Master thesis

Approximation function analysis for explicit optimal control of electric drives

In optimal control of electric drives one can either implicitly optimize the control input at discrete-time steps by solving the underlying control problem on a receding horizon or one can try to explicitly find a control policy function which directly maps measured states to control actions. The latter approach is often referred to as explicit optimal control and requires using approximation functions to address continuous (i.e., infinite) state and action spaces. Once the (approximate) optimal control policy has been found, the inference is typically much faster than in the implicit case where an online optimization process must be carried out at each controller cycle. Since controller decision time intervals are in the sub-millisecond range for electric drives, the fast online inference of explicit optimal control is a compelling feature. Here, the potential control policy approximation functions span a wide range of function classes, such as neural networks, Gaussian processes, or Laguerre polynomials [1]. The control policy can be learned both from data (e.g., reinforcement learning [2]) or based on an available plant model (differential predictive control [3]). In both cases, the topology of the approximation function plays a crucial rule in the performance of the control policy as well as the resulting numerical complexity during the training and inference phase. While the specific choice of an approximation function is often based on ad-hoc heuristics, the question of how to systematically select the best approximation function for a given control task remains largely open.

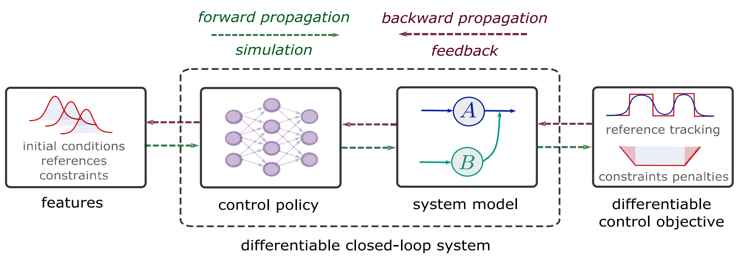

(a) Differential predictive control overview [3].

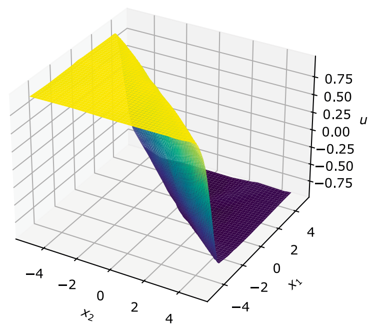

(b) Example policy mapping states to actions [3].

Key research questions:

- What policy approximation functions are best for explicit optimal control of electric drives?
- Can we automate the selection of the best approximation function for a given control task?
- What are the differences in terms of performance and computational complexity compared to (exact) implicit control approaches with established optimization solvers?

Necessary requirements:

- Finished course work on data or model-based optimal control
- Ideally: finished course work on electric drives
- Solid skills in scientific programming languages (e.g., Julia, JAX, PyTorch)

WP 1: Literature research **[3 weeks]**
Scanning the scientific literature for relevant publications and patents related to explicit optimal of electric drives and related fields is the first step. Also, getting familiar with typical constrained optimization problems in the drive control domain is part of this WP. Moreover, relevant (open-source) software work in the field (e.g., NeuroMANCER) should be considered. This also includes the identification of relevant keywords as part of the search strategy. Relevant work will be stored in a literature review software (e.g., JabRef) and summarized in the thesis.

WP 2: Machine control **[5 weeks]**
A typical use case for electric drive control with the need of long horizon optimal control are induction machines (IMs) due to their relatively large rotor time constant. This is a particular challenge in the domain of IM torque control [4], where implicit optimal control methods (in particular MPC) struggles to achieve the desired performance. The goal of this WP is to analysis the torque control problem of an IM and to develop an end-to-end differentiable model of the control problem. Hence, a software toolchain should be developed that allows training and evaluating different approximation functions for the control policy. The training data should be generated by solving the optimization problem for a set of random reference trajectories and different drive conditions.

WP 3: Policy function analysis **[5 weeks]**
While black box neural networks are often used as approximation functions for control policies, the question arises if more interpretable and especially numerically lightweight function topologies can be used for the control policy approximation. This WP aims to analyze different function classes (e.g., Gaussian processes, Laguerre polynomials) for the control policy approximation. The analysis should focus on the performance of the control policy as well as the numerical complexity during training and inference. A possible starting point could be to solve the underlying control problem in an implicit way to get a reference for the performance of the explicit control policy and an idea how the associated control policy could look like in terms of function topology.

WP 4: Features **[2 weeks]**
In order to reduce the complexity of the control policy approximation function, the state space selection could be enhanced by selecting a set of relevant features. This WP aims to identify a set of relevant features for the control policy approximation such that the overall numerical complexity of the control policy is reduced while remaining at the same performance level.

WP 5: Hyperparameters **[3 weeks]**
The explicit optimal control solvers (e.g., stochastic gradient descent as part of the DPC) as well as the approximation functions have a set of hyperparameters that need to be tuned. The hyperparameter optimization (HPO) can be automated using available open-source toolboxes and should also focus on the complexity vs. accuracy trade-off by identifying the Pareto front between these two objectives. Depending on the computational load of this WP, the HPO can be parallelized on a high-performance computing cluster.

WP 6: Embedded deployment **[4 weeks]**
Since a special focus of this thesis is on fast online inference, the trained policies should be deployed on an embedded system (e.g, a dSPACE rapid control prototyping platform). This requires the conversion of the previously learned policies to a format that can be executed on the embedded system. The performance of the embedded optimizer should be evaluated in terms of calculation speed and accuracy based on a software-in-the-loop scenario.

WP 7: Documentation **[3 weeks]**
All work packages should be reported in a structured way within the thesis. A LaTeX template should be used for this purpose: `https://github.com/IAS-Uni-Siegen/thesis_latex_template`. Writing instructions can be found within the provided template source files. Based on the previous empirical findings, conclusions should be drawn, and future research directions should be outlined.

Gantt chart

The planned timetable is shown in the Gantt diagram below.

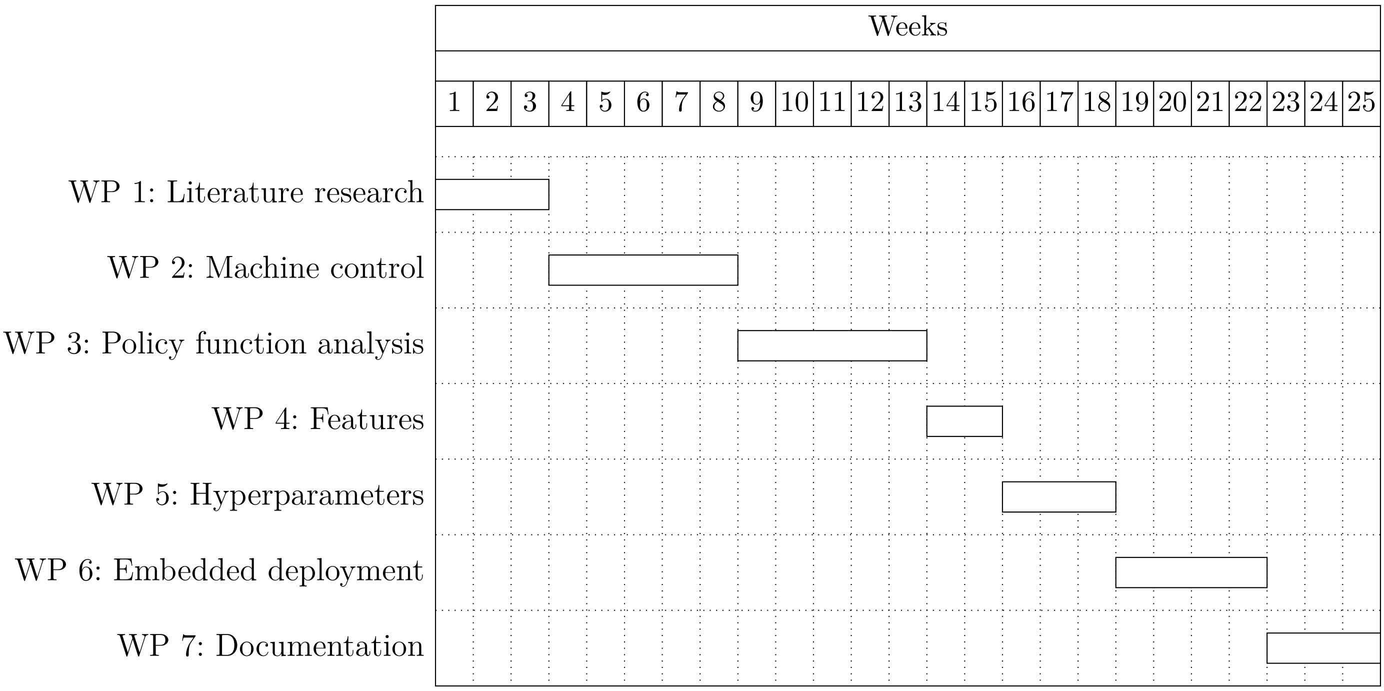

Figure 2: Gantt chart for the thesis.

References

[1] Y. Xu, S. Li, W. Zhang, G. Yu, and J. Zou, "Long-horizon constrained model predictive direct speed control for pmsm drives based on laguerre functions," *IEEE Transactions on Control Systems Technology*, vol. 32, no. 3, pp. 1002–1014, 2024.

[2] M. Schenke, W. Kirchgässner, and O. Wallscheid, "Controller design for electrical drives by deep reinforcement learning: A proof of concept," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 7, pp. 4650–4658, 2020.

[3] J. Drgoňa, A. Tuor, and D. Vrabie, "Learning constrained parametric differentiable predictive control policies with guarantees," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 54, no. 6, pp. 3596–3607, 2024.

[4] M. Becker, M. Stender, and O. Wallscheid, "Nonlinear efficiency-optimal model predictive torque control of induction machines," *Authorea Preprints*, 2023.