Stochastic Model Predictive Control Utilizing Bayesian Neural Networks

J. Pohlodek¹, H. Alsmeier¹, B. Morabito², C. Schlauch³, A. Savchenko¹, and R. Findeisen¹

Abstract—Integrating measurements and historical data can enhance control systems through learning-based techniques, but ensuring performance and safety is challenging. Robust model predictive control strategies, like stochastic model predictive control, can address this by accounting for uncertainty. Gaussian processes are often used but have limitations with larger models and data sets. We explore Bayesian neural networks for stochastic learning-assisted control, comparing their performance to Gaussian processes on a wastewater treatment plant model. Results show Bayesian neural networks achieve similar performance, highlighting their potential as an alternative for control designs, particularly when handling extensive data sets.

I. INTRODUCTION

Optimal, flexible, safe, and reliable control close to the feasibility boundary is becoming increasingly important to achieve many processes' economic, energy-efficient, and sustainable operation. Optimization-based control, such as model predictive control (MPC) [1], [2], is, in principle, well suited to achieve these objectives, as it allows the formulation of elaborate control goals while incorporating state and input constraints. However, the MPC performance heavily depends on the quality of the underlying process model [1], [3], as it is employed to predict the system's behavior and determine the optimal control action. This brings forward the question of counteracting inevitable model-plant mismatch and measurement uncertainties in real systems. Though nominal MPC exhibits inherent robustness properties [1], [4] through repeated optimization in the closed loop, uncertainties degrade its performance and can potentially lead to constraint violations or stability loss [1].

Learning-supported MPC approaches aim to reduce the uncertainty in model dynamics by augmenting it with datadriven parts. However, process uncertainty and stochasticity are often not considered in the MPC formulation if combined with learning approaches; one assumes that the learned part adjusts and compensates for these uncertainties.

In contrast, robust and stochastic MPC approaches [1], [5], [6] are explicitly designed for such a scenario, incorporating the information on model uncertainties directly into the MPC formulation. These formulations typically trade off some performance for guaranteeing constraint satisfaction under assumptions on the uncertainty, which is especially important in safety-critical scenarios.

Combining the two strategies — robust model predictive control strategies and learning — has been a subject of recent

³ Humboldt-Universität zu Berlin

work [7], [8], [9], [10], [11]. The learned model is often expected to provide some measure of the uncertainty, which is then delivered to the robust or stochastic MPC. Gaussian processes (GPs) have been widely employed in this role, as they explicitly yield a posterior variance along with the regression mean [12]. Though GPs are generally easy to tune, their computational complexity grows cubicle with the dataset size, restricting their application to relatively small sets [12]

An alternative is the use of Bayesian neural networks (BNNs) designed to provide a measure of uncertainty besides the regression [13]. Belonging to the family of deep learning models, it is computationally efficient, especially as most of the computational effort is spent offline during training. However, compared to other deep learning methods, BNNs are not yet as widely researched — it is generally challenging to infer the distributions analytically. Instead, one relies on posterior approximations [13].

In this work, we aim to determine the suitability of BNNs for (stochastic) predictive control applications, with a particular focus on comparing them to state-of-the-art GPs. BNNs have been used in combination with model predictive control, e.g., [14] employs BNNs in combination with a hierarchical MPC approach for the control of a surgical robot to reduce the uncertainty. Here, the kinematics and dynamics of a highly nonlinear robotic system were modeled with the help of BNNs. The uncertainty information provided by the BNN is used in a hierarchical MPC scheme, achieving superior performance. In [15], a learning-based adaptivescenario-tree model predictive control (MPC) approach is used to achieve probabilistic safety guarantees using BNNs to learn the model uncertainty.

We produced all results in this work using the opensource Python toolbox HILO-MPC¹ [16]. In a simple-to-use way, the toolbox allows combining (robust) predictive, and optimization-based control and estimation approaches with methods from machine learning, such as Gaussian processes and Bayesian neural networks.²

II. PROBLEM FORMULATION AND METHODS

We consider dynamic systems which can be described by

$$x_{k+1} = f(x_k, u_k) + B(d(x_k, u_k) + w_k).$$
(1)

Here, $x \in \mathbb{R}^{n_x}$ are the dynamical states, $u \in \mathbb{R}^{n_u}$ are the inputs to the system and $B \in \mathbb{R}^{n_x \times n_d}$ is a known

¹ Control and Cyber-Physical Systems Laboratory, Technical University of Darmstadt, email: rolf.findeisen@iat.tu-darmstadt.de

² Yokogawa Insilico Biotechnology GmbH, Stuttgart

¹https://www.ccps.tu-darmstadt.de/

research_ccps/hilo_mpc/

 $^{^2 {\}rm The}$ case study used as an example, including all code, will be made freely available in HILO-MPC.

matrix. The function $f: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$ describes the known part of the system dynamics, while the function $d: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$ describes an unknown or difficult to model effect. This unknown effect will be learned using specific machine learning algorithms. The variable $w \in \mathbb{R}^{n_d}$ is assumed to be zero-mean normally distributed process noise $w \sim \mathcal{N}(0, \Sigma^w)$ with the variance matrix $\Sigma^w = \text{diag}\left(\left[\sigma_1^2 \dots \sigma_{n_d}^2\right]\right)$. We focus on Gaussian processes and Bayesian neural networks to learn the unknown effect, which we briefly introduce in the following.

Gaussian Process: We briefly review the main concepts of Gaussian processes. For a more detailed introduction to GPs, we refer to [12]. A GP is a stochastic supervised machine learning algorithm used in regression and classification tasks. GPs are less prone to overfitting and naturally provide uncertainty measures on predictions. We formulate a regression task as a mapping $\psi \colon \mathbb{R}^{n_{\chi}} \to \mathbb{R} : \chi \to \varphi(\chi) + \nu$ with the input vector $\chi \in \mathbb{R}^{n_{\chi}}$, the output $\psi \in \mathbb{R}$ and the zero-mean normally distributed noise $\nu \sim \mathcal{N}(0, \sigma^2)$ affecting the output. The unknown function φ is assumed to be normally distributed $\varphi(\chi) \sim \mathcal{N}(m(\chi), k(\chi, \chi)),$ with the mean function $m \colon \mathbb{R}^{n_{\chi}} \to \mathbb{R}$ and covariance function $k \colon \mathbb{R}^{n_{\chi}} \times \mathbb{R}^{n_{\chi}} \to \mathbb{R}$. These functions are assumed to depend on hyperparameters, that are determined by GP training procedure, resulting in scalar-valued predictions. For simplicity, multiple outputs are handled as separate onedimensional GPs independently.

Bayesian Neural Network: As a second learning method, we consider Bayesian neural networks [17], which can also provide a measure of uncertainty similar to GPs. In contrast to GPs, computational complexity of BNNs does not grow with the number of data points. Unlike traditional feed-forward NNs, BNNs learn distributions over the output instead of single value predictions. To accomplish this task, the weights of conventional neural networks are replaced by distributions over the weight in each layer. Thus, BNNs extend the class of NNs to estimate posterior distributions instead of most likely values, cf. Fig. 1.

A common way of implementing the weights W as distributions is to assume that they are independently Gaussian distributed with zero mean and some arbitrary variance λ [13] as follows

$$p(\mathcal{W}) = \prod_{l=i}^{L} \prod_{i=0}^{V_l} \prod_{j=1}^{V_{l-1}+1} \mathcal{N}(w_{i,j}|0,\lambda^{-1}).$$
(2)

Here L denotes here the number of layers in the network and V_l the nodes in each layer l. The likelihoods with regard to the network weights are given by

$$p(y|\mathcal{W}, \mathcal{Z}) = \prod_{i=1}^{M} \mathcal{N}(y_i|d(z_i|\mathcal{W}), \gamma^{-1}),$$
(3)

where γ is the precision and \mathcal{Z} is the collection of nodes z_i . Now we can estimate the posterior distribution over the weights given the data via Bayes' rule

$$p(\mathcal{W}|\mathcal{D}) = \frac{p(y|\mathcal{W}, \mathcal{Z})p(\mathcal{W})}{p(y|\mathcal{Z})},$$
(4)



Fig. 1: a) Bayesian neural network with two features (inputs), one hidden layer, and one label (output). b) Single neuron.

where we can calculate the marginal likelihood via

$$p(y|\mathcal{Z}) = \int p(y|\mathcal{W}, \mathcal{Z}) \mathcal{W} d\mathcal{W}.$$
 (5)

In reality, the marginal likelihood and equation (5) cannot be calculated, it becomes intractable because of the nonlinearities in the network caused by the activation $d(\cdot, W)$. This results in the need to approximate the posterior distribution. This is a common challenge in the Bayesian inference domain, and different methods exist to approximate the needed posterior. The most popular are Markov chain Monte Carlo (MCMC), Laplace approximation (LA), probabilistic backpropagation (PBP), and variational inference (VI), for a survey see [18] and [19].

Nominal Model Predictive Control: As a comparison and baseline, we consider a nominal model predictive controller, which solves a finite-horizon optimal control problem at the sampling times [1], [2] based on the nominal model³:

$$\min_{\{u_k\}} J(\{x_k\}, \{u_k\})$$
(6a)

s.t.
$$x_{k+1} = f(x_k, u_k), \quad x_0 = \tilde{x}_j,$$
 (6b)

$$x_k \in \mathcal{X}, \quad u_k \in \mathcal{U},$$
 (6c)

where $J({x_k}, {u_k}) = \sum_{k=0}^{N} L(x_k, u_k) + E(x_k)$ is the cost function, N is the control horizon, $L: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$ is the stage cost, $E: \mathbb{R}^{n_x} \to \mathbb{R}$ is the terminal cost, \mathcal{X} and \mathcal{U} are the feasible input and state sets. The first input u_k^* of the obtained optimal input sequence is applied. The optimization problem is repeated at every sampling time updating x_k with the measured state \tilde{x}_i .

Stochastic Model Predictive Control: To exploit the stochastic uncertainty information of the model, we use stochastic MPC. Compared to the nominal case, we consider chance constraints, allowing for a violation of the constraints

³In principle the nominal MPC could also use the learned "nominal" model part, which would be the mean of the predicted state for case of the GP model. This is avoided here, due to space limitations.

with a certain probability [5]. Assuming a cost that contains the expected value of the now stochastic state and control variables minimizing a sequence of optimal policies $\Pi(x) =$ $\{\pi_0(x), ..., \pi_N(x)\}$ instead of the optimal open-loop inputs, one obtains the following stochastic optimal control problem

$$\min_{\Pi(x)} \mathbb{E}\left[J(\{x_k\}, \{u_k\})\right]$$
(7a)

s.t.
$$x_{k+1} = f_{\text{prob}}(x_k, u_k), \quad x_0 = \tilde{x}_j,$$
 (7b)

$$p(x_k \in \mathcal{X}) \ge p_x, \quad \forall k \in 0, \dots, N,$$
 (7c)

$$p(u_k \in \mathcal{U}) \ge p_u, \quad \forall k \in 0, \dots, N,$$
 (7d)

$$u_k = \pi \left(x_k \right). \tag{7e}$$

with f_{prob} defining the probabilistic model of the system, p_x and p_u are the probabilities with which the chance constraints are allowed to be violated and k describes the time steps of the discrete model.

Since the stochastic optimal control problem is generally intractable, we reformulate and approximate it to make it tractable and include the learning-based hybrid model. The problem's intractability results from the fact that we need to consider infinitely many state trajectories if we solve for an optimal sequence of policies Π .

The reformulations and approximations are based on the approximation presented in [9] and the theory of reachable sets presented in [8]. We omit the necessary assumptions for the sake of brevity and state the derived outcome: the system state, control input, and uncertainty dynamics, that is to be learned by a GP or a BNN, are jointly Gaussian distributed:

$$\mathcal{N}(\mu_k, \Sigma_k) = \mathcal{N}\left(\begin{bmatrix} \mu_k^x \\ \mu_k^u \\ \mu_k^d \end{bmatrix}, \begin{bmatrix} \Sigma_k^x \ \Sigma_k^{xu} \ \Sigma_k^{xd} \\ \Sigma_k^{ux} \ \Sigma_k^{u} \ \Sigma_k^{ud} \end{bmatrix} \right).$$
(8)

Furthermore, we assume that the chance constraints are tightening around the mean of our input and state variables μ^x and μ^u , which are defined by the assumed Gaussian distributions. With these assumptions, we can formulate of tractable stochastic optimal control problem of the following structure (as we only want to show the basic structure and due to limited space, we omit introducing all symbols)

$$\min_{\substack{\mu_k^u \\ \mu_k^u}} \quad \mathbb{E}\left[J(\{x_k\}, \{u_k\})\right] \tag{9a}$$

s.t.
$$\mu_{k+1}^{x} = \hat{f}(x_k, u_k, k) + B_d \mu_k^d$$
, (9b)

$$\Sigma_{k+1}^{x} = \left[\nabla \hat{f}\left(\mu_{k}^{x}, \mu_{k}^{u}\right) B_{d}\right] \Sigma \left[\nabla \hat{f}\left(\mu_{k}^{x}, \mu_{k}^{u}\right) B_{d}\right]^{\mathsf{T}}, \quad (9c)$$

$$\mu_0^x = \tilde{x}_j, \mu_k^d = \mu^d \left(\mu_k^x, \mu_k^u \right), \Sigma_0^x = 0,$$
(9d)

$$\Sigma_k = \text{according to eq. (8)},$$
 (9e)

$$\mu_k^x \in \mathcal{Z}(\Sigma_{k+1}^x), \, \mu_k^u \in \mathcal{V}(\Sigma_{k+1}^x), \, \forall k \in 0, \, \dots, \, N \quad (9f)$$

This reformulation and approximation allow fast solutions. It is implemented in HILO-MPC [16] allowing to use GP and BNN models. A more detailed discussion of the case of GPs can be found in [11].

TABLE I: Parameters used in the simulations [20].

parameters	values	parameters	values	parameters	values
K_d	0.0131	$\mu_{\rm max,con}$	0.9297	$\mu_{\max, mon}$	0.6275
Y	0.2116	B	0.4818	K_S	443.1
V	5	—	—	—	-

III. CASE STUDY: WASTEWATER TREATMENT PLANT

We want to compare the suitability of hybrid GP and BNN models for model-based control using stochastic MPC. To do so, a wastewater treatment plant is considered, as introduced in [20], which can be described by

$$\dot{X}(t) = \mu(X, S)X(t) - \frac{F(t)}{V}X(t) - K_dX(t) + w_X(t),$$

$$\dot{S}(t) = \frac{F(t)}{V}(S_f(t) - S(t)) - \frac{\mu(X, S)}{Y}X(t) + w_S(t).$$

Here, the states (X, S) denote the total biomass (mg/L) and the substrate concentrations (mg/L). The mapping $\mu(X, S)$ determines biomass growth rate, while K_d denotes its death rate. The in- and outflow rate is defined by F(t) (L/d) over the reactor volume V (L), and the substrate is fed with the concentration of $S_f(t)$ (mg/L). The dimensionless parameter Y denotes the substrate yield coefficient. Furthermore, normally distributed process noise is acting on each state, i.e., $w_X \sim \mathcal{N}(0, 1)$ and $w_S \sim \mathcal{N}(0, 1)$. The reactor is controlled by adjusting the flow rate directly, i.e., u(t) = F(t). The real specific growth rate μ is defined by the Contois equation $\mu_{\rm con}(X,S) = \frac{\mu_{\rm max,con}S}{BX+S}$, where $\mu_{\rm max,con}$ is the maximum growth rate and B is the kinetic saturation coefficient.

For the control model, the growth rate is assumed to be a Monod equation leading to a model-plant mismatch. $\mu_{\max, mon}S$ The Monod equation is given by $\mu_{
m mon}$ = $K_S + \overline{S}$, with the half-velocity constant K_S . Furthermore, we assume a mismatch in the process parameters K_S and $\mu_{\text{max,mon}}$ of +10% and -20%, respectively, from the nominal parameters. All parameters used in the true plant and the nominal controller model are given in Table I. The influent substrate concentration $S_{f}(t)$ is modeled as a time-varying parameter, as typically is the case in real world wastewater plants. For S_f , we assume it underlies a smooth and random but bounded disturbance $S_f(t) = 5500 + 100 \left(\sin \left(w_t t \right) + \sin \left(w_\pi \pi t \right) + \sin \left(w_e e t \right) \right),$ with e being Euler's number and $w_t \sim \mathcal{N}(0.3, 0.01) +$ $\mathcal{N}(-0.3, 0.01), w_{\pi} \sim \mathcal{N}(0.01, 0.01) + \mathcal{N}(-0.01, 0.01),$ $w_{\rm e} \sim \mathcal{N} \left(0.08, 0.01 \right) + \mathcal{N} \left(-0.08, 0.01 \right).$

All simulations were run on a MacBook Pro with an Apple M2 chip and 24 GB memory using the operating system macOS Ventura (version 13.2.1). CasADi [21], which is used internally by HILO-MPC, was installed with its most recent version (version 3.5.5). Additionally, the linear solver HSL_MA97 from the HSL package [22] was used.

Data Generation & Training: The true plant model and the nominal control model were used to generate the data for the training of both, the GP and the BNN. To do so, we created a nominal MPC design according to (6). The stage cost L and the terminal cost *E* are defined as $L(x_k, u_k) = ||x_k - x_{\text{ref}}||_{Q_s}^2 + ||u_k - u_{\text{ref}}||_{R_s}^2 + ||u_k - u_{k-1}||_{R_c}^2$, $E(x_k) = ||x_N - x_{\text{ref}}||_{Q_t}^2$, with the reference values $x_{\text{ref}} = [X_{\text{ref}} S_{\text{ref}}] = [x_{\text{ref}} S_{\text{ref}}]$ $\begin{bmatrix} 1046.28 & 101.615 \end{bmatrix}$, $u_{ref} = F_{ref} = 0.714286$, and the weights $Q_s = Q_t = \begin{bmatrix} 10 & 0 \\ 0 & 10 \end{bmatrix}$, $R_s = 1$, $R_c = 5 \cdot 10^3$. The reference of the flow rate F_{ref} translates to a hydraulic retention time of $\tau = 7 \,\mathrm{d}$, i.e., the average time a volume of wastewater will remain in a particular part of the plant, and lies within the range of hydraulic retention times referenced in [20]. The reference values of the biomass concentration $X_{\rm ref}$ and the substrate concentration $S_{\rm ref}$ are the steady states of the open loop simulation using the chosen reference flow rate F_{ref} . The initial conditions of the states were $X_0 = 0.2$ and $S_0 = 0$, and both states were constrained to be at least zero over the whole process time. Furthermore, the input was constrained to lie within the range $0 \le F(t) \le 2$. The sampling interval was $\Delta t = 0.125 \,\mathrm{d}$ and the length of the control horizon was N = 80, which translates to a time of 10 d. The influent substrate concentration was calculated using (III) and was kept constant over the control horizon.

We ran six closed loop simulation using this nominal MPC for a simulation time of $T_{\rm sim} = 70 \, \rm d$ for each simulation, which ensured that the steady state was reached in each simulation. Additionally, we sampled the reference states for each simulation from a uniform distribution around the actual reference states to gain a sufficient distribution of data in the state space region of interest $X_{\rm ref}$ ~ $\mathcal{U}(0.9X_{\text{ref}}, 1.1X_{\text{ref}}), \tilde{S}_{\text{ref}} \sim \mathcal{U}(0.9S_{\text{ref}}, 1.1S_{\text{ref}})$. Both states are assumed to be measured at all times and serve as the features of the machine learning models. The discrepancy between the true plant model and the nominal control model, the d part of (1)), is used as the label for the machine learning models. Overall, 2800 data points were assembled and split between training and test sets with the ratio 0.8 : 0.2. The data was scaled to zero-mean and standard deviation to improve training performance.

Gaussian Process: Since the training of a GP is computationally expensive, a sparser subset of training data was generated. This was done by iterating through all observations in the training data and disregarding all the observations that have an Euclidean distance to the current observation below a certain threshold. This way we were able to obtain more observations from regions that have been observed very little and disregard observations from well-observed regions in the state space. Since the data was normalized, the threshold was set to 0.2. The new training data set obtained using this threshold amounts to 92 data points.

As the model of the wastewater treatment plant has two states, we trained two GPs. The noise variance of each GP was fixed to $\sigma_X^2 = 0.1$ and $\sigma_S^2 = 0.01$, respectively. We chose a squared exponential kernel with automatic relevance detection as the covariance function for both GPs $k(\chi_i, \chi_j) = \sigma_f^2 \exp\left(-\frac{(\chi_i - \chi_j)^2}{2\ell^2}\right)$, where σ_f is the signal variance and ℓ are the length scales. The mean function was assumed to be zero for both GPs, as is typically done

TABLE II: Hyperparameters of the trained GPs.

	length scales ℓ	signal variance σ_f^2
GP_X	2.35425 2.91528	5.86936
GP_S	$\begin{bmatrix} 2.82233 & 2.44933 \end{bmatrix}$	10.493



Fig. 2: Heat map of the predictive standard deviation for the biomass X. The black dots are the training data points.

[23]. The hyperparameters of the trained GPs can be found in Table II. The length scales of both GPs are relatively close to each other, indicating that every input to the GPs is equally important. Fig. 2 shows the resulting predictive standard deviation for GP_X in the top row. Its model uncertainty is small in the region around the training data points. In regions where there are no observations available, it predicts a very high uncertainty. Fig. 3 shows the calibration curves of both GPs [24]. The calibration curve for GP_X is very close to the ideal calibration line, indicating reliable predictions of the model-plant mismatch d_X for the biomass concentration. This can also be shown by calculating the miscalibration area, i.e., the area between the curve and the ideal calibration line. The closer the calibration curve is to the perfect calibration line, the lower the miscalibration area, resulting in higher prediction reliability. The resulting miscalibration area



Fig. 3: Calibration curves, the ideal calibration is shown red.

TABLE III: Root-mean-square error (RMSE), negative log likelihood (NLL) and micalibration area of the trained GPs and BNNs. Bold indicates the lowest value for each metric.

	RMSE	NLL	Miscalibration area
GP_X	0.173	-0.277	0.117
BNN_X	0.165	0.478	0.107
GP_S	0.205	0.148	0.135
BNN_S	0.264	0.043	0.101

is listed in Table III. The calibration curve for the GP trained on the model-plant mismatch in the substrate concentration (GP_S) is a further away from the ideal calibration line, leading to a higher miscalibration area (see Table III). This indicates a lower prediction reliability of GP_S compared to GP_X . Two additional metrics, the root-mean-square error (RMSE) as well as the negative log-likelihood (NLL), are listed in Table III. The RMSE reflects the correctness of the predictions, and the NLL is another indicator of the predictions' reliability. Both metrics respectively show that GP_X has higher accuracy and reliability in the forecasts compared to GP_S , as both values are lower for GP_X .

1) Bayesian Neural Network: In this work, we used probabilistic backpropagation [25] to approximate the posterior distribution (4) of the BNN, which is a closed-form approximation based on the assumed density filtering approach [26]. Its full derivation is beyond the scope of this work, so we refer the reader to [25] for details. In contrast to our work, [14] and [15] employ variational inference to approximate the posterior distribution.

In probabilistic backpropagation, the prior precision λ and the noise precision γ are assumed to be Gamma distributed hyperpriors $p(i|\alpha_i, \beta_i) \sim \Gamma(\alpha_i, \beta_i)$ where α_i are the shape parameters and β_i are the inverse scale parameters. Here, we specified the weight hyperpriors with shape $\alpha_{\lambda} = 6$ and inverse scale $\beta_{\lambda} = 6$ for both outputs. The noise hyperpriors were $\alpha_{\gamma,X} = \beta_{\gamma,X} = 40$ and $\alpha_{\gamma,S} = \beta_{\gamma,S} = 6$, respectively. As with the GPs, the BNNs also need to be trained for each output individually. Both BNNs had one hidden layer with 50 nodes and the ReLU activation function [25]. The BNNs were trained for 10 epochs each.

Fig. 2 shows the resulting predictive distribution for BNN_X in the bottom row. Similar to GP_X , it predicts a very low uncertainty in the regions where the training data are located. Unlike GP_X , BNN_X centers its posterior on the best observed region, shows higher uncertainty in regions with fewer observations, but lower uncertainties in regions without observations. The calibration curves of both BNNs are also very close to the ideal calibration line, indicating high reliability in the predictions. Overall, BNNs and GPs show competitive performance, which is also reflected when comparing the metrics (RMSE, NLL, miscalibration area ---see Table III).

Open Loop Simulation of Hybrid Models: Given the trained GPs and BNNs, we generated the respective hybrid models. To illustrate their recovery of the model-plant mismatch, we compared open loop step response versus the true plant model and the nominal control model without the process noise. The result for a the step response is shown



in Fig. 4. Both hybrid models significantly improve the nominal control model, especially substrate concentration. The performance improvement compared to the nominal control model is substantial, while the mismatch in the steady states of the biomass concentration of a true model is negligible if process noise were considered.

Closed Loop Control: Finally, learning-supported stochastic MPC for both learning approaches is evaluated. Box constraints of the nominal MPC augmented with a time-varying constraint that ensures survival of the biomass despite disturbed substrate concentration S_f . This constraint is realistic for practical applications that assume model-plant mismatches, since drastic drops in S_f could wash out the biomass and drastic increases of S_f could lead to unmodeled behavior, e.g., cell aggregation. Hence, we demand the biomass concentration to be within once the steady state is reached $X(t) \in [(X_{ref} - 20)b_l, (X_{ref} + 20)b_u]$. For t = 30 dafter the steady state is reached, the parameters b_l and b_u are set to $b_l(t) = \tanh(0.1t + 0.01), b_u(t) = \tanh(0.1t + 1)$. For the stochastic MPC we reformulated the constraint into a chance constraint with an acceptance probability of 99%. In order to ease the computational burden, the control horizon was shortened to N = 32 time steps.



Fig. 5: Closed loop simulations.

Fig. 5 shows the closed loop behavior of the control designs. The difference between the approaches is mainly

present in the vicinity of the constraints. Disturbances in s_f lead to learning-supported MPC violating lower constraint, showing its inability to guarantee constraint satisfaction. In contrast, the learning-supported stochastic MPC designs satisfy the constraints.



Fig. 6: Comparison of the total computation time.

Fig. 6 compares the computation time for both stochastic MPC controllers. Multiple GPs were trained and run in the closed loop control setup, varying the number of training points. One can see the increase in total computation time of the corresponding stochastic MPC, and around 130 training data points it surpasses the computation time of the stochastic MPC using the BNN.

IV. CONCLUSION AND OUTLOOK

State-of-the-art learning-supported stochastic MPC requires uncertainty prediction, often using Gaussian processes. However, their computational complexity increases with data set size. We explored Bayesian neural networks (BNNs) for hybrid models in learning-supported stochastic MPC, comparing their performance to Gaussian processes. BNNs achieved similar performance, minimized model-plant mismatch in open-loop simulations, and proved effective in closed-loop simulations. BNNs offer a valuable alternative, efficiently handling large data sets. All results used the opensource Python toolbox HILO-MPC [16].

Future research will investigate BNNs for more complex models and provide strict stability and performance guarantees for specific model classes.

ACKNOWLEDGMENT

The authors acknowledge funding of the DIGIPOL project (Magdeburg, Saxony-Anhalt) funded in the EU-ERDF scheme and the KI-Embedded project funded by the BMWK.

REFERENCES

- J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed. Nob Hill Publishing, 2017.
- [2] R. Findeisen and F. Allgöwer, "An Introduction to Nonlinear Model Predictive Control," in 21st Benelux Meeting on Systems and Control, 2002, pp. 119–141.
- [3] M. G. Forbes, R. S. Patwardhan, H. Hamadah, and R. B. Gopaluni, "Model Predictive Control in Industry: Challenges and Opportunities," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 531–538, 2015, 9th IFAC Symp. on Advanced Control of Chemical Processes (ADCHEM).

- [4] R. Findeisen, T. Raff, and F. Allgöwer, "Sampled-Data Nonlinear Model Predictive Control for Constrained Continuous Time Systems," in Advanced Strategies in Control Systems with Input and Output Constraints, S. Tarbouriech, G. Garcia, and A. H. Glattfelder, Eds. Springer, 2007, pp. 207–235.
- [5] A. Mesbah, "Stochastic Model Predictive Control: An Overview and Perspectives for Future Research," *IEEE Cont. Syst. Mag.*, vol. 36, no. 6, pp. 30–44, 2016.
- [6] A. Mesbah, S. Streif, R. Findeisen, and R. D. Braatz, "Stochastic nonlinear model predictive control with probabilistic constraints," in *Amer. Cont. Conf. (ACC)*, 2014, pp. 2413–2419.
- [7] E. Bradford, L. Imsland, M. Reble, and E. A. del Rio-Chanona, "Hybrid Gaussian Process Modeling Applied to Economic Stochastic Model Predictive Control of Batch Processes," in *Recent Advances* in *Model Predictive Control: Theory, Algorithms, and Applications*, T. Faulwasser, M. A. Müller, and K. Worthmann, Eds. Springer, 2021, pp. 191–218.
- [8] L. Hewing and M. N. Zeilinger, "Stochastic Model Predictive Control for Linear Systems Using Probabilistic Reachable Sets," in *IEEE Conf. Dec. Cont. (CDC)*. IEEE, 2018, pp. 5182–5188.
- [9] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious Model Predictive Control Using Gaussian Process Regression," *IEEE Trans. Cont. Syst. Tech.*, vol. 28, no. 6, pp. 2736–2743, 2020.
- [10] M. Maiworm, D. Limon, J. M. Manzano, and R. Findeisen, "Stability of Gaussian Process Learning Based Output Feedback Model Predictive Control," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 455–461, 2018, 6th IFAC Conf. Nonlinear Model Predictive Control (NMPC).
- [11] B. Morabito, J. Pohlodek, L. Kranert, S. E. Ríos, and R. Findeisen, "Efficient and Simple Gaussian Process Supported Stochastic Model Predictive Control for Bioreactors using HILO-MPC," *IFAC-PapersOnLine*, vol. 55, no. 7, pp. 922–927, 2022.
- [12] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, ser. Adaptive Computation and Machine Learning. MIT Press, 2006.
- [13] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, "Hands-On Bayesian Neural Networks — A Tutorial for Deep Learning Users," *IEEE Comp. Int. Mag.*, vol. 17, no. 2, pp. 29–48, 2022.
- [14] F. Cursi, V. Modugno, L. Lanari, G. Oriolo, and P. Kormushev, "Bayesian Neural Network Modeling and Hierarchical MPC for a Tendon-driven Surgical Robot with Uncertainty Minimization," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2642–2649, 2021.
- [15] Y. Bao, K. J. Chan, A. Mesbah, and J. M. Velni, "Learning-based Adaptive-Scenario-Tree Model Predictive Control with Probabilistic Safety Guarantees Using Bayesian Neural Networks," in *Amer. Cont. Conf. (ACC).* IEEE, 2022, pp. 3260–3265.
- [16] J. Pohlodek, B. Morabito, C. Schlauch, P. Zometa, and R. Findeisen, "Flexible development and evaluation of machine-learning-supported optimal control and estimation methods via HILO-MPC," 2022. [Online]. Available: https://arxiv.org/abs/2203.13671
- [17] R. M. Neal, Bayesian Learning for Neural Networks, ser. Lecture Notes in Statistics. Springer, 1996.
- [18] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, and S. Nahavandi, "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," *Information Fusion*, vol. 76, pp. 243–297, 2021.
- [19] E. Daxberger, A. Kristiadi, A. Immer, R. Eschenhagen, M. Bauer, and P. Hennig, "Laplace Redux — Effortless Bayesian Deep Learning," in *Proce. Int. Conf. Neural Inf. Proc. Syst.* NIPS, 2021.
- [20] W. C. Hu, K. Thayanithy, and C. F. Forster, "A kinetic study of the anaerobic digestion of ice-cream wastewater," *Process Biochemistry*, vol. 37, no. 9, pp. 965–971, 2002.
- [21] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – a software framework for nonlinear optimization and optimal control," *Mathe. Prog. Comp.*, vol. 11, no. 1, pp. 1–36, 2019.
- [22] "HSL. A collection of Fortran codes for large scale scientific computation," https://www.hsl.rl.ac.uk.
- [23] J. Kocijan, R. Murray-Smith, C. E. Rasmussen, and B. Likar, "Predictive control with Gaussian process models," in *EUROCON 2003*, 2003, pp. 352–356.
- [24] V. Kuleshov, N. Fenner, and S. Ermon, "Accurate Uncertainties for Deep Learning Using Calibrated Regression," in *Proceedings of the* 35th International Conference on Machine Learning, vol. 80. PMLR, 2018, pp. 2796–2804.

- [25] J. M. Hernández-Lobato and R. P. Adams, "Probabilistic Backpropagation for Scalable Learning of Bayesian Neural Networks," in *Proc. 32nd Int. Conf. Mach. Lear.*, vol. 37. PMLR, 2015, pp. 1861–1869.
 [26] T. P. Minka, "A Family of Algorithms for Approximate Bayesian Inference," Ph.D. dissertation, Massachusetts Institute of Technology, 2001.
- 2001.