Learning and Adaptation in Model Predictive Controller Design with Constraints Handling Certifications

This research project will be centred on safe and data-driven decision-making for constrained dynamical systems, at the intersection of control, optimization, and learning. A particular interest will be placed on how learning-based methods can be integrated with model predictive control and safety supervision to enable practical deployment for reliable autonomy under severe uncertainty.

1 Background

Model predictive control (MPC) is an optimization-based technique that utilizes a mathematical model of a system to predict its future behaviour and determine optimal control decisions, while satisfying hard constraints on the system states and inputs [1]. Recent advances have been reported on MPC design for discrete-time linear time-invariant systems under hard state and input constraints, while relaxing two restrictive assumptions commonly made in the literature: exact model knowledge and full state availability. The *adaptive output feedback MPC* theory developed in [2, 3] considers the system model to have both parametric and additive uncertainties with known bounds. The framework proposes the integration of an adaptive observer [4–6] with a suitably reformulated MPC constrained optimal control problem to tackle the uncertainties in the model as well as the lack of state measurements. The observer [6] is designed to collect input-output data and provide simultaneous online estimates of unknown parameters and unmeasured states. These estimates are then incorporated into the suitably reformulated MPC optimization problem that explicitly utilizes the observer dynamics for prediction. The MPC optimization problem generates an input sequence together with a homothetic tube containing the estimated state trajectory. By augmenting this tube with sets that bound the state estimation error, an outer tube is obtained that is guaranteed to contain the true state trajectory. The resulting two-tube architecture ensures recursive feasibility, boundedness of all closed-loop signals, and robust exponential stability of the origin for both the plant and the observer. These guarantees are established using a suitably designed terminal set, computed as a maximal admissible robust positive invariant set [7] accounting for both parametric and additive uncertainties.

2 Challenges

A common feature in the adaptive output feedback MPC framework [2, 3] discussed above, as well as in other adaptive MPC methods [8–15], is their reliance on *quadratic stabilizability* assumption, which implicitly restrict the admissible parametric uncertainty. Quadratic stabilizability implies the existence of a common stabilizing feedback gain for all the elements in the parametric uncertainty set. While such assumption simplifies stability analysis and controller synthesis, they inherently limit the class of systems for which standard guarantees apply [16–19]. In parallel, a separate but equally important challenge in MPC for uncertain systems is the need for *constraint tightening* to ensure robust feasibility. When uncertainty is large, this tightening can become excessively conservative, often shrinking the feasible region to the point where meaningful control performance is no longer achievable. Furthermore, slowly varying parameters (due to aging, wear, etc.) and coarse instrumentation frequently render standard MPC formulations impractical when hard constraints must be enforced. In such regimes, insisting on optimal or near-optimal MPC behaviour from the outset is not realistic, and feasibility or resilient MPC need to be privileged, for example, by resorting to k-invariant sets [20] or inner-outer pairs of invariant sets [21] to address the feasibility.

3 Research objectives

The objectives of this project, aimed at addressing the challenges discussed in Sec. 2, are briefly mentioned below.

3.1 Robust MPC with adaptive tubes

This postdoctoral research project will explore the development of an MPC framework that will utilize robust MPC theories [22, 23] and adaptive tubes to relax conventional quadratic stabilizability requirements. Unlike existing robust tubes [22, 24], the proposed adaptive tube would improve its structure whenever there is a reduction in the parametric uncertainty through online learning from data, thereby allowing for better efficiency in decision-making. A primary focus will be to explicitly decouple safety assurances from optimality criteria during the learning phase. This direction will raise fundamental questions concerning invariant set characterization (e.g. by means of state-dependent bounds [25]), robustness metrics, and stability under time-varying uncertainty, particularly within constrained and partially observed systems. A further objective will be to translate these theoretical advances to practical platforms in domains such as robotics, power transmission networks, and building climate control. To ensure practical viability, the adaptive tube design will require rigorous guarantees of recursive feasibility and stability. To achieve these provable guarantees, the investigation will focus on concepts including dwell times for switched systems, data-driven construction of piecewise affine Lyapunov functions [26], and the use of k-invariant sets [23].

3.2 Safe MPC with a balance between learning, optimality, and computational efficiency

Starting from the observation that the existing literature on MPC—whether adaptive [2, 3, 9–12] or robust [8, 22, 24]—is insufficient as a standalone tool for systems with large initial uncertainty and strict, potentially state-dependent safety constraints, the work will focus on developing safe learning architectures that will integrate data-driven methods, such as learning-based models governed by barrier function constraints. The proposed framework will be designed to initiate learning conservatively, prioritizing absolute safety and constraint satisfaction while gradually improving model fidelity. This approach will allow the system to operate initially in a deliberately non-optimal yet provably safe regime, where learning and uncertainty reduction can proceed without aggressive constraint tightening. As model confidence improves and uncertainty is sufficiently reduced, the controller will progressively transition to a full-fledged MPC formulation with tighter performance objectives and an expanded region of operation.

The explored state space regions will inherit lower bounds on the uncertainty, while the regions seldom used will be associated with larger uncertainty sets. This concept will be further extended to handle systems with slowly-varying or periodically-varying parameters. Here, a barrier function (or positive invariance) will be employed as a supervisory layer to constantly monitor MPC performance, similar to fault detection schemes [27]. Whenever parameter variations threaten to drive the system towards unsafe operation, the supervisory controller will take corrective action. In such events, learning can be restarted within a non-optimal yet safe regime to regain sufficient model knowledge. For parameters that vary over sufficiently long durations, data-driven methods will also be investigated to learn and effectively predict the trends of the variations.

3.3 Observability and identifiability for adaptive output feedback control

Since adaptive observers are a useful tool in designing adaptive output feedback control for practical systems, in this work, we will focus on a foundational question concerning the relationship between observability and identifiability in adaptive observer design. The property of observability of a given system realization is widely explored and used in designing observers when the system parameters are known [28, 29]. Likewise, identifiability [30–32] of unknown parameters has been studied for the design of adaptive laws, assuming a certain structure on the system parameters and/or the availability of state measurements. However, in practical scenarios, exact model knowledge and state measurements are hard to obtain. This necessitates the design of suitable adaptive observers that can provide online simultaneous estimates of the states as well as the parameters. While the existing literature [4–6, 33, 34] strongly emphasizes observability, parameter identifiability is often treated implicitly. Although substantial literature on identifiability exists, its intricate relationship with observability will require further exploration for various systems. For multi-input multi-output systems, certain parameters may be structurally unidentifiable despite state observability, depending on system realization and excitation conditions. This distinction will have important implications for adaptive learning and control, particularly in scenarios where unknown terms do not lie within the assumed basis space or whenever the constraints (in particular, input constraints) are active. Consequently, a key objective will be to develop explicit characterizations of identifiability alongside observability. Building on these characterizations, the research will aim to design adaptive output feedback controllers that explicitly account for (and avoid) scenarios of partial or structural non-identifiability.

4 Timeline

The project will systematically advance this research agenda from foundational theory to implementable methodologies. The initial phase (one or two months) will be dedicated to learn the new concepts developed by the faculty at L2S. Following this, contributions are sought towards

- developing rigorous theoretical guarantees for safety-separated learning and control architectures along with the design of suitable state and parameter estimators,
- translating these guarantees into computationally tractable MPC, and
- validating the resulting methods on representative practical systems.

The structure of the fellowship is particularly well-suited to pursuing this progression without prematurely prioritizing performance over safety, enabling a balanced exploration of theory, algorithms, and experimental validation.

5 Broader impact

The broader impact of this project lies in enabling safe, reliable, and trustworthy decision-making under uncertainty. The combination of the three works described in Sec. 3 will yield an MPC framework that can be implemented across a wider class of real-life systems. By explicitly separating safety from optimality during learning, the proposed frameworks aim to reduce the risks associated with deploying optimization- and learning-based controllers in safety-critical environments like automotive or energy transmission systems.

6 Application procedure

6.1 Qualifications:

Suitable candidates should have:

- A Ph.D. in Control Systems, Electrical Engineering, Computer Science, or a closely related field.
- A strong background in control theory and optimisation, evidenced by publications in top-tier journals and conferences.
- Experience with model predictive control and its applications.
- Excellent programming skills for the simulation and implementation of control algorithms.
- Strong analytical and problem-solving abilities.
- Effective communication skills and the ability to work collaboratively in a multidisciplinary team.
- Familiarity with power systems, particularly transmission networks, is a plus.

6.2 Application Process:

Interested candidates should submit the following:

- A cover letter detailing their research experience and interest in the position.
- Transcript of academic records.
- A curriculum vitae (CV) highlighting relevant qualifications and publications.
- Contact information for at least two professional references.

Please send your application materials with object [Postdoc adaptation and learning MPC] to Sorin Olaru.

The position will be available subject to internal L2S validation of the financing. A National Security clearance is required, and the process takes approximately 2 months.

The postdoc will present the developments to both academic and industrial international meetings and conferences.

References

[1] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, LLC, 2017.

[2] A. Dey and S. Bhasin, "Adaptive output feedback MPC with guaranteed stability and robustness," *IEEE Transactions on Automatic Control*, vol. 70, no. 12, pp. 8345–8352, 2025.

[3] A. Dey, A. Dhar, and S. Bhasin, "Adaptive output feedback model predictive control," *IEEE Control Systems Letters*, vol. 7, pp. 1129–1134, 2023.

[4] G. Kreisselmeier, "Adaptive observers with exponential rate of convergence," *IEEE Transactions on Automatic Control*, vol. 22, no. 1, pp. 2–8, 1977.

[5] T. Suzuki, T. Nakamura, and M. Koga, "Discrete adaptive observer with fast convergence," *International Journal of of Control*, vol. 31, no. 6, pp. 1107–1119, 1980.

[6] A. Dey and S. Bhasin, "Adaptive observers for MIMO discrete-time LTI systems," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 8708–8713, 2023.

[7] A. Dey and S. Bhasin, "Computation of maximal admissible robust positive invariant sets for linear systems with parametric and additive uncertainties," *IEEE Control Systems Letters*, vol. 8, pp. 1775–1780, 2024.

[8] M. Bujarbaruah, U. Rosolia, Y. R. Stürz, and F. Borrelli, "A simple robust MPC for linear systems with parametric and additive uncertainty," in *American Control Conference*, pp. 2108–2113, 2021.

[9] S. Jafari Fesharaki, M. Kamali, and F. Sheikholeslam, "Adaptive tube-based model predictive control for linear systems with parametric uncertainty," *IET Control Theory & Applications*, vol. 11, no. 17, pp. 2947–2953, 2017.

[10] A. Dhar and S. Bhasin, "Indirect adaptive MPC for discrete-time LTI systems with parametric uncertainties," *IEEE Transactions on Automatic Control*, vol. 66, no. 11, pp. 5498–5505, 2021.

[11] M. Lorenzen, M. Cannon, and F. Allgöwer, "Robust MPC with recursive model update," *Automatica*, vol. 103, pp. 461–471, 2019.

[12] X. Lu and M. Cannon, "Robust adaptive model predictive control with persistent excitation conditions," *Automatica*, vol. 152, p. 110959, 2023.

[13] D. Tranos, A. Russo, and A. Proutiere, "Self-tuning tube-based model predictive control," in *American Control Conference*, pp. 3626–3632, 2023.

[14] J. Köhler, E. Andina, R. Soloperto, M. A. Müller, and F. Allgöwer, "Linear robust adaptive model predictive control: Computational complexity and conservatism," in *IEEE Conference on Decision and Control*, pp. 1383–1388, 2019.

[15] X. Lu and M. Cannon, "Robust adaptive tube model predictive control," in *American Control Conference*, pp. 3695–3701, 2019.

[16] P. P. Khargonekar, I. R. Petersen, and K. Zhou, "Robust stabilization of uncertain linear systems: quadratic stabilizability and H^∞ control theory," *IEEE Transactions on Automatic Control*, vol. 35, no. 3, pp. 356–361, 1990.

[17] P. Gahinet and P. Apkarian, "A linear matrix inequality approach to H^∞ control," *International Journal of Robust and Nonlinear Control*, vol. 4, no. 4, pp. 421–448, 1994.

[18] B. R. Barmish, "Necessary and sufficient conditions for quadratic stabilizability of an uncertain system," *Journal of Optimization theory and applications*, vol. 46, no. 4, pp. 399–408, 1985.

[19] I. Petersen, "Quadratic stabilizability of uncertain linear systems: Existence of a nonlinear stabilizing control does not imply existence of a linear stabilizing control," *IEEE Transactions on Automatic Control*, vol. 30, no. 3, pp. 291–293, 1985.

[20] S. Olaru, M. Soyer, Z. Zhao, C. E. T. Dórea, E. Kofman, and A. Girard, "From relaxed constraint satisfaction to p-invariance of sets," *IEEE Transactions on Automatic Control*, vol. 69, no. 10, pp. 7036–7042, 2024.

[21] R. Comelli, S. Olaru, and E. Kofman, "Inner–outer approximation of robust control invariant sets," *Automatica*, vol. 159, p. 111350, 2024.

[22] S. V. Raković, B. Kouvaritakis, R. Findeisen, and M. Cannon, "Homothetic tube model predictive control," *Automatica*, vol. 48, no. 8, pp. 1631–1638, 2012.

[23] Z. Zhao, A. Girard, and S. Olaru, "Nonlinear model predictive control based on K-step control invariant sets," *European Journal of Control*, vol. 80, p. 101040, 2024. 2024 European Control Conference Special Issue.

[24] S. Diaconescu, F. Stoican, B. D. Ciubotaru, and S. Olaru, "Elastic tube model predictive control with scaled zonotopic sets," *IEEE Control Systems Letters*, vol. 8, pp. 1343–1348, 2024.

[25] N. Athanasopoulos, E. Vlahakis, S. Olaru, and C. Townsend, "Equivalence of state dependent disturbances to piecewise polytopic affine dynamics." HAL preprint, hal-04493695v2, 2024.

[26] M. Tacchi, Y. Lian, and C. N. Jones, "Robustly learning regions of attraction from fixed data," *IEEE Transactions on Automatic Control*, vol. 70, no. 3, pp. 1576–1591, 2025.

[27] S. Olaru, J. D. Doná, M. Seron, and F. Stoican, "Positive invariant sets for fault tolerant multisensor control schemes," *International Journal of Control*, vol. 83, no. 12, pp. 2622–2640, 2010.

[28] D. G. Luenberger, "Observing the state of a linear system," *IEEE Transactions on Military Electronics*, vol. 8, no. 2, pp. 74–80, 1964.

[29] D. Boutat and G. Zheng, *Observability and Observer for Dynamical Systems*, pp. 1–29. Cham: Springer International Publishing, 2021.

[30] M. Grewal and K. Glover, "Identifiability of linear and nonlinear dynamical systems," *IEEE Transactions on Automatic Control*, vol. 21, no. 6, pp. 833–837, 1976.

[31] K. Godfrey and J. DiStefano, "Identifiability of model parameter," *IFAC Proceedings Volumes*, vol. 18, no. 5, pp. 89–114, 1985. 7th IFAC/IFORS Symposium on Identification and System Parameter Estimation, York, UK, 3-7 July.

[32] L. Ljung and T. Glad, "On global identifiability for arbitrary model parametrizations," *Automatica*, vol. 30, no. 2, pp. 265–276, 1994.

[33] Q. Zhang, "Adaptive observer for multiple-input-multiple-output (MIMO) linear time-varying systems," *IEEE Transactions on Automatic Control*, vol. 47, no. 3, pp. 525–529, 2002.

[34] A. Katiyar, S. B. Roy, and S. Bhasin, "Initial excitation based robust adaptive observer for MIMO LTI systems," *IEEE Transactions on Automatic Control*, 2022.