A DATA-DRIVEN APPROACH FOR OPTIMIZING DRY CLASSIFICATION GRINDING CIRCUIT EFFICIENCY

Author: Rupert Kirchner

Company: Cemtec Cement & Mining Technology GmbH

ABSTRACT

This research introduces a novel, data-driven method to optimize dry grinding classification circuits in mineral processing, leveraging advanced machine-learning techniques, specifically reinforcement learning. It pioneers the integration of digital twins, creating dynamic simulations that enhance operational efficiency, reduce energy consumption, and therefore promote sustainability. By implementing digital twins paired with reinforcement learning, the research establishes a system capable of real-time adaptation, significantly improving product quality and production rates. While this study focuses on cement clinker grinding, the developed approach demonstrates potential for adaptation and scaling across various mineral processing operations, setting a precedent for advanced control systems in the industry.

KEYWORDS

Machine Learning, Mineral Processing, Dry Grinding Classification Circuits, Operational Efficiency, Sustainable Comminution

INTRODUCTION

In the evolving landscape of the industrial sector, particularly in mineral processing, the quest for enhanced operational efficiency and technological advancement is ever-present. Key among these processes is the optimization of grinding circuit operations, critical for both production efficiency and product quality. Yet, achieving optimal performance in these circuits is a challenging endeavor due to their complexity and the dynamic interplay of variables involved.

This research explores the development of innovative methodologies and technological solutions, with a particular focus on the application of machine learning to improve dry grinding classification circuit operations across the mineral processing industry. Given the extensive application of these processes in cement production, it serves as a critical example. Cement grinding is notably energy-intensive, with electrical energy accounting for about 10% of the total energy consumption in cement production. Specifically, the electrical energy used in the cement-making process is approximately 95 to 110 kWh per ton of cement, with the clinker grinding stage alone consuming about 40% of this amount (Hosten & Fidan, 2012). Optimizing energy efficiency in such processes is not just a technical challenge but a critical environmental imperative due to the considerable energy demands and environmental implications associated with grinding processes.

Furthermore, the mineral processing industry faces challenges including the intricate nature of ore bodies, rapid shifts in market demands, and an emerging scarcity of skilled labor. These challenges often lead to reliance on increased operational safety margins, which can result in fluctuations in product quality and suboptimal plant performance. The reliance on human oversight in complex and variable processes like grinding circuits introduces potential errors and inefficiencies.

Against this backdrop, our research explores the transformative potential of data science and machine learning in mineral processing. The integration of digital twin technology and reinforcement learning in grinding circuits represents a significant technological leap. Digital twins allow for real-time monitoring and control, providing a platform for operational experimentation and optimization. When combined with reinforcement learning algorithms, these systems offer a dynamic, adaptive approach to process control, surpassing traditional methods.

The aim of this study is to create an autonomous control system for dry grinding circuits through a data-driven, machine learning approach. The research encompasses the identification of key variables, the construction of a digital twin for the grinding circuit, the design of a reinforcement-learning model, and discusses some initial thoughts about the development of an edge-computing layout for industrial implementation. In doing so, this research addresses the pressing need for energy efficiency, operational safety, and sustainability in the mineral processing industry.

By bridging the gap between theoretical research and practical application, this study not only contributes to the field of intelligent process control but also provides actionable solutions to enhance the efficiency and sustainability of mineral processing operations, leading the industry towards a more sustainable and more efficient future.

LITERATURE REVIEW

Introduction: The Necessity of Grinding Circuit Optimization

This research presents an in-depth exploration of dry grinding circuits in mineral processing, with a particular focus on a configuration that is frequently encountered in cement production. As demonstrated in Figure 1, the study concentrates on a specific setup: a two-compartment ball mill paired with a dynamic separator. This choice is based on the widespread use of this configuration in the industry and its significant potential for operational optimization.

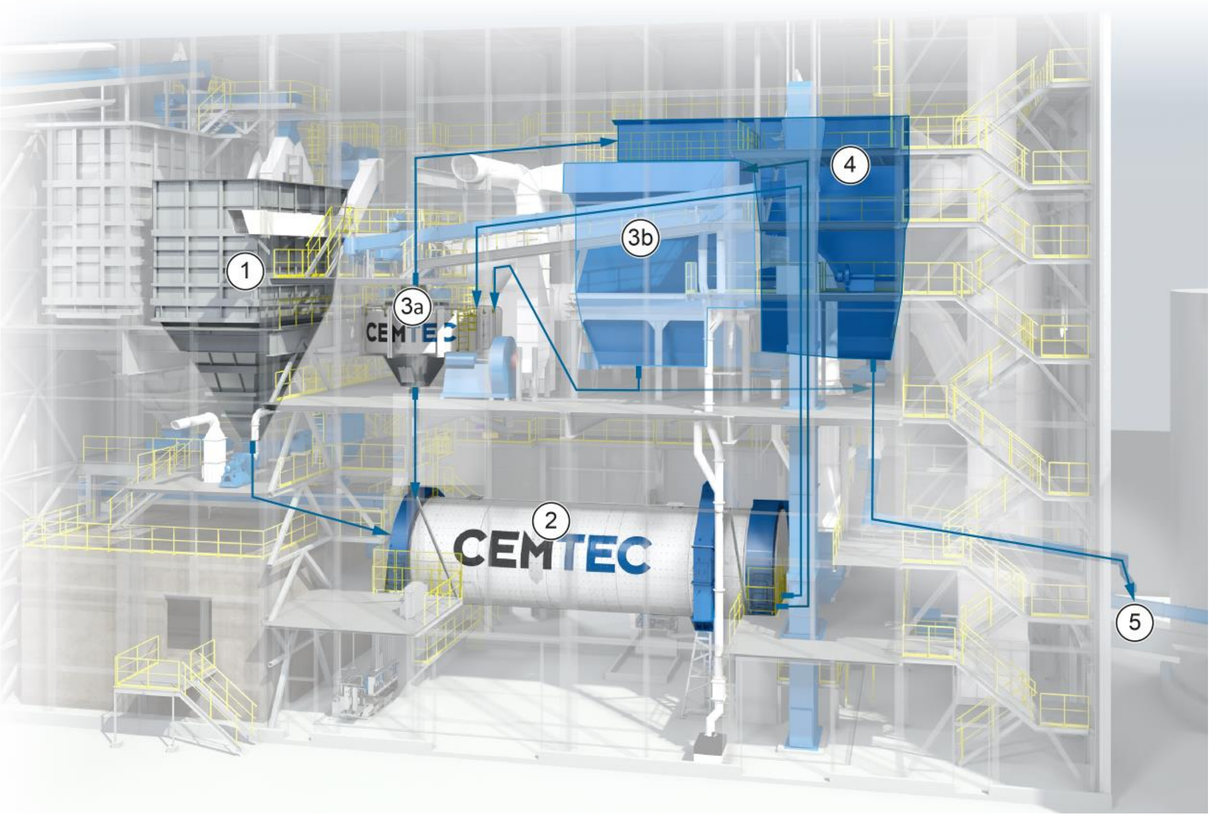

1) *Fresh feed bunker*
2) *Ball mill*
3) *a: dynamic separator; b: mill dedusting filter & fan*
4) *product filter*
5) *CEOPS particle size distribution measurement*

Figure 1: Schematic representation of a typical dry grinding circuit as considered in this research, emphasizing the specific configuration of a two-compartment ball mill paired with a dynamic separator.

The methodologies developed in this research, while tailored to this specific layout, are designed to be modular and adaptable. They could be extended to various circuit layouts in mineral processing, underlining the versatility of the proposed approach. This adaptability makes the research valuable not just for the specific grinding circuit under study but also for broader applications in mineral processing.

The study intentionally omits the analysis of pre-grinding equipment such as roller presses, which are commonly integrated into modern grinding circuits. Roller presses induce micro-fractures in the material, thereby reducing the energy required for grinding and enhancing overall energy efficiency. However, incorporating a roller press would require specific evaluations tailored to individual plants, a complexity beyond the scope of this study. This aspect of comminution technology, and its implications for energy efficiency in grinding circuits, is extensively discussed in the work of (Forssberg & Yanmin, 2003). Additionally, this research does not delve into other physical optimization methods like optimizing the liners or grinding media distribution, which can also significantly improve grinding efficiency. Instead, the focus is on optimizing operational parameters within the preexisting setup.

Evolution of Control Strategies in Grinding Circuits

The realm of control systems in grinding circuits has witnessed a significant evolution, marked by the transition from manual adjustments to more complex automated systems. This shift is reflective of the ongoing efforts to address the challenges posed by the intricate nature of grinding operations.

Proportional-Integral-Derivative (PID) Controllers

Historically, the workhorse of control systems in these circuits has been the Proportional-Integral-Derivative (PID) controllers. These systems have been extensively utilized due to their simplicity and ease of implementation. Studies by (Edwards, Vien, & Perry, 2002; Wei & Craig, 2009a) underline the widespread adoption of PID controllers in the industry. However, the inherent limitations of PID controllers, particularly their inability to adapt to the dynamic and complex nature of grinding processes, have been a point of concern.

Alternative Control Methodologies

In addition to PID controllers, there are alternative control methodologies that can be paired with them to enhance their capability, each with its unique strengths and challenges. The literature highlights several of these strategies:

- **Step Controllers:** Designed to maintain system variables within a specified range, making discrete adjustments in response to deviations from a set threshold. Their application is advantageous in systems not designed for continuous updates to the control value.
- **Fuzzy Logic Controllers:** These controllers use "If-Then" rules to manage uncertainties and approximate reasoning, suitable for complex systems where variables do not conform to strict binaries (Radha Krishna & Biswal, 2016; Zadeh, 1973).
- **Model Predictive Control (MPC):** MPC represents a significant advancement, utilizing dynamic models to predict and optimize future system behaviors (Qin & Badgwell, 2003; Rawlings & Mayne, 2009). Despite its capabilities, the computational intensity of MPC and the need for continuous model updates present practical challenges.

A comparative analysis of conventional control methods (such as PID, step and fuzzy-logic controllers) and advanced methods (like Model Predictive Control (MPC)) reveals the limitations of traditional approaches in adapting to the dynamic nature of grinding processes. The literature suggests that while conventional controllers are simple and widely used, they fall short in managing complex systems (Costea, et al., 2014; Pomerleau, Hodouin, Desbiens, & Gagnon, 2000). Advanced methods like MPC have historically faced challenges in computational intensity and model upkeep (Qin & Badgwell, 2003; Rawlings & Mayne, 2009), but recent advancements in computational hardware and machine learning technologies have begun to overcome these hurdles.

The increasing power of modern processors and the advent of state-of-the-art machine learning techniques have greatly reduced concerns around computational demands, making real-time analysis and model adjustment feasible even in complex industrial environments. Furthermore, machine learning's inherent adaptability offers a promising solution to the issue of model degradation, facilitating systems that can evolve and self-correct over time.

It is with these considerations that we recognize the value of MPC as a robust control strategy in our field, and simultaneously suggest that the integration of recent technological advancements can further enhance its effectiveness and reliability.

Shift Toward Sophisticated Strategies

Recent literature indicates a notable shift towards more sophisticated control strategies. This transition is driven by the need to address the complexity and variability inherent in grinding processes. The emergence of digitalization and data-driven technologies has paved the way for the integration of advanced machine learning algorithms and real-time analytics into these systems, heralding a new era of efficiency and adaptability in grinding circuit control (Ivezić & Petrović, 2003).

Digitalization and Automation: The Rise of Digital Twins

The concept of the digital twin originated from a NASA technology report in 2010 (Shafto, et al., 2010), defining it as an integrated multiphysics, multiscale simulation of a system that reflects its real-life counterpart. These sophisticated simulations facilitate predictive modeling and enable comprehensive testing of control strategies. Studies by (Cronrath, Aderiani, & Lennartson, 2019) and (García & Fernández, 2015) underscore the vital role of digital twins in revolutionizing behavioral control methodologies, particularly when paired with reinforcement learning. This approach has significant implications for the development of industrial controllers.

Digital twins act as a critical bridge between physical operations and digital analysis. They offer a unique advantage in understanding and managing the complex techno-socio-economic systems inherent in mineral processing. By providing a real-time, dynamic representation of physical systems, digital twins allow for enhanced decision-making and strategic planning in process control.

Machine Learning and Advanced Data Analysis

Machine learning, particularly neural networks, has emerged as a powerful tool for handling complex patterns in everyday situations (Rosenblatt, 1958; Taye, 2023). The literature reviews the application of various machine learning techniques in operations optimization, with a focus on LSTM (Long Short-Term Memory) networks and linear regression models for their ability to predict and analyze data (Hochreiter & Schmidhuber, 1997; Leonel, 2018). These methods have shown significant potential in enhancing control strategies by enabling predictive modeling and real-time analysis of processes.

Reinforcement Learning: A Paradigm Shift in Control Strategies

The application of reinforcement learning in grinding circuits represents a paradigm shift in control strategies. This machine learning technique, where an algorithm learns optimal decision-making through interactions with its environment, has been identified as a promising approach for dynamic and complex systems like grinding circuits (Sutton & Barto, 2018; Silver, et al., 2017; Conradie & Aldrich, 2001). The literature underscores its potential for autonomous control and continuous optimization of operational parameters.

Gaps and Future Directions in Grinding Circuit Control

In our investigation of this topic, we identified several gaps in current research, particularly in the practical application of advanced machine learning models for optimizing dry grinding circuits. Key challenges include the need for extensive data and computational resources, along with the integration of sophisticated models such as neural networks and reinforcement learning into real-world operations. Our review and analysis set a strong foundation for addressing these gaps, focusing on the development of scalable, efficient, and practical machine learning models for enhancing grinding circuit optimization.

METHODOLOGY FOR OPTIMIZING DRY GRINDING CLASSIFICATION CIRCUITS USING DATA-DRIVEN APPROACHES

Introduction to Methodology

The emergence of data-driven methodologies marks a new era in mineral processing, particularly in optimizing industrial processes. This project adopts a comprehensive approach, leveraging advanced data collection, machine learning, and reinforcement learning techniques, aimed at enhancing the operational efficiency and sustainability of grinding operations. This approach is in line with the study's focus on a specific dry grinding circuit configuration illustrated in the introduction (Figure 1).

Data Collection and Preliminary Analysis

The schematic below details the preliminary considerations and preparatory steps taken to ensure a thorough and effective data analysis, crucial for optimizing the grinding circuit operations.

Initial Process Analysis

A meticulous analysis was conducted to assess data quality and frequency, establishing a robust understanding of the existing grinding circuit. This critical examination laid the groundwork for future data-driven optimization strategies.

Data Types and Acquisition

Identification of key data types—measurement values, operational setpoints, and design characteristics—was crucial. This phase also covered pilot testing at Cemtec’s small-scale facility to test initial models and explore data interdependencies, setting the stage for subsequent scale-up to industrial settings.

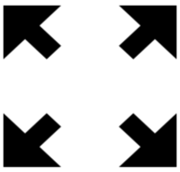

Scaling from Pilot Plant to Industrial Scale

The project initiated with pilot testing at Cemtec's facility, equipped with a grinding setup (1.2 x 3.2 meters, 35 kW drive, 1.0 tph throughput). This stage focused on developing and refining initial models, assessing data interdependencies under controlled conditions. Following pilot success, the study expanded to three industrial-scale plants in the cement and minerals sector to test model scalability and robustness. This expansion aims to validate models across diverse operational environments, ensuring adaptability and effectiveness.

Addressing Data Gaps

Recognizing infrequent analysis of particle size in industrial setups, the decision was made to implement the CEOPS (Cemtec Online Particle Measurement System). This allowed for real-time measurements and significantly enhanced dynamic process modeling capabilities.

Development of Data Collection Infrastructure

Edge Device Selection

An industrial-grade OnLogic Helix 511 fanless edge-PC, equipped with an Intel Core i5 processor and 32 GB of RAM, was integrated into the plant’s control system. This edge computing device, running Ubuntu Desktop 22.04 LTS and utilizing the OPC-UA communication protocol, served as a central repository for all operational data, crucial for the modeling process. It was selected for its capability to support advanced machine learning operations and potential future control algorithm deployment, ensuring seamless and secure data handling.

Data Integrity and Frequency

Data collection integrity and frequency were optimized by setting a one-minute sampling interval. This interval was chosen to balance the need for timely data against potential noise and fluctuations in measurements, with data being averaged over this period to mitigate short-term variability. This setup ensures the data reflects the dynamic nature of the grinding process accurately and reliably, supporting robust modeling and analysis.

Data Preparation and Feature Engineering

To ensure the effectiveness of the machine learning models, the data preparation and feature engineering phases were meticulously structured. These phases included:

- **Data Merging and Normalization:** Conducted a comprehensive merging and normalization process to align measurement values and operational setpoints from various sources, standardizing data formats and scales for analytical coherence.
- **Feature Selection and Cleaning:** Identified and selected key features impacting the grinding process, rigorously cleaning the dataset to remove inconsistencies and applying action smoothing techniques for real-time operational stability.
- **Feature Engineering:** Employed advanced techniques to transform raw data into formats optimized for machine learning, enhancing predictive power and capturing the complex dynamics of the grinding circuit.

Model Development and Validation

Model Segmentation

The analysis of the grinding circuit highlighted the feasibility of independently modeling the separator, given its non-limiting design and size. This insight led to a two-segment modeling strategy:

- **Dynamic Separator Model:** Utilizes LSTM networks combined with CNN layers to predict the particle size distribution (specifically the 80% passing size) based on operational parameters. This model effectively captures both time-series dependencies and spatial data, ensuring accurate predictions of the separator's performance.
- **Ball Mill Model:** This linear regression model predicts the reject mass flow rate, taking into account the time-dependent dynamics of the process. Key to this model is the inclusion of a time lag component that represents the delay between changes in operational setpoints (like feed rate and mill power) and their effects on output. This component is calibrated using the average material retention time in the mill, allowing the model to account for the gradual impacts of operational adjustments on the product size distribution.

This strategic segmentation enhances the specificity and efficacy of the analysis, ensuring a comprehensive understanding of each component's impact within the circuit.

Validation and Testing

The data-based digital twin models were rigorously evaluated using historical data, previously unseen during the training phase. For this purpose, the collected data was split into two sets: 80% for training and 20% for testing. This split allowed for a comprehensive assessment of the models' performance, ensuring they were tested against data reflecting various operational scenarios. This validation process was critical to ascertain the accuracy and reliability of the models in simulating real-world scenarios and to refine them for enhanced predictive capabilities.

Reinforcement Learning for Optimization

The development of a reinforcement learning environment is a critical step in training a reinforcement learning agent within a safe and controlled simulation. This environment defines the actions and observations within the grinding process, as well as a reward system to guide the agent towards process optimization. It also includes constraints to inform the agent of the limits of its actions, such as the maximum reject mass flow rate that can be conveyed within the circuit. A key aspect of this environment is the integration of a databased digital twin, modeled to replicate the grinding circuit's dynamics. This integration not only allows the reinforcement learning algorithm to interact, learn, and adapt based on simulated feedback but also ensures that the training is grounded in a realistic representation of the industrial process. The digital twin serves as a dynamic and high-fidelity model, providing the reinforcement learning agent with a rich, simulated context in which to develop its decision-making capabilities.

Simulation Environment Development

The reinforcement learning environment was created using Farama's Gymnasium library (Towers, et al., 2023), replicating the dynamics of the grinding circuit. This simulated environment is essential for training the reinforcement learning agent before deploying it in real-world systems.

Definition of Action and Observation Spaces

Both action and observation spaces in our approach are continuous, defining the range of possible interactions and feedback within the process.

- **Action Space:** Includes setpoints such as fresh material feed rate, mill dedusting fan speed, and separator fan and cage speeds.
- **Observation Space:** Encompasses feedback values from the process, including measured and target product size, mill power, feed material mixture, and reject material flow rate.

Reward Function Design

The reward function in reinforcement learning is fundamental for guiding the agent's decisions. For controlling a dry grinding circuit, we devised a reward function segmented into three objectives, each targeting a specific operational goal:

- **Regulation of Target Product Size:** The primary goal is to maintain particle size as close as possible to a predefined target value. A hysteresis value (typically 1 µm) creates a tolerance range around the target. The reward is calculated as 1,000 times the inverse of the absolute difference between observed and target product size, penalizing deviations outside the hysteresis range.
- **Maximizing Throughput:** The second goal is to maximize throughput while maintaining quality at the targeted level. The reward increases linearly with every added ton of throughput, penalizing any amount exceeding the maximum feed rate.
- **Reject Amount Constraint:** The third objective limits the maximum level of the reject amount, penalizing excesses to maintain material capacity within specific plant areas, like the dynamic separator or bucket elevator.

The cumulative reward combines these components, aligning them with the specific requirements of the grinding circuit. This alignment drives the reinforcement learning algorithm towards solutions that are efficient and compliant with operational constraints, mirroring the intended control strategies for the system.

Training and Evaluating Reinforcement Learning Algorithms

In this study, we utilized the open-source reinforcement learning library Stable-Baselines 3 (Raffin, et al., 2021). This library simplifies the implementation and testing of various reinforcement learning setups, providing a robust platform for our experiments.

Algorithm Selection and Parameterization

A critical step was the selection of reinforcement learning algorithms suitable for grinding circuit control. We evaluated both on-policy algorithms like Proximal Policy Optimization (PPO) and Actor-Critic (A2C), and off-policy algorithms including Soft Actor-Critic (SAC), Twin Delayed DDPG (TD3), and Deep Deterministic Policy Gradient (DDPG). This diverse range enabled a comprehensive comparison to determine the most effective algorithms for navigating the continuous action space and addressing the complexities of the circuit.

Reinforcement Learning Training Process

Algorithms underwent training within a digital twin environment, facing various simulated scenarios to refine strategies and optimize operational setpoints. The iterative learning process was guided by a reward function designed to align with key operational goals: regulating particle size, maximizing throughput, and controlling reject rates. This approach ensured that algorithms could adapt effectively to the dynamic conditions of the grinding process.

Evaluation of Learning Performance

The efficacy of each algorithm was assessed by its ability to meet control objectives within operational constraints. Performance evaluations focused on stability, efficiency, and adaptability to changes, using these metrics to identify strengths and weaknesses. Insights from this comprehensive evaluation aided in selecting the most suitable algorithms for practical deployment in the grinding circuit.

Further Evaluation in a Simulated Industrial Environment

This phase involved developing a sophisticated simulation environment based on virtual replicas of industrial-scale plants to replicate real-time operational conditions. The core objective was to integrate the refined reinforcement learning models into this environment, enabling their interaction and adaptation to various operational scenarios. The efficacy of the control strategies was rigorously tested by monitoring key performance indicators such as energy consumption, product quality, and operational stability. This evaluation phase was crucial for verifying the robustness and adaptability of the control strategies under diverse and changing conditions, preparing the models for real-world deployment.

RESULTS AND DISCUSSION

Digital Twin Model Predictions

The development and evaluation of digital twin models formed the cornerstone of this research, as detailed in our methodology. The process involved a systematic and iterative approach to model generation and selection. Two critical models were developed: one predicting the product size of the grind and the other estimating the reject rate. The selection of these models for the reinforcement learning phase was based on their performance in test datasets, particularly their ability to generalize predictions accurately.

To evaluate the efficacy of these models, a comparative analysis was performed. Predictions from the digital twin models were compared against actual measurements from a 1,000-minute operational period that included several recipe changes. This period was chosen deliberately to challenge the models under varying operational conditions, thus providing a comprehensive assessment of their predictive capabilities.

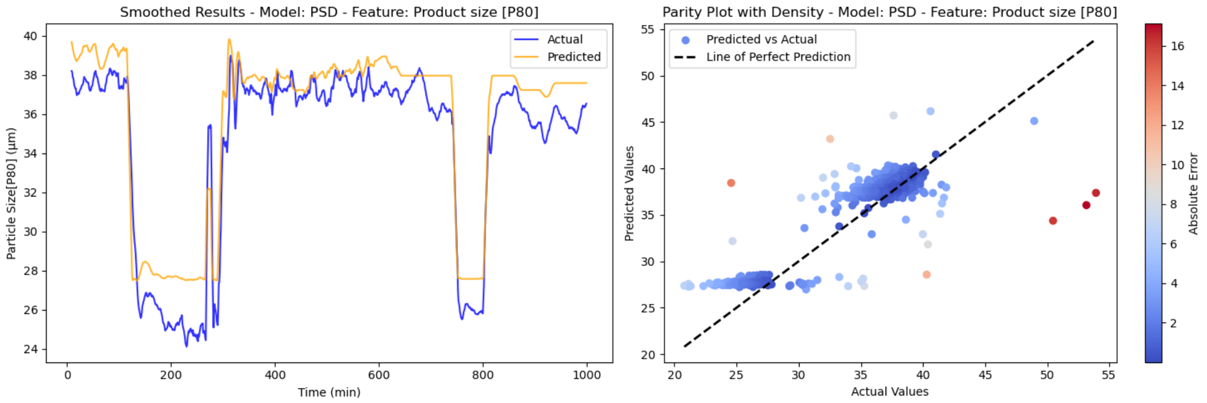

Figure 2: Comparison between the actual values and the prediction values of the digital twin for product size

Figure 3: Comparison between the actual values and the prediction values of the digital twin for reject amount

Figure 2 and Figure 3 illustrate a side-by-side comparison of the predicted values against the actual data for product size and reject rate. Despite minor discrepancies noted in the product size predictions over extended periods, the models demonstrated high accuracy for short-term predictions. This finding is particularly significant as it highlights the models' utility for real-time control applications, aligning with the goals set forth in the introduction and methodology chapters.

Prediction Metrics and Model Evaluation

The model evaluation utilized both individual and composite metrics (0.5 × Mean Squared Error + 0.5 × Standard Deviation). After 20 training iterations, for the product size model, a Mean Squared Error (MSE) of 5.31 was noted, with a Standard Deviation (STD) of 1.94 representing the variability in prediction accuracy compared to actual sensor readings. This resulted in a composite score of 3.63. In contrast, the reject amount model displayed an MSE of 71.81 and an STD of 8.47 against sensor data measurements, culminating in a higher composite score of 40.14. The greater standard deviation in the reject amount model likely reflects the generally higher inaccuracy of the sensors used to measure reject amounts. This factor should be considered when evaluating the prediction errors for this model. The training duration for each model, completed in just 36 minutes on the industrial edge device, underscores the practicality of these models for industrial applications.

The comprehensive evaluation of the digital twin models demonstrated their accuracy in predicting actual operational values within the plant. This accuracy is crucial for creating a reliable representation of the plant's behavior within a reinforcement learning architecture. The results affirm the potential of using digital twins as an integral part of an intelligent process control system in mineral processing, laying the groundwork for their effective incorporation in real-world industrial applications.

Reinforcement Learning Results

Policy Evaluation Results

The training of reinforcement learning algorithms was an intensive process, conducted over 400,000 epochs for on-policy algorithms and 200,000 epochs for off-policy methods. On-policy reinforcement learning algorithms, such as Proximal Policy Optimization (PPO), directly learn from and improve the policy that is used to make decisions, meaning they evaluate and improve the same policy that determines the action. In contrast, off-policy methods like Deep Deterministic Policy Gradient (DDPG) learn a policy different from the one used to generate the data. This allows off-policy methods to learn from past experiences stored in a replay buffer, potentially making them more data-efficient as they can reuse this information for multiple updates. On the edge device, chosen for its suitability for industrial application, the training of each model took approximately 1 to 1.25 hours. All models performed commendably, achieving mean rewards as high as 1,043, which is significant when compared to the maximum achievable reward of 1,400. Whereas this maximum reward is a theoretically calculated value, representing the ideal scenario where the particle size consistently meets the target size, the mill operates at the maximum allowable feed rate, and the amount of reject material remains within acceptable limits. This value serves as a benchmark for evaluating the performance of the reinforcement learning algorithms against the best possible outcome in controlled conditions. Figure 4 exemplifies these results, showcasing the SAC algorithm's training outcomes as a representative sample of all models.

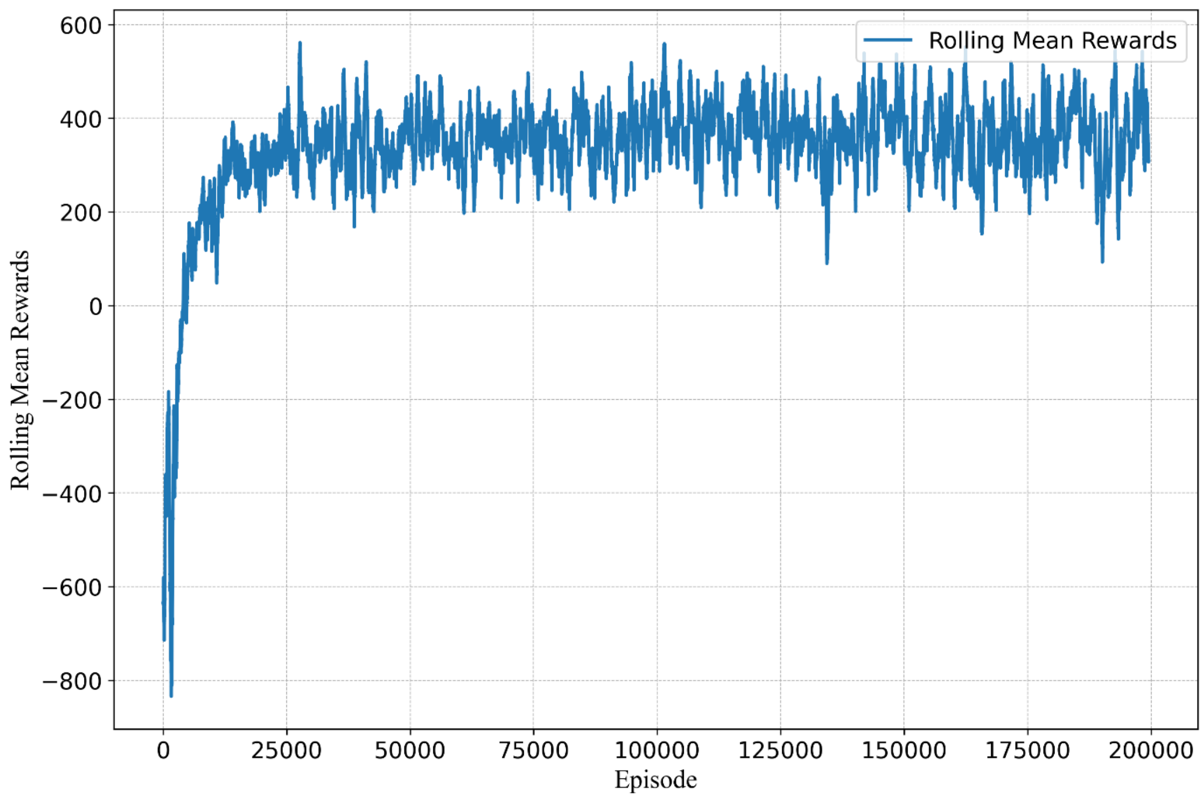

Figure 4: Reward curve for Soft Actor-Critic over 200,000 epochs, illustrating an initial surge and logarithmic growth, culminating in lower final rewards

The policy evaluation results are summarized in Table 1, highlighting the training time, mean reward, standard deviation, and overall score (0.8 × Mean Reward - 0.2 × Standard Deviation Reward) for each algorithm:

Table 1: Policy evaluation results

Algorithm	Training Time	Mean Reward ± Std Reward	Score
PPO	1hr 16min	816,739 ± 35,205	646,350
A2C	1hr 15min	780,319 ± 57,866	612,682
DDPG	1hr	1,019,611 ± 20,606	811,568
SAC	1hr 16min	1,027,977 ± 28,815	816,618
TD3	51min	966,746 ± 47,008	763,995

During training, every 50 timesteps (equivalent to 50 minutes), the operational values, akin to different recipes, were altered to test the flexibility of the control algorithms. The adaptability and efficiency of these algorithms, particularly in responding to changes in operational targets, are exemplified in Figure 5, which contrasts the target product size with the actual product size achieved.

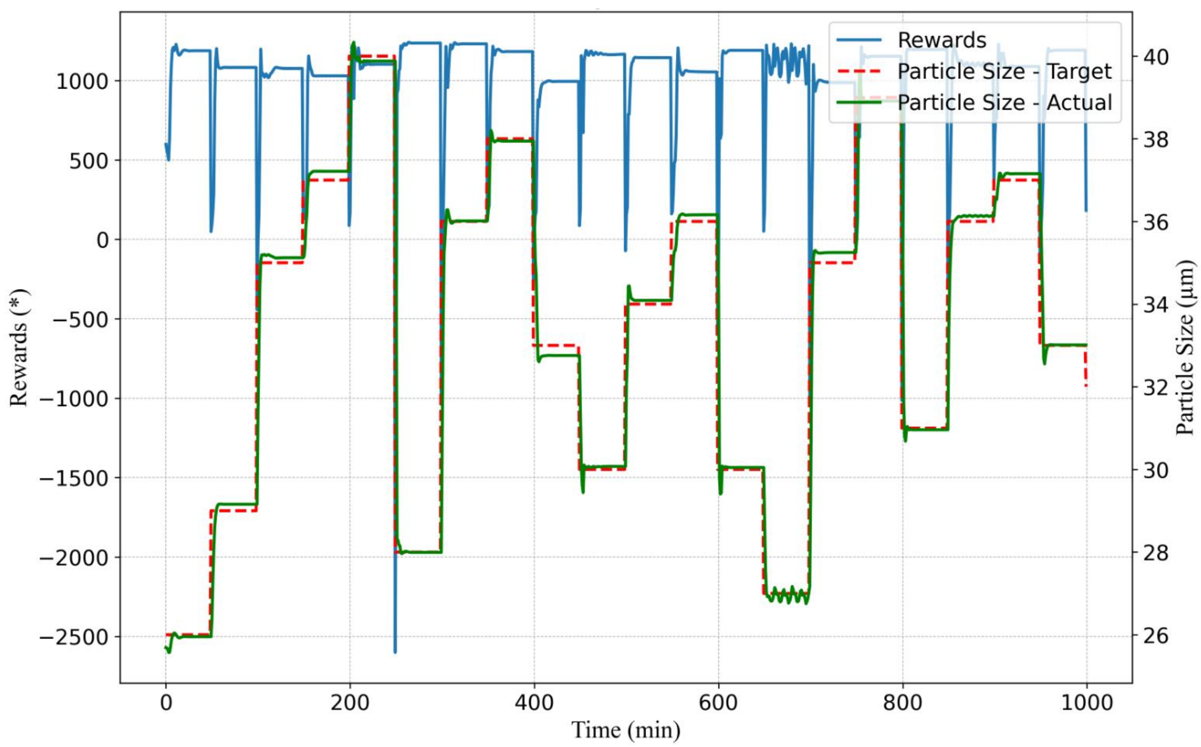

Figure 5: Dynamic Reward Curve and Regulation Quality for SAC Algorithm

Simulation Environment Results

In a customized simulation environment replicating industrial-scale plants, the Soft Actor-Critic (SAC) algorithm showcased the most promising results, particularly in terms of flexibility. One notable example from our trials involved a dramatic shift in target product size from 35 µm to 25 µm. The SAC algorithm responded swiftly, stabilizing new setpoints within 10 minutes, and aligning closely with expectations for real-world industrial applications. Within another 10 minutes, adjustments in feed rate and mill fan speed were optimized. Notably, the recirculating load increased from 55% to 120% during this change, consistent with theoretical expectations and practical observations for achieving a finer product size. Moreover, the feed rate to the mill was carefully modulated to keep the reject rate below the circuit's maximum capacity of 120 tph.

This efficient modulation of operational parameters by the SAC algorithm emphasizes its capability to swiftly adapt to new observations, leading to two significant benefits in the operation of grinding plants: improved product quality and enhanced energy efficiency. These benefits stem directly from the reinforcement agent's reward function, which, as outlined in the methodology chapter, is designed to constantly optimize these objectives. By efficiently adjusting to changes in operational conditions, the SAC algorithm ensures that product quality remains consistently high, regardless of variations in the grinding process. This adaptability is crucial for meeting stringent quality standards and responding to dynamic market demands.

Additionally, the focus on maximizing energy efficiency reflects a key aspect of sustainable mineral processing. The algorithm's ability to operate close to the circuit's limits, while maintaining optimal performance, demonstrates a significant advancement in reducing energy consumption. This not only aligns with the overarching goal of sustainable operations but also offers substantial cost savings and environmental benefits.

These results, visualized in Figure 6, confirm the efficacy of the reinforcement learning approach in enhancing grinding circuit operations. The SAC algorithm sets a benchmark for future implementations and research in the field, offering a data-driven pathway to higher product quality and energy efficiency in mineral processing.

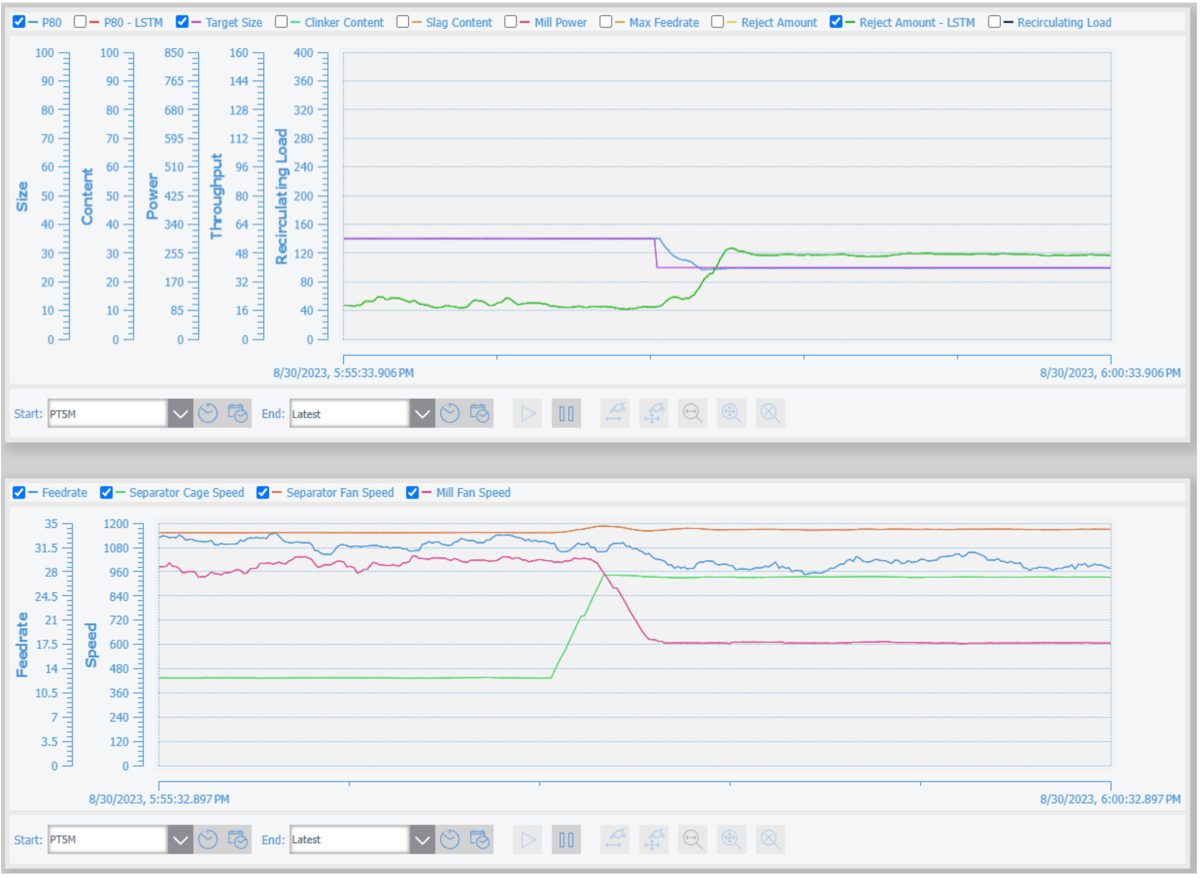

Figure 6: Trends showcasing changes in recirculating load and other operational parameters during a shift in target product size.

CONCLUSION AND FUTURE WORK

The research journey detailed in this research has been a transformative exploration into the autonomous control of dry grinding circuits in mineral processing, grounded in data-driven machine learning techniques. The research has successfully investigated complex technological paradigms and formulated innovative strategies and methodologies that are applicable in industrial settings.

Summary of Findings

The core findings of this research demonstrate the significant potential of online reinforcement learning, a branch of machine learning, for intelligent control in industrial processes:

- **Identification of Key Variables:** The research successfully identified critical operational parameters and control features, laying a robust foundation for subsequent modeling and control strategies.
- **Digital Twin Development:** The creation of a databased digital twin has been pivotal. This high-fidelity simulation mirrors the intricate dynamics of grinding processes, offering a platform for predictive analysis and operational experimentation without real-world interruptions.

- **Training Environment for Reinforcement Learning:** A specialized training environment was developed, tailored to the unique requirements of the grinding circuit. This environment's adaptability and scalability highlight its capability to handle diverse operational data effectively.
- **Training of Reinforcement Learning Models:** The research meticulously trained and evaluated both on-policy and off-policy reinforcement learning algorithms, which demonstrated high efficacy and robustness, indicating their substantial potential for real-world applications.

These contributions not only advance the field of mineral processing but also set a precedent for future research and development in intelligent industrial control systems. The exploration of advanced machine learning techniques, particularly in autonomous and adaptive control, opens up new avenues for innovation in mineral processing and related industrial applications.

Future Work

Looking ahead, the study outlines several promising avenues for further exploration and practical application:

- **Implementation in Industrial Scale Plant:** The immediate next step involves deploying the developed algorithm in a full-scale industrial plant. This crucial phase will validate the algorithm's practical efficacy and reliability in a real-world setting, providing valuable insights for further refinement.
- **Transfer Learning:** Investigating transfer learning to enhance data efficiency and model robustness. This approach facilitates rapid deployment of intelligent control systems across different plant setups, adapting previously learned knowledge to new environments with minimal need for retraining.
- **Grey-Box Models:** Exploring the integration of grey-box models, which combine physical laws with data-driven insights, offers a more flexible and efficient approach to modeling of digital twins. Such models can provide a deeper understanding of the underlying processes while still leveraging the power of machine learning.
- **Hybrid Implementation Strategy:** Developing a hybrid approach that merges cloud-based learning with localized edge computing. This strategy aims to enhance inter-plant collaboration and provide tailored solutions to each plant’s unique requirements, balancing centralized insights with local operational needs.

In summary, this research not only marks a significant advancement in the field of intelligent manufacturing but also sets a solid foundation for the future of mineral processing. The continued exploration and application of these innovative strategies promise to bring about substantial improvements in efficiency, sustainability, and adaptability within the industry. Emphasizing environmental impact, this approach has the potential to significantly reduce energy consumption and costs, contributing to more sustainable practices in mineral processing. The journey into these new territories is expected to yield novel solutions and deeper insights into complex industrial systems, further driving progress, and operational efficiency in the mineral processing sector.

REFERENCES

Conradie, A. V., & Aldrich, C. (2001, October). Neurocontrol of a ball mill grinding circuit using evolutionary reinforcement learning. *Minerals Engineering, 14*, 1277–1294. doi:10.1016/s0892-6875(01)00144-3

Costea, C. R., Silaghi, H., Gergely, E. I., Husi, G., Coroiu, L., & Nagy, Z. (2014, November). Approach of PID controller tuning for ball mill. *2014 International Symposium on Fundamentals of Electrical Engineering (ISFEE).* IEEE. doi:10.1109/isfee.2014.7050606

Cronrath, C., Aderiani, A. R., & Lennartson, B. (2019, August). Enhancing Digital Twins through Reinforcement Learning. *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE).* IEEE. doi:10.1109/coase.2019.8842888

Edwards, R., Vien, A., & Perry, R. (2002). Strategies for the instrumentation and control of grinding circuits. In A. L. Mular, D. N. Halbe, & D. J. Barratt (Eds.), *Mineral Processing Plant Design, Practice, and Control. Society for Mining Metallurgy and Exploration* (pp. 2130–2151).

Forssberg, E., & Yanmin, W. (2003). International overview and outlook on comminution technology. Retrieved from https://www.semanticscholar.org/paper/19f7ec5489308b20881de1fbbf07b078a105f5d1

García, J., & Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. doi:10.5555/2789272.2886795

Hochreiter, S., & Schmidhuber, J. (1997, November). Long Short-Term Memory. *Neural Computation, 9*, 1735–1780. doi:10.1162/neco.1997.9.8.1735

Hosten, C., & Fidan, B. (2012, May). An industrial comparative study of cement clinker grinding systems regarding the specific energy consumption and cement properties. *Powder Technology*, pp. 183-188. doi:https://doi.org/10.1016/j.powtec.2011.12.065

Ivezić, D. D., & Petrović, T. B. (2003, June). New approach to milling circuit control—robust inverse Nyquist array design. *International Journal of Mineral Processing, 70*, 171–182. doi:10.1016/s0301-7516(03)00022-x

Leonel, J. (2018). Perceptrons. *Perceptrons*. Retrieved from https://medium.com/@jorgesleonel/perceptrons-a22fb29facc4

Pomerleau, A., Hodouin, D., Desbiens, A., & Gagnon, É. (2000, March). A survey of grinding circuit control methods: from decentralized PID controllers to multivariable predictive controllers. *Powder Technology, 108*, 103–115. doi:10.1016/s0032-5910(99)00207-7

Qin, S. J., & Badgwell, T. A. (2003, July). A survey of industrial model predictive control technology. *Control Engineering Practice, 11*, 733–764. doi:10.1016/s0967-0661(02)00186-7

Radha Krishna, T. Y., & Biswal, K. (2016). Speed Control of Separately Excited DC Motor using Fuzzy Logic Controller. Bachelor's thesis, Electrical Engineering. (unpublished). *Speed Control of Separately Excited DC Motor using Fuzzy Logic Controller. Bachelor's thesis, Electrical Engineering. (unpublished)*. Rourkela, Odisha.

Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., & Dormann, N. (2021). Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*. Retrieved from https://www.semanticscholar.org/paper/e3fc5b5627af62ee6981a02090cf6bae368202d7

Rawlings, J. B., & Mayne, D. Q. (2009). *Model predictive control theory and design.* Nob Hill Pub.

Rosenblatt, F. (1958). Two Theorems of Statistical Separability in the Perceptron: (Project Para) Contract Nonr-2381(00). *Two Theorems of Statistical Separability in the Perceptron: (Project Para) Contract Nonr-2381(00)*.

Shafto, M., Conroy, M., Doyle, R., Glaessgen, E., Kemp, C., Lemoigne, J., & Wang, L. (2010, November). *Modeling, Simulation, Information Technology & Processing Roadmap - Technology Area 11.* Tech. rep., National Aeronautics and Space Administration. Retrieved from https://www.nasa.gov/pdf/501321main_TA11-MSITP-DRAFT-Nov2010-A1.pdf

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., . . . Hassabis, D. (2017, October). Mastering the game of Go without human knowledge. *Nature, 550*, 354–359. doi:10.1038/nature24270

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning An Introduction.* A Bradford Book.

Taye, M. M. (2023, April). Understanding of Machine Learning with Deep Learning: Architectures, Workflow, Applications and Future Directions. *Computers, 12*, 91. doi:10.3390/computers12050091

Towers, M., Terry, J. K., Kwiatkowski, A., Balis, J. U., De Cola, G., Deleu, T., . . . Younis, O. G. (2023, July). Gymnasium. *Gymnasium*. Zenodo. doi:10.5281/ZENODO.8127026

Wei, D., & Craig, I. K. (2009a, August). Economic performance assessment of two ROM ore milling circuit controllers. *Minerals Engineering, 22*, 826–839. doi:10.1016/j.mineng.2009.02.013

Zadeh, L. A. (1973). Outline of a New Approach to the Analysis of Complex Systems and Decision Processes. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-3*, 28–44. doi:10.1109/tsmc.1973.5408575