

Finite Horizon Control Design for Optimal Model Discrimination

Lars Blackmore and Brian Williams

Abstract—In many fault detection and system identification problems, it is essential to be able to discriminate between a number of competing models of a system based on observed system outputs. For example, in a fault detection scenario we may wish to determine whether a system is best modeled by a known nominal model, or a known failure model. The probability of detecting the true system model can be enhanced by design of the control inputs applied to the system. In this paper we present a method by which a finite sequence of control inputs is designed automatically in order to minimize an upper bound on the probability of model selection error between any two linear, discrete-time systems. We are able to solve this problem efficiently by showing that it is an instance of a Quadratic Program. In addition, linear equality and inequality constraints can be applied to the control inputs and expected system state. These constraints can be used to ensure that a certain task is fulfilled, make sure the system stays within a valid linearization region, or to guarantee safe operation. Experimental results for the case of an aircraft actuator failure scenario show that the method significantly reduces the upper bound on the probability of model selection error when compared to a manually generated sequence and a fuel-optimal sequence.

I. INTRODUCTION

In multiple-model (MM) fault detection it is necessary to select the most likely model from a finite set, given observations [1][2]. For example Hanlon et al. investigated the detection of an aircraft flight control actuator failure [3]. In this case detecting a fault becomes a problem of deciding whether the dynamic model describing the nominal behavior of the aircraft, or the dynamic model describing the actuator failure is most likely.

Previous authors have developed methods for deciding between models given a set of observations [3][4]. For this paper we assume a Bayesian decision rule, which selects the most likely model given the observations and a prior distribution over the models. The ability of a detection method to discriminate between different competing dynamic system models is highly dependent on the control inputs applied to the system. For example, in the control actuator failure case, if no control inputs are applied to the actuator, the responses of both the faulty and the nominal system model will be identical. A new method for designing system inputs that discriminate between the possible models in an optimal sense is presented in this paper.

Esposito et al. created a *persistent excitation* solution to the problem of model discrimination for two linear filters [5].

However, as Prasanth et al. [6] noted for the case of model parameter estimation, a persistent excitation approach [7][8] may not be feasible for a physical system such as a satellite, because control inputs typically need to be designed subject to state and control constraints. They noted that a finite horizon optimization approach is more appropriate, and pose the problem as a Model Predictive Control (MPC) problem.

We extend the work of [5] and [6] by describing a method by which finite, optimized sequences of control inputs can be designed, subject to control and state constraints, in order to discriminate between two discrete-time dynamic linear models. While our method is limited to the case of two models, discrimination between two models is useful for binary hypotheses such as whether the system is in either a nominal or a failure mode, and is also a step towards the design of control inputs for discrimination between more than two models.

Prior work in the field of experiment design [9][10] has suggested a number of different criteria for the design of experiments for model discrimination. In this paper, consistent with a Bayesian approach to model selection, the aim is to minimize the probability of model selection error by the Bayes-optimal decision rule, known as the *Bayes Risk* [11]. Assuming a 0-1 loss function for the model selection task, the decision-theoretic optimal design is the one that minimizes the probability of model selection error [12].

We use an upper bound on the probability of model selection error, the Battacharyya bound, to create a tractable optimization criterion for model discrimination. Then we pose the problem of designing a finite sequence of control inputs to minimize this bound, subject to constraints, as a finite horizon trajectory design problem. Lastly, we show that in the case of linear constraints this is an example of a concave Quadratic Program. Prior work in the field of optimization allows concave Quadratic Programs to be solved efficiently [13][14][15][16].

The result of this work is a new algorithm that generates a finite sequence of control inputs that minimize an upper bound on the probability of model selection error. These sequences are designed subject to state and control constraints. The algorithm can be used to ensure that a given task, defined in terms of constraints on the expected state, is fulfilled while optimally detecting failures. We present results for the example of determining whether an aircraft's elevator actuator has failed. Compared to a typical sequence designed by a human, and a sequence optimized to minimize fuel consumption, the method dramatically reduces the upper bound on the probability of model selection error.

This work is supported by NASA Award NNA04CK91A

Lars Blackmore is a PhD student, Massachusetts Institute of Technology, Cambridge, MA 02139 larsb@mit.edu

Brian Williams is an associate professor, Massachusetts Institute of Technology, Cambridge, MA 02139 williams@mit.edu

A. Aircraft Fault Detection Scenario

In this paper, the detection of elevator actuator failure for an aircraft is used as a motivating example. The discrete time approximation to the longitudinal dynamics of the aircraft, linearized about the trim state, is shown in Fig. 1. In the model selection task, we must determine which

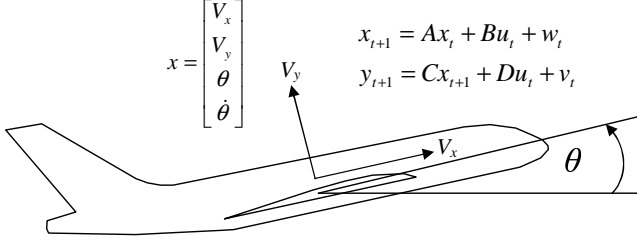


Fig. 1. Discrete-time aircraft model linearized about the trim state. Here x_t is the state of the system at time step t while y_t is the observed output of the system at time t , which is taken to be the pitch rate $\dot{\theta}$. The input is denoted u_t , and is taken to be the elevator angle. The terms w_t and v_t are the *process noise* and *observation noise*. The noise at any time step is assumed to be independent of the noise at any other time step, and w_t and v_t are independent of each other with normal distributions $\mathcal{N}(0, Q)$ and $\mathcal{N}(0, R)$ respectively. The initial state of the system is also assumed to be normally distributed.

system described by the system matrices $\{A, B, C, D\}$ best models the data. Here the situation under consideration has two candidate models. Under Hypothesis 0, the system is described by $\{A_0, B_0, C_0, D_0\}$ while under Hypothesis 1, the system is described by $\{A_1, B_1, C_1, D_1\}$.

For the aircraft in Fig. 1, we may wish to determine whether the elevator actuator is faulty. In the faulty case, the B matrix is zero, indicating that the input (commanded elevator angle) has no effect on the system. Intuitively, one might carry out an experiment where, at the trim state, a large elevator angle would be commanded. In the case of a working actuator the effect on the system would be significant, while in the case of a faulty actuator the commanded input would have no effect. The resulting observations would therefore reveal which of the two hypotheses is correct. This paper presents an online, optimized technique for designing experiments of this type.

II. HYPOTHESIS SELECTION AND BAYES RISK

Here we assume models are selected by Bayesian hypothesis selection. Restricting our attention to selection between two models, Bayesian hypothesis selection can be expressed as follows:

Select H_0 if $p(H_0|\mathbf{y}, \mathbf{u}) > p(H_1|\mathbf{y}, \mathbf{u})$, else select H_1 .

Using Bayes' rule, this selection is given by:

Select H_0 if $p(\mathbf{y}|H_0, \mathbf{u})p(H_0) > p(\mathbf{y}|H_1, \mathbf{u})p(H_1)$, else select H_1 .

The terms $p(H_0)$ and $p(H_1)$ correspond to prior probabilities of the two hypotheses. These can be calculated in a number of different ways: there may be explicit knowledge about how *a priori* likely the different hypotheses are, or the prior can represent the belief state created by an estimator.

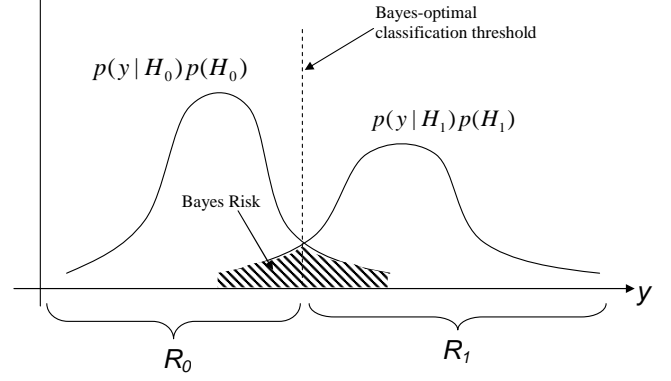


Fig. 2. Selection between two models given an observation y and a prior. In general Bayesian selection between two hypotheses yields a threshold that splits the possible observations into two sets. If the observation y falls into set \mathfrak{R}_0 then the classifier selects H_0 , whereas if the observation falls into set \mathfrak{R}_1 then the classifier selects H_1 . Even with Bayes optimal selection there is a finite probability of error given by the *Bayes Risk*, denoted by the shaded region.

The Bayesian selection rule minimizes the likelihood of selecting an incorrect hypothesis given the available information. As shown in Fig. 2, the Bayesian optimal classifier has a finite probability of selecting the incorrect hypothesis, known as the *Bayes Risk*. The Bayes risk is given by:

$$\begin{aligned} P(\text{error}) &= P(\mathbf{y} \in \mathfrak{R}_1, H_0|\mathbf{u}) + P(\mathbf{y} \in \mathfrak{R}_0, H_1|\mathbf{u}) \\ &= P(\mathbf{y} \in \mathfrak{R}_1|H_0, \mathbf{u})P(H_0) + P(\mathbf{y} \in \mathfrak{R}_0|H_1, \mathbf{u})P(H_1) \\ &= \int_{\mathfrak{R}_1} p(\mathbf{y}|H_0, \mathbf{u})P(H_0)d\mathbf{y} + \int_{\mathfrak{R}_0} p(\mathbf{y}|H_1, \mathbf{u})P(H_1)d\mathbf{y} \quad (1) \end{aligned}$$

Since the Bayes Risk is the probability of error when using the optimal classifier, we would like to optimize our control inputs to the system to minimize this measure.

III. THE ALGORITHM

A. Minimizing the Bayes Risk

The idea behind the design of control inputs for model selection is that while the probability of error cannot be reduced beyond the Bayes Risk by selection of the classification threshold, the Bayes Risk itself is affected by the control inputs to the system. Hence by selecting the inputs to the system, we can reduce the Bayes Risk and therefore significantly reduce the probability of error of the Bayesian classifier. The effect of input choice on the Bayes Risk is illustrated in Fig. 3.

B. Input Design as Trajectory Optimization

A key observation is that the design of control inputs to minimize the Bayes risk is in fact a problem of optimal trajectory design for dynamic systems. In typical optimal trajectory design problems, an optimized sequence of control inputs is designed so that a system passes through a sequence of predicted states that minimize some cost function (for example the time taken to reach a goal state). In the model discrimination problem, we would like to design an optimal sequence of control inputs so that the system passes through

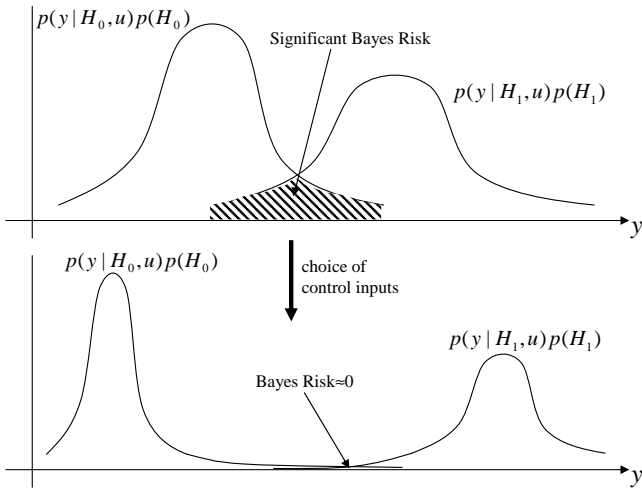


Fig. 3. Graph showing $p(y|H_0, u)$ and $p(y|H_1, u)$ for two different choices of u . In the upper figure, the predicted distributions overlap significantly, leading to a large Bayes risk. In the lower figure, a different selection of u has separated the two distributions, meaning that when the observation y is made, the correct hypothesis can be selected with high confidence. The Bayes risk is very low, meaning that the probability of error is very low.

a trajectory of state *distributions* that minimize the probability of model selection error (Bayes Risk).

Optimal trajectory design for linear, discrete-time systems can be posed as a quadratic program in the case of a cost function that is quadratic in the decision variables. A quadratic program involves the minimization of a quadratic function subject to linear constraints. For the trajectory design problem, the linear nature of the state update equation shown in Fig. 1 means that linear state constraints correspond to linear constraints on the decision variables. Efficient algorithms exist for the solution of quadratic programs. The main contribution of this project is to show that the model discrimination problem can be posed as an optimal trajectory design problem and solved using a quadratic program. This is described in the following sections.

1) *Cost Function for Optimal Input Design:* In the model discrimination problem, the objective is to minimize the probability of error of Bayesian model selection, known as the Bayes Risk. However the Bayes Risk given in (1) is not in closed form and can only be calculated using numerical integration. It is, however, possible to *bound* the Bayes Risk in closed form, and doing so allows the use of Quadratic Programming which is reliable and efficient.

The *Battacharyya Bound* [11] provides an upper bound on the Bayes Risk and is given by the integral:

$$P(\text{error}) \leq \sqrt{P(H_0)P(H_1)} \int \sqrt{p(y|H_0)p(y|H_1)} dy \quad (2)$$

If the distributions are Gaussian such that $p(y|H_0)$ has mean μ_0 and variance Σ_0 , and $p(y|H_1)$ has mean μ_1 and variance Σ_1 , the above value can be calculated without the need for integration:

$$P(\text{error}) \leq \sqrt{P(H_0)P(H_1)} \exp\{-k\},$$

where:

$$k = \frac{1}{4}(\mu_1 - \mu_0)^T [\Sigma_0 + \Sigma_1]^{-1} (\mu_1 - \mu_0) + \frac{1}{2} \ln \frac{|\frac{\Sigma_0 + \Sigma_1}{2}|}{\sqrt{|\Sigma_0||\Sigma_1|}}. \quad (3)$$

Since the logarithm is a monotonically increasing function, the value of x that optimizes $f(x)$ is also the value that optimizes $\ln[f(x)]$. We therefore take the logarithm of the Battacharyya bound for Gaussian distributions to yield the following cost function:

$$J = \frac{1}{2} \ln[P(H_0)P(H_1)] - \frac{1}{2} \ln \frac{|\frac{\Sigma_0 + \Sigma_1}{2}|}{\sqrt{|\Sigma_0||\Sigma_1|}} - \frac{1}{4}(\mu_1 - \mu_0)^T [\Sigma_0 + \Sigma_1]^{-1} (\mu_1 - \mu_0) \quad (4)$$

Minimizing this cost function will therefore minimize an upper bound on the probability of error when using Bayesian selection to decide between two models based on a vector of observations y .

2) *Error Probability for a Finite Horizon:* Rather than selecting inputs at one moment in time, we would like to plan a finite sequence of inputs in order to minimize the upper bound on the probability of error. This form of planning is known as *finite horizon* planning. In this case, if the horizon is of length k , we are concerned with a sequence of observations $y_{t+1} \dots y_{t+k}$ and a sequence of inputs $u_t \dots u_{t+k-1}$. The derivation in III-B.1 extends readily to a finite horizon. Defining:

$$\mathbf{y} = [y_{t+1}^T \quad y_{t+2}^T \quad \dots \quad y_{t+k}^T]^T$$

$$\mathbf{u} = [u_t^T \quad u_{t+1}^T \quad \dots \quad u_{t+k-1}^T]^T \quad (5)$$

The new objective is to minimize an upper bound on the probability of error of a Bayes optimal classifier that makes its decision based on all of the observations $y_{t+1} \dots y_{t+k}$, by designing the sequence of control inputs $u_t \dots u_{t+k-1}$.

Due to uncertainty in the initial state and noise, the future observations $y_{t+1} \dots y_{t+k}$ are random variables, which we denote $Y_{t+1} \dots Y_{t+k}$. Under the assumptions in Section I-A, Y_{t+i} is normally distributed given a sequence of inputs \mathbf{u} and given a model (H_0 or H_1). We now define $\mu_{t+i,h}$ and $\Sigma_{t+i,h}$ for time steps $i = 1, \dots, k$ and hypotheses $h = 0, 1$ such that:

$$p_{Y_{t+i}}(y|H_0, \mathbf{u}) = \mathcal{N}(\mu_{t+i,0}, \Sigma_{t+i,0})$$

$$p_{Y_{t+i}}(y|H_1, \mathbf{u}) = \mathcal{N}(\mu_{t+i,1}, \Sigma_{t+i,1}) \quad (6)$$

Then the vector of all observations $Y = [Y_{t+1}^T \dots Y_{t+k}^T]^T$ is a vector of normally distributed random variables given a sequence of inputs and a hypothesis. We define μ_h and Σ_h for $h = 0, 1$ to be the mean and covariance of the vector of all observations such that:

$$p_Y(y|H_0, \mathbf{u}) = \mathcal{N}(\mu_0, \Sigma_0)$$

$$p_Y(y|H_1, \mathbf{u}) = \mathcal{N}(\mu_1, \Sigma_1) \quad (7)$$

From the above definitions the distribution of Y_h is given by:

$$\mu_h = [\mu_{t+1,h}^T \dots \mu_{t+k,h}^T]^T \quad (8)$$

$$[\Sigma_h]_{i,j} = E\left[\left([Y]_i - [\mu_h]_i\right)\left([Y]_j - [\mu_h]_j\right) \middle| H_h\right] \quad (9)$$

Here $[\cdot]_i$ denotes the value at index i into the vector, and similarly $[\cdot]_{i,j}$ denotes the value at index (i, j) into the matrix.

Having defined μ_0, μ_1, Σ_0 and Σ_1 , the Battacharyya bound given in (3) provides an upper bound for the probability of error when using the entire sequence of observations from time $t+1$ to time $t+k$. Hence the cost function in (4) applies to the finite horizon formulation given in this section. The problem of designing a sequence of inputs to minimize this cost function is addressed in the following sections.

3) Predicting the Distribution of Future Observations:

Given a certain hypothesis, the system equations shown in Fig. 1 are fully known. Hence the distributions $p(\mathbf{y}|H_0, \mathbf{u})$ and $p(\mathbf{y}|H_1, \mathbf{u})$ can be calculated. This section gives the result for $p(\mathbf{y}|H_0, \mathbf{u})$; an identical method applies for $p(\mathbf{y}|H_1, \mathbf{u})$.

For some sequence of inputs $\mathbf{u}_t \dots \mathbf{u}_{t+k-1}$ the system will pass through a sequence of states $\mathbf{x}_t \dots \mathbf{x}_{t+k}$. In order to minimize (2), we must determine how the distribution of $\mathbf{x}_t \dots \mathbf{x}_{t+k}$, and hence $p(\mathbf{y}|H_0, \mathbf{u})$, depends on this sequence of inputs. Given a sequence of inputs, the initial state \mathbf{x}_t and a system model, the system equations can be applied recursively to derive the following equation relating the observation at time step $t+i$ to the initial state, the inputs, and the noise:

$$\begin{aligned} \mathbf{y}_{t+i,0} = & C_0 \sum_{l=0}^{i-1} A_0^{i-l-1} (B_0 \mathbf{u}_{t+l} + \mathbf{w}_{t+l}) \\ & + C_0 A_0^i \mathbf{x}_t + D_0 \mathbf{u}_{t+i-1} + \mathbf{v}_{t+i-1} \end{aligned} \quad (10)$$

Given a distribution for the initial state of the system $\mathcal{N}(\hat{x}_0, P)$, the mean μ_0 and covariance Σ_0 as defined in (8) and (9) can be calculated. The results are given here for a system where $\mathbf{y}_t \in \mathbb{R}^n$.

Define:

$$i = n(p-1) + q, \quad (11)$$

where:

$$1 \leq q \leq n \quad 1 \leq p \leq k \quad p, q \in \mathbb{Z}. \quad (12)$$

Then following from (5) and (10):

$$\begin{aligned} [\mu_0]_i = & [\mu_{t+p,0}]_q \\ = & \left[C_0 A_0^p \hat{\mathbf{x}}_0 + C_0 \sum_{l=0}^{p-1} A_0^{p-l-1} B_0 \mathbf{u}_{t+l} + D_0 \mathbf{u}_{t+p-1} \right]_q \end{aligned} \quad (13)$$

Defining $j = n(r-1) + s$ in the same manner as (11), the expression for the covariance is:

$$\begin{aligned} [\Sigma_0]_{i,j} = & [R'(p, r) + C_0 A_0^p P_0 A_0^{Tr} C_0^T]_{q,s} \\ & + \left[\sum_{l=0}^{m-1} C_0 A_0^{(p-l-1)} Q A_0^{T(r-l-1)} C_0^T \right]_{q,s} \end{aligned} \quad (14)$$

where $m = \min\{i, j\}$ and:

$$R'(p, r) = \begin{cases} R & p = r \\ 0 & p \neq r. \end{cases} \quad (15)$$

Under the assumptions mentioned in Section I-A, these equations give an expression for the belief state $p(\mathbf{y}|H_0, \mathbf{u})$ in terms of a sufficient statistic for Y , namely the mean and variance of the distribution. These equations have two important properties:

- 1) The equation for the mean of the predicted distribution of Y is linear in the control inputs \mathbf{u} .
- 2) The covariance of the predicted distribution of Y is not a function of the control inputs \mathbf{u} .

These two properties are important because they mean that the cost function is a quadratic function of the control inputs as will be shown in Section III-C.

C. Cost as a Quadratic Function of Control Variables

Examining (4) it can be seen that the only term that depends on the inputs \mathbf{u} is the term involving μ_1 and μ_0 , since neither the prior nor the covariance depend on the input. The other terms need therefore not be considered when optimizing J with respect to the inputs. The term to be minimized is shown here:

$$J' = -(\mu_1 - \mu_0)^T [\Sigma_0 + \Sigma_1]^{-1} (\mu_1 - \mu_0) \quad (16)$$

Now that we have explicit expressions for Σ_0 and Σ_1 from (14) that do not depend on the control inputs \mathbf{u} , and since the means μ_1 and μ_0 are linear in the control inputs, this equation can be written in terms of a quadratic with regard to the input variables:

$$J' = \mathbf{u}^T H \mathbf{u} + \mathbf{f}^T \mathbf{u} + \mathbf{g} \quad (17)$$

The quantities H, \mathbf{f}^T and \mathbf{g} can be calculated explicitly. The Battacharyya bound on the Bayes Risk therefore yields a cost function that is quadratic in the control inputs.

Since the covariance matrices Σ_0 and Σ_1 are positive definite, the cost function given by (16) is a concave function of $(\mu_1 - \mu_0)$. From (13), both μ_1 and μ_0 are linear functions of \mathbf{u} , and hence the cost in (17) is a concave function of the control inputs \mathbf{u} . This concavity makes the cost function particularly tractable for optimization and guarantees that a global optimum can be found in bounded time [13].

Although the two terms in (4) that do not involve μ_0 and μ_1 can be neglected when optimizing (4), these terms do affect the value of the optimum that optimization can achieve. A trivial example is the case where the prior for one of the hypotheses is zero. In this case the upper bound on the Bayes Risk is zero regardless of the observations.

D. Linear Constraints

As noted by Prasanth et al., a powerful aspect of the finite horizon optimization formulation is that both equality and inequality constraints can be placed on the trajectory design [6]. In the model discrimination problem there are a number of constraints on the design that can be expressed as linear

functions of the control inputs; for example, constraining the expected state of the system at any point in the horizon for a particular hypothesis. This follows from the following equation (shown for Hypothesis 0):

$$\begin{aligned} \mathbf{x}_{t+i,0} &= \sum_{l=0}^{i-1} A_0^{i-l-1} (B_0 \mathbf{u}_{t+l} + \mathbf{w}_{t+l}) + A_0^i \mathbf{x}_t \\ E[\mathbf{x}_{t+i}|H_0] &= \sum_{l=0}^{i-1} A_0^{i-l-1} (B_0 \mathbf{u}_{t+l}) + A_0^i \hat{\mathbf{x}}_0 \end{aligned} \quad (18)$$

Constraints on the expected mean of the system state can be used to:

- 1) Ensure that a certain task, defined in terms of the expected system state, is fulfilled
- 2) Ensure that the mean of system stays within a ‘safe’ operating region or within a valid linearization region
- 3) Ensure that the system ends the experiment in the same region as it started

In addition linear constraints can be placed on the control inputs directly, for example $u_{min} < u_{t+i} < u_{max}$ modeling actuator limits, or those of the type $\sum_i |\mathbf{u}_{t+i}| < fuel$ that limit total control effort over the horizon. Such constraints allow the user to trade off the cost of the control effort against the corresponding reduction in the probability of model selection error.

E. Summary

We have shown that the problem of designing a sequence of control inputs in order to minimize an upper bound on the probability of error of Bayesian model selection can be posed as a finite-horizon trajectory design problem, and have shown that this problem has a cost function that is quadratic in the decision variables. In addition, we are able to place a number of linear constraints on the control variables in order to model control or state constraints. Hence the active model discrimination problem can be posed as a Quadratic Program and can be solved efficiently using existing methods.

IV. SIMULATION

This section describes results from a number of trajectory design tasks for the aircraft elevator failure scenario. In this scenario the ability to detect a fault is critical, and depends heavily on the control inputs issued. In addition, actuator saturation, linearization about a trim state and safety considerations mean that the ability to constrain the designed inputs and expected state is critical.

Although the designed trajectories would ideally be compared in terms of the true Bayes Risk, the cost of computing this value through numerical integration means that here the designed trajectories are compared in terms of the upper bound on the probability of model selection error, the Battacharyya bound.

A. Constrained Inputs

The algorithm was used to design a sequence of control inputs in order to choose between two models for the aircraft shown in Fig. 1. According to Hypothesis 0 the aircraft has a

working elevator actuator, and according to Hypothesis 1 this actuator is broken. The inputs in the latter case have no effect on the system, as described in section I-A. The maximum allowable elevator angle is $\pm 0.25rad$ and the horizon length is 40 time steps, or 20 seconds. Fig. 4 shows the sequence of inputs designed by the algorithm along with the trajectories of the expected observations conditioned on Hypothesis 0 and Hypothesis 1. The Battacharyya bound for the generated sequence is 0.0029, meaning that the probability of model selection error is at most 0.29%.

B. Constrained State and Inputs

Although the control sequence for this scenario produces a low bound on the Bayes Risk, the resulting motion of the aircraft closely resembles an unstable oscillation, with steadily increasing magnitude. This is readily addressed by placing constraints on the expected mean of the system state \mathbf{x} . Fig. 5 shows results for an identical scenario, except with the additional constraints that at the end of the horizon, $|E[\dot{\theta}]| \leq 0.25rad/s$ and $|E[\theta]| \leq 0.25rad$ for both H_0 and H_1 . This time the generated control sequence induces a trajectory for the mean of the system state that ends up within these bounds at the end of the horizon; and with a Battacharyya bound that is only slightly greater than for the unconstrained case at a value of 0.0031.

C. Manually Generated Identification Sequence

When performing system identification for the longitudinal dynamics of an aircraft, a pilot often issues elevator inputs that form a pulse or doublet pattern [17]. Fig. 6 shows the expected observations for such a sequence with the same actuator limits as for the optimized sequences. The Battacharyya bound for this sequence is 0.073, more than twenty times the bound for the optimized sequence in Fig. 5.

Note that the human-generated doublet sequence is similar to the first portion of the sequence optimized for model discrimination. However the optimized sequence in Fig. 5 is, by comparison, able to reduce the bound on the probability of error dramatically while guaranteeing that the final state of the system is bounded.

D. Model Discrimination during Altitude Change Maneuver

The ability to constrain the state of the system means that this method can be used to optimize model discrimination during a specified maneuver. Fig. 7 shows results for a maneuver where the aircraft performs a change in altitude. The two plots compare the maneuver designed to minimize the probability of model selection error to that designed to minimize fuel consumption. A fuel-optimal design is typical for finite horizon path planning with unmanned air vehicles.

The Battacharyya bound for the fuel-optimal case is 0.13 whereas the bound for the model discrimination optimized case is 0.0053. This example demonstrates the significant improvement in fault detection that can be achieved by using control inputs designed specifically for model discrimination, rather than those optimized with regard to some other parameter and employing only passive model selection.

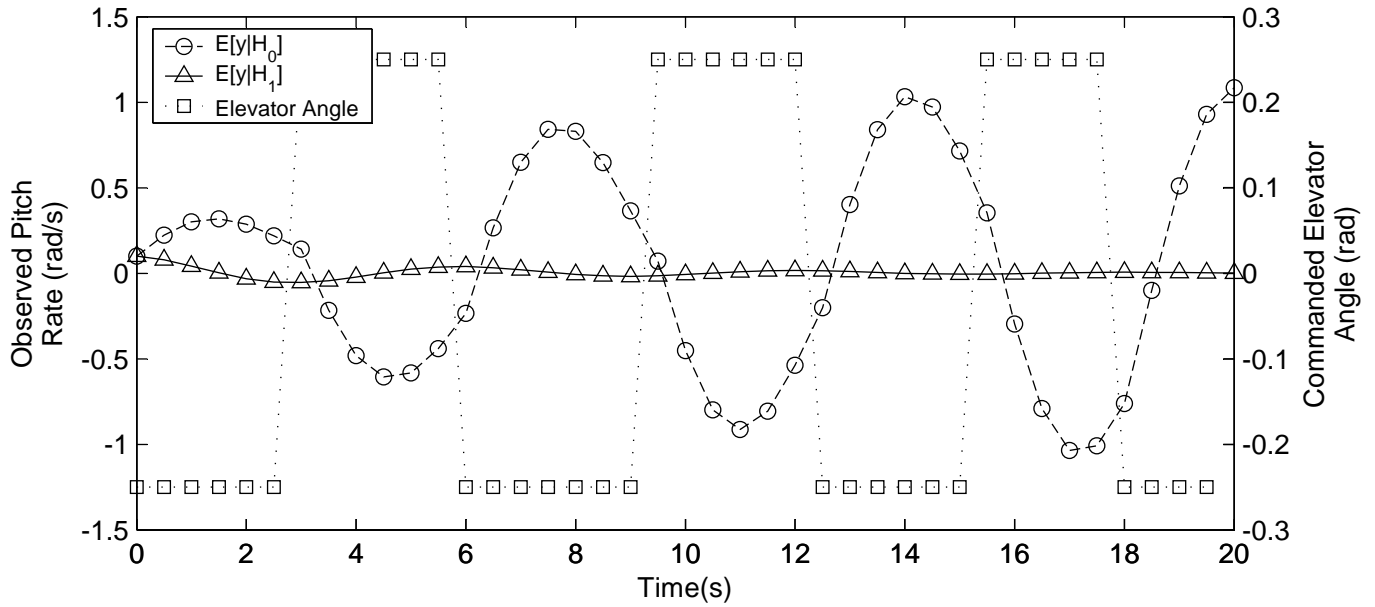


Fig. 4. Optimized control inputs and predicted observation means for aircraft with $|\mathbf{u}|_{max} = 0.25rad$. The algorithm has arrived at a solution that gives a sequence of inputs at the furthest extremes of the allowable range, alternating between periods of $u = 0.25rad$ and periods of $u = -0.25rad$. Notice also that the period of this sequence is close to that of the *short period oscillation* mode of the aircraft seen in the unforced sequence μ_1 . Hence the control sequence effectively drives the aircraft as far as possible over the time horizon in order to reduce the ambiguity between the two hypotheses.

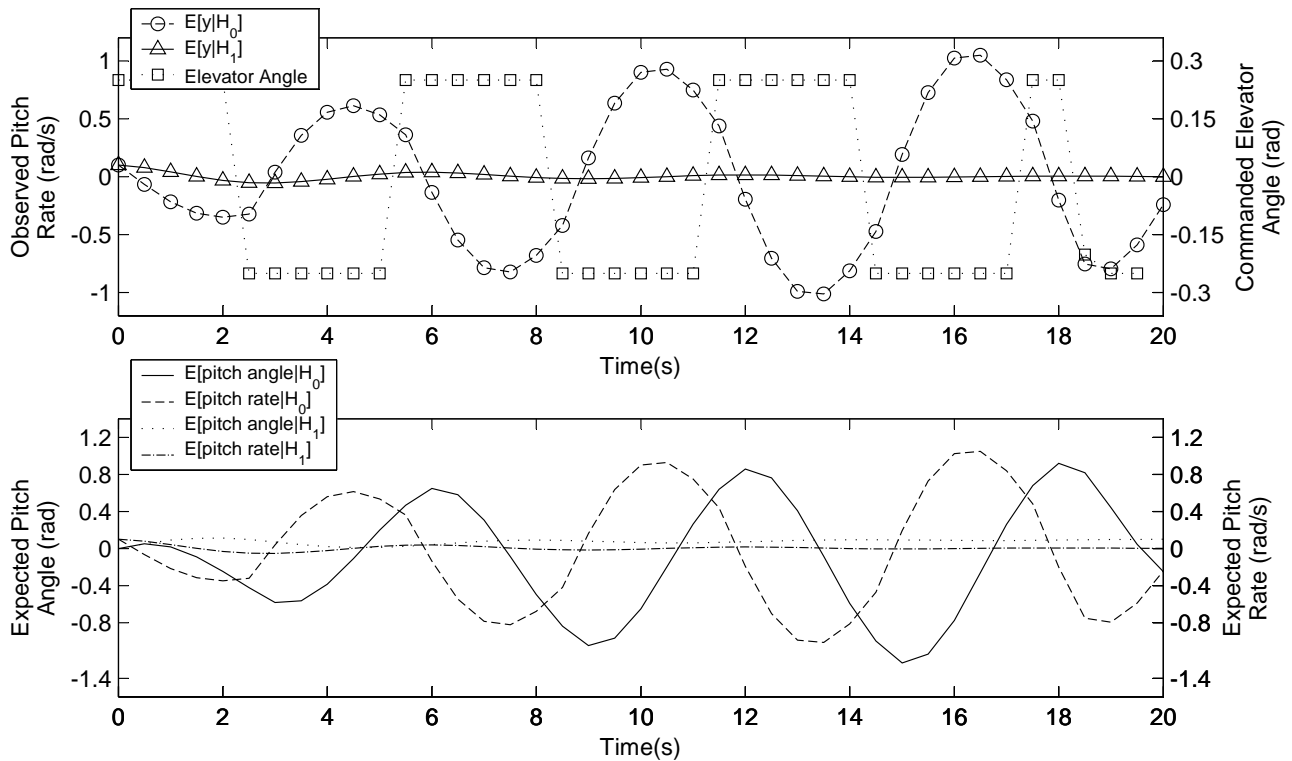


Fig. 5. Optimized control inputs, expected observation, pitch angle and pitch rate for aircraft with expected final state constrained by $|E[\dot{\theta}]| \leq 0.25rad/s$ and $|E[\theta]| \leq 0.25rad$ for both H_0 and H_1 .

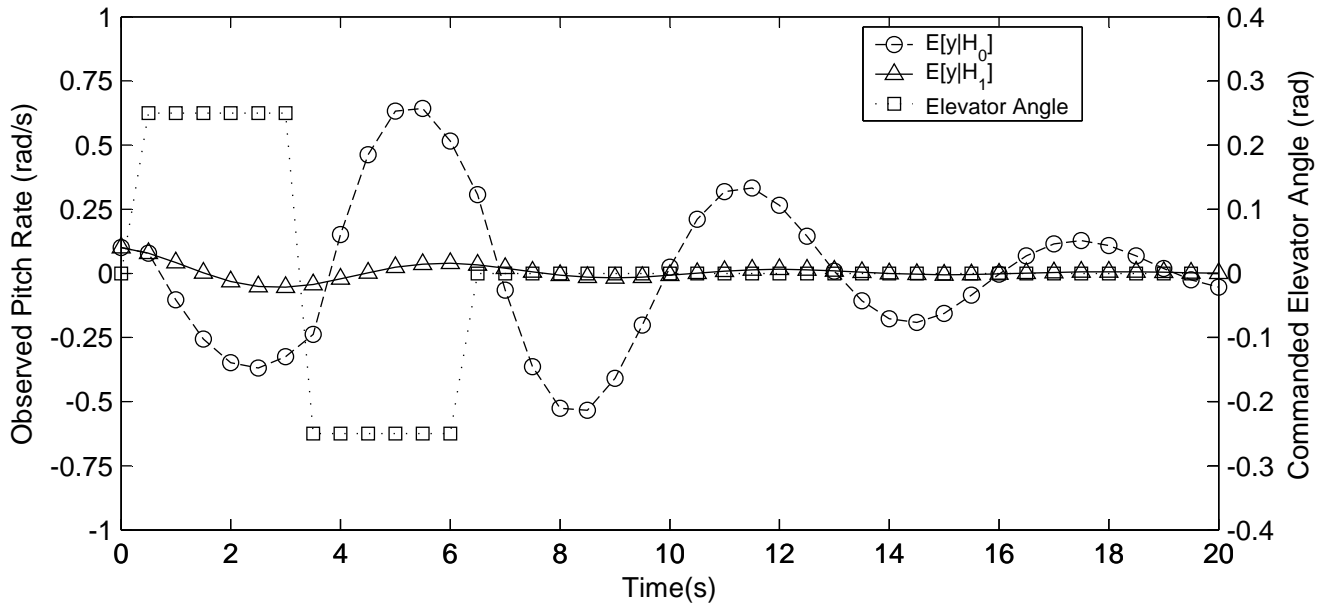


Fig. 6. Control input, expected observation, pitch angle and pitch rate for manually generated doublet sequence.

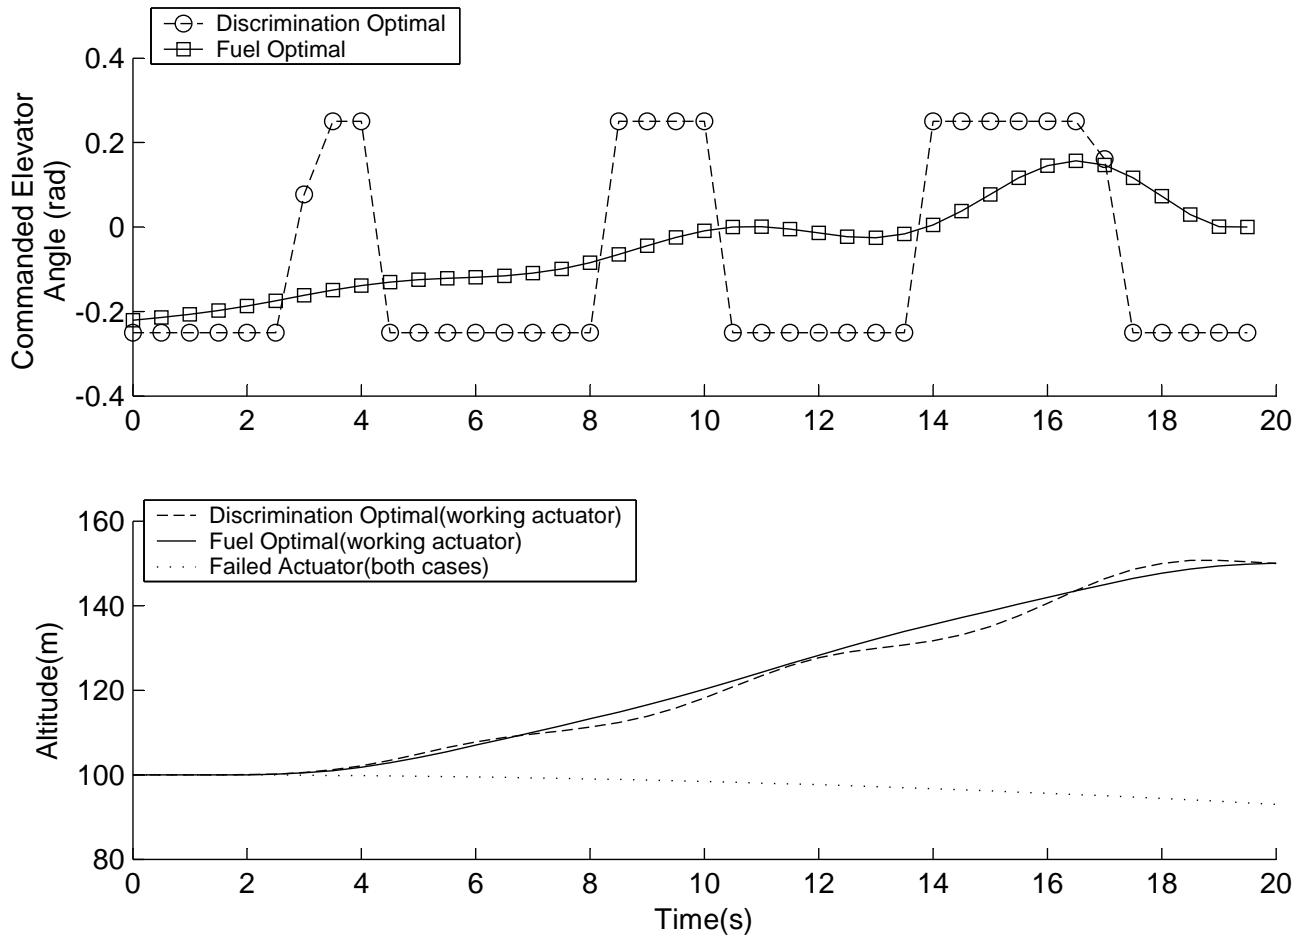


Fig. 7. Comparison of altitude increase maneuvers optimized for fuel consumption and model discrimination. Conditioned on a working actuator, the end state of the aircraft is constrained to have altitude=150m and zero vertical velocity. No constraints are placed on the state given a broken actuator.

V. DISCUSSION

In this section, some properties of the new method are discussed. Firstly, the new method has an additional capability that was not demonstrated in this paper. The expected system state can be constrained conditioned on *either* model being the true one. This means that, for example, control inputs can be constrained so as to guarantee safe operation or task completion under either hypothesis. Note that the ability to apply these constraints does not guarantee that a feasible solution exists.

In addition the method can be applied to a slightly different problem formulation than that described in the previous sections. In this alternative formulation, we would like to constrain the probability of model selection error to be below an acceptable threshold while minimizing some other cost, such as fuel consumption. In this formulation the constraint on the probability of error gives rise to quadratic constraints, and the new problem can be solved using Quadratically Constrained Quadratic Programming. This formulation is, however, less tractable than the simpler Quadratic Programming formulation.

The algorithm presented in this paper has three main limitations. Firstly, while the method aims to reduce the Bayes Risk for model selection, in fact it is only an upper bound on this value that is minimized. Furthermore, there are no guarantees of the tightness of this bound. In many cases, however, an optimized trajectory that yields a low upper bound on the Bayes Risk will be an acceptable solution. Second, the method is limited to the case of discrimination between two models. In a fault detection scenario, for example, there may be more than two competing models, corresponding to different failure modes of the system. Future work will extend the method to the case of multiple models. Lastly, the method is restricted to linear systems. While linearization may be used to solve this problem in certain cases, future work will investigate model discrimination for non-linear dynamic systems.

VI. CONCLUSION

This paper presents a new algorithm for model discrimination that poses the problem as a finite horizon trajectory design problem and minimizes an upper bound on the probability of model selection error. This problem is an example of a concave Quadratic Program, and hence can be solved efficiently by methods that are guaranteed to converge to the

global optimum in bounded time. The Quadratic Programming formulation allows arbitrary linear constraints to be placed on the control inputs and expected system state.

Simulation results showed that compared to a typical human-generated control sequence, and a control sequence optimized with regard to fuel consumption, the new method can significantly reduce the upper bound on the probability of model selection error.

REFERENCES

- [1] A. Pouliezios and G. Stavrakakis, *Real Time Fault Monitoring of Industrial Processes*. Boston: Kluwer Academic, 1994.
- [2] J. Chen and R. Patton, *Robust Model-Based Fault Diagnosis for Dynamic Systems*. Boston: Kluwer Academic, 1999.
- [3] P. Hanlon and P. Maybeck, "Multiple-model adaptive estimation using a residual correlation kalman filter bank," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 36, pp. 393–406, Apr. 2000.
- [4] A. Willsky, "A survey of design methods for failure detection in dynamic systems," *Automatica*, vol. 12, pp. 601–611, 1976.
- [5] R. Esposito and M. Schumer, "Probing linear filters - signal design for the detection problem," *IEEE Transactions on Information Theory*, vol. IT-16, 1970.
- [6] R. Prasanth, J. Amin, K. J. S. Seereeram, R. Mehra, D. Bayard, and F. Hadaegh, "Predictive control approach to maneuver design for autonomous formation flying sensor calibration," in *Proc. AIAA Conference and Exhibit on Guidance, Navigation and Control*, Rhode Island, Aug. 2004.
- [7] R. Mehra, "Optimal input signals for parameter estimation in dynamic systems: Survey and new results," *IEEE Transactions on Automatic Control*, vol. 19, Dec. 1974.
- [8] L. Ljung, *System Identification: Theory for the user*. New York: Prentice Hall, 1999.
- [9] V. Fedorov, *Theory of Optimal Experiments*. New York: Academic Press, 1972.
- [10] P. Hill, "A review of experimental design procedures for regression model discrimination," *Technometrics*, vol. 20, Feb. 1978.
- [11] R. Duda, P. Hart, and D. Stork, *Pattern Classification (2nd ed.)*. Wiley Interscience, 2000.
- [12] K. Felsenstein, "Optimal bayesian design for discrimination among rival models," *Computational Statistics and Data Analysis*, vol. 14, pp. 427–436, 1992.
- [13] P. Pardalos and J. Rosen, "Methods for global concave minimization: A bibliographic survey," *Society for Industrial and Applied Mathematics*, SIAM Review, Sept. 1986.
- [14] J. Rosen, "Global minimization of large-scale constrained concave quadratic problems by separable programming," *Mathematical Programming*, vol. 34, pp. 163–174, 1986.
- [15] I. Bomze, "A global optimization algorithm for concave quadratic programming problems," *SIAM Journal on Optimization*, vol. 3, pp. 826–842, 1993.
- [16] B. Kalantari, "Large scale global minimization of linearly constrained concave quadratic functions and related problems," Ph.D. dissertation, Univ. of Minnesota, Minneapolis, Minnesota, 1984.
- [17] J. Jang and C. Tomlin, "Longitudinal stability augmentation system design of the stanford dragonfly uav using a single gps receiver," in *Proceedings of the AIAA GNC conference*, Austin, Texas, Aug. 2003.