Safe Active Learning of Cerebrospinal Fluid Dynamics

Fabian Flürenbrock
ETH Zurich
ffluerenb@ethz.ch

Johannes Köhler
ETH Zurich
jkoehle@ethz.ch

Marianne Schmid Daners
ETH Zurich
marischm@ethz.ch

Melanie Zeilinger
ETH Zurich
mzeilinger@ethz.ch

Abstract

This paper presents a safe active learning framework for a clinically relevant class of nonlinear systems with time-varying and uncertain parameters. The framework aims to provide a systematic trade-off between three competing clinical objectives: regulation of physiological variables to a safe zone, learning of patient-specific parameters, and minimization of the medical intervention. To address these challenges, we integrate the covariance propagation of a Kalman filter used for patient parameter estimation into an optimization-based control algorithm and enforce a desired estimation accuracy by introducing a soft constraint on the predicted covariances. We demonstrate the potential of the safe active learning framework for healthcare applications in a case study on cerebrospinal fluid dynamics. Our proposed method improves patient monitoring and shunt therapy for the neurological condition hydrocephalus by doubling the parameter estimation accuracy while requiring less than half the rate of intervention compared to standard approaches.

1 Introduction

Healthcare increasingly relies on time series data to guide both diagnosis and therapy of disease. Physiological signals, however, are noisy, highly patient-specific, and subject to changing dynamics over time. To deliver safe and effective interventions, medical systems such as artificial pancreas or anesthesia machines must balance three competing objectives: (i) regulating physiological variables within a predefined zone to ensure patient safety and therapeutic efficacy, (ii) accurately learning uncertain and time-varying patient parameters to support both closed-loop control and clinical decision-making, and (iii) achieving these goals with minimal intervention to reduce discomfort and adverse effects [20]. This introduces a fundamental trade-off between regulation, exploration, and actuation. We aim to address this trade-off by developing a safe active learning framework based on model predictive control (MPC), which regulates a dynamic system to a safe zone while minimizing input variations necessary to achieve and maintain a desired parameter estimation accuracy.

MPC is an advanced control strategy for constrained nonlinear systems that computes control inputs by predicting and optimizing the system's future trajectory based on a dynamical model [32]. While traditional MPC is typically designed to regulate the system to a specific set point or track a reference trajectory [24], for many applications a specific preferred operating point or reference trajectory is not known. In such cases, zone MPC can be applied to instead steer the system toward a predefined zone where all states are considered equally valuable [14]. This formulation enables aggressive control actions when the system is outside the target zone but reduces control interventions once inside, making the controller less sensitive to disturbances and noise. Zone MPC has been proven successful in healthcare applications such as chronic kidney disease [29] and diabetes [15, 23].

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Learning from Time Series for Health.

Learning accurate dynamical system models from time series requires sufficiently informative data to be contained in the observed signals. Optimal experiment design provides a principled approach to design input trajectories that collect informative data for learning system dynamics [12, 31] and has been widely applied in practice, e.g., in robotics [37] and biomedical engineering [7]. Typically, system dynamics are learned offline prior to controller design and deployment. In healthcare, however, patient dynamics often change over time and therefore require closed-loop parameter adaptation. As time series data collected during closed-loop operation is often not informative enough to precisely learn the underlying parameters, the system needs to be actively excited to reduce uncertainty [17, 27].

In this work, we integrate optimal experiment design into a zone MPC scheme to develop a safe active learning framework that systematically balances the trade-off between regulation, exploration, and actuation. Specifically, we integrate the covariance propagation equations of a Kalman filter used for parameter estimation into a zone MPC scheme, and control the accuracy of the patient's parameter estimates to a desired level by introducing a suitable soft constraint on the predicted covariances. This formulation enables a structured prioritization of competing objectives: the system shall first be regulated to the predefined safe zone, then actively excited within the zone until the desired parameter accuracy is reached, and finally, the least invasive control input necessary to sustain this accuracy should be applied. To highlight the potential of the proposed safe active learning framework for healthcare applications, we conduct a simulation-based case study on cerebrospinal fluid (CSF) dynamics and demonstrate how the framework can be used to improve patient monitoring and shunt therapy for the neurological condition hydrocephalus [9]. An extended version of this paper, with an additional application of the proposed method to CSF infusion studies used for the diagnosis of hydrocephalus, can be found in [8].

2 Related work

A substantial body of work on safe active learning for dynamical systems exists in both reinforcement learning (RL) [35] and MPC. In RL, active learning is often achieved through random excitation strategies that are undesirable for applications with humans in the loop, though there is an increasing amount of work that explicitly incorporate safety constraints into RL [16]. While these safe RL methods have shown strong empirical results, the resulting models and decisions are generally given by black box neural network functions, which makes them difficult to interpret and trust for clinicians. In contrast, the physiological states and parameter estimates from clinically established patient models can support both patient monitoring and clinical decision making processes [21]. This motivates the use of MPC schemes grounded in physiological models, which provide both safety guarantees and clinically meaningful insights. However, most existing active learning MPC approaches excite the system independently of the current parameter uncertainty, potentially leading to unnecessary or undesired interventions. See also Appendix D for a more detailed discussion.

3 Methods

We consider the following class of nonlinear discrete-time dynamical systems, which can represent CSF dynamics for hydrocephalus (see Appendix E) as well as other physiological processes,

$$x(k+1) = \Phi(x(k), u(k))\theta(k) + w(k), \tag{1}$$

with state $x(k) \in \mathbb{R}^{n_x}$, input $u(k) \in \mathbb{R}^{n_u}$, known non-linear features $\Phi(x(k), u(k)) \in \mathbb{R}^{n_x \times n_\theta}$, unknown time-varying parameters $\theta(k) \in \mathbb{R}^{n_\theta}$, and additive noise $w(k) \in \mathbb{R}^{n_x}$. For the state $x(k)$, we consider the desired zone $\mathbb{X} = \{x \in \mathbb{R}^{n_x} \mid x_{\text{min}} \leq x \leq x_{\text{max}}\}$ with element-wise inequalities and $x_{\text{min}},\ x_{\text{max}} \in \mathbb{R}^{n_x}$. We further consider hard constraints on the input $u(k) \in \mathbb{U} \subset \mathbb{R}^{n_u}$ and the input rate $\Delta u(k) = u(k) - u(k-1) \in \Delta\mathbb{U} \subset \mathbb{R}^{n_u}$. The evolution of the unknown time-varying system parameters is modeled as a random walk

$$\theta(k+1) = \theta(k) + v(k), \tag{2}$$

where $v(k) \in \mathbb{R}^{n_\theta}$ is a noise term that represents the parameter drift. We assume that $w(k)$ and $v(k)$ are i.i.d. Gaussian random variables with $w \sim \mathcal{N}(0,\, \Sigma_w)$ and $v \sim \mathcal{N}(0,\, \Sigma_v)$, respectively. Additionally, we assume that a prior distribution of the parameter $\theta_0 \sim \mathcal{N}(\hat{\theta}_0, \Sigma_0)$ is given and that noise-free measurements of the state $x(k)$ are available.

The goal of this work is to achieve the following list of ranked objectives:

1. The state $x(k)$ is regulated to the desired zone $\mathbb{X}$.
2. A desired estimation accuracy for the system parameters $\theta(k)$ is achieved and maintained.
3. The input change $\|\Delta u(k)\|$ is minimized.

To achieve the outlined objectives, we design a safe active learning framework by integrating Kalman filtering, optimal experiment design, and zone MPC. This is done in two steps. First, the covariance propagation equations of the Kalman filter (see Appendix B) are integrated into the optimization problem of a zone MPC formulation (see Appendix C). This enables the MPC to predict the future covariance matrices $\Sigma_{i=0:N-1}$ of the patient parameters and directly assess the effect of the inputs $u_{i=0:N-1}$ on the uncertainty of the future parameter estimates, similar to [18, 3]. Second, to enforce a desired parameter estimation accuracy, a desired covariance matrix $\bar{\Sigma} \succ 0$ is introduced as an upper bound. A computationally tractable comparison between the predicted covariance and the upper bound is implemented using a function $\xi(\Sigma_i)$ that maps the covariance matrix to a scalar value. In optimal experiment design, common choices include the trace (A-optimal) and the determinant (D-optimal) [12, 31]. The desired parameter accuracy is encoded with a soft constraint $\xi(\Sigma_i) \leq \xi(\bar{\Sigma}) + \eta_i$ to ensure feasibility, and the slack variable $\eta_i \geq 0$ is minimized with a user-defined scaling $s > 0$ to incentivize active learning until the desired parameter estimation accuracy is achieved. The resulting active learning zone MPC is described by the following nonlinear program. At each time step k, given the previous input $u(k-1)$, the measured state $x(k)$, the estimated parameter mean $\hat{\theta}(k)$ and covariance $\Sigma(k)$, the following optimization problem is solved

$$
\begin{aligned}
\min_{\substack{x,u,\epsilon,\eta \\ S,K,\Sigma}} \quad & \sum_{i=0}^{N-1} \|Q\epsilon_i\|_1 + \|\Delta u_i\|_R^2 + s\eta_i && \text{(3a)} \\
\text{s.t.} \quad & \forall\, i = 0, \ldots, N-1, \\
& \Phi_i = \Phi(x_i, u_i), && \text{(3b)} \\
& S_i = \Phi_i(\Sigma_i + \Sigma_v)\Phi_i^\top + \Sigma_w, && \text{(3c)} \\
& K_i = (\Sigma_i + \Sigma_v)\Phi_i S_i^{-1}, && \text{(3d)} \\
& \Sigma_{i+1} = (\mathbb{I} - K_i\Phi_i)(\Sigma_i + \Sigma_v), && \text{(3e)} \\
& x_{i+1} = \Phi_i\hat{\theta}(k), && \text{(3f)} \\
& \Delta u_i = u_i - u_{i-1}, && \text{(3g)} \\
& (u_i, \Delta u_i) \in \mathbb{U} \times \Delta\mathbb{U}, && \text{(3h)} \\
& x_{\min} - \epsilon_i \leq x_i \leq x_{\max} + \epsilon_i, && \text{(3i)} \\
& \xi(\Sigma_i) \leq \xi(\bar{\Sigma}) + \eta_i, && \text{(3j)} \\
& \epsilon_i \geq 0,\ \eta_i \geq 0, && \text{(3k)} \\
& x_0 = x(k),\ u_{-1} = u(k-1),\ \Sigma_0 = \Sigma(k), && \text{(3l)}
\end{aligned}
$$

where N is the prediction horizon, the constants $Q \succ 0$, $R \succ 0$, and $s > 0$ are user-chosen weights prioritizing the different objectives of the MPC formulation, and the variables S and K are the innovation covariance and Kalman gain of the Kalman filter, respectively.

4 Experiments and results

Hydrocephalus is a neurological condition in which disturbed CSF dynamics lead to physical and mental impairment [9]. While the assessment of CSF dynamics is typically performed using infusion studies [5], therapy is based on shunt systems that drain excessive CSF out of the cerebral ventricular system [1]. In this simulation study, we demonstrate how smart shunt systems [6] capable of CSF drainage control can achieve both objectives simultaneously by utilizing the proposed safe active learning framework. To this end, we perform a $T = 60\,\text{min}$ long shunt control simulation with sampling time $\Delta_t = 1\,\text{s}$, starting from the initial state $x(0) = 13$. The initial patient parameters are uniformly sampled from a physiological range (see Appendix E) and a random walk with parameter drift covariance $\Sigma_v = \mathbb{I} \cdot 10^{-6}$ is performed. For the active learning zone MPC (3), the prediction horizon is chosen as $N = 10$ and the desired estimation accuracy is defined as $\xi(\bar{\Sigma}) = \text{tr}(\bar{\Sigma}) = 10^{-3}$.

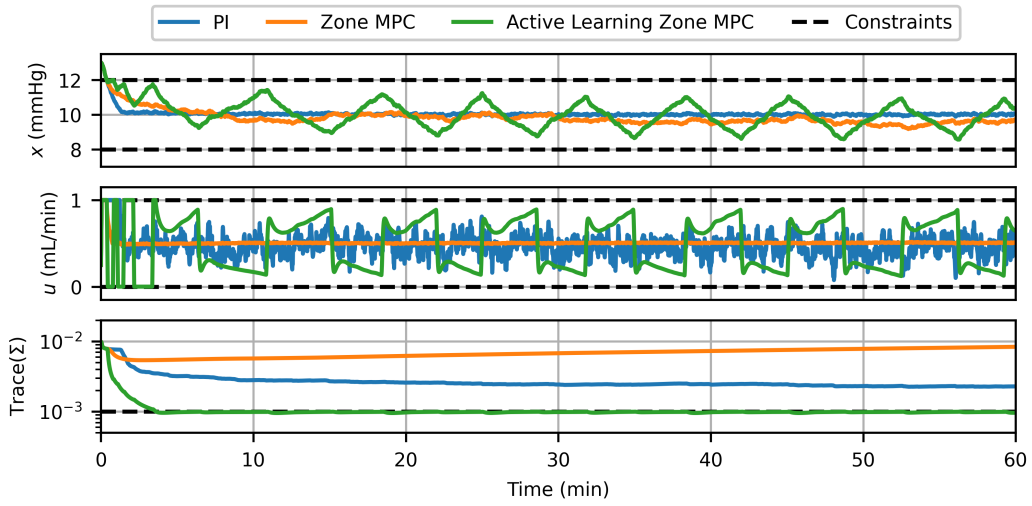

Figure 1: Closed-loop trajectories comparing proportional-integral (PI) control, standard zone model predictive control (MPC), and the proposed active learning zone MPC. Panels show cerebrospinal fluid (CSF) pressure (top), CSF drainage rate (middle), and parameter estimation uncertainty (bottom). Dashed lines indicate constraints on the state, input, and parameter estimates.

Table 1: Results of the hydrocephalus case study

Approach	$t_{x \notin \mathbb{X}}$	Δu_{rms}	$\hat{\theta}_{\text{rmse}}$
PI	24 s	0.0600	0.0335
Zone MPC	24 s	0.0079	0.0537
Active Learning Zone MPC	25 s	0.0267	0.0163

The scalar weights of the cost function are chosen as $Q > s > R$, prioritizing zone regulation first, estimation accuracy second, and input rate minimization last. The desired state zone is defined as $\mathbb{X} = [8, 12]$, covering a range of physiological pressure values [26]. As for medical safety reasons shunt systems must avoid CSF backflow, the input and input rate constraints are defined as $\mathbb{U} = [0, 1]$ and $\Delta\mathbb{U} = [-0.25, 0.25]$, respectively. Two methods are used for comparison. First, a standard zone MPC with the same constants and constraints as the active learning zone MPC. Second, a proportional-integral (PI) controller, which has been the predominant control approach for automated CSF drainage systems [6, 22]. The PI controller is designed to regulate the state to the center of the desired state zone and the control inputs are saturated to enforce the input and input-rate constraints. Both MPC problems are implemented in Python using CasADi [2] and solved with IPOPT [36]. Mean computation times for the zone MPC and active learning zone MPC are $4.4\,\text{ms}$ and $34.2\,\text{ms}$, respectively, which is well below the sampling time $\Delta_t = 1\,\text{s}$.

The closed-loop trajectories of the simulation study are sown in Figure 1. All approaches apply maximum input first to regulate the state to the zone or reference. Once inside the desired zone, the active learning zone MPC switches between maximum and minimum input to excite the CSF dynamics. After around $4\,\text{min}$, the desired estimation accuracy is achieved and the rate of input change reduced to maintain this accuracy. While the standard zone MPC quickly converges to a stable drainage rate after entering the desired zone, the PI controller conducts frequent input changes to counteract adverse effects of the noise on the set-point regulation. These results are quantitatively reflected in Table 1. The root mean square (RMS) value for the input change is the lowest for the standard zone MPC and the highest for the PI controller. The best parameter estimation performance in terms of the root mean square error (RMSE) is achieved by the proposed active learning zone MPC. In comparison to the predominantly used PI controller, the active learning zone MPC achieves more than double the estimation accuracy with less than half the input variation. The time outside the desired safe zone $t_{x \notin \mathbb{X}}$ is negligible for all approaches.

5 Conclusion and future work

This paper introduces a safe active learning framework that combines zone MPC, optimal experiment design, and Kalman filtering. The proposed method provides a structured trade-off between the three competing clinical objectives of physiological state regulation, accurate patient parameter estimation, and minimal medical intervention. Its potential for healthcare applications was demonstrated in a case study on hydrocephalus shunt therapy, where minimal actuation effort was used to regulate CSF pressure to a safe zone while continuously generating informative time series data to learn accurate patient parameters. While the proposed method assumes noise-free state measurements, this assumption is restrictive in clinical practice. Integrating joint state and parameter estimation for CSF dynamics [7] into the framework remains an important open problem to facilitate clinical application. Future work may also investigate the proposed methods for different model descriptions of the CSF dynamics [30] and their deployment on hardware for in vitro testing [10].

Acknowledgments

This work was funded by the NCCR Automation (51NF40 180545) of the Swiss National Science Foundation and the Janina Hug memorial fund.

References

[1] Kamran Aghayev, Sheikh M. A. Iqbal, Waseem Asghar, Bunyad Shahmurzada, and Frank D. Vrionis. Advances in CSF shunt devices and their assessment for the treatment of hydrocephalus. *Expert review of medical devices*, 2021.

[2] Joel A.E. Andersson, Joris Gillis, Greg Horn, James B. Rawlings, and Moritz Diehl. CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 2019.

[3] Angelo D. Bonzanini, Ali Mesbah, and Stefano Di Cairano. Perception-aware model predictive control for constrained control in unknown environments. *Automatica*, 2024.

[4] Jeremi Chabros, Michal M. Placek, Ka Hing Chu, Erta Beqiri, Peter J. Hutchinson, Zofia Czosnyka, Marek Czosnyka, Alexis Joannides, and Peter Smielewski. Embracing uncertainty in cerebrospinal fluid dynamics: A Bayesian approach to analysing infusion studies. *Brain and Spine*, 2024.

[5] Anders Eklund, Peter Smielewski, Iain Chambers, Noam Alperin, Jan Malm, Marek Czosnyka, and Anthony Marmarou. Assessment of cerebrospinal fluid outflow resistance. *Medical & Biological Engineering & Computing*, 2007.

[6] Fabian Flürenbrock, Leonie Korn, Dominik Schulte, Anthony Podgoršak, Joris Chomarat, Janina Hug, Tiago Hungerland, Caroline Holzer, David Iselin, Luca Krebs, Rosina Weiss, Markus F. Oertel, Lennart Stieglitz, Miriam Weisskopf, Mirko Meboldt, Melanie N. Zeilinger, and Marianne Schmid Daners. VIEshunt: towards a ventricular intelligent and electromechanical shunt for hydrocephalus therapy. *Fluids and Barriers of the CNS*, 2025.

[7] Fabian Flürenbrock, Simon Muntwiler, Leonie Korn, Marianne Schmid Daners, and Melanie N. Zeilinger. Model-based Estimation of Ventricular Cerebrospinal Fluid Volume. *Conference on Control Technology and Applications (CCTA)*, 2024.

[8] Fabian Flürenbrock, Johannes Köhler, Marianne Schmid Daners, and Melanie N. Zeilinger. Zone Model Predictive Control with Active Learning and Applications to Cerebrospinal Fluid Dynamics. *Conference on Decision and Control (CDC)*, 2025.

[9] Gary L. Gallia, Daniele Rigamonti, and Michael A. Williams. The diagnosis and treatment of idiopathic normal pressure hydrocephalus. *Nature Clinical Practice Neurology*, 2006.

[10] Manuel Gehlen, Vartan Kurtcuoglu, and Marianne Schmid Daners. Patient specific hardware-in-the-loop testing of cerebrospinal fluid shunt systems. *IEEE Transactions on Biomedical Engineering*, 2016.

[11] Hasmet Genceli and Michael Nikolaou. New Approach to Constrained Predictive Control with Simultaneous Model Identification. *AIChE Journal*, 1996.

[12] Michel Gevers and Lennart Ljung. Optimal Experiment Designs with Respect to the Intended Model Application. *Automatica*, 1986.

[13] A. H. González, A. Ferramosca, G. A. Bustos, J. L. Marchetti, M. Fiacchini, and D. Odloak. Model predictive control suitable for closed-loop re-identification. *Systems and Control Letters*, 2014.

[14] Alejandro H. González and Darci Odloak. A stable MPC with zone control. *Journal of Process Control*, 2009.

[15] Benyamin Grosman, Eyal Dassau, Howard C. Zisser, Lois Jovanovič, and Francis J. Doyle III. Zone Model Predictive Control: A Strategy to Minimize Hyper-and Hypoglycemic Event. *Journal of Diabetes Science and Technology*, 2010.

[16] Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, and Alois Knoll. A review of safe reinforcement learning: Methods, theories, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

[17] Ivar Gustavsson, Lennart Ljung, and Torsten Söderström. Identification of Processes in Closed Loop - Identifiability and Accuracy Aspects. *Automatica*, 1977.

[18] Tor A. N. Heirung, Birger Erik Ydstie, and Bjarne Foss. Dual adaptive model predictive control. *Automatica*, 2017.

[19] Rudolf E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 1960.

[20] Mohammad Javad Khodaei, Nicholas Candelino, Amin Mehrvarz, and Nader Jalili. Physiological Closed-Loop Control (PCLC) Systems: Review of a Modern Frontier in Automation. *IEEE Access*, 2020.

[21] Dong-Joo Kim, Hakseung Kim, Young-Tak Kim, Byung C. Yoon, Zofia Czosnyka, Kun-Woo Park, and Marek Czosnyka. Thresholds of resistance to csf outflow in predicting shunt responsiveness. *Neurological Research*, 2015.

[22] Inga Krause, Sebastian Hahne, Marian Walter, Sabine Linke, Klaus Radermacher, Sebastian Antes, Michael Kiefer, Regina Eymann, Wolf-Ingo Steudel, and Steffen Leonhardt. Eine neue automatisierte externe Liquor-Drainage. *Automation in Medical Engineering*, 2010.

[23] Justin J. Lee, Ravi Gondhalekar, and Francis J. Doyle. Design of an artificial pancreas using zone model predictive control with a Moving Horizon State Estimator. *Conference on Decision and Control (CDC)*, 2014.

[24] Daniel Limon, Ignacio Alvarado, Teodoro Alamo, and Eduardo F. Camacho. MPC for tracking piecewise constant references for constrained linear systems. *Automatica*, 2008.

[25] Xiaonan Lu and Mark Cannon. Robust adaptive model predictive control with persistent excitation conditions. *Automatica*, 2023.

[26] Jan Malm, Bo Kristensen, Thomas Karlsson, Markku Fagerlund, and Jan Ekstedt. The Predictive Value of Cerebrospinal Fluid Dynamic Tests in Patients With the Idiopathic Adult Hydrocephalus Syndrome. *Archives of Neurology*, 1995.

[27] Ian R. Manchester, Kenneth S. Andersson, Anders Eklund, and Anton S. Shiriaev. Identifiability of the parameters of a nonlinear model of the cerebrospinal fluid system. *IFAC Proceedings*, 2007.

[28] Anthony Marmarou, Kenneth Shulman, and Roberto M. Rosende. A nonlinear analysis of the cerebrospinal fluid system intracranial pressure dynamics and. *Journal of Neurosurgery*, 1978.

[29] Jayson McAllister, Zukui Li, Jinfeng Liu, and Ulrich Simonsmeier. Erythropoietin Dose Optimization for Anemia in Chronic Kidney Disease Using Recursive Zone Model Predictive Control. *IEEE Transactions on Control Systems Technology*, 2019.

[30] Jonas Ohnemus, Fabian Flürenbrock, Anthony Podgoršak, Melanie N. Zeilinger, and Marianne Schmid Daners. CSFsim: A Simulation Framework for Cerebrospinal Fluid Dynamics and Hydrocephalus Shunt Systems. *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2025.

[31] Luc Pronzato. Optimal experimental design and some related control problems. *Automatica*, 2008.

[32] James B. Rawlings, David Q. Mayne, and Moritz M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2 edition, 2017.

[33] Simo Särkkä. *Bayesian filtering and smoothing*. Cambridge University Press, 2013.

[34] Oscar A. Z. Sotomayor, Darci Odloak, and Lincoln F. L. Moro. Closed-loop model re-identification of processes under MPC with zone control. *Control Engineering Practice*, 2009.

[35] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 2nd edition, 2018.

[36] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 2006.

[37] Andrew D. Wilson, Jarvis A. Schultz, and Todd D. Murphey. Trajectory synthesis for fisher information maximization. *IEEE Transactions on Robotics*, 2014.

[38] Eva Žáčeková, Samuel Prívara, Zdeněk Váňa, Jiří Cigler, and Lukáš Ferkl. Dual control approach for zone model predictive control. *European Control Conference (ECC)*, 2013.

A Notation

For a vector $x \in \mathbb{R}^n$ and a positive semi-definite matrix $Q \in \mathbb{R}^{n\times n}$, we abbreviate $\|x\|_Q^2 = x^\top Q x$. We write the p-norm with $p \in \{1, 2, \infty\}$ of the vector x as $\|x\|_p$ and refer to the n-th element of x as $x_{[n]}$. For a square matrix $A \in \mathbb{R}^{n\times n}$, the trace and determinant of A are denoted as $\mathrm{tr}(A)$ and $\det(A)$, respectively. We use $\mathbb{I}$ to describe an identity matrix of appropriate size. A normally distributed variable z with mean μ and covariance matrix Σ is denoted by $z \sim \mathcal{N}(\mu,\, \Sigma)$.

B Kalman filtering

Since the considered dynamical system (1) is linear in the unknown parameters $\theta(k)$, learning the system dynamics reduces to a parameter estimation problem. Given the Gaussian noise $w(k)$ and $v(k)$, a time-varying Kalman filter [19] can be designed using the random walk (2) as a process model and the dynamical system (1) as a measurement model. The posterior distribution $\theta(k) \sim \mathcal{N}(\hat{\theta}(k), \Sigma(k))$ of the system parameters is computed via the following Kalman filter recursions

$$\Phi(k) = \Phi(x(k), u(k)), \tag{4a}$$
$$S(k) = \Phi(k)(\Sigma(k) + \Sigma_v)\Phi(k)^\top + \Sigma_w, \tag{4b}$$
$$K(k) = (\Sigma(k) + \Sigma_v)\Phi(k)S(k)^{-1}, \tag{4c}$$
$$\hat{\theta}(k{+}1) = \hat{\theta}(k) + K(k)(x(k{+}1) - \Phi(k)\hat{\theta}(k)), \tag{4d}$$
$$\Sigma(k{+}1) = (\mathbb{I} - K(k)\Phi(k))(\Sigma(k) + \Sigma_v), \tag{4e}$$

where $S(k)$ is the innovation covariance, $K(k)$ is the Kalman gain, $\hat{\theta}(k{+}1)$ is the updated parameter estimate, and $\Sigma(k{+}1)$ is the updated parameter covariance. The prior distribution $\theta(0) \sim \mathcal{N}(\hat{\theta}(0), \Sigma(0))$

of the system parameters can be chosen such that the true initial parameter lies in the following confidence ellipsoid

$$\Theta_p = \{\theta \mid (\theta - \hat{\theta}(0))^\top \Sigma(0)^{-1} (\theta - \hat{\theta}(0)) \leq \chi^2_{n_\theta}(p)\} \tag{5}$$

with probability threshold $0 < p < 1$ and Chi-squared distribution $\chi^2_{n_\theta}(p)$. A detailed overview of Bayesian filtering and estimation can be found in [33].

C Zone model predictive control

A standard zone MPC [14] without active learning for system (1) is defined by the following optimization problem

$$\begin{aligned}
\min_{x,\,u,\,\epsilon} \quad & \sum_{i=0}^{N-1} \|Q\epsilon_i\|_1 + \|\Delta u_i\|_R^2 && \text{(6a)}\\
\text{s.t.} \quad & \forall\, i = 0, \ldots, N-1, \\
& x_{i+1} = \Phi(x_i, u_i)\hat{\theta}(k), && \text{(6b)}\\
& \Delta u_i = u_i - u_{i-1}, && \text{(6c)}\\
& (u_i, \Delta u_i) \in \mathbb{U} \times \Delta\mathbb{U}, && \text{(6d)}\\
& x_{\min} - \epsilon_i \leq x_i \leq x_{\max} + \epsilon_i, && \text{(6e)}\\
& \epsilon_i \geq 0,\ x_0 = x(k),\ u_{-1} = u(k-1), && \text{(6f)}
\end{aligned}$$

where N is the prediction horizon length and ϵ is a slack variable quantifying the violations of zone constraints on the system state. The slacks ϵ are penalized with $Q \succ 0$ using a 1-norm to impose an exact penalty cost on these violations, while $R \succ 0$ weights the quadratic cost on the input rate. The described zone MPC follows a certainty-equivalence approach, where the mean parameter estimate $\hat{\theta}(k)$ is used to predict the future system trajectory. The zone MPC operates in a receding horizon fashion, i.e., at each time step k, the optimization problem is solved using the previous input $u(k-1)$, the current state measurement $x(k)$ and parameter estimate $\hat{\theta}(k)$, but only the first element of the optimized input sequence $u(k) = u_0^\star$ is applied to system (1). Due to the soft constraints, the problem is recursively feasible.

D Active learning in model predictive control

Within the context of MPC, the classical approach to collect informative data during closed-loop operation is to integrate a constraint in the optimization problem that enforces the input trajectory to be persistently exciting [11, 25]. A more direct approach to active learning MPC integrates the parameter estimator equations into the MPC constraints, facilitating the prediction of the covariance of the future parameter estimates and providing a direct assessment of the control input's effect on the parameter estimation uncertainty [18, 3]. However, instead of using a slacked information constraint that allows to directly define a desired parameter estimation accuracy as in our proposed active learning zone MPC, these approaches directly penalize the predicted covariance matrices in the objective function, requiring more involved hyperparameter tuning for the MPC cost function.

Few works exist that specifically address the problem of active learning in zone MPC. In [34], excitation is enforced by adding a dithering signal in the cost function of the steady-state target optimizer used for the zone MPC. In [13], a zone MPC is proposed that regulates the system into an invariant set, inside which a persistently exciting input is safely applied. In [38], a two-stage approach adapts the nominal zone MPC trajectory to maximize the smallest eigenvalue of the incremental information matrix, subject to a bound on the cost increase. A drawback of these approaches is that in contrast to our proposed active learning zone MPC, the active system excitation is performed independently of the current information about the system parameters, potentially resulting in unnecessary or even undesired control interventions.

E Cerebrospinal fluid dynamics

The description of the CSF dynamics in this work is based on the well-established Marmarou model [28]. The CSF system is modeled as a scalar system, where the change of CSF Volume $\dot{V}_{\mathrm{csf}}$ depends on the constant natural CSF formation Q_{form}, the natural CSF absorption Q_{absorb}, the externally applied drainage or infusion of CSF Q_{ex}, and the physiological CSF fluctuations Q_{phy}, such that

$$\dot{V}_{\mathrm{csf}}(t) = Q_{\mathrm{form}} - Q_{\mathrm{absorb}}(t) - Q_{\mathrm{ex}}(t) + Q_{\mathrm{phy}}(t). \tag{7}$$

Assuming that P_{csf} always lies above the constant dural sinus pressure P_{ds}, the natural absorption of CSF is defined as

$$Q_{\mathrm{absorb}}(t) = \frac{P_{\mathrm{csf}}(t) - P_{\mathrm{ds}}}{R_{\mathrm{out}}}, \tag{8}$$

where R_{out} is the CSF outflow resistance. If no external input is applied, i.e., $Q_{\mathrm{ex}} = 0$, the CSF pressure will converge to a resting pressure P_{r} at which the natural formation of CSF equals the natural absorption and it holds that

$$Q_{\mathrm{form}} = \frac{P_{\mathrm{r}} - P_{\mathrm{ds}}}{R_{\mathrm{out}}}. \tag{9}$$

The CSF pressure is modeled as an exponential function

$$P_{\mathrm{csf}}(t) = P_0 \exp(k_{\mathrm{el}} V_{\mathrm{csf}}(t)), \tag{10}$$

where k_{el} is the cerebral elastance and P_0 a constant. Finally, the change in CSF pressure over time can be described as

$$\begin{aligned}
\dot{P}_{\mathrm{csf}} &= \frac{\partial P_{\mathrm{csf}}}{\partial V_{\mathrm{csf}}} \frac{\partial V_{\mathrm{csf}}}{\partial t} && \text{(11a)}\\
&= P_0 \exp(k_{\mathrm{el}} V_{\mathrm{csf}}) k_{\mathrm{el}} \dot{V}_{\mathrm{csf}} && \text{(11b)}\\
&= \frac{k_{\mathrm{el}} P_{\mathrm{r}}}{R_{\mathrm{out}}} P_{\mathrm{csf}} - \frac{k_{\mathrm{el}}}{R_{\mathrm{out}}} P_{\mathrm{csf}}^2 - k_{\mathrm{el}} P_{\mathrm{csf}} Q_{\mathrm{ex}} + k_{\mathrm{el}} P_{\mathrm{csf}} Q_{\mathrm{phy}}, && \text{(11c)}
\end{aligned}$$

where the time argument (t) has been omitted for brevity. After defining the state $x = P_{\mathrm{csf}}$ and input $u = Q_{\mathrm{ex}}$, Euler discretization with sampling time Δ_t is performed to derive the nonlinear discrete-time system model

$$\begin{aligned}
x(k{+}1) &= x(k) + \Delta_t \dot{x}(k) && \text{(12)}\\
&= \underbrace{\begin{bmatrix} x(k) & -x(k)^2 & -x(k)u(k) \end{bmatrix}}_{\Phi(x(k),u(k))} \underbrace{\begin{bmatrix} 1 + \Delta_t \frac{k_{\mathrm{el}} P_{\mathrm{r}}}{R_{\mathrm{out}}} \\ \Delta_t \frac{k_{\mathrm{el}}}{R_{\mathrm{out}}} \\ \Delta_t k_{\mathrm{el}} \end{bmatrix}}_{\theta} + \underbrace{\Delta_t k_{\mathrm{el}} x(k) Q_{\mathrm{phy}}}_{w(k)}
\end{aligned}$$

with unknown parameters θ but known features $\Phi(x_k, u_k)$. There exists a one-to-one mapping to compute the physiological parameters k_{el}, R_{out} and P_{r} from the technical parameters θ via the following equations

$$k_{\mathrm{el}} = \frac{\theta_{[3]}}{\Delta_t}, \quad R_{\mathrm{out}} = \frac{\theta_{[3]}}{\theta_{[2]}}, \quad P_{\mathrm{r}} = \frac{\theta_{[1]} - 1}{\theta_{[2]}}. \tag{13}$$

Table 2: Variables of the cerebrospinal fluid dynamics model

Variable	**Name**	**Unit**
Q_{form}	Formation rate	$\mathrm{mL/min}$
Q_{absorb}	Absorption rate	$\mathrm{mL/min}$
Q_{ex}	Drainage or infusion rate	$\mathrm{mL/min}$
Q_{phy}	Physiological disturbance rate	$\mathrm{mL/min}$
V_{csf}	Cerebrospinal fluid volume	mL
P_{csf}	Cerebrospinal fluid pressure	mmHg
P_{ds}	Dural sinus pressure	mmHg
P_{r}	Resting pressure	mmHg
R_{out}	Outflow resistance	$\mathrm{mmHg\,min/mL}$
k_{el}	Cerebral elastance	$1/\mathrm{mL}$

Table 3: Range of physiological parameters for simulations

Parameter	**Mean**	**Std. Dev.**	**Range**	**Source**
R_{out}	10.45	2.03	[6.47 14.43]	[4]
k_{el}	0.33	0.08	[0.17 0.49]	[4]
P_{r}	14.63	0.57	[13.50 15.75]	[26]