

DATA-DRIVEN ROBUST CONTROL DESIGN: UNFALSIFIED CONTROL *

Michael G. Safonov[†]
Dept. of Electrical Engineering
Univ. of Southern California
Los Angeles, CA 90089-2563
USA

Abstract

Feedback control systems for aerospace applications must maintain 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 also be capable of learning and adapting when accumulating data indicates that previous models must be abandoned and that existing control strategies must be changed. We present recent developments that address the need for data-driven design methods well suited to situations in which available mathematical models are poor or unreliable. These innovative data-driven design methods, collectively known as unfalsified control theory, facilitate the creation of robust control systems that learn, discover and evolve in real time in order to rapidly and reliably compensate for the effects of battle damage, equipment failures and other changing circumstances. Potential applications include aircraft stability augmentation systems, highly maneuverable aircraft design, missile guidance systems, and precision pointing and tracking systems.

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

Arthur Conan Doyle, Sherlock Holmes

1 Introduction

The robust multivariable control theory that has evolved over the past quarter century includes methods based on the H_∞ μ/K_m -synthesis, and BMI/LMI/IQC theories. The robust control theory offers a major improvement over earlier algebraic and optimal control methods. It has enabled the design of controllers with greater tolerance of uncertainty in system model and, hence, increased reliability. Commercial computer-aided control synthesis tools like those introduced by Chiang and Safonov [20, 21], Balas et al. [22] and Gahinet et al. [23] have made robust control synthesis routine, and because of this aerospace and industrial applications have now become commonplace. Further, on-going improvements based on LMI/IQC robust control problem formulations are continuing to expand the range of problems that can be cast and solved within the robust control framework (e.g., [24]).

*Research supported in part by AFOSR F49620-98-1-0026 and F49620-01-1-0302.

[†]email msafonov@usc.edu; web <http://routh.usc.edu>

Yet, despite the assurances of greater uncertainty tolerance and better reliability, the existing H_∞ μ/K_m -synthesis, and BMI/LMI/IQC techniques for robust control design have an Achilles heel: They are introspective theories. They derive their conclusions based on assumed prior knowledge of models and uncertainties. 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.

Data-driven design tools are needed to make the overall robust control design process more complete and reliable. Ideally, these tools 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. Recent thrusts in this direction are control-oriented identification theory and [25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44] and, more recently, unfalsified control [45, 46, 47, 48, 49, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 14, 16] While both theories are concerned with the difficult problem of assimilating real-time measurement data into the otherwise introspective process of robust control design, the unfalsified control theory is a particular interest because it directly and precisely characterizes the control design implications of experimental data.

2 Data-Driven Robust Control Design

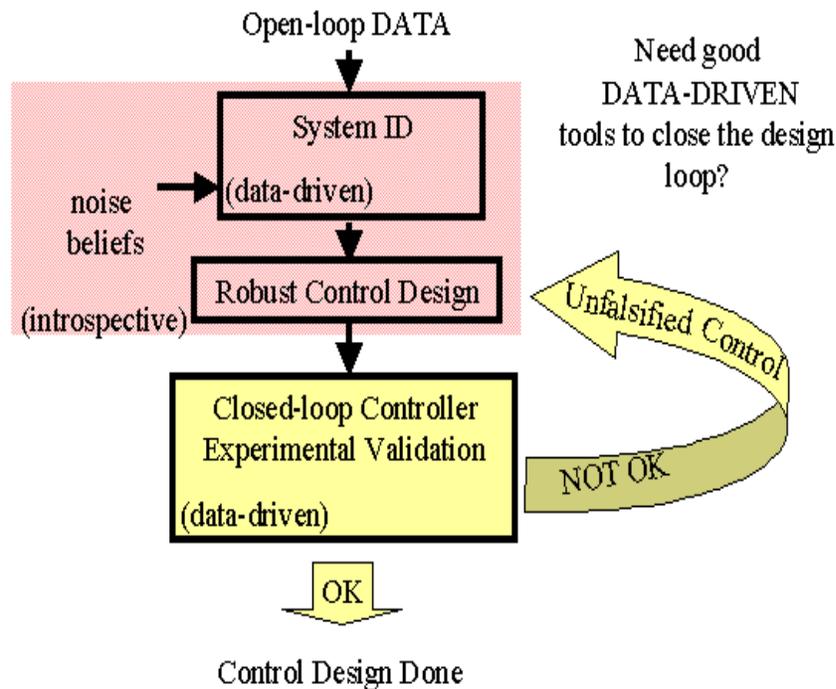


Figure 1: The data-driven theory of unfalsified control closes data-driven portion of the design loop by focusing squarely and precisely on the control design implications of data.

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 [49].

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 [50, 51]. The theory concerns the feedback control configuration in Figure 2. 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.

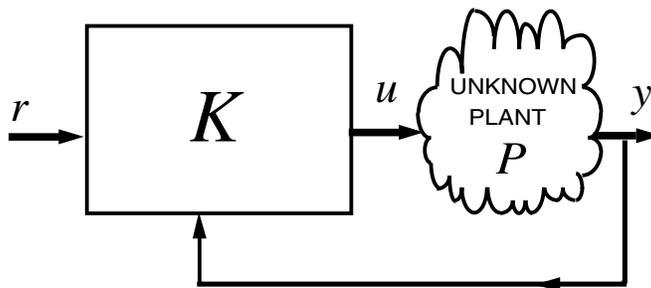


Figure 2: Feedback control system.

The 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 [49] *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 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

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

And, if $J(r, y, u)$ is a given loss-function that we wish to be non-positive, then the performance specification \mathcal{T}_{spec} would be simply the set

$$\mathcal{T}_{spec} = \left\{ (r, y, u) \mid J(r, y, u) \leq 0 \right\}. \quad (1)$$

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 [49]

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

where P_τ is the familiar time-truncation operator of input-output stability theory (cf. [52, 53]), 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 \mathcal{T}_{spec} .

Unfalsified Control Theorem [49] *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*

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

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 $(r_0, y_1, u_1) \in \mathcal{T}_{spec}$. 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} , \mathbf{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 [7].

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 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. [54, 55]). 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. 3.

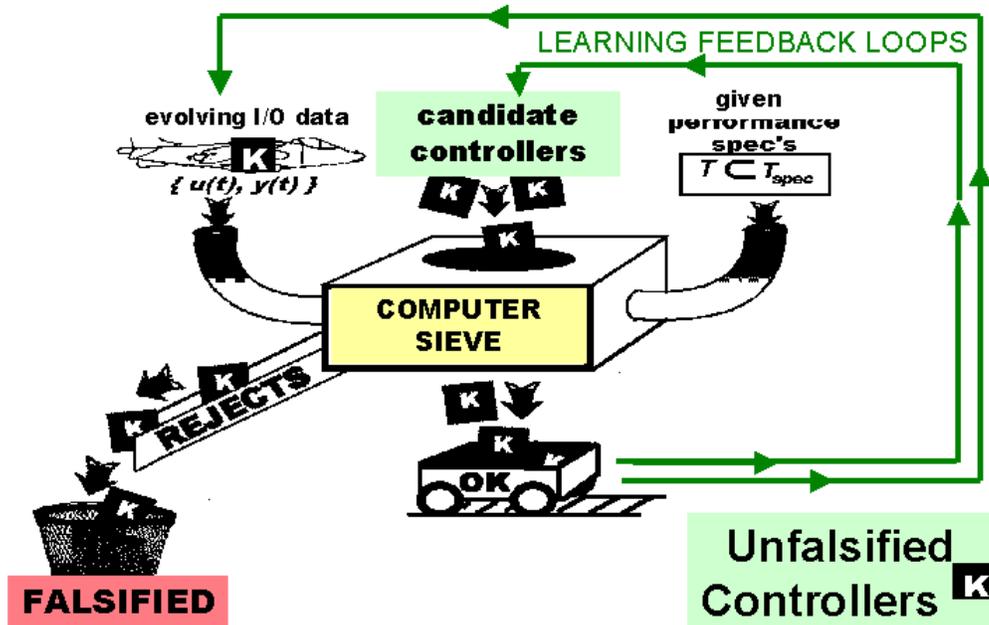


Figure 3: The unfalsification process mathematically sifts controllers to find those that are consistent with both performance goals and physical data. It plays the role of a ‘supervisor’ that chooses one of the currently unfalsified controllers to put in the aircraft’s control loop.

The step from the **Unfalsified Control Theorem** 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 [56, 57], 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 [56]. The reason is that, unlike the theory in [56], 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 (2) holds, these restrictions on \mathcal{T}_{spec} and \mathbf{K} are sufficient to permit the unfalsified set to be evaluated in real-time via the following conceptual algorithm.

Algorithm (Recursive Adaptive Control)

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 \mathcal{T}_{spec} consisting of the set of triples (r, y, u) satisfying the inequalities

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

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.

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

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

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 Algorithm 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.*

Conceptual Issues

“Heavier-than-air flying machines are impossible.”

Lord Kelvin, President, British Royal Society, 1895

A typical initial response from a knowledgeable academic researcher might be to dismiss unfalsified control theory out of hand as a sort of mathematical chicanery. The claim that unfalsified control permits control design without a plant models has tended to be regarded as too outlandish to be taken seriously. Certainly it is true that unfalsified control theory has its limitations — and that the theory needs improvement. But, such objections to unfalsified control are fallacious — based on intuition derived from inappropriate analogies. Thoughtful control theorists have been genuinely surprised and impressed by the simplicity and power of the unfalsified control theory as a mathematical basis for explaining feedback and learning, as a practical method for designing more reliable adaptive controllers, and as a data-driven technique for off-line tuning of non-adaptive feedback control gains.

Following are some typical examples of fallacious objections to unfalsified control:

1. *Unfalsified control seems unacceptably weak in its conclusion of mere unfalsification, given that familiar theories of control seem to offer stronger predictions like global stability and optimality derived deductively through the analysis of models even without the aid of validating data.*

However, this objection fails to recognize the fundamental distinction between conclusions deduced from model or assumptions and conclusions obtained via a mathematical analysis of experimental data: Beliefs about the validity of models and assumptions, and therefore any conclusions based in whole or in part on mathematical models, are not necessarily scientific truths — they might be falsified by future physical data. Unfalsified control augments introspective model-based robust control design methods by providing a quantitative methodology for closing the loop on the control design process when, at the experimental validation stage, the model-based robust control design proves to be unsatisfactory.

2. *Unfalsified control theory incorporates no sensor noise models, and therefore must perform poorly.*

In unfalsified control, control system performance criteria are framed directly in terms of observed variables. This runs counter to established intuition for some control theorists who have grown overly comfortable with traditional control problem formulations that characterize physical measurements as noise-corrupted observations of the unseen internal ‘reality’ of a ‘true’ model. Given this tradition of regarding models and noise beliefs as having more truth content than physical observations, it has been easy to succumb to the temptation to assume that the models and their noise are the ‘true’ explanation of observed physical data. The fact is that unfalsified control does accommodate noise quite well when suitably ‘soft’ performance criteria are employed. And, moreover, the performance criteria used in unfalsified control are quite flexible, and may even be associated with, and derived from, traditional stochastic noise hypotheses.

3. *Unfalsified adaptive controllers are claimed to be quick and sure-footed in discovering good control gains, even for non-minimum-phase plants. This is too good to be true, and so must be false.*

This erroneous belief apparently arises from knowledge that popular model reference adaptive control schemes are relatively sluggish and have been proved to fail for non-minimum-phase plants. But, unfalsified control is not model reference control and does not suffer its limits. It is fast and sure-footed because, unlike other adaptive schemes such as model reference control, unfalsified control theory is based on a precise analysis of the mathematical constraints induced by (1) performance criteria, (2) physical data and (3) the control law.

4. *Unfalsified adaptive controllers make use of the inverse of the controller transfer function, so they must be sensitive to plant model error and noise.*

Apparently some control theorists are confused by the thought of controller inversion, since it triggers unrelated memories concerning known difficulties with controllers that rely on inverting the plant itself. In any case, the controller inversion is not an essential part of theory but merely one of several conceivable ways to perform the computations that are necessarily associated with any logically correct test of consistency a controller hypothesis against performance goals and physical data.

5. *Anyone who claims to be able to design a controller without a mathematical model of the plant must be a charlatan, ergo unfalsified control must be the product of charlatans.* A seasoned reviewer of our paper [49] put this objection very eloquently, saying

“Modern science is model-based. If to abandon models is not to abandon mathematical science...”

In such arguments, one may perhaps glimpse elements of the conflict between the introspective belief-driven methods of ancient Platonic science and the observation-driven methods of post-Galileo experimental science. Such fallacious views are based on the knowledge that established control theories (like the theories of Plato) are heavily introspective, relying on models and assumptions – much like modern control theorists must rely on augmenting evidence with models and assumptions in order to deduce such things as stability or optimality. Platonism notwithstanding, the unfalsified control method makes it possible to devise dependable data-driven methods for discovering good control designs without models via careful mathematical analysis of the logical implications of experimental observation alone. They also fail to notice that models play a key role in the unfalsified control method, but that the models used in unfalsified control are models of candidate controller hypotheses, not plants. Indeed, a careful comparison of unfalsified control theory and system identification theory shows that they are conceptually the same, except that in unfalsified control one identifies controller models (not plant models) and performance criteria involve closed-loop control errors (not open-loop plant model errors).

The foregoing are representative of fallacious criticisms that have been levied by experienced nay sayers who were convinced that unfalsified control is a heavier-than-air theory that could not possibly fly. They were wrong. Unfalsified control theory is taking off. It is proving to be a legitimate and useful vehicle for data-driven control design, even if the initial flights have been somewhat ungainly and short.

Design Studies

The 1895 declaration of the Lord Kelvin notwithstanding, the Wright brothers flew in 1903. As for unfalsified control theory, accumulating case study evidence is likewise proving the intuition of experienced nay sayers to be wrong. Over the past three years several design studies have confirmed the theoretical expectation that unfalsified control can be useful 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. These initial design studies have helped us to better understand the potential of the unfalsified control theory, as well as limitations of the current theory.

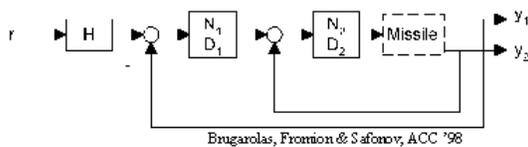
Missile Autopilot

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 [9]. Figure 4 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 model reference adaptive methods.

Robot Manipulator Arm

We used to unfalsified methodology to adaptively tune the parameters of a nonlinear ‘computed-torque’ controller for a robot manipulator arm [58, 5] — see Fig fig:robot. The arm proved to be capable of a quick and reliable control response despite large and sudden variations in load mass.

- Learns control gains
- Adapts quickly to compensate for damage & failures
- Superior performance



Unfalsified adaptive missile autopilot:

- discovers stabilizing control gains as it flies, nearly instantaneously
- maintains precise sure-footed control

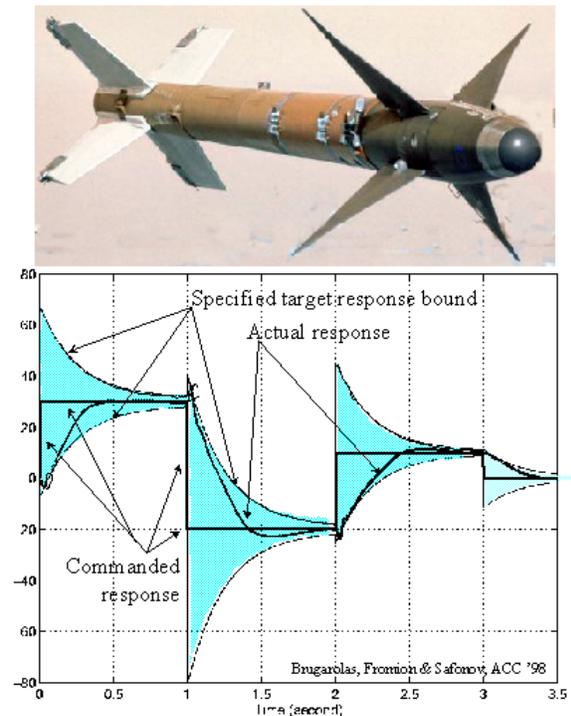


Figure 4: A data-driven unfalsified missile controller would have abilities to adaptively discover solutions in real-time to compensate for sudden in-flight changes and damage.

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.

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 [41] 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 [59] and used the theory with good result for data-driven discovery of good control gains for a laboratory control problem involving two spring-connected masses. Also, Collins and Fan [42] 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 [43, 60] based on the unfalsified control methodology.

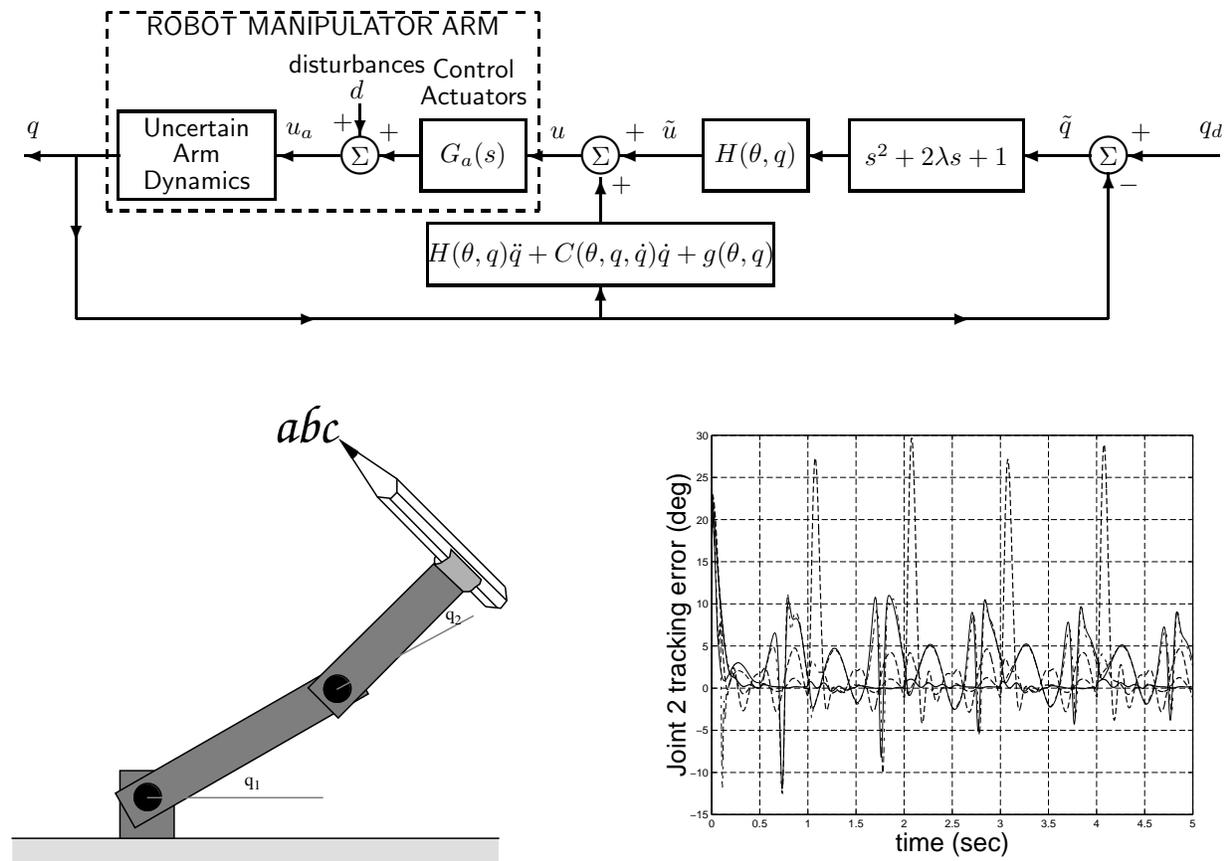


Figure 5: Unfalsified control produced superior results for a nonlinear two-link robot manipulator subject to uncertain dynamics, noisy disturbances and abrupt changes in load mass. The two sluggishly smooth traces large amplitude signals in the plot are with a conventional adaptive controller used to 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.

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 [10] — see Fig 6. The controller was capable of sifting through a bank of candidate controllers in real-time, stabilizing an open-loop unstable plant without knowledge of the plant model despite sensor noise and without noticeable transients.

3 Discussion

In our unfalsified control theory, 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

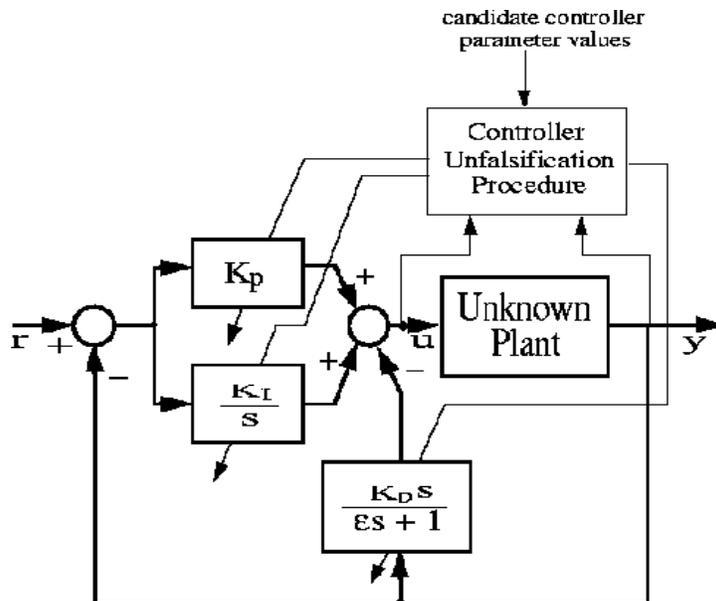


Figure 6: In one study, we designed an adaptive ‘universal controller’ having a PID structure and based on the unfalsified control theory. Simulations using MATLAB Simulink showed that the adaptive unfalsification loops were so fast that the controller was able to stabilize an unstable plant without prior plant knowledge and without appreciable transients.

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 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.

4 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* information flows from models, noisy data, and more. The contribution of *unfalsified control* theory has been to close the loop on the adaptive and robust control design processes by developing data-driven methods to complement traditional model-based methods for the design of robust control systems. The theory strengthens the foundations of control theory, paving the way for carefully reasoned extensions of feedback control theory into the realm of intelligent and learning systems. Unfalsified control theory lays the foundation for a more complete and rigorous understanding of feedback that focusing squarely on the quantitative design implications of physical data, thereby enabling the design of robust control systems with the ability to autonomously enhance both their robustness and their performance by exploiting real-time data as it unfolds. Such designs will be better able to compensate for uncertain and time-varying effects, battle damage, equipment failures and other changing circumstances.

5 Acknowledgment

I thank my students Tom Tsao, Tom Brozenec, Fabricio Cabral, Paul Brugarolas, Rengrong Wang, and Ayanendu Paul for their invaluable assistance in developing the theory and examples described in this paper.

- [1] C.-H. Huang, P. A. Ioannou, J. Maroulas, and M. G. Safonov. The design of strictly positive real systems using constant output feedback. *IEEE Trans. Autom. Control*, AC-44(3):569–573, March 1999.
- [2] M. G. Safonov. Robust control. In J. G. Webster, editor, *Encyclopedia of Electrical and Electronics Engineering*, volume 18, pages 592–602. Wiley, NY, 1999.
- [3] M. Mesbahi, M. G. Safonov, and G. P. Papavassilopoulos. Bilinearity and complementarity in robust control. In L. El Ghaoui and S. Niculescu, editors, *Advances in Linear Matrix Inequality Methods in Control*, pages 269–292. SIAM, Philadelphia, PA, 2000.
- [4] M. G. Safonov. Robust control, stability margin. In P. M. Pardalos and C. A. Floudas, editors, *Encyclopedia of Optimization*, volume 5, pages 44–49. Kluwer, Boston, MA, June 2001.
- [5] 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, 2001.
- [6] T. F. Brozenec and M. G. Safonov. Controller identification. In *Proc. American Control Conf.*, pages 2093–2097, Albuquerque, NM, June 4–6, 1997.
- [7] M. G. Safonov and F. B. Cabral. Fitting controllers to data. *Systems and Control Letters*, 43(4):299–308, July 23, 2001.
- [8] 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.

- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] R. L. Kosut and M. G. Safonov. Editorial, adaptive control, with confidence. *Int. J. Adaptive Control and Signal Processing*, 15(5):429, August 2001.
- [14] 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.
- [15] 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, New York.
- [16] P. Brugarolas and M. G. Safonov. A data driven approach to learning dynamical systems. In *Proc. IEEE Conf. on Decision and Control*, Las Vegas, NV, December 10–13, 2002, to appear. IEEE Press, New York.
- [17] V. Kulkarni, S. K. Bohacek, and M. G. Safonov. Robustness of interconnected systems with controller saturation and bounded delays. In *Proc. American Control Conf.*, pages 3206–3210, San Diego, CA, June 2–4, 1999.
- [18] J. H. Ly, R. Y. Chiang, K. C. Goh, and M. G. Safonov. LMI multiplier K_m/μ -analysis of the Cassini spacecraft. *Int. J. Robust and Nonlinear Control*, 8(2):155–168, February 1998. Special Issue on Robust Control Applications.
- [19] V. Kulkarni, S. Bohacek, and M. G. Safonov. Stability issues in hop-by-hop rate-based congestion control. In *Proc. Allerton Conference on Communication, Control and Computing*, Allerton House, Monticello, IL, September 23–25, 1998. Coordinated Science Laboratory, University of Illinois, Urbana-Champaign, IL.
- [20] R. Y. Chiang and M. G. Safonov. *Robust-Control Toolbox*. MathWorks, South Natick, MA, 1988.
- [21] R. Y. Chiang and M. G. Safonov. *Robust Control Toolbox*. MathWorks, Natick, MA, 1988 (Ver. 2.0, 1992).
- [22] G. J. Balas, J. C. Doyle, K. Glover, A. Packard, and R. Smith. *μ -Analysis and Synthesis Toolbox (μ -Tools)*. MathWorks, Natick, MA, 1991.
- [23] P. Gahinet, A. Nemirovskii, A. Laub, and M. Chilali. *LMI Control Toolbox*. MathWorks, Natick, MA, 1995.

- [24] G. Dullerud and f. Paganini. *A Course Robust Control Theory: A Convex Approach*. Springer-Verlag, New York, 1999.
- [25] 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.
- [26] 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.
- [27] R. L. Kosut. Adaptive control via parameter set estimation. *Int. Journal of Adaptive Control and Signal Processing*, 2:371–399, 1988.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] P. M. Makila. Robust control-oriented identification. In *Proceedings of IFAC symposium on system identification*, 1997.
- [36] P. M. Makila. On autoregressive models, the parsimony principle, and their use in control-oriented system identification. *International Journal of Control*, 68, 1997.
- [37] L. Giarre, B. Z. Kacewicz, and M. Milanese. Model quality evaluation in set-membership identification. *Automatica*, 33, 1997.
- [38] L. Giarre, M. Milanese, and M. Taragna. H_∞ identification and model quality evaluation. *IEEE Transaction on Automatic Control*, 42, 1997.

- [39] S. R. Venkatesh and M. A. Dahleh. Multivariable system identification: A control perspective. In *Proceedings of the IFAC Symposium on Identification*, 1996.
- [40] 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.
- [41] 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.
- [42] 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.
- [43] 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.
- [44] 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.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] 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.
- [49] M. G. Safonov and T. C. Tsao. The unfalsified control concept and learning. *IEEE Trans. Autom. Control*, AC-42(6):843–847, June 1997.
- [50] 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.
- [51] T. Mitchell. *Machine Learning*. McGraw-Hill, Inc., 1997.
- [52] I.W. Sandberg. On the L_2 -boundedness of solutions of nonlinear functional equations. *Bell System Technical Journal*, 43(4):1581–1599, July 1964.
- [53] 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.

- [54] 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.
- [55] 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.
- [56] 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.
- [57] 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.
- [58] 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.
- [59] 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.
- [60] 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.