Self-Confidence in Autonomous Systems
Papers from the AAAI 2015 Fall Symposium

Trusting Learning Based Adaptive Flight Control Algorithms

Maximilian Mühlegg and **Florian Holzapfel**
Institute of Flight System Dynamics
Technische Universität München
85748 Garching bei München, Germany
maximilian.muehlegg@tum.de
florian.holzapfel@tum.de

Girish Chowdhary
School of Mechanical and Aerospace Engineering
Oklahoma State University
Stillwater, OK, 74078
girish.chowdhary@okstate.edu

Introduction

Autonomous unmanned aerial systems (UAS) are envisioned to become increasingly utilized in commercial airspace. In order to be attractive for commercial applications, UAS are required to undergo a quick development cycle, ensure cost effectiveness and work reliably in changing environments. Learning based adaptive control systems have been proposed to meet these demands. These techniques promise more flexibility when compared with traditional linear control techniques. However, no consistent verification and validation (V&V) framework exists for adaptive controllers. The underlying purpose of the V&V processes in certifying control algorithms for aircraft is to build trust in a safety critical system. In the past, most adaptive control algorithms were solely designed to ensure stability of a model system and meet robustness requirements against selective uncertainties and disturbances. However, these assessments do not guarantee reliable performance of the real system required by the V&V process. The question arises how trust can be defined for learning based adaptive control algorithms. From our perspective, self-confidence of an adaptive flight controller will be an integral part of building trust in the system. The notion of self-confidence in the adaptive control context relates to the estimate of the adaptive controller in its capabilities to operate reliably, and its ability to foresee the need for taking action before undesired behaviors lead to a loss of the system. In this paper we present a pathway to a possible answer to the question of how self-confidence for adaptive controllers can be achieved. In particular, we elaborate how algorithms for diagnosis and prognosis can be integrated to help in this process.

Adaptive Flight Control

Adaptive control algorithms aim at making a system behave like a-priori chosen dynamics. This is achieved by approximating the modeling uncertainty using a weighted combination of known basis functions. The underlying assumption is that there exists an optimal set of gains that best captures the modeling uncertainty for a particular structure of the adaptive element. If the gains do converge to their ideal values and the uncertainty is entirely canceled, the closed loop system matches the chosen model and recovers its properties. Most algorithms try to achieve this goal by adjusting the parameters based on the minimization of a quadratic cost. A Lyapunov based stability proof ensures boundedness of all system signals (see e.g. (Tao 2003)).

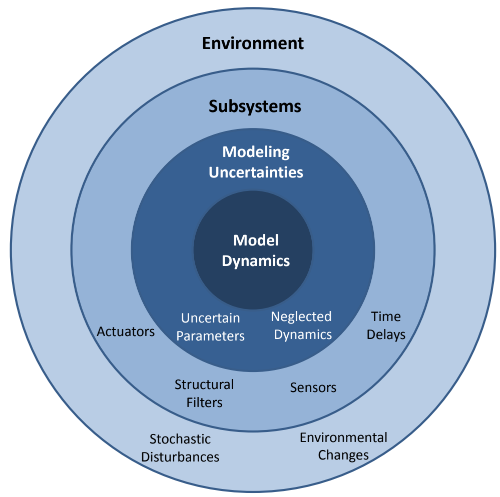

Figure 1: Important areas for adaptive control design.

Aspects Preventing Successful Verification and Validation of Adaptive Flight Control Algorithms

Verification and Validation are key processes in the certification of products. They are intended to show a systems' compliance with requirements and specifications. In the context of control algorithms, this includes guarantees about the performance and behavior of a controlled system. Thus, from an abstract point of view it is intended to generate trust in a successful operation.

Several shortcomings prevent a successful V&V required in certification of adaptive flight control algorithms. First of all, guarantees for the convergence of the adaptive weights for systems subject to disturbances do not exist. In particular, in the case that the uncertainty is approximated by an over-determined control structure, weight trajectories cannot be predicted. Even though the Lyapunov based stability ensures boundedness of all system signals, these bounds are often utterly conservative and physically unreasonable and still might cause a system to crash due to physical constraints. Furthermore, most adaptive control developments only take a part of the system, shown in Fig. 1, into account; in particular, subsystems and effects from the environment are often neglected. Hence, the stability proof might not even hold when transported to an experimental system.

A crucial point in the V&V process of linear flight controllers are analytically derived performance and robustness guarantees in the form of metrics such as phase and gain margin. Unfortunately, such universal metrics do not exist for nonlinear controllers. Hence, an a-priori analytical evaluation of the adaptive controller is not possible with the current state-of-art.

Role of Diagnosis and Prognosis for Health Monitoring of Learning Based Adaptive Control Algorithms

There seems to be an emerging consensus between various authors that if an a-priori analytical evaluation of the adaptive controller is not possible, monitoring and health assessment during its run-time will be an integral part of verifying adaptive controllers in the future (see e.g. (Jacklin 2008)). Aligned with this direction of thought, we propose a pathway to monitoring and health assessment inclusive of two steps, diagnostics and prognosis.

Diagnostics involve quantifying a belief in the current uncertainty approximation capabilities of the adaptive controller and identify undesired behavior, such as excessive weight oscillations or parameter drift, from onboard available data. Online monitoring approaches of the control algorithm can help gauge self-confidence in its current performance. Existing approaches rely on data driven techniques such as Bayesian regression. They try to determine whether the adaptive weights have converged and the predictive covariance is sufficiently small (Schumann and Yan Liu 2007), or if a significant change in the control output from previously observed behavior occurred (Szukalski, Mühlegg, and Johnson 2015). Using the data set D and the current state z, the latter is achieved by inferring a probability P_ζ from a predictive distribution $p(y|z,D)$ whether the adaptive output y is within a ζ threshold of the most probable output y_{MP}:

$$P_\zeta = \int_{-\infty}^{y_{MP}+\zeta} p(y|z,D)dy - \int_{-\infty}^{y_{MP}-\zeta} p(y|z,D)dy. \quad (1)$$

Markovian models are used for formal verification purposes and pose an interesting extension to the online monitoring algorithms (Cámara 2013).

The second step deals with prognosticating future behavior. The goal is to answer the question whether the system still abides stability and performance requirements for a fixed future time horizon given the current circumstances.

One approach adopts the strategy of model predictive control (Wang 2009): the model used for the control design is employed along diagnostic results in order to predict future states. Each region is associated with a probability of occurrence, thus allowing for a statistical quantification of how likely a state limit transition is.

Recently, Kullback-Leibler (KL) divergence has attracted increased attention as a metric, which measures how the belief in a system changes over time. The underlying idea is to measure the difference between the posterior P and the prior distribution Q in a Bayesian setting (Axelrod 2015):

$$D_{KL}(P\|Q) = \int_{-\infty}^{\infty} p(x) ln \frac{p(x)}{q(x)} dx. \quad (2)$$

Therefore, it also presents an interesting approach to try and predict future changes in the uncertainty estimation.

Neither diagnostics nor prognosis alone will be able to allow assertions about the health of the controller. Rather, a measure of confidence will arise from a combination of both areas. Additionally, the intrinsic properties of different adaptive control methods can help build self-confidence. A guaranteed convergence of the adaptive weights to a close vicinity of their true values can be related to a reduction in the predictive covariance (see e.g. (Chowdhary, Mühlegg, and Johnson 2014)). Also the technique itself might already supply some sort of self-confidence (see e.g. Bayesian techniques in (Chowdhary et al. 2014)). Our ongoing work includes extensions to these techniques such that either a measure of the learning confidence can be leveraged by the system to navigate to safe regions of the state space or switch controllers, ultimately reducing the probability of a system loss.

References

Axelrod, A. M. 2015. Learning to exploit time-varying heterogeneity in distributed sensing using the information exposure rate. Master's thesis, Oklahoma State University.

Cámara, J. 2013. *Assurances for self-adaptive systems: Principles, models, and techniques*, volume 7740. Berlin and New York: Springer.

Chowdhary, G.; Kingravi, H. A.; How, J. P.; and Vela, P. A. 2014. Bayesian nonparametric adaptive control using gaussian processes. *IEEE Transactions on Neural Networks and Learning Systems* PP(99).

Chowdhary, G.; Mühlegg, M.; and Johnson, E. 2014. Exponential parameter and tracking error convergence guarantees for adaptive controllers without persistency of excitation. *International Journal of Control* 87(8):1583–1603.

Jacklin, S. 2008. Closing the certification gaps in adaptive flight control software. In *AIAA Guidance, Navigation, and Control Conference*.

Schumann, J., and Yan Liu. 2007. Tools and methods for the verification and validation of adaptive aircraft control systems. In *IEEE Aerospace Conference*, 1–8.

Szukalski, R.; Mühlegg, M.; and Johnson, E. 2015. Consistency monitoring of adaptive controllers using bayesian linear regression. In *American Control Conference*, 177–182.

Tao, G. 2003. *Adaptive control design and analysis*. Hoboken and N.J: Wiley-Interscience.

Wang, L. 2009. *Model Predictive Control System Design and Implementation Using MATLAB®*. London and Heidelberg: Springer.