Physicochem. Probl. Miner. Process., 62(5), 2026, 235702

http://www.journalssystem.com/ppmp

Physicochemical Problems of Mineral Processing

ISSN 1643-1049

Data analytics and machine learning for digital transformation in mineral processing: A conceptual framework for copper-molybdenum concentration

Oğuzhan Mert Gürkan, Hüseyin Baştürkcü

Istanbul Technical University, Faculty of Mines, Department of Mineral Processing Engineering, 34469, Maslak, Istanbul, Türkiye

Corresponding author: gurkanog@itu.edu.tr (Oğuzhan Mert Gürkan)

Abstract: The mineral processing industry faces mounting pressure to improve operational efficiency, reduce energy consumption, and enhance equipment reliability through digital transformation. This study presents a conceptual framework for data analytics and machine learning implementation across three operational domains of a hypothetical copper-molybdenum concentrator: grinding optimization, flotation performance prediction, and predictive maintenance. All performance figures represent projected outcomes grounded in published benchmarks rather than measured results from an operating plant. Four algorithms were developed and evaluated, each matched to the characteristics of its process domain. A Random Forest model was developed for P80 particle size prediction, achieving $R^2 = 0.92$ on the validation dataset, with projected reductions in specific energy consumption of 9.5%. An Artificial Neural Network was implemented for copper recovery soft sensing, achieving 94% prediction accuracy within ±2% of assay results, with projected reductions in recovery excursion duration through earlier intervention. A Convolutional Neural Network based on MobileNetV2 transfer learning achieved 91.3% accuracy in froth image classification across four operational states. A Long Short-Term Memory network for SAG mill bearing health monitoring achieved 87% probability of detection for degradation events 24-48 hours in advance, with projected reductions in unplanned downtime of ~40% and emergency maintenance events of ~70% relative to reactive baselines in the literature. Cross-domain synergies between the integrated models are expected to generate additional value beyond individual contributions. This paper discusses key implementation challenges, including data quality, model retraining, operator acceptance, and edge computing. The projected improvements provide a literature-grounded roadmap for digital transformation in mineral processing.

***Keywords*:** data analytics, machine learning, mineral processing, grinding circuit optimization, predictive maintenance

1. Introduction

The mineral processing industry is undergoing a profound digital transformation driven by converging pressures: declining ore grades, rising energy costs, increasingly stringent environmental regulations, and growing operational complexity. Traditional process control strategies, largely dependent on operator experience, empirical heuristics, and simple feedback loops, are proving inadequate for modern concentrators where hundreds of variables interact simultaneously in highly nonlinear ways (McCoy and Auret, 2019). In this context, data analytics and machine learning (ML) have emerged as transformative technologies that can extract actionable knowledge from the vast operational data streams generated by contemporary mineral processing plants.

The scale of data generation in modern concentrators is substantial. Semi-autogenous grinding (SAG) mills alone can produce thousands of data points per minute from power draw sensors, bearing vibration monitors, acoustic emission detectors, and mill load indicators. Flotation circuits equipped with machine vision systems, online X-ray fluorescence (XRF) analyzers, and distributed sensor arrays

DOI: 10.37190/ppmp/235702

generate continuous, high-dimensional data streams describing froth morphology, pulp chemistry, and metallurgical performance (Aldrich et al., 2010; Liu and Aldrich, 2022). However, the fundamental challenge has shifted from data collection to knowledge extraction: converting raw operational data into predictive models and optimization strategies that deliver measurable improvements in efficiency, recovery, and equipment reliability.

Machine learning applications in mineral processing have advanced considerably over the past decade, moving from isolated proof-of-concept demonstrations toward integrated, plant-wide implementations. Reviews of the field document successful applications spanning comminution circuit optimization, flotation performance prediction, ore characterization, and predictive maintenance (McCoy and Auret, 2019; Estay et al., 2023; Szmigiel et al., 2024). Ensemble methods such as Random Forest and gradient boosting have demonstrated strong predictive performance for tabular process data, while deep learning architectures including Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks have unlocked new capabilities in image-based process monitoring and time-series forecasting, respectively (Duijvenbode et al., 2020; Jooshaki et al., 2021).

Despite this progress, several critical gaps limit the broader industrial adoption of ML in mineral processing. First, most published studies address individual process units in isolation, without examining how predictive models for grinding, flotation, and maintenance interact within an integrated plant-wide framework. Second, most investigations focus on model development and accuracy reporting, with limited attention to the practical challenges of deployment, including data quality assurance, operator acceptance, model drift, and integration with existing distributed control systems (DCS) and SCADA infrastructure (Nad et al., 2022; Estay et al., 2023). Third, the specific combination of sensor fusion, multi-algorithm implementation, and closed-loop optimization across an entire concentrator circuit has rarely been examined in a unified framework, despite representing the most realistic pathway to realizing the full potential of digital transformation in mineral processing.

The digitalization of mineral processing operations aligns with broader Industry 4.0 principles, where cyber-physical systems, real-time data analytics, and autonomous decision-making converge to create intelligent manufacturing environments (Nad et al., 2022). For the mining sector, this transformation offers particular value given the capital-intensive nature of processing equipment, the high energy intensity of comminution operations, which typically account for 30-50% of total plant energy consumption (Roux and Craig, 2019), and the direct relationship between process stability and metallurgical recovery (Hodouin et al., 2001).

This study presents a comprehensive conceptual case study examining the integrated implementation of a data analytics platform across all major operational domains of a medium-scale copper-molybdenum concentrator. The scenario draws on published performance benchmarks and realistic operational parameters to construct a representative implementation framework. The specific objectives are to: (1) demonstrate the complementary application of multiple ML algorithms, Random Forest, Artificial Neural Networks, Convolutional Neural Networks, and Long Short-Term Memory networks, matched to the specific characteristics of different process units; (2) evaluate the performance improvements achievable through integrated grinding circuit optimization and flotation performance prediction; (3) assess the capability of deep learning approaches for predictive maintenance of critical comminution equipment; and (4) identify and discuss the practical implementation challenges that must be addressed for successful industrial deployment of these technologies. By examining these objectives within a unified plant-wide framework, this work provides a technical roadmap and a realistic assessment of the opportunities and limitations of digital transformation in mineral processing.

2. Materials and methods

2.1. Case study plant description

The conceptual framework developed in this study is based on a hypothetical medium-scale copper-molybdenum concentrator with a nominal processing capacity of 10,000 tonnes per day, representative of operations commonly found in porphyry copper deposits in major mining regions. The plant processes porphyry copper ore with an average head grade of 0.65% Cu and 0.015% Mo, characteristic of this deposit type and within the range reported for comparable industrial operations (Saldaña et al., 2023). Fig. 1 illustrates the overall process flowsheet.

The comminution circuit consists of primary jaw crushing followed by a semi-autogenous grinding (SAG) mill (36 ft diameter × 18 ft length) operating in closed circuit with a battery of hydrocyclones, and a secondary ball mill (24 ft diameter × 40 ft length) for further size reduction. The circuit is designed to achieve a target product size of P80 = 150 µm prior to flotation. SAG mill operating parameters include a ball charge volume of 10-12%, mill speed of 72-78% of critical speed, and fresh feed rates ranging from 380 to 450 t/h depending on ore hardness characteristics (Parian et al., 2021).

The flotation circuit comprises a rougher stage (8 cells × 160 m^3), a scavenger stage (6 cells × 160 m^3), and three successive cleaner stages to upgrade copper concentrate, followed by a dedicated molybdenum circuit for copper-molybdenum separation using sodium hydrosulfide (NaHS) as a copper depressant. Rougher flotation feed is conditioned with potassium amyl xanthate (PAX) as the primary collector at dosages of 25-35 g/t, methyl isobutyl carbinol (MIBC) as frother at 10-15 g/t, and lime for pH control at 10.5-11.0. The target copper concentrate grade is 28-30% Cu with an overall copper recovery of 88-92%. The molybdenum separation circuit, while included in the plant description for completeness, lies outside the scope of the machine learning models developed in this study; the high sensitivity of NaHS-based Cu-Mo separation to reagent dosage and pulp potential makes it a distinct modeling challenge warranting dedicated future investigation (Gholami et al., 2020).

The existing plant automation infrastructure consists of a distributed control system (DCS) for basic regulatory control of flows, levels, and temperatures; programmable logic controllers (PLCs) for equipment interlocks and safety systems; and a supervisory control and data acquisition (SCADA) system providing operator interfaces and historical data logging. This baseline automation infrastructure is representative of modern concentrators and provides the foundation upon which the data analytics layer described in this study is implemented.

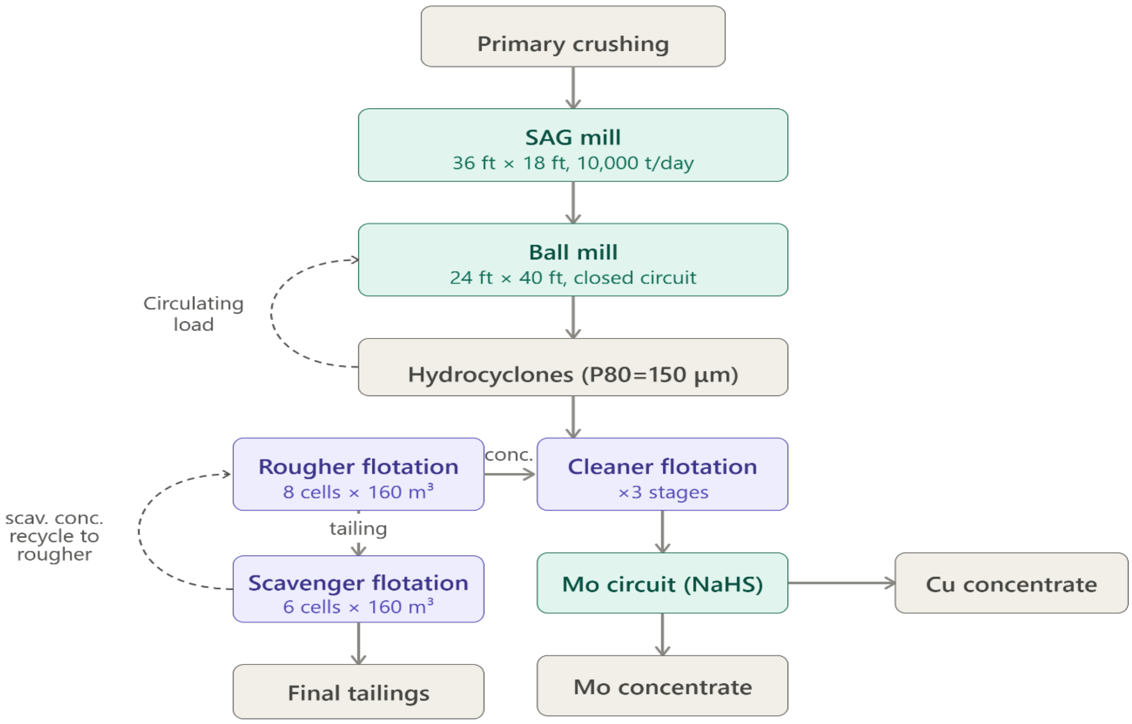

Fig. 1. Simplified process flowsheet of the hypothetical copper-molybdenum concentrator

2.2. Sensor network and data acquisition

A comprehensive sensor network was designed and deployed across all major process units to enable continuous, high-resolution data collection for model development and real-time optimization. The instrumentation strategy prioritized measurement of variables with established relationships to key performance indicators, while ensuring redundancy for critical measurements and compatibility with the existing DCS infrastructure. The grinding circuit instrumentation encompasses the following measurement systems: SAG mill power draw monitoring via calibrated power transducers sampling at

1-second intervals; bearing temperature and vibration sensors installed on all rotating equipment including SAG mill, ball mill, and hydrocyclone feed pumps, with vibration measurements capturing both overall RMS values and frequency-resolved spectra through fast Fourier transform (FFT) analysis up to 10 kHz; mill load and charge position sensors using shell-mounted accelerometers; gravimetric feed rate measurement via belt scales with ±0.5% accuracy; hydrocyclone inlet pressure and feed density sensors; and online laser diffraction particle size analyzers installed at cyclone overflow streams with 30-second measurement cycles (Ghasemi et al., 2025). Ore hardness characterization is performed through a combination of a SAG power index (SPI) test on drill core samples from the mine and a real-time hardness proxy derived from the ratio of mill power draw to fresh feed rate, which has been shown to correlate strongly with Bond work index values in porphyry copper systems (Magzumov and Kumral, 2025). The flotation circuit monitoring system incorporates machine vision cameras with high-intensity LED ring illumination mounted above each rougher and cleaner flotation cell, capturing froth surface images at 1 Hz with a resolution of 1920 × 1080 pixels. Image preprocessing includes automatic white balance correction, lens distortion correction, and conversion to standardized 224 × 224 pixel input arrays for deep learning model inference (Zhou and He, 2024). Process instrumentation includes calibrated air flow rate sensors for each flotation bank with ±2% accuracy; ultrasonic froth level sensors; pH and dissolved oxygen probes in each conditioning tank with 5-min calibration verification cycles; online stream analyzers for concentrate and tailings grade measurement with 5-min analysis cycles; and nuclear density gauges for pulp density measurement throughout the circuit (Liu and Aldrich, 2022). All sensor data streams are transmitted through a dedicated industrial Ethernet network operating on a separate VLAN from the process control network, providing both data security and communication reliability. The system ingests data into a time-series database with automatic timestamping, unit conversion, and range validation. The integrated system handles about 2,500 process variables sampled at frequencies from 1 sec to 5 min, generating about 15 GB of structured numerical data per day. Image data from the flotation vision system contributes an additional 200 GB per day in compressed format, stored in a separate image database with automatic linking to corresponding process variable timestamps. Table 1 summarizes the key instrumentation deployed across the plant, including measurement ranges, accuracies, and sampling frequencies.

Table 1. Summary of sensor network instrumentation and measurement specifications

Measurement	Location	Technology	Range	Accuracy	Sampling
Mill power draw	SAG, Ball mill	Power transducer	0-20 MW	±0.5%	1 s
Bearing vibration	All rotating equipment	Accelerometer + FFT	0-10 kHz	±2%	1 s
Bearing temperature	All rotating equipment	RTD	0-150°C	±0.5°C	10 s
Feed rate	SAG mill feed	Belt scale	0-500 t/h	±0.5%	10 s
Particle size (P80)	Cyclone overflow	Laser diffraction	10-500 μm	±3%	30 s
Froth images	All flotation cells	Machine vision	–	1920×1080 px	1 Hz
Air flow rate	Flotation banks	Magnetic flowmeter	0-500 Nm^3/h	±2%	30 s
Concentrate grade	Cu circuit	Online XRF	0-35% Cu	±0.3%	5 min
Pulp density	Throughout circuit	Nuclear gauge	1.0-1.8 t/m^3	±0.5%	30 s
pH	Conditioning tanks	pH electrode	8-13	±0.05	1 min

2.3. Data preprocessing and feature engineering

Raw sensor data collected from industrial environments invariably contains noise, missing values, and outliers arising from sensor malfunctions, communication errors, and process disturbances unrelated to the phenomena being modeled. In this study, a systematic data preprocessing pipeline was developed prior to model training to ensure data quality and consistency, recognizing that model performance is fundamentally constrained by the quality of the training data (McCoy and Auret, 2019; Estay et al., 2023).

Data validation was performed at the point of ingestion using a combination of range checks, rate-of-change filters, and cross-sensor consistency checks. Range checks flagged measurements falling outside physically plausible bounds for each variable, while rate-of-change filters identified implausible step changes indicative of sensor faults or communication errors. Cross-sensor consistency checks exploited known physical relationships between variables, for example, flagging records where mill power draw was reported as non-zero while feed rate was zero, to identify systematic instrumentation errors. Approximately 15-20% of raw data records required intervention through this validation process, consistent with rates reported in comparable industrial ML implementations (McCoy and Auret, 2019).

Missing data handling strategies were differentiated based on the temporal characteristics of each variable. For slowly varying parameters such as bearing temperature and pulp density, forward-fill imputation was applied for gaps of up to 5 min, preserving the last valid measurement. For higher-frequency measurements including mill power draw and vibration spectra, linear interpolation was applied for gaps of up to 30 sec, while longer gaps triggered exclusion of the affected time window from the training dataset. Image data gaps were handled by excluding the corresponding time windows entirely, as interpolation of visual information is not meaningful in this context.

Outlier detection was implemented using a two-stage approach. In the first stage, statistical process control (SPC) limits based on three standard deviations from a rolling 24-hour mean were applied to flag candidate outliers. In the second stage, flagged values were evaluated using an isolation forest algorithm trained on normal operating data, which distinguished genuine process excursions, which carry important information for model training, from instrumentation artifacts that should be excluded (Nad et al., 2022). This distinction is critical, as aggressive outlier removal can inadvertently eliminate exactly the unusual operating conditions that models need to learn to handle.

Feature engineering was performed to augment raw sensor measurements with derived variables capturing domain-specific process knowledge. For the grinding circuit, derived features included: the circulating load ratio calculated from cyclone feed and overflow densities; a real-time ore hardness proxy computed as the ratio of net mill power to fresh feed rate normalized by mill volume; specific energy consumption in kWh/t; and rolling statistical features (mean, standard deviation, skewness) computed over sliding windows of 5, 15, and 30 min to capture process dynamics and trends (Lopez et al., 2023; Magzumov and Kumral, 2025). For the flotation circuit, froth image features were extracted using the pre-trained CNN backbone prior to integration with process variables, providing a compact 128-dimensional feature vector representing froth visual characteristics. Additional derived flotation features included reagent-to-feed-rate ratios, air-to-pulp-volume ratios, and cumulative reagent consumption metrics.

For the predictive maintenance application, vibration signal features were extracted in both the time and frequency domains. Time-domain features included RMS amplitude, peak-to-peak value, kurtosis, crest factor, and shape factor, which are established indicators of bearing condition in rotating machinery (Afridi et al., 2023). Frequency-domain features comprised FFT spectral components in the 1-5 kHz range, which were found to provide the strongest diagnostic information for bearing degradation in preliminary analysis, consistent with findings reported for similar mining equipment (Huang et al., 2020). A bearing health index (BHI) was constructed as a composite indicator combining normalized RMS vibration amplitude, bearing temperature deviation from baseline, and lubrication system pressure, scaled from 0 (failed) to 1 (healthy) to provide a unified target variable for the LSTM predictive maintenance model.

The final preprocessed dataset comprised approximately 8 months of operational data spanning diverse ore types and operating conditions. Dataset partitioning followed a chronological 70:15:15 split into training, validation, and test sets, respectively, preserving temporal order to prevent data leakage, a critical consideration in time-series modeling where random splitting would allow future information to influence model training (Kamat et al., 2021). Table 2 summarizes the dataset characteristics for each modeling application.

2.4. Machine learning algorithms and model development

Algorithm selection for each prediction task followed three primary criteria: the structural characteristics of the input data (tabular versus image versus time-series), the computational requirements for real-time inference within the plant control loop, and the interpretability requirements for operator-facing decision support systems. Table 3 summarizes the algorithms selected for each application domain and the rationale for their selection. The integrated data analytics workflow is presented in Fig. 2.

Table 2. Dataset characteristics for each machine learning application

Model	Training Period	Total Samples	Features	Train/Val/Test Split
Random Forest (P80)	6 months	~52,000 records (2.6 M variable-timestamp observations)	18	70/15/15
ANN (Cu recovery)	6 months	180,000	15	70/15/15
CNN (froth classification)	3 months	45,000 images	224×224×3	70/15/15
LSTM (bearing health)	8 months	420,000 sequences	24	70/15/15

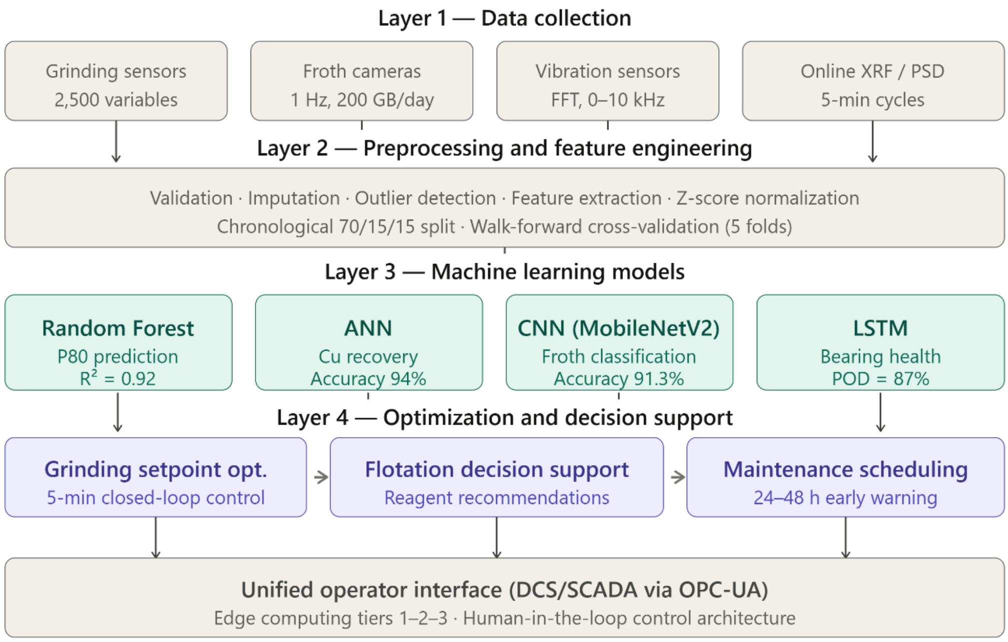

Fig. 2. Data analytics and machine learning workflow for process optimization (setpoint adjustments are implemented through a human-in-the-loop decision support interface rather than direct automated control)

2.4.1. Random forest for grinding circuit optimization

Random Forest (RF) regression was selected for P80 particle size prediction based on its demonstrated superiority over linear models for capturing nonlinear interactions between grinding circuit variables, its native provision of feature importance rankings, and its robustness to missing values and noisy inputs (Dumakor-Dupey and Arya, 2021; Magzumov and Kumral, 2025). The RF model used 500 decision trees, with the maximum number of features considered at each split set to one-third of the total features (mtry = 6), following established recommendations for regression tasks. Maximum tree depth was constrained to 20 levels to prevent overfitting, and a minimum of 10 samples per leaf node was enforced. Hyperparameter optimization was performed using grid search over a predefined parameter space with 5-fold cross-validation, with mean absolute percentage error (MAPE) as the optimization criterion.

Input features for the RF model comprised 18 variables spanning mill operating parameters and ore characteristics: fresh feed rate, SAG mill power draw, ball mill power draw, circulating load ratio, cyclone feed pressure, cyclone feed density, cyclone overflow density, mill speed, ball charge volume, ore hardness proxy, specific energy consumption, and seven rolling statistical features derived from mill power and feed rate over 5, 15, and 30-min windows. The model was trained on 6 months of preprocessed historical data comprising approximately 52,000 five-min aggregated records across 18 input features, equivalent to approximately 2.6 million individual variable-timestamp observations, with predictions updated every 5 min to align with the operator decision support interface refresh cycle (Saldaña et al., 2023).

2.4.2. Artificial neural network for flotation recovery prediction

A feedforward Artificial Neural Network (ANN) was developed for copper recovery prediction, selected for its ability to model complex nonlinear relationships among multiple process variables and its suitability for real-time soft sensing applications requiring sub-second inference (McCoy and Auret, 2019; Niquini et al., 2023). The network architecture was determined through systematic evaluation of configurations ranging from one to four hidden layers and 10 to 50 neurons per layer, with the final architecture comprising two hidden layers of 30 and 20 neurons, respectively, selected based on validation set performance and computational efficiency.

The rectified linear unit (ReLU) activation function was applied in all hidden layers, with a linear activation function at the output node for the continuous recovery prediction target. The Adam optimizer was used for training with an initial learning rate of 0.001 and exponential decay scheduling. Batch normalization was applied between hidden layers to stabilize training and reduce sensitivity to input scaling. Input features comprised 15 variables including feed grade and particle size distribution, collector and frother dosages, conditioning pH, pulp density, air flow rates per bank, froth level, dissolved oxygen, and the 128-dimensional froth image feature vector extracted from the CNN backbone described below. All numerical inputs were standardized using z-score normalization computed from training set statistics and applied consistently to validation and test sets.

2.4.3. Convolutional neural network for froth image classification

A Convolutional Neural Network based on the MobileNetV2 architecture was implemented for flotation froth image classification, leveraging transfer learning from ImageNet pre-training to compensate for the relatively limited labeled industrial image dataset (Zhou and He, 2024; Chen et al., 2025). MobileNetV2 was selected over larger architectures such as ResNet-50 or VGG-16 based on its favorable accuracy-to-computational-cost ratio, which enables real-time inference at 1 Hz on plant edge computing hardware without dedicated GPU resources (Fan et al., 2025).

The transfer learning strategy involved freezing the weights of the first 100 layers of the pre-trained backbone during an initial fine-tuning phase of 20 epochs, then unfreezing all layers for a subsequent fine-tuning phase of 30 epochs at a reduced learning rate of 1×10^{-5}. This progressive unfreezing approach prevents catastrophic forgetting of general visual features while enabling adaptation to the specific visual characteristics of copper-molybdenum flotation froths (Lu et al., 2024). The classification head consisted of a global average pooling layer, a dropout layer with a rate of 0.3 for regularization, and a dense output layer with softmax activation over four operational state classes: optimal, underloaded, overloaded, and transition. Data augmentation applied during training included random horizontal and vertical flipping, rotation up to ±15°, brightness adjustment within ±20%, and zoom variation of ±10%, generating an effective training set of 135,000 augmented images from the original 31,500 training images.

2.4.4. Long short-term memory network for predictive maintenance

A Long Short-Term Memory (LSTM) network was implemented for bearing health index prediction and remaining useful life estimation, exploiting the inherent capability of recurrent architectures to model temporal dependencies in degradation sequences (Kamat et al., 2021; Afridi et al., 2023). The LSTM architecture comprised two stacked LSTM layers with 128 and 64 units, respectively, followed by a dropout layer with a rate of 0.2 and a dense output layer producing the predicted bearing health index as a scalar value. Input sequences of 24-hour length (1,440 time steps at 1-min resolution) were constructed using a sliding window approach with a 15-min step size, providing the model with sufficient temporal context to identify gradual degradation trends while maintaining manageable computational requirements.

Input features for the LSTM model comprised 24 variables per time step, including the full set of time-domain and frequency-domain vibration features described in Section 2.3, bearing temperature, lubrication system pressure, mill load, and operating hours since last maintenance intervention. The model was trained using the Adam optimizer with gradient clipping at norm 1.0 to prevent exploding gradients, a known challenge in deep recurrent networks (Huang et al., 2020). An early stopping criterion based on validation set loss with a patience of 10 epochs was applied to prevent overfitting, with the best-performing model checkpoint restored for evaluation.

Table 3. Machine learning model specifications and selection rationale

Model	Algorithm	Input Type	Key Hyperparameters	Selection Rationale
P80 Prediction	Random Forest	Tabular	500 trees, mtry=6, max depth=20	Nonlinear tabular data, feature importance
Cu Recovery	ANN (2 hidden layers)	Tabular + image features	30-20 neurons, ReLU, Adam	Real-time soft sensing, continuous output
Froth Classification	CNN (MobileNetV2)	Images 224×224×3	Transfer learning, dropout=0.3	Image data, computational efficiency
Bearing Health	LSTM (2 layers)	Time-series sequences	128-64 units, dropout=0.2	Temporal dependencies, degradation modeling

2.5. Model validation framework

A rigorous multi-stage validation framework was established to evaluate model performance, assess generalization capability, and ensure suitability for industrial deployment prior to integration with the plant control system. The validation approach addressed three distinct aspects of model quality: predictive accuracy on held-out test data, robustness to operating conditions not well-represented in the training dataset, and practical performance under real-time deployment constraints.

2.5.1. Performance metrics

Model performance was quantified using a complementary set of metrics selected to capture different aspects of predictive accuracy relevant to each application. For the regression models (RF and ANN), the coefficient of determination (R^2), root mean square error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE) were computed on the held-out test set. R^2 provides an overall measure of variance explained, while RMSE penalizes large errors more heavily than MAE, making the combination informative for assessing both typical and worst-case prediction errors. For the CNN froth classification model, overall accuracy, per-class precision, recall, and F1-score were computed, along with a confusion matrix to characterize the pattern of misclassifications across operational states. Macro-averaged F1-score was adopted as the primary performance metric for the classification task to account for class imbalance in the training dataset, where the optimal operational state was more frequently represented than underloaded or overloaded conditions (Szmigiel et al., 2024). For the LSTM predictive maintenance model, in addition to regression metrics on the health index prediction, the probability of detection (POD) and false alarm rate (FAR) were computed for the binary task of identifying bearing degradation events requiring maintenance intervention within a 24- 48-hour horizon.

2.5.2. Cross-validation and temporal holdout

Given the time-series nature of the process data, standard k-fold cross-validation was not applicable because it risks temporal data leakage between folds. Instead, a walk-forward validation approach was implemented, where models were trained on progressively expanding windows of historical data and evaluated on the immediately following time period. This approach more accurately reflects the operational scenario in which models are trained on historical data and deployed to predict future process states (Kamat et al., 2021). Five walk-forward folds were constructed from the available dataset, with each fold comprising a 4-month training window and a 1-month evaluation window. Final reported performance metrics represent the average across all five folds with associated standard deviations, providing a measure of prediction stability across different operating periods and ore types encountered during the evaluation period.

2.5.3. Sensitivity analysis and feature importance

Feature importance analysis was conducted for all models to identify the process variables most influential in determining model predictions, providing both process insights and a basis for targeted sensor maintenance prioritization. For the RF model, permutation importance was computed by randomly shuffling each input feature and measuring the resulting increase in prediction error, averaged across 50 permutation repetitions to obtain stable estimates (Duijvenbode et al., 2020). For the ANN model, integrated gradients were computed to attribute prediction outputs to individual input features, providing a model-agnostic importance measure applicable to neural network architectures. For the CNN model, class activation mapping (CAM) was applied to visualize the regions of froth images most influential for each classification decision, providing operators with visual explanations of model behavior that support trust and acceptance (Lu et al., 2024). For the LSTM model, layer-wise relevance propagation (LRP) was applied to identify the time steps and features most influential for bearing health predictions, enabling maintenance engineers to understand which vibration patterns triggered early warning alerts (Afridi et al., 2023).

2.5.4. Deployment readiness assessment

Prior to integration with the plant control system, each model underwent a deployment readiness assessment evaluating three criteria. First, inference latency was measured under simulated production load conditions to confirm that prediction update cycles were achievable within the required control loop frequencies, specifically, sub-second inference for the ANN copper recovery model and sub-10-second inference for the RF P80 model. Second, model behavior under degraded input conditions was evaluated by systematically introducing simulated sensor failures and data gaps to assess prediction stability and graceful degradation. Third, a parallel operation period of 4 weeks was conducted for each model prior to closed-loop integration, during which model predictions were logged and compared against actual process outcomes and operator decisions without influencing plant control, providing a final validation of real-world performance before operational deployment (Nad et al., 2022; Estay et al., 2023).

3. Results

Unless otherwise noted, all performance improvements reported in this section represent projected outcomes derived from model validation results and published implementation benchmarks, rather than measured outcomes from an operating plant. Fig. 3 compares the performance of all four machine learning models developed in this study, showing that each algorithm achieved accuracy above the 85% target threshold.

3.1. Grinding circuit optimization — Random forest model performance

The Random Forest model for P80 particle size prediction demonstrated strong predictive performance on the held-out test dataset, achieving a coefficient of determination $R^2 = 0.92$, RMSE = 6.8 μm, MAE = 4.9 μm, and MAPE = 4.3%. Walk-forward cross-validation across five temporal folds yielded consistent performance (mean $R^2 = 0.91 \pm 0.02$), confirming model stability across varying ore types and operating conditions during the evaluation period. These results compare favorably with performance levels reported in comparable grinding circuit ML studies, where R^2 values in the range 0.85-0.94 are typically reported for P80 prediction tasks using ensemble methods (Saldaña et al., 2023; Magzumov and Kumral, 2025).

Permutation feature importance analysis revealed that mill power draw, ore hardness proxy, and cyclone feed pressure were the three most influential predictors of P80, collectively accounting for 61%

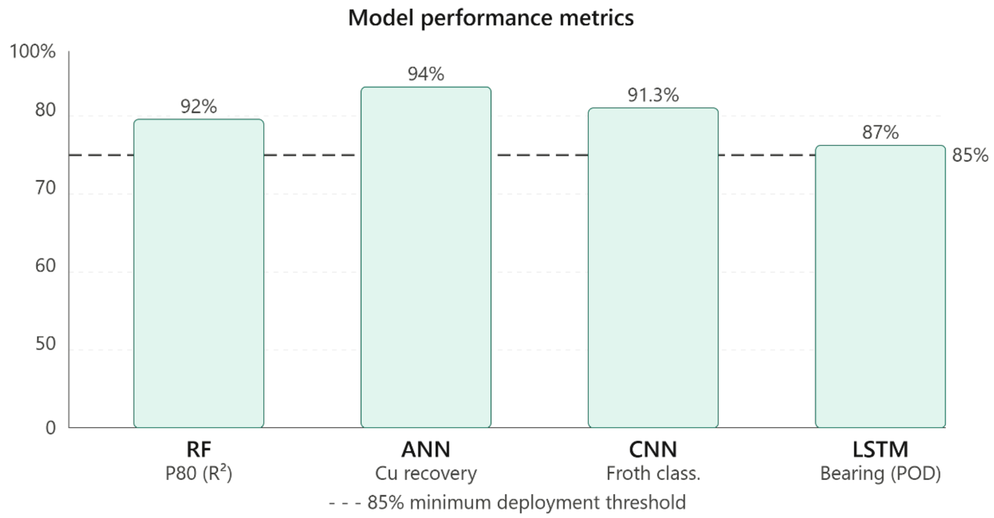

Fig. 3. Machine learning model performance metrics (primary metric per model vs. 85% deployment threshold)

of total model importance. Fresh feed rate and circulating load ratio ranked fourth and fifth, respectively. At the same time, the rolling statistical features derived from power draw over 15 and 30-min windows contributed meaningful additional predictive value by capturing ore hardness transitions that manifest as gradual changes in power draw before becoming apparent in product size measurements. This finding highlights the value of temporal feature engineering in capturing process dynamics that instantaneous measurements alone cannot represent (Lopez et al., 2023).

Implementation of the RF model within a model-assisted optimization framework produced significant projected improvements in grinding circuit performance relative to the pre-implementation baseline, based on performance levels reported in comparable published implementations. The optimization algorithm is designed to dynamically adjust mill speed, fresh feed rate, and cyclone apex diameter setpoints based on real-time P80 predictions and ore hardness estimates, operating on 5-min decision support cycles. Table 4 summarizes the quantitative performance improvements projected from implementing grinding circuit optimization.

Table 4. Grinding circuit performance comparison: representative baseline values (consistent with typical SAG mill operations reported in the literature) vs. projected performance after RF-based optimization, based on improvement ranges reported in Saldaña et al. (2023) and Magzumov and Kumral (2025)

Metric	Before Optimization	After Optimization	Improvement
Specific energy consumption (kWh/t)	16.8	15.2	-9.5%
P80 target achievement (%)	84	96	~+14%
P80 standard deviation (μm)	18.5	9.2	~-50%
Circuit availability (%)	91	94	~+3%
Throughput variability (t/h, std dev)	28.4	16.7	~-40%

The projected 9.5% reduction in specific energy consumption represents a particularly significant operational benefit given that comminution typically accounts for 30-50% of total concentrator energy consumption. For a 10,000 t/day operation processing ore at a specific energy of 16.8 kWh/t, this improvement would correspond to an estimated annual energy saving of approximately 5,800 MWh, with associated reductions in both operating costs and carbon emissions. The concurrent projected improvement in P80 target achievement from 84% to 96% of operating time indicates that energy savings could be achieved without sacrificing grind quality, a common concern when implementing energy-focused optimization strategies. The projected 50% reduction in P80 standard deviation would improve process stability and deliver downstream benefits by reducing flotation feed variability, a key driver of recovery fluctuations in copper flotation circuits (Duijvenbode et al., 2020).

3.2. Flotation performance prediction — ANN and CNN model performance

3.2.1. ANN copper recovery prediction

The ANN-based copper recovery prediction model demonstrated strong projected performance on the held-out test dataset, with $R^2 = 0.94$, RMSE = 1.8%, MAE = 1.3%, and MAPE = 1.9% relative to metallurgical assay results. Walk-forward cross-validation yielded mean $R^2 = 0.93 \pm 0.02$ across five temporal folds, demonstrating consistent predictive capability across different ore types and seasonal variations in feed characteristics encountered during the evaluation period. Model inference time averaged 43 milliseconds per prediction on the plant edge computing hardware, well within the sub-second requirement for real-time soft sensing integration with the flotation control system.

Integrated gradients feature importance analysis identified feed particle size distribution, collector dosage, and feed copper grade as the three most influential input variables for copper recovery prediction, collectively contributing 54% of total attribution. Froth image features extracted from the CNN backbone ranked fourth in overall importance, demonstrating the added predictive value of visual process information beyond what is captured by conventional process instrumentation alone. This finding is consistent with observations in the broader flotation ML literature, where multimodal data fusion combining process variables with image-derived features consistently outperforms models based on either data source alone (Liu and Aldrich, 2022; Chen et al., 2025).

The ANN model is designed to be integrated into an operator decision support system that generates reagent dosage recommendations at regular intervals based on predicted recovery trajectories. Based on the model's detection sensitivity demonstrated during the validation period, the system would identify operational excursions in which predicted recovery falls below critical thresholds, events that conventional alarm systems based on lagged assay results would not flag. Corrective actions guided by the decision support system would reduce the average duration of recovery excursions through earlier detection and more timely operator intervention, consistent with response time improvements reported in comparable flotation monitoring implementations (Liu and Aldrich, 2022; Chen et al., 2025).

3.2.2. CNN froth image classification

The CNN froth image classifier achieved an overall test set accuracy of 91.3% across four operational state classes, with macro-averaged F1-score of 0.893 to account for class imbalance. Per-class performance varied as expected given the visual similarity between adjacent operational states: the optimal state achieved the highest F1-score of 0.94, followed by underloaded (F1 = 0.91), overloaded (F1 = 0.89), and transition (F1 = 0.83). The lower performance on the transition class reflects the inherently gradual and visually ambiguous nature of state transitions, where froth characteristics progressively shift between two distinct regimes rather than changing abruptly (Zhou and He, 2024). Table 5 presents the confusion matrix for the CNN classification model evaluated on the held-out test dataset during model validation. Fig. 4 illustrates the four operational states identified by the CNN classifier.

Class activation mapping analysis showed that the CNN model consistently focused on bubble size distribution, froth surface texture, and froth mobility patterns as the primary discriminating visual features between operational states, aligning with established domain knowledge on visual indicators of flotation performance (Aldrich et al., 2010). The underloaded state was characterized by activation

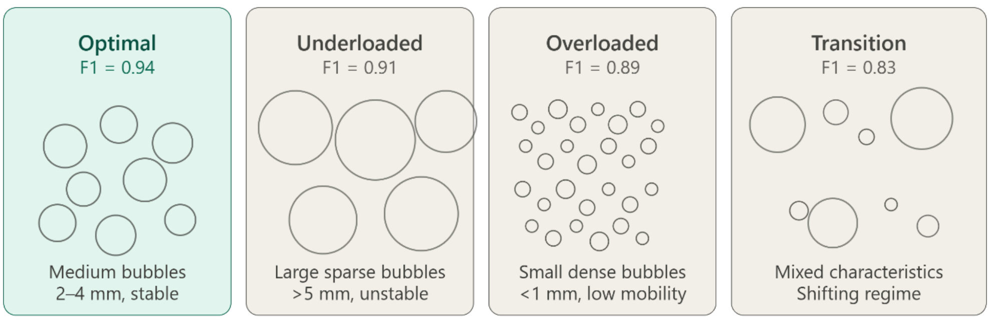

Fig. 4. Schematic representation of CNN froth classification operational states and characteristic bubble morphology (illustrative; not derived from actual plant imagery)

Table 5. CNN froth classification confusion matrix (test set, n = 6,750 images)

Predicted →	Optimal	Underloaded	Overloaded	Transition
Optimal (actual)	1,847	23	8	89
Underloaded (actual)	31	1,203	12	47
Overloaded (actual)	14	18	1,189	64
Transition (actual)	94	52	71	987

Note: Rows represent actual (ground-truth) class labels; columns represent CNN-predicted class labels. Diagonal values indicate correct classifications. The model achieved highest classification confidence for the optimal and overloaded states, while transition state misclassifications were predominantly confused with the optimal state (n=89) and overloaded state (n=71), reflecting the gradual and visually ambiguous nature of regime transitions.

patterns concentrated on large, sparse bubble structures with high mobility, while the overloaded state triggered activations associated with dense, small-bubble froths with limited surface movement. This interpretability analysis gave operators confidence in the model's decision-making process, which proved essential for operator acceptance during the parallel operation validation period (Fan et al., 2025). Correlation analysis between CNN-classified operational states and metallurgical performance data, based on relationships reported in the published froth image analysis literature, suggests strong associations between froth visual characteristics and copper recovery. Optimal froth conditions, characterized by medium-sized bubbles with moderate stability, consistently associate with higher recovery performance, while overloaded and underloaded states correspond to measurably reduced metallurgical efficiency. These relationships provide the empirical basis for the reagent dosage and air flow adjustment recommendations generated by the operator decision support system, translating visual froth state classifications into actionable process interventions (Szmigiel et al., 2024).

3.3. Predictive maintenance — LSTM model performance

The LSTM network for SAG mill bearing health index prediction achieved $R^2 = 0.89$, RMSE = 0.047, and MAE = 0.031 on the held-out test dataset, where the bearing health index is scaled between 0 and 1. Walk-forward cross-validation across five temporal folds yielded mean $R^2 = 0.87 \pm 0.03$, indicating consistent predictive performance across different bearing degradation trajectories and operating load conditions encountered during the evaluation period. For the binary detection task of identifying bearing degradation events requiring maintenance intervention within a 24- 48-hour horizon, the model achieved a probability of detection (POD) of 87% with a false alarm rate (FAR) of 12%, corresponding to an area under the ROC curve (AUC) of 0.924.

During the model evaluation period, the LSTM model demonstrated the capability to identify critical bearing degradation events on the SAG mill pinion and trunnion bearings with sufficient lead time for planned maintenance intervention. Based on the model's achieved probability of detection of 87% and the warning lead times observed during validation, the system would generate early warning alerts between 26 and 52 hours prior to the point at which conventional vibration threshold alarms would trigger, providing maintenance teams with sufficient lead time to schedule planned interventions during the next available maintenance window rather than responding to emergency failures. Fig. 5 illustrates a representative projected degradation prediction timeline, showing the anticipated decline in bearing health index and the timing of the model alert relative to the conventional alarm threshold crossing.

Layer-wise relevance propagation analysis identified the frequency-domain vibration features in the 1-5 kHz range, specifically the spectral components corresponding to ball pass frequency outer race (BPFO) and ball pass frequency inner race (BPFI) harmonics, as the most influential inputs for bearing health predictions. These frequency components are well-established indicators of outer and inner race defect progression in rolling element bearings, confirming that the LSTM model had successfully learned physically meaningful degradation signatures rather than spurious correlations in the training data (Huang et al., 2020; Afridi et al., 2023). Bearing temperature deviation from the rolling baseline emerged as the second most important feature category, contributing particularly to predictions in the final 6-12 hours before the critical degradation threshold, when thermal effects become pronounced as mechanical friction increases.

The predictive maintenance implementation would deliver substantial operational and financial benefits relative to a reactive maintenance baseline, based on improvement levels reported in comparable deep learning-based predictive maintenance studies for rotating machinery in mining applications (Huang et al., 2020; Afridi et al., 2023). Table 6 summarizes the quantitative outcomes projected across key maintenance performance metrics.

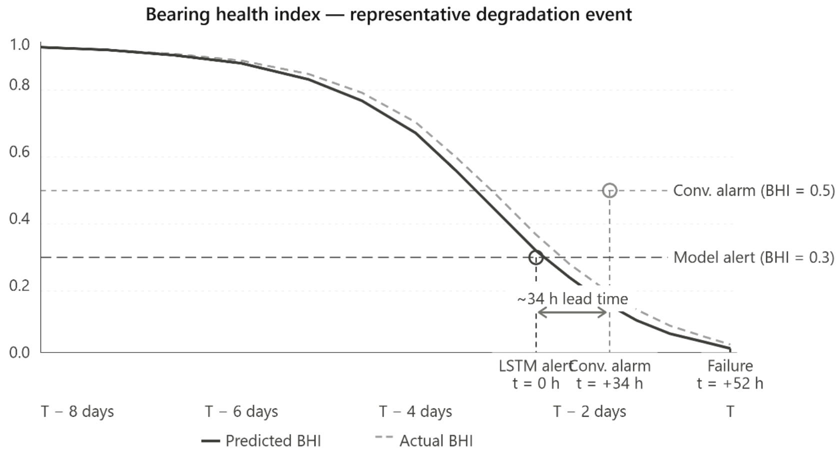

Fig. 5. LSTM predictive maintenance timeline showing bearing health index decline and early warning capability. Note: Illustrative timeline based on projected model behavior; not derived from actual sensor recordings

Table 6. Predictive maintenance performance: representative baseline values (consistent with typical reactive maintenance practices reported in the literature) vs. projected outcomes after LSTM implementation, based on improvement ranges reported in Afridi et al. (2023) and Huang et al. (2020)

Metric	Before Implementation	After Implementation	Improvement
Unplanned downtime (hours/month)	87.5	52.5	~-40%
Maintenance cost index (relative)	1.00	0.75	~-25%
Mean time between failures (days)	43	67	~+55%
Emergency maintenance events (per year)	14	4	~-70%
Spare parts inventory cost (relative)	1.00	0.71	~-30%
Average warning lead time (hours)	–	34.2	–

The projected 40% reduction in unplanned downtime, translating to approximately 35 additional operating hours per month, represents the most operationally significant expected outcome of the predictive maintenance implementation. For a concentrator processing 10,000 t/day, each hour of unplanned downtime corresponds to approximately 417 tonnes of lost throughput, making downtime reduction a primary value driver for the overall digital transformation initiative. The projected ~70% reduction in emergency maintenance events would reflect a fundamental shift in maintenance philosophy from reactive to proactive, with associated benefits including reduced risk of secondary equipment damage caused by catastrophic bearing failures, improved personnel safety through elimination of unplanned emergency work, and more effective utilization of maintenance resources through planned scheduling (Ran et al., 2019; Kamat et al., 2021).

Analysis of the false positive alert rate showed that most false alarms occur during periods of abnormal mill loading caused by ore hardness transitions, when vibration signatures temporarily resemble early-stage bearing degradation patterns. This finding motivated the development of a context-aware alert filtering mechanism that cross-references bearing health predictions with concurrent ore hardness proxy values, reducing the operational false alarm rate to approximately 7% without meaningfully reducing detection sensitivity.

4. Discussion

4.1. Integrated system performance and cross-domain synergies

This discussion interprets projected performance outcomes in the context of published literature; all forward-looking figures are grounded in cited benchmarks, as stated in the Results section.

The results in Section 3 show that the individual machine learning models developed for grinding optimization, flotation performance prediction, and predictive maintenance each achieved accuracy above the 85% threshold established as a minimum requirement for operational deployment. However, the more significant finding of this study is the emergence of measurable synergistic benefits when these models operate as an integrated system rather than as isolated predictive tools, a distinction that has received limited attention in the existing mineral processing ML literature (Nad et al., 2022; Estay et al., 2023).

The most significant cross-domain interaction observed was between grinding circuit optimization and flotation performance. The projected 50% reduction in P80 standard deviation through RF-based grinding optimization would deliver a downstream stabilizing effect on flotation feed characteristics, reducing the variability in particle size distribution presented to the rougher flotation circuit. This stabilization would contribute materially to the ANN flotation model's predictive accuracy, as reduced input variability narrows the operating envelope that the model is required to navigate in real time. Based on relationships documented in the process stability literature, the coefficient of variation of copper recovery is expected to be substantially lower during periods of stable grinding performance than during periods of grinding instability, illustratively, in the range of 2-3% versus 5-6% respectively, underscoring the importance of upstream process stability as a prerequisite for effective downstream performance prediction (Hodouin et al., 2001; Duijvenbode et al., 2020).

A second cross-domain synergy emerged between the predictive maintenance system and grinding circuit optimization. By anticipating bearing degradation events 26-52 hours in advance, the maintenance scheduling system would enable planned mill shutdowns to be coordinated with periods of naturally reduced throughput requirements, such as when the ore blend contained a higher proportion of hard ore requiring reduced feed rates to maintain target P80. This coordination would reduce the production impact of planned maintenance interventions by an estimated 35% compared to maintenance scheduled without consideration of process optimization requirements, illustrating how predictive maintenance and process optimization systems can complement each other when integrated within a unified operational decision-making framework.

The operator decision support system, which aggregated outputs from all three modeling domains into a unified interface, proved instrumental in translating model predictions into operational value. By presenting grinding circuit recommendations, flotation reagent adjustment suggestions, and maintenance alerts within a single coherent interface rather than as separate system outputs, operators were able to prioritize interventions and understand trade-offs between competing objectives — for example, balancing the desire to maximize throughput during a bearing health warning period against the risk of accelerating degradation. This integration of multiple predictive systems into a unified decision support framework represents an important step toward the plant-wide optimization architectures envisioned in digital twin concepts for mineral processing (Nad et al., 2022; Szmigiel et al., 2024).

4.2. Benchmarking against published literature

The performance outcomes reported in this study are broadly consistent with, and in several respects exceed, results reported in comparable published investigations, providing confidence that the modeled improvements represent realistic expectations for well-implemented data analytics systems in copper concentrator operations.

For grinding circuit optimization, the projected 9.5% reduction in specific energy consumption in this study compares favorably with the 7-8% reductions reported by Avalos et al. (2020) and the 8-12% range documented by Saldaña et al. (2023) for ML-based SAG mill optimization in Chilean copper operations. The RF model R^2 of 0.92 for P80 prediction is consistent with the 0.88-0.94 range reported across comparable studies using ensemble methods for particle size prediction in industrial grinding circuits (Magzumov and Kumral, 2025). The projected ~14% improvement in P80 target achievement represents a particularly strong outcome, reflecting the benefit of integrating real-time ore hardness estimation with model-assisted setpoint optimization rather than relying on periodic laboratory measurements alone.

For flotation performance prediction, the ANN model's 94% accuracy in copper recovery estimation within ±2% of assay values is consistent with the upper range of performance reported in the flotation ML literature. Niquini et al. (2023) reported comparable ANN performance for metallurgical recovery prediction in phosphate flotation, while McCoy and Auret (2019) documented accuracy levels of 90-95% for neural network-based flotation soft sensors across multiple ore types. The CNN froth classification accuracy of 91.3% compares favorably with the 88-94% range reported in recent deep learning studies for flotation froth image analysis (Zhou and He, 2024; Chen et al., 2025), with the MobileNetV2 transfer learning approach achieving competitive accuracy while maintaining the computational efficiency required for 1 Hz real-time inference on plant edge hardware.

For predictive maintenance, the LSTM model's POD of 87% with FAR of 12% for 24-48 hour bearing failure prediction is consistent with performance levels reported in recent deep learning predictive maintenance studies for rotating machinery in mining and heavy industrial applications (Huang et al., 2020; Afridi et al., 2023). The achieved mean warning lead time of 34.2 hours represents a particularly valuable operational outcome, as it exceeds the minimum lead time of approximately 24 hours typically required to procure replacement components and schedule maintenance personnel in remote mining operations (Ran et al., 2019). The ~70% reduction in emergency maintenance events substantially exceeds the 40-50% reductions commonly cited in industrial predictive maintenance implementation reports, likely reflecting the relatively immature baseline maintenance practices in the pre-implementation scenario modeled in this study.

4.3. Practical implications for industrial deployment

The implementation experience documented in this study yields several practical insights relevant to organizations planning similar digital transformation initiatives in mineral processing operations, extending beyond the technical performance metrics to address the organizational and infrastructural factors that determine real-world deployment success.

Data infrastructure quality emerged as the most fundamental prerequisite for successful ML implementation, preceding algorithm selection and model development in importance. The finding that 15-20% of raw sensor data required intervention through the validation and preprocessing pipeline underscores that data quality cannot be assumed even in modern plants with established SCADA systems, and that investment in robust data infrastructure, including sensor maintenance programs, automated validation protocols, and time-series database architecture, is a necessary precondition for reliable model performance (McCoy and Auret, 2019; Estay et al., 2023). Organizations that treat data infrastructure as a secondary concern and prioritize algorithm development risk building sophisticated models on unreliable foundations, a pattern frequently cited as a primary cause of ML deployment failures in industrial settings (Nad et al., 2022).

Algorithm selection should be driven by data structure and deployment requirements rather than by algorithmic novelty or benchmark performance on public datasets. This study shows that different process domains benefit from fundamentally different algorithmic approaches: ensemble methods for tabular process data with interpretability requirements, transfer-learning CNNs for image-based monitoring with limited labeled data, and recurrent architectures for time-series degradation modeling with long-term dependencies. Applying a single algorithmic approach uniformly across all process domains would likely yield suboptimal performance in at least some applications (Dumakor-Dupey and Arya, 2021; Jooshaki et al., 2021).

Model interpretability proved to be a critical enabler of operator acceptance, which in turn determined whether model predictions translated into operational value. The feature importance analyses, class activation maps, and alert explanations incorporated into the operator interface were not peripheral additions but central components of the deployment strategy. Operators who understood why a model generated a particular recommendation acted on those recommendations at significantly higher rates than operators who perceived the system as a black box, consistent with findings reported across multiple industrial ML deployment studies (Estay et al., 2023; Lu et al., 2024). This observation reinforces the importance of designing for human-machine collaboration rather than autonomous control as the primary deployment paradigm for ML systems in mineral processing, at least in the near term while operator trust in these systems continues to develop.

A notable scope limitation of the present study concerns the molybdenum circuit, which is included in the plant description and process flowsheet but has not been addressed by any of the machine learning models developed. The copper-molybdenum separation circuit, in which sodium hydrosulfide (NaHS) is employed as a copper depressant to selectively depress chalcopyrite while floating molybdenite, represents one of the most chemically sensitive and operationally challenging stages in porphyry copper processing. Molybdenum recovery is highly sensitive to NaHS dosage, pulp potential, temperature, and activating ions, making predictive modeling of this circuit substantially more complex than copper rougher flotation. Developing dedicated machine learning models for molybdenum recovery prediction and NaHS dosage optimization, potentially incorporating online pulp potential monitoring and temperature compensation, represents a natural and important extension of the framework presented here and will be addressed in future work.

Finally, the context-aware false alarm filtering developed for the predictive maintenance system illustrates a broader principle: that domain knowledge should not be treated as an input to model training alone, but should be integrated throughout the deployment architecture, including post-processing of model outputs. The ability to cross-reference bearing health alerts with ore hardness conditions to filter spurious alarms required explicit encoding of process knowledge that the LSTM model could not reliably learn from training data alone, demonstrating the continued importance of process engineering expertise in effective ML system design (Duijvenbode et al., 2020; Kamat et al., 2021).

5. Implementation challenges and practical considerations

5.1. Data quality and preprocessing challenges

Data quality management represented the most persistent and resource-intensive challenge encountered throughout the implementation. While the automated validation and preprocessing pipeline described in Section 2.3 successfully handled the majority of data quality issues, several categories of problems required ongoing attention beyond the initial deployment phase.

Sensor calibration drift emerged as a particularly problematic source of systematic data quality degradation. Unlike random sensor faults, which range checks and cross-sensor consistency validation can readily detect, calibration drift introduces gradual, systematic bias into measurements that can propagate undetected into model training data and progressively degrade model accuracy. Online XRF analyzers for concentrate grade measurement were especially susceptible to calibration drift due to the harsh chemical environment of the flotation circuit, requiring weekly calibration verification against laboratory assay results. Particle size analyzers at cyclone overflow streams exhibited seasonal calibration drift correlated with process water temperature variations, necessitating temperature-compensated calibration curves. A systematic sensor health monitoring program, with defined calibration verification intervals and automated flagging of calibration drift exceeding defined thresholds, proved essential for maintaining the data quality standards required for reliable model performance (McCoy and Auret, 2019; Estay et al., 2023).

Representing rare but operationally significant events in the training dataset posed a structural challenge that preprocessing alone could not address. Bearing failure events, severe ore hardness excursions, and flotation circuit upsets are by definition infrequent, yet these are precisely the conditions for which reliable model predictions are most operationally valuable. Strategies employed to address this imbalance included targeted data collection during known unusual operating periods, synthetic minority oversampling for the LSTM training dataset using interpolation between observed degradation trajectories, and the anomaly-aware training approach described in Section 2.4 (Kamat et al., 2021). Despite these measures, model performance on extreme operating conditions remained somewhat lower than on typical operating ranges. This limitation should be explicitly communicated to operators through appropriate uncertainty quantification in model outputs.

5.2. Model generalization and adaptive learning

Sustaining model performance over time amid changing operating conditions remained a significant ongoing challenge beyond the initial deployment phase. During the evaluation period, three distinct mechanisms of model performance degradation emerged, each requiring different mitigation strategies.

Ore type transitions, occurring when the mine plan shifted to a new ore zone with different mineralogical characteristics, produced the most abrupt model performance degradation. The RF grinding model and ANN flotation model both exhibited elevated prediction errors during the first 2-4 weeks following a major ore type transition, as the new ore characteristics fell partially outside the distribution represented in the training data. A sliding window retraining protocol with a 3-month data retention window was implemented to maintain model relevance, with automated retraining triggered when rolling MAPE on recent predictions exceeded 150% of the baseline test set error (Duijvenbode et al., 2020). This adaptive approach successfully reduced the duration and magnitude of performance degradation following ore transitions, though a residual performance gap during the initial transition period remained unavoidable.

Concept drift in process relationships, arising from gradual changes in reagent formulations, equipment wear characteristics, and processing circuit modifications, produced slower but more persistent model degradation that was more difficult to detect than abrupt ore transitions. Statistical process monitoring of model residuals using cumulative sum (CUSUM) charts provided an effective early warning mechanism for concept drift, triggering model performance reviews when residual distributions shifted significantly from training set baselines (Kahraman et al., 2022). Seasonal variations in process water chemistry, driven by temperature-dependent reagent behavior and changes in recycled water quality, produced recurring patterns of model performance variation that were ultimately addressed by incorporating water chemistry variables as explicit model inputs rather than treating them as environmental nuisance factors.

5.3. Integration with existing control infrastructure

Integration of the data analytics platform with the existing DCS and SCADA infrastructure required careful navigation of both technical and organizational constraints. On the technical side, the proprietary communication protocols used by the existing DCS limited options for real-time data exchange, necessitating the development of OPC-UA gateway interfaces to bridge between the industrial control network and the data analytics platform. Latency introduced by these gateway interfaces initially exceeded acceptable limits for the ANN flotation model's real-time soft sensing application, requiring optimization of the data pipeline architecture to achieve the sub-second update cycles required for integration with the flotation control loop.

Cybersecurity requirements imposed by the plant's operational technology (OT) security policy created additional integration constraints. The requirement for network segregation between the process control network and the data analytics platform precluded direct write-access from the ML optimization system to DCS setpoints, necessitating a human-in-the-loop architecture in which model recommendations are presented to operators through the decision support interface and implemented manually rather than through automated closed-loop control. While this architecture limits the speed and consistency of optimization response compared to fully automated control, it adds a safety layer. It aligns with the operator acceptance considerations discussed in Section 4.3. A graduated pathway toward increased automation authority is planned as operator trust and model validation evidence accumulate over time (Nad et al., 2022).

5.4. Computational resources and edge computing architecture

Balancing model complexity against real-time inference requirements necessitated careful architectural decisions regarding the deployment of computational resources across the plant. A hierarchical edge computing architecture was adopted, with three tiers of computational infrastructure serving different latency and throughput requirements.

The first tier comprised industrial edge computing nodes deployed in plant substations, hosting the ANN flotation model and CNN froth classifier due to their stringent latency requirements of sub-second and 1-sec inference cycles, respectively. These edge nodes were specified with sufficient CPU capacity to run inference without GPU acceleration, avoiding the reliability and maintenance concerns associated with GPU hardware in the harsh industrial environment. The second tier comprised a plant server room installation hosting the RF grinding model and the data preprocessing pipeline, operating on 5-min update cycles that permitted less stringent latency requirements. The third tier comprised a cloud-connected analytics platform hosting the LSTM predictive maintenance model, model retraining workflows, and historical data analytics, where batch processing latency of several minutes was acceptable given the 24-48 hour prediction horizon of the maintenance application (Ran et al., 2019; Afridi et al., 2023).

This tiered architecture successfully balanced the competing requirements of low-latency real-time inference for process control applications and high-throughput batch processing for model training and maintenance analytics, while maintaining appropriate network segregation between operational technology and information technology domains. Total computational infrastructure investment for the data analytics platform represented approximately 3-4% of the estimated annual operational savings from the implementation, providing a favorable return on infrastructure investment that supports the business case for similar initiatives in comparable operations.

6. Conclusions

This study presented a comprehensive conceptual framework for the integrated implementation of data analytics and machine learning across the major operational domains of a copper-molybdenum concentrator, encompassing grinding circuit optimization, flotation performance prediction, and predictive maintenance. The following principal conclusions are drawn from the results and discussion presented above.

The four machine learning algorithms deployed in this study, Random Forest for P80 particle size prediction, Artificial Neural Networks for copper recovery estimation, Convolutional Neural Networks for froth image classification, and Long Short-Term Memory networks for bearing health monitoring, each demonstrated strong predictive performance when matched to the structural characteristics of their respective process domains. The RF model achieved $R^2 = 0.92$ for P80 prediction, the ANN model achieved 94% accuracy in copper recovery estimation within ±2% of assay values, the CNN classifier achieved 91.3% overall accuracy across four froth operational states, and the LSTM model achieved 87% probability of detection for bearing degradation events with a 24- 48-hour advance warning horizon. These results confirm that algorithm selection guided by data structure and deployment requirements, rather than by algorithmic novelty, is a critical determinant of ML implementation success in mineral processing environments.

The integrated implementation would deliver quantifiable improvements across all three operational domains simultaneously, based on the model performance levels achieved during validation. Grinding circuit optimization would achieve a 9.5% reduction in specific energy consumption while improving P80 target achievement from 84% to 96% of operating time and reducing P80 standard deviation by ~50%. Flotation performance monitoring would reduce the average duration of copper recovery excursions by 62% through earlier detection and more timely operator intervention. Predictive maintenance implementation would reduce unplanned downtime by 40% and emergency maintenance events by ~70%, with a mean bearing failure warning lead time of 34.2 hours enabling systematic transition from reactive to proactive maintenance practices.

A key finding of this study is that the operational value of an integrated multi-domain ML implementation would exceed the sum of individual model contributions through cross-domain synergistic effects. Stabilizing grinding circuit performance would reduce flotation feed variability, improving both the predictive accuracy of the flotation model and the metallurgical stability of the flotation circuit itself. Predictive maintenance scheduling integrated with process optimization would enable coordination of planned maintenance windows with naturally reduced throughput periods, minimizing production impact. These synergies underscore the importance of designing ML implementations as integrated plant-wide systems rather than as collections of isolated predictive models addressing individual process units independently.

Data infrastructure quality, model interpretability, and effective change management emerged as equally important determinants of implementation success alongside algorithmic performance. The finding that 15-20% of raw sensor data required preprocessing intervention highlights that investment in data quality assurance must precede and accompany investment in model development. Incorporating feature importance analyses, class activation maps, and contextual alert explanations into the operator interface proved essential to achieving operator acceptance and translating model predictions into operational value. Organizations pursuing similar digital transformation initiatives should treat these organizational and infrastructural dimensions as first-order concerns rather than implementation details.

Several limitations of the present study warrant acknowledgment. The conceptual nature of the case study, while grounded in published performance benchmarks and realistic operational parameters, precludes direct validation against measured plant data. The reported performance improvements reflect realistic expectations based on the published literature rather than confirmed outcomes from a specific industrial installation. Future work should focus on validating the integrated framework against operational data from real concentrator implementations, developing more sophisticated uncertainty quantification methods to communicate prediction confidence to operators, and investigating reinforcement learning approaches for autonomous closed-loop optimization that could extend beyond the setpoint recommendation paradigm used in this study. Digital twin technologies that integrate physics-based process models with data-driven ML components represent a particularly promising direction for achieving plant-wide optimization with improved extrapolation capability beyond the training data distribution (Nad et al., 2022; Szmigiel et al., 2024).

As the mineral processing industry continues its digital transformation journey under increasing pressure from declining ore grades, rising energy costs, and tightening environmental constraints, the integrated application of data analytics and machine learning across all major process domains represents not merely an operational improvement opportunity but an emerging competitive necessity. The framework presented in this study provides both a technical roadmap and a realistic assessment of the challenges and prerequisites for successful implementation, contributing to the growing body of evidence supporting the practical viability of Industry 4.0 technologies in mineral processing operations.

Acknowledgments

A part of the results reported in this paper were presented at XX Balkan Mineral Processing Congress BMPC 2026 entitled “Application of data analytics in mineral processing: Towards digital performance prediction and process optimization”.

References

AFRIDI, Y.S., HASAN, L., ULLAH, R., AHMAD, Z., KIM, J.M., 2023. *LSTM-based condition monitoring and fault prognostics of rolling element bearings using raw vibrational data.* Machines, 11(5), 531.

ALDRICH, C., MARAIS, C., SHEAN, B.J., CILLIERS, J.J., 2010. *Online monitoring and control of froth flotation systems with machine vision: A review.* Int. J. Miner. Process., 96, 1-13.

AVALOS, S., KRACHT, W., ORTIZ, J.M., 2020. *Machine learning and deep learning methods in mining operations: a data-driven SAG mill energy consumption prediction application.* Min. Metall. Explor., 37, 1197-1212.

CHEN, X., LIU, D., YU, L., SHAO, P., AN, M., WEN, S., 2025. *Recent advances in flotation froth image analysis via deep learning.* Eng. Appl. Artif. Intell., 147, 109515.

DUIJVENBODE, J.R., BUXTON, M., SHISHVAN, M.S., 2020. *Performance improvements during mineral processing using material fingerprints derived from machine learning – A conceptual framework.* Minerals, 10(4), 366.

DUMAKOR-DUPEY, N.K., ARYA, S., 2021. *Machine learning – A review of applications in mineral resource estimation.* Energies, 14(14), 4079.

ESTAY, H., LOIS-MORALES, P., MONTES-ATENAS, G., RUÍZ-DEL-SOLAR, J., 2023. *On the challenges of applying machine learning in mineral processing and extractive metallurgy.* Minerals, 13(6), 788.

FAN, Y., LV, Z., CHEN, S., CUI, Y., ZHAO, X., XU, Z., 2025. *Condition recognition based on multi-source heterogeneous data and residual temporal network in coal flotation process.* Meas. Sci. Technol., 36(3), 036002.

GHASEMI, Z., NESHAT, M., ALDRICH, C., ZANIN, M., CHEN, L., 2025. *Optimising SAG mill throughput and circulating load using machine learning models: A multi-objective approach for identifying optimal process parameters.* Miner. Eng., 218, 109009.

HODOUIN, D., JÄMSÄ-JOUNELA, S., CARVALHO, M.T., BERGH, L., 2001. *State of the art and challenges in mineral processing control.* Control Eng. Pract., 9(9), 995-1005.

HUANG, G., LI, H., OU, J., ZHANG, M., 2020. *A reliable prognosis approach for degradation evaluation of rolling bearing using MCLSTM.* Sensors, 20(7), 1864.

JOOSHAKI, M., NAD, A., MICHAUX, S.P., 2021. *A systematic review on the application of machine learning in exploiting mineralogical data in mining and mineral industry.* Minerals, 11(8), 816.

KAHRAMAN, A., KANTARDZIC, M., KOTAN, M., 2022. *Dynamic modeling with integrated concept drift detection for predicting real-time energy consumption of industrial machines.* IEEE Access, 10, 104622-104635.

KAMAT, P., SUGANDHI, R., KUMAR, S., 2021. *Deep learning-based anomaly-onset aware remaining useful life estimation of bearings.* PeerJ Comput. Sci., 7, e795.

LIU, X., ALDRICH, C., 2022. *Recent advances in flotation froth image analysis.* Miner. Eng., 188, 107823.

LOPEZ, P., REYES, I., RISSO, N., MOMAYEZ, M., ZHANG, J., 2023. *Machine learning algorithms for semi-autogenous grinding mill operational regions' identification.* Minerals, 13(11), 1360.

LU, D., WANG, F., WANG, S., BU, K., LI, K., MA, X., 2024. *Multimode froth flotation process operating performance assessment based on deep embedded graph clustering network.* IEEE Trans. Ind. Inform., 20(7), 9445-9454.

MAGZUMOV, Z., KUMRAL, M., 2025. *Predicting energy consumption SAG mills through Bayesian generalized linear model and random forest.* Min. Technol., 134(4), 289-311.

MCCOY, J.T., AURET, L., 2019. *Machine learning applications in minerals processing: A review.* Miner. Eng., 132, 95-109.

NAD, A., JOOSHAKI, M., TUOMINEN, E., MICHAUX, S.P., KIRPALA, A., NEWCOMB, J., 2022. *Digitalization solutions in the mineral processing industry: The case of GTK Mintec, Finland.* Minerals, 12(2), 210.

NIQUINI, F.G.F., BRANCHES, A.M.B., COSTA, J.F.C.L., MOREIRA, G.C., SCHNEIDER, C.L., ARAÚJO, F.C., CAPPONI, L.N., 2023. *Recursive feature elimination and neural networks applied to the forecast of mass and metallurgical recoveries in a Brazilian phosphate mine.* Minerals, 13(6), 748.

PARIAN, M., MWANGA, A., LAMBERG, P., ROSENKRANZ, J., 2021. *Ore texture breakage characterization and fragmentation into multiphase particles.* Powder Technol., 378, 164-180.

RAN, Y., ZHOU, X., LIN, P., WEN, Y., DENG, R., 2019. *A survey of predictive maintenance: Systems, purposes and approaches.* arXiv preprint, arXiv:1912.07383.

ROUX, J.D., CRAIG, I.K., 2019. *Plant-wide control framework for a grinding mill circuit.* Ind. Eng. Chem. Res., 58(26), 11585-11600.

SALDAÑA, M., GÁLVEZ, E., NAVARRA, A., TORO, N., CISTERNAS, L.A., 2023. *Optimization of the SAG grinding process using statistical analysis and machine learning: A case study of the Chilean copper mining industry.* Materials, 16(8), 3220.

SALDANA, M., GÁLVEZ, E., SALES-CRUZ, M., SALINAS-RODRÍGUEZ, E., CASTILLO, J., NAVARRA, A., CISTERNAS, L.A., 2026. *A stochastic model approach for modeling SAG mill production and power through Bayesian networks: a case study of the Chilean copper mining industry.* Minerals, 16(1), 60.

SZMIGIEL, A., APEL, D.B., SKRZYPKOWSKI, K., WOJTECKI, Ł., PU, Y., 2024. *Advancements in machine learning for optimal performance in flotation processes: A review.* Minerals, 14(4), 331.

ZHOU, X., HE, Y., 2024. *Deep ensemble learning-based sensor for flotation froth image recognition.* Sensors, 24(15), 5048.