Stochastic data-driven NMPC for partially observable systems using Gaussian processes: a mineral flotation case study

Yicong Wang * **Antonio del Río Chanona** *
Paulina Quintanilla **

* *Department of Chemical Engineering, Imperial College London, United Kingdom(e-mail: a.del-rio-chanona@imperial.ac.uk)*
** *Department of Chemical Engineering, Brunel University London, United Kingdom (e-mail: paulina.quintanilla@brunel.ac.uk)*

Abstract: This paper presents a nonlinear model predictive control (NMPC) strategy using Gaussian Processes (GPs) to control a froth flotation process under partial observability. The GP state-space model predicts future states for both observable and latent variables, using available data, while incorporating the probability distribution of these predictions into an optimization problem. This improves robustness against measurement noise and process disturbances and evaluates the impact of feed particle size, a typical process disturbance. We assessed the framework's ability to maintain optimal process performance across varying operating conditions. The results demonstrate that the proposed GP-MPC framework improves process efficiency, even with frequent changes in particle size and measurement noise, confirming its potential for online control of partially observable systems.

Keywords:
Model Predictive Control; Gaussian Processes; Froth Flotation;Process Control

1. INTRODUCTION

Froth flotation is a critical process in mineral processing to separate valuable minerals from gangue. The process is highly sensitive to disturbances, such as variations in feed particle size, which are decisions originating from upstream grinding operations (Gaudin et al., 1942; Quintanilla et al., 2025). Controlling these upstream conditions is challenging, significantly affecting the floatability of valuable minerals and gangue (Hu, 2014; Quintanilla et al., 2023a).

Although advanced control strategies, including Model Predictive Control (MPC), have been extensively studied across various processes in recent years, their application in mineral processing remains limited. Surveys suggest that it could greatly benefit the mineral processing industry (Olivier and Craig, 2017). However, implementing MPC in froth flotation has been constrained by the complexity of system dynamics required for accurate modeling (Quintanilla et al., 2021b). Traditional control methods, such as proportional-integral-derivative (PID) and rule-based approaches, have struggled to manage the process's inherent complexities, including nonlinear dynamics, disturbances, and measurement noise, while hybrid systems relying on static models have failed to adequately capture the uncertainties of the process (Maldonado et al., 2007; Putz and Cipriano, 2015). Additionally, high computational costs have hindered real-time applications.

Given the system's nonlinear nature and complexity, a data-driven approach offers notable advantages (Bradford et al., 2020). Gaussian Process (GP) models provide probabilistic state predictions, allowing for more informed control decisions. This enables the GP-MPC framework to dynamically balance exploration (reducing uncertainty) and exploitation (optimizing control actions) based on real-time measurements. As a result, this approach is particularly effective in maintaining robust control performance in the presence of disturbances and measurement noise.

Most GP-based MPC applications assume full observability, where all relevant state variables can be directly measured (Bradford et al., 2019; McAllister and Rasmussen, 2017; Park et al., 2022). However, this assumption does not hold in real-world froth flotation processes, where key variables, such as the mass of each mineralogical class, are challenging to measure online (González and Quintanilla, 2024). Given the importance of these variables in determining the economic objective of the control strategy, it is crucial to develop a partially observable model. The model should be capable of predicting latent variables by using prior predictions, measurements of observable variables, and the proposed control actions to effectively tackle this challenge.

This study introduces a Nonlinear MPC (NMPC) framework that employs Gaussian Processes (GPs) to model and control the froth flotation process under partial observability. The control actions are optimized to enhance the economic performance of the process while also accounting for system uncertainties and disturbances. Our approach overcomes the limitations of previous studies by integrating stochastic process modelling through GPs, which explicitly captures prediction uncertainty. Previous studies have identified the optimal particle size range for flotation as 20 to 150 µm (Gaudin et al., 1942). Here, we present a case study evaluating the performance of a flotation bank controlled by the proposed GP-MPC framework under 5% measurement noise and particle size disturbances ranging from 80 to 140 μm. Our results demonstrate that even under frequent particle size changes, the GP model accurately predicts partially observed (latent) state variables while effectively minimizing the economic objective. This confirms the robustness of the proposed control strategy.

2. METHODOLOGY

The control strategy employs a GP-based state-space model to predict the evolution of the flotation process. This model accounts for both observable and latent (unobservable) state variables, with the latter being predicted based on historical data using a partially observable GP model.

2.1 Partially observable Gaussian process state-space model

Discrete-time MPC is advantageous for real-world applications, as it enables the controller to update information at specific intervals rather than continuously. This discrete approach is ideal for digital implementations where control strategies are executed at fixed time steps, ensuring compatibility with digital processors and sampled data systems while also lowering computational demands. In this context, we consider the nonlinear discrete-time system described by:

$$\mathbf{x}_{k+1} = \mathbf{f}_d(\mathbf{x}_k, \mathbf{u}_k) + \boldsymbol{\epsilon}_k \tag{1}$$

Here, $\mathbf{x}_k \in \mathbb{R}^{n_x}$ denotes the state vector, $\mathbf{u}_k \in \mathbb{R}^{n_u}$ is the control input, $\boldsymbol{\epsilon}_k$ represents process noise, and $\mathbf{f}_d : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$. The GP model provides probabilistic predictions for the mean and variance of the state at the next time step, incorporating these probability estimates into control decisions.

To model dynamic systems, we employ a GP as a state-space representation of the system. For ease of notation, we group the variables as $\boldsymbol{\xi}_k = [\mathbf{x}_k^T, \mathbf{u}_k^T]^T$ such that $\boldsymbol{\xi} \in \mathcal{X} \subseteq \mathbb{R}^{n_x+n_u}$. Using a dataset of n_d samples, we construct the GP state-space model, which can be interpreted as a distribution over functions as $f_\psi(\boldsymbol{\xi}) \sim \mathcal{GP}(m_\psi(\boldsymbol{\xi}), k_\psi(\boldsymbol{\xi}, \boldsymbol{\xi}'))$

The model is defined by a mean function $m_\psi(\boldsymbol{\xi}) : \mathcal{X} \to \mathbb{R}$ and a covariance function $k_\psi(\boldsymbol{\xi}, \boldsymbol{\xi}') : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ each yielding a scalar output. Both functions map control and state inputs $\boldsymbol{\xi}$ to a single output value. To generate n_x outputs (e.g., one for each future state), we require n_x separate GPs, with one GP assigned per state. These n_x GPs are then combined to model the system dynamics:

$$\mathbf{f}_\Psi(\boldsymbol{\xi}) = \begin{cases} f_{\psi_1}(\boldsymbol{\xi}) \sim \mathcal{GP}(0, k_{\psi_1}(\boldsymbol{\xi}, \boldsymbol{\xi}')) \\ \vdots \\ f_{\psi_{n_x}}(\boldsymbol{\xi}) \sim \mathcal{GP}(0, k_{\psi_{n_x}}(\boldsymbol{\xi}, \boldsymbol{\xi}')) \end{cases} \tag{2}$$

We can then combine all parameter vectors into $\Psi = [\psi_1, ..., \psi_{n_x}]$ and all covariance functions into $\mathbf{k}_\Psi(\cdot,\cdot) = [k_{\psi_1}(\cdot,\cdot), ..., k_{\psi_{n_x}}(\cdot,\cdot)]$ to describe the concatenation of all kernel functions into $\mathbf{f}_\Psi(\boldsymbol{\xi}) \sim \mathcal{GP}(0, \mathbf{k}_\Psi(\boldsymbol{\xi}, \boldsymbol{\xi}'))$. This GP-based representation captures the system dynamics within a state-space model. However, in practical applications, not all state variables are measurable online. We therefore categorize them into observed variables $\mathbf{x}_{obs}$ and latent variables $\mathbf{x}_{latent}$, where $\mathbf{x}^{(i)} = \{\mathbf{x}_{latent}^{(i)}, \mathbf{x}_{obs}^{(i)}\}_{i=l}^{n_d}$, with l defined as the window size. To address this, the GP model was trained offline using data in which all state variables (both observed and latent) were assumed available. This assumption holds in practice because latent variables, such as mineral concentrations, can be measured offline through laboratory assays. These offline measurements were used to cross-validate the physics-based flotation model, ensuring its accuracy in representing the system dynamics. The training dataset was generated from the validated physics-based model. During real-time operation, the GP-MPC compensates for partial observability by inferring latent variables from observed state measurements and past observations, thereby ensuring the framework's practical applicability under realistic operating conditions.

2.2 GP-based model predictive control (GP-MPC)

For online MPC, the GP model takes the historical observed values from the previous l steps as inputs ($\mathbf{x}_{k,history}$) and generates predictions for the next time step ($\mathbf{x}_{k+1}$), which includes both observed and estimated latent states:

$$\mathbf{x}_{k+1} = \hat{\boldsymbol{f}}_d(\mathbf{x}_k, \mathbf{x}_{k,history}, \mathbf{u}_k) \tag{3}$$

where $\mathbf{x}_{k,\text{history}} = [\mathbf{x}_{obs,k}, \dots, \mathbf{x}_{obs,k-l}]$ and l defines the time-window size. This model enables more accurate state estimation and control under partial observability. Here, the GP model predicts future states and outputs of the process, incorporating both mean and variance predictions directly into the optimization. At each discrete time step k, an optimization problem is solved to determine the optimal sequence of control actions that minimize a cost function, which includes terms for economic performance, control effort, and predicted variance, thereby balancing performance with robustness to uncertainty. The objective is to minimize:

$$J = \sum_{k=0}^{N-1} \big(w_1 \mathbf{z}_k^T \mathbf{Q}_1 \mathbf{z}_k + w_2 (\Delta \mathbf{u}_k)^T \mathbf{Q}_2 (\Delta \mathbf{u}_k) + w_3 \mathbf{u}_k^T \mathbf{R} \mathbf{u}_k + w_4 \boldsymbol{\sigma}_k^T \mathbf{Q}_3 \boldsymbol{\sigma}_k \big) \tag{4}$$

where $\mathbf{z}_k$ represents the performance or economic term to be minimized, $\Delta \mathbf{u}_k$ denotes the change in control input, $\mathbf{u}_k^T \mathbf{R} \mathbf{u}_k$ captures the control input magnitude, and $\boldsymbol{\sigma}_k^T \mathbf{Q}_3 \boldsymbol{\sigma}_k$ represents the predicted variance. Each term in the objective function is penalized with a corresponding weight w_i. The GP-MPC optimization problem is thus formulated as:

$$\begin{aligned} \min_{\mathbf{u}_{0:N}} \quad & J(\mathbf{z}_{1:N}, \mathbf{u}_{0:N-1}, \boldsymbol{\sigma}_{1:N}) \\ \text{subject to:} \quad & \\ & \boldsymbol{\xi}_k = [\boldsymbol{\mu}_k^T, \mathbf{u}_k^T]^T \\ & \mathbf{z}_{k+1} = \mathcal{J}_z(\boldsymbol{\mu}_{k+1}) \\ & \boldsymbol{\mu}_{k+1}, \boldsymbol{\sigma}_{k+1} \leftarrow \mathbf{f}_\Psi(\boldsymbol{\mu}_{k,\text{obs}}, \mathbf{u}_k) \\ & \boldsymbol{\mu}_{0,\text{obs}} = \mathbf{x}_{\text{obs}} \\ & \mathbf{u}_{\min} \leq \mathbf{u}_k \leq \mathbf{u}_{\max} \\ & \text{for } k = 0, \dots, N-1 \end{aligned} \tag{5}$$

Here, $\boldsymbol{\xi}_k$ represents the concatenated vector of the predicted state mean $\boldsymbol{\mu}_k$ and the control input $\mathbf{u}_k$ at each time step k, with N denoting the control horizon. The term $\mathbf{z}_k = \mathcal{J}_z(\boldsymbol{\mu}_k)$ represents a mapping of the predicted

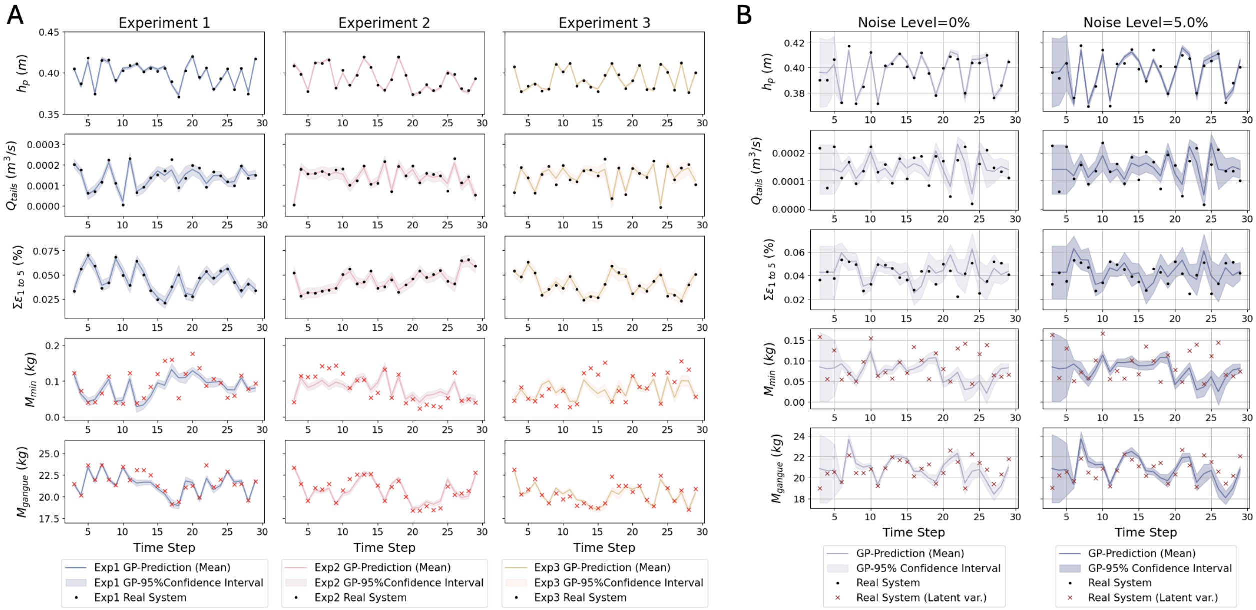

Fig. 1. (A) Training Performance of the Partially Observable GP Model; (B) Cross-validation performance of the GP model under different noise levels

state mean $\boldsymbol{\mu}_k$ to an economic or performance-related variable, used in the objective function to assess the system's performance over the control horizon. Additionally, $\boldsymbol{\mu}_{k+1}$ and $\boldsymbol{\sigma}_{k+1}$ denote the predicted mean and variance of the output, respectively, and $\mathcal{GP}(m_\psi(\boldsymbol{\xi}_k), k_\psi(\boldsymbol{\xi}_k, \boldsymbol{\xi}'_k))$ describes the Gaussian process distribution. $\mathbf{u}_{\min}$ and $\mathbf{u}_{\max}$ define the bounds for the control inputs.

3. CASE STUDY: MINERAL FROTH FLOTATION

3.1 Problem set-up

The control strategy is implemented in a mineral froth flotation process, with the primary objective of maximizing metallurgical recovery while maintaining concentrate grade. We used a physics-based dynamic model, developed by Quintanilla et al. (2021b) as a differential and algebraic equation system to act as the real system. The model was validated at a laboratory scale (Quintanilla et al., 2021a, 2023b). The flotation model includes nine states: mineral masses (M_{min}, M_{gangue}), gas holdup for each bubble size class (Mesa et al., 2022) ($\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \varepsilon_5$), pulp height ($h_p$), and tailings flow rate (Q_{tails}). There are only three observed variables: h_p, Q_{tails}, and the total gas hold up ($\sum_{b=1}^{5} \varepsilon_b$), while the rest are considered latent (unobservable) variables. These observed and latent variables align with what is commonly found on the industrial scale with the existing instrumentation. We used the typical industrial-scale manipulated variables, which are pulp height and air flow rate (Q_{air}) setpoints.

3.2 Implementation details

The GP model was trained on six sets of simulation data generated in MATLAB R2022b, using a time horizon of 30 steps and a sampling interval of 5 minutes, as this interval is sufficient to reach a new steady state under the operating conditions and tank dimensions used in this study (Quintanilla et al., 2021a, 2023c). The control action bounds in the simulation were set to $h_p \subseteq [0.37, 0.42]$ and $Q_{air} \subseteq [9 \cdot 10^{-4}, 3 \cdot 10^{-3}]$. The GP model was then implemented in Python 3.8.19 with a Radial Basis Function (RBF) kernel and optimized with 10 hyperparameters, using a multistart approach. The entire GP-MPC framework was developed in Python, integrating MATLAB for process simulation and Python for model training and control execution.

For MPC implementation, the control optimisation was performed using Sequential Least Squares Programming (SLSQP). The economic term $C_{1,tails}$ is calculated as defined in Eq.(6) at the final step of the control horizon (5 time steps). The MPC weight coefficients w_1, w_2, w_3, w_4 in the cost function were carefully tuned to balance economic performance, control effort, and variance minimization. Improper weighting can result in excessive fluctuations in control inputs or slow response times. Based on tuning, the selected weight coefficients were set to $[w_1, w_2, w_3, w_4] = [500, 0.05, 0.01, 0.2]$.

$$C_{1,tails} = \frac{M_{min}}{h_p(1 - \sum_{b=1}^{5} \varepsilon_b) A_{cell}} \tag{6}$$

where M_{min} represents the mineral masses, h_p is the pulp height, ε_{1-5} denote bubble size classes, and A_{cell} is the cross-sectional area of the cell.

4. RESULTS AND DISCUSSION

4.1 GP-model construction

To assess the model's fit, particularly the alignment of predicted distributions with the true data distribution, we used the Negative Log Predictive Density (NLPD) metric, a standard for evaluating GP models (GPyTorch Team, 2021). Figure 1(A) shows the training performance of the partically obervable GP model. The NLPD values for all observed variable fittings are between -9.78 (for Q_{tails}) and -2.49 (for h_p), indicating excellent model

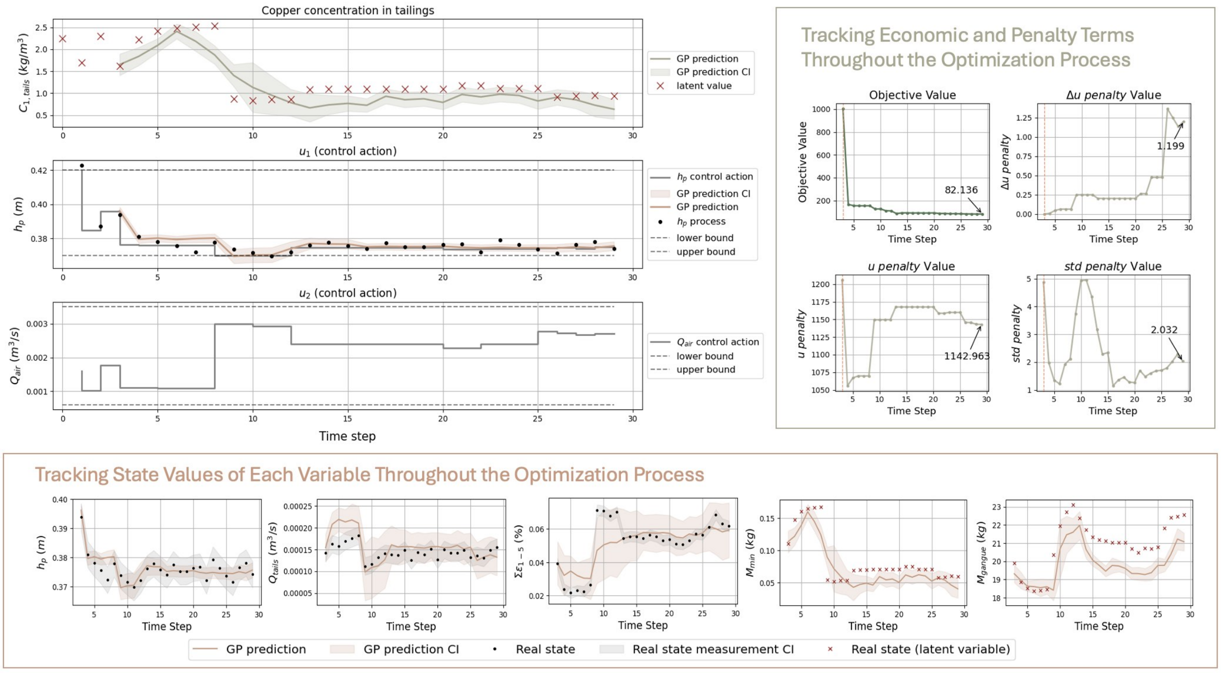

Fig. 2. GP-MPC Optimization with 5% measurement noise without external disturbances

performance, as nearly all state values fall within the model's confidence intervals (CI), which are relatively narrow. For latent variables, the NLPD values remain close to zero, with an average NLPD of 4.14 for M_{min} and 5.38 for M_{gangue}, reflecting a well-calibrated model that achieves high accuracy with appropriately estimated uncertainty.

For cross-validation of the fitted GP model, Figure 1(B) shows that the actual state values of observed variables consistently lie within the model's CI, with a moderate CI size, indicating no signs of overfitting. Regarding latent variables, the model aligns well with the real system trends. Although a slight deviation is observed in M_{min} between steps 22 and 26, the error does not accumulate over subsequent predictions; from step 27 onward, the predictions realign closely with the actual system values. This ability to recover alignment is particularly valuable when operating under various disturbances.

Furthermore, increasing the level of measurement noise does not significantly affect model performance, likely due to the partially observable nature of the model, which inherently enhances robustness against external measurement noise. The results confirm that the GP-MPC framework can reliably operate with up to 5% measurement noise. In this work, process noise $\boldsymbol{\epsilon}_k$ is assumed to be zero-mean Gaussian, i.e., $\boldsymbol{\epsilon}_k \sim \mathcal{N}(0, \sigma^2 I)$. This assumption is widely used in modeling chemical processes, as Gaussian noise effectively represents measurement uncertainties and unmodeled process variations.

4.2 GP-MPC performance

The proposed GP-MPC framework is designed to achieve reliable tracking by optimizing control actions based on state predictions while explicitly accounting for uncertainty.

The GP-MPC was initially tested without disturbances to assess its baseline performance. Since M_{min}, a latent variable, is required to calculate $C_{1,tails}$ (the economic term), the control feedback on $C_{1,tails}$ also relies on model predictions, introducing additional challenges. Nevertheless, as shown in Figure 2, the GP model successfully captures the trends of all state variables.

For observed variables, the CIs of the real state measurements, affected by measurement noise, are either equal to or narrower than the CIs of the model predictions, indicating stable performance under noise. For latent variables, the model performs particularly well in predicting M_{min}, as all state values fall within or very close to the prediction CI. Although there is a minor deviation in the prediction of M_{gangue}, this deviation remains minimal at approximately $+5\%$ of the state value, while the overall trend remains aligned. These results confirm the reliability of the GP model, establishing a strong foundation for MPC implementation.

The controller demonstrates substantial effectiveness, beginning to minimize the economic term within the first five time steps once the GP-MPC completes its initial data collection window and initiates control actions. During the first three time steps, a fixed control action is applied while collecting measurements of observable variables. Active control begins at the fourth step, using real-time measurements of observable variables and past predictions of latent variables as model inputs. The economic term is successfully minimized, reaching a steady-state value of approximately 0.9 kg/m^3, closely aligning with the predicted mean of around 0.7 kg/m^3. In addition, to evaluate the practical feasibility of our GP-MPC framework, we measured the optimization time at each control step. The optimization consistently took less than 30 seconds per 5-minute interval, confirming suitability for real-time implementation.

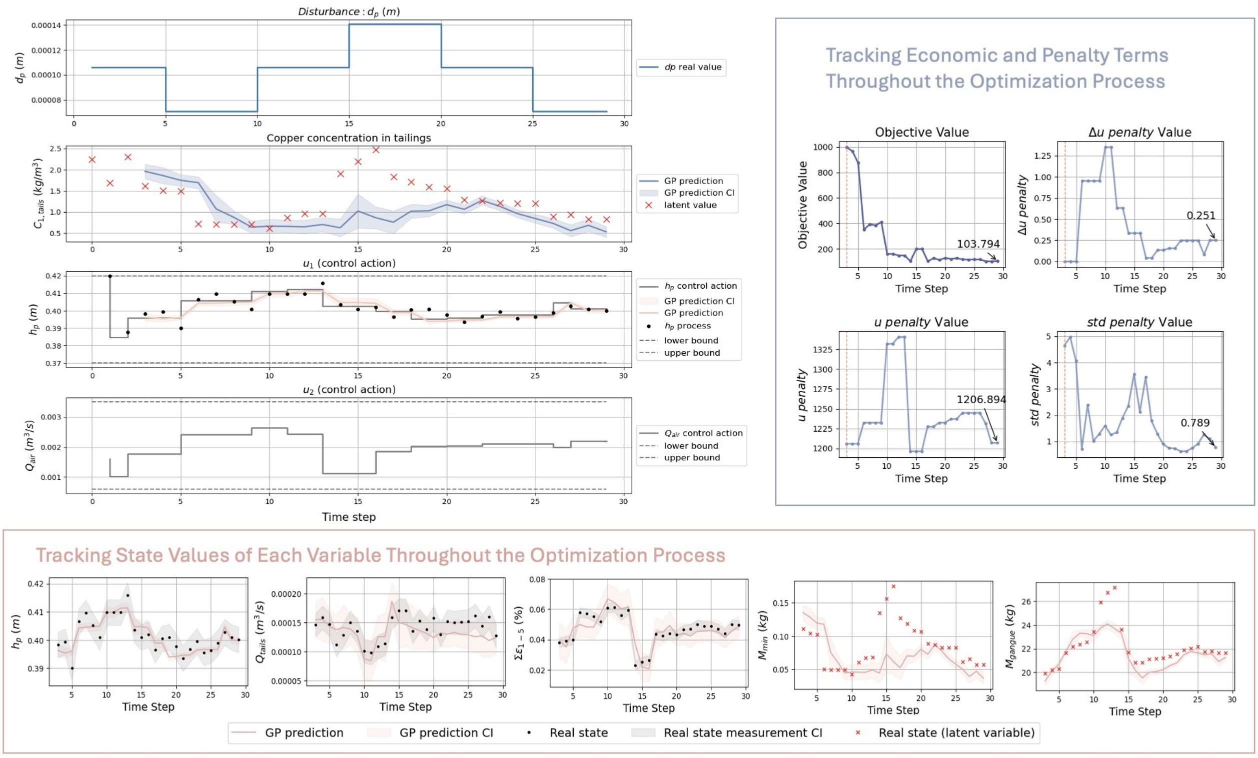

Fig. 3. GP-MPC optimization with 5% measurement noise under particle size disturbances (Case 1)

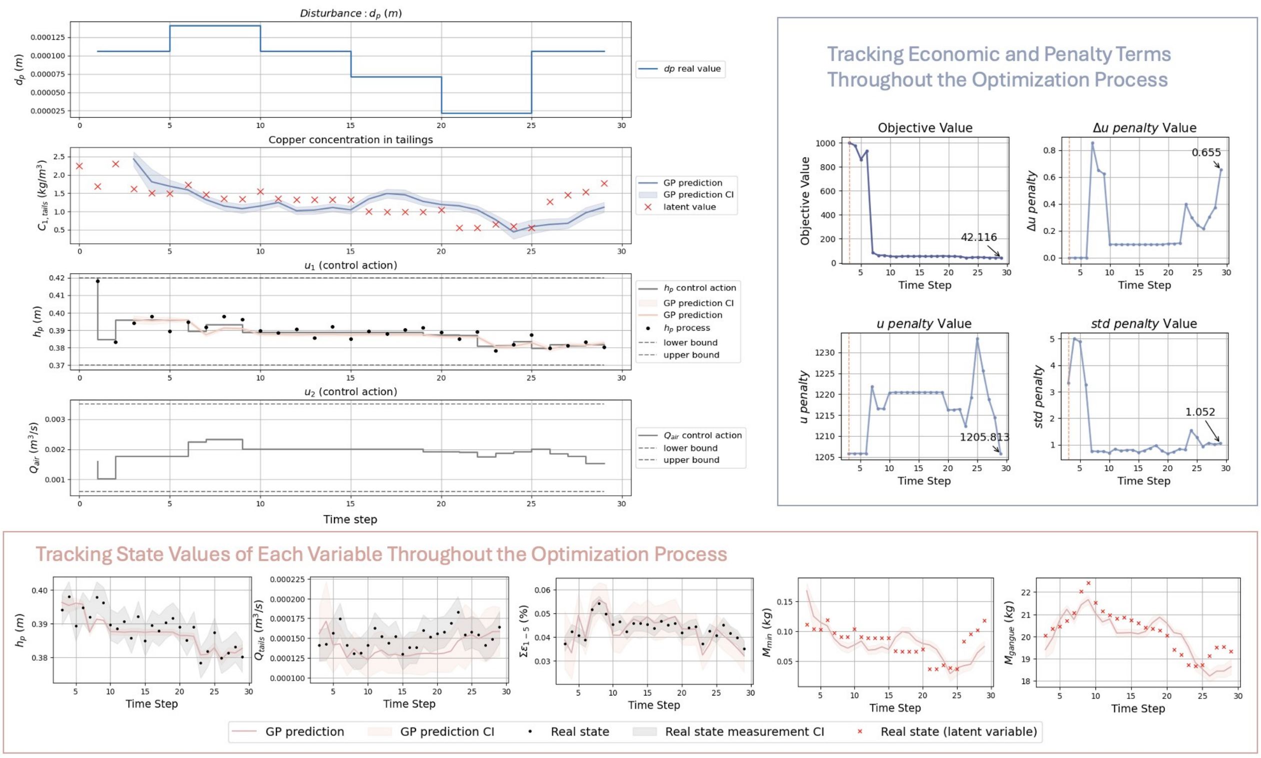

Fig. 4. GP-MPC optimization with 5% measurement noise under particle size disturbances (Case 2)

4.3 GP-MPC: Disturbance handling

The GP-MPC was then tested under both 5% measurement noise and two randomly generated sequences of particle size disturbances, each differing from the conditions in the training dataset. Changes in particle size affect floatability (Hu, 2014), potentially impacting system dynamics and reducing model prediction accuracy. In Case 1, M_{min} and $C_{1,tails}$ show minor deviations in alignment between steps 14 and 20 (see Figure 3); similarly, in Case 2, these variables exhibit slight misalignments between steps 26 and 29 (see Figure 4). Despite these minor deviations, the model consistently aligns with the trends of most observed and latent variables. Additionally, the GP-MPC responds quickly, effectively minimizing the economic term in both cases.

For control variables, neither control variable exceeded or leverage its boundary limits. The control action for h_p aligned well with its set trajectory without frequent, drastic adjustments, demonstrating stable and effective control performance. When comparing economic and penalty terms between disturbed and non-disturbed cases, the objective function is minimized effectively across all cases, with no significant differences attributable to disturbances. Interestingly, the variance penalty term is further minimized in the disturbed case compared to the non-disturbed case. This could indicate a potential risk where the model overfits by narrowing the confidence interval of state predictions, potentially compromising reliability.

5. CONCLUSIONS

This study demonstrates the potential of a Gaussian Process-based Model Predictive Control (GP-MPC) framework for managing a froth flotation process under disturbances and partial observability. The GP model effectively handles both observed and latent variables by incorporating prediction uncertainty into the control strategy, enabling robust decision-making even when key state variables cannot be directly measured. The framework was tested across various scenarios, including feed particle size disturbances and measurement noise. It is demonstrated to maintain stable control while minimizing economic objectives and adapting to dynamic conditions. These results suggest that GP-MPC is a promising, adaptable approach for complex, nonlinear processes in mineral processing, where reliable performance under uncertainty is essential.

REFERENCES

Bradford, E., Imsland, L., and del Rio-Chanona, E.A. (2019). Nonlinear model predictive control with explicit back-offs for gaussian process state space models. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 4747–4754. doi:10.1109/CDC40024.2019.9029443.

Bradford, E., Imsland, L., Zhang, D., and del Rio Chanona, E.A. (2020). Stochastic data-driven model predictive control using gaussian processes. *Computers & Chemical Engineering*, 139, 106844. doi: https://doi.org/10.1016/j.compchemeng.2020.106844. URL https://www.sciencedirect.com/science/article/pii/S0098135419313080.

Gaudin, A., Schuhmann, R., and Schlechten, A. (1942). Flotation kinetics. ii: The effect of size on the behavior of galena particles. *Journal of Physical Chemistry*, 46(8), 902–910. doi:10.1021/j150422a013.

González, R.A. and Quintanilla, P. (2024). Grey-box recursive parameter identification of a nonlinear dynamic model for mineral flotation. In *2024 10th International Conference on Control, Decision and Ianformation Technologies (CoDIT)*, 2967–2972. doi:10.1109/CoDIT62066.2024.10708161.

GPyTorch Team (2021). Metrics for gaussian process models. https://docs.gpytorch.ai/en/v1.8.1/examples/00_Basic_Usage/Metrics.html.

Hu, W. (2014). *Flotation Circuit Optimisation and Design.* Ph.D. thesis, Imperial College London. http://spiral.imperial.ac.uk/handle/10044/1/24805.

Maldonado, M., Desbiens, A., Del Villar, R., and Quispe, R. (2007). Towards the optimization of flotation columns using predictive control. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, 12(1), 75–80. doi:10.3182/20070821-3-CA-2919.00011.

McAllister, R. and Rasmussen, C.E. (2017). Data-efficient reinforcement learning in continuous state-action gaussian-pomdps. *Advances in Neural Information Processing Systems*, 30.

Mesa, D., Quintanilla, P., and Reyes, F. (2022). Bubble Analyser - An open-source software for bubble size measurement using image analysis. *Minerals Engineering*, 180. doi:10.1016/j.mineng.2022.107497.

Olivier, L.E. and Craig, I.K. (2017). A survey on the degree of automation in the mineral processing industry. *2017 IEEE AFRICON: Science, Technology and Innovation for Africa, AFRICON 2017*, 404–409. doi:10.1109/AFRCON.2017.8095516.

Park, S.S., Park, Y.J., Min, Y., and Choi, H.L. (2022). Online gaussian process state-space model: Learning and planning for partially observable dynamical systems. *International Journal of Control, Automation and Systems*, 20(2), 601–617.

Putz, E. and Cipriano, A. (2015). Hybrid model predictive control for flotation plants. *Minerals Engineering*, 70, 26–35. doi:10.1016/j.mineng.2014.08.013.

Quintanilla, P., Navia, D., Neethling, S., and Brito-Parada, P. (2023a). Evaluation of Changes in Feed Particle Size within an Economic Model Predictive Control Strategy for Froth Flotation. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, 56(2), 2317–2322. doi:10.1016/j.ifacol.2023.10.1200.

Quintanilla, P., Fernández, F., Mancilla, C., Rojas, M., and Navia, D. (2025). Digital twin with automatic disturbance detection for an expert-controlled sag mill. *Minerals Engineering*, 220, 109076. doi:https://doi.org/10.1016/j.mineng.2024.109076. URL https://www.sciencedirect.com/science/article/pii/S0892687524005053.

Quintanilla, P., Navia, D., Moreno, F., Neethling, S.J., and Brito-Parada, P.R. (2023b). A methodology to implement a closed-loop feedback-feedforward level control in a laboratory-scale flotation bank using peristaltic pumps. *MethodsX*, 10, 102081.

Quintanilla, P., Navia, D., Neethling, S.J., and Brito-Parada, P.R. (2023c). Economic model predictive control for a rougher froth flotation cell using physics-based models. *Minerals Engineering*, 196, 108050. doi:10.1016/J.MINENG.2023.108050.

Quintanilla, P., Neethling, S.J., Mesa, D., Navia, D., and Brito-Parada, P.R. (2021a). A dynamic flotation model for predictive control incorporating froth physics. part ii: Model calibration and validation. *Minerals Engineering*, 173, 107190. doi:10.1016/j.mineng.2021.107190.

Quintanilla, P., Neethling, S.J., Navia, D., and Brito-Parada, P.R. (2021b). A dynamic flotation model for predictive control incorporating froth physics. Part I: Model development. *Minerals Engineering*, 173, 107192. doi:10.1016/j.mineng.2021.107192.