Sixteenth Yale Workshop on Adaptive and Learning Systems, New Haven, CT, June 5-7, 2013, pp 16-21.

Adaptive Control: Fooled by False Assumptions?

Michael G. Safonov
University of Southern California
Los Angeles, CA 90089-2563
Email: msafonov@usc.edu

***Abstract*—A mismatch between prior modeling assumptions and reality can fool an adaptive controller design into persistently preferring a destabilizing controller even which the instability is patently obvious to the eyes of the most casual observer. To eliminate all possibility of being thusly fooled, the assumption-free unfalsified control concept was introduced two decades ago and has since been developed to the point where it now provides a unifying overarching theoretical framework for understanding the relationships, benefits and weakness of various adaptive control methods. We review how the theory allows one to parsimoniously sift the individual elements of accumulating evidence and experimental data to determine precisely which of elements falsify a given performance level and very briefly discuss recent research on cost-detectable fading-memory cost-functions for time-varying plants.**

I. Introduction

"Nothing is what it seems. False assumptions take hold through constant repetition, blowing away the more complicated evidence in front of our eyes.

Steve Richards [28], 2012

Assumptions and models have always played an important role in decision and control theory. Modern optimal control and Bayesian estimation theory revolve around plant and process models. The models are characterized by assumed structures and, in the case of uncertain models, the uncertainty itself is typically characterized by further assumptions about its behavior and, perhaps, its statistical properties. In statistical learning theory and as well as in some approaches to adaptive control theory, the powerful mathematical tools of Bayesian probability theory and optimal control theory have been brought to bear to produce *a priori* guarantees of convergence and stability — *provided that the true plant conforms to the modeling assumptions*.

Of course, there is a problem with assumptions: When assumptions fail to hold, so do conclusions that rest on those assumptions. If the only theoretical tools one has for evaluating the implications of raw data rest on assumptions, then one may be too easily tempted to ignore the possibility that assumptions may be wrong. When this happens we can sometimes be fooled into accepting a design or even an entire design methodology that is fatally flawed from the outset — like the adaptive controller which is reputed to have caused the fatal 1967 crash [10] of the experimental X-15 aircraft (Fig. 1) or the more generic failures of certain 'standard assumptions' of adaptive control theory illustrated by the Rohr's counterexamples [29].

Fig. 1. An X-15 aircraft with an adaptive controller crashed in 1967.

Despite the ever present danger of being fooled by wrong assumptions, there is a consensus that assumptions are unavoidable. As the late MIT professor Fred Schweppe (cf. [34]) was fond of telling his estimation theory students, "You have to assume something to know something." This premise, which it turns out is actually false, seems nevertheless to be tacitly accepted by many researchers in learning theory, estimation and control. Assumptions can certainly be productive to form hypotheses about what classes of learning algorithms might work, but it is quite another matter to design a learning system that fails to respond intelligently when future experimental data is inconsistent with prior assumptions. Yet, this seems to be an unfortunate characteristic of most (but not all) textbook algorithms for learning, estimation and adaptive control.

II. Assumption-free Learning

"It ain't so much the things we don't know that get us into trouble. It's the things we know that just ain't so."

Josh Billings [20], c. 1870

Engineers continue to be confronted with situations where assumptions required by control and estimation theory either do not hold or are not possible to verify without invoking further assumptions. Yet, it remains the general consensus that some assumptions will always be necessary. Everyone 'knows' that modern control theory is built on models and modeling assumptions, so it seems impossible to imagine that this could change. After all, the oft-quoted maxim 'all observation is theory laden' seems to deny even the existence of uninterpreted raw data.[1]

In fact, it turns out to be fairly simple to formulate and solve estimation and control problems without prior assumptions via a fairly straightforward application of the scientific 'curve-fitting' methods to raw uninterpreted data, whenever such data is available — e.g., Gauss-Newton curve-fitting methods (cf. Bertsekas [6, Sect. 1.5]) which select a parameter vector x of function $h(x, y)$ to approximately fit input-output data pairs $(y_i, z_i) i = 1, \ldots, n$ by computing the optimal value x^* minimizing a cost function of the form

$$f(x) \triangleq \sum_{i=1}^{n} \|(z_i - h(x, y_i)\|^2. \quad (1)$$

The computation of the optimal parameter vector x^* using such curve-fitting methods requires no assumption about whether the 'true' system is or is not exactly modeled by some value of the parameter vector x, and (unlike Bayesian estimation) it requires no noise models and no assumed prior probabilities. The optimal value of the cost $f(x^*)$ is the accuracy of the model fit in the sense that the model $z = g(x^*, y)$ is *unfalsified* by the data at accuracy level $f(x^*)$.

Here we may find fertile ground for a shift towards a *data-driven* paradigm of estimation and control. Such a shift might occur as engineers are increasingly confronted with problems in which either priors do not hold or there is a need to adapt in real-time, as when systems change in unforeseen ways due to wear, damage, or evolving operating conditions. To do this reliably, we must have methods that respond correctly when new data falsifies our prior assumptions; that is, we need *scientific* methods.

A. Popper's Logic of Science

> *"There can be no ultimate statements in science: there can be no statements in science which cannot be tested, and therefore none which cannot in principle be refuted, by falsifying some of the conclusions which can be deduced from them."*
>
> *Karl Popper [26], 1934*

What is science, and how is it distinguished from pseudo-science and quackery? Philosopher Karl Popper [26] studied these questions in depth, examining in detail the processes and practices employed by the best scientists like Newton and Einstein. In the end, Popper concluded that an empirical theory (i.e., a theory concerning the material world, as contrasted with a mathematical theorem) may be properly called a *scientific theory* only if it passes three tests, which we may loosely summarize as follows:

1) *Falsifiability.* Data from some conceivable future experimental outcome must be in principle capable of falsifying (i.e., invalidating) the theory.
2) *Simplicity.* The theory must be parsimonious, which is to say it must not assert extra conditions beyond those that are necessary to explain observable data.
3) *Validation.* The theory must be thoroughly tested by surviving aggressive experimental efforts designed to seek falsifying empirical evidence (i.e., data).

The falsifiability test was Popper's hallmark contribution to the definition of science. It addresses the empirical fact that scientific knowledge does not always advance. It explains that scientific theories are always tentative, remaining forever subject to possible falsification by new information. The simplicity test is basically the classical Occam's razor principle, but it also implies Newton's [23] famous *hypotheses non fingo* principle ('make no feigned hypotheses') prohibiting experimentally unverifiable prior assumptions about the 'true' internal workings and/or motivations of the process that generates the observed data.[2] The validation test is what distinguishes a scientific theory from a mere *scientific hypothesis*.

B. The Role of Probability Theory

> *"It's turtles all the way down!"*
>
> *Stephen Hawking [14], 1988*

As noted by Popper [26], Finetti [8], Kalman [18], [19] and Taleb [40], [41], Bayesian probability theory is on inherently shaky ground from a scientific perspective. There is an 'infinite regress' problem associated with the application of Bayes' rule,

$$p(X|D) = \frac{P(D|X)P(X)}{P(D)} = \frac{P(D,X)}{P(D)}.$$

To estimate the posterior probability $p(X|D)$ of X given data D using Bayes' rule, we evidently must first have an estimate of the prior probability $P(X, D)$. So, any attempt to estimate a probability rests on having an estimate of a prior probability, thus leading to an infinite regress — like the infinite stack of turtles of a primitive cosmology (cf. Hawking [14]). Unless we terminate the regress at some level either by simply assuming a prior probability or by actively intervening to 'shuffle the deck' as is done in RCT experiments,[3] the probability $P(X|D)$ cannot be computed. Because of this infinite regress problem, probability statements are generally experimentally unfalsifiable, and therefore by Popper's criterion Bayesian estimation theory may not be a scientific theory. Unlike RCT methods, Bayesian estimation

[1] The maxim 'all observation is theory laden' can be traced to philosopher Hanson [13], who was referring to situations similar to human telescopic observation where the distant thing said to be 'observed' is in fact a human re-interpretation of the unprocessed raw data. Hanson did not deny the existence or potential accessibility of uninterpreted raw data.

[2] An interesting albeit, to some, possibly disturbing consequence of this is that any theory for system identification or adaptive control that begins with prior assumptions about the structure or form of either the 'true plant' (LTI, parameters, order, etc.) or the 'true noise' (Gaussian, unknown-but-bounded, etc.) may not be a scientific theory by Popper's definition.

[3] In the *randomized controlled trial (RCT)* experimental method of Fisher [11], the experimenter intervenes by randomly assigning members of the sample set to test and control groups prior to performing an experiment in order to create a uniform prior probability distribution of potentially confounding factors, even when these factors are unknown. RCT should not be confused with the assumption-based Monte Carlo "randomization" methods like those described in Tempo *et al.* [42], in which one *assumes* a uniform prior probability on an *assumed* sample space. According to Pearl [25, pp. 410–418], RCT is the only known statistical method for reliably testing causal hypotheses.

and system identification problem formulations usually assume both a 'true system' satisfying prior assumptions about its form and internal structure, and noises with assumed prior probabilities.

The problem of infinite regress is not isolated to probabilistic problem formulations. Infinite regress issues can also arise with unknown-but-bounded estimation methods favored by robust control theorists that begin with prior assumption that the 'true plant' lies in a predetermined set with a given *a priori* known internal structure and, perhaps, noises or other uncertain elements that satisfy *a priori* uncertainty bounds (e.g., [24], [12]). Any prior assumptions about the internal structure of the 'true plant' violate the *hypothesis non fingo* principle of Isaac Newton [23]. To correct this, we need data-driven black-box methods; i.e., methods in which only observed data is regarded as known, without any prior assumptions about the properties or internal workings of the so-called 'true system' or measurement noises.

III. The unfalsification paradigm

Heavier-than-air flying machines are impossible.

Lord Kelvin [43], President, British Royal Society, 1895

Paradigm shifts are controversial, even risky. To completely eliminate prior assumptions in estimation and control is to develop a assumption-free methodology, and perhaps thereby even to risk being tarred with the label of simpleton or even charlatan. Let us tentatively and cautiously take this risk and see where it leads.

What we are seeking is a data-driven theory, i.e., a theory in which no prior assumptions corrupt our interpretation of results. One implication of this is that the performance criterion by which we judge a candidate controller or estimator must be expressed directly in terms of the raw noise-corrupted past input-output data collected from the outputs of sensors attached to the real plant. And, since with a finite amount of past plant input-output data we cannot know for sure without further assumptions how plant would respond to other as yet untried input signals, or even that it would respond exactly the same way in the future, the best we can hope to conclude from raw data analysis is that the particular past observed behavior is, or is not consistent with our performance goals. In other words, the best data-driven theory of control and estimation can hope to conclude that a given candidate controller or estimator is as yet unfalsified by the past data.

So to have proper scientific data-driven theory, we must limit our aims to judging performance goals expressed as functions of raw data. And, after we examine any candidate estimator or controller, we must be careful to limit the claims to unfalsification of performance, and resign ourselves to the fact that traditional *a priori* guarantees of asymptotic stability are unattainable without gratuitously introducing additional assumptions to constrain future data in ways that nature may not. That is, the strongest guarantee that a scientific data-driven theory of estimation or control can logically offer is that given controller's or estimator's performance is as yet unfalsified by past data. But, this may be enough when the data is information rich (e.g., persistently exciting), so that bad controllers and estimators will be quickly and efficiently falsified and discarded.

It turns out that all this is not hard, and has in fact been done without much fanfare for a very long time.

A. Unfalsified estimation theory

"Probability does not exist."

Bruno de Finetti [8], 1974

For the case of data-driven estimator design, a simple example of a data driven 'curve-fitting' solution is provided by the classical *weighted least-squares* algorithms that select model parameters by minimizing a cost function similar to (1)(e.g., [34], [6]). Without unneeded the presumption of a 'true' plant or noise model, weighted least-squares estimators solve the problem of computing optimal model parameter estimates such that the input-output behavior of the models minimizes an arbitrary quadratic function of the difference between the raw experimental input-output data and the model output [34]. For some problems, the solution even takes the exact form of a Kalman filter. But of course weighted-least-squared is not Kalman filtering theory, as it does not require access to the 'true plant' or to probabilistic noise models, nor does it claim to be able to *a priori* predict future error-covariance. More significantly unlike the Kalman filtering theory, the data-driven least square curve fitting solutions parsimoniously make no use of prior assumptions and, consequently, make none of the purely assumption-driven claims of the Kalman filtering theory to have precisely pre-computable Gaussian statistics. Such unfalsifiable claims of access to certain knowledge about future statistics goes beyond limits of what is logically knowable from data alone, violating Popper's simplicity requirement for a scientific theory as well as *hypotheses non fingo* principle of Newton [23]. In fact, this seems to be one of the reasons why Kalman [18], [19] now questions both Bayesian probability and the probabilistic interpretation of his Kalman filter equations.

Willems [48] and Smith [37] provide further examples of assumption-free curve-fitting-type formulations of estimation problems. Explicitly citing Popper [27] as motivation, Willems defined an unfalsified model as one whose graph contains all currently available plant data, and he introduced the term most powerful unfalsified model (MPUM) to describe an unfalsified model that offers the best (in a particular sense) fit to the data. Smith proposed optimally curve-fitting past plant input-output data to minimize a robust-control-oriented 'LFT' error criterion that admits a curve-fitting cost function with terms penalizing multiplicative errors among others.

Essentially, these data-driven methods involve nothing more than curve-fitting plant and controller models to raw data so as to optimize 'unfalsified' performance levels by minimizing a curve-fit-error performance criterion. Because the process of curve-fitting raw data to models requires no prior assumptions about plant or noise, such methods are wholly incapable of being confounded by prior assumptions and are therefore inherently robust.

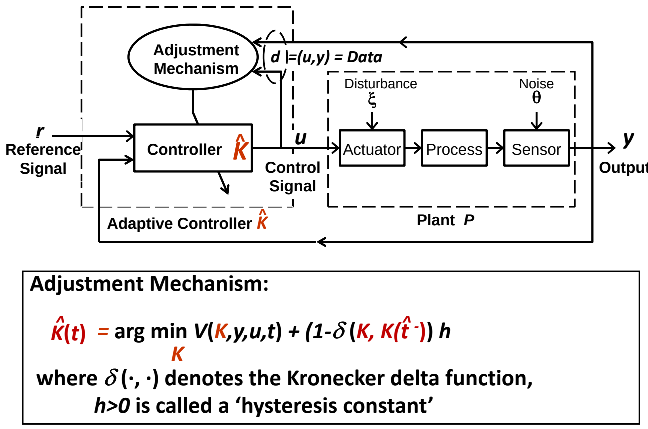

Fig. 2. In adaptive control it is desirable that the cost-function $V(K, y, u, t)$ be *cost-detectable*; otherwise, the adjustment mechanism can be destabilizing [47], [38].

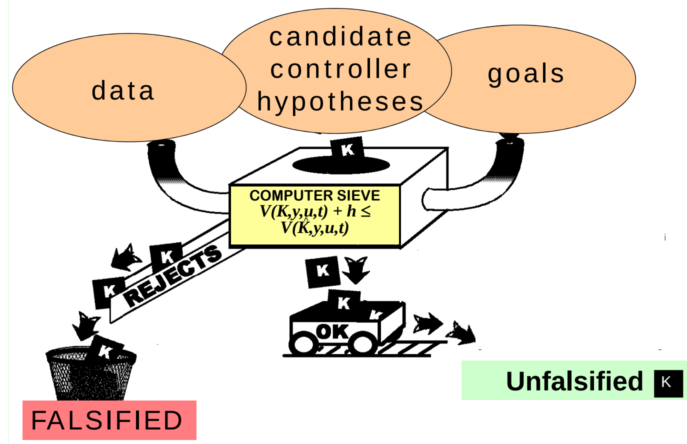

Fig. 3. An unfalsification process used in unfalsified control design is essentially a data-driven sieve. The sieve requires three types of input: (1) goals, (2) candidate controllers and (3) data. Controllers are sifted to find those that are consistent with both performance goals and physical data. 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.

B. Unfalsified Control Theory

> *"Just have lots of ideas and throw away the bad ones."* — *Linus Pauling [1], 1935*

It is almost as easy to formulate the problem of directly finding closed-loop feedback controller parameters as a data-driven curve-fitting problem formulation as it is to do so for open-loop estimator design problems [32], [30]. The problem is essentially a controller identification problem, which may be reliably solved by choosing controllers that minimize any suitable cost-function expressing ones performance goals as a causal function of raw plant input-output data, uninterpreted by plant/noise models or other prior beliefs. Further, a most basic requirement is that this cost function should be *cost-detectable* [47], [38], which is to say the cost must remain finite if, only if, the data does not falsify a controllers ability to stabilize the plant.

As first proved by Wang-Paul-Stefanovich-Safonov [46], [47, Thm. 2], by simply substituting a cost-detectable cost function together into a now classic Morse-Mayne-Goodwin [22] adaptive hysteresis switching algorithm, one can prevent model-mismatch instability.

Theorem 1 (Safe Adaptive Control —[46], [47], [38]): Consider the adaptive feedback control system in Figure 2. Suppose that the adaptive control problem is feasible in the sense that there exists at least one candidate controller $K \in K$ that stabilizes the plant. If the cost-function $V(K, y, u, t)$ is monotone in t, the controllers $K \in K$ are minimum-phase and the(V, K) is cost-detectable, then the adaptive system converges to a stabilizing $K \in K$ after at most finitely many controller switches. □

The safe adaptive control theorem says that merely substituting a cost-detectable cost function into the classic Morse-Mayne-Goodwin hysteresis switching algorithm is sufficient to prevent model-mismatch instabilities, thereby ensuring that no mismatch between prior modeling assumptions and reality can fool the adaptive control logic into persistently preferring a destabilizing candidate controller over a stabilizing one. The algorithm is essentially a data-driven sieve that rejects/falsifies controllers $K \in K$ whose cost is not at least $h > 0$ lower than the cost of the currently active controller $\hat{K}(t)$, as depicted in Figure 3. Using a cost-detectable cost function $V(K, y, u, t)$ in the unfalsification sieve solves the problem posed by the Rohrs counterexamples [29] and perhaps even could help to prevent problems like the fatal 1967 failure of the adaptively controlled X-15 aircraft [10]. For the unity negative feedback adaptive control system in Figure 2 with minimum-phase candidate controllers $K(s) \in K$, an example of a cost-detectable cost function is [32], [33], [47], [38]

$$V(K, y, u, t) = \sup_{\tau \le t} \frac{\|W_1 e_K\|_\tau^2 + \|W_2 u\|_\tau^2}{\|r_K\|_\tau^2 + \epsilon} \tag{2}$$

where $\epsilon > 0$ is a small constant, $\|x\|_\tau$ denotes the L_2-norm of the signal x truncated at time τ, and e_K and r_K are $K(s)$-dependent 'fictitious' signals, defined as the signals e and r that would have generated the data (y, u) had the controller K been in the loop when the data was collected. For minimum-phase $K(s)$, they are computed via the formula $e_K \triangleq K^{-1}u$, $r_K \triangleq e_K + y$.[4]

The above cost function $V(K, y, u, t)$ is an unfalsified lower-bound on the standard weighted-sensitivity cost function of robust control [31], [36], viz.

$$\left\| \begin{bmatrix} W_1(I + PK)^{-1} \\ W_2K(I + PK)^{-1} \end{bmatrix} \right\|_\infty .$$

The theory of unfalsified control, including its application to the design of inherently robust assumption-free adaptive control systems is now developing rapidly, including several recent

[4] If $K^{-1}(s)$ is unstable, then we may substitute $e_K = D_K u$ and $r_K = N_K y + e_K$ where $K(s) = D^{-1}(s)N(s)$ is any stable left-coprime matrix fraction description of $K(s)$ [21].

application studies, e.g., [39], [44], [35]. Significant recent progress includes alternative cost functions with windowed and fading memory suitable for adaptive control of plants that may vary so greatly with time that the controller must be switched to maintain stability and performance [3], [4], [17], [16], [15], [5]. Interestingly, there has also been recent success in developing cost-detectable cost functions that reflect plant-model-based goals [2].

IV. Discussion

> *"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."*
>
> *A. C. Doyle [9], 1891*

> *"It may be so, there is no arguing against facts and experiments."* *Issaac Newton [7], 1726*

Audiences of mathematical system theorists at past talks that I have given on unfalsified control usually express discomfort, and occasionally even outrage, at the thought that plant models and noise assumptions do not seem to play a direct role in the unfalsified control theory. A common refrain is that 'science has always relied on models and assumptions and to throw these away is to throw away valuable information'. A slight problem with that view of science is that prior assumptions, mathematical models and even basic physical laws like $f = ma$ and $e = mc^2$ are regarded by scientists as mere theories, always subject to possible refutation by future experiments data. Indeed, it is precisely when some prior assumptions or some models are falsified by new data that the need for learning and adaptation arises.

The fact is that models can, and usually do, have a pivotal role in the application of unfalsified control theory just as they do in most of engineering. In unfalsified control applications, plant models are normally used in the synthesis of the candidate controllers that are subsequently fed to the assumption-free unfalsification algorithm. In particular, the set of candidate controllers K is typically created by applying traditional model-based and assumption-based control synthesis methods. Each candidate controller in the set is designed to meet performance goals for a member of a set of possible 'true' plant and 'true' noise models. These model-based candidate controller designs are then fed to the unfalsified control algorithm, which separately performs its assumption-free and plant-model-free checks for falsification of each controller using only *raw* plant input-output data.[5] Each unfalsified controller that emerges from this process has passed two tests: (1) a purely model-driven test for robust performance with one or more of the hypothesized candidate plant and noise models that was used for the controller's initial design, and (2) a purely data-driven test for whether the controller's performance is falsified by past input-output data. The sole role of data-driven second test of unfalsification is to allow adaptive feedback to robustly correct any performance problems that may arise when some of the assumed plant or noise models are later falsified by new data. A key point here is that information in prior models is not discarded at all, but is instead fully exploited in designing a set of candidate controllers to pass the model-based first test.

[5]Using raw data in formulating the cost-functions used to evaluate performance is essential. It avoids the need to invoke the peculiarly circular logic of introducing unrealistic Santa Claus assumptions to 'prove' stability of adaptive systems that in fact are not—e.g., as exemplified by the Rohr's counterexamples [29].

V. Conclusion

> *"The task of science is to stake out the limits of the knowable and to center the consciousness within them."* *Rudolf Virchow [45], 1849*

In the present paper, I have argued that there may be paradigm shift now in the making as assumption-laden theories of learning and adaptive control theory face potentially falsifying empirical tests. Of particular concern are assumptions of prior knowledge of uncertainty structure and bounds, leading the questions like, "How should controllers and estimators change when these assumptions are falsified by new data?" The next revolution in estimation and control could be a discomforting shift towards assumption-free methods for the design of learning systems and adaptive control algorithms based on Popper's unfalsification paradigm. The unfalsified control paradigm allows one to clearly and sharply separate the implications prior assumptions from those of data, which is a logical necessity if one is to avoid being fooled by prior assumptions.

Acknowledgment

I wish to express my sincere appreciation to the many colleagues and students whose contributions have helped shape the development of the unfalsified control concept that is the essential element of the plant-assumption-free formulation of the adaptive control problem presented in this paper. Former students and postdocs who have made direct contributions to the theory or its application include Michael W. Chang, Shin-Young Cheong, Seung-Gyun Cheong, Margareta Stefanovic, Ayanedu Paul, Crisy Wang, Fabricio Cabral, Tom Brozenec, Myungsoo Jun, Vincent Fromion, Paul Brugarolas, Claudia Manuelli, Hardy Siahaan and especially Tom Tsao who helped lay the initial foundations of the theory. I also thank Robert Kosut, Brian Anderson, Lennart Ljung, Michel Gevers, Mike Athans, Charlie Rohrs, Masami Saeki, Petar Kokotovic, Petros Ioannou, Karl Astrom, Edoardo Mosca, Bob Narendra, Jan Willems, Jeroen van Helvoort, Martin Steinbuch, Huiyu Jin whose work and words helped to shape my perspective on the adaptive control problem and to focus my thoughts on the need for a plant-assumption-free theory for safe adaptive control.

References

[1] J. Angier (exective producer), "Linus Pauling: Crusading scientist," *Transcript of broadcast of NOVA*, vol. 417, 1977. [Online]. Available: http://osulibrary.oregonstate.edu/specialcollections/coll/pauling/bond/audio/1977v.66-ideas.html

[2] S. Baldi, G. Battistelli, E. Mosca, and P. Tesi, "Multi-model unfalsified adaptive switching supervisory control," *Automatica*, vol. 46, no. 2, pp. 249–259, 2010.

[3] G. Battistelli, J. Hespanha, E. Mosca, and P. Tesi, “Unfalsified adaptive switching supervisory control of time varying systems,” in *Proc. IEEE Conf. on Decision and Control*, Shanghai, China, December 16-18, 2009, pp. 805 –810.

[4] ——, “Model-free adaptive switching control of uncertain time-varying plants,” in *IFAC World Congress*, vol. 18, Part 1, Milan, Italy, August 28 – September 2, 2011, pp. 1273–1278.

[5] ——, “Model-free adaptive switching control of time-varying plants,” *IEEE Trans. Autom. Control*, vol. 58, no. 5, pp. 1208–1220, 2013.

[6] D. Bertsekas, *Nonlinear Programming (2nd ed.)*. Belmont, MA: Athena Scientific, 1999.

[7] D. Brewster, *Memoirs of the life, writings, and discoveries of Sir Isaac Newton*. Edinburgh, Scotland: T. Constable and Co., 1855.

[8] B. de Finetti, *Theory of probability: A critical introductory treatment*. NY: Wiley, 1974.

[9] A. C. Doyle, *A Scandal in Bohemia*. London: Penguin, 2001, first published in the Strand Magazine, July 1891.

[10] Z. Dydek, A. Annaswamy, and E. Lavretsky, “Adaptive control and the NASA X-15-3 flight revisited,” *IEEE Control Systems Magazine*, vol. 30, no. 3, pp. 32–48, 2010.

[11] R. A. Fisher, *The Design of Experiments*. Edinburgh: Oliver and Boyd, 1934.

[12] M. Gevers, X. Bombois, B. Codrons, G. Scorletti, and B. D. Anderson, “Model validation for control and controller validation in a prediction error identification framework—Part I: Theory,” *Automatica*, vol. 39, no. 3, pp. 403 – 415, 2003.

[13] N. R. Hanson, *Patterns of Discovery*. Cambridge: Cambridge University Press, 1958.

[14] S. Hawking, *A Brief History of Time*. New York: Bantam, 1988.

[15] H. Jin and M. G. Safonov, “Unfalsified adaptive control: Controller switching algorithms for nonmonotone cost functions,” *Int. J. Adaptive Control and Signal Processing*, vol. 26, no. 8, p. 692704, August 2012. [Online]. Available: http://dx.doi.org/10.1002/acs.2265

[16] H. Jin, M. Chang, and M. G. Safonov, “A fading memory data-driven algorithm for controller switching,” in *Proc. IEEE Conf. on Decision and Control and European Control Conference*, Orlando, FL, December 12-15, 2011, pp. 6097–6103.

[17] ——, “Unfalsifying dominant pole locations using a fading memory cost function,” in *Proc. IFAC World Congress*, vol. 18, Part 1, Milan, Italy, August 28 – September 2, 2011, pp. 1291–1295.

[18] R. E. Kalman, “Randomness reexamined,” *Modeling, Identification and Control*, vol. 15, no. 3, pp. 141–151, 1994.

[19] ——, “Randomness and probability,” *Mathematica Japonica*, vol. 41, no. 1, pp. 41–58, 1995, in memory of E. R. Caianiello.

[20] R. Keyes, *The Quote Verifier*. New York: St. Martin’s Press, 2006.

[21] C. Manuelli, S. G. Cheong, E. Mosca, and M. G. Safonov, “Stability of unfalsified adaptive control with non SCLI controllers and related performance under different prior knowledge,” in *Proc. European Control Conf.*, Kos, Greece, July 2–5, 2007, pp. 702–708. [Online]. Available: http://routh.usc.edu/pub/safonov/safo06g.pdf

[22] A. S. Morse, D. Q. Mayne, and G. C. Goodwin, “Applications of hysteresis switching in parameter adaptive control,” *IEEE Trans. Autom. Control*, vol. 37, no. 9, pp. 1343–1354, September 1992.

[23] I. Newton, *Philosophiae Naturalis Principia Mathematica (2nd Ed.)*, Cambridge, England, 1713.

[24] B. M. Ninness and G. C. Goodwin, “Rapprochement between bounded-error and stochastic estimation theory,” *International Journal of Adaptive Control and Signal Processing*, vol. 9, no. 1, pp. 107–132, 1995.

[25] J. Pearl, *Causality, Second Edition*. Cambridge: Cambridge University Press, 2009.

[26] K. R. Popper, *The Logic of Scientific Discovery*. Routledge, London, 1959, english translation of K.R. Popper *Logik der Forschung*. 1934.

[27] ——, *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge, London, 1963.

[28] S. Richards, “Don’t be fooled by the power of false assumptions,” *The Independent*, January 17, 2012, accessed 3/11/2012. [Online]. Available: http://www.independent.co.uk/voices/commentators/steve-richards/steve-richards-dont-be-fooled-by-the-power-of-false-assumptions-6290557.html

[29] 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*, vol. AC-30, no. 9, pp. 881–889, September 1985.

[30] M. G. Safonov, “Focusing on the knowable: Controller invalidation and learning,” in *Control Using Logic-Based Switching*, A. S. Morse, Ed. Berlin: Springer-Verlag, 1996, pp. 224–233. [Online]. Available: http://dx.doi.org/10.1007/BFb0036098

[31] M. G. Safonov and R. Y. Chiang, “CACSD using the state-space L^∞ theory–A design example,” *IEEE Trans. Autom. Control*, vol. AC-33, pp. 477–479, 1988.

[32] 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. Berlin: Springer-Verlag, 1995, pp. 196–214.

[33] M. G. Safonov and T. C. Tsao, “The unfalsified control concept and learning,” *IEEE Trans. Autom. Control*, vol. AC-42, no. 6, pp. 843–847, June 1997.

[34] F. C. Schweppe, *Uncertain Dynamic Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1973.

[35] H. B. Siahaan, H. Jin, and M. G. Safonov, “An adaptive PID switching controller for pressure regulation in drilling,” in *Proc. IFAC Workshop on Automatic Control in Offshore Oil and Gas Production (ACOOG 2012)*, Trondheim, Norway, May 31 - June 1, 2012, pp. 90–94.

[36] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control*. New York: Wiley, 1996.

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

[38] M. Stefanovic and M. G. Safonov, *Safe Adaptive Control: Data-driven Stability Analysis and Robust Synthesis*. Berlin: Springer-Verlag, 2011, lecture Notes in Control and Information Sciences, Vol. 405. [Online]. Available: http://dx.doi.org/10.1007/978-1-84996-453-1

[39] M. Steinbuch, J. van Helvoort, W. Aangenent, B. de Jager, and R. van de Molengraft, “Data-based control of motion systems,” in *Proc. IEEE Conference on Control Applications*, Toronto, Canada, August 28-32, 2005, pp. 529–534.

[40] N. Taleb, *Fooled by Randomness, The Hidden Role of Chance in Life and in the Markets*. New York: Random House, 2005.

[41] ——, “The a priori problem of observed probabilities,” 2007, accessed June 2, 2012. [Online]. Available: http://www.fooledbyrandomness.com/central.pdf

[42] R. Tempo, F. Dabbene, and G. Calafiore, *Randomized Algorithms for Analysis and Control of Uncertain Systems*. Berlin: Springer-Verlag, 2005.

[43] W. Thompson (Lord Kelvin), “Heavier-than-air flying machines are impossible,” in *The Experts Speak: The Definitive Compendium of Authoritative Misinformation*, C. Cerf and V. Navasky, Eds. New York, NY: Pantheon, 1984, p. 238.

[44] J. van Helvoort, A. de Jager, and M. Steinbuch, “Data-driven controller unfalsification with analytic update applied to a motion system,” *IEEE Trans. on Control Systems Technology*, vol. 16, no. 6, pp. 1207–1217, 2008.

[45] R. Virchow, “Der mensch (on man),” in *Einheitsbestrebungen in der wissenschaftlichen Medicin*, Berlin, 1849, also, *Disease, Life and Man — Selected Essays of Rudolf Virchow* (trans. L. J. Rather), Stanford University Press, Stanford, CA, pp. 67–70, 1958.

[46] R. Wang, A. Paul, M. Stefanovic, and M. G. Safonov, “Cost-detectability and stability of adaptive control systems,” in *Proc. IEEE Conf. on Decision and Control*, Seville, Spain, December 12-15, 2005, pp. 3584–3589. [Online]. Available: http://dx.doi.org/10.1109/CDC.2005.1582718

[47] ——, “Cost-detectability and stability of adaptive control systems,” *Int. J. Robust and Nonlinear Control*, vol. 17, no. 5-6, pp. 549–561, 25 March - April 2007, special Issue: Frequency-domain and Matrix Inequalities in Systems and Control Theory. Dedicated to the 80th Birthday of V. A. Yakubovich.

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