The Unfalsified Control Concept and Learning

Michael G. Safonov and Tung-Ching Tsao

Abstract— **Without a plant model or other prejudicial assumptions, theory is developed for identifying control laws which are consistent with performance objectives and past experimental data — possibly before the control laws are ever inserted in the feedback loop. The theory complements model-based methods such as H-infinity robust control theory by providing a precise characterization of how the set of suitable controllers shrinks when new experimental data is found to be inconsistent with prior assumptions or earlier data. When implemented in real time, the result is an adaptive switching controller. An example is included.**

.

I. Introduction

Commenting on the limits of the knowable, philosopher Karl Popper [1] said, "The Scientist ... can never know for certain whether his theory is true, although he may sometimes establish ... a theory is false." Discovery in science is a process of elimination of hypotheses which are falsified by experimental evidence. The present paper concerns how Popper's falsification concept may be applied to the development of a theory for discovering good controllers from experimental data without reliance on feigned hypotheses or prejudicial assumptions about the plant, sensors, uncertainties or noises.

Closely related previous efforts [2]–[10] to develop carefully reasoned methods for incorporating experimental data into the control design process have centered on an indirect two-step decomposition of the problem involving (i) identifying of a plant model and uncertainty bounds, then (ii) designing a robust controller for the uncertain model. These methods have the desirable property that they lead to controllers that are *not-demonstrably-unrobust*, based solely on the available experimental evidence. However, such methods fall short of the goal of providing an *exact* mathematical characterization of the class of not-demonstrably-unrobust controllers. The problem is that assumed uncertainty-structures lead to uncertain models that upper-bound — but do not exactly characterize — plant behaviors observed in the data. This results in over-designed controllers whose ability to control is unfalsified not only by the observed experimental data, but also by other behaviors of which the plant is incapable.

In the present paper, we dispense with assumptions about uncertainty structure and tackle the problem of directly identifying the set of controllers which, based on experimental data alone, are not-demonstrably-unrobust. Our approach is to look for inconsistencies between the constraints on signals each candidate controller would introduce if inserted in the feedback loop and the constraints with associated performance objectives and constraints implied by the experimental data. When inconsistencies exist among these constraints, then the controller is said to be falsified; otherwise it is unfalsified. In essence, it is a feedback generalization of the open-loop model validation schemes used in control-oriented identification. The result is a paradigm for direct identification of not-demonstrably-unrobust controllers which we call the *unfalsified control concept*.

A preliminary version of our unfalsified control work appeared in [11]. Connections with model validation are further examined in [12]. Reference [13] offers a non-technical discussion of some of the broader implications of our theory with regard to an age old philosophical debate on the relative merits of observation verses prior knowledge in science.

II. Learning

Consider the feedback control system in Figure 1. Our goal is to determine a control law K for a plant P so that the closed-loop system response, say T, satisfies a specification requiring that, for all command inputs $r \in \mathcal{R}$, the triple (r, y, u) be in a given specification set T_{spec}. The need for learning (i.e., controller identification) arises when the plant and, hence, the solutions (r, y, u) are either unknown or are only partially known and one wishes to extract information from measurements which will be helpful in selecting a suitable control law K. Learning takes place when the available experimental evidence enables one to falsify a hypothesis about feedback controller's ability to meet a performance specification. Formally, we introduce the following definition.

Definition : A controller $K \in \mathbf{K}$ 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**. □

The future is never certain and even the past may be imprecisely known. We consider the situation that arises when experimental observations give incomplete information about (u, y). More precisely, we suppose that all that can be deduced from measurements available at time $t = \tau$ is a set, say M_{data}, containing the actual plant input-output pair (u, y). For example, in this case of perfect past measurements $(u_{data}(t), y_{data}(t))$ for $t \leq \tau$, the set M_{data} becomes

$$M_{data} = \left\{ \ (u, y) \in \mathcal{U} \times \mathcal{Y} \ \middle| \ \mathsf{P}_\tau \begin{bmatrix} (u - u_{data}) \\ (y - y_{data}) \end{bmatrix} = 0 \ \right\} \qquad (1)$$

and P_τ is the familiar time-truncation operator of input-output stability theory [14], [15], viz.,

$$[\mathsf{P}_\tau x](t) \triangleq \begin{cases} x(t), & \text{if } t \leq \tau \\ 0, & \text{if } t > \tau \end{cases}$$

It is convenient to associate with the set M_{data} the following embedding in $\mathcal{R} \times \mathcal{Y} \times \mathcal{U}$.

Definition (Measurement Information Set) The set

$$P_{data} \triangleq \left\{ \ (r, y, u) \in \mathcal{R} \times \mathcal{Y} \times \mathcal{U} \ \middle| \ (u, y) \in M_{data} \ \right\}$$

is called the **measurement information set** at time τ. □
The significance of P_{data} is it that, if one were given *only* experimental measurement information (viz. $(u, y) \in M_{data}$), then the strongest statement one could make about the triple (r, y, u) is that $(r, y, u) \in P_{data}$. Of course, stronger statements are possible when a known control law or a partially known plant model introduces further constraints on the triple (r, y, u). The present paper is primarily concerned with statements which can be made without the plant model information.

The unfalsified control problem may be formally stated follows.

Problem 1 (Unfalsified Control) Given

a) a measurement information set $P_{data} \subset \mathcal{R} \times \mathcal{Y} \times \mathcal{U}$,

This work was supported in part by the U.S. Air Force Office of Scientific Research under grants F49620-92-J-0014 and F49620-95-I-0095.

Michael G. Safonov is with Dept. of Electrical Engineering—Systems, University of Southern California, Los Angeles, CA 90089–2563 USA. Phone 1-213-740-4455. FAX 1-213-740-4449. Email msafonov@usc.edu.

T. C. Tsao was with University of Southern California, Los Angeles, CA. He is with Spectrum Astro, Inc., 3601 Aviation Blvd., Manhattan Beach, CA 90266 USA.

b) a performance specification set $T_{spec} \subset \mathcal{R} \times \mathcal{Y} \times \mathcal{U}$, and
c) a class $\mathbf{K}$ of admissible control laws,

determine the subset of $\mathbf{K}_{OK}$ of control laws $K \in \mathbf{K}$ whose ability to meet the specification $(r, y, u) \in T_{spec}\ \forall r \in \mathcal{R}$ is not falsified by the measurement information P_{data}. □

We consider the plant P and the controller K to be relations in $\mathcal{R} \times \mathcal{Y} \times \mathcal{U}$ (cf. [15], [16]):

$$P \subset \mathcal{R} \times \mathcal{Y} \times \mathcal{U}, \quad K \subset \mathcal{R} \times \mathcal{Y} \times \mathcal{U}$$

That is, both the plant P and the controller K are regarded as a locus or "graph" in $\mathcal{R} \times \mathcal{Y} \times \mathcal{U}$, viz.,

$$P = \left\{ \ (r, y, u) \in \mathcal{R} \times \mathcal{Y} \times \mathcal{U} \ \middle| \ y = Pu \ \right\}$$

$$K = \left\{ \ (r, y, u) \in \mathcal{R} \times \mathcal{Y} \times \mathcal{U} \ \middle| \ u = K \begin{bmatrix} r \\ y \end{bmatrix} \ \right\}.$$

On the other hand, experimental information P_{data} from a plant corresponds to *partial* knowledge of the constraint P — viz., an "interpolation constraint" on the graph of P.

When all that is known about a plant P is measurement information P_{data}, the following theorem provides a basis for the solution to Problem 1.

Theorem 1 (Unfalsified Control) Consider Problem 1. 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 (u_1, y_1) such that

$$(r_0, y_1, u_1) \in P_{data} \cap K \cap T_{spec}. \tag{2}$$

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 (2) holds. ■

Theorem 1 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:

- Theorem 1 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.
- Theorem 1 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 Section V illustrates.

III. Adaptive Control

The step from Theorem 1 to adaptive control is, conceptually at least, a small one. Simply choosing as the current control law one that is not falsified by the past data produces 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 [17], [18], 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 [17]. The reason is that, unlike the theory in [17], unfalsified control theory efficiently eliminates broad classes of controllers before they are ever inserted in the feedback loop. The main difference between unfalsified 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 (1) 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 algorithm.

Algorithm 1 (Recursive Adaptive Control)
Input:
- A finite set $\mathbf{K}$ of m candidate dynamical controllers $K_i(r, y, u) = 0$, $(i = 1, \ldots, m)$ 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 the inequalities

$$\int_0^{k\Delta t} \tilde{T}_{spec}(r(t), y(t), u(t), t)\, dt \leq \ 0, \quad \forall k = 1, \ldots, n.$$

Initialize:
set $k = 0$, set $\widehat{i} = m$;
for $i = 0 : m$, set $s(i) = 1$, set $\tilde{J}(i) = 0$, end.
Procedure:
while $\widehat{i} > 0$;
 $k = k + 1$;
 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;

$\widehat{i} = \max\{\ i \mid s(i) > 0\ \}$;
end.

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

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

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

We stress 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 Algorithm 1 by *fixing* $\tau = \tau_0$ and regarding $t - \tau_0$ as the deviation from the current time. The result is a 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.

IV. Practical Considerations

Practical application of Theorem 1 requires that one have characterizations of the sets $\mathbf{K}$ and T_{spec} which are both simple and amenable to computations. In the control field, experience has shown that linear equations and quadratic cost functions often lead to tractable problems. A linear parameterization of the set $\mathbf{K}$ of admissible control laws is possible by representing each $K \in \mathbf{K}$ as a sum of filters, say $Q_i(s)$, so that the control laws $K \in \mathbf{K}$ are linearly parameterized by an unspecified vector $\theta \in \mathbb{R}^n$; i.e.,

$$K_\theta = \{\ (r, y, u) \mid K_\theta(r(s), y(s), u(s)) = 0\ \} \qquad (3)$$

where the argument (s) indicates Laplace transformation and

$$K_\theta(r(s), y(s), u(s)) = \theta' Q(s) \begin{bmatrix} r(s) \\ y(s) \\ u(s) \end{bmatrix} \ \forall \theta \in \mathbb{R}^n. \qquad (4)$$

Thus,

$$\mathbf{K} = \bigcup_{\theta \in \mathbb{R}^n} \{K_\theta\}.$$

The performance specification set T_{spec} might be selected to be a collection of quadratic inequalities, expressed in terms of L_2 inner-products weighted by a given transfer function matrix, say $T_{spec}(s)$; e.g.,

$$T_{spec} = \text{sector}(T_{spec}) \triangleq$$
$$\{\ (r, y, u) \in L_{2e} \mid \langle z_1, z_2 \rangle_\tau \ \le 0 \ \forall \tau \in [0, \infty)\ \} \qquad (5)$$

where

$$\langle z_1, z_2 \rangle_\tau \triangleq \langle \mathsf{P}_\tau z_1, \mathsf{P}_\tau z_2 \rangle_{L_2[0,\infty)}$$

and

$$\begin{bmatrix} z_1(s) \\ z_2(s) \end{bmatrix} = T_{spec}(s) \begin{bmatrix} r(s) \\ y(s) \\ u(s) \end{bmatrix}.$$

Here, as elsewhere, the argument (s) indicates Laplace transformation. Note that sets of the type sector(T_{spec}) have been studied by Safonov [16]; such sets are a generalization of the Zames-Sandberg [14], [15] L_{2e} conic sector

$$\text{sector}(a, b) \triangleq \{\ (u, y) \mid \langle y - au, y - bu \rangle_\tau \ \le 0 \ \forall \tau \ge 0\ \}.$$

The example in the following section illustrates some of the foregoing ideas.

V. Example

In this section we describe a simulation study involving an adaptive unfalsified control design based on Algorithm 1. Simulation data is generated in closed-loop operation by the following time-varying model:

$$P = (G)\,(1 + \Delta(s)) \qquad (6)$$

where $\Delta(s) = \frac{2s^2+2s+10}{s^2+2s+100}$ and G is an unstable time-varying system with its input u and output x satisfying $\frac{dx}{dt}(t) = (1 + 0.5\sin 8t)x(t) + (1 + 0.5\sin 20t)u(t)$. Note that the control design itself is entirely model-free in that the control design proceeds without any specific information about the above plant model other than past input/output data. The simulation shows that the algorithm converges after a finite number of switches to a linear time-invariant control law in the unfalsified set $\mathbf{K}_{OK}$.

The performance specification set T_{spec} is taken to be the set of $(r, y, u) \in L_{2e} \times L_{2e} \times L_{2e}$ which for all $\tau \ge 0$ satisfy the inequality (cf. eqn. (5))

$$\|w_1 * (y - r)\|^2_{L_2[0,\tau]} + \|w_2 * u\|^2_{L_2[0,\tau]} \le \|r\|^2_{L_2[0,\tau]}. \qquad (7)$$

It says that the error signal $r - y$ and the control signal u should be "small" compared to the command signal r; the dynamical "weights" w_1 and w_2 determine what is small. If the plant were linear time-invariant, it would be equivalent to the familiar H^∞ weighted mixed-sensitivity performance criterion (e.g., [19], [20])

$$\left\| \begin{bmatrix} W_1 S \\ W_2 K S \end{bmatrix} \right\|_\infty \le 1$$

where $S \triangleq 1/(1 + PK)$ is the Bode sensitivity function.

In (7), $w_1(t)$ and $w_2(t)$ are the impulse responses of stable minimum phase weighting transfer functions $W_1(s)$ and $W_2(s)$ respectively, "$*$" means convolution, r is the input reference signal, y is the plant output signal, and u is the control signal.

In keeping with (3)–(4), the set $\mathbf{K}$ of admissible controllers is chosen to be

$$0 = \theta' Q(s) \begin{bmatrix} r(s) \\ y(s) \\ u(s) \end{bmatrix} \qquad (8)$$

where $\theta \in \mathbb{R}^5$,

$$Q(s) = \begin{bmatrix} 0 & 0 & H(s) \\ 0 & H(s) & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix} \quad \text{and} \quad H(s) = \frac{.5}{s + .5}.$$

See Figure 2. Without loss of generality, we *chose* $\theta_5 = 1$. Note that there is not any special motivation for the particular forms of $Q(s)$ and $H(s)$ given above, save that they happen to be such that for each admissible control gain vector θ, it easy to solve (8) uniquely for the corresponding past $\mathsf{P}_\tau r$ of $r(t)$ in terms of θ and the past plant data data $u(t), y(t)$. Also, we chose not to consider more than five adjustable control gains θ_i, to ensure that the computations are tractable.

The simulation was conducted as follows. At each time τ a control law in $K_{\hat{\theta}(\tau)} \in \mathbf{K}$ was connected to the simulation model (6) where $\hat{\theta}(\tau)$ denotes the value of θ associated with the controller in use at time τ. The value of $\hat{\theta}(\tau)$ was held constant until such time as it is falsified by the past data $(\mathsf{P}_\tau u, \mathsf{P}_\tau y)$, then it is switched to a point near the geometric center of the current unfalsified controller parameter set. The controller state was left unchanged at these switching times. The following were used in the simulation:

- $W_1(s) = \frac{s+3.5}{s+.35}$, $W_2(s) = \frac{.01}{(s+1)^3}$
- Unit step command: $r(t) = 1 \; \forall t \geq 0$
- Initial conditions at time $\tau = 0$ are all zero.
- Set intersections are performed every 0.2 sec.
- We arbitrarily restricted our search for unfalsified control laws as follows:

$$\hat{\theta}_1^2 + \hat{\theta}_2^2 \leq 100^2, \quad |\hat{\theta}_3| \leq 350, \quad 100 \leq \hat{\theta}_4 \leq 700$$

- We arbitrarily initialized: $\hat{\theta}(0) = [0, 0, 0, 400, 1]^T$

The simulation was carried out by using Simulink. The results are shown in Figure 3. From the plots, one can see that the parameter vector switches three times (at 0.2, 2, 82.4 sec), and despite the time-varying perturbations the vector converges within finite time to the steady-state value: $\lim_{t\to\infty} \hat{\theta}(t) = [2.0009, -31.471, -218.75, 250.00, 1]^T$. The convergent nominal steady-state command-to-error transfer function $T_{r-y,r}$, with nominal plant $\frac{1}{s-1}\frac{2s^2+2s+10}{s^2+2s+100}$, satisfies $|T_{r-y,r}(j\omega)| < |W_1^{-1}(j\omega)|, \forall \omega$. The initial overshoot of plant output is due to the instability of the open loop system and that the first control parameter updating can only occur after performing the first set intersection operation at the time 0.2 sec. Although the performance specification is achieved after about 72 sec, the ratio

$$\frac{1}{\|r\|_{L_2[0,t]}} \sqrt{\|W_1(y-r)\|^2_{L_2[0,t]} + \|W_2 u\|^2_{L_2[0,t]}}$$

approaches one rapidly, which means good transient behavior. Further details about finite time parameter convergence and performance of unfalsified control systems can be found in [10].

Figure 4 shows the evolution of the unfalsified controller parameter set. Since a set in $\mathbb{R}^5$ cannot be visualized, the projection of the set to $[\theta_1, \theta_2, \theta_3]$-space with $\theta_4 = 260$ and $\theta_5 = 1$ is shown.

VI. Conclusions

Departing from the traditional practice in mathematical control theory of identifying assumptions required to prove desired conclusions, we have focused attention instead on what is knowable from experimental data alone without assuming anything about the plant. This does not mean that we believe that plant models ought to be abandoned. But our results do demonstrate that there is merit in asking what information and conclusions can be derived without them.

In our unfalsified control theory, decisions about which control laws are suitable are made based on 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, and so forth) provide mechanizations of this prior selection and narrowing process. Unfalsified control takes over where the model-based methods leave off, providing a mathematical framework for 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 concept embodied in Theorem 1 is of importance to adaptive control theory because 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, controllers need not be actually inserted in the feedback loop to be falsified. This is important because it means that adaptive unfalsified controllers should be significantly less susceptible to poor transient response than adaptive 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.

References

[1] K. R. Popper, *Conjectures and Refutations: The Growth of Scientific Knowledge*, Routledge, London, 1963.

[2] R. L. Kosut, "Adaptive calibration: An approach to uncertainty modeling and on-line robust control design," in *Proc. IEEE Conf. on Decision and Control*, Athens, Greece, December 10–12, 1986, pp. 455–461, IEEE Press, New York.

[3] R. L. Kosut, M. K. Lau, and S. P. Boyd, "Set-membership identification of systems with parametric and nonparametric uncertainty," *IEEE Trans. on Automatic Control*, vol. AC-37, no. 7, pp. 929–941, July 1992.

[4] K. Poolla, P. Khargonekar, J. Krause A. Tikku, and K. Nagpal, "A time-domain approach to model validation," in *Proc. American Control Conf.*, Chicago, IL, June 1992, pp. 313–317, IEEE Press, New York.

[5] J. J. Krause, G. Stein, and P. P. Khargonekar, "Sufficient conditions for robust performance of adaptive controllers with general uncertainty structure," *Automatica*, vol. 28, no. 2, pp. 277–288, March 1992.

[6] R. Smith and J. Doyle, "Model invalidation — a connection between robust control and identification," *IEEE Trans. on Automatic Control*, vol. AC-37, no. 7, pp. 942–952, July 1992.

[7] R. Smith, "An informal review of model validation," in *The Modeling of Uncertainty in Control Systems: Proc. of the 1992 Santa Barbara Workshop*, R. S. Smith and M. Dahleh, Eds., pp. 51–59. Springer-Verlag, New York, 1994.

[8] M. M. Livestone, M. A. Dahleh, and J. A. Farrell, "A framework for robust control based model invalidation," in *Proc. American Control Conf.*, Baltimore, MD, June 29–July 1, 1994, pp. 3017–3020, IEEE Press, New York.

[9] J. C. Willems, "Paradigms and puzzles in the theory of dynamical systems," *IEEE Trans. on Automatic Control*, vol. AC-36, no. 3, pp. 259–294, March 1991.

[10] T. C. Tsao, *Set Theoretic Adaptor Systems*, Ph.D. thesis, University of Southern California, May 1994, Supervised by M. G. Safonov.

[11] M. G. Safonov and T.-C. Tsao, "The unfalsified control concept: A direct path from experiment to controller," in *Feedback Control, Nonlinear Systems and Complexity*, B. A. Francis and A. R. Tannenbaum, Eds., pp. 196–214. Springer-Verlag, Berlin, 1995.

[12] T. F. Brozenec and M. G. Safonov, "Controller invalidation," in *Proceedings IFAC World Congress*, San Francisco, CA, USA, July 19–23, 1996.

[13] M. G. Safonov, "Focusing on the knowable: Controller invalidation and learning," in *Control Using Logic-Based Switching*, A. S. Morse, Ed. Springer-Verlag, Berlin, 1996.

[14] I.W. Sandberg, "On the L_2-boundedness of solutions of nonlinear functional equations," *Bell System Technical Journal*, vol. 43, no. 4, pp. 1581–1599, July 1964.

[15] 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. on Automatic Control*, vol. AC-15, no. 2, pp. 228–238, April 1966.

[16] M. G. Safonov, *Stability and Robustness of Multivariable Feedback Systems*, MIT Press, Cambridge, MA, 1980.

[17] M. Fu and B. R. Barmish, "Adaptive stabilization of linear systems via switching control," *IEEE Trans. on Automatic Control*, vol. AC-31, no. 12, pp. 1097–1103, December 1986.

[18] A. S. Morse, "Supervisory control of families of linear set-point controllers — Part II: Robustness," in *Proc. IEEE Conf. on Decision and Control*, New Orleans, LA, December 13–15, 1995, pp. 1750–1755, IEEE Press, New York.

[19] R. Y. Chiang and M. G. Safonov, *Robust-Control Toolbox*, Mathworks, South Natick, MA, 1988.

[20] J. M. Maciejowski, *Multivariable Feedback Design*, Addison-Wesley, Reading, MA, 1989.

List of Figures

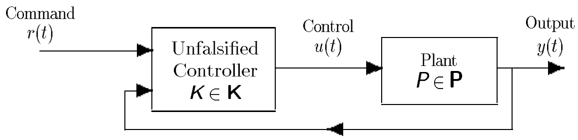

Fig. 1. Learning control system.

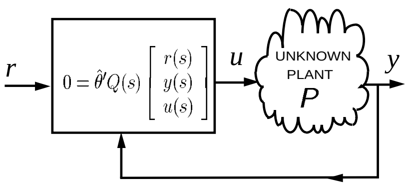

Fig. 2. Control structure.

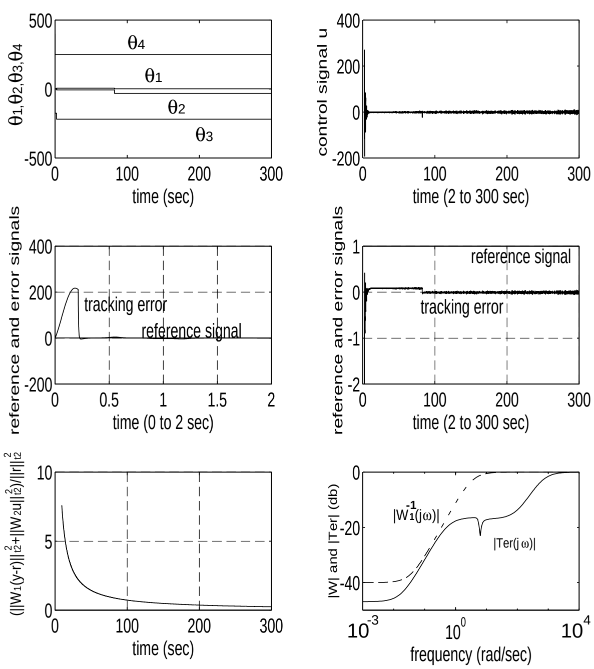

Fig. 3. Simulation results.

p]

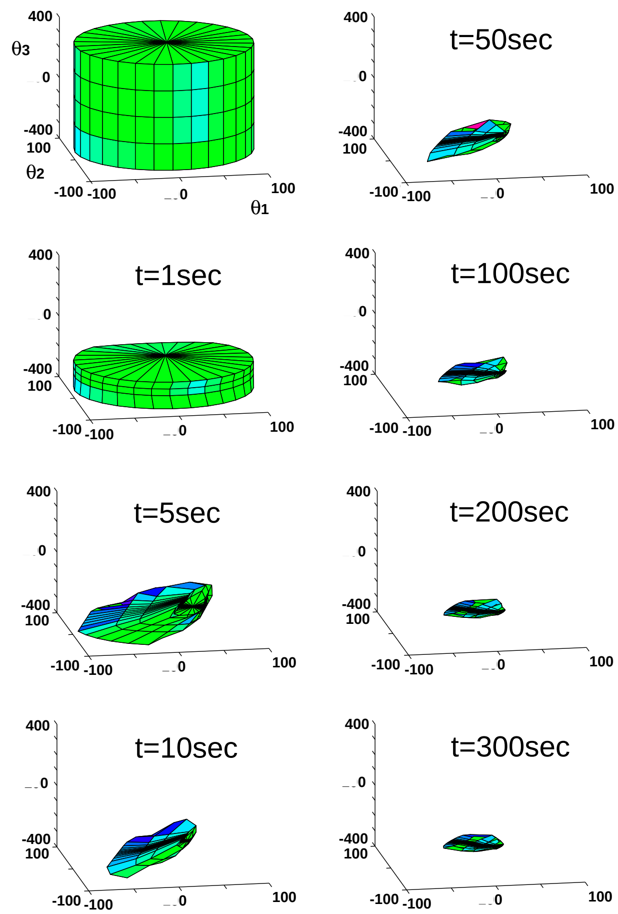

Fig. 4. Evolution of unfalsified controller parameter set.