Sustainable Earth Resources Communications

Article

AI-driven optimization under uncertainty for mineral processing operations

William Xu[1,2], Amir Eskanlou[2], Mansur Arief[3], David Zhen Yin[2], Jef K. Caers[2,*]

[1] Materials Science & Engineering, Stanford University, Stanford, CA 94305, USA
[2] Earth & Planetary Sciences, Stanford University, Stanford, CA 94305, USA
[3] Aeronautics & Astronautics, Stanford University, Stanford, CA 94305, USA

* Correspondence: jcaers@stanford.edu

Received: November 11, 2025
Revised: December 17, 2025
Accepted: December 23, 2025
Published: December 25, 2025

Cited as:
Xu, W., Eskanlou, A., Arief, M., Yin, D. Z., Caers, J. K. AI-Driven Optimization under Uncertainty for Mineral Processing Operations. Sustainable Earth Resources Communications, 2025, 1(2): 100-112.
https://doi.org/10.46690/serc.2025.02.07

Abstract: The global capacity for mineral processing must expand rapidly to meet the demand for critical minerals, which are essential for building the clean energy technologies necessary to mitigate climate change. However, the efficiency of mineral processing is severely limited by uncertainty, which arises from both the variability of feedstock and the complexity of process dynamics. To address this uncertainty, the current approach to designing and operating mineral processing circuits emphasizes process stability and control, relying on limited and/or indirect empirical tests, deterministic methods, and expert intuition. Yet a significant portion of valuable minerals is lost in waste streams, translating to millions of dollars of lost revenue and greater potential for environmental damage. To optimize mineral processing circuits under uncertainty, we introduce an AI-driven approach that formulates mineral processing as a Partially Observable Markov Decision Process. We demonstrate the capabilities of this approach in handling both feedstock uncertainty and process model uncertainty to optimize the operation of a simulated, simplified flotation cell as an example. We show that by integrating the process of information gathering (i.e., uncertainty reduction) and process optimization, this approach has the potential to consistently perform better than traditional approaches at maximizing an overall objective, such as net present value. We highlight the power of this approach in scenarios where the dynamics of the system, and subsequently the relationship between the inputs (e.g., feedstock composition, flotation operation settings) and desired outputs (e.g., recovery and grade), are not well known. Our methodological demonstration of this optimization-under-uncertainty approach for a synthetic case provides a mathematical and computational framework for later real-world applications, with the potential to improve both the laboratory-scale design of experiments and industrial-scale operation of mineral processing circuits without any additional hardware.

Keywords: Mineral processing; process optimization; flotation; critical minerals

1. Introduction

In 2024, the average global temperature surpassed the 1.5 °C threshold set by the UN Intergovernmental Panel on Climate Change (IPCC) for the first time in recorded history (World Meteorological Organization, 2024). The IPCC's most recent report is clear: the consequences of human-caused climate change are immense and already taking place, and a clean energy transition is necessary to cut carbon emissions and mitigate these consequences (Lee et al., 2023). To build the requisite clean energy technologies in such a short timeframe will require rapid sourcing of vast quantities of critical minerals (International Energy Agency, 2021). At the same time, many countries have expressed geopolitical and national security concerns regarding critical mineral supply chains—especially for refining and processing capacity, which is heavily concentrated in China.

Yandy Scientific Press

ISSN: 3104-977X

Mineral processing, a key component to the sourcing of critical minerals, faces increasing difficulties stemming from declining ore quality and a growing imperative to improve environmental performance. Additionally, mineral processing is energy- and water-intensive (International Energy Agency, 2021), and so as an industry, it must work towards improving the sustainability of operations to ensure consistency with the eventual goal of combating climate change. Improving the efficiency of mineral processing can serve the dual goal of reducing waste and resource usage while increasing production and thereby revenue.

The efficiency of mineral processing is severely limited by uncertainty from both the variability of feedstock and the complexity of process dynamics (Amini, 2017; Amini and Noble, 2021; Koermer and Noble, 2025; Koermer, 2022). As Bascur (2019) states, "a critical problem in the process of ore extraction is the variability of the different elements that constitute the ore". Traditionally, this uncertainty is addressed with process control and operational intelligence, which go hand-in-hand. Process control seeks to minimize variations in output by adjusting control parameters in response to input variations, and operational intelligence aims to collect information via real-time sensors to inform process control and optimization (Bascur, 2019; Concha A and Bascur, 2024).

The proportional-integral-derivative (PID) control scheme is still the most commonly used process control technique today, but it is only sufficient for single-loop systems, which have one controlled and one manipulated variable. Indeed, despite decades of research into advanced control, Hodouin (2011) and Shean and Cilliers (2011) both emphasize that PID controllers continue to dominate industrial mineral processing, with limited measurable improvements in performance. Even where advanced multivariable or predictive control has been introduced, it is often constrained by the reliability of sensors, model uncertainty, and the difficulty of tuning nonlinear systems (Jovanović and Miljanović, 2015). These reviews collectively highlight that while model predictive control (MPC) has become the de facto "advanced" option, its effectiveness depends heavily on accurate, deterministic models and consistent process dynamics, which is an unrealistic representation of industrial flotation or grinding circuits. As a result, the traditional approach relies on a mix of expert intuition, empirical testing, extensive data collection, and deterministic optimization and control methods (Jiang et al., 2017).

Notably, Hodouin (2011) and Shean and Cilliers (2011) each call for a more holistic or hierarchical view of process optimization, integrating sensors, observers, controllers, and optimizers. Yet, even in these "optimization" frameworks, optimization remains subordinate to control: setpoints are tuned to achieve a target grade or recovery, rather than directly optimizing the operation itself. In practice, this means the underlying objective functions—whether metallurgical or economic—are treated as supervisory layers above fixed control architectures, rather than as part of a unified decision process.

As Jovanović and Miljanović (2015) note, the result is an architecture that can stabilize the process but struggles to adapt optimally when ore characteristics or process conditions shift unpredictably.

Accordingly, the most common approach to operating a mineral processing plant is to view it as a problem of control first and optimization second, even though the overarching goal of the plant is an optimization problem (i.e., maximizing economic profit, sustainability, safety, etc.). We will study the alternative: approaching process operation as a problem of optimization first and foremost, with subproblems of control. In this view, the goal is not to control variations, as process control seeks to do, but rather to optimize the process while accounting for variations—in other words, to leverage uncertainty rather than fight it. The potential value of optimization is undeniable (Bascur, 2019; Ding et al., 2012; Hodouin et al., 2001), and the limitations of current control frameworks suggest that a probabilistic, decision-theoretic formulation may be necessary to achieve it.

A few projects have demonstrated the value of considering uncertainty in the optimization of mineral processing. Välikangas et al. (2025) used sensitivity analysis and uncertainty propagation to understand the influence of feedstock variability and inform data collection. Koch and Rosenkranz (2020) and Amini (2017, 2021) presented stochastic approaches that outperformed deterministic methods at designing mineral processing circuits. Jiang et al. (2017), Koermer and Noble (2025), and Koermer (2022) used reinforcement learning (RL) and machine learning (ML) to determine optimal operating conditions given unknown process dynamics at steady-state. This body of prior work forms a strong basis for considering uncertainty in optimizing mineral processing.

It is important to note that this paper focuses on decision-making and optimization under uncertainty, which is what RL is designed to do, rather than data-driven modeling, which is the goal of ML. While there has been a growing body of work applying AI to mineral processing, these efforts have been almost entirely focused on ML (e.g., improving empirical process models, predicting metallurgical outcomes from sensor data) rather than RL (McCoy and Auret, 2019; Bai et al., 2025). Although model-based RL (which we employ) can incorporate ML for improved process modeling, the core objective is to learn operational policies that optimize performance over time, not to generate predictive models. To date, the mineral processing literature lacks a framework for explicitly integrating uncertainty reduction with optimization, particularly one that accounts for uncertainty arising from both feedstock variability and process complexity.

In this work, we aim to show that mineral processing operations can be framed as a problem of optimization under uncertainty, and outline the features of this approach. We then develop a mathematical formulation of a simplified flotation cell that incorporates both feedstock uncertainty and process uncertainty to inform optimization over time via data collection. We use synthetic scenarios to demonstrate the capability of this framework for optimizing the operation of a flotation cell in comparison to PID and MPC approaches, particularly in cases of significant feedstock and process uncertainty. This paper

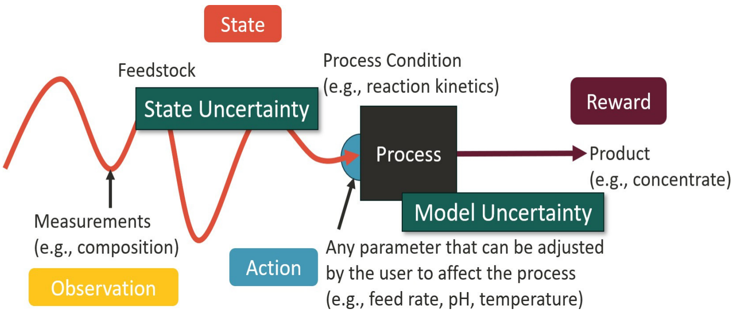

Fig. 1. Bare-bones framing of a mineral process (e.g., flotation) with a variable feedstock and complex process framed in terms of decision-making under uncertainty.

serves as a demonstration of a comprehensive mathematical approach to optimizing mineral processing under uncertainty, rather than attempting to claim that this approach "performs better" than existing approaches. Having highlighted the potential for this approach in various synthetic scenarios, we will discuss its potential application to real-world test cases.

2. Features of an optimization under uncertainty approach

2.1. Framing a mineral process

Any mineral process can be viewed as a variable feed stream passing through a complex process to turn into a product. As our goal is optimization rather than to accurately describe the intricacies of a given process, we model any mineral process as a system with uncontrolled inputs, outputs, and control parameters that change how inputs translate into outputs (see Fig. 1).

To put this in the language of decision-making under uncertainty, the inputs and the process dynamics can be conceptualized as the states of the system, control parameters as actions, the outputs as the reward, and any measurements taken to better ascertain the conditions of the process as observations.

In this framework, the key uncertainties, feedstock variability and process complexity, can be classified as state uncertainty and model uncertainty, respectively (labeled with teal boxes in Fig. 1). Feedstock variability makes the feedstock composition and mineralogy, or the state, uncertain, since we cannot measure every aspect of the feedstock at every point in time and space. Process complexity means we cannot know exactly how inputs (states and actions) translate into the output (reward), so the model we use to describe this causal relationship is uncertain.

2.2. The mathematical framework

The principal mathematical framework for decision-making under uncertainty problems is called a Partially Observable Markov Decision Process (POMDP). A POMDP models a sequence of actions in the real world at the same time as information is gathered. This approach is now common in AI applications such as aircraft collision avoidance, self-driving cars, and robotics, and is similar to the AI used in chess and other games (Xiang and Foo, 2021; Xiang et al., 2021).

Mathematically, a POMDP is defined by a tuple $\langle S, A, O, T, R, Z, \gamma \rangle$, where S is the state space, A is the action space, O is the observation space, T is the transition function, R is the reward function, Z is the observation function, and γ is the discount factor (Arief et al., 2025). The POMDP framework builds upon the Markov Decision Process (MDP) framework, a broad approach to sequential decision making in stochastic environments that forms the basis for most reinforcement learning. In all MDPs, an agent makes decisions (referred to as actions) at discrete timesteps, which influence how the system transitions from one state to the next. The current state of the system and action taken at a given timestep result in a reward that the agent receives, which represents the optimization objective. The way that actions are chosen based on the current state or observations of the system is called a policy. An intelligent agent decides a policy by learning from interacting with the system over time (Kochenderfer et al., 2022).

The POMDP framework assumes that the true state of the system cannot be known. Instead, the agent holds a belief over possible states, which is represented as a probability distribution informed by indirect, incomplete, and/or noisy observations. More information about MDPs and POMDPs can be found in the textbook Algorithms for Decision Making (Kochenderfer et al., 2022).

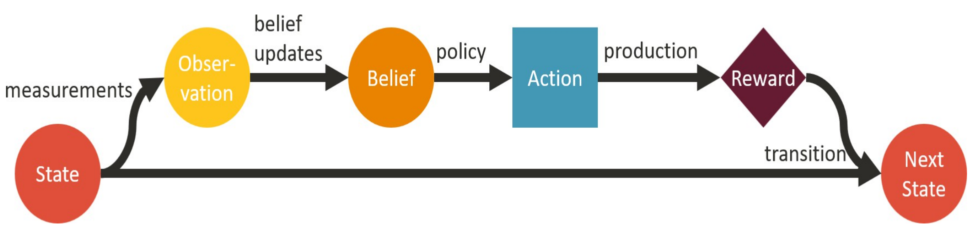

Fig. 2. Simplified diagram depicting the key components of a POMDP and how they progress at each timestep.

The diagram in Fig. 2 shows the key components of a POMDP at a given timestep, revealing its sequential nature. The labels given in Fig. 1 for a generic mineral process map directly onto this framework.

2.3. Belief: A stochastic representation of an uncertain state

The belief is the likelihood of being in a given state and is typically represented by a probability distribution over states. Although the belief is not technically part of the POMDP formulation itself, it is a crucial component of decision-making under uncertainty and how POMDP solvers navigate problems.

The AI literature uses belief instead of probability as a nomenclature to identify uncertainty.

The intelligent agent forms a belief of the state based on measurements that it receives. As shown in Fig. 3, the belief evolves over time as more measurements are collected. Uncertainty, represented by variance values, is captured in the belief's nature as a probability distribution rather than a discrete quantity or set of quantities.

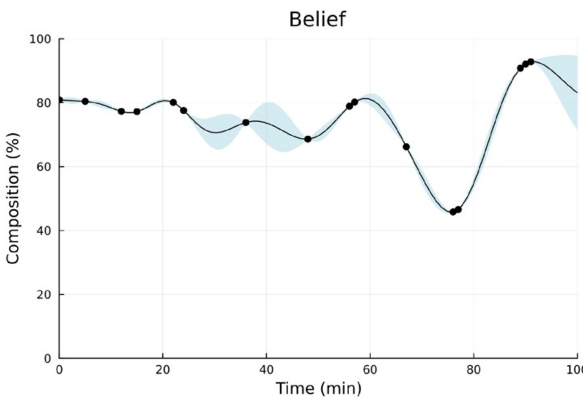

Fig. 3. An example of a belief at the end of a simulation, with occasional measurements (black dots) informing the belief and its associated uncertainty (blue regions).

2.4. Reward: The optimization objective

The reward function describes the optimization objective. Value judgments from experts are necessary to define what the ultimate goal of the optimization should be. As shown in Fig. 4, the reward captures the inherent tradeoff between the cost of measuring and the cost of not having information.

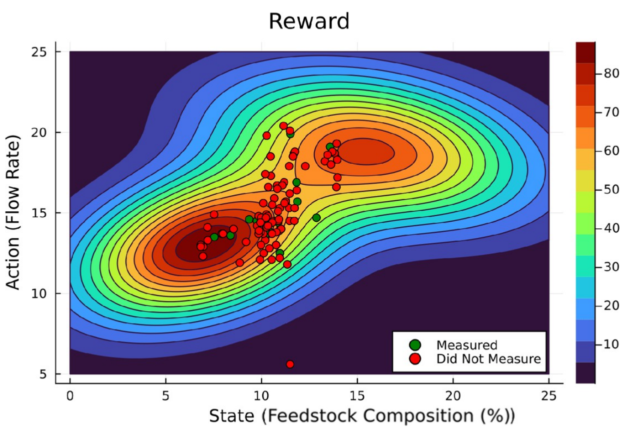

Fig. 4. An example of a reward surface (reward as a function of state and action) showing how the intelligent agent balances a cost associated with taking measurements with the cost of choosing a poor action. Actions are taken as the state fluctuates in time, with red dots representing times when measurements were not taken, and green dots representing times when measurements were taken.

3. Mathematical formulation of a flotation cell

Now that we have established the general concept of applying an optimization-under-uncertainty approach to mineral processing, we can formulate the operation of a flotation cell as a POMDP. A simple formulation is shown in Fig. 5, labeling a few components of a flotation cell under this framework.

For now, we represent the flotation cell as a batch process, where one batch is processed at each timestep. This allows for straightforward experimental validation at the bench scale. The formulation can be adapted for a continuous process, as is typical at the industrial scale. The values used throughout our flotation cell formulation and implementation are meant to roughly reflect typical values for phosphate flotation as an example.

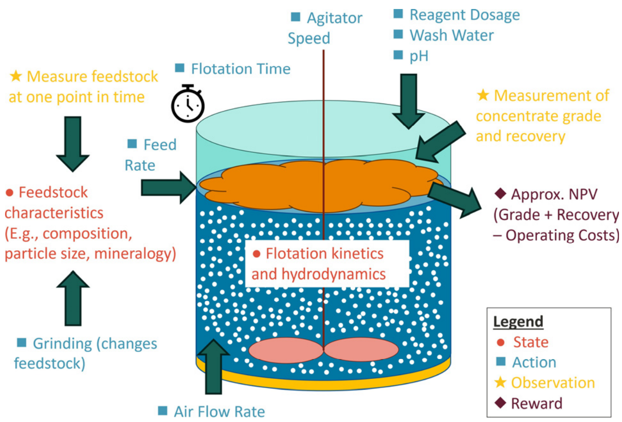

Fig. 5. A simple POMDP formulation of a flotation cell. Additional possible state variables and control parameters beyond the scope of the formulation in this paper are included as examples.

3.1. State

The state is represented by the following variables:

- Feedstock composition $c \in [0.0, 42.2]$ (%)
- Concentrate recovery $r \in [0.0, 100.0]$ (%)
- Concentrate grade $g \in [0.0, 42.2]$ (%)
- Timestep T

For simplicity and clarity of presentation, we represent the feedstock characteristics with one variable, the average composition for a given batch. Feed and concentrate grade can only achieve a maximum of 42.2% because they represent P2O5 grade, and pure fluorapatite (the primary phosphate-bearing mineral) has a P2O5 grade of 42.2%. The state includes time mostly as a technicality, since the transition function (which is a function of the state and action) depends on time.

In this case, there is no explicit component of the state that corresponds to the "flotation kinetics" as depicted in

Fig. 5 or "process condition" as depicted in Fig. 1. The concentrate recovery and grade implicitly capture these components, as discussed later in Section 4, which is why they are used as components of the state instead. If, instead of average composition, a size-by-liberation matrix were used to represent the feedstock, then the liberation and size-dependent kinetic rate constant would become an internal "process condition" state variable.

3.2. Actions

Actions are control parameters that can be adjusted to change the operating conditions of the flotation cell, as well as the decision to make measurements.
For simplicity and clarity of presentation, we choose two control parameters as our action set.

- Flotation time $t \in [5.0, 15.0]$ (min)
- Air flow rate $f \in [50, 150]$ (L/hr)
- Measure feedstock

For the model uncertainty tests in Section 7.1, the Measure feedstock action is always set to true. Measure feedstock becomes a choice between true and false for the feedstock uncertainty tests in Section 7.2.

3.3. Transition function

The transition function describes how the current state transitions to the next state as a function of the current state and action. In other words, it corresponds to a forward model that describes how the system changes in time.

The transition model is described by a combination of the following:

- the simple kinetic model (see Section 4)
- a stochastic representation (e.g., a Gaussian process) of the feedstock composition fluctuating in time
- stochastic representations (e.g., Gaussian process) of the errors between the kinetic model and the true grade and recovery

Note that in this formulation, transition probabilities are only dependent on the state, not the action. Also, we consider actions to have a deterministic effect. In other words, if we were to know the state, then choosing an action would deterministically result in a given reward. Transition uncertainty can be introduced by making actions stochastic (in other words, imprecise).

3.4. Observations

The observations are:

- Average feedstock composition (can be null)
- True recovery and grade

For the model uncertainty tests in Section 7.1, full observations of the state are received at every timestep, so the implementation technically reduces to an MDP (no state uncertainty, only transition uncertainty). For the feedstock uncertainty tests in Section 7.1, observations of feedstock composition are only received when the measure feedstock action is taken. Otherwise, no information about the feedstock is collected at that timestep.

3.5. Observation function

The observation function is the likelihood of an observation given a state. For the scope of this paper, since the goal is more to demonstrate the approach, the observation function is just a delta function (i.e., an exact observation). The observation returned at each timestep is simply the true state.

3.6. Reward

Here, we consider the reward to be an approximation of the net present value (NPV) of the process. We use the Moroccan phosphate industry (i.e., the OCP Group) as an example. The specific formula for the reward defined in Eq. 1 uses back-of-the-envelope estimates for the current production of phosphate concentrate as a function of recovery, the price as a function of grade, and the operating costs (U.S. Geological Survey, 2024; World Bank, 2025). The operating cost formula in Eq. 2 is not intended to reflect realistic operating costs, but rather is designed to create a global optimum from the tradeoff between grade, recovery, and operating costs. Examples of the reward at a fixed feedstock composition depicted in Fig. 8 exhibit this tradeoff.

$$reward = \frac{50g\ [\$/t] \cdot 35r\ [Mt/yr]}{100\ [timestep/yr]} - OPEX\ \left[\frac{\$M}{timestep}\right] \quad (1)$$

$$OPEX = \frac{1}{2}t + \frac{1}{50}f\ \left[\frac{\$M}{timestep}\right] \quad (2)$$

4. Simple flotation model

We lay out a simple flotation model to describe the system to be optimized. This model of the system corresponds to the transition function in a POMDP formulation.

In flotation, recovery and grade are the key performance metrics, and there is a natural tradeoff between the two. (As recovery approaches 100%, concentrate grade approaches feed grade, and as concentrate grade approaches 100%, recovery approaches 0%. A grade-recovery curve then essentially forms a Pareto front.)

We model the black box mechanical flotation cell with empirically-inspired equations. The goal is to capture broad relationships between the inputs, control parameters, and outputs, rather than to be accurate.

The instantaneous recovery r and the instantaneous concentrate grade g are both reported in percentage (i.e., ranging from 0 to 100),

$$r(k,t,f) = 100\frac{kt}{1+kt}\frac{f}{f+10} \quad (3)$$

$$g(c,k,t,f) = c\left[1 + \left(1 - \frac{c}{42.2}\right)\left(1 - \frac{\exp(-kt/10)}{1+\exp(4-0.04f)}\right)\right] \quad (4)$$

where c is the feedstock composition (i.e., feed grade) in percentage, k is the flotation rate constant in min^{-1} (set to 1), t is the flotation time in minutes, and f is the air flow rate in L/hr.

The exact equations are mostly arbitrary, with numbers that very roughly correspond to phosphate flotation. The grade equation, in particular, is set up to refer to P_2O_5 grade. They are designed to highlight the tradeoffs present in flotation cells in a simple fashion, as can be seen in the plots of this simple kinetic model in Figs. 6a and 7a.

Although we are aware of much more sophisticated models, we intentionally choose a simplistic model of flotation to more clearly demonstrate and display the features of a POMDP approach.

We do not expect our simple kinetic model to accurately capture flotation dynamics, but rather to serve as a reasonable first guess or prior. We can capture all inaccuracies in our model, as well as any inaccuracies in any measurement of the true grade and recovery in an error function. Then, we can represent the true grade or recovery as the sum of the kinetic model and some stochastic error function, as depicted in Fig. 6 and 7. The true grade and recovery result in a true reward function as well, shown in Fig. 8.

For the synthetic cases in this paper, the "true" error function is generated stochastically to produce a ground truth grade and recovery that represent "reality". Just like in real life, this ground truth is unknown to the intelligent agent seeking to optimize the flotation cell, but can be explored through measurements.

5. Belief update

As introduced in Section 2.3, the intelligent agent updates the belief to learn the true grade and recovery and improve upon our prior model of the system. In the flotation problem, the belief is represented by:

- Gaussian process of feedstock composition
- Gaussian processes of the grade and recovery error functions

Uncertainty is represented stochastically with Gaussian processes. Actions (i.e., setting the air flow rate and flotation time) are chosen at each point in time as the feedstock composition fluctuates. The measured grade and recovery then inform the updated belief (examples shown in Fig. 9), which helps improve decision-making. The Gaussian processes in the belief are updated by sequentially refitting them to include the new data.

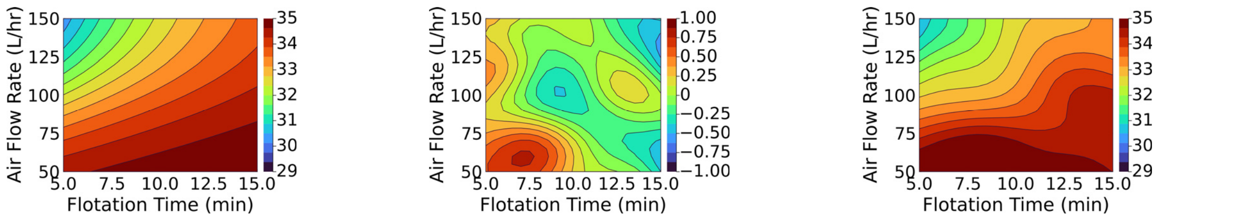

(a) Kinetic model (b) Error function (c) True grade

Fig. 6. An example of the grade (%) as a function of the actions (air flow rate and flotation time) at a fixed feedstock composition.

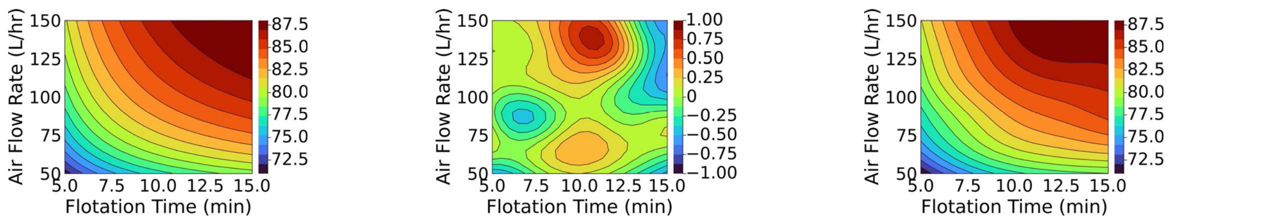

(a) Kinetic model (b) Error function (c) True recovery

Fig. 7. An example of the recovery (%) as a function of actions (air flow rate and flotation time) at a fixed feedstock composition.

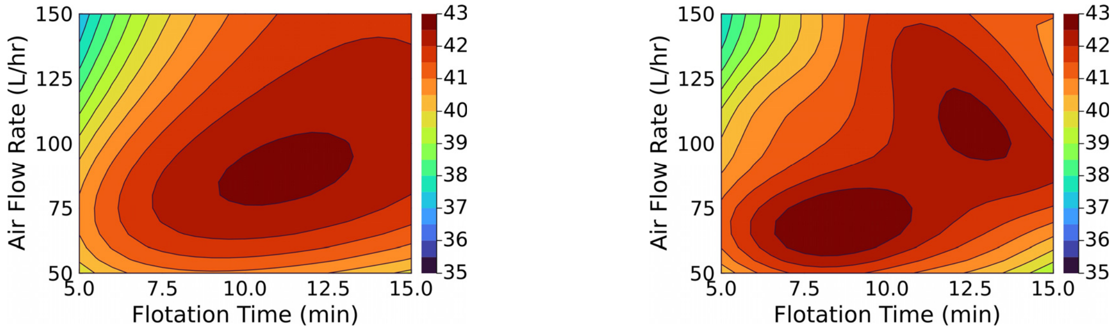

(a) Kinetic model (b) True reward

Fig. 8. Examples of the reward (NPV) as a function of actions (air flow rate and flotation time) at a fixed feedstock composition.

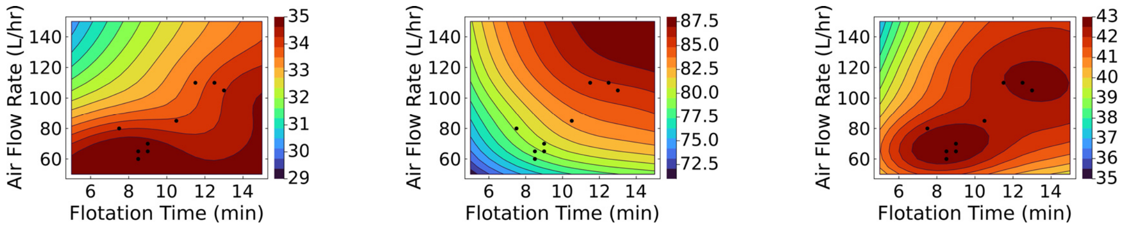

(a) Grade belief (%) (b) Recovery belief (%) (c) Reward belief (NPV)

Fig. 9. Beliefs represented by Gaussian processes that are progressively updated as new data is collected over time. Black dots represent collected data.

6. Simulation setup

We investigate the extent to which the performance of different approaches is affected by feedstock (i.e, state) and process (i.e., model) uncertainty. We do so by running simulations of 100 timesteps, which represent 100 flotation batches processed over one year. All simulations were run on a 13th Gen Intel Core i9 using Windows 11.

6.1. Establishing the baseline

In a POMDP framework, a given approach to choosing actions based on the current belief is called a policy. Control and optimization algorithms can be considered types of policies. As stated in Section 1, two commonly used deterministic methods, PID and MPC, are used as a frame of reference. Although MPC is typically used as a control method, we implement MPC using an optimization approach with the goal of maximizing the reward, not just the grade and recovery, to serve as a direct comparison to the POMDP approach. Here, MPC uses the simple kinetic model (prior) as a fixed model throughout the simulation, unless otherwise stated.

6.2. Performance metric

Since we are aiming for the goal of process optimization, the reward (proxy for NPV, established in Section 3.6) is used as the metric of comparison.

6.3. POMDP solver

To solve the problem we have now formulated as a POMDP, we use a well-established online solver called Partially Observable Monte Carlo Planning (POMCP) (Silver and Veness, 2010). To determine a policy, POMCP uses Monte Carlo tree search (MCTS), a common algorithm for deciding a course of action from many possible futures. The most well-known application of MCTS is in gameplaying AI.

POMCP requires a discrete action space, so the flotation time is discretized with a step size of 0.5, and the air flow rate is discretized with a step size of 5.0. These values reflect the approximate sensitivity of equipment in laboratory and industrial applications, and there is limited practical use in varying the flotation time by less than 30 seconds and the air flow rate by less than 5 L/hr. However, we do explore the impact of the action space sizing on the performance of POMCP in Appendix D Table D1.

For applications to continuous action spaces and more expansive state and action spaces in general, we recommend the use of POMCPOW, which has been well-established to address large state and action spaces (Sunberg and Kochenderfer, 2018).

7. Demonstration of optimization-under-uncertainty approach

7.1. Optimization under model uncertainty

To evaluate optimization under model uncertainty, we consider three cases of differing degrees of model uncertainty: where the simple flotation model has high, medium, and low accuracy. The model accuracy reflects how closely the kinetic model matches reality. Examples are plotted in Fig. 10.

The median results (over 100 simulations) in Table 1 show that although MPC performs better than the POMDP approach when the model is accurate, its performance lags behind the POMDP approach as the model accuracy decreases. Table A1. shows that in the low accuracy scenario, there are even cases where MPC performs worse than a PID controller.

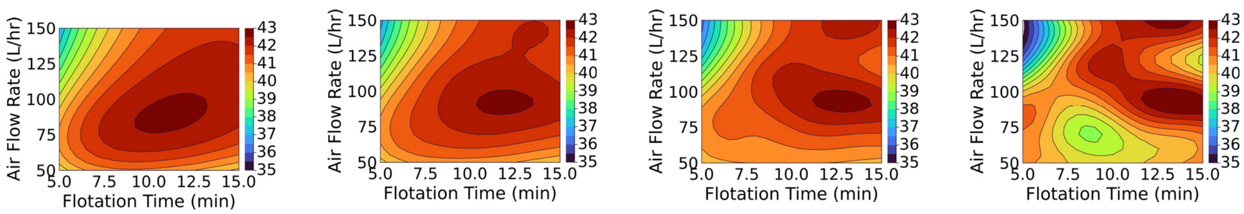

(a) Kinetic model (b) High accuracy (c) Medium accuracy (d) Low accuracy

Fig. 10. Reward functions (NPV) of varying degrees of similarity to the kinetic model.

https://doi.org/10.46690/serc.2025.02.07

Table 1. Median performance of MPC and POMDP approaches relative to PID controller baseline when varying model accuracy (i.e., increasing model uncertainty).

Model accuracy	Model predictive control			POMDP approach		
	High	Med	Low	High	Med	Low
rel. recovery [Δ%]	-3.6	-3.8	-4.3	-3.1	-3.3	-3.1
rel. grade [Δ%]	+0.4	+0.5	+0.9	+0.4	+0.7	+1.9
rel. reward [Δ$M/yr]	**+119**	**+126**	**+126**	**+95**	**+129**	**+283**

7.1.1. Effect of feedstock variability

Now, we consider the effect of feedstock variability on optimization under model uncertainty when the feedstock is fully known (i.e., measured at every timestep). We test different variances of the feedstock composition under a set of grade and recovery surfaces for which, when the feedstock composition is constant, MPC and the POMDP approach have near-equivalent performance. This corresponds roughly to the medium model accuracy case (see Fig. 10). (A log variance of -3 can be considered near-constant feedstock, as can be seen in the plots of sample feedstock composition signals in Appendix B Fig. B1.)

As can be seen in Fig. 11, as the feedstock variance increases, the relative reward of the POMDP approach increases. Detailed results (shown in Table C1) indicate that decreasing the feedstock composition correlation length also seems to lead to an increase in the relative reward of the POMDP approach. However, this only occurs at high variance, and the effect is less pronounced.

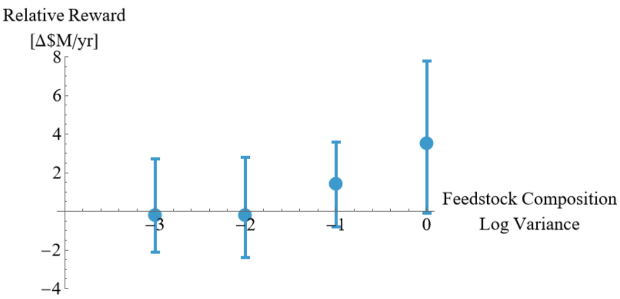

Fig. 11. Median reward of POMDP approach relative to MPC. Error bars represent 20th and 80th percentiles. Log correlation length is fixed at 2.

7.2. Optimization under state uncertainty

Next, we consider the independent effect of state (i.e., feedstock) uncertainty. We fix the grade and recovery surfaces and choose a high-variance feedstock composition signal. Then, the number of feedstock measurements across the simulation is varied. Taking fewer measurements corresponds to higher state uncertainty. A high, medium, and low model accuracy scenario is considered.

As shown in Fig. 12, for a high variability feedstock, the POMDP approach improves at a faster rate than MPC as the number of measurements (i.e., information gathered) increases. The POMDP approach is never able to outperform MPC at high model accuracy (consistent with the results in Section 7.1), nor in any case with zero measurements. However, after a certain number of measurements in the medium and low accuracy cases, the POMDP approach is able to surpass MPC. For the medium accuracy case, crossover occurs at n = 30, and for the low accuracy case, crossover occurs at n = 3 (where n = number of measurements).

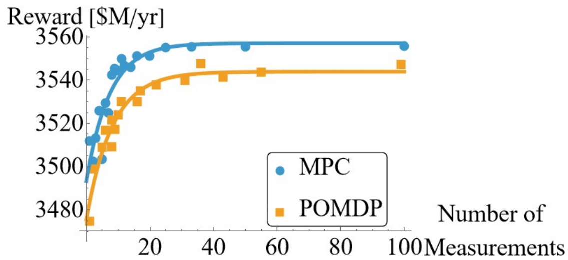

(a) High model accuracy

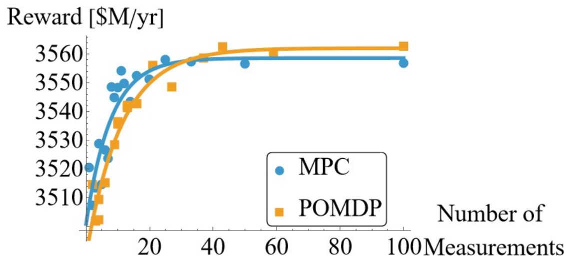

(b) Medium model accuracy

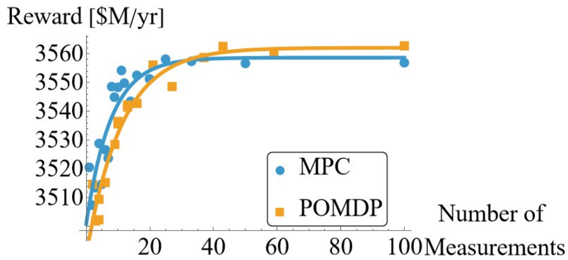

(c) Low model accuracy

Fig. 12. Performance of MPC and POMDP approaches for a high-variance feedstock composition signal and three fixed grade and recovery surfaces, varying the number of measurements. Lines are exponential curves of best fit.

7.3. Uncertainty reduction

Finally, we analyze the capacity of different approaches to reduce uncertainty. To put MPC and the POMDP approach on even footing, MPC is given the same learning capacity as the POMDP approach. In other words, an adaptive MPC approach is implemented, where the same Gaussian processes used to form the belief of the

POMDP are used to update the model that the adaptive MPC approach uses. We now conduct the same test as in Section 7.2, for another set of fixed grade and recovery surfaces with medium model accuracy and high-variance feedstock composition.

As shown in Fig. 13, adaptive MPC does not perform significantly better than MPC, even in a scenario where the POMDP approach surpasses MPC. Fig. 14 shows the Gaussian process models of the system at the end of the simulation when measurements are always taken. Points represent samples taken over the simulation. Fewer points shown indicate more repeated sampling. Evidently, adaptive MPC explores a much smaller space than the POMDP approach.

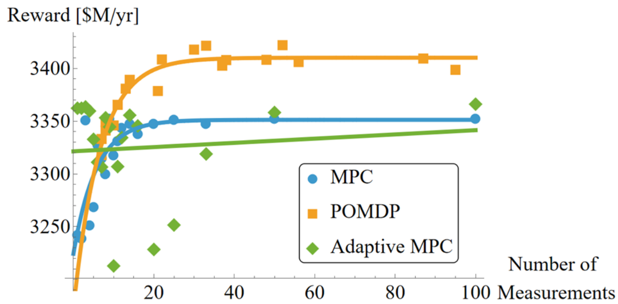

Fig. 13. Performance of MPC, POMDP, and adaptive MPC for a medium model accuracy and high-variance feedstock compositional signal case, varying the number of measurements. Lines are curves of best fit. Since the data for adaptive MPC is noisy, a linear rather than exponential fit is used.

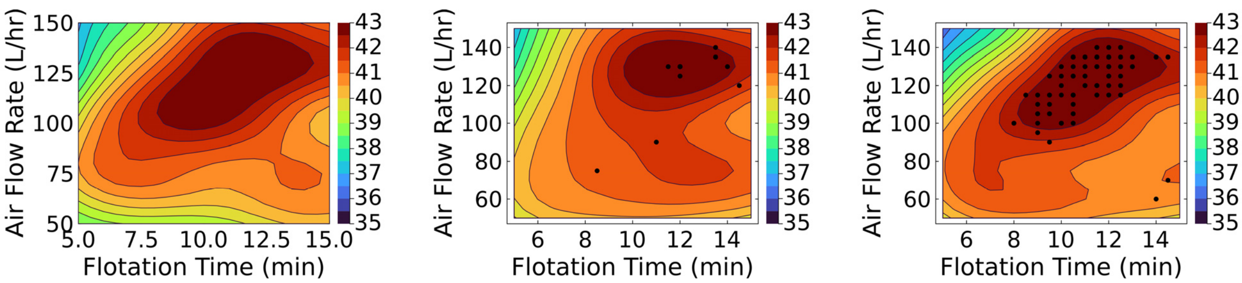

(a) True reward (b) Adaptive MPC (c) POMDP approach

Fig. 14. Comparison of the Gaussian process predictions for the reward function (NPV) with the true reward function. Adaptive MPC exhibits more repeated sampling, while the POMDP approach explores more and achieves a more accurate process model.

8. Discussion

8.1. Analysis of results

The test cases presented reveal that an optimization-under-uncertainty approach using a POMDP framework is more successful at handling cases of significant model uncertainty, adjusting to significant feedstock variability under model uncertainty, and utilizing limited information under state (i.e., feedstock) uncertainty.

In Section 7.1, we see that for MPC, a high-quality deterministic optimization algorithm whose performance depends heavily on the quality of the model (i.e., having an accurate picture of the system), performance correlates strongly with the accuracy of the model. In the high accuracy case, the predicted optimal region largely overlaps with the actual optimal region, so using the kinetic model suffices to inform optimal decision-making. As a result, MPC is not only sufficient, but it outperforms the POMDP approach.

But as soon as the model deviates from reality, MPC struggles, while a POMDP approach consistently performs well, even when the flotation model has low accuracy. In the medium accuracy case, there is some overlap in optimal regions, but there is now a new optimal region that is not captured by the kinetic model. In the low accuracy case, there is almost no overlap in optimal regions, so relying on the kinetic model would result in suboptimal decision-making.

(Note that both MPC and the POMDP approach consistently find solutions that result in worse recoveries but better grades than the PID controller. This is because these two approaches are both seeking to optimize the reward, which takes into account operating costs, while the PID controller just seeks to maintain a high grade and recovery setpoint regardless of other factors.)

Even with relatively high model accuracy, small inaccuracies in the model can compound when the feedstock has high variability, as shown in Section 7.1.1. The POMDP approach is able to adjust and better account for feedstock variability by developing a more accurate model over time. These results imply that, in this case, having more accurate information about the process model is more important than having perfect information about the feedstock in optimizing performance—in other words, model uncertainty matters more than feedstock uncertainty.

The importance of model uncertainty over feedstock uncertainty is further supported by the results in Section 7.2. The POMDP approach almost always achieves a greater reward than MPC in the low model accuracy case, while the opposite is true in the high model accuracy case, regardless of the level of feedstock uncertainty. If feedstock uncertainty had a greater influence on performance, we would expect to see a crossover of the curves for all depicted model accuracies. However, feedstock uncertainty is still relevant, which is particularly clear in the medium model accuracy case (see Fig. 12b). The greater rate of improvement for the POMDP approach as the number of measurements increases indicates that it can translate additional feedstock information into a greater reward increase than MPC can. In other words, a POMDP approach "learns" faster than MPC as more data is acquired.

The faster learning rate of the POMDP approach, thanks to its explicit integration of uncertainty reduction into the optimization (balancing exploration with exploitation), is especially evident in the results shown in Section 7.3. Figs. 13 & 14 together clearly demonstrate that although adaptive MPC is able to learn a more accurate process model over time, its capacity to balance exploration with exploitation is limited, hindering its performance in the long run.

8.2. Summary of POMDP advantages

The advantages of the POMDP approach are best summarized by comparing it to existing approaches in order of complexity, and highlighting how it incorporates more holistic and forward-looking features.

1. Direct search (or other deterministic approach) vs. POMDP: POMDP performs global optimization in time, not one-shot optimization.

2. Adaptive MPC vs. POMDP: POMDP looks at longer horizons and performs stochastic optimization.

3. Robust optimization (probabilities encoded in constraints, or any other stochastic optimization approach) vs. POMDP: POMDP incorporates exploration for uncertainty reduction to inform future decision-making.

Even though robust/stochastic optimization approaches do consider uncertainty, they do not explicitly contain a mechanism to reduce uncertainty over time. So, although a comparison to a robust/stochastic optimization approach is not presented here, Section 7.3 still demonstrates the inherent capability of a POMDP approach to perform uncertainty reduction, which improves performance in the long run.

8.3. Implications for real systems

These results emphasize the power of the POMDP in handling applications with high degrees of uncertainty, especially when the model of the system (i.e., the description of process dynamics) has low accuracy. Note that although these results are for simple test cases, we would expect that in more complex, real systems, there would be significant model uncertainty, which a POMDP approach is better equipped to handle than deterministic optimization. This approach has immediate relevance for potential application in improving the design-of-experiments of bench-scale flotation cells as well as in optimizing the operation of industrial-scale flotation cells without the need for any retrofitting. And although we use the example of a flotation cell in this work, the framework can be adapted to any process unit, or scaled up for optimization of a flotation circuit, an entire mineral processing circuit, an integrated mine and processing plant, and even an entire vertical mineral supply chain.

Beyond improving existing operations, using a POMDP approach could aid in the design and optimization of versatile, highly adaptable mineral processing plants that were previously not possible due to feedstock variability and process complexity. Such a processing plant could even obviate the need for blending, as it could adjust operational settings for a wide range of possible feedstocks.

Lastly, the use of a solver based on Monte Carlo tree search ensures that the optimization decisions it makes are interpretable. The goal is to aid mineral processing experts in making decisions, rather than to take over decision-making with AI altogether.

9. Conclusions

We have demonstrated that mineral processing can be framed as a problem of optimization-under-uncertainty, presenting a mathematical formulation of a simplified flotation cell using the POMDP framework. A range of synthetic test cases demonstrates the utility of the POMDP approach compared to deterministic approaches like MPC, especially in cases with significant feedstock (i.e., state) and process (i.e., model) uncertainty. Through belief updating, the POMDP formulation is designed to incorporate both feedstock (i.e., state) and process (i.e., model) uncertainty when performing optimization, which enables it to handle conditions of significant uncertainty more readily than deterministic methods such as MPC. Thus, an optimization-under-uncertainty approach is particularly well-suited for optimizing mineral processing.

Our work has presented the following advancements:

1. Mineral processing can be framed as a problem of optimization-under-uncertainty, as demonstrated by our mathematical formulation of a simplified flotation cell.

2. Framing the ultimate goal as optimization, rather than control, is better suited to handling uncertainty in mineral processing. MPC's performance over PID alone emphasizes this.

3. The representation of an unknown state and model via a belief, and the integration of real-time data collection into process optimization via belief updating, is fundamental to how a POMDP approach models uncertainty and the reduction of uncertainty over time. In other words, an intelligent agent learns a more accurate model of process dynamics and estimation of feedstock variability over time to improve process optimization.

4. Synthetic test cases confirm that in scenarios with significant feedstock (i.e., state) and process (i.e., model) uncertainty, a POMDP approach performs better than deterministic approaches like MPC.

Future work is needed to apply this approach to real-world test cases. The nearest-term application could be for the design of experiments of bench-scale flotation. The formulation and code as presented in this paper could be directly applied, along with a few tweaks to add complexity, such as a more specific, well-developed flotation model, the inclusion of feedstock characteristics beyond an average composition, a larger set of control parameters, and a more case-specific reward function. Similarly, this work could be readily applied to optimizing the operation of an industrial-scale flotation cell, with similar tweaks, as well as swapping out flotation time for feed rate to consider a continuous flotation process. The more fruitful application would be in the design and operation of mineral processing circuits, which could apply the same approach, but would require the development of a new mathematical formulation. We hope that this will inspire future work to improve the efficiency and sustainability of industrial-scale mineral processing facilities.

Author Contributions

W.X.: conceptualization, methodology, software, investigation, validation, visualization, writing—original draft preparation, writing—reviewing and editing; A.E.: conceptualization, methodology, writing—reviewing and editing; M.A.: software, validation; D.Z.Y.: supervision, writing—reviewing and editing; J.K.C.: supervision, conceptualization, methodology, writing—reviewing and editing. All authors have read and agreed to the published version of the manuscript.

Funding

This research was funded through the Mineral-X Industrial Affiliates program, which is supported by affiliate members KoBold Metals, Bidra VC, Ero Copper, Fleet Space Technologies, Ideon Technologies, and Xcalibur Smart Mapping.

Data Availability Statement

Data and code provided upon request.

Acknowledgments

The authors would like to thank Yale Zhang (Hatch) for valuable discussions.

Conflict of interest

The authors declare no conflict of interest.

Use of AI and AI-assisted Technologies

During the preparation of this work, the authors used Github Copilot to generate well-established functions and for general debugging, and ChatGPT to assist in conducting the literature review and polish the writing. After using these AI tools, the authors reviewed and edited the content as needed and took full responsibility for the content of the published article.

Appendix A

Detailed results of the relative reward for the low accuracy case described in Section 7 are presented in Table A1. The p50 (median) values are the same as the values presented in Table 1.

Table A1. Comparison of reward between different control and optimization approaches.

Policy	Relative Reward (Δ$M/yr)		
	p20	p50	p80
PID Controller (Baseline)	-	-	-
Model Predictive Control (MPC)	-14	**+126**	+258
Online Solver (DMU approach)	+188	**+283**	+381

Appendix B

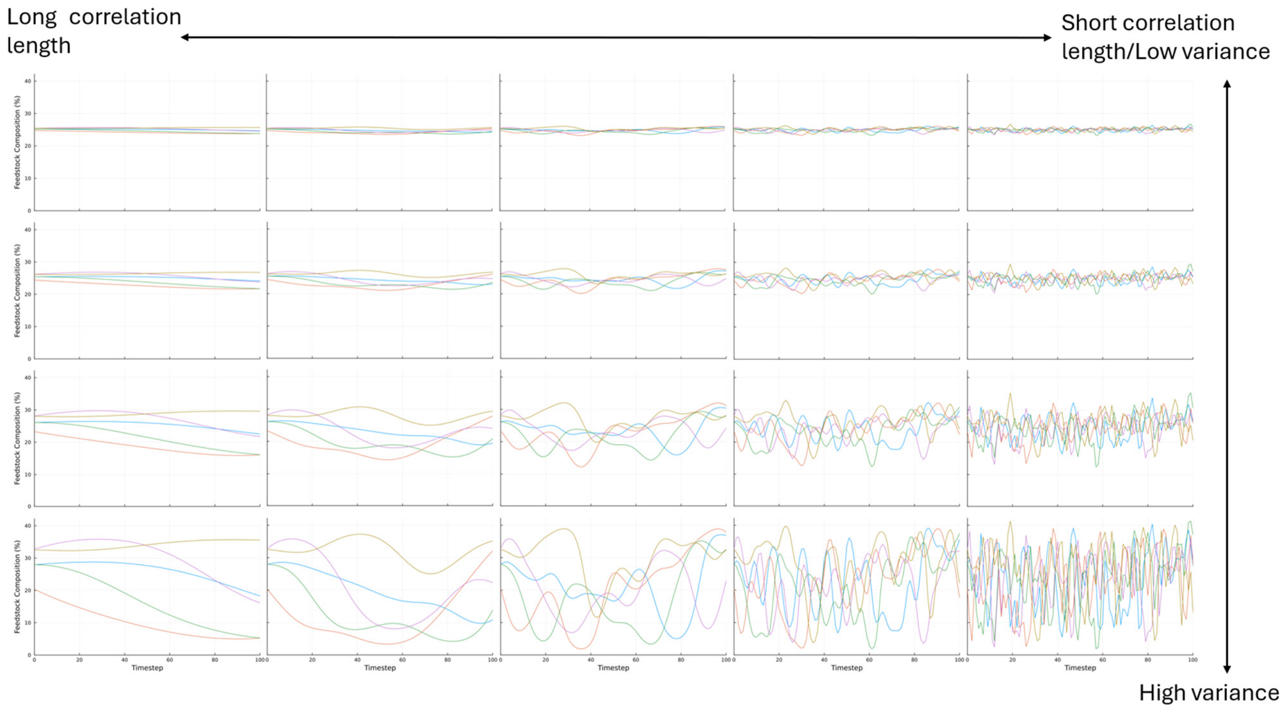

Fig. B1. Five example feedstock compositions over time with different variances and correlation lengths. The log variance increases from -3 to 0 from top to bottom, and the log correlation length decreases from 4 to 0 from left to right. The specific curves are not important; the intent is to show the change in the shape of the curves.

Appendix C

Detailed Results for Feedstock Variability (Detailed results for Section 7.1.1):

Table C1. Reward of POMDP approach relative to MPC (in Δ$M/yr) when varying feedstock composition correlation length and variance.

Log	Feedstock Composition Log Variance											
Corr.	-3.0			-2.0			-1.0			0.0		
Len.	p20	p50	p80	p20	p50	p80	p20	p50	p80	p20	p50	p80
4.0	-6.0	**0.1**	5.4	-5.8	**-1.3**	5.5	-3.5	**0.8**	6.5	-2.5	**1.6**	9.8
3.0	-3.8	**-0.2**	3.7	-3.7	**0.3**	2.9	-2.3	**1.2**	4.5	-1.7	**3.0**	9.5
2.0	-2.1	**-0.2**	2.7	-2.4	**-0.2**	2.8	-0.8	**1.4**	3.6	-0.1	**3.5**	7.8
1.0	-2.0	**0.1**	1.8	-2.5	**-0.0**	1.6	-0.3	**1.5**	3.9	0.9	**4.4**	6.6
0.0	-1.4	**0.2**	1.9	-1.5	**0.1**	1.2	0.0	**1.7**	3.1	1.3	**3.8**	5.6

Appendix D

The size and granularity of the action space (i.e., control parameters) affect the results, since a larger action space cannot be explored as efficiently by a Monte Carlo tree search algorithm. In this paper, we use an action grid spacing of [0.5, 5.0] (i.e, consider flotation times with step size 0.5 minutes and air flow rates with step size 5.0 L/hr) for all results, since it best reflects the most fine-grained control settings that are still realistic. Additional testing in Table D1 shows that for larger action spaces (i.e., smaller step sizes), more model uncertainty is necessary for the POMDP approach to perform better than MPC.

Table D1. Reward of POMDP approach relative to MPC (in Δ$M/yr) when varying action space granularity at different levels of grade and recovery error variance.

Action space	Grade and recovery log variance			
Step size	-3.0	-2.0	-1.0	0.0
[0.1, 1.0]			-9	+94
[0.25, 2.5]	-14	-4	+21	+127
[0.5, 5.0]	-5	+9	+78	

References

Amini, S. H. Optimization of mineral processing circuit design under uncertainty. Morgantown, West Virginia University, 2017.

Amini, S. H., Noble, A. Design of cell-based flotation circuits under uncertainty: A techno-economic stochastic optimization. Minerals, 2021, 11: 459.

Arief, M., Alonso, Y., Oshiro, C., et al. Managing geological uncertainty in critical mineral supply chains: A POMDP approach with application to us lithium resources, arXiv 2025, arXiv:2502.05690. Available online: https://arxiv.org/abs/2502.05690 (accessed on December 2025).

Bai, Z., Gao, P., Chu, M., et al. Artificial intelligence of mineral processing process: A review of research progress. Journal of Environmental Chemical Engineering, 2025, 13(5): 118313.

Bascur, O. Process control and operational intelligence, in SME Mineral Processing and Extractive Metallurgy Handbook, Society for Mining, Metallurgy, and Exploration (SME), edited by R. C. Dunne and S. K. Kawatra, pp. 277–316, 2019.

Concha A, F., Bascur, O. A. The Engineering Science of Mineral Processing: A Fundamental and Practical Approach, CRC Press, 2024.

Ding, J., Chai, T., Wang, H., et al. Knowledge-based global operation of mineral processing under uncertainty. IEEE Transactions on Industrial Informatics, 2012, 8: 849–859.

Hodouin, D. Methods for automatic control, observation, and optimization in mineral processing plants. Journal of Process Control, 2011, 21: 211–225.

Hodouin, D., Jämsä-Jounela, S.-L., Carvalho, M., et al. State of the art and challenges in mineral processing control. Control Engineering Practice, 2001, 9: 995–1005.

International Energy Agency (IEA). (2021). The role of critical minerals in clean energy transitions, 2021. Available online: https://www.iea.org/reports/the-role-of-critical-minerals-in-clean-energy-transitions (accessed on December 2025).

Jiang, Y., Fan, J., Chai, T., et al. Lewis, Data-driven flotation industrial process operational optimal control based on reinforcement learning. IEEE Transactions on Industrial Informatics, 2017, 14: 1974–1989.

Jovanović, I., Miljanović, I. Contemporary advanced control techniques for flotation plants with mechanical flotation cells–A review. Minerals Engineering, 2015, 70: 228–249.

Koch, P.-H., Rosenkranz, J. Sequential decision-making in mining and processing based on geometallurgical inputs. Minerals Engineering, 2020, 149: 106262.

Kochenderfer, M. J., Wheeler, T. A., Wray, K. H. Algorithms for decision making. MIT press, 2022.

Koermer, S. C. Bayesian methods for mineral processing operations. Blacksburg, Virginia, Virginia Polytechnic Institute and State University, 2022.

Koermer, S., Noble, A. (2025). Optimization of a metallurgical process with uncertain dynamics. Available online: https://skoermer.github.io/media/Koermer%20Scott%20MPD%20contest%202021.pdf (accessed on December 2025).

Lee, H., Calvin, K., Dasgupta, D., et al. (2023). Synthesis report of the IPCC sixth assessment report (AR6). Available online: https://www.ipcc.ch/report/ar6/syr/downloads/report/IPCC_AR6_SYR_SPM.pdf (accessed on December 2025).

McCoy, J. T., Auret, L. Machine learning applications in minerals processing: A review. Minerals Engineering, 2019, 132: 95–109.

Shean, B., Cilliers, J. A review of froth flotation control. International Journal of Mineral Processing, 2011, 100: 57–71.

Silver, D., Veness, J. (2010) Monte-Carlo planning in large POMDPs. Available online: https://proceedings.neurips.cc/paper/2010/file/edfbe1afcf9246bb0d40eb4d8027d90f-Paper.pdf (accessed on December 2025).

Sunberg, Z., Kochenderfer, M. Online algorithms for POMDPs with continuous state, action, and observation spaces. Proceedings of the International Conference on Automated Planning and Scheduling, 2018, 28(1): 259–263.

U.S. Geological Survey. (2024). Phosphate rock. Available online: https://pubs.usgs.gov/periodicals/mcs2024/mcs2024-phosphate.pdf (accessed on January 2024).

Välikangas, H., Ohenoja, M., Brochot, S., et al. Evaluation of model uncertainty propagation in mineral process flowsheet designs. Scandinavian Simulation Society, 2025, 456–463.

World Bank. (2025). World Bank Commodities Price Data. Available online: https://thedocs.worldbank.org/en/doc/18675f1d1639c7a34d463f59263ba0a2-0050012025/related/CMO-Pink-Sheet-April-2025.pdf (accessed on 2 April 2025).

World Meteorological Organization (WMO). (2025). WMO confirms 2024 as warmest year on record at about 1.55 °C above pre-industrial level, 2025. Available online: https://wmo.int/news/media-centre/wmo-confirms-2024-warmest-year-record-about-155degc-above-pre-industrial-level (accessed on 10 January 2025).

Xiang, X., Foo, S. Recent advances in deep reinforcement learning applications for solving partially observable markov decision processes (POMDP) problems: Part 1—fundamentals and applications in games, robotics and natural language processing. Machine Learning and Knowledge Extraction, 2021, 3: 554–581.

Xiang, X., Foo, S., Zang, H. Recent advances in deep reinforcement learning applications for solving partially observable markov decision processes (POMDP) problems part 2—applications in transportation, industries, communications and networking and more topics. Machine Learning and Knowledge Extraction, 2021, 3: 863–878.