Systems & Control Transactions

Research Article - Peer Reviewed Conference Proceeding
ESCAPE 35 - European Symposium on Computer Aided Process Engineering
Ghent, Belgium. 6-9 July 2025
Jan F.M. Van Impe, Grégoire Léonard, Satyajeet S. Bhonsale,
Monika E. Polańska, Filip Logist (Eds.)

PSE PRESS

Probabilistic Model Predictive Control for Mineral Flotation using Gaussian Processes

Victor Dehon[a], Paulina Quintanilla[b], Antonio Del Rio Chanona[a]

[a] Department of Chemical Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom
[b] Department of Chemical Engineering, Brunel University of London, Uxbridge, UB8 3PH, United Kingdom
Corresponding Authors: paulina.quintanilla@brunel.ac.uk; a.del-rio-chanona@imperial.ac.uk

ABSTRACT

Recent advancements in machine learning and time series analysis have opened new avenues for improving predictive control in complex systems such as mineral flotation. Techniques leveraging multivariate predictive control in mineral flotation have seen significant progress in recent years. However, challenges in developing an accurate dynamic model that encapsulates both the pulp and froth phases have hindered further advancements. Now, with a readily available model containing equations that describe the physics of flotation froths, an opportunity for novel control strategies presents itself. In this study, a Gaussian Process (GP) Model Predictive Control (MPC) strategy is proposed to integrate uncertainty quantification directly into the control framework. By leveraging the probabilistic nature of GP models, this approach captures process variability and adapts dynamically to new data, ensuring continuous refinement of the GP model within the MPC strategy. Unlike previous implementations where model parameters remained static, this methodology updates the GP model in real time, allowing for improved decision-making in the face of process uncertainty. The GP model was trained, optimized with JAX, evaluated using relevant metrics, and implemented as a surrogate within the MPC framework. The results demonstrate the capability of the GP model to accurately represent process dynamics while minimising prediction errors. Moreover, incorporating uncertainty reduction through standard deviation minimisation in the objective function enhances both control performance and system robustness. This paper sets the basis for the potential of using GP-MPC to enhance both the accuracy and robustness of mineral froth flotation control.

Keywords: Mineral Flotation, Model Predictive Control, Gaussian Processes, Machine Learning

INTRODUCTION

Machine learning has rapidly become one of the most exciting and revolutionary tools in modern industrial processes [1]. Leveraging data-driven models, such as Gaussian Processes (GPs), offers sophisticated methods to tackle the inherent complexity and nonlinearity of chemical processes. Mineral flotation is recognised as a critical separation technique in the mining industry, playing a vital role in extracting value from low-grade and complex ore bodies [2]. This process utilises differences in surface properties to selectively separate valuable minerals from the gangue. The process can be fine-tuned by adjusting various parameters such as air flow rate, chemical dosages, and pulp level to optimise recovery and grade (purity) of the desired minerals. However, the introduction of non-linear dynamics, complex interactions, and instability arising from multi-phase operation introduce challenges to achieving effective control [3]. GPs are particularly valuable in these contexts because they provide a flexible, probabilistic approach to modelling complex systems, allowing for both accurate predictions and uncertainty quantification. This makes them well-suited for optimising and controlling industrial processes where uncertainty and variability are significant challenges. Optimal control and optimisation in mineral froth flotation processes have made significant progress in recent times [5,6]. In mineral froth flotation, key control variables, such as air flow rate and pulp level, play a crucial role in maintaining process stability and optimising performance. These parameters directly influence the dynamics of both the pulp and froth phases, affecting the efficiency of the separation process. The process is of particular interest to researchers due to marginal improvements in efficacy yielding substantial economic benefits as well as helping with global efforts of meeting increasing demand effectively.

Problem Description

A common control strategy in flotation is Model Predictive Control (MPC), widely recognised as one of the most efficient for multi-variable processes. Its development for mineral flotation was previously limited by the lack of a dynamic model accurately representing the complex froth flotation interactions [3]. Earlier MPC implementations relied on empirical models capturing only the pulp phase, neglecting the froth phase, which significantly impacts flotation performance [5]. However, due to the contributions made by Quintanilla et al. [5,6], a physics-based model has been proposed, developed, and calibrated. This opens the possibility of applying state-of-the-art probabilistic modelling to a control strategy for this mineral flotation process. Building on the success of using this process model to implement an economic-model predictive controller [4], this paper aims to design a novel GP-MPC control strategy for the optimal operation of a mineral froth flotation process. By integrating GPs into the MPC framework, this strategy aims to achieve higher recovery rates and better product quality, ultimately contributing to more sustainable and economically viable mineral processing practices.

BACKGROUND

Gaussian Processes

Gaussian Processes (GPs) have become a powerful tool in various fields, including process control [16]. A key advantage is their ability to capture residual uncertainty from plant-model mismatches, useful in control systems where understanding the confidence of predictions is crucial. Unlike traditional mathematical optimisation, which assumes exact function evaluations, GPs model observations as stochastic, providing a probabilistic system representation. Although GPs in control for a mineral flotation process are not well documented, the use of GP as a surrogate model are present in the control community. GPs alongside MPC first started appearing in the early 2000's with Murray Smith et al [7]. The inclusion of information about the trust in the model depending on the region, enabled a more nuanced exploration of the control space. Overall, this departure from the traditional deterministic view creates a novel opportunity for optimal control by leveraging uncertainty to refine dynamic models, improving robustness. Given the process's inherent nonlinearities [6], GP-based probabilistic modeling holds strong potential for mineral flotation.

Model Predictive Control for Mineral Flotation

Flotation is a key recovery process in the mineral industry. Studying this process is crucial due to its versatility and efficiency in processing large ore volumes, where even small improvements can yield significant economic benefits [3]. Attempts to introduce a robust and effective control strategy for the mineral flotation process have been made [8,9]. However, in previous work, modelling for control purposes was challenging. Most models have only focused on the pulp phase rather than the froth phase [5]. A comprehensive phenomenological model has been created by Quintanilla et al [5,6]. The model involves several innovative features, encompassing not only the representation of the pulp phase but also the inclusion of variables pertinent to the froth phase, such as froth stability, bursting rate, and air recovery. The equations in the model facilitate the estimation of critical variables associated with the key performance indicators: grade (product purity) and recovery. Such variables could then be seamlessly integrated into a dynamic optimisation objective function formulation for an MPC strategy. The availability of this dynamic model enables novel control strategies to be attempted. While success in stochastic data-driven MPC using GPs and froth flotation MPC is documented, applying GPs and MPC to mineral flotation will be, to the best of the authors' knowledge, the first of its kind.

METHODOLOGY

Gaussian Process Training and Implementation

As previously outlined, GP's offer a flexible and nonparametric approach to modelling processes. In the context of MPC, this GP will serve as a surrogate model that predicts the future states of the process given current and past inputs. The structure of the GP model is designed to handle the multi-input, multi-output nature of the froth flotation process with multiple training datasets provided by the synthetic flotation model. This training data consisted of 9 state variables and 2 control variables. Before training the GP model, all input data were standardised to have zero mean and unit variance. This preprocessing step prevents any one input feature from dominating the prediction due to scale differences. A negative-log-likelihood (NLL) objective function was deployed in the training process to optimise the hyperparameters of the GP model. The negative log likelihood is given by:

$$NLL(\boldsymbol{\theta}) = -\log P\left(\boldsymbol{Y} \mid \boldsymbol{\theta}\right) \tag{1}$$

NLL is a standard objective function employed in hyper-parameter optimisation for GP's due to its versatility and being directly related to the model's predictive performance. In this GP formulation, a radial basis function (RBF) kernel was selected with respective hyper-parameters acquired through the maximisation of the marginal likelihood for each parameter (or equivalently, the minimisation of the negative log likelihood). With an 11-dimensional input space, a multi-start strategy improved exploration. These samples were passed to a stochastic gradient-based optimisation algorithm, specifically the Adam algorithm, to generate optimal hyperparameters, as seen in Algorithm 1. The selection of this solver was due to its computational efficiency as well as its capabilities in handling noisy and large volumes of data. An area of concern for this iterative process was the length of time taken to arrive at a suitable solution due to the computational complexity, nonconvexity, and taxing matrix operations conducted in the GP formulation procedure. To speed up this hyper-parameter optimisation procedure, the JAX python library was deployed. Notable advantages such as vectorised mapping, just-in-time compilation through XLA [10] and automatic differentiation allows for efficient computation of gradients in regard to the GP log-likelihood objective function. Due to the repeated nature of complicated matrix operations such as the inversion of the optimal covariance matrix, JAX's GPU acceleration capabilities can dramatically speed up the central matrix calculations in the GP computations.

Algorithm 1. JAX GP – Creation & Inference

1. **Input:** Xdat, Ydat, Kernel, multi_hyper, var_out
2. **Initialise:** Normalise Xdat and Ydat
3. Set bounds for optmisation procedure
4. Generate initial points using Sobol sequences
5. **For** i in initial points & GP output:
   a. **Define:** $-\log p(X \mid \theta)$ *where* $\theta = \{\iota_{opt}, \sigma^2_{f,opt}, \sigma^2_{n,opt}\}$
   b. Compute $\nabla_\theta NLL(\theta \mid X)$
   c. Update parameters using SGD: $\theta_{t+1} = \theta_t - \eta \nabla_\theta Loss(\theta_t)$
   d. **If** Max iterations or convergence criteria met **Then:**
      i. Select optimal solution θ_i^* with smallest $-\log p(X_i{}^* \mid \theta^*{}_i)$ value **end if**
   e. **End for**
6. Computer optimal coveriance matrix: $cov_{ij} = \sigma_f^2(-\frac{1}{2}dist_{ij})$
7. Computer inverse of covariance matrix: K^{-1}
8. Computer covariance between the test and training points: k
9. Computer predicted mean: $\mu_* = k^T K^{-1} Y$
10. Compute predicted variance: $\sigma_*^2 = max(0, \sigma_f^2 - k^T K^{-1} Y)$
11. **Output:** $\theta_{opt} = \{\iota_{opt}, \sigma^2_{f,opt}, \sigma^2_{n,opt}\}$, μ_*,: σ_*^2

Gaussian Process Model Predictive Control

The objective function contains 3 objectives. The first is related to the concentration of valuable materials lost in the tailings (Equation 2) due to the primary goal of flotation being to optimise the separation of valuable minerals. Secondly, the control effort (Equation 3) explicitly incorporates the cost of actuator movements, which is important due to controls such as pulp height (h_p) being controlled via a valve. Finally, the standard deviation penalty is introduced to discourage the MPC from making decisions on uncertain estimates and rather to guide it towards more reliable and robust control strategies (Equation 4).

$$J_{loss} = \sum_{k=0}^{N-1} Loss_k \quad (2)$$

$$J_{control} = \sum_{k=1}^{N-1} (u_k - u_{k-1})^2 \quad (3)$$

$$J_{std} = \sum_{k=0}^{N-1} \sigma^2_{GP_k} \quad (4)$$

The three equations are then put altogether with weighting constant values of 50,1 and 10 respectively to form a multi-objective function to be minimised by an SLSQP method.

The GP-MPC approach leverages the predictive capabilities of GP models to enhance the performance of MPC by providing more accurate and probabilistic fore-casts of the system's behaviour. The combination of both should enable better adaptation to dynamic conditions and improve overall system robustness. The use of a sur-rogate model involves the creation of simplified models that approximate the behaviour of complex systems. This approach reduces the computational burden typically as-sociated with solving the optimisation problem at each MPC iteration, as it allows for faster evaluations of the system's future states without sacrificing accuracy. This efficiency is crucial in this mineral froth flotation applica-tion where computational resources and time are limited. In the context of this mineral froth flotation process MPC framework, the trained GP model will be employed at each time step to predict the future states of the system for the entire prediction horizon, allowing the MPC to evaluate the impact of different control actions over this horizon. These GP-generated predictions, which include both mean estimates and associated uncertainties, are fed into the MPC's objective function. The MPC then optimises the control inputs based on these predictions to achieve desired future outcomes while considering uncertainties. The GP can be updated online as new data becomes available. As the system evolves, the updated models are then used to predict the next set of future states, ensuring that the control strategy remains relevant and effective as the system dynamics change. The overall process of training the GP, simulating the surrogate within the MPC framework, calculating the optimal control and finally updating with the relevant process model states as depicted in Figure 1.

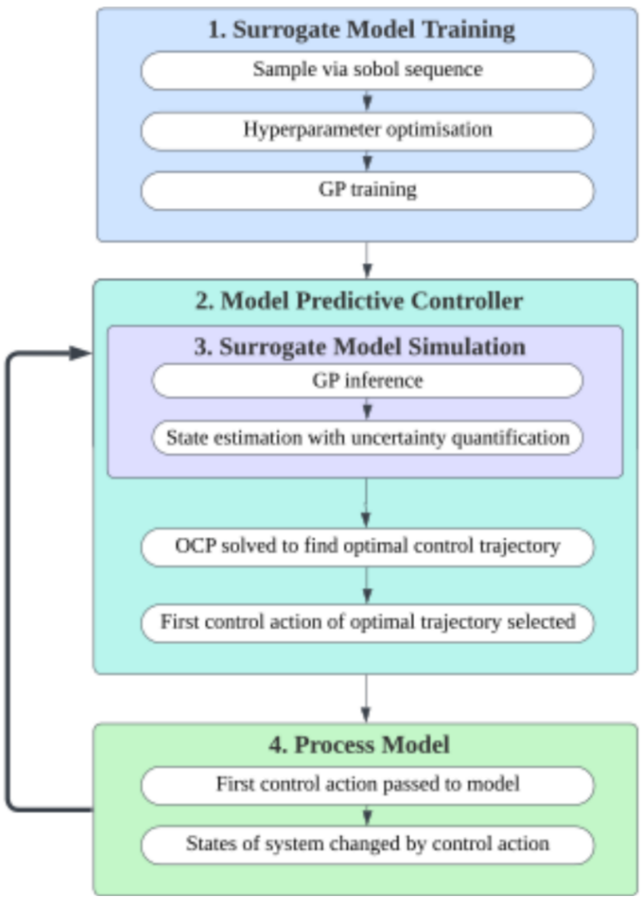

Figure 1 - GP-MPC Flow Diagram

RESULTS AND DISCUSSIONS

GP Training and Validation

Evaluating the effectiveness of the GP surrogate implementation was done with a training and validation dataset. It is important to note that this is a state-space GP model, where the states are: mineral (M_{min}), gangue masses (M_{gangue}), gas holdup classes ($\varepsilon_1 - \varepsilon_5$), pulp level ($h_p$) and tailings flow rate (Q_{tails}). Both datasets were generated using the physics-based process model from [5,6]. Various statistical measures were used to evaluate the performance of the GP surrogate model, including negative log predictive density (NLPD), mean absolute percentage (MAPE), and root mean squared percentage error (RMSPE). Before the calculations of these statistical parameters, the states were normalised to avoid problems due to differences in the order of magnitude of the different states. Figure 2 displays the GP predictions closely following the underlying system data. It was observed that the testing results show slightly higher error rates than the training data, which is expected. The ϵ's are slightly less accurate and confident, with the GP predictions deviating slightly from the real data and a wider confidence region. Confidence regions remain consistent from training to testing, with estimations related to Qtails, h_p, M_{gangue} and M_{min} showing the highest confidence in estimations due to a tighter standard deviation. Overall, from graphical interpretation, the model seems to fit the training data well, with predictions closely following the experimental data points. Tables 1 and 2 reflect a similar conclusion drawn from the graphs, with a low RMSPE and MAPE.

Table 1: GP Without JAX Testing Metric Results

Experiment No.	RMSPE	NLPD	MAPE
Experiment 1	0.15%	-5.47	0.09%
Experiment 2	0.24%	-5.29	0.14%
Experiment 3	0.36%	-4.91	0.20%

A low NLPD is also seen, portraying good probabilistic predictions and uncertainty quantification. The metrics show consistent performance across all three testing experiments, with little variation, suggesting robustness in the GP model. Results from the JAX-enhanced GP formation are seen in Figure 2 which also illustrates similar prediction patterns and confidence intervals across the training and testing dataset. Quantitatively, the performance metrics for the JAX-GP in Table 2 demonstrates comparability with the non-JAX GP. Some differences can be observed from the JAX-GP and non-JAX GP, such as estimation of the states M_{min} and Q_{tails} in the testing data, which are slightly less accurate in the JAX-GP, with a smoother mean being estimated. The estimations of states M_{min}, h_p and ϵ's are confident with an observable tight standard deviation. The most notable improvement was observed in compilation time, with a reduction of 33% when run on an Apple M2 chip running on macOS 14.4.1. This speedup in compilation, coupled with the preservation of model accuracy, underscores the potential of JAX to enhance the efficiency of GP implementations.

Table 2: GP with JAX testing metric results

Experiment No.	RMSPE	NLPD	MAPE
Experiment 4	0.40%	-4.78	0.26%
Experiment 5	0.40%	-4.79	0.28%
Experiment 6	0.63%	-4.17	0.44%

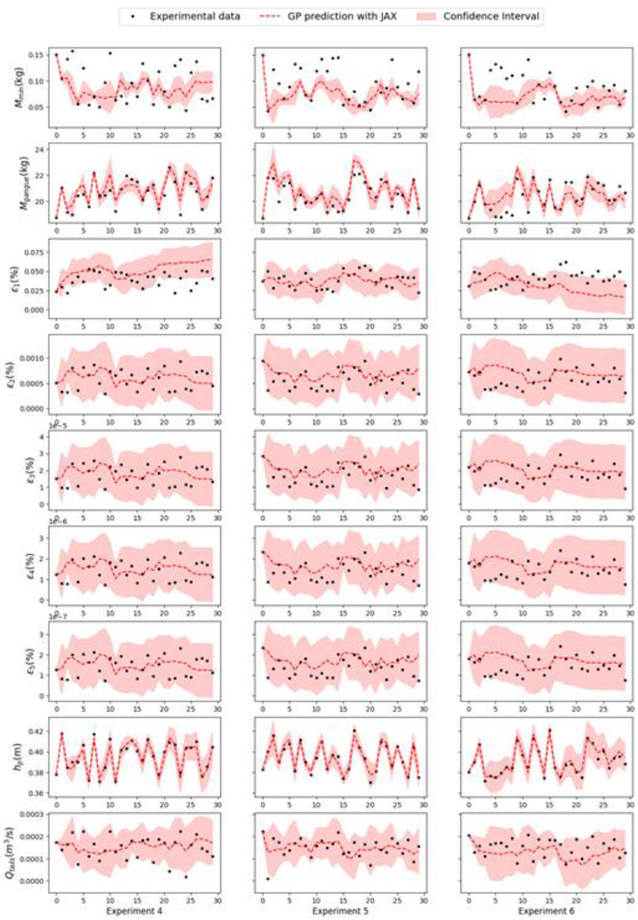

Figure 2 – GP with JAX results illustrating the evolution of system states and predictions. The black dots represent the experimental data from the process, the red lines indicate the GP predictions, and the red shadows depict two standard deviations from the GP model. These states are assumed to be observable within the MPC framework.

GP-MPC Control Performance

The effectiveness of the overall control strategy was assessed using MAE, RMSE, and ISE metrics, comparing the desired setpoint with the actual control actions. Table 3 depicts these metrics, showing that the control mechanism appears to be effective in maintaining the desired levels for both h_p and Q_{air}. Additionally, the evaluation considered the average concentration of valuable material lost in the tails. The dataset fed was unseen by the GP and MPC. Figure 3 demonstrates the GP estimates for the h_p (pulp height) state evolution with associated confidence intervals as well as the resulting optimal control action for h_p control. The GP predictions follow closely the process, with tight uncertainty region suggesting a well-calibrated surrogate model for the process. Moreover, Figure 4 visualises the air flow rate (Q_{air}) control strategy derived by the MPC strategy with an overall trend of increasing air flow over time, which typically promotes increased bubble production and mineral recovery in froth flotation.

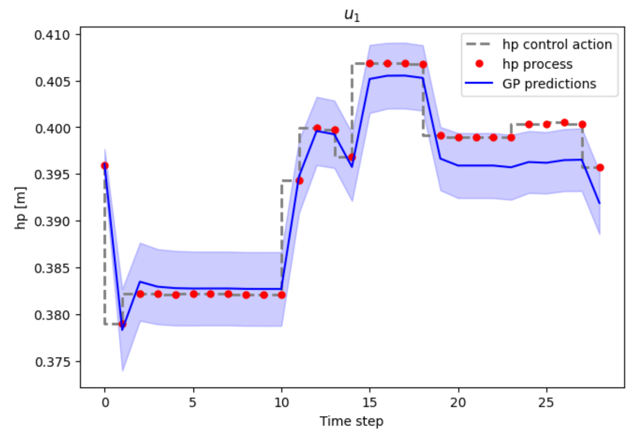

Figure 3 – Pulp height (h_p) control output from GP-MPC

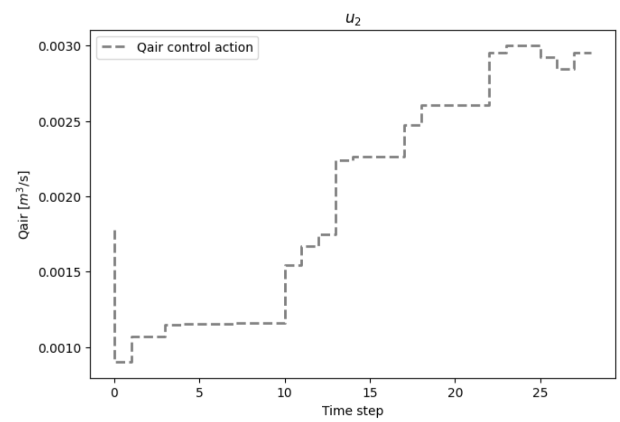

Figure 4 - Q_{air} control output from GP-MPC

Table 3: Non-JAX GP with MPC metric values

Metric	h_p [m]	Q_{air} [m^3/s]
MAE	0.012431	0.000758
RMSE	0.015232	0.000957
ISE	0.006728	0.000027
Mean Loss in Tails (kg/m^3)	0.630612	

Overall, the control strategy balances stability and responsiveness, initially maintaining constant action (time steps 2-10) and introducing significant adjustments around steps 12 and 20 to manage the trade-off between metallurgical recovery and grade. Lower pulp heights, which yield higher grades but reduce recovery, are dynamically adjusted alongside air flowrate Q_{air} to optimise flotation performance. This aligns with findings on peak air recovery and froth stability from prior studies [5], highlighting the importance of Q_{air} in achieving optimal recovery. Furthermore, Figure 4 demonstrates a stepped approach rather than continuous increase, suggesting the controller is allowing the flotation system to stabilise at each air flow level before making further adjustments. The strategy demonstrated efficiency, with low mean material loss in tailings (0.630612 kg/m³). Control metrics, including low MAE, RMSE, and ISE values, confirm precise and robust system performance, validating the effectiveness of the GP-MPC approach.

CONCLUSION

MPC strategies have proven effective in mineral froth flotation, improving recovery and grade. However, many models focus on the pulp phase and struggle with the froth phase. With a more accurate process model for both phases now available, an opportunity arises for GPs to be applied, leveraging confident decisions. MPC with GPs for mineral froth flotation has not been widely explored. The findings from the implemented models demonstrated that the surrogate model generates accurate and certain inferences on unseen datasets with low amounts of associated error. Moreover, the possibility of accelerating the surrogate model creation was outlined with the implementation of JAX. The MPC strategy utilised the GP surrogate to generate an optimal control trajectory. Most notably, the mean loss in tails was measured at a low 0.630612 kg/m³. The performance is further evidenced by examining the control accuracy, with both h_p and Q_{air} parameters exhibiting low MAE and RMSE values, indicating precise control overall. These findings reinforce the benefits of leveraging the uncertainty quantification and probabilistic attributes of the GP model alongside an MPC control strategy. Being the first of its kind to apply GP-MPC to a mineral froth flotation process, this work drives forward the objective of multivariate predictive control of flotation processes, providing an alternative probabilistic approach to enhancing mineral recovery and grade. In addition, the improvements noted in this study related to reduced compilation outline the possibility for an accurate and fast control strategy.

REFERENCES

1. McCoy JT, Auret L. Machine learning applications in minerals processing: A review. Minerals Engineering. 2019;132:95–109.
2. Wang B, Peng Y. The effect of saline water on mineral flotation--A critical review. Minerals Engineering. 2014;66:13–24.
3. Quintanilla P, Neethling SJ, Brito-Parada PR. Modelling for froth flotation control: A review. Minerals Engineering. 2021;162:106718.
4. Quintanilla P, Navia D, Neethling S, Brito-Parada P. Experimental implementation of an economic model predictive control for froth flotation. In: Computer Aided Chemical Engineering. Elsevier; 2024. p. 1759–64.
5. Quintanilla P, Neethling SJ, Navia D, Brito-Parada PR. A dynamic flotation model for predictive control incorporating froth physics. Part I: Model development. Miner Eng 173:107192 (2021).
6. Quintanilla P, Neethling SJ, Mesa D, Navia D, Brito-Parada PR. A dynamic flotation model for predictive control incorporating froth physics. Part II: Model calibration and validation. Minerals Engineering. 2021;173:107190.
7. Girard A, Rasmussen CE, Quinonero Candela J, Murray-Smith R. Gaussian process priors with uncertain inputs & application to multiple-step ahead time series forecasting. In: Advances in Neural Information Processing Systems (NIPS) 15. MIT Press (2003).
8. Wan K, Xu D, Cai Y. Data-driven-based model predictive control of reagents addition for tungsten flotation process. In: Proceedings of the 2021 China Automation Congress (CAC), IEEE, pp. 1629-1634 (2021).
9. Cubillos FA, Lima EL. Identification and optimizing control of a rougher flotation circuit using an adaptable hybrid-neural model. Miner Eng (1997).
10. Google. JAX quickstart guide. Accessed: 2024-08-24. Available at: https://jax.readthedocs.io/en/latest/quickstart.html (2023).