PERSPECTIVES ON STOCHASTIC PREDICTIVE CONTROL WITH AUTONOMOUS MODEL ADAPTATION FOR MODEL STRUCTURE UNCERTAINTY

Tor Aksel N. Heirung and Ali Mesbah*

Department of Chemical and Biomolecular Engineering, University of California, Berkeley, CA 94720

Abstract

Integrated stochastic optimal control and system learning to simultaneously reduce parametric and model structure uncertainty can create new avenues for achieving high-performance operation of uncertain systems using model predictive control. This paper presents a generic framework for stochastic optimal control with integrated (control-oriented) model structure adaptation, and discusses general solution methods and key research issues associated with the framework. The potential advantages of the proposed framework include autonomous maintenance of model predictive controllers as well as handling of multiple system models under closed-loop conditions, for example, for fault-tolerant control applications and dealing with systems with intrinsically uncertain dynamics.

Keywords

Stochastic optimal control; Dual control; Model structure uncertainty

Introduction

This paper explores the notion of integrated learning and model predictive control (MPC) for stochastic systems with uncertain model structures. The position taken by this work is motivated by the problem of parametric model uncertainty in model-based control, which has inspired important research directions in the fields of stochastic adaptive control (Wittenmark, 1975b) and identification for control (Gevers, 2005; Hjalmarsson, 2005). The techniques in these fields generally rely on model adaptation in closed loop to reduce model uncertainty under normal operation. Most model-based control design approaches are based on the *certainty equivalence* (CE) principle (Bar-Shalom and Tse, 1974), which involves separately designing the controller and performing parameter estimation such that merely the point estimates of the parameters are used in the controller as if the estimates were exact. Except for specific cases, adaptive controllers based on the CE principle are suboptimal since they do not account for model uncertainty (Åström and Wittenmark, 1995). This can potentially result in severe problems such as parameter drift, bursting, and loss of stability (Wittenmark, 1975b).

When system uncertainty affects the control performance, the uncertainty is *control relevant.* A *cautious controller* explicitly incorporates the current knowledge of uncertainty into the control design, so that increased system uncertainty can lead to more cautious control (Åström and Wittenmark, 1995). One consequence of this approach is that large uncertainties may cause vanishingly small control inputs, which in turn reduce the amount of information generated in the closed loop and further increase the model uncertainty, exacerbating the problem. This phenomenon can be attributed to the accidental learning in CE and cautious controllers, both of which are *passively adaptive feedback* control approaches.

On the other hand, *actively adaptive feedback* controllers have a probing feature, the purpose of which is to actively learn about the system (Tse and Bar-Shalom, 1973). Probing the system is instrumental in handling model uncertainty when the uncertainty is *reducible*, that is, when the control inputs affect not only the system states but also the future uncertainty of the states. This is known as the *dual effect* of control inputs (Bar-Shalom and Tse, 1974). The dual control problem (Feldbaum, 1961) involves solving a stochastic optimal control problem (OCP) that relies on a single model structure with reducible, control-relevant parametric uncertainty. Dual control relies on Bellman's principle of optimality (Bellman, 1957) to seek an optimal synergy between probing the system for model improvement (i.e., reducing parametric uncertainty) and optimal control based on the current uncertain model.

This paper adopts the dual control paradigm for

*Corresponding author: *mesbah@berkeley.edu*

model structure uncertainty in the context of model predictive control. The proposed dual control framework relies on competing model structures to account for the existence of multiple model hypotheses. Including multiple models into an OCP offers the flexibility of using different model structures when a system transitions from one mode of behavior to another (e.g., due to system faults), or when the model structure is not known a priori. In addition, online model structure adaptation is likely to reduce the modeling effort during controller development since a more complex system model (e.g., valid over a wider operating range) can be replaced with a set of relatively simpler models in the OCP. This is particularly advantageous in light of the costs associated with system modeling, which can account for up to 80 % of resource expenditure in the design of model-based controllers (Sun et al., 2013). MPC with autonomous model adaptation can sustain high-performance operation of complex systems over extended times with possibly no user intervention, a critical consideration for reliable operation of high-precision and safety-critical systems.

Recent years have witnessed growing interest in the notion of adaptive control based on multiple models (e.g., see Narendra and Han (2011) and the references therein). The multiple-model adaptive control methods commonly entail switching between the models, meaning the control inputs are at any given time computed using one model without active discrimination between the competing models. In this work, the proposed dual control framework for MPC with multiple models is intended to actively probe the system for reducing future parametric and model structure uncertainty, while achieving the best attainable control performance with the most appropriate model structure at present. For a general class of stochastic systems described by a set of uncertain models, this paper presents the stochastic optimal control problem with dual effect for controlling the states and the future uncertainty of the states. The connection between the stochastic OCP and stochastic MPC (SMPC) is briefly discussed in light of the recursion of the Bellman equation in stochastic dynamic programming (DP). We then outline some of the key challenges associated with solving the stochastic OCP with dual effect, and discuss solution methods that can result in tractable SMPC approaches with autonomous model adaptation.

Notation. $\mathsf{P}[\,\cdot \mid A]$ denotes conditional (joint) probability distribution given A. $\mathrm{E}[\,\cdot \mid A]$ denotes conditional expected value given A. $\Pr[B]$ denotes the probability of event B.

The Dual Control Problem with Multiple Models

Consider a general dynamical system that is described by a set of n_m discrete-time model structures $m^{[i]}$ (e.g., due to model structure uncertainty, or using different model structures for different modes of system behavior)

$$m^{[i]}\ : \bar{x}_{t+1}^{[i]} = f^{[i]}(\bar{x}_t^{[i]}, u_t, w_t^{[i]}, \theta_t^{[i]}), \quad \forall\, i \in \mathbb{Z}_{[1,n_m]}, \qquad (1)$$

where t is the time index; the superscript $[i]$ denotes the model index; $\bar{x}_t$ denotes the states of each model structure, with x being the true system states; u_t denotes the system inputs; θ denotes the model parameters; w_t is a sequence of stochastic process noise with known statistics; and f represents the (possibly nonlinear) model equations. Let $\mathcal{Y}_k$ denote the system observations at sampling time k. Define the *hyperstate* ξ_k as the conditional joint probability distribution of the system states given the observations $\mathcal{Y}_k$, i.e., $\xi_k \triangleq \mathsf{P}[x_k|\mathcal{Y}_k]$. In fact, the hyperstate is a state variable that provides a statistical representation of the system uncertainty.

Given ξ_k at sampling time k, the stochastic OCP with multiple models (1) can be stated as

$$\min_{\boldsymbol{\pi}}\ J_k(\xi_k) = \mathrm{E}\left[\sum_{j=k}^{k+N-1} \ell_j(\bar{x}_j^{[i]}, \pi_j) + \ell_{k+N}(\bar{x}_{k+N}^{[i]}) \middle| \xi_k\right] \qquad (2a)$$

$$\text{s.t.: multiple models (1)}, \qquad \forall\, t \in \mathbb{Z}_{[k,k+N-1]} \qquad (2b)$$

$$\pi_t \in \mathbb{U}, \qquad \forall\, t \in \mathbb{Z}_{[k,k+N-1]} \qquad (2c)$$

$$\Pr[g^{[i]}(\bar{x}_t^{[i]}) \leq 0] \geq \beta^{[i]}, \qquad \forall\, t \in \mathbb{Z}_{[k+1,k+N]}, \qquad (2d)$$

where $\boldsymbol{\pi}$ denotes a vector of N-stage control policy, each element of which is defined by $u_t = \pi_t(\cdot)$; ℓ_j denotes the stage-wise cost at the jth stage of control; and (2d) represents state chance constraints with g being a (possibly) nonlinear function and $1-\beta$ the maximum admissible probability of constraint violation.

According to the principle of optimality, the optimal cost function $J_k^*(\xi_k)$ (i.e., *cost-to-go*) must satisfy the recursion of the *Bellman equation* (Bellman, 1957)

$$J_j^*(\xi_j) \triangleq \min_{\pi_j} \mathrm{E}\left[\ell_j(\bar{x}_j^{[i]}, \pi_j) + J_{j+1}^*(\xi_{j+1})|\xi_j\right], \qquad \forall j \in \mathbb{Z}_{[k,k+N-1]}, \qquad (3)$$

with the initial condition $J_{k+N}^*(\xi_{k+N}) = \ell_{k+N}(\bar{x}_{k+N})$. Eq. (3) implies that the expected control cost at each stage k is evaluated through accounting for the system uncertainty ξ_j at future stages $j \geq k$. Thus, the control inputs will not only influence the system states but also the state uncertainty, suggesting that the control inputs have dual effect.

The concept of recursion of the Bellman equation (3) is closely related to the choice of the control policy $\boldsymbol{\pi}$ in the stochastic OCP (2) (Tse and Bar-Shalom, 1973). The dual effect of the control inputs largely relies on incorporating the (probabilistic) knowledge of future system observations into $\boldsymbol{\pi}$. Using a *closed-loop control policy* in (2), which considers the knowledge that the loop is closed throughout the horizon $[k,\ k+N-1]$, allows for anticipation of the future system observations through $(N-1)$-step recursion in (3). The causal anticipation of future observations will enable the closed-loop control policy to actively probe the system for new information. Hence, in addition to affecting the system states, closed-loop control policies will affect the uncertainty of future states (i.e., dual effect). Simultaneously, closed-loop control policies will intelligently exercise *caution* in accounting for uncertainty since the control inputs know that future system observations will be available and can inform corrective actions. Another key property of closed-loop control policies is that the active learning occurs to the extent dictated by the closed-loop control performance. This implies that closed-loop control policies intrinsically seek control-oriented model adaptation.

On the contrary, a *feedback control policy* uses only the prior system observations $\mathcal{Y}_k$, gathered up until k, independently of the knowledge of future observations.[1] This implies that a feedback control policy does not involve recursion of the Bellman equation (3), which from a computational standpoint is appealing for solving (2) but it is suboptimal since the optimality of the control cost over the future stages is neglected (Bertsekas, 2005). In contrast to closed-loop policies that are actively adaptive, feedback control policies are passively adaptive in that they do not enable active learning of the system since they do not anticipate the future observations (Tse and Bar-Shalom, 1973).

Approximate Solution Methods for the Dual Control Problem

Stochastic DP is a natural approach to solving the dual control problem (Feldbaum, 1961), including the extension to the case of model structure uncertainty in (2). However, the *curse of dimensionality* (Bellman, 1957) generally renders DP intractable for problems of practical interest since the computational requirements grow exponentially with the state space. Broadly speaking, approximate solutions to the original dual control problem for parametric uncertainty include (Filatov and Unbehauen, 2000): (i) *implicit* dual control that involves direct approximation of the recursion of the Bellman equation (3), and (ii) *explicit* dual control that involves reformulation of the dual control problem to a tractable (stochastic) optimal control problem with some form of dual control effect. Implicit dual control essentially relies on numerical approximation of the dynamic programming problem through approximating the solution to the Bellman equation only in limited regions of the hyperstate ξ_j (e.g., Lee and Lee, 2009; Bayard and Schumitzky, 2010). This yields a control policy that is a function of approximate cost-to-go functions. However, approximate DP approaches typically involve recursion of the Bellman equation over N stages, which can be computationally formidable. This has motivated explicit approaches to dual control, which seek to recast the dual control problem as a tractable OCP that does not entail recursion of the Bellman equation but the control inputs still retain some form of cautious and probing features (Wittenmark, 1975a). In the context of MPC, most common approaches to explicit dual control under parametric uncertainty involve adding probing effect to the control inputs. Generally speaking, this is done either by directly adding perturbation signals to the control inputs (e.g., Tanaskovic et al., 2014; Marafioti et al., 2014), or by reformulating the OCP such that it incorporates some measure of model uncertainty over the future stages (e.g., Larsson et al., 2013; Heirung et al., 2015). The addition of the probing effect to control inputs inevitably results in control performance loss due to undesired system perturbations. However, the improved system learning facilitated by probing is expected to lead to better control performance over the future control stages and, as a result, decrease the overall performance loss that would otherwise be incurred due to large system uncertainty in case no learning takes place.

Inspired by explicit dual control approaches for parametric uncertainty in the context of predictive control, we propose reformulating the stochastic OCP (2) with dual effect to an OCP with integrated input design for discrimination between the competing model structures (1). This will lead to a SMPC framework with the endogenous capability of learning about the model structure uncertainty (e.g., see Mesbah (2016) for the general description of SMPC). Input design for model discrimination seeks to ensure that outputs of the competing models are sufficiently far apart so that system measurements can invalidate all but one model to the largest extent possible (Atkinson and Cox, 1974). For the stochastic models (1), model discrimination hinges on quantifying the distance between outputs of the competing models in a probabilistic sense. Various measures exist for quantifying the distance between probability distributions (Gibbs and Su, 2002). To this end, a natural choice is the Kolmogorov distance (Kolmogorov, 1933), which is the maximum difference between two cumulative distribution functions (cdfs); the Kolmogorov distance thus converges to the maximum value of one as the two corresponding probability distributions approach complete separation. A key issue in optimal control with multiple models is ensuring the feasibility of the optimization problem. Requiring all models be feasible on the entire prediction horizon can lead to feasibility issues and overly conservative control. One approach to addressing the latter issue is to propagate the least likely model(s) on a shorter horizon than that of the most likely model(s). Alternatively, each model can have an associated sequence of control inputs, the first elements of which are identical (i.e., nonanticipativity in stochastic programming).

The performance of an SMPC approach with inte-

[1] The celebrated open-loop optimal feedback (OLOF) control result, which is fundamental to MPC, belongs to this class of control policies (Dreyfus, 1963).

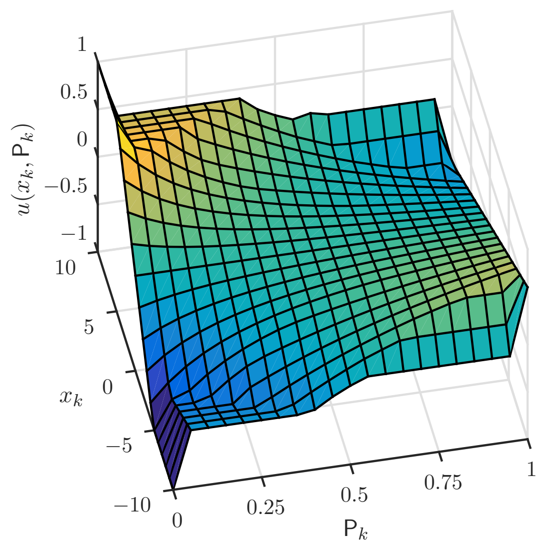

Figure 1. The control inputs computed by the SMPC approach with integrated input design for active model discrimination in relation to the probability of model validity and observed system state.

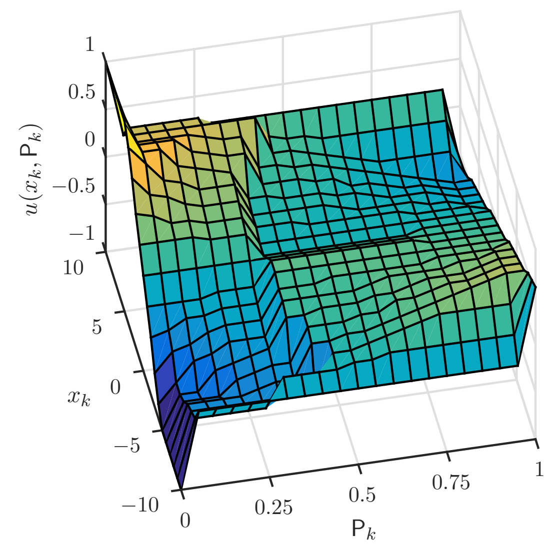

Figure 2. The control inputs computed by solving the stochastic OCP (2) using a closed-loop control policy in relation to the probability of model validity and observed system state.

grated input design for active model discrimination is demonstrated for a simple case: two competing scalar, discrete-time models with no parametric uncertainty, with one of the models being the true representation of the system. Under the effect of a Gaussian stochastic process, one model is stable and has a high positive input gain while the other is unstable, with a low negative input gain. The probability of validity of the first model structure P_k is updated based on Bayes' rule using the system observations $\mathcal{Y}_k$ at each sampling time. The Kolmogorov distance is used to enable model separation one sampling time into the future, with the two models having separate input sequences over the control horizon except for the first element, which is the same for both sequences. For a standard quadratic stage cost and a prediction horizon of $N = 7$, the performance of the SMPC approach with integrated input design is compared to that of (numerically) solving the stochastic OCP (2) subject to the two models using a closed-loop control policy. Notice that the latter approach requires solving the Bellman equation (3) for DP over the N stages, whereas the SMPC approach does not involve recursion of (3).

Figures 1 and 2 show the control input profiles computed by the SMPC approach with integrated input design and the stochastic OCP (2) with a closed-loop control policy, respectively. The input profiles are plotted as a function of the observed system state and the probability of validity of the first model. The similarity of the control input profiles designed by the two approaches suggests that the SMPC approach with integrated input design is able to adequately approximate the N-stage recursion of the Bellman equation (3). While solving (3) recursively based on quantization of the space of state x_k and probability P_k is computationally expensive and must be performed offline, the SMPC approach computes the control input for every pair (x_k, P_k) online, several orders of magnitude faster. Monte Carlo simulation of the true system with the control input profile obtained by the SMPC approach verifies that the system performance, as measured by the average cost function, is almost identical to what results from using the input profile obtained via solving the stochastic OCP (2) with a closed-loop control policy.

The performance of the SMPC approach with integrated input design for active model discrimination remains to be verified for practically-sized systems. To this end, the next section outlines some of the key challenges associated with an explicit dual control approach to solving the stochastic OCP (2).

Key Challenges in SMPC with Integrated Input Design for Active Model Discrimination

Efficient uncertainty propagation

Efficient propagation of probabilistic model uncertainties and stochastic system disturbances through the competing models (1) poses a key challenge to solving the stochastic OCP (2). Arbitrary polynomial chaos (aPC) can provide a computationally efficient method for sample-based approximation of the hyperstate (Oladyshkin and Nowak, 2012; Paulson et al., 2017). In aPC, each stochastic state is approximated by an expansion of orthogonal polynomial basis functions, which are defined based on moments of probabilistic uncertainties (without requiring the knowledge of full distributions). A key feature of aPC is its ability to propagate correlated random variables. The aPC expansions can be used as a surrogate for the nonlinear system models to perform Monte Carlo simulations efficiently since the basis functions are computed offline for different uncertainty realizations.

Updating the probability of validity of models based on system observations

As new system observations become available in the course of operation, the confidence in each model must be updated to reflect the degree to which the model agrees with the closed-loop data. An intuitive representation of confidence in each model is the conditional probability that each model is "true" given the current observations. Bayesian estimation techniques can provide a statistical framework for quantifying the confidence in each model structure (e.g., see Kuure-Kinsey and Bequette, 2009).

Model Invalidation

Similar to the issue of identifiability in parameter estimation (Ljung, 1999), system observations should be sufficiently informative for model invalidation. In other words, the closed-loop data generated by the controller should allow invalidating the competing models that cannot describe the data adequately. Verifying whether model invalidation is possible given the system observations poses a key theoretical challenge.

Practical Significance of Model Structure Uncertainty Handling in SMPC

The ability to handle multiple model structures in SMPC can lead to unprecedented opportunities in stochastic optimal control of modern engineering applications. Two important application areas are outlined below. In a wide range of technical systems, there is structural uncertainty in system dynamics. As an example, consider a reaction where an enzyme and a substrate are combined to form an enzyme-substrate complex, which is then converted into the final product. Here, an optimal experiment can be designed offline for probabilistic discrimination of the competing model hypotheses for the reaction mechanisms (Streif et al., 2014). A different approach involves designing a controller using both models to operate the system in such a way that the controller can decide, with a high degree of certainty, which model best describes the system at any given time. Such a controller can detect changes in reaction kinetics and adapts its underlying model by deciding whether the system is better described and controlled with a different model. SMPC with active model discrimination also enables using a collection of simpler models (relative to a complex model that is valid over a wide operating range) and switching between models as the system transitions from one operating mode to another. This can circumvent the need for developing complex models, which is often resource intensive.

Another application is SMPC with active fault diagnosis (Heirung and Mesbah, 2017). It is often not known a priori when a system fault will occur, if at all, and with what likelihood. However, models of potential system faults and failures may be available. In this case, the closed-loop data collected at each sampling time can be used to improve the confidence regarding which fault scenario is currently taking place. However, the mere presence of a feedback controller in the loop may mask the consequences of a fault, and hamper fault detection and isolation based on nominal closed-loop data. Moreover, this masking may in certain cases allow the faults to worsen and, for example, potentially compromise system safety. Some standard approaches to fault detection and isolation (Blanke et al., 2006) share features with model discrimination (Ashari et al., 2011). These approaches are generally based on offline input design (e.g., see Mesbah et al., 2014 and the references therein). A controller that intelligently operates the system to better reveal faults as they occur can ensure safer operation and overall better closed-loop control performance. This involves actively exploring and searching for faults, that is, performing online fault detection and isolation with autonomous adaptation of the system models.

References

Ashari, A.E., Nikoukhah, R., and Campbell, S.L. Auxiliary signal design for robust active fault detection of linear discrete-time systems. *Automatica*, 47(9):1887–1895, 2011.

Åström, K.J. and Wittenmark, B. *Adaptive Control*. Dover Publications, New York, 1995.

Atkinson, A.C. and Cox, D.R. Planning experiments for discriminating between models. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(3):321–348, 1974.

Bar-Shalom, Y. and Tse, E. Dual effect, certainty equivalence, and separation in stochastic control. *IEEE Transactions on Automatic Control*, 19(5):494–500, 1974.

Bayard, D.S. and Schumitzky, A. Implicit dual control based on particle filtering and forward dynamic programming. *International Journal of Adaptive Control and Signal Processing*, 24(3):155–177, 2010.

Bellman, R. *Dynamic Programming*. Princeton University Press, New Jersey, 1957.

Bertsekas, D. Dynamic Programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11(4-5):310–334, 2005.

Blanke, M., Kinnaert, M., Lunze, J., and Staroswiecki, M. *Diagnosis and Fault-Tolerant Control*. Springer-Verlag, Berlin, 2006.

Dreyfus, S.E. Some types of optimal control of stochastic systems. Technical report, The Rand Corporation, 1963.

Feldbaum, A.A. Dual-control theory. I. *Automation and Remote Control*, 21(9):874–880, 1961.

Filatov, N.M. and Unbehauen, H. Survey of adaptive dual control methods. *IEE Proceedings - Control Theory and Applications*, 147(1):118–128, 2000.

Gevers, M. Identification for control: From the early achievements to the revival of experiment design. *European Journal of Control*, 11(4-5):335–352, 2005.

Gibbs, A.L. and Su, F.E. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002.

Heirung, T.A.N. and Mesbah, A. Stochastic nonlinear model predictive control with active model discrimination for online fault detection. In *Proceedings of IFAC World Congress*. Submitted, Toulouse, France, 2017.

Heirung, T.A.N., Ydstie, B.E., and Foss, B. Dual MPC for FIR systems: Information anticipation. In *Proceedings of Advanced Control of Chemical Processes*, pages 1034–1039. Whistler, Canada, 2015.

Hjalmarsson, H. From experiment design to closed-loop control. *Automatica*, 41(3):393–438, 2005.

Kolmogorov, A.N. Sulla determinazione empirica di una legge di distribuzione. *Giornale dell'Istituto Italiano degli Attuari*, 4:83–91, 1933.

Kuure-Kinsey, M. and Bequette, B.W. Multiple model predictive control of nonlinear systems. In L. Magni, D.M. Raimondo, and F. Allgöwer, editors, *Nonlinear Model Predictive Control*, Lecture Notes in Control and Information Sciences, pages 153–165. Springer-Verlag, Berlin, 2009.

Larsson, C.A., Annergren, M., Hjalmarsson, H., Rojas, C.R., Bombois, X., Mesbah, A., and Modén, P.E. Model predictive control with integrated experiment design for output error systems. In *Proceedings of European Control Conference*, pages 3790–3795. Zürich, Switzerland, 2013.

Lee, J.M. and Lee, J.H. An approximate dynamic programming based approach to dual adaptive control. *Journal of Process Control*, 19(5):859–864, 2009.

Ljung, L. *System Identification. Theory for the User*. Prentice-Hall, New Jersey, 1999.

Marafioti, G., Bitmead, R., and Hovd, M. Persistently exciting model predictive control. *International Journal of Adaptive Control and Signal Processing*, 28(6):536–552, 2014.

Mesbah, A. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6), 2016.

Mesbah, A., Streif, S., Findeisen, R., and Braatz, R.D. Active fault diagnosis for nonlinear systems with probabilistic uncertainties. In *Proceedings of IFAC World Congress*, pages 7079–7084. Cape Town, South Africa, 2014.

Narendra, K.S. and Han, Z. The changing face of adaptive control: The use of multiple models. *Annual Reviews in Control*, 35(1):1–12, 2011.

Oladyshkin, S. and Nowak, W. Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion. *Reliability Engineering and System Safety*, 106:179–190, 2012.

Paulson, J.A., Buehler, E.A., and Mesbah, A. Arbitrary polynomial chaos for uncertainty propagation of correlated random variables in dynamic systems. In *Proceedings of IFAC World Congress*. Submitted, Toulouse, France, 2017.

Streif, S., Petzke, F., Mesbah, A., Findeisen, R., and Braatz, R.D. Optimal experimental design for probabilistic model discrimination using polynomial chaos. In *Proceedings of IFAC World Congress*, pages 4103–4109. Cape Town, South Africa, 2014.

Sun, Z., Qin, S.J., Singhal, A., and Megan, L. Performance monitoring of model-predictive controllers via model residual assessment. *Journal of Process Control*, 23(4):473–482, 2013.

Tanaskovic, M., Fagiano, L., Smith, R., and Morari, M. Adaptive receding horizon control for constrained MIMO systems. *Automatica*, 50(12):3019–3029, 2014.

Tse, E. and Bar-Shalom, Y. An actively adaptive control for linear systems with random parameters via the dual control approach. *IEEE Transactions on Automatic Control*, 18(2):109–117, 1973.

Wittenmark, B. An active suboptimal dual controller for systems with stochastic parameters. *Automatic Control Thoery & Applications*, 3(1):13–19, 1975a.

Wittenmark, B. Stochastic adaptive control methods: A survey. *International Journal of Control*, 21(5):705–730, 1975b.