Evaluation of Models for Interaction Probability in Autonomous Monitor and Control Environments

Robert Bissonette and Samir Sbai
NIOSH Spokane Mining Research Division (SMRD)

ABSTRACT

A critical component of implementing autonomous systems is the ability to use information from a robust sensory and information system to make sound, predictable, and repeatable decisions to assure safety. This means that any such system would need to evaluate what objects (people, equipment, boulders, etc.) of interest exist in that environment (with relevant data about identity, trajectory, characteristics, etc.), project a future continuum of those objects (including considerations for variabilities), and make control intervention choices to minimize risk due to exposures in the present operational area. The components of an overall system would include: a) sensor and data fusion, b) probability projection of future states, c) generation and evaluation of alternate futures, and d) implementation of best course of action to reduce risk. The National Institute for Occupational Safety and Health (NIOSH) Spokane Mining Research Division (SMRD) is studying this framework for applications for autonomous mining equipment. This paper focuses on using information from the sensor fusion engine to generate a continuous projection of future and alternate states (b and c above) incorporating the factors of confidence, accuracy, tolerance, and other variabilities. We refer to this as Machine Situational Awareness (MSA).

INTRODUCTION

In this paper, a technical approach for safety assessment of interaction probability in autonomous mining environments is proposed for mobile, stationary, and hybrid equipment. This research was informed by preliminary evaluations done in the NIOSH pilot project Assured Autonomy Safety Intervention System Technology (AASIST) and Machine Situational Awareness (MSA) research as part of the Conveyor Safety project. We note that ambiguity exists as to how autonomous safety is defined in academic/scientific literature as well as in industrial practice [1]. For the scope of this work, autonomous safety (as we define it) is proposed as Assured Autonomy (AA), the ability of autonomous equipment to operate safely in the presence of humans, as manually operated equipment would.

The current state of autonomy in mining applications is typically the “autonomous by isolation” phase [2]. The progression to “full autonomy” is projected to be difficult and elusive because there lacks an advancement in real-time risk assessment and perception that would enable that transition. Currently, mine sites are designed to accommodate autonomous equipment without exposing humans to the autonomous operation. For haul trucks, this requires that a mine be conceived or reconstructed to have exclusion zones, wide separated roads, robust administrative and electronic systems to monitor and control access to those exclusion zones, as well as other requirements. These are costly undertakings that serve to deprive smaller mining operations of the safety and productivity that could be offered by autonomous operation. Looking forward 5 to 10 years, if autonomy in mining is going to be the rule rather than the exception, systems will need to exist that allow autonomous equipment to operate in the presence of humans. This effort mirrors the collaborative robot (CoBot) development that has taken place in the manufacturing sector for the past several years [3].

AA has been studied extensively (and with mixed results) [4] in the automotive sector and there is a great deal of useful research applicable to mining, but mining has some special conditions that may make this safety supervision considerably less complicated. Relative to the automotive sector, mining occurs in a relatively static and limited topography (surface and underground), all the players (equipment, features, personnel, wildlife) are known, rules of engagement are well established, and the environment is generally sterile (e.g., no new intruders are present in the mine environment). This greatly limits the scope of developing new digital twin components and topography after initial integration, as well as being able to clearly define expected rules of engagement among players. Given this, any models used for object recognition (in the sensor fusion portion) can be thoroughly validated before being put into operation.

The Economics of Safety

Safety, including assured autonomy, ultimately boils down to risk/reward. If there was absolutely no tolerance for risk, the mining industry would not exist [5]. What ultimately drives safety in a mine environment is economics and social responsibility. Social responsibility means placing high value on individuals: their time, their health (mental and physical), their contribution, and their futures. We conclude that any framework must include a means to model that ethos.

About Artificial Intelligence (AI) and Machine Learning (ML)

There has been a series of well-known safety failures associated with artificial intelligence (AI) and machine learning (ML) **[6]**. These failures almost always factor to a set of circumstances for which the AI/ML was not trained for and thus caused the system to not respond correctly. The problem can be summarized by the lack of ability to validate these methods because the internal logic is not transparent. Tasks like object recognition would be very difficult to accomplish without AI/ML, but we propose that the limited number of objects in a mine comprise discrete datasets that are small enough to be validated. The same is not true for open scenarios and environments where nearly infinite input combinations are being processed **[4]**. Because of that, we conclude that AI/ML is appropriate for front-end activity like object identification and behavioral prediction, but not for real-time risk analysis.

The Ideal Solution

The ideal solution would constantly evaluate environments and activities to maximize reward in economic and social terms. We, as motorists, do that every time we drive a car. We subconsciously make decisions on questions such as, “Is driving another 1 mph faster worth the marginally higher risk?”. We ask that question until we are comfortable with the balance we have chosen between efficiency and safety [7]. We process millions of such questions without even knowing it, and thus the world moves. We think the framework we propose herein would accomplish that without using AI/ML in the risk recognition, assessment, and decision portion of a safety intervention system.

A wish list of features of such a framework would be comprised of many things, including the following:

- Sophisticated long-range perception
- Robust and accurate short-range perception
- Pragmatic processing of risk/reward scenarios
- Rapid evaluation of alternate futures that minimize risk (probability x cost)
- Predictable, consistent, accountable responses to hazards
- The ability to operate indefinitely without rest.

Some of these features loosely reflect the way humans deal with the emergence of hazards in everyday life. Others reflect improvements on human limitations. Humans are very adept at preserving their health and wealth while engaged in activities that would seem very dangerous to an outside observer (such as in Figure 1). Anyone that has driven in one of the highly populated cities in this world can appreciate just how good humans are at effortlessly negotiating multiple risks concurrently [6]. Humans also have limits and shortcomings that a good safety intervention system could overcome.

Figure 1. Everyday hazard in a congested traffic area. (ShutterStock-1140426830)

RESEARCH METHODOLOGY

The research for the effort described in this document will involve: a) identifying the framework components (algorithms, data processing, probability schemes, etc.) that are most likely to succeed, b) integrating those components into a framework, c) testing function against predicted performance in simulation, d) testing the framework in a small-scale “sandbox” environment and d) ultimately collaborating with industry partners for full-scale testing in real environments. This is an iterative process in which segments of the framework are replaced or modified as a result of failing to meet expected performance metrics.

REAL-TIME RISK ASSESSMENT (RA) MODEL TYPES [8]

Single Behavior Threat Metrics (SBTMs)

An STBM assumes excellent accuracy and known intentions. It is intended for "perfect world" and simple conditions. It can be used for robotic installations where position is always known accurately and there are few and known interactors. Discrete responses and interactions are programmed to address very limited scenarios that create hazards. These are not characteristics of most mining equipment and environments.

Optimization Methods

Optimization is a predecessor to ML. Known mathematical and functional relationships can be modeled and adjusted to reliably predict system behavior and act to prevent hazards. This approach works where limited physical and/or chemical interactions are well understood, and external influences are not present. The system is modeled with response parameters that are actively adjusted to create the desired outcome. This approach works very well for processing plants, chemical manufacturing, and industries where variability is limited, and results are determinant.

Formal Methods

Formal methods of real-time risk assessment require extensive programming for complex circumstances, which is basically impossible in open environments (where all players are not known). This method requires considering every condition and has been the default approach in some cases but has failed in mobile autonomy because there are always scenarios, dependencies, and interactions that have not been considered. With this approach, it is very difficult to deal with variances and instabilities because it is based on foreseen events that involve a fixed set of interactors. This approach can only deal with what it has been programmed to deal with and cannot evaluate risk, only respond to a given set of circumstances. RA must be proactively done before programming.

Probabilistic Methods

Probabilistic methods for a dynamic environment involve creating probability fields that are cast along an educated projection of the near future to identify risks. This method can involve a large number of players (variables) and relies on a sound resolution of variance (the inverse of confidence). If graphic monitoring is not required, the calculations (processing demand) can be quite small. Our research during development work for this project helped establish this approach as most appropriate for real-time risk assessment.

Data-Driven Methods

These methods require utilization of machine learning and artificial intelligence. Large amounts of experience from the related environment are used to try to develop a trustable response to all possible scenarios. This method has been used with considerable success in relatively static environments but fails when a previously unexperienced set of circumstances causes an indeterminant response. These methods can and are used in object recognition, but do not seem workable for real-time risk assessment.

HOW WE GET THERE – SAFETY INTERVENTION FRAMEWORK

Perception[1]

There is considerable effort underway in the automotive, manufacturing, and mining sectors to develop robust sensory systems that give a system considerably more time to change the future in terms of safety [9]. The ultimate goal is to collect accurate and timely data on the environment and players (mobile and stationary objects, including humans) within that environment. This means the proper use of good data and rapid resolution of conflicting data. Sources could include GPS, V2V (Vehicle to vehicle), 4D radar, Lidar, vision systems, laser range detectors, central control system, etc. Relevant environmental data would also be collected, such as temperature, humidity, dust, ambient light, road conditions, as well as equipment status and condition data.

Track Forecasting

This portion of the framework would contain path prediction, calculation and modification of variances, application of rules, and process prioritization. This loosely reflects the process by which we humans absorb our environment and project the path and behavior of ourselves as well as other players (people, cars, flying objects, etc.). We use this as an analogy to visualize how a probability field risk assessment system would work. It is specific to how we negotiate our daily lives without running into things or have things run into us. Our brains assume players will act in a particular way (follow rules) until they don't, at which point we prioritize our attention (processing power) on the perceived future events that pose a risk to us or our property. We subconsciously evaluate possible problems based on what we know about the players and their current behavior and assign appropriate uncertainty or confidence to our assessment. We instantaneously calculate alternate responses and choose the one that carries the least risk. Those that have experienced driving in any number of nonwestern cities (see Figure 1) can likely appreciate how good humans are at this. A good system would have that innate sense of what is coming, what to pay attention to, what the options are, and which options reduce risk the most.

There is a school of thought that proposes that trying to emulate human perception and processing is not the best approach for safety intervention [9]. That may end up being recognized as true, but for now, human-like emulation is a lofty and worthy goal. The industrial world has relied on it for centuries, and so far, no technology matches humans ability to visually discriminate.

Real-time Risk Assessment

Human brains operate mostly in the low-level sensorimotor control system while driving a vehicle [5], which means that we drive in an "automatic" mode. We think this is indicative of just how

---

[1] This section is not specifically part of the topic herein but is an important part of the associated research.

fundamental this process will eventually become for autonomous equipment. We see that the approach outlined herein is a step in that direction. Theoretically our model has the benefits of AI, but is predictable, traceable, repeatable, and is relatively cheap in terms of computing effort.

Our concept relies on one easily calculable data point for each object pair in the equipment environment. The largest portion of processing power for the overall MSA framework will be in the "perception (sensor and data fusion)" portion of the framework that is not covered herein.

Intervention[1]

Humans and animals have parallel control systems, both conscious and subconscious [6]. The conscious mind controls actions congruent with decisions we make (where to go, what to do). The subconscious mind intercedes when it perceives danger and elicits an instinctual response. This is a sound model on which a safety intervention system would be modeled. We briefly explore the paradigm by which the wishes of the control system and the safety intervention system "negotiate" a response without slowing down the intervention, but it is not specifically part of the topic herein.

THE PROBABILITY KERNEL

Path Projection

We can only gather data about position and derivatives (speed, acceleration, rate of change of acceleration) in the current timestamp and use them to project our "best guess." For each refresh cycle, the kinematic model is updated and recast based on current information. The projected path as a function of position and time (thus, also speed) is the smoothed integration of position derivatives for displacement (typically a 3rd degree polynomial function). This approach applies to any object in motion, but for mobile equipment the ground profile will affect the motion in a predictable way, thus it is also included for its influence.

Path projection recognizes that some objects abide by rules (e.g., rules of the road) or attempt to do so. For example, a haul truck drives within 3 standard deviations (3σ) of an assigned track, unless there is something wrong. The path for that player would be predicted as such with the associated variance (a function of σ). But, if it is observed deviating from that track by more than 2σ, a much wider variance is assigned to it (through a broadened probability field). The result would typically be a raised risk of interaction and thus an elevated attention in terms of building confidence in the behavior of that object (more data points, attempts at affirmation of data, etc.).
For each identified object, there is an established dataset for "normal" operation: dimensions, weight, turn radius, determinacy, maximum and typical acceleration and deceleration, etc. These parameters are used to confine variance for that object.

Probability Field

Our only actual measurement of position, speed, etc. would be in the current timestamp (t=0). The probability field surrounding a haul truck is relatively constrained, and variance is generally attributed to factors such as sensor accuracy and historical variance. As we project position and motion into the future, there is less confidence in those values. If an object is recognized as a road follower (typically abides by rules), then the side-to-side variance will be relatively small unless it deviates from its track. Confidence will diminish as a function of time into the future because we are only projecting what will be the position and derivatives. For a haul truck, this can be envisioned as a long slender field that surrounds the truck as it moves down the road, growing longer along the axis of travel the farther into the future we project. If the truck deviated beyond an established limit from its assigned track, that field would immediately widen dramatically and "level out" to create a standard probability distribution based on a larger variance.

Probability Field Interactions

If the position of the fields is cast along the projected paths, their interaction (overlap) can be found as a function of time (Fig. 2). Although these interactions can be very subtle and small, they are observable and calculable. Indeed, we are not interested in the gradients or the distribution, only in the peak of the waveform that appears along these interactions. These peaks are a function of the standard distribution of each objects position and the line segment that connects them. This greatly reduces the required calculation cost of monitoring and projecting multiple objects. Primarily, a controller would only be interested in projected interaction between itself and other objects, but in fact could monitor all interactions within its perceptible environment.

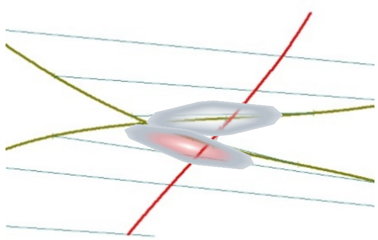

Green – projected path
Red – probability "ridge"
Blue – line segments between players at increments of time
Clouds – probability fields at nearest passing

Figure 2. Probability field interaction.

Calculation

An iterative algorithm would observe only the peak interaction probability as a function of time. The apex would occur at the time that the players are projected to pass closest to one another. This point is the highest probability of interaction and ultimately the only value that we actively monitor. See Figure 2 for the calculation of total probability (peak value x combined vehicle width). The risk of that possible interaction is calculated as the probability times the projected cost if contact occurs. That cost is an engineered, predetermined function of speed, angle of incidence, value of the equipment, and other factors). For each possible interaction, a risk is calculated (probability x cost). Example: $P=10^{-6}$, C=100,000, Risk = 10. It is important to note that the calculation must include social and environmental impact in some manner. That is an ethics discussion.

Alternate Futures

Each significant risk is monitored and updated, and highest risks are re-evaluated for futures that include a range of actions that the monitored equipment can undertake (slow down, speed up, turn right, turn left, alert other equipment, etc.) to reduce those risks and evaluate against the cost of those actions (slowing down costs production, turning wears tires, driving on the shoulder increases risk, etc.). Thousands of scenarios can be calculated in a few milliseconds, and the total risk can be minimized to determine a recommended action. This is necessarily an iterative evaluation that would curve fit the results from nonlinear progression (more calculation near the risk peak).

Intervention

Since this system is a parallel processor, there needs to be a negotiation between the safety intervention system and the primary equipment controller. This is not a specific topic of this paper, but we mention it because it is an important consideration in integrating safety intervention into an autonomous control system. It is conceivable that a system could "make suggestions" based on economic evaluation that could be ignored by the main controller. That would contrast with high-risk, emergency commands that would require immediate action. In general, it is likely that this interface would have varying levels of intervention, but this would require a very high reliability/confidence, and thus is deserving of significant consideration and study.

NEXT STEPS

This research effort will need to focus on the development of logic and logistics to determine if the concept is viable and feasible. Because of the extreme liability sensitivity of this area, validation and supporting research will be as important as the approach itself. We are actively soliciting feedback and collaboration with industry stakeholders in an effort to catalyze a consensus solution that elevates the paradigm for all those with a vested interest in the future of mining. We recognize the extremely competitive nature of the autonomous equipment providers in this sector. NIOSH's position as a nonregulatory government research organization will allow a synergy of coordination that will provide improved health and safety for mining generations to come. The logical steps to the research are:

Validation of Concept

Validation will necessarily be an iterative process by which we continually build confidence, identify and address gaps, and integrate the framework into successively more complicated equipment to progressively press for failure. Industry partners can help us identify scenarios that have been challenging to them to be introduced as edge cases, as well as work with NIOSH to implement confidence building, largescale testing.

Simulation

Simulation is not a substitute for real-world application, but it can potentially prevent expensive, time-consuming failures that might otherwise occur in the physical world. We are evaluating CARLA and ROS as open-source software packages to complete this work. We would expect to do a vast majority of development during this simulation cycle. The intent would be to simulate haul trucks in a environment that reflects a real-world mine.

Small-Scale Test

This would be a first step into the physical world in which we would test major software components integrated with real-world hardware. It would be a proof-of-concept effort applied in a safe, scaled environment. NIOSH/SMRD has a machine safety laboratory in Spokane, WA, that would provide an excellent opportunity to conduct such testing.

“Sand Box” Testing

It will be necessary to test in a variable environment that closely reflects the conditions found in full-scale mining, without exposing personnel to unvalidated equipment in terms of safety. This activity is yet to be planned or conceived, but one could imagine small-scale electric vehicles negotiating in closed quarters that reflects the kind of conditions and environment that could be found in most mines.

Real-World Testing

This phase would be years in the future, but early partnership with players in the mining autonomy space is critical to the eventual culmination of testing on full-scale equipment in a genuine mining environment. This supports our approach of not focusing on the specifics of this testing, but rather developing sound relationships with stakeholders that share the vision of a future where autonomy assures safety for the miners and maintenance personnel that it operates alongside.

CONCLUSION

Real-time risk assessment and determined action is only one of the critical elements of a safety intervention system, but it is the kernel of having an autonomous system that can be trusted to operate in the proximity of humans. This is not a simple or trivial problem, but it is one of the most important areas of research (in our opinion) because safety and social responsibility is ever increasingly at the core of what allows mining to provide the materials that the future will demand. As near surface and easily extracted mineral deposits are depleted, mining will be driven deeper and into harsher environments [10]. Autonomy is the only foreseeable economically viable solution to those ever more hazardous conditions for current and future mine workers.

The research discussed in this paper is an effort to identify and rectify risk before it happens, in the context of autonomous mining equipment. It will be important for NIOSH to work closely with mining industry stakeholders to further develop this framework and address safety concerns that may be introduced with autonomous equipment.

DISCLAIMER

The findings and conclusions in this paper are those of the authors and do not necessarily represent the official position of the National Institute for Occupational Safety and Health (NIOSH), Centers for Disease Control and Prevention (CDC). Mention of any company name or product does not constitute endorsement by NIOSH/CDC.

REFERENCES

[1] Computing Community Consortium - Catalyst, "Assured Autonomy Path Toward Living With Autonomous Systems We Can Trust," Computing Community Consortium - Catalyst, Washington, 2020.

[2] Automomous Systems Sub-Committee, "Guideline for the Implementation of Autonomous Systems in Mining," Global Mining Guidelines Group (GMG), 2019.

[3] A. Townshend, "Standards guide the expansion of collaborative robots," 4 April 2022. [Online]. Available: https://www.controldesign.com/motion/robotics/article/11289077/standards-guide-the-expansion-of-collaborative-robots. [Accessed 14 2 2023].

[4] L. Chang and L. Dormehl, "6 self-driving car crashes that tapped the brakes on the autonomous revolution," 22 June 2018. [Online]. Available: https://www.digitaltrends.com/cool-tech/. [Accessed 5 March 2023].

[5] T. Macpherson, M. Matsumoto, H. Gomi and et. al., "Parallel and hierarchical neural mechanisms for adaptive and," *Neural Networks,* vol. 144, pp. 507-521, 2021.

[6] A. S. Mueller, J. B. Cicchino and D. S. Zuby, "What humanlike errors do autonomous vehicles need to avoid to maximize safety?," *Journal of Safety Research*, pp. 1-4, 2020.

[7] J. Magliano, "Why Are Teen Brains Designed for Risk-taking?," Psychology Today, 9 June 2015. [Online]. Available: https://www.psychologytoday.com/. [Accessed 5 March 2023].

[8] J. Dahl, G. Rodrigues de Campos, C. Olsson and J. Fredriksson, "Collision Avoidance: A Literature Review on," *IEEE TRANSACTIONS ON INTELLIGENT VEHICLES,* vol. VOL. 4, no. NO. 1 March, pp. 101-113, 2019.

[9] W. Lihua and J. Kang-Hyun, "Deep Learning-Based Perception Systems for Autonomous Driving: A Comprehensive Survey," March 2022. [Online]. Available: https://www.researchgate.net/publication. [Accessed 5 March 2023].

[10] F. Charles, "Some Challenges of Deep Mining," *Engineering*, vol. 3, no. 4, pp. 527-537, 2017.

into consideration time periods as variables and integrating them as part of the analysis, providing a better overview of the forecast plan completion with respect to past experiences and evaluating best alternative to handle future events to minimize losses or maximize profits, resulting in better truck assignments, minimum shovel idle times and overall improved Net Present Value of the operation at the end of the time horizon.

8:20 am

Training/testing Mining Truck Drivers for Proximity Awareness Through Multiplayer Virtual Reality Game

A. Kamran Pishhesari, J. Dahl, E. Marsh, J. Sattarvand, F. Harris Jr.; University of Nevada Reno, Reno, NV, United States

Based on the mine safety and health administration (MSHA) reports, mining truck accidents are among the significant causes of fatalities in the mining industry. The major causes of haulage accidents can be divided into non-driving and driving factors, such as truck-related sudden failures and operator performance. Numerous parameters affect trucks' safe operation, but situational awareness plays a significant role. Due to the distinct structure of mining trucks, making them efficient for mining purposes, the operator has a very limited field of view (FOV), making them prone to accidents resulting in losing lives. The cabin's position blocks the truck's right and rear sides, creating several blind points around it. Aiming to test/train the operator's awareness of the surrounding, a proximity detection system is developed through a multiplayer virtual reality (VR) application to investigate the human factor in equipment collisions. The operators were examined and monitored when asked to conduct what exactly happened in the Marigold mine truck accident without being aware of the scenario. After receiving sound, vibration, and visual alarm warnings, their reaction and performance were analyzed. The platform is developed using the Unity game engine for HTC VIVE virtual reality goggles. An Omni-direction treadmill and a motion simulator platform are utilized for the player and dump truck operators. Fifteen students were monitored when driving the dump truck receiving/not receiving proximity detection alarms. None of them repeated the accident when receiving proximity alarms. Also, they found the visual alarm the most comfortable while driving the truck.

8:40 am

Evaluation of Models for Interaction Probability in Autonomous Monitor and Control Environments

R. Bissonette and S. Sbai; NIOSH, Spokane, Washington, United States

A critical component of implementing autonomous control and/or monitoring systems is realistic, human-like perception of the operational environment. This means that any such system would need to gather robust and accurate data about its environment, evaluate what objects of interest exit in that environment (with relevant data about identity, trajectory, characteristics, etc.), project a future continuum of those objects (including considerations for variabilities), and make control choices to minimize risk cost due to exposures in the present operational area. The components of such system are a) sensors and fusion b) probability projection of future states of objects and the environment, c) generation of alternate futures and d) determination of best action to reduce risk cost. The system needs to take data with different time stamps with variability parameters and create probability clouds into the near future, evaluate how those clouds interact and identify areas of risk. This paper focuses on using information from the sensor fusion engine to generate a continuous projection of future state incorporating the factors of confidence, accuracy, tolerance, and other variabilities.

9:00 am

Applications of Digital Twin Technology in Productivity Optimization of Mining Operations

J. Sattarvand and M. Ghahramanieisalou; University of Nevada Reno, Reno, Nevada, United States

As new technologies are introduced to the mining industry, the challenges related to their safe assimilation and the potential changes they bring to mining techniques necessitate adjustments to the current operations. Predictive simulations are critical to understanding unforeseen scenarios and shifting costly changes from the operational stage to the design. The Digital Twin methodology has gained a lot of attention in recent years. Many industries, from manufacturing to engineering and even social sciences, are adopting this approach to understand better the complex systems they are working with. By considering the enormous expenses and risks related to a pilot project, DT could replace that by acting as a prototype in which realistic tests and simulations are carried out with close to zero cost. DT includes a data hub, simulation and analysis tools, and visualization platforms to enable appropriate designs and monitoring plans focusing on unknown areas. A few examples of DT's application in mining include fleet management, Mine-to-Mill optimization (especially D&B), geotechnical digital twin (gDT), etc. When developing regulations, testing the efficiency of the mine layout, design, and magnitude of the restricted area and its access points, human-machine interactions, risk identification, and management processes, examining the efficacy of control systems, and monitoring, DT could help speed up (or completely replace) the process in the physical world.

smeapcom.org

SME | APCOM

Intelligent Mining: Innovation, Vision, and Value

June 25-28, 2023 | Rapid City, South Dakota