

DATA-DRIVEN ROBUST CONTROL DESIGN: UNFALSIFIED CONTROL

Michael G. Safonov
Department of Electrical Engineering
University of Southern California
Los Angeles, CA 90089-2563
U.S.A.

msafonov@usc.edu

ABSTRACT

Aerospace applications require precise control despite uncertain operating conditions and unanticipated circumstances such as battle damage. These systems must be designed to perform robustly, despite uncertain design models and difficult to analyze nonlinear effects. They must be capable of learning and adapting when accumulating data indicates that previous models must be abandoned and that existing control strategies must be changed. Data-driven design methods, collectively known as un-falsified control theory, facilitate the creation of robust control systems that learn, discover and evolve in real time in order to rapidly switch controller gains to compensate for the effects of battle, equipment failures, and other changing circumstances. Applications studies will be presented that include adaptive robot arm control and missile control.

“I have devised seven separate explanations, each of which would cover the facts as far as we know them. But which of these is correct can only be determined by fresh information which we shall no doubt find waiting for us.”

Sherlock Holmes
Arthur Conan Doyle

1.0 INTRODUCTION

Though the robust multivariable control theory that has evolved over the past quarter century offers a major improvement over earlier algebraic and optimal control methods, it cannot produce reliable control designs unless reliable prior bounds on plant uncertainty are available. This applies to methods based on the H_∞ μ/K_m -synthesis, and BMI/LMI/IQC theories [1]–[6]. These robust control design methods all have an Achilles heel: They are dependent of the premise that uncertainty models are reliable, and they offer little guidance in the event that experimental data either invalidates prior knowledge of uncertainty bounds or, perhaps, provides evidence of previously unsuspected patterns in the data. That is, the standard H_∞ μ/K_m -synthesis, and BMI/LMI/IQC robust control techniques fail in the all too common situation in which prior knowledge is poor or unreliable.

To correct this, reliable *data-driven* adaptive design techniques are needed. Ideally, these techniques should incorporate mechanisms for evaluating the design implications of each new experimental data point, and for directly integrating that information into the mathematics of the robust control design process to allow methodical update and re-design of control strategies so as to accurately reflect the

implications of new or evolving experimental data. Examples of recent thrusts in this direction are indirect controller tuning/adaptation methods based on control-oriented identification theory and [7]–[26] and, more recently, related direct methods that bypass plant identification based on controller unfalsification [27]–[55]. While both control-oriented identification theory and unfalsified control theory are concerned with the difficult problem of assimilating real-time measurement data into the otherwise introspective process of robust control design, the unfalsified control approach is a particular interest because it directly and precisely characterizes the control design implications of experimental data.

2.0 DATA-DRIVEN ROBUST CONTROL

Validation — or more precisely *unfalsification* — of hypotheses against physical data is the central aspect of the process of scientific discovery. This validation process allows scientists to sift the elegant tautologies of pure mathematics in order to discover mathematical descriptions of nature that are not only for logically self-consistent, but also consistent with physically observed data. This data-driven process of validation is also a key part engineering design. Successful engineering design techniques inevitably arrive at a point where pure introspective theory and model-based analyses must be tested against physical data. But, in control engineering in particular, the validation process is one that has been much neglected by theoreticians. Here, the theory tying control designs to physical data has for the most part focused on pre-control-design ‘system identification’. Otherwise, the mathematization of the processes of post-design validation and re-design has remained relatively unexplored virgin territory. In particular, a satisfactory quantitative mathematical theory for direct feedback of experimental design-validation data into the control design process has been lacking, though this seems to be changing with the recent introduction of a theory of unfalsified control [31].

2.1 Theory: Validation and Unfalsification

Unfalsified control is essentially a data-driven adaptive control theory that permits learning based on physical data via a process of elimination, much like the candidate elimination algorithm of Mitchell [56, 57]. The theory concerns the feedback control configuration in Figure 1. As always in control theory, the goal is to determine a control law K for the plant P such that the closed-loop system response, say T , satisfies given specifications. Unfalsified control theory is concerned with the case in which the plant is either unknown or is only partially known and one wishes to fully utilize information from measurements in selecting the control law K . In the theory of unfalsified control, learning takes place when new information in measurement data enables one to eliminate from consideration one or more candidate controllers.

As indicated in Fig. 2 three elements that define the unfalsified control problem are (1) plant measurement data, (2) a class of candidate controllers, and (3) a performance specification, say T_{spec} , consisting of a set of admissible 3-tuples of signals (r, y, u) . More precisely, we have the following.

Definition [31] *A controller K is said to be **falsified** by measurement information if this information is sufficient to deduce that the performance specification $(r, y, u) \in T_{spec} \forall r \in \mathcal{R}$ would be violated if that controller were in the feedback loop. Otherwise, the control law K is said to be **unfalsified**. \square*

To put plant models, data and controller models on an equal footing with performance specifications, these like T_{spec} are regarded as sets of 3-tuples of signals (r, y, u) — that is, they are regarded

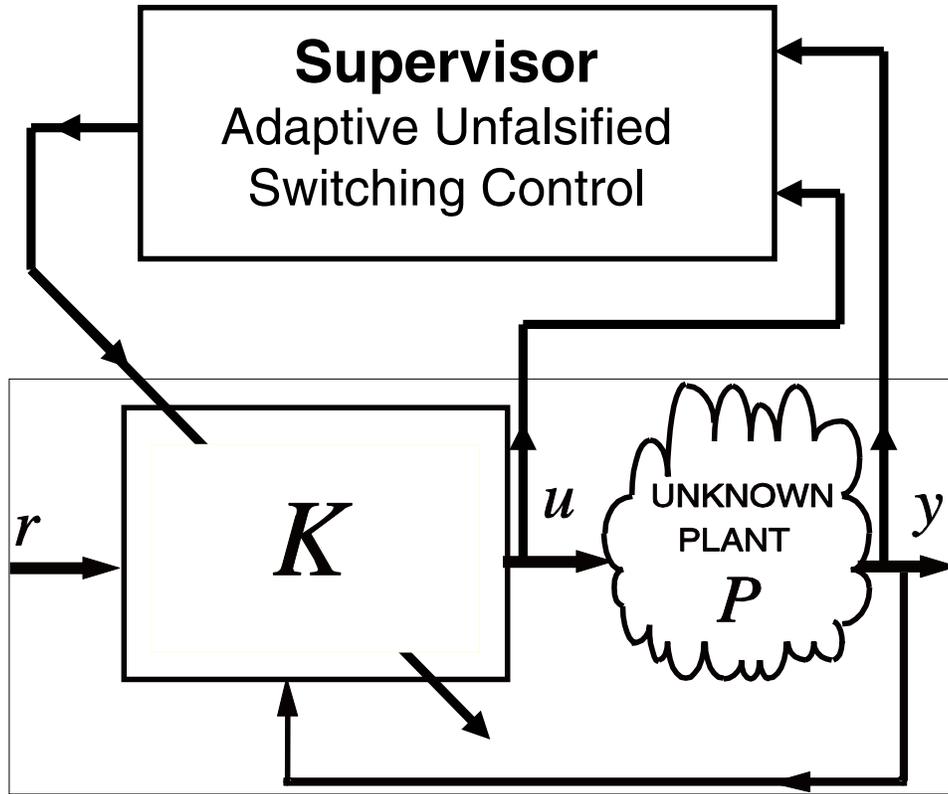


Figure 1: An unfalsified adaptive controller has a two-tier structure, consisting of a adaptive supervisor and a conventional controller K . The supervisor monitors plant data (u, y) for evidence that would falsify candidate controllers. If the currently active controller K becomes falsified by data, then an as yet unfalsified controller is switched into the loop to replace it.

as relations in $\mathcal{R} \times \mathcal{Y} \times \mathcal{U}$. For example, if $P : \mathcal{U} \rightarrow \mathcal{Y}$ and $K : \mathcal{R} \times \mathcal{Y} \rightarrow \mathcal{U}$ then

$$\begin{aligned}
 P &= \{ (r, y, u) \mid y = Pu \} \\
 K &= \left\{ (r, y, u) \mid u = K \begin{bmatrix} r \\ y \end{bmatrix} \right\}.
 \end{aligned}$$

And, if $J(r, y, u, t)$ is a given loss-function that we wish to be non-positive for all time t , then the performance specification T_{spec} would be simply the set

$$T_{spec} = \{ (r, y, u) \mid J(r, y, u, t) \leq 0 \forall t \}. \quad (1)$$

On the other hand, experimental information from a plant corresponds to *partial* knowledge of the plant P . Loosely, data may be regarded as providing a sort of an “interpolation constraint” on the graph of P — i.e., a ‘point’ or set of ‘points’ through which the infinite-dimensional graph of dynamical operator P must pass.

Typically, the available measurement information will depend on the current time, say τ . For example, if we have complete data on (u, y) from time 0 up to time $\tau > 0$, then the *measurement*

information is characterized by the set [31]

$$P_{data} \triangleq \left\{ (r, y, u) \in \mathcal{R} \times \mathcal{U} \times \mathcal{Y} \mid P_{\tau} \begin{bmatrix} (u - u_{data}) \\ (y - y_{data}) \end{bmatrix} = 0 \right\} \quad (2)$$

where P_{τ} is the familiar time-truncation operator of input-output stability theory (cf. [58, 59]), viz.,

$$[P_{\tau}x](t) \triangleq \begin{cases} x(t), & \text{if } 0 \leq t \leq \tau \\ 0, & \text{otherwise.} \end{cases}$$

The main result of unfalsified control theory is the following theorem which gives necessary and sufficient conditions for past open-loop plant data P_{data} to falsify the hypothesis that controller K can satisfy the performance specification T_{spec} .

Unfalsified Control Theorem [31] *A control law K is unfalsified by measurement information P_{data} if, and only if, for each triple $(r_0, y_0, u_0) \in P_{data} \cap K$, there exists at least one pair (\hat{u}_0, \hat{y}_0) such that*

$$(r_0, \hat{y}_0, \hat{u}_0) \in P_{data} \cap K \cap T_{spec}. \quad (3)$$

Proof: With controller K in the loop, a command signal $r_0 \in \mathcal{R}$ could have produced the measurement information if, and only if, $(r_0, y_0, u_0) \in P_{data} \cap K$ for some (u_0, y_0) . The controller K is unfalsified if and only if for each such r_0 there is at least one (possibly different) pair (u_1, y_1) which also could have produced the measurement information with K in the loop and which additionally satisfies the performance specification $(r_0, y_1, u_1) \in T_{spec}$. That is, K is unfalsified if and only if for each such r_0 , condition (3) holds. \square

The **Unfalsified Control Theorem** constitutes a mathematically precise statement of what it means for experimental data and a performance specification to be inconsistent with a particular controller. It has some interesting implications:

- The **Unfalsified Control Theorem** is nonconservative; i.e., it gives “if and only if” conditions on K . It uses all the information in the past data — and no more. It provides a mathematically precise “sieve” which rejects any controller which, based on experimental evidence, is demonstrably incapable of meeting a given performance specification.
- The **Unfalsified Control Theorem** is “model free”. No plant model is needed to test its conditions. There are no assumptions about the plant.
- Information P_{data} which invalidates a particular controller K need not have been generated with that controller in the feedback loop; it may be open loop data or data generated by some other control law (which need not even be in \mathbf{K}).
- When the sets P_{data} , K and T_{spec} are each expressible in terms of equations and/or inequalities, then falsification of a controller reduces to a minimax optimization problem. For some forms of inequalities and equalities (e.g., linear or quadratic), this optimization problem may be solved analytically, leading to procedures for direct identification of controllers — as the example in [34].

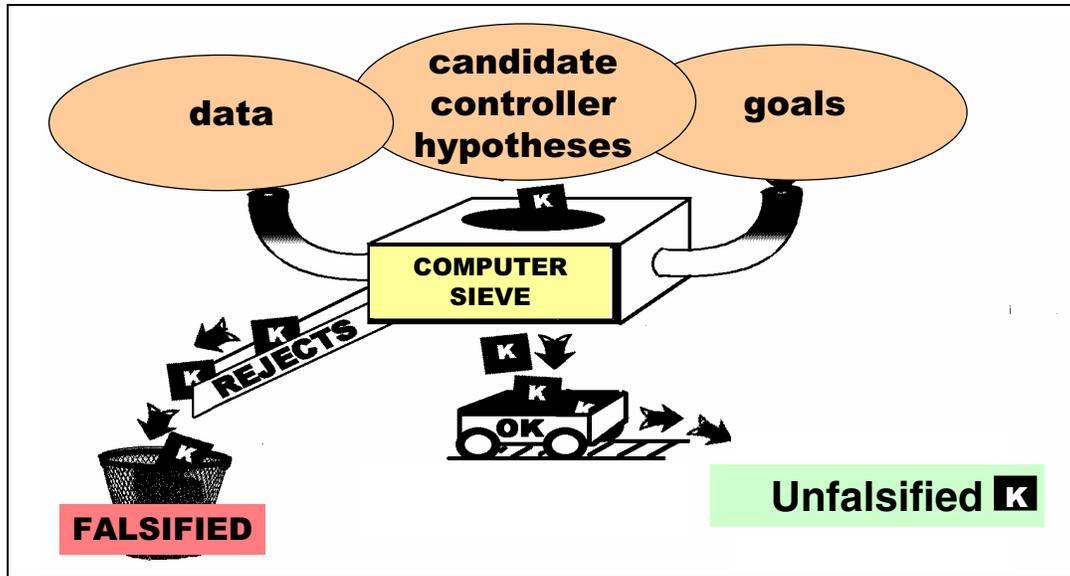


Figure 2: An unfalsification process is used supervisory control design. The process requires three types of input (1) goals, (2) candidate controllers and (3) data. Controllers are sifted to find those that are consistent with both performance goals and physical data. No plant models are required while the process is running, though a plant model can be useful for prior selections of the candidate controllers and the performance goal.

- Given data (u_0, y_0) and a candidate controller K , the r_0 's satisfying the conditions of the **Unfalsified Control Theorem** are called the fictitious reference signals. When K has a causal inverse, the r_0 is uniquely determined by (u_0, y_0) and a candidate controller K ; that is, there exists a causal function $\tilde{r}(K, u_0, y_0)$ such that

$$r_0 = \tilde{r}(K, u_0, y_0).$$

The function $\tilde{r}(K, u_0, y_0)$ is called the *fictitious reference signal* [27, 31]; and is closely related to the *virtual* reference signal of [50].

- In adaptive control the fact that the supervisor chooses a controller means, at least implicitly, that there is a real-valued data-driven cost function $V(K, y_0, u_0, t)$ such that at any give time τ the active controller is cost-minimizing

$$\hat{K}(y_0, u_0, t) = \arg \min_{K \in \mathbf{K}} V(K, y_0, u_0, t). \tag{4}$$

In the case of unfalsified control, the cost function V is determined from the cost $J(r, y, u, t)$ and evaluating it with r equal to the fictitious reference signal $\tilde{r}(K, u_0, y_0)$:

$$V(K, y_0, u_0, t) = J(\tilde{r}(K, u_0, y_0), y_0, u_0, t). \tag{5}$$

2.2 Data-Driven Learning and Adaptive Control

The unfalsified control theorem says simply that controller falsification can be tested by computing an intersection of certain sets of signals. A noteworthy feature of the unfalsified control theory is that

Table 1: Recursive Adaptive Control Algorithm

INPUT:

- A finite set \mathbf{K} of m candidate dynamical controllers $K_i(r, y, u) = 0$, ($i = 1, \dots, m$) each having the causal-left-invertibility property that $r(t)$ is uniquely determined from $K_i(r, y, u) = 0$ by past values of $u(t), y(t)$.
- Sampling interval Δt and current time $\tau = n\Delta t$;
- Plant data $(u(t), y(t))$, $t \in [0, \tau]$;
- Performance specification set T_{spec} consisting of the set of triples (r, y, u) satisfying for all $k = 1, \dots, n$

$$\int_0^{k\Delta t} \tilde{T}_{spec}(r(t), y(t), u(t), t) dt \leq 0.$$

INITIALIZE:

- set $k = 0$, set $\hat{i} = m$;
- for $i = 0 : m$, set $s(i) = 1$, set $\tilde{J}(i) = 0$, end.

PROCEDURE:

- while $\hat{i} > 0$;
 - $k = k + 1$;
 - for $i = 1 : m$;
 - if $s(i) > 0$;
 - for each $t \in [(k-1)\Delta t, k\Delta t]$;
 - solve $K_i(r, y, u) = 0$ for $r(t)$;
 - (note that $r(t)$ exists and is unique since K_i has the causal-left-invertibility property)
 - end;
 - $\tilde{J}(i) = \tilde{J}(i) + \int_{(k-1)\Delta t}^{k\Delta t} \tilde{T}_{spec}(r(t), y(t), u(t), t) dt$;
 - if $\tilde{J}(i) > 0$, set $s(i) = 0$, end;
 - end;
 - end;
 - $\hat{i} = \max \{ i \mid s(i) > 0 \}$;
- end.

a controller need not be in the loop to be falsified. Broad classes of controllers can be falsified with open-loop plant data or even data acquired while other controllers were in the loop. Adaptive control is achieved within this framework by using the unfalsification process as the key element of a supervisory controller (cf. [60, 61]). The supervisor switches an unfalsified controller into the feedback loop whenever the current controller in the loop is amongst those falsified by observed plant data — see Fig. 2. The supervisor chooses as the current control law one that is not falsified by the past data, resulting in a control law that is adaptive in the sense that it learns in real time and changes based on what it learns.

Like the controllers of [62, 63], this approach to adaptive real-time unfalsified control leads to a sort of “switching control.” Controllers which are determined to be incapable of satisfactory performance are switched out of the feedback loop and replaced by others which, based on the information in past data, have not yet been found to be inconsistent with the performance specification. However, adaptive unfalsified controllers generally would not be expected to exhibit the poor transient response associated with switching methods such as [62]. The reason is that, unlike the theory in [62], unfalsified control theory efficiently eliminates broad classes of controllers before they are ever inserted in the

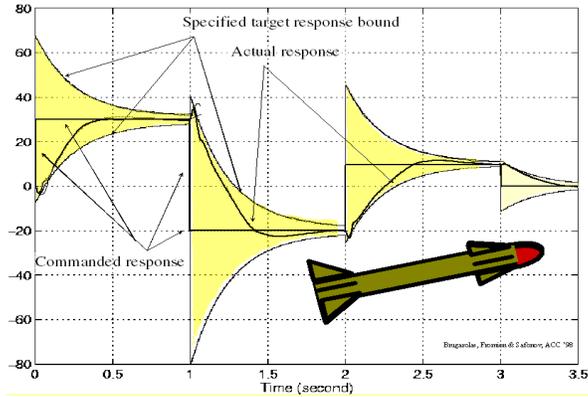


Figure 3: A data-driven unfalsified missile controller would have abilities to adaptively discover solutions in real-time to compensate for sudden in-flight changes and damage.

feedbackunfalsified control and other adaptive methods is that in unfalsified control one evaluates candidate controllers objectively based on experimental data alone, without prejudicial assumptions about the plant.

While, in principle, the unfalsified control theory allows for the set \mathbf{K} to include continuously parameterized sets of controllers, restricting attention to candidate controller sets \mathbf{K} with only a finite number of elements can simplify computations. Further simplifications result by restricting attention to candidate controllers that are “causally-left-invertible” in the sense that, given a $K \in \mathbf{K}$, the current value of $r(t)$ is uniquely determined by past values of $u(t), y(t)$. When (2) holds, these restrictions on T_{spec} and \mathbf{K} are sufficient to permit the unfalsified set to be evaluated in real-time via the following conceptual in Table 1.

This algorithm returns for each time the least index \hat{i} for which $K_{\hat{i}}$ is unfalsified by the past plant data. Real-time unfalsified adaptive control is achieved by always taking as the currently active controller

$$\hat{K} \triangleq K_{\hat{i}}$$

provided that the data does not falsify all candidate controllers. In this latter case, the algorithm terminates and returns $\hat{i} = 0$.

It is important to note that while the above algorithm is geared towards the case of an integral inequality performance criterion T_{spec} and a finite set of causally-left-invertible K_i 's, the underlying theory is, in principle, applicable to arbitrary non-finite controller sets \mathbf{K} and to hybrid systems with both discrete and continuous time elements.

Comment *If the plant is slowly time-varying, then older data ought to be discarded before evaluating controller falsification. This may be effected within the context of our the Recursive Adaptive Control Algorithm described in Table 1 by fixing $\tau = \tau_0$ and regarding $t - \tau_0$ as the deviation from the current time. The result is an algorithm which only considers data from moving time-window of fixed duration τ_0 time-units prior to the current real-time. In this case the unfalsified controller set \mathbf{K}_{OK} no longer shrinks monotonically as it would if τ were increasing in lockstep with real-time.*

2.3 Design Studies

Design studies have confirmed the theoretical expectation that supervisory controllers that are designed based on the logic of unfalsification can be effective in closing the outer data-driven loop on the control design process. Unfalsified control theory has proved effective in applications involving both off-line controller gain tuning and in real-time adaptive control design studies. Following is a brief description of some of the design studies that have helped us to better understand the potential of the unfalsified control theory, as well as limitations of the current theory.

2.3.1 Missile Autopilot

One design study that we conducted involved using an unfalsified controller to robustly discover PID controller gains for an adaptive missile autopilot ‘on the fly’ in real-time [36]. Figure 3 summarizes the results of the missile design. In all trials, the response of the adaptive loops was swift and sure-footed — in stark contrast to what would be expected from traditional quasi-static adaptive methods (e.g., standard model reference adaptive control).

2.3.2 Universal PID Controller

One application of the theory involved implementing a PID-based adaptive ‘universal’ controller implemented as MATLAB Simulink block based on the unfalsified theory [37] — see Fig 4. The controller sifts through a bank of candidate controllers in real-time, identifying which of the 30 controllers is unfalsified with respect to an inequality performance goal of ‘mixed-sensitivity’ type, viz.,

$$J(r, y, u, t) \leq 1 \text{ for all } t$$

where

$$J(r, y, u, t) = \max_{\tau \leq t} \frac{\|w_1 * (r - y)\|_{\tau}^2 + \|w_2 * u\|_{\tau}^2 - \sigma^2 \tau}{\|r\|_{\tau}^2 + \rho} \quad (6)$$

and w_1 and w_2 are ‘weighting’ filters and ρ and σ are constants chosen by the designer based on control bandwidth and robustness considerations – exactly as in standard mixed-sensitivity robust control design. Initially, there are 30 candidate PID gain combinations, indicated on the vertical axis. Unfalsified controllers are indicated by the horizontal traces. When the currently active controller indicated by the bold trace becomes falsified, then one of the as yet unfalsified controllers is switched into the loop to replace it. Notice that adaptive supervisor loop is so fast that the controller is able to stabilize the open-loop unstable plant without prior plant knowledge and without appreciable transients. The unfalsification procedure always discovers a stabilizing controller, The supervisory controller designed via unfalsified control quickly discovers a stabilizing candidate controller for the open-loop-unstable plant. The unfalsified control theory assures the procedure will always converge to a stabilizing controller that meets the performance goals, provided only that the initial candidate controller set contains at least one such controller; *no other prior knowledge of the plant is necessary to assure convergence*. In particular, the unfalsification logic in the supervisor is guaranteed to stabilize and meet performance goals without any of the ‘standard assumptions’ on the plant (e.g., [64, 65]) — viz., without minimum phase, linearity, bounds on relative-degree, knowledge of the sign of the plant high frequency gain, or persistence of excitation.

2.3.3 Robot Manipulator Arm

We used to unfalsified methodology to adaptively tune the parameters of a nonlinear ‘computed-torque’ controller for a robot manipulator arm [66, 32] — see Fig 5. The arm proved to be capable of a quick

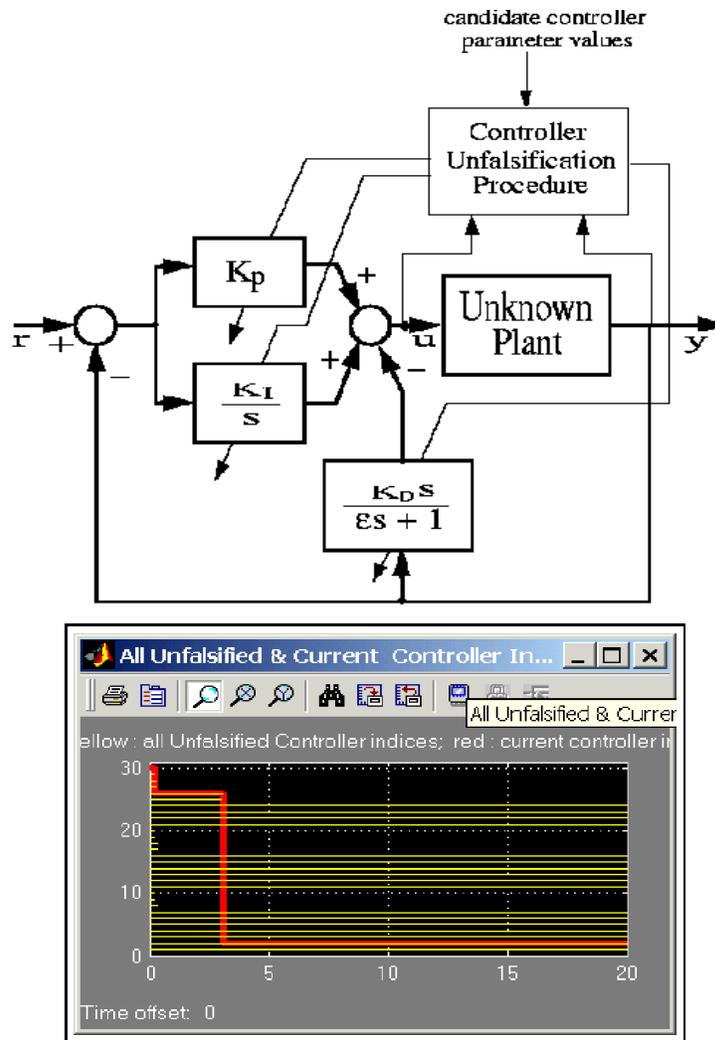


Figure 4: Here we see a MATLAB Simulink time-history for an adaptive ‘universal PID controller’.

and reliable control response despite large and sudden variations in load mass. Again, the controller performed with precision, despite noise, dynamical actuator uncertainties and without prior knowledge of the plant model or its parameters. Results for the robot design were surefooted and precise, with the controller maintaining an order of magnitude more precise control than a similar model-reference adaptive controller during widely fluctuating manipulator load variations; the controller was also more robust in that it was capable of maintaining precise control even during load variations that destabilized a similarly structured model-reference adaptive controller.

2.3.4 Industrial Process Control

Although very few researchers other than ourselves have as yet examined unfalsified control methods, those who have taken this step have predictably confirmed the effectiveness of unfalsified control methods in several industrial process control applications. For example, Kosut [38] examined unfalsified controller for direct data-driven off-line control gain tuning under the assumption of a noise-free linear-time-invariant plant. Woodley, How and Kosut [67] and used the theory with good result for

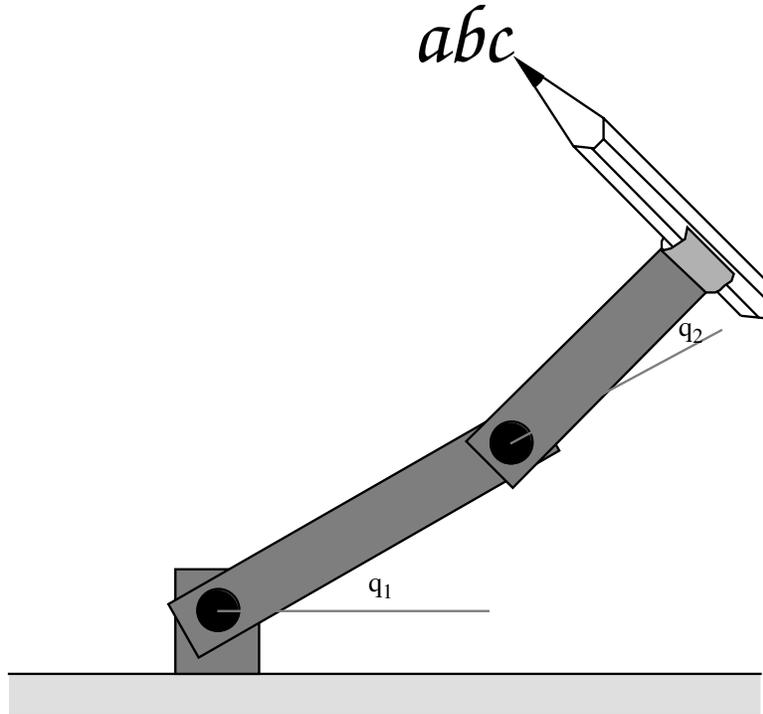


Figure 5: Unfalsified control produced superior results for a nonlinear two-link robot manipulator subject to uncertain dynamics, noisy disturbances and abrupt changes in load mass. The two sluggishly smooth traces large amplitude signals in the plot are with a conventional adaptive controller used to adjust control gain-vector $\theta(t)$, and the two very low amplitude traces are for the unfalsified controller. The unfalsified controller had a much quicker, sure-footed and precise response without increased control effort.

data-driven discovery of good control gains for a laboratory control problem involving two spring-connected masses. Also, Collins and Fan [39] successfully used the unfalsified control methodology in a run-to-run setting to tune gains off-line in an industrial weigh-belt feeder control design study. More recently, there have been some promising adaptive control applications to machine control by Razavi and Kurfess [40, 68] based on the unfalsified control methodology.

3.0 STABILITY AND CONVERGENCE

When modeling assumptions about the plant fail to hold, there is the possibility that badly designed adaptive algorithms can fail to stabilize — even when the adaptive control problem is theoretically feasible in the sense that one of the candidate controllers $K \in \mathbf{K}$ is stabilizing. The problem is that in most cases, ‘proofs’ of stability in adaptive control only hold when there is no mismatch between model assumptions and the true plant. Well-know standard assumptions in adaptive control include upper-bounds on plant order, assumption the plant is minimum phase and has no time-delays, or that the plant is ‘sufficiently close’ to one of several presumed prior plant models. There are a number of studies illustrating how adaptive systems can fail in the presence of a mismatch between assumptions and reality [69, 70]. A typical result of model mismatch instability is shown in Figure 6 which illustrates the consequences of model mismatch instability for a typical multi-model adaptive (MMAC) system. But,

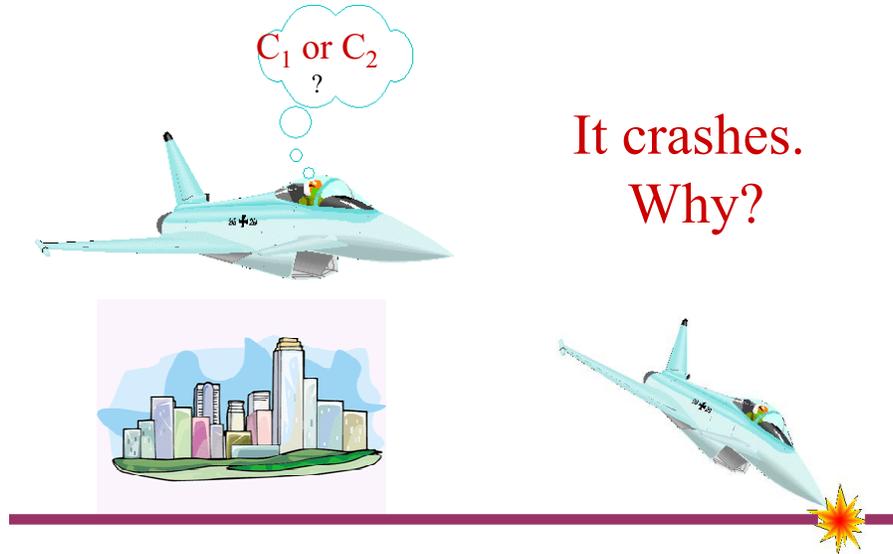


Figure 6: Without cost-detectability, model mismatch can cause an adaptive controller to switch a destabilizing controller C_2 into the loop and keep it, even when the original stabilizing controller C_1 was working well [70].

though currently popular adaptive algorithms are susceptible to model-mismatch instability, optimal unfalsified adaptive design designs with a suitable cost-detectability properties as well as some lesser known early adaptive methods robustly avoid model mismatch instability, provided that the unfalsified cost function $J(r, y, u, t)$ is chosen to have a property called *cost detectability* [72, 73].

For example, the multi-controller adaptive control (MCAC) switching algorithms and Martensson [71] and of Fu and Barmish [62] are robust against model-mismatch. They require essentially only feasibility to assure convergence. They work robustly in the presence of model mismatch because they assign a cost penalty to destabilizing controllers that tends to infinity for destabilizing candidate controllers. This is the essential feature of cost-detectability. Cost-detectability is a feature not typically present for most currently popular adaptive algorithms, and it explains why they are susceptible to model mismatch instability. Without the cost detectability property, adaptive control algorithms are in general not able to reliably distinguish stable and unstable behavior when model-mismatch exceeds certain thresholds. While, unfortunately, the early cost-detectable algorithms of [71, 62] are too slow for applications requiring real-time adaptive stabilization, this is not true of unfalsified adaptive control.

Definition 3.1 (Cost Detectability) A cost function $V(K, u, y, t)$ is said to have the cost detectability property if it has all the following attributes:

1. $V(K, u, y, t)$ is a monotone increasing function of time t for all K, u, y ,
2. When K is not stabilizing and the data u, y are not stable, then $\infty = \lim_{t \rightarrow \infty} V(K, u, y, t)$.
3. When K is stabilizing, $\lim_{t \rightarrow \infty} V(K, u, y, t)$ is uniformly bounded for all (u, y, t) , (even if the data u, y are unstable signals).

So-called ‘ L_{2e} -gain related’ cost functions like the ‘mixed-sensitivity’ cost (6) generically assure the

cost detectability property for unfalsified cost function [72, 73]

$$V(K, y, u, t) = J(\tilde{r}(K, u, y), y, u, t).$$

The plant does not need to be minimum phase, nor does it need to satisfy any other standard assumptions. So, with L_{2e} -gain related cost functions, destabilizing controllers are never retained when a stabilizing candidate controller is available, irrespective of model mismatch and irrespective of whether *standard assumptions* or other prior beliefs about the plant fail to hold. Model-mismatch instability cannot occur. The adaptive system is stable whenever the adaptive control problem is feasible. [72, 73]. Moreover unlike the early cost-detectable adaptive algorithms of Martensson [71] and of Fu and Barmish [62], unfalsified adaptive control systems adapt with optimally rapidity. Simulation studies [37] demonstrate that unfalsified adaptive systems can adapt fast enough to do real-time stabilization of open-loop unstable plants in cases where the measurement signal to noise ratio is not too great.

Using the *hysteresis switching lemma* [74], Stefanovic et al. proved the following.

Convergence and Stability [72] *A switched sequence of controllers $K(t_i)$ ($i = 1, 2, \dots$) that minimize the current unfalsified cost $V(K, u, y, t)$ at each switch-time t_i will stabilize the plant P if the cost $V(K, u, y, t)$ cost detectability property. Proof: As illustrated in Figure 7, cost detectability assures that the cost-minimizing controller tends towards one with finite cost, which implies stability when the cost has the cost-detectability property. See [72] for more a detailed proof. \square*

While Morse et al. were able to demonstrate the cost-detectability property (which they called ‘tunability’) for the adaptive methods that they examined only by introducing assumptions on the plant, the results of [72, 73] demonstrated one can directly design an unfalsified cost $V(K, u_0, y_0, t)$ to be cost detectable without regard to plant assumptions. Cost-detectably unfalsified adaptive control selectors quickly, reliably and robustly sift candidate controllers without assumptions on the plant itself.

In designs like [37], we have demonstrated that the unfalsified control approach with an L_{2e} -gain related cost function converges quickly and reliably in real time to a stabilizing controller that robustly achieves specified performance goals, often converging within a fraction of an unstable plant’s fastest unstable time constant. This speed of adaptive response means that “bursting phenomena” that plague conventional slow adaptive systems do not occur, even in the absence of persistently exciting disturbance signals. Because our unfalsified adaptive systems perform reliably irrespective of plant model mismatch, they have the potential to reliably achieve rapid real-time failure recovery for battle-damaged aircraft and similar systems.

4.0 DISCUSSION

In the unfalsified control approach to supervisory control, decisions to adapt are data-driven. Determination of which candidate control laws are suitable are made based on experimental evidence, i.e., the actual values of sensor output signals and actuator input signals. In this process the role, if any, of plant models and of probabilistic hypotheses about stochastic noise and random initial conditions is entirely an a priori role: These provide concepts which are useful in selecting the class \mathbf{K} of candidate controllers and in selecting achievable goals (i.e., selecting T_{spec}). The methods of traditional model-based control theories (root locus, stochastic optimal control, Bode-Nyquist theory, H_∞ robust control and so forth) provide mechanizations of this prior selection and narrowing process. Unfalsified control takes over where traditional model-based methods leave off, providing a mathematical framework for

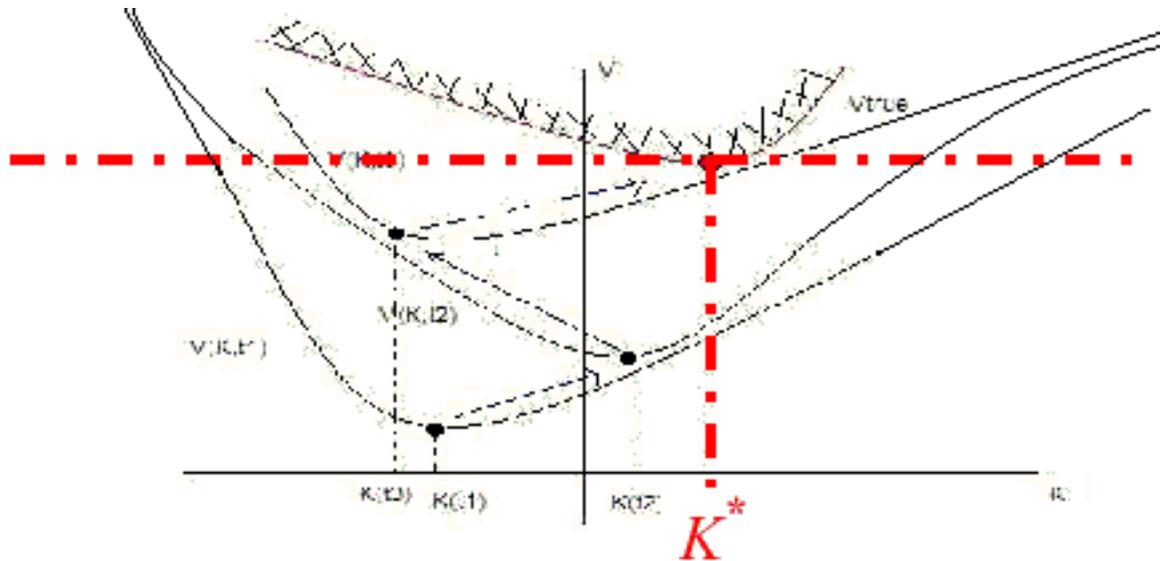


Figure 7: L_{2e} -gain related cost functions have the cost-detectability property: They are monotone in t , uniformly bounded for stabilizing K and unbounded for destabilizing K . Cost-minimizing adaptive laws robustly converge irrespective of the size of plant model-mismatch [72, 73].

determining the proper consequences of experimental observations on the choice of control law. In effect, the theory gives one a model-free mathematical “sieve” for candidate controllers, enabling us (i) to precisely identify what of relevance to attaining the specification T_{spec} can be discovered from experimental data alone and (ii) to clearly distinguish the implications of experimental data from those of assumptions and other prior information.

The **Unfalsified Control Theorem** explains the learning mechanisms of adaptive control theory. It provides an exact characterization of what can, and what cannot, be learned from experimental data about the ability of a given class of controllers to meet a given performance specification. A salient feature of the theory is that the data used to falsify a class of control laws may be either open-loop data or data obtained with other controllers in the feedback loop. Consequently, large classes of candidate controllers are falsified by even a few experimental samples of plant input output data. Candidate controllers need not be actually inserted in the feedback loop to be falsified. This is important because it means that adaptive unfalsified controllers will be significantly less susceptible to poor transient response than adaptive learning algorithms which require inserting controllers in the loop one-at-a-time to determine if they are unsuitable.

A noteworthy feature of the unfalsified control theory is its flexibility and simplicity of implementation. Controller falsification typically involves only real-time integration of algebraic functions of the observed data, with one set of integrators for each candidate controller. The theory may be readily applied to nonlinear time-varying plants, as well as to linear time-invariant ones.

5.0 CONCLUSION

As robust control theory has matured, a key challenge has been the need for a more flexible theory that provides a unified basis for representing and exploiting *evolving* real-time data. The role of *unfalsified control* is to close the loop on the adaptive and robust control design processes by developing precise data-driven methods to implement supervisory control schemes that can optimally exploit the

information in data to enhance robustness and improve performance. The types of control designs will be better able to compensate for uncertain and time-varying effects, battle damage, equipment failures and other changing circumstances.

A salient feature of unfalsified adaptive control systems is that they are robustly convergent irrespective of model mismatch, provided that the problem is feasible and one uses a cost function $J(r, y, u, t)$ with the cost-detectability property. No plant assumptions are required to assure convergence. Safe adaptive control is assured whenever the adaptive control problem is feasible. Adaptation is rapid and reliable because unfalsified control makes optimal use of information in measurements data to eliminate destabilizing controllers efficiently and quickly, selecting only controllers that are optimal with respect to data.

“It is a capital mistake to theorize before one has data. Insensibly one begins to twist facts to suit theories instead of theories to suit facts.”

Sherlock Holmes
Arthur Conan Doyle

6.0 ACKNOWLEDGMENT

I thank my current and former students T.C. Tsao, T. Brozenec, F. Cabral, P. Brugarolas, M. Jun, R. Wang, M. Stefanovic and A. Paul for their invaluable assistance in developing the theory and examples described in this paper.

7.0 REFERENCES

- [1] R. Y. Chiang and M. G. Safonov. *Robust-Control Toolbox*. MathWorks, South Natick, MA, 1988.
- [2] J. M. Maciejowski. *Multivariable Feedback Design*. Addison-Wesley, Reading, MA, 1989.
- [3] G. J. Balas, J. C. Doyle, K. Glover, A. Packard, and R. Smith. *μ -Analysis and Synthesis Toolbox (μ -Tools)*. MathWorks, Natick, MA, 1991.
- [4] S. Skogestad and I. Postlethwaite. *Multivariable Feedback Control*. Wiley, New York, 1996.
- [5] R. S. Sánchez-Peña and M. Sznaier. *Robust Systems Theory and Applications*. Wiley, New York, 1998.
- [6] G. Dullerud and f. Paganini. *A Course Robust Control Theory: A Convex Approach*. Springer-Verlag, New York, 1999.
- [7] R. L. Kosut. Adaptive calibration: An approach to uncertainty modeling and on-line robust control design. In *Proc. IEEE Conf. on Decision and Control*, pages 455–461, Athens, Greece, December 10–12, 1986. IEEE Press, New York.
- [8] R. L. Kosut. Adaptive uncertainty modeling: On-line robust control design. In *Proc. American Control Conf.*, pages 245–250, Minneapolis, MN, June 1987. IEEE Press, New York.
- [9] R. L. Kosut. Adaptive control via parameter set estimation. *Int. Journal of Adaptive Control and Signal Processing*, 2:371–399, 1988.
- [10] R. Smith and J. Doyle. Model invalidation — a connection between robust control and identification. In *Proc. American Control Conf.*, pages 1435–1440, Pittsburgh, PA, June 21-23, 1989. IEEE Press, New York.

- [11] J. M. Krause. Stability margins with real parameter uncertainty: Test data implications. In *Proc. American Control Conf.*, pages 1441–1445, Pittsburgh, PA, June 21–23, 1989. IEEE Press, New York.
- [12] R. L. Kosut, M. K. Lau, and S. P. Boyd. Set-membership identification of systems with parametric and nonparametric uncertainty. *IEEE Trans. Autom. Control*, AC-37(7):929–941, July 1992.
- [13] K. Poolla, P. Khargonekar, J. Krause A. Tikku, and K. Nagpal. A time-domain approach to model validation. In *Proc. American Control Conf.*, pages 313–317, Chicago, IL, June 1992. IEEE Press, New York.
- [14] J. J. Krause, G. Stein, and P. P. Khargonekar. Sufficient conditions for robust performance of adaptive controllers with general uncertainty structure. *Automatica*, 28(2):277–288, March 1992.
- [15] R. Smith. An informal review of model validation. In R. S. Smith and M. Dahleh, editors, *The Modeling of Uncertainty in Control Systems: Proc. of the 1992 Santa Barbara Workshop*, pages 51–59. Springer-Verlag, New York, 1994.
- [16] R. L. Kosut. Uncertainty model unfalsification: A system identification paradigm compatible with robust control design. In *Proc. IEEE Conf. on Decision and Control*, volume 1, pages 3492–3497, New Orleans, LA, December 13–15, 1995. IEEE Press, New York.
- [17] P. M. Makila. Robust control-oriented identification. In *Proceedings of IFAC symposium on system identification*, 1997.
- [18] P. M. Makila. On autoregressive models, the parsimony principle, and their use in control-oriented system identification. *International Journal of Control*, 68, 1997.
- [19] L. Giarre, B. Z. Kacewicz, and M. Milanese. Model quality evaluation in set-membership identification. *Automatica*, 33, 1997.
- [20] L. Giarre, M. Milanese, and M. Taragna. H_∞ identification and model quality evaluation. *IEEE Transaction on Automatic Control*, 42(2), February 1997.
- [21] S. R. Venkatesh and M. A. Dahleh. Multivariable system identification: A control perspective. In *Proceedings of the IFAC Symposium on Identification*, 1996.
- [22] S. R. Venkatesh and M. A. Dahleh. On learning the input-output behavior of nonlinear fading memory systems from finite data. *Int J. Robust and Nonlinear Control*, 2000. To appear.
- [23] P. V. Zhivoglyadov, R. H. Middleton, and M. Fu. Localization based switching adaptive control for time-varying discrete-time systems. *IEEE Trans. Autom. Control*, AC-45(4):752–755, April 2000.
- [24] P. V. Zhivoglyadov, R. H. Middleton, and M. Fu. Further results on localization-based switching adaptive control. *Automatica*, 37(2):257–263, February 2001.
- [25] M. Milanese and L. Giarre. Optimality, approximation, and complexity in set membership H_∞ identification. *IEEE Trans. Autom. Control*, AC-47(10), October 2002.
- [26] A. Garulli, L. Giarre, and G. Zappa. Identification of approximated Hammerstein models in a worst-case setting. *IEEE Trans. Autom. Control*, AC-47(12), December 2002.
- [27] M. G. Safonov and T.-C. Tsao. The unfalsified control concept: A direct path from experiment to controller. In B. A. Francis and A. R. Tannenbaum, editors, *Feedback Control, Nonlinear Systems and Complexity*, pages 196–214. Springer-Verlag, Berlin, 1995.

- [28] M. G. Safonov and T. C. Tsao. The unfalsified control concept and learning. In *Proc. IEEE Conf. on Decision and Control*, pages 2819–2824, Lake Buena Vista, FL, December 14–16, 1994. IEEE Press, New York.
- [29] M. G. Safonov. Unfalsified control: A direct path from experiment to controller. In *1995 AFOSR Workshop on Dynamics and Control*, Minneapolis, MN, June 5-7, 1995.
- [30] T.-C. Tsao and M. G. Safonov. Adaptive robust manipulator trajectory control — An application of unfalsified control method. In *Proc. Fourth International Conference on Control, Automation, Robotics and Vision*, Singapore, December 3–6, 1996.
- [31] M. G. Safonov and T. C. Tsao. The unfalsified control concept and learning. *IEEE Trans. Autom. Control*, AC-42(6):843–847, June 1997.
- [32] T. C. Tsao and M. G. Safonov. Unfalsified direct adaptive control of a two-link robot arm. *Int. J. Adaptive Control and Signal Processing*, 15:319–334, May 2001.
- [33] T. F. Brozenec and M. G. Safonov. Controller identification. In *Proc. American Control Conf.*, pages 2093–2097, Albuquerque, NM, June 4–6, 1997.
- [34] M. G. Safonov and F. B. Cabral. Fitting controllers to data. *Systems and Control Letters*, 43(4):299–308, July 23, 2001.
- [35] F. B. Cabral and M. G. Safonov. Fitting controllers to data. In *Proc. American Control Conf.*, pages 589–593, Philadelphia, PA, June 24-26, 1998.
- [36] P. B. Brugarolas, V. Fromion, and M. G. Safonov. Robust switching missile autopilot. In *Proc. American Control Conf.*, pages 3665–3669, Philadelphia, PA, June 24-26, 1998.
- [37] M. Jun and M. G. Safonov. Automatic PID tuning: An application of unfalsified control. In *Proc. IEEE CCA/CACSD*, pages 328–333, Kohala Coast–Island of Hawaii, HI, August 22-27, 1999.
- [38] R. L. Kosut. Iterative unfalsified adaptive control: Analysis of the disturbance free case. In *Proc. American Control Conf.*, pages 566–570, San Diego, CA, June 2–4 1999. IEEE Press, New York.
- [39] E. G. Collins and C. Fan. Automated PI tuning for a weigh belt feeder via unfalsified control. In *Proc. IEEE Conf. on Decision and Control*, pages 785–789, Phoenix, AZ, December 1999. IEEE Press, New York.
- [40] T. R. Kurfess and H. A. Razavi. Real time force control of a hydraulic drive using model reference unfalsification concepts and learning. Technical report, Mechanical Engineering Dept., Georgia Institute of Technology, December 1999.
- [41] H. A. Razavi and T. R. Kurfess. Real time force control of a grinding process using unfalsification and learning concept. In *Proc. Intl. Mechanical Engineering Congress and Exposition*, Orlando, FL, November 5–10, 2000.
- [42] J. W. Polderman. Sequential continuous time adaptive control: A behavioral approach. In *Proc. IEEE Conf. on Decision and Control*, Sydney, Australia, December 12–15, 2000. IEEE Press, New York.
- [43] M. G. Safonov. Recent advances in robust control, feedback and learning. In S. O. R. Moheimani, editor, *Perspectives in Robust Control*, pages 283–296. Springer-Verlag, London, 2001.
- [44] T. F. Brozenec, T. C. Tsao, and M. G. Safonov. Controller validation. *Int. J. Adaptive Control and Signal Processing*, 15(5):431–444, August 2001.

- [45] R. L. Kosut and M. G. Safonov. Editorial, adaptive control, with confidence. *Int. J. Adaptive Control and Signal Processing*, 15(5):429, August 2001.
- [46] M. G. Safonov. Unfalsified control: A behavioral approach to learning and adaptation. In *Proc. IEEE Conf. on Decision and Control*, pages 2682–2685, Orlando, FL, December 4–7, 2001. IEEE Press, New York.
- [47] M. Jun and M. G. Safonov. Controller parameter adaptation algorithm using unfalsified control theory and gradient method. In *Proc. IFAC World Congress*, Barcelona, Spain, July 21–26, 2002. Pergamon Press, Oxford.
- [48] P. Brugarolas and M. G. Safonov. A data driven approach to learning dynamical systems. In *Proc. IEEE Conf. on Decision and Control*, pages 4670–4675, Las Vegas, NV, December 10–13, 2002. IEEE Press, New York.
- [49] D. Angeli and E. Mosca. Lyapunov-based switching supervisory control of nonlinear uncertain systems. *IEEE Trans. Autom. Control*, AC-47(3):500–505, March 2002.
- [50] T. Agnoloni and E. Mosca. Controller falsification based on multiple models. *Int. J. Adaptive Control and Signal Processing*, 17(2):163–177, March 2003.
- [51] A. Paul and M. G. Safonov. Model reference adaptive control using multiple controllers & switching. In *Proc. IEEE Conf. on Decision and Control*, Maui, HI, December 9–12, 2003. IEEE Press, New York.
- [52] F. B. Cabral and M. G. Safonov. Unfalsified model reference adaptive control using the ellipsoid algorithm. In *Proc. IEEE Conf. on Decision and Control*, Maui, HI, December 9–12, 2003. IEEE Press, New York.
- [53] F. B. Cabral and M. G. Safonov. Fitting controllers to data: The MIMO case. Technical report, EE Systems Dept., Univ. of Southern Calif., March 1, 2003.
- [54] T.-C. Tsao, T. Brozenec, and M. G. Safonov. Unfalsified adaptive spacecraft attitude control. In *Proc. AIAA Guidance, Navigation and Control Conf.*, Austin, TX, August 11–14, 2003.
- [55] M. Saeki. Unfalsified control approach to parameter space design of PID controllers. In *Proc. IEEE Conf. on Decision and Control*, Maui, HI, December 9–12, 2003, submitted. IEEE Press, New York.
- [56] T. Mitchell. Version spaces: A candidate elimination approach to rule learning. In *Proc. Fifth Int. Joint Conf. on AI*, pages 305–310, Cambridge, MA, 1977. MIT Press.
- [57] T. Mitchell. *Machine Learning*. McGraw-Hill, Inc., 1997.
- [58] I.W. Sandberg. On the L_2 -boundedness of solutions of nonlinear functional equations. *Bell System Technical Journal*, 43(4):1581–1599, July 1964.
- [59] G. Zames. On the input–output stability of time-varying nonlinear feedback systems — Part I: Conditions derived using concepts of loop gain, conicity, and positivity. *IEEE Trans. Autom. Control*, AC-11(2):228–238, April 1966.
- [60] A. S. Morse. Supervisory control of families of linear set-point controllers — Part I: Exact matching. *IEEE Trans. Autom. Control*, AC-41(10):1413–1431, October 1996.
- [61] D. Borrelli, A. S. Morse, and E. Mosca. Discrete-time supervisory control of families of two-degrees-of-freedom linear set-point controllers. *IEEE Trans. Autom. Control*, AC-44(1):178–181, January 1999.

- [62] M. Fu and B. R. Barmish. Adaptive stabilization of linear systems via switching control. *IEEE Trans. Autom. Control*, AC-31(12):1097–1103, December 1986.
- [63] A. S. Morse. Supervisory control of families of linear set-point controllers — Part II: Robustness. In *Proc. IEEE Conf. on Decision and Control*, pages 1750–1755, New Orleans, LA, December 13–15, 1995. IEEE Press, New York.
- [64] K. S. Narendra and A. M. Annaswamy. *Stable Adaptive Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [65] P. Ioannou and J. Sun. *Stable and Robust Adaptive Control*. Prentice-Hall, Englewood-Cliffs, NJ, 1994.
- [66] T.-C. Tsao and M. G. Safonov. Unfalsified direct adaptive control of a two-link robot arm. In *Proc. IEEE CCA/CACSD*, Kohala Coast–Island of Hawaii, HI, August 22–27, 1999.
- [67] B. R. Woodley, J. P. How, and R. L. Kosut. Direct unfalsified controller design — solution via convex optimization. In *Proc. American Control Conf.*, pages 3302–3306, San Diego, CA, June 2–4 1999. IEEE Press, New York.
- [68] H. A. Razavi and T. R. Kurfess. Force control of a reciprocating surface grinder using unfalsification and learning concept. *Int. J. Adaptive Control and Signal Processing*, 15(5):503–518, August 2001.
- [69] C. E. Rohrs, L. Valavani, M. Athans, and G. Stein. Robustness of adaptive control algorithms in the presence of unmodeled dynamics. *IEEE Trans. Autom. Control*, AC-30(9):881889, September 1985.
- [70] R. Wang, M. Stefanovic, and M. G. Safonov. Unfalsified direct adaptive control using multiple controllers. In *Proc. AIAA Guidance, Navigation and Control Conf.*, Providence, RI, August 16–19, 2004.
- [71] B. Mrtenson. The order of any stabilizing regulator is sufficient a priori information for adaptive stabilization. *Systems and Control Letters*, 6(2):87–91, February 1985.
- [72] M. Stefanovic, W. Wang, and M. G. Safonov. Stability and convergence in adaptive systems. In *Proc. American Control Conf.*, Boston, MA, June 30 – July 2, 2004. IEEE Press, New York.
- [73] Michael G. Safonov, Margareta Stefanovic, and Ayanendu Paul. Safe adaptive switching through an infinite controller set: Stability and convergence. In *Proc. IFAC World Congress*, Prague, Czech Republic, July 4–8, 2005, to appear. Pergamon Press, Oxford.
- [74] A. Stephen Morse, David Q. Mayne, and Graham C. Goodwin. Applications of hysteresis switching in parameter adaptive control. *IEEE Trans. Autom. Control*, 37(9):1343–1354, September 1992.