outlines the data exploration approach.

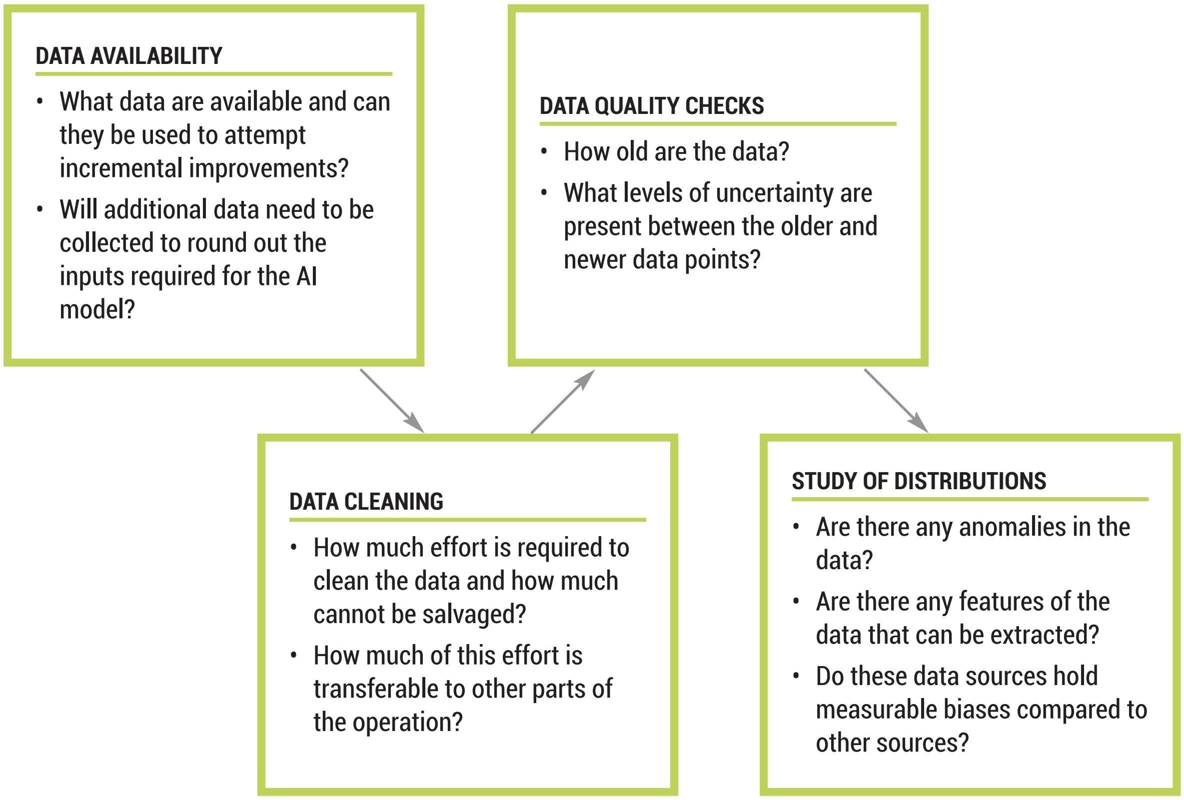

Figure 6. Data Exploration Approach

3. Consider data organization and collection systems

An AI implementation is only as good as the data it can access, so it is important to consider the systems required to handle existing and incoming data. This work includes the organization and data collection protocols as needed; these are outlined in Figure 7.

For the first AI project, up to 80% of the effort can go into collecting and cleaning the data (Ruiz, 2017). However, this task is highly transferable and multiple processes can benefit from the results. Once an efficient system is in place, the data collection and organization for future projects would be minimized.

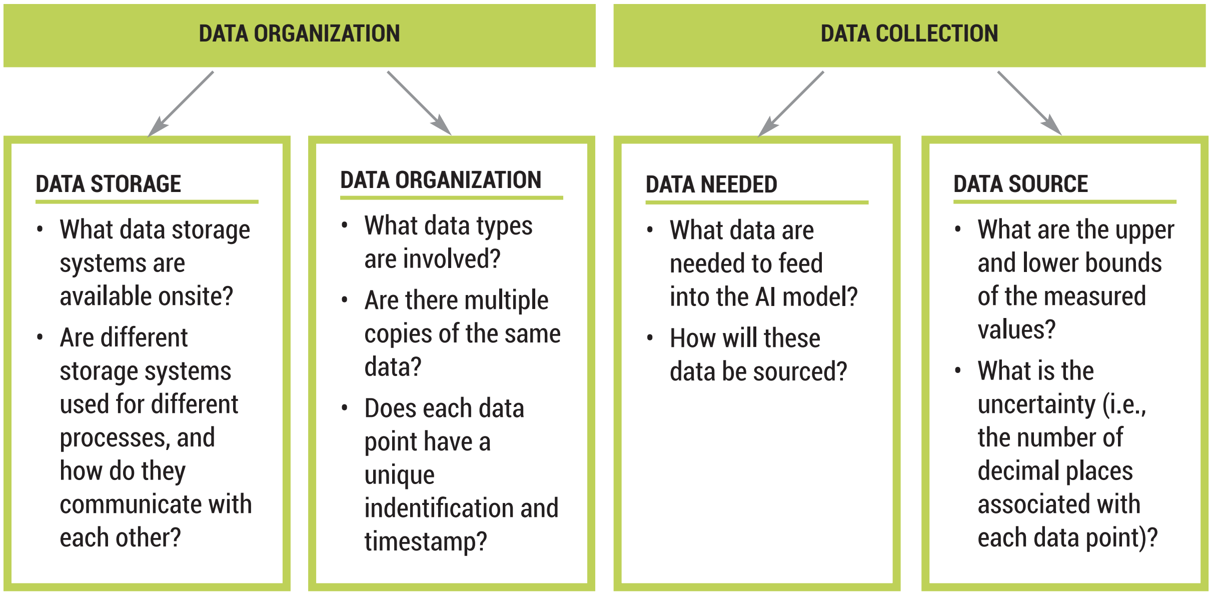

Figure 7. Data Organization and Collection Process

4. AI pilot project

The AI project can begin after a thorough exercise in data organization and and an analysis of the model inputs and/or the data in order to measure for accuracy. Figure 8 outlines an iterative and agile process that would occur throughout an AI pilot project. If multiple tools are developed at the same time, then these processes would occur in parallel.

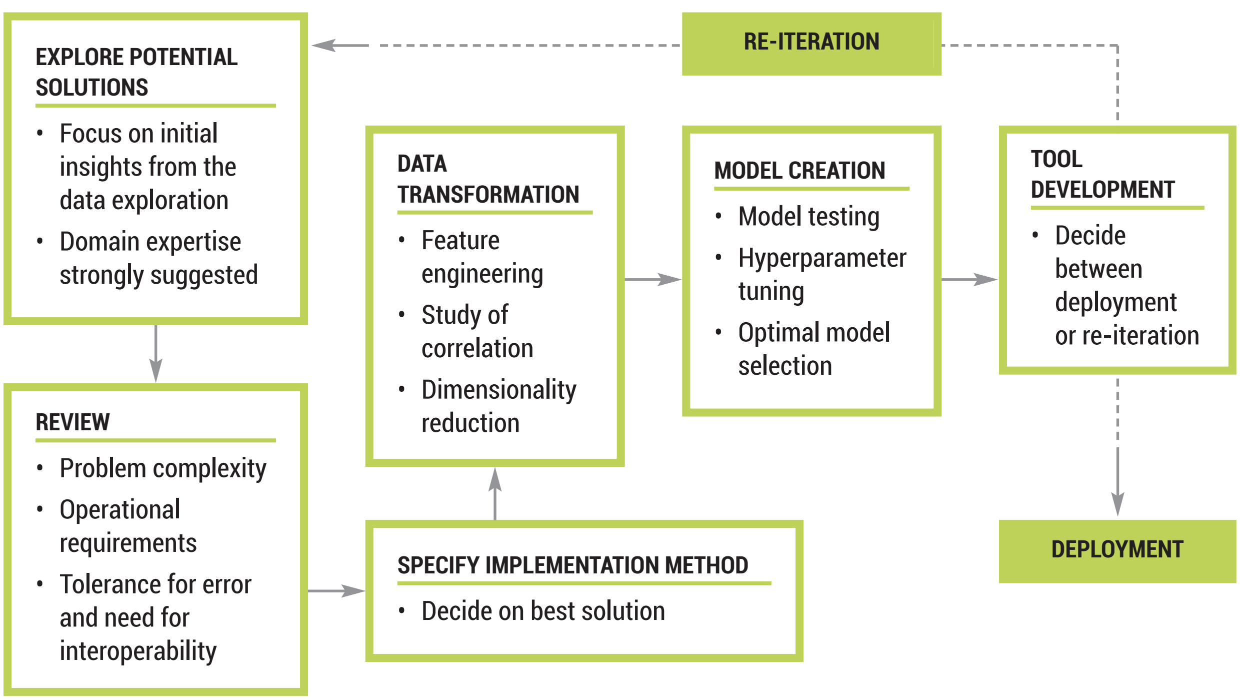

Figure 8. AI Pilot Project Process

5. Validate findings

The AI project should be assessed for its technical performance (i.e., accuracy), business value achieved/projected, transferability to other processes, and potential ease of integration/roll out. Domain experts and AI specialists should convene for the steps outlined in Figure 9 to ensure there is a well-rounded assessment and clear communication with senior management. Once the project is implemented, it is important to periodically review the model or analysis and rerun it with new data to ensure the outcomes remain valid.

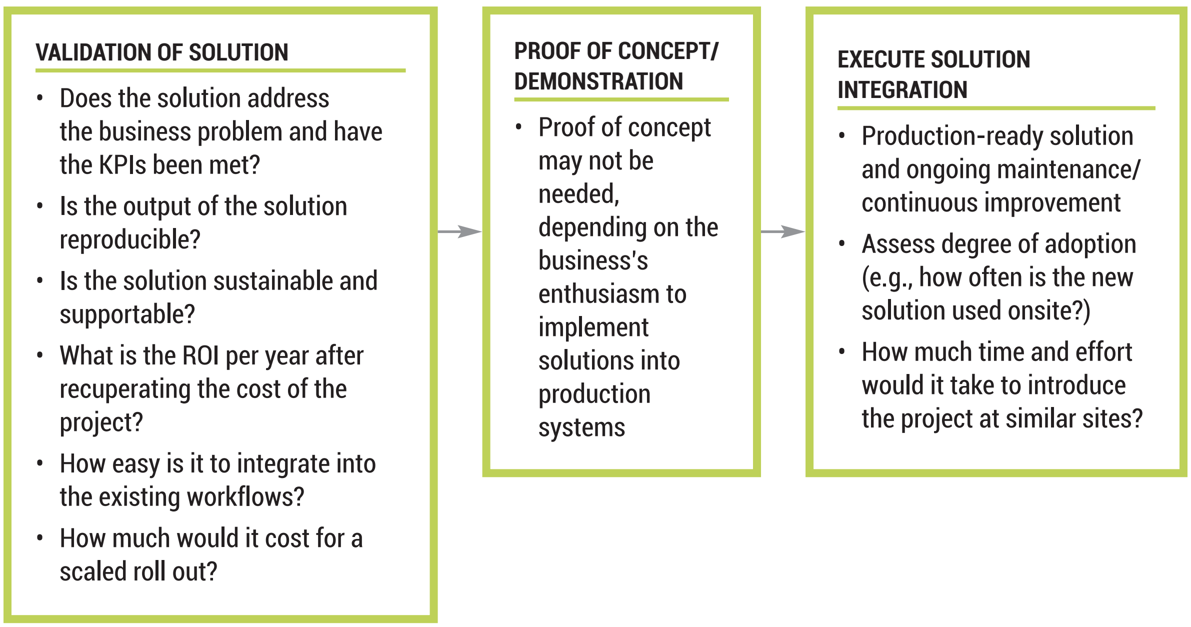

Figure 9. Validation Process

3.3 BUILDING AN AI TEAM

One question that organizations will face early on is whether to build internal or use external resources. Table 2 identifies some factors organizations could consider when deciding which type of resources to use.

Table 2. Deciding Factors for Building Internal or External AI Teams

INTERNAL	EXTERNAL
An internal data team regularly analyzes multiple process-based datasets	Additional data are required from outside sources (e.g., data for training a model)
There is internal domain expertise on the process across multiple sites	Expensive equipment either needs to be rented or purchased to collect additional data
An organized and consolidated data storage system is in place	There is no consolidated and/or easily queried data storage system
The problem is highly dependent on other on-site processes (i.e., the solution would therefore be unique to one or two sites)	The solution can be rolled out to more than three sites (i.e., the solution is generalized for multiple operations)
There is an internal process and data maintenance team and there is training offered on the AI solution	There is no internal data maintenance team when the solution goes to production across one or more sites
Solution implementation is not time-sensitive	Business value from the solution can be gleaned immediately and/or it will directly affect other processes
Other non-AI improvements are needed before applying AI	The process needs an AI/machine learning/hardware implementation to be improved
There is an internal audit system for a well-rounded view of problem	External domain experts primarily have a view of the process with which they are involved

Each organization will lie on a spectrum between internal and external resources but there will be some similarities between approaches:

- **For understanding the problem and the available data:** Internal resources that understand the problem at multiple sites would be ideal because they would have a deep understanding of the domain and be able to draw insights from the data. For example, a geologist with a background in statistics and an internal data specialist would form a great team.

- **For data auditing processes:** External or semi-external entities such as consultants are ideal because they can leverage their experience with similar projects across many similar organizations in order to assess data quality and identify high-level dependencies.
- **For building accurate, robust, and generalizable models:** Experts with access to varied datasets are often needed, so this process is often more easily outsourced to external vendors that bring the right expertise and experience.

It would be beneficial to have a team member to liaise between the domain expert and the AI specialist in order to ensure that all insights are clearly interpreted so that the collaboration runs smoothly and the insights can be readily communicated with key stakeholders such as those in senior management positions. This person would be responsible for:

- Ensuring the project scope, progress, and results are communicated effectively.
- Developing and maintaining project documentation for reference in future AI implementations.

3.4 BENEFITS OF AN AI PROJECT

Carefully considering the processes outlined earlier in this section can help the organization achieve the following benefits from their AI project:

- **Safety:** Many AI projects, such as those that use advanced autonomous systems, remove people from dangerous situations. Applications such as predictive maintenance, object detection, and fatigue and other monitoring systems can also be applied to prevent safety incidents before they occur.
- **Productivity:** Many AI projects improve productivity by eliminating time-consuming and repetitive tasks humans used to do. They can also improve planning and decision-making processes and lead to more productive operations.
- **Sustainability:** The business and operational improvements that AI can enable can help mining operations be sustainable in the long term.
- **Reliability:** By harnessing real-time plant information and maintenance data, the reliability of assets can be increased.

Section 4 provides some examples of AI applications that can help to achieve these benefits.

3.5 CHALLENGES AND RISKS OF AN AI PROJECT

There are many challenges and risks to manage when integrating AI into an existing mining operation.

Industry and Stakeholder Priorities

Industry and stakeholder priorities may not always align with implementing AI. Industries such as mining that have been around for a long time might be wary of adopting AI due to the importance placed on established manual processes and the upheaval implementing an AI-based innovation would have on these.

In terms of stakeholder priorities, on-site stakeholders in supervisory positions may be reluctant to embrace an AI innovation project due to the risk of assigning others whose regular duties are outside of the project's domain, especially if they will be involved in the AI system's continued maintenance. In particular, if the team is mostly internal, then relying on a small team of AI experts could introduce the risk of resource overload and mismanagement.

Technical Risk

Implementing AI into long-established systems and processes involves technical risk throughout the entire process. As covered in Section 3.1, it is essential to have a detailed technical assessment of how the AI would improve a target process. If this assessment is not sufficiently diligent, then its weaknesses could lead to:

- Biases in the model
- Unnoticed feedback alerts
- Misuse of the AI protocol
- Model suboptimality

Useful resource

"AI and risk management: innovating with confidence," a report by the Deloitte Centre for Regulatory Strategy (2018) covers risk management for AI innovations in the financial services industry. There are many parallels with risk management in mining.

Challenges with maintaining data storage and transfer systems could lead to incomplete data inputs, and this could adversely affect the AI model's output. On the hardware side, being unable to assess measurement uncertainty and maintain the equipment would result in suboptimal, biased, and incomplete data inputs into the model (Deloitte Centre for Regulatory Strategy, 2018, p. 8).

Further, if the data infrastructure required to maintain and scale the AI technology were not sufficient, then it will could result in system lags and be unable to handle large data inputs.

Ethics and Trust

Ethics are an increasing concern when it comes to implementing AI, especially concerning bias and privacy.

AI applications are not as objective as many people think because algorithms can incorporate the biases of their developers. AI has already been found to exacerbate existing human biases in areas as diverse as hiring, retail, security and criminal justice (Hao, 2019, para. 1). For example, in a recent study, Raji and Buolamwini (2019) found that some facial recognition technologies have a tendency to mistake darker-skinned women for men. The potential for misidentifying employees is something to keep in mind when implementing AI based on data from sources such as closed-circuit television (CCTV).

In terms of privacy, AI technologies can aggregate data about people that they would not want shared with third parties. For example, wearables such as fitness trackers and smartwatches can provide exceedingly intimate portraits of their owners (see Maddox, 2015). Such technologies are being increasingly used in the workplace, and those implementing them need to ensure they are using personal data ethically and are transparent about it.

Governments have started to build legal frameworks to address some of the privacy issues around AI adoption. The European Union has already implemented the General Data Protection Regulation (GDPR), a law that has a direct consequence on AI technologies via data regulation. This protection framework provides rules and benchmarks for adequate privacy and security practices. The Personal Information Protection and Electronic Documents Act (PIPEDA) is the Canadian equivalent for data privacy. Regulations such as these affect how organizations share and store data.

Useful resource

The following declarations on responsible AI are great resources for companies that want to implement AI in an ethical way:

- *The Montreal Declaration for a Responsible Development of Artificial Intelligence* (2018)
- *The Toronto Declaration: Protecting the rights to equality and non-discrimination in machine learning systems* (Amnesty International and Access Now, 2018)

Privacy International (2019) also provides a good overview of data protection.

"In our surveys and our work with hundreds of clients, we've seen that AI initiatives face formidable cultural and organizational barriers. But we've also seen that leaders who at the outset take steps to break down those barriers can effectively capture AI's opportunities."

(From "Building the AI-Powered Organization" by Tim Fountaine, Brian McCarthy, and Tamim Saleh. Harvard Business Review, July-August 2019)

People and Culture

The personnel and cultural commitment required for short- and long-term AI adoption cannot be underestimated. A key risk associated with short-term adoption is the potential for comprehension gaps between site experts and AI specialists. These gaps may result in the AI model's output not meeting the stakeholder's expectations. Furthermore, issues with distributing and managing resources thoroughly could result in risks such as requiring more resources than intended or stopping the project before completion.

In the long term, a lack of cultural commitment could lead to challenges, so it is imperative to address any cultural shifts that are needed in order to employ the technology. Transparently addressing insecurities about job loss and then encouraging a culture that accommodates the use of the AI model will be necessary to avoid the risk of low levels of buy-in and negative perceptions associated with the change.

Common Standards and Platforms

There is also the risk of choosing the wrong technology or platform and having to revisit some of the initial steps. Identifying common standards and platforms is important for mitigating this risk and preventing increased vendor dependency, especially in situations where parts of data platforms malfunction.

Organizations face the decision of whether to use either cloud- or on-premise-based data infrastructures or whether to choose aspects of both. Features of these, based on those described by Hale (2017), are summarized in Table 3.

A data framework that consists of both types of infrastructures can be created to suit immediate and future needs. Careful management by IT and data engineering personnel—both internal and external—is expected since these systems are needed to integrate an AI solution into an operation.

Table 3. Cloud-Based and On-Premise Data Infrastructures

CLOUD	ON-PREMISE
Monthly or annual subscription	One-time perpetual license
Operating expenditure (small amounts over the product's lifetime)	Capital expenditure (large upfront investment) and requires space to house hardware installation
Hosted on vendor's servers with online access	Installed on operation's servers and/or computers and requires internal IT personnel
Ability to scale exponentially with the use of technology	Requires significant expenditure to scale; risk of system lags if not well planned
Security in hands of vendor; may not be compliant with all of the standards relevant to organizational activities	Security is tailored to the organization's regulations and activities; if planned correctly, air gap security is achievable
Less customizability for specialized tasks and protocols	Tailor-made customization for specific tasks; organization has greater control of design
Short and seamless implementation time and regular system updates operated by vendors	Long implementation time and lack of installation/update support

4

Applications of AI

This section provides some examples of the application of AI to demonstrate some of the ways in which AI is used in mining. As discussed in earlier sections, benefiting from all of these applications requires quality data and sufficient planning.

Process Optimization

Process optimization is an application of AI that harnesses existing data and analyzes variances and trends in order to predict future outcomes and provide decision support. These techniques can be implemented across the mining value chain and can range in scope from total mine optimization to single process optimization and can focus on achieving long- or short-term goals. Some examples in mining are:

- **Corrective mine planning systems:** Mine planning data including deviations are input into a model that analyzes trends and identifies affected areas. The output is then analyzed to develop an integrated and optimized plan.
- **Ore reconciliation and optimization systems:** AI-enabled ore reconciliation and optimization systems streamline the process of comparing and analyzing variances between disparate sources of estimated and actual data in order to understand, forecast, and optimize value.
- **Portfolio optimization systems:** Historical data from the operation, similar operations, and external markets and operational planning data are input into a model. The resulting algorithm makes inferences between these factors in order to predict the measurable financial operational functions.

Predictive Maintenance

Data on machine health are collected from sensors and historical component failure records are input into the model. Predictive maintenance models produce realistic predictions for when potential component failures or malfunctions may occur and provide estimates for when maintenance should be performed in order to avoid potential failures and malfunctions. This can reduce unscheduled downtime and improve safety and productivity.

Safety Monitoring

Some of the principles of predictive maintenance can also be applied to solving other safety and environmental problems. For example:

- **Tailings modelling:** Data from tailings sensors can create a model of the tailings pile; the technical engineer could analyze the resulting data in order to prevent leakage, contamination, or problems with water treatment processes.
- **Geotechnical monitoring systems:** Data inputs from geotechnical sensors, seismic alerts, dashboards, response plans, and historical seismic reports are analyzed in order to provide decision support and mitigate risk in activities such as stress management, drilling, blasting, and cave propagation.

GMG is looking to build a library of case studies that will provide insight into the variety of processes that can be enhanced through AI and the potential benefits derived from it.

Do you have a case study you'd like to share? Contact GMG Managing Director Heather Ednie, hednie@gmggroup.org

Machine Vision

Machine vision technologies use image-based capture to monitor activities. Algorithms are then used to process data and provide insights that support decision-making. Examples include:

- **Safety systems for autonomous vehicles:** Machine vision systems on autonomous vehicles can identify obstacles and unsafe conditions and use these insights to react accordingly.
- **Decision support in drilling and blasting:** Machine vision systems on buckets or shovels can be used to analyze fragmentation and other parameters that can be used to optimize drilling and blasting processes.
- **Monitoring and recording repetitive tasks:** Machine vision systems can be applied to repetitive tasks such as load-haul-dump cycles to automatically record data that can be used to optimize these processes.

Useful resource

A great place to explore innovative applications of AI in mining is by reading through the finalists of the Disrupt Mining Challenge **(https://disruptmining.com/)**

REFERENCES

Algorithmia. (2019). The Roadmap to Machine Learning Maturity. Retrieved from https://blog.algorithmia.com/navigating-the-machine-learning-roadmap/

Amnesty International & Access Now. (2018). The Toronto Declaration: Protecting the rights to equality and non-discrimination in machine learning systems. Retrieved from https://www.accessnow.org/the-toronto-declaration-protecting-the-rights-to-equality-and-non-discrimination-in-machine-learning-systems/

Atlassian (2019). *The Agile Coach.* Retrieved from https://www.atlassian.com/agile

Bostrom, N. (2014). *Superintelligence*. New York: Oxford University Press.

Chapman, P, Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., Wirth, R. (1999). The CRISP-DM User Guide. CRISP-DM Consortium. Downloaded from ftp://ftp.software.ibm.com/software/analytics/spss/support/Modeler/Documentation/14/UserManual/CRISP-DM.pdf

Deloitte Centre for Regulatory Strategy. (2018). AI and Risk Management – Innovating with Confidence. Retrieved from https://www2.deloitte.com/global/en/pages/financial-services/articles/gx-ai-and-risk-management.html

Dhanuka, V. (2016). Hortonworks Big Data Maturity Model. *Hortonworks.* Retrieved from *http://hortonworks.com/wp-content/uploads/2016/04/Hortonworks-Big-Data-Maturity-Assessment.pdf*

Domingos, P. (2015). *The Master Algorithm*. Hachette Group.

Fountaine, T., McCarthy, B., & Saleh, T. (2019, July-August). Building the AI-Powered Organization. *Harvard Business Review*. Retrieved from https://hbr.org/2019/07/building-the-ai-powered-organization

Goodfellow, I., Bengio, Y., Courville, A. (2018). *Deep Learning*. MIT Press.

Hale, Z. (2017, October 4). Cloud ERP vs. On-Premise ERP. *Software Advice*. Retrieved from https://www.softwareadvice.com/resources/cloud-erp-vs-on-premise/

Halper, F., & Krishnan, K. (2013). TDWI big data maturity model guide. Retrieved from https://tdwi.org/whitepapers/2013/10/tdwi-big-data-maturity-model-guide.aspx

Hao, Karen (2019, February 4). This is how AI bias really happens—and why it's so hard to fix. *MIT Technology Review*. Retrieved from https://www.technologyreview.com/s/612876/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/

Humphrey, W. (1988). Characterizing the Software Process: A Maturity Framework. *IEEE Software* 5(2). doi: 10.1109/52.2014

IBM (n.d.). The Four V's of Big Data (n.d.). [Infographic] Retrieved from https://www.ibmbigdatahub.com/infographic/four-vs-big-data

International Organization for Standardization (2018a). *Information technology — Big data reference architecture — Part 2: Use cases and derived requirements* (Standard No. ISO/IEC TR 20547-2:2018). Retrieved from https://www.iso.org/standard/71276.html

Maddox, T. (2015, October 7). The dark side of wearables: How they're secretly jeopardizing your security and privacy. *TechRepublic*. Retrieved from https://www.techrepublic.com/article/the-dark-side-of-wearables-how-theyre-secretly-jeopardizing-your-security-and-privacy/

Microsoft (2017). TDSP (Team Data Science Process). Microsoft. Downloaded from https://docs.microsoft.com/en-au/azure/machine-learning/team-data-science-process/overview

The Montreal Declaration for a Responsible Development of Artificial Intelligence (2018). Retrieved from https://www.montrealdeclaration-responsibleai.com/

Ng, A. (2018). The AI Transformation Playbook. Landing AI. Retrieved from https://landing.ai/ai-transformation-playbook/

Ng, A. (2019a). AI for Everyone. MOOC offered through Coursera. Retrieved from https://www.coursera.org/learn/ai-for-everyone

Ng, A. (2019b, February 6). How to Choose Your First AI Project. *Harvard Business Review*. Retrieved from https://hbr.org/2019/02/how-to-choose-your-first-ai-project

Ng, A. (2019c) Machine Learning. MOOC offered through Coursera. Retrieved from https://www.coursera.org/learn/machine-learning

Nott, C. (2014, August 15). Big Data & Analytics Maturity Model. *IBM Big Data & Analytics Hub*. Retrieved from https://www.ibmbigdatahub.com/blog/big-data-analytics-maturity-model

Privacy International (2019). Data Protection. Retrieved from https://privacyinternational.org/topics/data-protection

Raji, I.D. & Buolamwini, J. (2019). Actionable Auditing: Investigating the Impact of Publicly Naming Biased Performance Results of Commercial AI Products, presented at the Association for the Advancement of Artifical Intelligence Conference on AI Ethics and Society, Honalulu, Hawaii, January 2019. Retrieved from https://dam-prod.media.mit.edu/x/2019/01/24/AIES-19_paper_223.pdf

Ruiz, A. (2017, September 26). The 80/20 data science dilemma. *InfoWorld*. Retrieved from: https://www.infoworld.com/article/3228245/the-80-20-data-science-dilemma.html

Schank, R.C. (1987). What is AI, Anyway. *AI Magazine*. 8(4). DOI: https://doi.org/10.1609/aimag.v8i4.623

Schmarzo, W. (2016, March 22). Big Data Business Model Maturity Index Guide. *InFocus: Dell Technologies Services & Dell Digital Blog*. Retrieved from https://infocus.dellemc.com/william_schmarzo/big-data-business-model-maturity-index-guide/

The Software Development Process. (n.d.). Retrieved from https://sce2.umkc.edu/BIT/burrise/pl/software-process/

World Economic Forum (2018). The Future of Jobs Report 2018. Retrieved from https://www.weforum.org/reports/the-future-of-jobs-report-2018

APPENDIX A: RELATED WORK

The work being done in the GMG AI in Mining Working Group is not taking place in a vacuum. Many other organizations are working on AI-related guidance and standardization. These projects could be leveraged in the future.

The International Electrotechnical Commission (IEC) and the International Organization for Standardization (ISO) set up a joint committee (ISO/IEC JTC 1/SC 42) to carry out standardization activities for artificial intelligence (https://www.iso.org/committee/6794475.html). They have published three big data standards (ISO 2018a, ISO 2018b, and ISO 2018c). They are currently working on several others, including a Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML) and Artificial intelligence - Concepts and terminology.

The IEEE Standards Association have similarly set up working groups over the last few years. Active standards projects include:

- P7008 – Standard for ethically driven nudging for robotic, intelligent, and autonomous systems (https://standards.ieee.org/project/7008.html)
- P7009 – Standard for fail-safe design of autonomous and semi-autonomous systems (https://standards.ieee.org/project/7009.html)
- P7010 – Well-being metrics standard for ethical artificial intelligence and autonomous systems (https://standards.ieee.org/project/7010.html)

Other GMG Initiatives

Several other GMG Working Groups intersect with AI, including:

- The Autonomous Mining Working Group
- The Interoperability Working Group
- The Data Access and Usage Working Group

The Partnership on AI (www.partnershiponai.org) is a multistakeholder organization founded by leaders in the tech industry. It aims to develop and share best practices, advance public understanding, provide an open and inclusive platform for discussion and engagement, and consider and support socially beneficial applications. They currently have three working groups:

1. Safety Critical AI
2. Fair, Transparent and Accountable AI
3. AI, Labour and the Economy

APPENDIX B: EDUCATIONAL RESOURCES

There are many resources available that will help further AI knowledge. In this section we focus on online resources, but there may also be colleges and universities offering courses in your local area.

Massively Open Online Courses

Massively open online courses (MOOCs) have exploded over the last few years. They provide an easily accessible source of educational material. Some popular platforms that offer courses on AI, machine learning, data science are listed in Table B1.

Table B1. Selection of MOOC Platforms that Offer Courses on AI-Related Fields

NAME	DESCRIPTION
Coursera (coursera.org)	Offers a mixture of free and paid courses in collaboration with universities and industry
Udacity (udacity.com)	Offers a mixture of free and paid courses in collaboration with universities and industry
EDX (edx.org)	Offers university style courses, not for profit
MIT OpenCourseWare (ocw.mit.edu)	Site where MIT makes materials from their courses, including video lectures, freely available
Fast.ai (fast.ai)	Technical introduction to machine learning and deep learning from Jeremy Howard, University of San Francisco
The School of AI (https://www.theschool.ai)	Slightly less formal approach to teaching AI concepts
Khan Academy (https://www.khanacademy.org)	High school / entry-level university topics, available in multiple languages
Kadenze (https://www.kadenze.com)	Offers some technology courses, mostly specializing in courses geared towards art, music, and creative technology
DataCamp (https://www.datacamp.com)	Offers paid and free courses for those specializing in programming languages for data science
Kaggle (https://www.kaggle.com)	Platform for entering machine learning competitions, finding curated data sets, and learning about machine learning

Postgraduate Studies

Many colleges and universities now offer postgraduate certificates, diplomas, and degrees, some of which are online. A selection of these are listed in Table B2.

Table B2. Selection of Postgraduate Studies Offered in AI-Related Fields

NAME	DESCRIPTION
Harvard (Extension School)	Data Science Certificate (https://www.extension.harvard.edu/academics/professional-graduate-certificates/data-science-certificate) Master of Science in Data Science (https://www.extension.harvard.edu/academics/graduate-degrees/data-science-degree)
Georgia Institute of Technology	Master of Science in Computer Science, online with AI and machine learning specialization options (http://www.omscs.gatech.edu)
UC Berkeley	Master of Information and Data Science, online (https://datascience.berkeley.edu/)
Colombia	Master of Science in Data Science (https://datascience.columbia.edu/master-of-science-in-data-science)
Glasgow University	Master of Science in Data Analytics, online (https://www.gla.ac.uk/schools/mathematicsstatistics/postgraduate/analytics/)

GLOBAL MINING GUIDELINES GROUP

GMG

Innovation through Collaboration