

CIRP Conference on Manufacturing Systems (CIRP CMS)

Joint Parameter and State-Space Bayesian Optimization: Using Process Expertise to Accelerate Manufacturing Optimization

 Saksham Kiroirwal^a, Julius Pfrommer^a, Jürgen Beyerer^{a,b}

{saksham.kiroirwal, julius.pfrommer, juergen.beyerer}@iosb.fraunhofer.de

^aCognitive Industrial Systems, Fraunhofer IOSB, Karlsruhe, Germany

^bKarlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Abstract

Bayesian optimization (BO) is a powerful method for optimizing black-box manufacturing processes, but its performance is often limited when dealing with high-dimensional multi-stage systems, where we can observe intermediate outputs. Standard BO models the process as a black box and ignores the intermediate observations and the underlying process structure. Partially Observable Gaussian Process Networks (POGPN) model the process as a Directed Acyclic Graph (DAG). However, using intermediate observations is challenging when the observations are high-dimensional state-space time series. Process-expert knowledge can be used to extract low-dimensional latent features from the high-dimensional state-space data. We propose POGPN-JPSS, a framework that combines POGPN with Joint Parameter and State-Space (JPSS) modeling to use intermediate extracted information. We demonstrate the effectiveness of POGPN-JPSS on a challenging, high-dimensional simulation of a multi-stage bioethanol production process. Our results show that POGPN-JPSS significantly outperforms state-of-the-art methods by achieving the desired performance threshold twice as fast and with greater reliability. The fast optimization directly translates to substantial savings in time and resources. This highlights the importance of combining expert knowledge with structured probabilistic models for rapid process maturation.

© 2026 The Authors. Published by Elsevier B.V.

 This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

Peer-review under responsibility of the scientific committee of the 2025 CIRP CMS

Keywords: Bayesian Optimization; Partially Observable Gaussian Process Network; Manufacturing Optimization; Joint Parameter and State-Space Optimization

1. Introduction

The goal of a manufacturing process is to increase the throughput and reduce expenses and operational bottlenecks [1]. A common bottleneck is unoptimized or immature processes, which lead to billions of dollars in annual costs, high resource consumption, or waste or lost production time [2]. In many industrial settings, the primary objective is to reach a performance threshold—not necessarily perfection—in the shortest possible time. Achieving the threshold faster also minimizes resource consumption and opportunity costs.

Figure 1 shows a multi-stage bioethanol production process as described by [3, 4, 5]. We discuss the bioethanol production process in detail in Section 4. Conventionally, such a system is modeled as a black-box $\mathbf{P}$ that takes controllable inputs $\mathbf{x}$ and yields a product, whose performance is quantified by a met-




Figure 1: A multi-stage bioethanol production process $\mathbf{P}$ with three subprocesses: $P^{(1)}$, $P^{(2)}$, and $P^{(3)}$, each with controllable inputs $\mathbf{x} = (\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)})$. Using sensors, we are also able to observe the state-space $\mathbf{S}^{(1)}$, $\mathbf{S}^{(2)}$ and $\mathbf{S}^{(3)}$ of the subprocesses respectively. The final process performance, $\mathbf{y}^{(V)}$, is measured at the objective node V .

ric $\mathbf{y}^{(V)}$, such as the space-time yield (STY) of ethanol [3, 6]. Bayesian optimization (BO) is a powerful black-box optimization technique widely used to maximize performance metrics in manufacturing [7, 8, 9, 10].

However, modeling a process network as a simple black box can be suboptimal, as it ignores prior knowledge of its internal structure. As illustrated in Figure 1, this process can be represented as a causal Directed Acyclic Graph (DAG) $\mathcal{A}$, where

nodes $\{P^{(v)}\}_{v \in \mathcal{V}}$ correspond to individual subprocesses. We assume the DAG $\mathcal{A}$ is topologically ordered, with the objective node $P^{(V)}$ being the final node. Each subprocess $P^{(v)}$ receives controllable inputs $\mathbf{x}^{(v)}$ and the outputs $\tilde{\mathbf{f}}^{\mathcal{P}a(v)}$ from its parent subprocesses $\mathcal{P}a(v)$ to produce its own output $\tilde{\mathbf{f}}^{(v)}$. We are able to observe the outputs $\tilde{\mathbf{f}}^{(v)}$ as noisy observations $\mathbf{y}^{(v)}$ using sensors. However, we often observe high-dimensional state-space data $\mathbf{S}^{(v)}$ within each subprocess as intermediate observations. For example, during bioethanol production, we can measure temperature and pH in the bioreactor over time.

Gaussian Process Networks (GPNs) introduced by [11, 12], later extended by [13] as Partially Observable GPN (POGPN), showed that BO can be improved by replacing the single-task GP surrogate model with a POGPN and incorporating prior knowledge of causal structure and intermediate observations $\mathbf{g} = (\mathbf{y}^{(v)})_{v \in \mathcal{V}}$ of $\mathcal{A}$. While powerful, both GPNs and POGPNs are not designed to handle high-dimensional state-space observations $\mathbf{S}^{(v)}$, such as the time-series data common in modern manufacturing, which remains a key challenge.

Separately, the Joint Parameter and State-Space (JPSS) modeling proposed by [14] showed how process expert knowledge can be used to extract low-dimensional latent features $\mathbf{h}^{(v)}$ from the high-dimensional state-space $\mathbf{S}^{(v)}$ to create a Gaussian Process Network (GPN). We present the following contributions:

- Importance of process expert knowledge in BO and introduce POGPN-JPSS, a novel framework for optimizing complex processes with high-dimensional intermediate observations.
- Empirically show that POGPN-JPSS can rapidly mature the multi-stage bioethanol production process.

Section 2 provides the background information on the Gaussian Processes and Bayesian Optimization. Section 3 discusses the related work in detail. The described methods are used for the multi-stage bioethanol production process, described in Section 4. Sections 5 and 6 describe the experiments and results. Section 7 discusses the results and provides the conclusion and outlook.

2. Background

We consider a manufacturing process $\mathbf{P}$ composed of multiple subprocesses $\{P^{(v)}\}_{v \in \mathcal{V}}$. When the process is evaluated at a given input configuration $\mathbf{x}_n = (\mathbf{x}_n^{(v)})_{v \in \mathcal{V}}$, we observe intermediate outputs $\mathbf{g}_n = (\mathbf{y}_n^{(v)})_{v \in \mathcal{V}}$ and a final objective output $\mathbf{y}_n^{(V)}$. A dataset of N evaluations is denoted as $\mathcal{D} = \{\mathbf{x}_n, \mathbf{g}_n, \mathbf{y}_n^{(V)}\}_{n=1}^N$.

2.1. Stochastic Variational Inference for Gaussian Processes

A Single-Task Gaussian Process (STGP) is a Bayesian non-parametric model that defines a distribution over functions $f(\cdot)$ [15]. For a given input $\mathbf{x}_n$, the function value $f_n^{(V)}$ is assumed to be drawn from a Gaussian distribution $p(f_n^{(V)}|\mathbf{x}_n) = \mathcal{N}(\mu(\mathbf{x}_n), \sigma(\mathbf{x}_n, \mathbf{x}_n'))$, where $\mu(\cdot)$ is the mean function and $\sigma(\cdot, \cdot)$ is the covariance function. We also assume that the likelihood

$p(\mathbf{y}_n^{(V)}|f_n^{(V)})$ of observing $\mathbf{y}_n^{(V)}$ is Gaussian. An STGP treats the entire process $\mathbf{P}$ as a single black box, and hyperparameters are learned by maximizing the marginal log-likelihood (MLL)

$$\mathcal{L}_{\text{GP}} = \sum_{n=1}^N \log \mathbb{E}_{p(f_n^{(V)}|\mathbf{x}_n)} [p(\mathbf{y}_n^{(V)}|f_n^{(V)})]. \quad (1)$$

However, the cubic computational complexity of STGP inference makes it impractical for large datasets or hierarchical models [15, 16]. The Stochastic Variational Gaussian Process (SVGP) [16] makes inference scalable by introducing a set of J inducing points $\mathbf{u}$, which are function values at chosen inducing locations $\mathbf{Z}$. These inducing points are summarized by a variational distribution $q(\mathbf{u}) = \mathcal{N}(\mathbf{a}_u, \mathbf{B}_u)$. For input $\mathbf{x}_n$, SVGP approximates the true GP posterior with a variational distribution $q(f_n^{(V)}) = \int q(f_n^{(V)}|\mathbf{u})q(\mathbf{u})d\mathbf{u}$. The GP hyperparameters and variational parameters $(\mathbf{Z}, \mathbf{a}_u, \mathbf{B}_u)$ are then optimized by maximizing the Evidence Lower Bound (ELBO) as

$$\mathcal{L}_{\text{ELBO}}^{\text{SVGP}} = \sum_{n=1}^N \mathbb{E}_{q(f_n^{(V)})} [\log p(\mathbf{y}_n^{(V)}|f_n^{(V)})] - \beta \text{KL}(q(\mathbf{u})||p(\mathbf{u})). \quad (2)$$

2.2. Gaussian Process Networks (GPNs)

To model the structured processes, Gaussian Process Network (GPN) [11] represents the process $\mathbf{P}$ as a Directed Acyclic Graph (DAG) $\mathcal{A}$. Each node $v \in \mathcal{V}$ in the graph corresponds to a subprocess $P^{(v)}$ and is modeled as an independent GP. A key assumption in GPNs is that the noisy observations $\mathbf{y}^{\mathcal{P}a(v)}$ from parent nodes are used as direct inputs to their child nodes, thus modeling the process as a network of GPs, $\mathcal{M}_{\mathbf{x} \rightarrow \mathbf{g} \rightarrow \mathbf{y}^{(V)}}$. Each node's GP is trained on its respective input-output pairs, $(\mathbf{x}^{(v)}, \mathbf{y}^{\mathcal{P}a(v)})$ and $\mathbf{y}^{(v)}$, by maximizing the MLL as in (1). Each node can be represented as a tuple

$$P^{(v)} = \langle p(\mathbf{f}^{(v)}|\mathbf{x}^{(v)}, \mathbf{y}^{\mathcal{P}a(v)}), p(\mathbf{y}^{(v)}|\mathbf{f}^{(v)}) \rangle. \quad (3)$$

2.3. Partially Observable Gaussian Process Networks (POGPNs)

The GPN's assumption that noisy parent observations $\mathbf{y}^{\mathcal{P}a(v)}$ are the direct inputs to child nodes can be limiting, as it fails to distinguish between the true underlying process output and measurement noise [13]. The Partially Observable Gaussian Process Network (POGPN), developed by Kiroriwal et al. [13], provides a more general and robust framework. As shown in Figure 1, it is not the noisy observation $\mathbf{y}^{\mathcal{P}a(v)}$ that becomes the input but rather the latent output $\tilde{\mathbf{f}}^{\mathcal{P}a(v)}$. Using this as the key idea, POGPN models each subprocess (node) as an SVGP, where each node can be represented as a tuple

$$P^{(v)} = \langle p(\mathbf{f}^{(v)}|\mathbf{x}^{(v)}, \tilde{\mathbf{f}}^{\mathcal{P}a(v)}), p(\mathbf{y}^{(v)}|\mathbf{f}^{(v)}) \rangle. \quad (4)$$

Due to POGPN's hierarchical structure, exact inference is intractable. Instead, POGPN uses doubly stochastic variational inference (DSVI) [17] to jointly optimize the entire network. Kiroriwal et al. [13] formulate the ELBO, $\mathcal{L}_{\text{ELBO}}^{\text{POGPN}}$, as

$$\underbrace{\sum_{v \in \mathcal{V}} \sum_{n=1}^{N^{(v)}} \mathbb{E}_{q(\mathbf{f}_n^{(v)})} [\log p(\mathbf{y}_n^{(v)} | \mathbf{f}_n^{(v)})]}_{\text{LL loss}} - \beta \underbrace{\sum_{v \in \mathcal{V}} \text{KL}(q(\mathbf{u}^{(v)}) || p(\mathbf{u}^{(v)}))}_{\text{KL loss}}, \quad (5)$$

where $\mathbf{u}^{(v)}$ are the inducing points for the SVGP at each node. Inspired by [18], [13] also propose a Predictive Log Likelihood (PLL) for POGPN where the ‘KL loss’ remains the same as in (5) and the ‘LL loss’ term is defined as

$$\text{LL}_{\text{PLL}} = \underbrace{\sum_{v \in \mathcal{V}} \sum_{n=1}^{N^{(v)}} \log \mathbb{E}_{q(\mathbf{f}_n^{(v)})} [p(\mathbf{y}_n^{(v)} | \mathbf{f}_n^{(v)})]}_{\text{Predictive Log-Likelihood}}. \quad (6)$$

2.4. Bayesian Optimization: Expected Improvement (EI)

In Bayesian optimization, Expected Improvement (EI) is a popular acquisition function. Given the best-observed output, $y_\star^{(V)}$, the EI at a candidate point $\mathbf{x}_c$ is the expected value of the improvement over the best observed output and is defined as $\text{EI}(\mathbf{x}_c) = \mathbb{E}_{p(\mathbf{y} | \mathbf{x}_c, \mathcal{D})} [\max(0, y_c^{(V)} - y_\star^{(V)})]$. When the model’s predictive posterior is Gaussian, $\mathcal{N}(a_{p(\mathbf{f}_c^{(V)})}, B_{p(\mathbf{f}_c^{(V)})})$, EI has a closed-form solution [7] as

$$\text{EI}(\mathbf{x}) = B_{p(\mathbf{f}_c^{(V)})} \Pi \left(\frac{a_{p(\mathbf{f}_c^{(V)})} - y_\star^{(V)}}{B_{p(\mathbf{f}_c^{(V)})}} \right), \quad \Pi(\eta) = \phi(\eta) + \eta \Phi(\eta), \quad (7)$$

where $\eta = (a_{p(\mathbf{f}_c)} - y_\star) / B_{p(\mathbf{f}_c)}$, and $\Phi(\cdot)$ and $\phi(\cdot)$ are the Cumulative Distribution Function (CDF) and Probability Density Function (PDF) of the standard normal distribution, respectively. Numerical issues of EI were first analyzed by [19], and they proposed Log Expected Improvement (LogEI) as $\text{LogEI}(\mathbf{x}) = \log \Pi(\eta) + \log B_{p(\mathbf{f}_c^{(V)})}$ as a robust and stable alternative to EI.

2.5. Inducing Point Allocator (IPA)

The placement of inducing locations significantly impacts the performance of an SVGP [20]. Initialization strategy Greedy Variance Reduction (GVR), introduced by [20], iteratively selects points from the observed inputs $\mathbf{X}$ that maximally reduce the model’s variance: $\mathbf{z}_j = \arg \max_{\mathbf{z} \in \mathbf{X}} \sigma_{j-1}(\mathbf{z})$, where $\sigma_{j-1}^2(\mathbf{z}) = \sigma(\mathbf{z}, \mathbf{z}) - \sigma(\mathbf{Z}_{1:j-1}, \mathbf{z})^T \sigma(\mathbf{Z}_{1:j-1}, \mathbf{Z}_{1:j-1})^{-1} \sigma(\mathbf{Z}_{1:j-1}, \mathbf{z})$.

Moss et al. [21] showed that GVR is not optimal for Bayesian optimization using SVGP and introduced a quality function $\phi(\cdot)$ into the selection criterion $\mathbf{z}_j = \arg \max_{\mathbf{z} \in \mathbf{X}} \phi(\mathbf{z}) \sigma_{j-1}(\mathbf{z})$. For improved BO performance, they proposed Greedy Improvement Reduction (GIR) that uses an EI-inspired quality function $\phi(\mathbf{x}_n) = y_n^{(V)} - y_\star^{(V)}$.

3. Related Work

Bayesian optimization (BO) with a Gaussian Process (GP) [15] surrogate model is an effective methodology for manufacturing processes, with applications ranging from chemical synthesis [10] to mechanical engineering [22, 23]. Research in this direction can be broadly categorized into three main categories: (1) development of robust acquisition functions, (2) hi-

erarchical surrogate models using process structure and intermediate data, and (3) improving the scalability of standard GPs to handle high-dimensional inputs.

Works by [10, 23, 9] belong to the first category, in which the authors develop robust acquisition functions either using trust region or other numerically stable techniques to apply existing BO methods for optimizing complex processes. In particular, TuRBO [9] and LogEI [19] acquisition functions have emerged as a consistent, robust choice across diverse BO applications.

In the second category, motivated by the fact that many manufacturing processes consist of interconnected subprocesses, Gaussian Process Networks (GPNs) [11, 12, 24] model the system as a DAG, using intermediate observations from each subprocess to improve the surrogate’s accuracy and, consequently, optimization performance. For instance, causal BO with GPNs was applied to optimize fertilizer dosage in farming with respect to the trade-off between crop growth and soil health [12]. However, as shown by [13], GPNs assume that the GPs at each node can be trained conditionally independently, resulting in suboptimal inference. To address this, the Partially Observable GPN (POGPN) uses variational inference to learn a latent distribution of the process network [13]. Results show that POGPN can improve BO performance by incorporating partial observations during both training and inference. However, both GPNs and POGPNs are limited in their ability to handle high-dimensional data, such as sensor time series.

A third direction of research by [25] aims to enhance standard single-task GP for high-dimensional inputs by using dimensionally-scaled priors for the GP kernel and likelihood. Combined with the robust LogEI [19] acquisition function, the results from [25] provide robust results for not only high-dimensional input processes. While the method is effective, it relies on a direct input-output mapping, and it does not incorporate the high-dimensional intermediate observations $\mathbf{S}^{(v)}$ from subprocesses.

Joint Parameter and State-Space (JPSS) modeling was introduced to handle high-dimensional state-space observations $(\mathbf{S}^{(v)})_{v \in \mathcal{V}}$ by using expert knowledge to extract low-dimensional features $\mathbf{h} = (\mathbf{h}^{(v)})_{v \in \mathcal{V}}$ for a GPN [14]. However, the original JPSS framework was proposed using a standard GPN, thus inheriting its limitations in uncertainty propagation. This paper integrates POGPN inference with JPSS to optimize process networks, such as the multi-stage bioethanol production process, with high-dimensional intermediate observations.

4. Motivating Example

The bioethanol production process, illustrated in Fig. 1, uses a seed train fermentation technology [4, 5, 3]. This is done by scaling up a microbial culture in a small laboratory flask to a large scale in three stages $\{P^{(v)}\}_{v \in \{1,2,3\}}$. This scaling up is called a seed train because of its sequential nature. This gives rise to a tightly coupled, multi-stage dynamical system in which the performance of the early stages critically determines the final outcome.




Figure 2: Observed state-space dynamics from a single simulation run of the multi-stage bioethanol production process.

The whole process, denoted by $\mathbf{P}$, is governed by an input vector with a dimensionality of $D_{\mathbf{x}} = 26$, denoted by $\mathbf{x}$, with dimensionalities $D_{\mathbf{x}^{(1)}} = 6$, $D_{\mathbf{x}^{(2)}} = 10$, and $D_{\mathbf{x}^{(3)}} = 10$, respectively. The input variables represent operating conditions, including substrate feed rate profiles, linear temperature and pH profiles, and stage durations. The objective function is to optimize the controllable input variables $\mathbf{x}$ to achieve maximum final process productivity $y^{(V)}$, quantified as Space-Time Yield (STY) [3]. STY, defined as the amount of ethanol produced per unit of time and reactor volume, penalizes inefficiently long stages and thus drives the optimization toward overall process efficiency.

The sequence begins with $P^{(1)}$ (*Shake Flask*), which involves activating the biomass originally obtained from a lab stock using a 2L flask. The next step is $P^{(2)}$ (*Seed Fermenter*), which cultivates the activated biomass in a 500 L bioreactor operating in fed-batch mode to initiate a dense cell culture ideal for production. The objective is to allow healthy cell growth by precisely manipulating factors such as temperature, pH, and substrate addition to avoid inhibitory effects. Then, final ethanol production occurs in $P^{(3)}$ (*Production Fermenter*), a 20000 L bioreactor, which uses biomass from culturing to break down a substrate on a larger scale to produce ethanol.

Process Dynamics and Control: The dynamics of each subprocess are governed by a set of ordinary differential equations (ODEs) that describe the evolution of concentration of biomass (s_{biomass}), ethanol (s_{ethanol}), and substrate ($s_{\text{substrate}}$), and the culture volume (s_{volume}) [26] as

$$\begin{aligned} \frac{ds_{\text{biomass}}}{dt} &= (\mu_{\text{growth}} - k_d)s_{\text{biomass}} - \frac{F_{\text{in}}}{s_{\text{volume}}}s_{\text{biomass}} \\ \frac{ds_{\text{ethanol}}}{dt} &= q_{\text{ethanol}}s_{\text{biomass}} - \frac{F_{\text{in}}}{s_{\text{volume}}}s_{\text{ethanol}} \\ \frac{ds_{\text{substrate}}}{dt} &= \frac{F_{\text{in}}}{s_{\text{volume}}}(s_{\text{in}} - s_{\text{substrate}}) - q_{\text{substrate}}s_{\text{biomass}} \\ \frac{ds_{\text{volume}}}{dt} &= F_{\text{in}}, \end{aligned} \quad (8)$$

where k_d is the death rate, F_{in} is the feed rate, s_{in} is the inlet substrate concentration, s_{volume} is the volume of the complete culture in the reactor, and μ_{growth} is the specific biomass growth rate. The specific ethanol production rate q_{ethanol} is described by the Luedeking-Piret model, and the specific substrate consumption rate $q_{\text{substrate}}$ follows the Pirt model, which ac-

counts for substrate used for growth, production, and cell maintenance [27, 28].

The Optimization Challenge: The bioethanol seed train $\mathbf{P}$ is a challenge to optimize because of several factors, such as:

- **Complex Biological Kinetics:** The dynamics are non-linear and involve growth- and non-growth-related product formation (Luedeking-Piret kinetics) [27], substrate consumption related to cell maintenance (Pirt model) [28], and strong inhibitory effects due to high substrate concentration (Andrews model) [29] and inhibitory effects due to high ethanol concentration (Ghose-Tyagi model) [30].
- **Critical Trade-offs Across Stages:** The optimizer has to make trade-offs between multiple objectives. For instance, minimizing seed stages reduces overall process time but can result in low-quality inoculum, thereby negatively affecting final productivity [4]. A very aggressive feeding mode to enhance productivity can lead to substrate and product inhibitions, causing retardation in the final process.
- **Tightly Coupled Dynamics:** The decisions that are made within the early stages impact other stages. This complicates the problem significantly because it is non-trivial to optimize variables within a control system at the holistic level.

These factors make the bioethanol seed train a highly suitable and demanding benchmark for assessing structure-aware optimization algorithms. An example of the observed state-space from different stages of the bioethanol production process is shown in Figure 2. A threshold STY required to make this pilot-scale simulation economically viable can be 0.5. Importantly, this state-space representation provides information on how long each of these processes is supposed to take. Notably, the optimizer must also determine the optimal duration for each subprocess— $P^{(1)}$, $P^{(2)}$, and $P^{(3)}$ —which can range from 6-48 hours, 12-96 hours, and 48-168 hours, respectively. Given that each simulation represents a very time-intensive real-world process, fast optimization of such a time-intensive process is of utmost importance, as it can save a lot of resources and time and hence opportunity costs as well.¹

5. Methodology

As explained in Section 4, optimizing the bioethanol production process is a significant challenge. In this section, we evaluate the performance of our proposed POGPN-JPSS approach, which integrates Partially Observable Gaussian Process Networks (POGPN) [13] with Joint Parameter and State Space (JPSS) representation [14]. We compare our method against three state-of-the-art models: a standard GPN [11], a high-dimensional single-task GP (STGP) [25], and a standard

¹ The code for the simulation is available at https://github.com/Sam4896/seed_train_bioethanol_sim




Figure 3: Histograms of key process observations from 50 random simulation runs. The observations shown are the final biomass $s_{\text{biomass}}^{(1)}$, ethanol $s_{\text{ethanol}}^{(1)}$ concentrations from $P^{(1)}$, the final ethanol concentrations $s_{\text{ethanol}}^{(2)}$ and $s_{\text{ethanol}}^{(3)}$ from $P^{(2)}$ and $P^{(3)}$, and the final Space-Time Yield (STY) $y^{(V)}$ for the entire process.




Figure 4: The Directed Acyclic Graph (DAG) for the bioethanol process is shown alongside the Bayesian optimization performance of different surrogate models. The plots show the mean and standard deviation of the best-observed Space-Time Yield (STY) with an optimization budget of 100 iterations. It compares the performance of POGPN-PLL (ours) with existing state-of-the-art methods.

SVGP [21]. All optimizations are driven by the robust Log Expected Improvement (LogEI) acquisition function [19].

The state-space observations $\mathbf{S}^{(1)}$, $\mathbf{S}^{(2)}$, and $\mathbf{S}^{(3)}$ from each stage are high-dimensional multivariate time-series. We reduce these to low-dimensional features using the JPSS approach, incorporating process-expert knowledge to identify features essential to subsequent stages and the final objective. The final objective in this case is to optimize STY, which depends on the total ethanol produced, the process volume, and the total duration. For $P^{(1)}$, our objective is to optimize a healthy inoculum to be introduced into the seed fermenter. Both biomass and ethanol concentrations are key indicators of culture health, as ethanol can inhibit growth. Therefore, we extract the final biomass and ethanol concentrations, $\mathbf{h}^{(1)} = (s_{\text{biomass}}^{(1)}, s_{\text{ethanol}}^{(1)})$, as the intermediate features. For $P^{(2)}$ and $P^{(3)}$ our objective is to optimize ethanol production, so we use the final ethanol concentrations, $h^{(2)} = s_{\text{ethanol}}^{(2)}$ and $h^{(3)} = s_{\text{ethanol}}^{(3)}$, as the respective intermediate observations. STY of the entire process depends on the total ethanol output, process volume, and total duration. The DAG used for our structured models, as shown in Figure 4.

To better understand this optimization challenge, we show the histogram plots of the discussed intermediate observations for 50 randomly performed simulation runs in Figure 3. Although the observations $\mathbf{h}^{(1)}$ are relatively spread out, outputs observed from the other subprocess and the STY are mainly concentrated close to zero and have a long tail. This distribution is particularly suggestive of regions corresponding to high-performing STY being relatively sparse in the 26-dimensional input space.

6. Benchmark Results

We compare the performance of POGPN-JPSS (using both ELBO and PLL objectives) against state-of-the-art models: GPN [11], STGP [25], and SVGP-PLL [21] with PLL objective [18]. All models were implemented in BoTorch [31] with a Matern-5/2 kernel, a Gaussian likelihood, and the dimensionally-scaled priors proposed by [25]. For POGPN, the number of inducing points per node was set to the number of training data points for that node. Inducing points for the objective node were initialized using Greedy Improvement Reduction (GIR) [21], while all other nodes used Greedy Variance Reduction (GVR) [20]. We used the Adam optimizer [32] with a learning rate of 0.02 and an exponential moving average stopping criterion. All observations were log-transformed and standardized before model training to improve numerical stability and the exploration-exploitation trade-off. Each optimization had a budget of 100 iterations, starting with an initial dataset of 50 randomly sampled points, and all experiments were repeated 5 times to report the average performance.

As depicted in Figure 4, POGPN-PLL is able to achieve an optimal solution for bioethanol production much faster and more robustly than any other state-of-the-art methods. POGPN-PLL provides a higher STY value with reduced variability. Most importantly, POGPN-PLL reaches the 0.5 STY threshold in 50 iterations, whereas other state-of-the-art approaches require nearly 100 iterations. Given the duration of the subprocess, as discussed in Section 4, this accelerated optimization results in an average savings of 300-350 days of production time, representing a substantial reduction in opportunity costs, materials, and energy use, while increasing overall productivity.

7. Discussion and Outlook

We show that the Partially Observable Gaussian Process Network - Predictive Log-Likelihood (POGPN-PLL), when combined with the Joint Parameter and State-Space (JPSS), significantly outperforms state-of-the-art models in optimizing the complex, multi-stage bioethanol production process. Faster, more robust convergence to the optimal solution can be attributed to POGPN-PLL's ability to learn the problem's latent structure. It is achieved by creating a process network that uses process expert knowledge to extract low-dimensional features from the high-dimensional state space. The rapid optimization translates directly into economic benefits, including reduced opportunity costs and savings in materials and energy.

Process expert knowledge in the form of feature extraction enables the model to exploit high-dimensional sensor data and the process structure, leading to faster, more robust convergence to the optimal solution. POGPN is a promising framework for Bayesian optimization of real-world industrial processes, where time and material efficiency are essential.

Future work could investigate the performance gap between the PLL and the standard Evidence Lower Bound (ELBO) objective for POGPN, which may be attributed to how each objective learns the process stochasticity. It would also be interesting to explore feature-extraction and causal-graph-discovery methods combined with POGPN-JPSS for Bayesian optimization, thereby extending POGPN's applicability to systems where the structure is not well-understood.

The research in this paper was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 459291153 (FOR5399).

References

- [1] E. M. Goldratt, J. Cox, D. Whitford, The goal, HighBridge, 2006.
- [2] IDS-INDATA, [The real cost of manufacturing downtime \(2020–2025\): Sector insights & forecast](https://idsindata.co.uk/manufacturing-downtime-costs-and-forecasting/), Tech. rep., accessed: 2025-08-26 (2025). URL <https://idsindata.co.uk/manufacturing-downtime-costs-and-forecasting/>
- [3] B. P. Facility, [Fermentation technology: Scale-up and seed train operations](https://bpf.eu/fermentation/), industrial guideline and case example (2023). URL <https://bpf.eu/fermentation/>
- [4] T. Hernández Rodríguez, R. Pörtner, B. Frahm, Seed train optimization for suspension cell culture, in: BMC Proceedings, Vol. 7, Springer, 2013, p. P9.
- [5] G. F. Institute, [Scaling up animal cell culture: Technical challenges and opportunities for biomanufacturing](https://gfi.org/resource/scaling-up-animal-cell-culture/), technical white paper (2020). URL <https://gfi.org/resource/scaling-up-animal-cell-culture/>
- [6] O. J. Sánchez, C. A. Cardona, Trends in biotechnological production of fuel ethanol from different feedstocks, *Bioresource Technology* 99 (13) (2008) 5270–5295. doi:10.1016/j.biortech.2007.11.013.
- [7] P. I. Frazier, A tutorial on bayesian optimization, arXiv preprint arXiv:1807.02811 (2018).
- [8] S. Sano, T. Kadowaki, K. Tsuda, S. Kimura, Application of bayesian optimization for pharmaceutical product development, *Journal of Pharmaceutical Innovation* 15 (3) (2020) 333–343.
- [9] Q. Liang, L. Lai, Scalable bayesian optimization accelerates process optimization of penicillin production, in: NeurIPS 2021 AI for Science Workshop, 2021.
- [10] B. J. Shields, J. M. Stevens, J. Li, M. Parasram, F. N. Damani, J. I. M. Alvarado, J. Janey, R. P. Adams, A. Doyle, Bayesian reaction optimization as a tool for chemical synthesis, *Nature* 590 (2021) 89 – 96. doi:10.1038/s41586-021-03213-y.
- [11] R. Astudillo, P. Frazier, Bayesian optimization of function networks, *Advances in neural information processing systems* 34 (2021) 14463–14475.
- [12] V. Aglietti, X. Lu, A. Paleyes, J. González, Causal bayesian optimization, in: International Conference on Artificial Intelligence and Statistics, PMLR, 2020, pp. 3155–3164.
- [13] S. Kiroriwal, J. Pfrommer, J. Beyerer, Partially observable gaussian process network and doubly stochastic variational inference, arXiv preprint arXiv:2502.13905 (2025).
- [14] S. Kiroriwal, J. Pfrommer, H. Mende, R. H. Schmitt, J. Beyerer, Joint parameter and state-space modelling of manufacturing processes using gaussian processes, in: 2024 IEEE 22nd International Conference on Industrial Informatics (INDIN), 2024, pp. 1–6.
- [15] C. E. Rasmussen, Gaussian processes in machine learning, in: Summer school on machine learning, Springer, 2003, pp. 63–71.
- [16] J. Hensman, N. Fusi, N. D. Lawrence, Gaussian processes for big data, arXiv preprint arXiv:1309.6835 (2013).
- [17] H. Salimbeni, M. Deisenroth, Doubly stochastic variational inference for deep gaussian processes, *Advances in neural information processing systems* 30 (2017).
- [18] M. Jankowiak, G. Pleiss, J. Gardner, Parametric gaussian process regressors, in: International conference on machine learning, PMLR, 2020, pp. 4702–4712.
- [19] S. Ament, S. Daulton, D. Eriksson, M. Balandat, E. Bakshy, Unexpected improvements to expected improvement for bayesian optimization, *Advances in Neural Information Processing Systems* 36 (2023) 20577–20612.
- [20] D. R. Burt, C. E. Rasmussen, M. Van Der Wilk, Convergence of sparse variational inference in gaussian processes regression, *Journal of Machine Learning Research* 21 (131) (2020) 1–63.
- [21] H. B. Moss, S. W. Ober, V. Picheny, Inducing point allocation for sparse gaussian processes in high-throughput bayesian optimisation, in: International Conference on Artificial Intelligence and Statistics, PMLR, 2023, pp. 5213–5230.
- [22] S. Chung, C.-H. Chou, X. Fang, R. A. Kontar, C. Okwudire, A multi-stage approach for knowledge-guided predictions with application to additive manufacturing, *IEEE Transactions on Automation Science and Engineering* 19 (2022) 1675–1687. doi:10.1109/tase.2022.3160420.
- [23] A. Brasington, J. Halbritter, R. Wehbe, R. Harik, Bayesian optimization for process planning selections in automated fiber placement, *Journal of Composite Materials* 56 (2022) 4275 – 4296. doi:10.1177/00219983221129010.
- [24] S. Sussex, A. Makarova, A. Krause, Model-based causal bayesian optimization, arXiv preprint arXiv:2211.10257 (2022).
- [25] C. Hvarfner, E. O. Hellsten, L. Nardi, Vanilla bayesian optimization performs great in high dimensions, arXiv preprint arXiv:2402.02229 (2024).
- [26] U. Salihu, U. Usman, A. Abubakar, G. Mansir, Effect of ph and temperature on bioethanol production: Evidences from the fermentation of sugarcane molasses using *saccharomyces cerevisiae*, *Dutse Journal of Pure and Applied Sciences* 8 (4b) (2022) 9–16.
- [27] R. Luedeking, E. L. Piret, A kinetic study of the lactic acid fermentation. batch process at controlled ph, *Journal of Biochemical and Microbiological Technology and Engineering* 1 (4) (1959) 393–412.
- [28] S. Pirt, The maintenance energy of bacteria in growing cultures, *Proceedings of the Royal Society of London. Series B. Biological Sciences* 163 (991) (1965) 224–231.
- [29] J. F. Andrews, A mathematical model for the continuous culture of microorganisms utilizing inhibitory substrates, *Biotechnology and Bioengineering* 10 (6) (1968) 707–723. doi:10.1002/bit.260100602.
- [30] T. Ghose, R. Tyagi, Rapid ethanol fermentation of cellulose hydrolysate. ii. product and substrate inhibition and optimization of fermentor design, *Biotechnology and Bioengineering* 21 (8) (1979) 1401–1420.
- [31] M. Balandat, B. Karrer, D. Jiang, S. Daulton, B. Letham, A. G. Wilson, E. Bakshy, Botorch: A framework for efficient monte-carlo bayesian optimization, *Advances in neural information processing systems* 33 (2020) 21524–21538.
- [32] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).