ENHANCING THE GUIDANCE, NAVIGATION, AND CONTROL OF AUTONOMOUS PARAFOILS USING MACHINE LEARNING METHODS

Lukas Hewing[(1)], Dennis Gramlich[(2)], Chris Verhoek[(3)], Joost Veenman[(1)], Rafael Polonio[(1)], Carlos Ardura[(1)], Roland Tóth[(3)], Christian Ebenbauer[(2)], Carsten Scherer[(4)], Valentin Preda[(5)]

[(1)] *SENER Aeroespacial, C/ Severo Ochoa, 4 (PTM) – 28760 Tres Cantos – Madrid, Spain,*
{lukas.hewing, joost.veenman, rafael.polonio, carlos.ardura}@aeroespacial.sener
[(2)] *RWTH Aachen University, Kopernikusstr. 16, 52074 Aachen, Germany,*
{dennis.gramlich, christian.ebenbauer}@ic.rwth-aachen.de
[(3)] *Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands,*
{c.verhoek, r.toth}@tue.nl
[(4)] *University of Stuttgart, Pfaffenwaldring 5a, 70569 Stuttgart, Germany,*
carsten.scherer@mathematik.uni-stuttgart.de
[(5)] *European Space Agency - ESTEC, Keplerlaan 1, PO Box 299, 2200 AG Noordwijk, The Netherlands, valentin.preda@esa.int*

ABSTRACT

In this paper, we give a brief account of the results obtained in the ESA technology development activity *Artificial Intelligence for GNC Systems*[1]. The aim of the study was to research (knowledge generation), assess the feasibility of, and industrialize (knowledge integration) machine learning techniques for guidance, navigation, and control (GNC) systems. Within the scope of the activity a wealth of different results were produced which now warrant a recapitulation and critical evaluation. The wide scope of results is aligned with the explorative essence of the project paired with developments of novel technologies whose real application to GNC systems is at a low TRL (Technology Readiness Level). The output includes results in three main research directions, namely i) Bayesian optimization for automated GNC tuning and worst-case analysis, ii) Robust optimization-based guidance, and iii) Learning-based model augmentation. These were demonstrated on a variety of problems, including a challenging autonomous parafoil landing scenario derived from the Space Rider mission.

1 INTRODUCTION

It is well-known that artificial intelligence (AI) and machine learning (ML) techniques have developed into a transformative force across many industries. Their industrial adaption in aerospace guidance, navigation and control (GNC) systems, however, has been rather limited to date. One reason for this is that ML tools are still perceived as complex and intransparent in the space industry such that clear benefits that outweigh these disadvantages need to be demonstrated. Nevertheless, the impact of the technology can no longer be ignored and its potential application to GNC systems is receiving a lot of interest lately.

[1]This activity was carried out and led by SENER Aeroespacial, Madrid, Spain, in collaboration with the Eindhoven University of Technology, Eindhoven, The Netherlands, the RWTH Aachen University, Aachen, Germany, and the University of Stuttgart, Stuttgart, Germany. The activity was funded by the European Space Agency (ESA), contract no. 4000133595/20/NL/CRS.

The adoption of ML techniques can have clear benefits: Smart algorithms are capable of processing previously unimaginable quantities of data and learning algorithms have shown prediction accuracies that far surpass traditional models while simultaneously promising to streamline the development process by requiring fewer human interventions. Nevertheless, the technology entails risks and uncertainties that have to be accounted for to enable robust and safe decision-making (see, e.g, [5], [16], [17]). It is therefore of paramount importance to perform critical analyses of the use of AI technologies to plan for the next generation GNC systems due to their often immediate effects on mission safety.

In this paper, we provide an overview of the results obtained in an ESA technology development activity called "Artificial Intelligence for Guidance, Navigation and Control" (AI4GNC) in which the goal was to investigate the potential of ML methods to enhance the performance and robustness of aerospace GNC systems. As main outcome, it was demonstrated how a combination of ML methods reinforced with the well-established theoretical foundation from of robust control and system identification can be used to systematically achieve this goal for a representative study case derived from the Space Rider mission. In particular, we investigated several complementing technologies at different hierarchical levels in the GNC and its design process and demonstrated the gained advantages. Here we focused on three major aspects, namely: i) Bayesian optimization for automated GNC tuning and worst-case analysis, ii) Robust optimization-based guidance, and iii) Learning-based model augmentation.

Bayesian optimization for GNC tuning: We investigated the use of Bayesian optimization (BO) techniques. BO has recently shown enormous success for data-efficient stochastic black-box optimization, for instance for hyper-parameters of ML algorithms [9]. It is also a prime technique researched for the automated tuning of complex control systems [6], [11]. In the present activity, we developed a GNC auto-tuner tool that utilizes BO techniques to efficiently tune high-level GNC parameters for complex natural language constraints or objectives formulated in terms of temporal logic expressions. The use of BO enables a data-efficient stochastic black-box optimization of several key GNC parameters using a small number of (simulation) experiments. We further demonstrate that it is straightforward to employ the techniques in an antagonistic fashion leading to an effective worst-case analysis tool. Our results show how such temporal logic-constrained BO can be efficiently used to improve system performance, explore parameter interdependencies and provide valuable insights to support the tuning of complex GNC systems.

Robust optimization-based guidance: Within the guidance layer, a robust trajectory planning technique based on on-board optimization was developed, which can take different sources of uncertainty into account. The planning method relies on a novel extension of differential dynamic programming (DDP) using standard tools from robust control (IQCs, LFTs) to formulate a sequence of semi-definite programs (SDPs) to find feedback & feedforward policies that efficiently steer the system despite the adversarial action of uncertainties. Extensive evaluations have shown a clear hierarchy of achieved performances: (nominal) optimization-based guidance outperforms the baseline solution, while the novel robust variant shows the strongest performance.

Learning-based model augmentation: Finally, at the intersection of guidance and controls, we employed data-driven system identification techniques to capture closed-loop system behaviour, with the aim to improve the higher-level planning and guidance algorithms. Such models are cumbersome to derive from first principles since flight software, including lower-level controllers, actuator saturation, and similar effects are part of the loop. Alternatively, neural networks can be trained with efficient deep-learning-based system identification methods that augment an idealized baseline model, i.e., learning-based model augmentation. This was shown to effectively reduce residual errors, while extending the region of validity compared to alternative methods, thereby providing an accurate system description to higher-level planning algorithms.

The remainder of this paper describes the carried-out activities and the results obtained during the study. In Section 2, we first summarize the main study case, which is based on the Space Rider mission. Subsequently, in Section 3, 4, and 5, we present a selected overview of results obtained on the topics of Bayesian optimization, Robust optimized guidance, and Learning-based model augmentation, respectively. The paper is concluded with a summary in Section 6.

2 BENCHMARK DEFINITION

To develop and validate the considered ML techniques, we consider a challenging autonomous parafoil landing scenario. This benchmark is inspired by the Space Rider mission, whose GNC is currently being developed by SENER and comprises a detailed simulation environment derived from the one used for the Space Rider mission. The environment was adjusted and extended within the AI4GNC activity, for example, to allow for full spatial distribution of winds (3D Wind) while the benchmark GNC architecture and functionality remained largely unchanged. This is based on a fictitious landing in the Yosemite Valley exhibiting strong spatial variations of the wind and allowing to stress the developed algorithms. Figure 1 shows an illustration.

The GNC system is in charge of guaranteeing an accurate landing under safe touch-down conditions. In the scope of this study, the system was required to (be able) to land within 150m from the specified target, compensating for the dispersions accumulated during the hypersonic and supersonic re-entry phases, while ensuring a soft landing ($<$ 3 m/s vertical velocity at impact) with no lateral loads guaranteeing the re-usability of the module. Throughout the project all developments were compared to the existing baseline GNC solution representing a GNC system of current industrial practice.

The baseline GNC implements a waypoint-based guidance strategy for the first parts of the flight under parafoil but uses a heuristic path planning algorithm for the final part [7]. Based on these waypoints, the approach and landing phase is therefore split into several legs imposing different guidance profiles as illustrated in Figure 1.

3 BAYESIAN OPTIMIZATION & TEMPORAL LOGIC

The goal of the investigation into *Bayesian optimization* (BO) for GNC tuning was to showcase the use of automated exploration of parameters of GNC systems such as to i) improve the performance of the resulting GNC software, and ii) gain engineering insights into complex parameter dependencies – all while (iii) automating the process and thereby reducing engineering effort and trading it off with computation power and large scale simulations. This automated exploration was therefore carried out using parametric batch-style simulations in the form of traditional Monte Carlo (MC) campaigns, representing a noisy evaluation of the GNC performance, while a few select parameters were chosen (optimized) by the Bayesian optimization algorithm in order to maximize the expected performance. The study had several features differentiating it from previous investigations and tools for global optimization more commonly used, for instance,

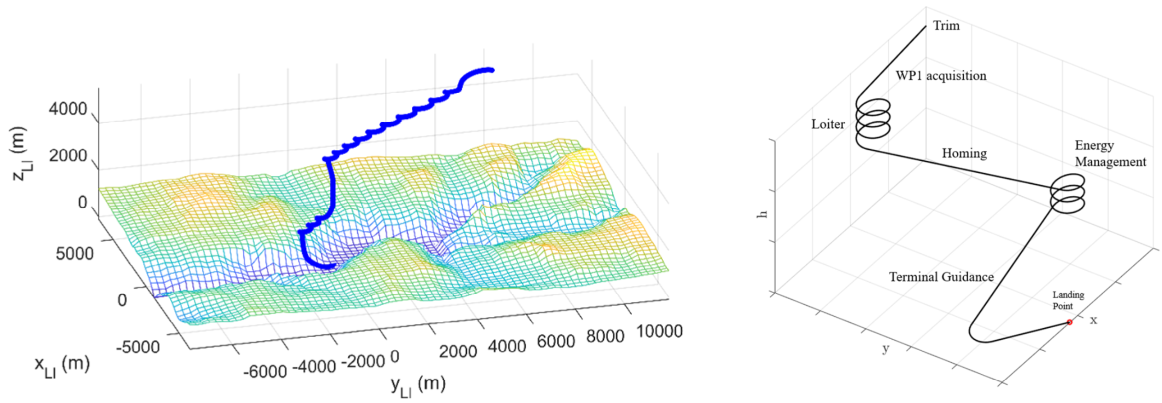

Figure 1: The right plot shows the relevant phases of the descent & landing considered in this study as a Benchmark GNC problem. The left plot shows a simulation shot of a trajectory in the extended scenario – landing in Yosemite Valley under the influence of spatial winds. Note the drift due to winds in the realized trajectory with respect to the plan on the right which is computed in wind-compensated coordinates.

- **Noisy simulations/function evaluations.** This corresponds to the typical main scenario in validation campaigns, in which the correct functionality of the GNC software is investigated under random perturbations and disturbances. Bayesian optimization is particularly suited to deal with such noisy optimizations, although noise remains a large challenge as demonstrated within this project. Due to this, it provides an intuitive interface towards established MC campaigns that GNC engineers are familiar with.
- **Temporal Logic Constraints.** Constraints in a global optimization tool can be used to optimize performance while ensuring that some other requirements are not violated for the optimizer found. While many global optimization tools do not consider such constraints, Bayesian optimization has been extended in this direction, including noisy evaluation of such constraints similar to the objective function. Naturally arising constraints often directly relate to requirements, which can typically be precisely and relatively easily expressed as *temporal logic* (TL) expressions, this constituted a second focal point of the activity in which it was demonstrated that temporal logic expressions can be effectively included in the constrained optimization making use of the *robustness degree* [1].
- **Interpretability & Engineering Insights** A final focus of the development was to create a tool which, besides optimizing the GNC performance, can provide some engineering insights over the parameter landscape. As such, we understand the developed tools as design tools within an overall process involving human engineers who the optimization tool informs, rather than dictates a result to. Aspects of this are found in plotting functionality of resulting surrogate Gaussian process models, as well as the interpretability of optimized hyperparameters of such models – large kernel lengthscales, for instance, indicate that the corresponding feature has little influence on the outcome, while the estimated noise levels or computed likelihoods give an indication of whether the data can be explained well by the selected features.

Various numerical examples were considered. These can be roughly grouped into three distinct categories: i) proof of concepts and algorithm performance analysis on simple examples presented as part of the tradeoff analysis, ii) application to the benchmark problem for the baseline and novel GNC developed within the project, and, finally, iii) the application in antagonistic optimization for worst-case-analysis. A selected overview is given next.

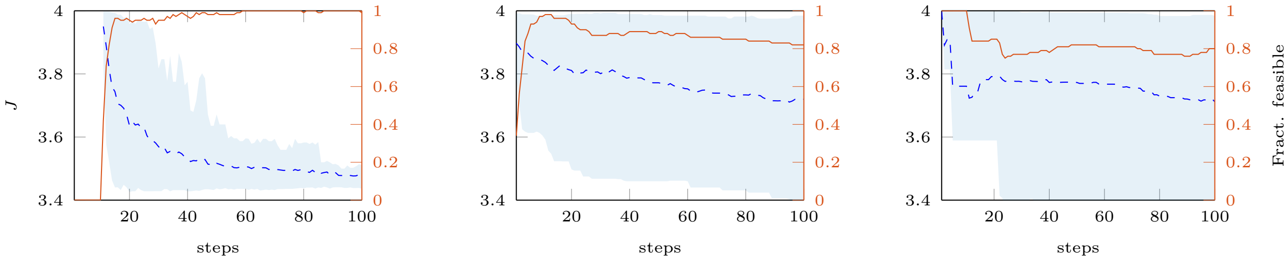

(a) Bayesian Optimization (b) Random (c) Surrogate optimization

Figure 2: Mass-Spring-Damper constrained results for different algorithms. Blue dashed line is the average over different optimization runs of the resulting performance of the optimizer, with shaded 10% and 90% quantiles, while the solid red line represents the fraction of respective optimizers leading to constraint satisfaction.

3.1 Algorithm Performance Analysis

In order to allow for a solid statistical foundation, the performance of the BO algorithm was primarily evaluated on simplified examples which can be simulated much more rapidly than the full-scale benchmark simulation. By an evaluation of the algorithm's performance, we here mean a repeated optimization using the BO algorithm, creating a statistic of its performance. This is necessary since the simulation itself contains random elements meaning that different realizations of random variables can lead to different algorithmic decisions and outcomes. The overall performance of an algorithm, in that case, can therefore only be judged by considering its overall statistics on the task, which, however, requires repeated optimization in order to generate these statistics. A second complication in evaluating the performance of the algorithm is that the true expected value given a certain parameterization is typically unknown and has to similarly be approximated by sampling. Due to this, the performance analysis study comprised several million simulation shots in order to paint a clear picture of its statistical behaviour.

The considered examples are a mass-spring-damper system and an inverted pendulum from the open-AI gym environment [2]. The exercise included a parameter study for different settings of the algorithm as well as a comparison to competing ones implemented in the Matlab global optimization toolbox. Figure 2 shows a representative example of such an analysis for the mass spring damper system, in which each optimization was repeated 100 times, and 10% and 90% quantiles, as well as the average performance, are shown. The BO algorithm with the associated *GncTuner* toolbox performed well and showed little dependency on its exact parameterization, e.g., selection of kernel or acquisition functions. What did prove to be crucial though, was the overall precise formulation as an optimization problem which was done including logarithmic domain or objective definitions as well as outlier treatment. In addition, it was demonstrated that constraints complicate the optimization task, but that the BO algorithm can effectively optimize also in this case including temporal logic specifications. These have been shown to effectively avoid chattering behaviour in controller switching for the inverted pendulum in [14].

3.2 GNC Tuning on Benchmark

The GNC autotuning tool was then successfully applied to the parafoil landing benchmark scenario, tuning high-level GNC parameters such as mode-switch conditions, waypoint placements and reference trajectory parameterizations, demonstrating the applicability of BO on

[2] `www.gymlibrary.dev`

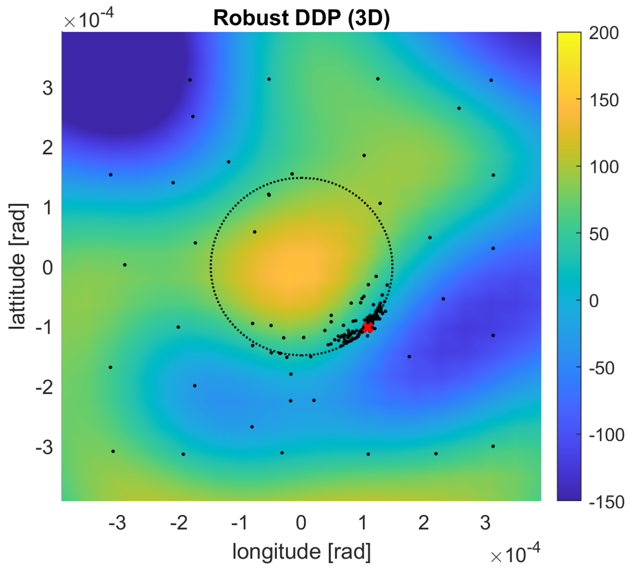

Figure 3: Worst-Case analysis using Bayesian optimization results of initial condition for terminal guidance phase for the 3D wind case using 200 simulations. The covered domain corresponds to roughly ± 2 kilometres around the previously identified optimum. Black dots show evaluated parameterizations, the red cross is the identified worst-case and the dashed line illustrates the radius identified as robustly satisfying the requirement.

industrially relevant examples. Nevertheless, it becomes apparent here that high noise intensities and susceptibility to outliers can present significant challenges in automated optimization. In particular, this holds true for higher-dimensional optimization tasks, considering, e.g., six parameters being optimized simultaneously. In these cases, the found ‘optimizers´ were generally satisfactory but likely correspond to locations near (local) minima. As such, with growing noise intensity and dimensionality of the problem, the reliability of the algorithm is observed to decrease, even though no catastrophic failures were observed and returned results were always satisfactory.

Furthermore, it was demonstrated that all relevant requirements for the optimization could be expressed as *temporal logic* expressions and were directly included as a constraint in the form of the robustness degree. Open toolboxes are available from which the necessary robustness degree can be efficiently and easily computed [3]. Nevertheless, for many of the relevant constraints, the machinery of temporal logic could feel "over the top": Most of the relevant conditions could be formulated rather easily directly "by inspection". This is because they are typically of a relatively simple nature such that the powerful machinery of temporal logic may not be necessary.

Finally, the *GncTuner* tool was used in conjunction with the optimization-based guidance developments by tuning function parameters, which is a typical time-consuming manual tuning task. Here the tool was successfully used to inform the hand-tuning decision, in particular significantly improving on the original tuning of the constraint penalty parameters. It was thereby demonstrated that it can be a useful tool in the design and exploration of typically hard to chose parameterizations. However, care needs to be taken in the design of these optimization problems in order to ensure the validity of the found optimizers, i.e., it needs to be ensured that the optimization is well-posed and decidable within a reasonable amount of simulation experiments. Early results corresponding to the preliminary design phase have been reported in [14].

3.3 Worst-case Analysis on Benchmark

Finally, the *GncTuner* tool was successfully used to carry out a worst-case analysis. This was demonstrated in two formulations. The first is a straightforward antagonistic optimization, that is, within a domain of a certain size, uncertain parameters are optimized such as to yield the worst-possible outcome. The second is powered by constrained BO and aims to find a robustness radius, that is, the radius within which a certain level of performance can be maintained.

Figure 3 shows an example of such a robustness analysis, in which the task was to find the maximum perturbation radius of the space craft position w.r.t. the second waypoint in the GNC architecture maintaining a certain performance threshold. That means, the optimization aims to find the maximum radius in which the actual space craft position can be dispersed from the optimal one (the second waypoint) while maintaining acceptable landing performance. As the figure indicates, the landing performance is highly dependent on the direction of the perturbation, and the optimization identifies the worst-case direction leading to the final estimate of the robustness radius. Finally, the worst-case analysis can be judged in similar terms as the original optimization/parameter tuning task, that is, it finds more reliable results in reasonably low-dimensional settings with controllable influence of noise. Nevertheless, the algorithms (under suitable parameterization) provide good exploration and have not failed dramatically in any of the carried-out experiments.

4 ROBUST OPTIMIZATION-BASED GUIDANCE

The second major technical development in the AI4GNC project was to fuse methods from robust control engineering and optimization-based trajectory planning to take robustness into account on a higher level within the overall GNC architecture. Robust control engineering is already a conventional design approach in aerospace engineering as well as for autonomous vehicles but is primarily concerned with local robustness of the linear control system – for instance guaranteeing adequate tracking performance of a given reference trajectory under uncertainties in the system. However, for the generation of reference trajectories, which would, e.g., enable adequate higher-level system performance under such uncertainties hardly any extensions of robust control theory exist so far.

This alone justifies the interest in integrating techniques from robust control into optimization-based trajectory generation problems as it could promise improving overall system performance by tailoring reference generation, i.e., guidance functionality, to the uncertainties that affect the system. However, we consider robust trajectory generation to be especially relevant in the context where ML techniques are used to generate the dynamic system models or parts thereof. Indeed, if a model used for trajectory generation is generated by learning methods, then such a nonlinear model typically has large model uncertainty outside its data domain, reducing its applicability there. In these setups, it is then particularly interesting to employ uncertainty descriptions to discourage planning algorithms from exploiting apparent dynamics which have not been sufficiently explored. Here, robust trajectory generation is a prime candidate to improve the safety of ML methods in such a control architecture by taking these uncertainties into account and avoiding uncertain regions of the state space, if necessary.

To address these tasks we have developed a new method called *robust Differential Dynamic Programming* (robust DDP). As the name suggests, this method is an integration of techniques form robust control into the trajectory optimization method *Differential Dynamic Programming* [12]. This integration was done step by step as follows.

- **Primal robust DDP:** The first algorithm that was derived provides solutions for the robust trajectory optimization problem of nonlinear systems through an iterative procedure. These systems are expressed in the fashion of the usual $M-\Delta$ structure established in robust control as so-called *nonlinear generalized plants*, nonlinear systems in feedback interconnection with uncertain components. The algorithm then makes use of sequential linearizations to solve the robust trajectory optimization problem, requiring the solution of a sequence of small semi-definite programs (SDPs). As an outcome, the robust DDP generates an optimized reference trajectory as well as a time-varying linear feedback gain to track the reference. This first formulation of robust DDP utilizes primal synthesis SDPs, which are only convex for uncertain disturbances, but not for model uncertainties.
- **Dual robust DDP:** A core issue of the primal robust DDP formulation is that it only supports uncertain disturbances (like uncertain wind) and not model uncertainties. This is because in the latter case the SDPs necessary in the sequential convex programming scheme of the robust DDP are not convex. This limitation was overcome by establishing a *dual formulation* of the robust DDP. This allows for a convex formulation of both uncertain disturbances and model uncertainties.
- **Primal robust DDP with accelerated solver for online optimization:** As a final algorithmic development, the numerical performance of Robust DDP was addressed and successfully improved. One of the driving factors of the computational complexity of Robust DDP is the solution of a sequence of small LMI optimization problems in each sequential convex programming iteration. This is expensive if using off-the-shelf LMI solvers, which is why a *custom solver* has been developed leading to significant computational performance improvements. On top of that, we improved the line-search and regularization techniques for robust DDP.

Some of these results have been reported in [12]. Furthermore, along with the development of the theory for robust DDP, a series of numerical studies were carried out. This was done on two different simulation environments for simulating parafoil landing scenarios: The *Comparison Study* and the *Robust Performance Studies* have been carried out with standard kinematic parafoil models and hand-designed model uncertainties, while the *Benchmark Results* were generated using the simulation environment from SENER with the uncertainties present in this simulator corresponding to the standard Monte Carlo campaign.

4.1 Comparison Study

For our robust performance studies, we utilized a 6-DOF kinematic parafoil model which has been described in the literature in [2]. In the comparison study, we validate the adequacy of Differential Dynamic Programming as a trajectory generation tool for this model. This is done by comparing the landing accuracy for the 6-DOF model to the landing accuracy established in [8] for the same model. In this study, DDP showed superior performance.

4.2 Robust Performance Studies

Studies investigating the robust performance gain with robust DDP were also carried out. A key result that was shown in the simulations is that the robust trajectory generation chooses different paths than nominal DDP in order to reduce the effects of the uncertainty on the objective. This is exemplified in Figure 4, where it can be seen that the robust trajectories avoid areas of high uncertainty.

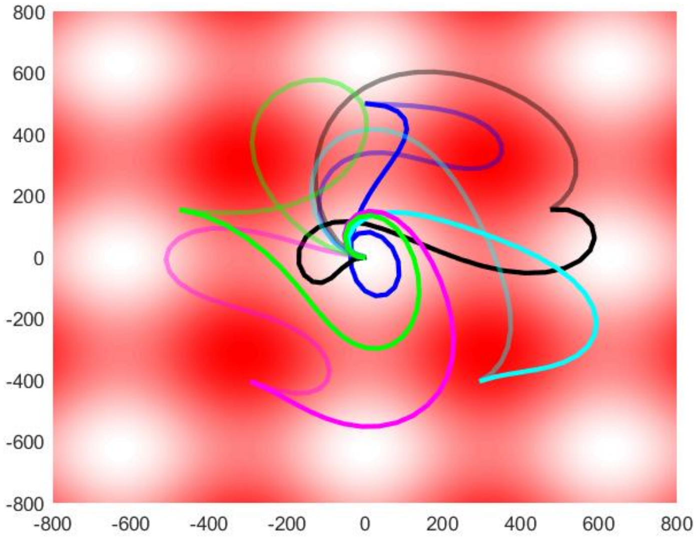

Figure 4: Illustration of robust and nominal trajectories generated by robust DDP for wind field uncertainty. Red shading indicates the spatially varying uncertainty, bold lines are robust trajectories for various initial conditions and transparent lines are nominal trajectories for the same initial conditions.

To demonstrate the potential of robust optimization for trajectory generation, our robust performance studies consider a vast number of simulations in nine distinct setups. These include landing scenarios with wind uncertainties with a custom wind field as suggested in [4], and the 3D wind field used in the SENER parafoil simulator. Furthermore, uncertainty in the parafoil velocity, input delays and combinations of model and wind uncertainties were studied. Generally, robust DDP improved the performance subject to these uncertainties. In addition to robust performance, the computation time of robust DDP was studied. Since this computation time was generally much longer than the computation time in nominal DDP, a custom solver has been developed, which reduced the computational demand by an order of magnitude.

4.3 Benchmark Results

The true envisioned challenge was to test the developed algorithms within the main parafoil simulator, which covers the parafoil landing scenarios including aerodynamic effects and various uncertain parameters. In its default setup, the simulator considers a "one-dimensional" wind field (mapping from only height to zonal and meridional winds). The considered baseline GNC solution to which we compared our results to was developed for this scenario. Here nominal DDP achieved comparable performances if compared to the baseline solution.

Subsequently, this benchmark was repeated with robust DDP instead of nominal DDP. In addition, a new scenario with a 3D wind field (dependence on all 3 spatial coordinates, instead of only height) and terrain constraints was introduced. In this second benchmark study, robust DDP outperforms the baseline on the 1D wind field while showing difficulties with the 3D wind field due to insufficient tuning. A final benchmark study was carried out on a cluster of 27 computers and therefore includes a much larger test set. Furthermore, nominal DDP and robust DDP were further tuned such that a clear ranking in performance was established: Robust DDP outperforms nominal DDP and nominal DDP outperforms the baseline. This can for example be seen in Table 1, which shows performance metrics for the 1D wind scenario and different intensities of the wind uncertainty for all three algorithms.

Table 1: Landing accuracy statistics of robust DDP, nominal DDP and the baseline guidance of the SENER simulator. The statistics are computed for a landing scenario with a 1D wind field and different intenstities of the wind uncertainty.

algorithm	wind error	90% quantiles	mean accuracy	median accuracy
Baseline	0m/s	27.8517	27.5154	16.1469
Baseline	$\frac{2.5}{3}$m/s	39.9353	34.9835	17.594
Baseline	$\frac{5}{3}$m/s	97.5034	59.1223	21.4595
Baseline	$\frac{7.5}{3}$m/s	185.8093	108.3531	26.7415
DDP	0m/s	24.7212	20.2993	11.7073
DDP	$\frac{2.5}{3}$m/s	38.3248	28.8612	14.4037
DDP	$\frac{5}{3}$m/s	82.8809	52.1305	20.4581
DDP	$\frac{7.5}{3}$m/s	145.9398	94.1244	27.9901
robust DDP	0m/s	21.4472	19.8503	11.6286
robust DDP	$\frac{2.5}{3}$m/s	24.3172	21.7991	12.7784
robust DDP	$\frac{5}{3}$m/s	49.698	38.3896	16.2578
robust DDP	$\frac{7.5}{3}$m/s	127.2579	80.0665	22.0501

Table 2: Landing statistics comparing robust DDP with a GP augmented model to nominal DDP with the unicycle model.

Statistics	GP augmented guidance	nominal guidance
90% Terminal Cost Quantile	0.31043	1.4633
Terminal Cost Median	0.060813	0.24776
Terminal Cost Mean	0.1607	0.69561
90% Position Error Quantile	15.2541	21.5877
Position Error Median	7.2427	8.8074
Position Error Mean	8.2827	11.0035
90% Heading Error Quantile	4.6454	19.4048
Heading Error Median	0.77642	6.4613
Heading Error Mean	2.0768	9.0401

4.4 Augmented ML-models

Finally, also the utilization of robust DDP with ML-enhanced dynamics models for the parafoil was suggested. As example, a 4 DOF parafoil model was augmented with a Gaussian process and this augmented model was integrated into robust DDP. The augmented 4-DOF model was trained to match the 6-DOF model from [2] and the variance output of the GP model has been used to define the uncertainty for robust DDP. (Note that such an uncertainty encourages the optimization algorithm to not leave the data distribution on which the model has been generated.) This robust DDP guidance based on the augmented GP model was compared to nominal DDP with the 4DOF model in Table 2. Here, the GP augmentation improved the performance of the guidance significantly effectively exploiting the learning behaviour together with its uncertainty description.

5 LEARNING-BASED MODEL AUGMENTATION

The third major technical development in the AI4GNC project was to use data – from a high-fidelity simulation environment or flight data – to improve baseline models of the flight dynamics by ML techniques. Such models can serve as the basis of developing high-performance and reliable next-generation GNC systems. In ‘conventional’ GNC designs, the models of the flight dynamics for guidance and control are often relatively simple, directly derived from first principle laws of physics and simplified aerodynamic relations and therefore represent somewhat idealized nominal characteristics, which are often inaccurate to represent the real, personal characteristics of each vehicle. On top of that, modelling from first principles gets prohibitively complex, in particular when considering (partially) closed-loop behaviours of the system, i.e., when aiming to describe the system when potentially complex software is in the loop. However, exactly such closed-loop behaviours are highly relevant for (autonomous) higher-level decision-making or guidance, presenting a significant challenge for which data-driven and ML-powered system identification can provide deployable solutions.

In most applications, a priori models based on first principles or engineering insights are readily available and provide valuable information for GNC design. The approach explored in this study is, hence, to combine such prior knowledge of the system with powerful machine-learning techniques to *enhance* the resulting model in terms of accuracy and usability, while maintaining or even increasing its interpretability. For this reason, a learning-based model-augmentation framework has been proposed within the project and studied on the available benchmarks.

The key idea within the proposed approach is to combine an existing model of the system with static or dynamic nonlinear mappings in the form of artificial neural networks (ANNs). The interconnection that realizes the learning-based augmentation of the existing model is in the form of a *linear fractional representation* (LFR), which opens up the possibility to systematically analyze the resulting augmented models using, e.g., robust techniques including *integral quadratic constraints* (IQCs). Moreover, as the LFR form is a widely used representation for *linear parameter-varying* (LPV) systems – which is a commonly used framework for control synthesis and analysis of complex nonlinear systems such as aerospace vehicles – we can combine multiple existing powerful tools by this ML component. Hence, the envisioned outcome was a flexible ML-based grey-box identification framework which can later be used for analysis and control. To build up this framework, we addressed the following aspects:

- **State-of-the-art LPV identification and control:** The LPV framework combines linear time-invariant (LTI) theory with surrogate models to tackle modelling of, control design for, and analysis of complex nonlinear systems. We applied the current state-of-the-art methods in a simulation setting to demonstrate the strength and capabilities of the framework.
- **Dynamic and static ANNs**: One of the key elements was to use ML tools to improve current GNC capabilities. We used universal approximators in terms of ANNs, both static and dynamic structures, to complete baseline dynamic models in terms of augmentation.
- **Model-augmentation framework:** In order to combine the two aforementioned aspects in a systematic way, we developed a learning-based model-augmentation approach parallel to the project. We have also implemented our developed methods in the deepSI toolbox[3], which is a flexible software stack built for ANN-based model learning.
- **Accurate open-source 6DoF/12DoF simulator:** In order to easily generate high-quality data for learning, we developed a simulator for *Generic Parafoil Return Vehicles* (GPRVs), based on advanced flight dynamics model that describes the 12 DoF motion of the vehicle and its parafoil attachment including relative motion of these bodies with respect to each other and flexibility of the tension lines. A flight controller for the GPRV has been also designed in terms of a reference-tracking LPV controller, which tracks the trajectory that is generated by high-level path planners such as the one implemented in the SENER simulator.

Next, we discuss the numerical studies and results that have been accomplished to demon-

[3] `github.com/GerbenBeintema/deepSI`

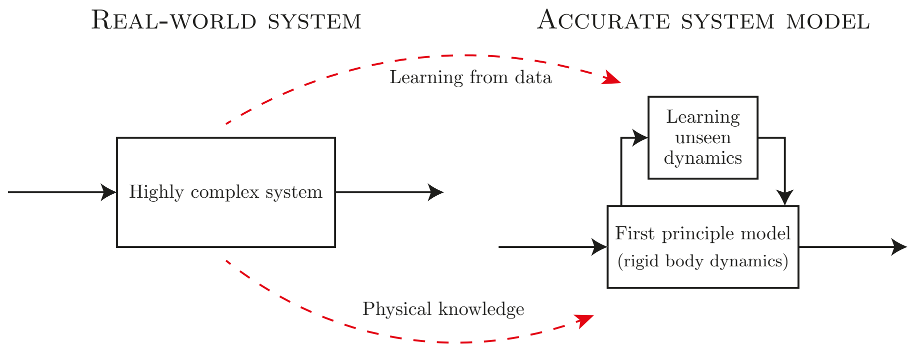

Figure 5: Learning-based model-augmentation concept based on an LFR-ANN structure.

strate the efficiency of the proposed tools. These can be roughly divided into three objectives; 1) Demonstration of state-of-the-art LPV techniques on a Control Moment Gyroscope (CMG), 2) Demonstration of the model-augmentation capabilities on the GPRV, 3) Towards the integration of model-augmentation-based models in guidance algorithms.

5.1 Demonstration of state-of-the-art LPV techniques on the CMG

Within this study case, we have shown that, even without powerful ML tools, it is possible to achieve highly accurate models and high-performance control of a complex system such as a CMG. For this, we used a high-fidelity simulation model of a CMG with white measurement noise. We have showcased and compared several LPV identification approaches to model the CMG dynamics. Based on the identified models, we designed LPV motion controllers for reference tracking with the CMG. We also designed LPV controllers based on analytical models derived based on the implemented dynamics in the simulator. For both approaches, we achieved similar, high tracking and disturbance rejection performance with the designed controllers, showing that data-based modelling is capable to substitute first-principles models in terms of performance of the controllers designed on these models, while identification can even capture from data unknown, difficult to model dynamic components.

5.2 Demonstration of the model-augmentation capabilities on the GPRV

In the next study case, we have demonstrated our developed learning-based model-augmentation framework based on the LFR-ANN concept. Figure 5 depicts a generalized structure of this framework. The core estimation approach, capable to provide augmentation in an LFR-ANN form, is called SUBNET, which allows state-of-the-art learning tools for efficient training in terms of a subspace-encoder for state estimation, automatic data normalisation and model scaling, early stopping, batch training, etc. Here we used a linearized version of the complex GPRV dynamics as a baseline model to be augmented by static and dynamic ANN components. The resulting models showed excellent performance. Next to the demonstration of the techniques on the GPRV, we also gave some preliminary results on learning-based model augmentation applied to the CMG. The framework is built up such that with the LFR-based representations, we can extract the activation functions to perform robust analysis, which is a promising future direction and is seen as a key element of the next generation of control design and analysis toolchains within the aerospace sector. Further details are found in [15], [10], [13].

5.3 Towards the integration of model-augmentation-based models in the guidance algorithms

In view of using the augmented models in the GNC structure, we also considered an augmentation study where the baseline unicycle model used in the current guidance level of the studied GPRV is augmented in an LFR-ANN form to describe the closed-loop flight dynamics of the vehicle. Based on data, this would allow characterising the additional dynamics of the controlled system that are not part of this ideal motion model and could be potentially used to characterise the uncertainty of the ideal model in the guidance.

Although we obtained good results, several challenges had to be overcome. For considering only flight trajectories over a long horizon, we observed that the ANN-based augmentation is likely to become over-trained such that it learns the trajectory itself, instead of the dynamics. For a shorter horizon, augmentation of the unicycle model in combination with minor excitation of the dynamics data does not reveal the complex underlying dynamics of the true vehicle, which yielded undesired training results. To obtain informative data on the dynamics, 15 % perturbations of the nominal dynamics were used and the data was segmented into several estimation windows. This allowed in a velocity-based training setting to achieve good results in terms of model-augmentation, showing that proper experiment design resulting in informative data together with proper choice of the estimation horizons is essential in identification and model augmentation in general and despite of the current state of automation of learning processes, still requires expert choices from the user to unlock their true potential.

6 SUMMARY

In summary, we look back on a highly successful technology development activity in which the systematic application of machine learning (ML) techniques to GNC system was investigated. Figure 6 recounts an early scheme developed within the project; it shows a preliminary development plan of the AI-enhanced GNC systems, including a retrospective analysis of realized components during the project. Altogether, the evaluated techniques have great potential to increase the capabilities of next-generation GNC systems, in particular in directions of increased performance and autonomy, as well as to streamline the GNC design process and decrease overall design effort. Specifically, for the three main focal points, we conclude the following:

- **Bayesian Optimization (BO) & Temporal Logic.** BO can be a useful tool for GNC design and analysis. It has shown good performance and high flexibility in its combination with temporal logic expressions, which allowed expressing all relevant requirements within the project for direction consideration in the optimization. While the reliability of the tool seems to be insensitive to many parameter choices, it is somewhat dependent on the exact problem formulation and can lose reliability in higher dimensional spaces.
- **Optimization-based Robust Guidance.** Robust Differential Dynamic Programming is a powerful approach for generating trajectories for nonlinear generalized plants (nonlinear M-Δ structures) that take the uncertainty into account in the trajectory generation. This has resulted in a novel sequential convex programming method based on solving a sequence of LMI optimization problems. This has shown to consistently improve the robust performance of trajectory generation over the nominal case in an wide variety of scenarios. The results provide a promising outlook on the integration of ML models in GNC systems since robust DDP can take into account the uncertainties of learned models.
- **Learning-Based Model-Augmentation.** The results produced during this activity

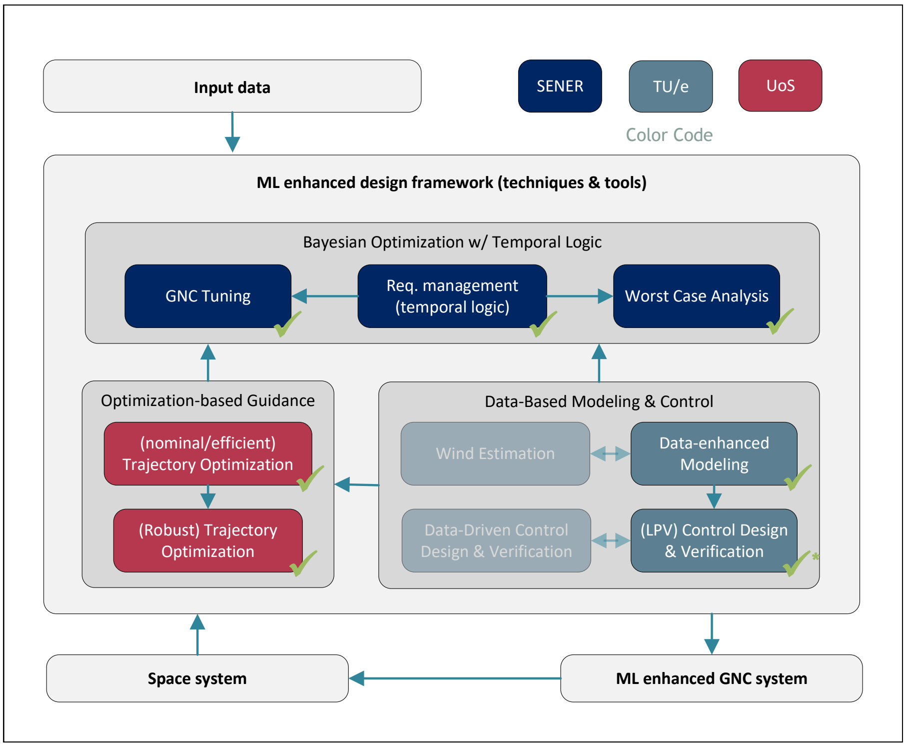

Figure 6: Review of ML-enhanced design framework. Green ticks indicate completed topics within the study, while greyed-out components were ultimately not realized to a significant extent.

have demonstrated that learning-supported system identification in terms of learning-based model augmentation is a valuable asset for capturing potentially complex system dynamics. This includes the case where data is produced from high-fidelity simulations. While the potential of ML-based techniques was demonstrated, generally more established techniques, such as LPV or LTI identification, can provide similarly good performance depending on the operating range of the system to be captured. Additionally, the performance of the learning-based model augmentation can depend significantly on choices in the model structure and the quality of the data present. Automation of the former and, for the latter, providing useful guidelines for aerospace applications, are currently seen as important development directions for deployment of these tools in aerospace engineering.

REFERENCES

[1] G. E. Fainekos and G. J. Pappas, "Robustness of temporal logic specifications for continuous-time signals", *Theoretical Computer Science*, vol. 410, no. 42, pp. 4262–4291, 2009.

[2] B. J. Rademacher, P. Lu, A. L. Strahan, and C. J. Cerimele, "In-Flight Trajectory Planning and Guidance for Autonomous Parafoils", *Journal of Guidance, Control, and Dynamics*, vol. 32, no. 6, pp. 1697–1712, Nov. 2009. DOI: 10.2514/1.44862.

[3] Y. Annpureddy, C. Liu, G. Fainekos, and S. Sankaranarayanan, "S-TaLiRo: A tool for temporal logic falsification for hybrid systems", in *Tools and Algorithms for the Construction and Analysis of Systems*, Springer Berlin Heidelberg, 2011, pp. 254–257.

[4] S. Yang, N. Wei, S. Jeon, R. Bencatel, and A. Girard, "Real-time optimal path planning and wind estimation using gaussian process regression for precision airdrop", in *2017 American control conference (ACC)*, IEEE, 2017, pp. 2582–2587.

[5] A. Gleave, M. Dennis, C. Wild, N. Kant, S. Levine, and S. Russell, "Adversarial Policies: Attacking Deep Reinforcement Learning", *Proceedings of ICLR*, 2020.

[6] M. Neumann-Brosig, A. Marco, D. Schwarzmann, and S. Trimpe, "Data-efficient autotuning with bayesian optimization: An industrial control study", *IEEE Tr. on Control Systems Technology*, vol. 28, pp. 730–740, 2020. DOI: `10.1109/TCST.2018.2886159`.

[7] A. Figueroa-González, F. Cacciatore, and R. Haya-Ramos, "Landing guidance strategy of space rider", *Journal of Spacecraft and Rockets*, vol. 58, no. 4, pp. 1220–1231, 2021.

[8] A. Leeman, V. Preda, I. Huertas, and S. Bennani, "Autonomous parafoil precision landing using convex real-time optimized guidance and control", *arXiv:2105.00715*, 2021.

[9] R. Turner, D. Eriksson, M. McCourt, *et al.*, "Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020", in *NeurIPS 2020 Competition and Demonstration Track*, 2021.

[10] M. de Lange, "Linear parameter-varying control for the atmospheric flight of a parafoil guided return vehicle", M.S. thesis, Eindhoven University of Technology, 2022.

[11] L. P. Fröhlich, C. Küttel, E. Arcari, L. Hewing, M. N. Zeilinger, and A. Carron, "Contextual tuning of model predictive control for autonomous racing", in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.

[12] D. Gramlich, C. W. Scherer, and C. Ebenbauer, "Robust differential dynamic programming", in *IEEE Conference on Decision and Control (CDC)*, 2022.

[13] M. H. de Lange, C. Verhoek, V. Preda, and R. Tóth, "LPV Modeling of the Atmospheric Flight Dynamics of a Generic Parafoil Return Vehicle", in *Proc. of the 5th IFAC Workshop on Linear Parameter-Varying Systems (LPVS)*, vol. 55, 2022, pp. 37–42.

[14] R. Polonio, C. Ardura, V. Preda, and L. Hewing, "Bayesian optimization with temporal logic constraints for automated gnc tuning", in *European Conference for Aeronautics and Space Sciences (EUCASS)*, 2022.

[15] C. Verhoek, G. I. Beintema, S. Haesaert, M. Schoukens, and R. Tóth, "Deep-Learning-Based Identification of LPV Models for Nonlinear Systems", in *Proc. of the 61st IEEE Conference on Decision and Control (CDC)*, 2022, pp. 3274–3280.

[16] T. Wang, A. Gleave, T. Tseng, N. Belrose, *et al.*, "Adversarial Policies Beat Superhuman Go AIs", *arXiv preprint arXiv:2211.00241*, 2023.

[17] X. Wu, W. Guo, H. Weil, and X. Xing, "Adversarial Policy Training against Deep Reinforcement Learning", *In the Proceedings of the* 30^{th} *USENIX Security Symposium*, 2023.