In A. S. Morse, editor, *Control Using Logic-Based Switching*, pages 224–233. Springer-Verlag, Berlin, 1996.

Focusing on the Knowable

Controller Invalidation and Learning

Michael G. Safonov*†

The task of science is to stake out the limits of the knowable and to center the consciousness within them.

Rudolf Virchow — Berlin, 1849

Abstract

Modern and classical control theories fail to clearly distinguish conclusions based on data from those that are a consequence of assumptions. The "unfalsified control concept" is offered as a possible solution to this dilemma. The concept provides a set-theoretic characterization of control-relevant information which is based solely on experimental data and is untainted by plant modeling assumptions or related prejudices. It forms the basis for an objective analysis of adaptation and learning, leading directly to practical model-free adaptive control algorithms. The basic idea is to invalidate, or falsify, controllers off-line using previously collected experimental plant data. The theory may be viewed as complementing traditional plant-model-based theories of control by providing needed additional guidance when real-time data fails to corroborate modeling assumptions.

Motivation and Issues

Describing the scientific method, Nobel laureate mathematician-philosopher Bertrand Russell [1] once said,

> *"The essential matter is an intimate association of hypothesis and observation. The Greeks were fruitful in hypotheses, but deficient in observation. Aristotle, for example, thought that women have fewer teeth than men, which he could not have thought if he had had a proper respect for observation."*

Like Aristotle [2] and many theoreticians since, we control theorists are too easily tempted to build grand, but fanciful theories based on assumptions which may fail to hold. We may rightly claim to prove stability, robustness and even

*Phone 1-213-740-4455. FAX 1-213-740-4449. Email safonov@bode.usc.edu.

†Research supported in part by AFOSR grants F49620-92-J-0014 and F49620-95-I-0095.

optimality, but only subject to assumptions that our models are valid to within conjectured tolerances or that our uncertainties conform to assumed probability distributions. Unfortunately, proofs of stability and optimality based on mathematical models all too often have little predictive value when applied to physical systems. The problem here is that mathematical theories of control typically give insufficient attention to the implications of possible future observations which may be at odds with assumptions. The Achilles heel of modern control theory has been this habit of "proof by assumption."

Even robust control theory fails to adequately address the question of what to do when subsequently obtained experimental data does not corroborate assumptions about uncertainty bounds, though model-validation theory provides a conservative first step in the right direction (see [3] and references therein). An essential question is:

> *How can we separate the control-relevant information in experimental data from that in our assumptions?*

Answering this question is key to understanding feedback control and learning. It turns out that it is not necessarily difficult to formulate this question in precise mathematical terms, though curiously traditional control theory problem formulations have apparently failed to do so.

We need to focus on the data. Unlike mathematics which is constrained only by one's imagination in creating self-consistent theorems and assumptions, the theories of science are more rigidly constrained to conform to experimental data — or at least they should be. While control theory has brought the full power of mathematics to bear on the question of how to proceed from assumptions about the physical world to conclusions about stability robustness and even optimality, the same attention has not been brought to bear on the critical question of how experimental data and performance goals constrain the choice of controller. The answer to this question is precisely what is needed to put feedback control theory on a sound scientific footing. To find the answer, it will be necessary to develop problem formulations that focus attention squarely on what of control relevance can be deduced from experimental data alone.

> *To start, put aside the dogma that control analysis must begin with a model.*

Of course, there can be no denying that models are useful in control design. They help us to formulate hypotheses about suitable controller structures by distilling apparently simple patterns from complex data. But we must recognize that models may also be misleading, since they can cause us to draw conclusions which we could not draw from the data alone. Models are essentially an aggregate of our a priori knowledge and prejudices. These prejudices are only partially removed when we augment a plant model with bounded uncertainties (Δ's), stochastic noises, or unknown parameters. Uncertain models are still models. And they still reflect prejudices.

A proper understanding of feedback and learning requires us to carefully examine the control-relevant information in the data so that we may better recognize situations in which our modeling assumptions are mistaken. To do so, we must temporarily cleanse our minds of models and other prejudicial assumptions about the plant to be controlled, for only then can we hope to develop a scientifically sound analysis of the proper role of experimental data. That is, we need to find mathematical problem formulations which allow us to clearly distinguish experimentally derived information from that which is a consequence of our modeling assumptions.

But, we must reconcile ourselves to the fact faced by all scientists that there are severe limitations on what can be logically concluded from data alone. To quote Sir Karl Popper [4, p. 114], noted philosopher and author of numerous works on the logical basis of scientific discovery,

> *The Scientist ... can never know for certain whether his theory is true, although he may sometimes establish ... a theory is false.*

No matter how many past experiments have corroborated a scientific theory, there always remains the possibility that a new experiment may disprove the theory. Consequently, a scientific theory remains forever a tentative conjecture as scientists seek, but never fully succeed, to validate the theory through diligent efforts to devise experiments which may falsify the theory. The role of mathematics in science is to test the logical self-consistency of various components of a scientific theory, as well as to assist in performing the computations needed to test whether a theory is consistent with experimentally observed facts. And, we must be prepared to abandon theories and models which are inconsistent with data. As Isaac Newton said when confronted with experimental data which seemed to refute his theory [5],

> *"It may be so, there is no arguing against facts and experiments."*

Like a scientific theory, a control law's ability to meet a given performance specification may be regarded as a conjecture to be subjected to validation by experimental data. And, as with a scientific theory, one must accept that in the absence of prejudicial assumptions the most that one can hope to learn about a control law from experimental data alone is that the conjecture is false, i.e., that the control law cannot meet a given performance specification. For example, without prejudicial assumptions about the plant it is not logically possible from data alone to deduce that a given controller stabilizes the plant. But, using only past data it turns out that it is sometimes possible to show that a candidate feedback controller is *not* capable of meeting certain quantitative specifications that are closely related to closed-loop stability, even without actually trying the controller on the real plant. It also turns out to be true that even a very small amount of open-loop experimental plant data can often be used to draw sweeping conclusions about the unsuitability of large classes of feedback controllers.

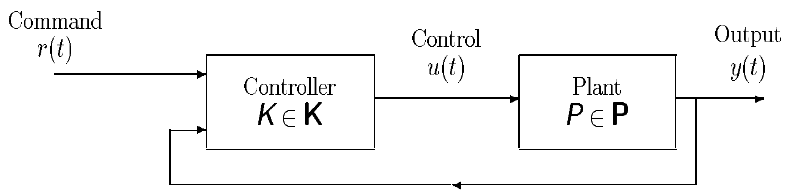

Figure 1: Feedback control system.

A Simple Example

Consider the system in Figure 1. Suppose that nothing is known about the plant P, save that an experiment has been performed during which some otherwise unknown nonlinear dynamical feedback controller was in place. Further, suppose that during this experiment the control input to the plant was observed to be $u_{data}(t) = 1$ and the experimentally observed output was $y_{data}(t) = -1$ over the same time-interval $t \in [0,1]$. Although this particular closed-loop data is consistent with the simple plant model $y = -u$, another simple model matching the data is $y \equiv -1$. The actual plant that generated this data might even be a complicated time-varying, nonlinear, dynamical system and subject to disturbances and noise such as

$$y + \dot{y} = -2u^3(t) + 1 - \dot{u}(t) sin(t) cos(u^2(t)),$$

though a good scientist would never opt for such a complex model without some corresponding complexity in the data u_{data}, y_{data}. We know only the experimental data. With this and nothing more, we may still conclude that if another feedback law $u = k(r - y)$ were to have been in the loop at the time of the experiment and if the hypothetical command input

$$r(t) = \tilde{r}(t,k) \stackrel{\Delta}{=} -1 + \frac{1}{k} \tag{1}$$

had been applied, then this would have resulted in the previously observed experimental data being exactly reproduced and, moreover, the resultant closed-loop error $\tilde{e}(t) \stackrel{\Delta}{=} \tilde{r}(t,k) - y(t)$ would have been $\tilde{e}(t) = 1/k$. Furthermore,

$$\sup_{\substack{r \neq 0 \\ \tau}} \frac{||e||_{L_2[0,\tau)}}{||r||_{L_2[0,\tau)}} \geq |\tilde{e}|/|\tilde{r}| = 1/|k-1|. \tag{2}$$

We thus see that even without a plant model or any assumptions whatsoever about the plant, experimental data alone is sufficient to allow one to logically conclude that if $|k - 1| < 0.1$ then the closed-loop system fails to meet a quantitative stability specification requiring closed-loop L_{2e}-gain less than 10. With additional experimental data, one might be able to identify other values of k that are inconsistent with this L_{2e}-gain specification, if there are any. And, if one collected further data while the plant was being subjected to a rich class of "persistently exciting" inputs, then one might hope eventually to arrive at a situation in which practically all inconsistent values of k had been eliminated or "falsified." Through this process of elimination, the set of "good" values of the feedback gain k would be identified.

The foregoing example suggests that by setting control system performance goals that are only slightly less ambitious than stability or robustness, it still may be possible to develop scientifically sound algorithms for designing good feedback control laws directly from past data — without recourse to a priori plant models or other prejudicial assumptions about the plant. To do this scientifically, we must focus squarely on what is knowable from the data alone. The key is to ask not that stability and robustness be proved, but to ask only that the experimental data not be inconsistent with stability and robustness — i.e., ask only that the system be *not-provably-unrobust*.

Unfalsified Control

In 1620 Sir Francis Bacon [6] identified the problem of scientific discovery as follows:

> *"Now what the sciences stand in need of is a form of induction which shall analyse experience and take it to pieces, and by a due process of exclusion and rejection lead to an inevitable conclusion."*

While nearly four centuries of scientific thought have not fulfilled Bacon's wish for the certainty of inevitable conclusions about the laws that govern the physical world, the scientific method continues to revolve around a process of exclusion and rejection of hypotheses that are falsified by observations [4]. In collaboration with my students T. Tsao and T. Brozenec, I have recently examined the application of the falsification principle to the problem of identifying the set of not-provably-unrobust controllers in the absence of plant models or related assumptions [7, 8, 9, 10]. Our work may be regarded as an attempt to simplify, streamline and reduce the conservativeness associated with works (e.g., [11, 12, 3]) which strive to achieve this same goal indirectly via a marriage of model validation and robust control theory.

We have considered systems of the general form in Figure 1, adopting the input-output perspective introduced by George Zames and Irwin Sandberg in their celebrated papers on stability [13, 14, 15, 16, 17]. And, like Zames [16], Rosenbrock [18], Safonov [19, 20], Willems [21] and others, we cut directly to the core of the problem by throwing away the plant state and thinking of the plant as simply a collection of input-output pairs $(u, y) \in \mathcal{U} \times \mathcal{Y}$; i.e., a "graph" consisting of the points (u, y) in the "$\mathcal{U} \times \mathcal{Y}$-plane". Then, treating each candidate controller's ability to meet specifications as a hypothesis to be validated against past data, we may perform a computation to test the candidate control laws for inconsistencies between the three components that define our data-only control design problem, viz.,

I. experimental input-output data from the plant u_{data}, y_{data},

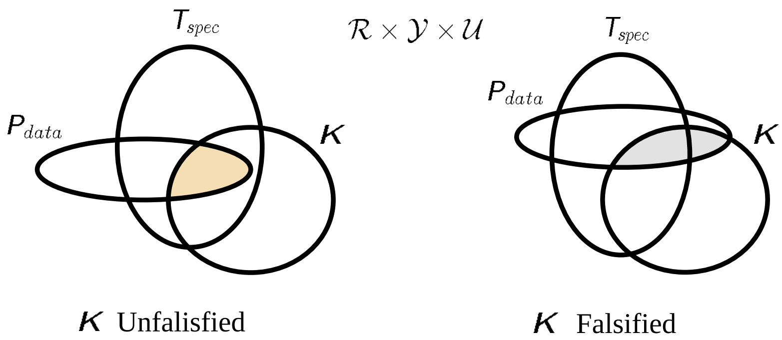

Figure 2: A controller $K(s)$ is *unfalsified* by the data if $\boldsymbol{P}_{data} \cap \boldsymbol{K} \subset \boldsymbol{T}_{spec}$, otherwise $\boldsymbol{K}$ is falsified.

II. constraints on (r, y, u) imposed by the control law $u(s) = K(s) \begin{bmatrix} y(s) \\ r(s) \end{bmatrix}$, and

III. a performance specification $(r, y, u) \in \boldsymbol{T}_{spec}$.

We call this idea the *unfalsified control concept*, since the "validated" control laws that emerge from this process are precisely those whose ability to meet the performance specification is not falsified by experimental data.

These three components of unfalsified control problem each define a corresponding subset of the space $\mathcal{R} \times \mathcal{Y} \times \mathcal{U}$, viz.,

$$\boldsymbol{P}_{data} \left\{ \ (r, y, u) \ \middle| \ u = u_{data}, y = y_{data} \ \right\} \subset \mathcal{R} \times \mathcal{Y} \times \mathcal{U},$$

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

$$\boldsymbol{T}_{spec} \subset \mathcal{R} \times \mathcal{Y} \times \mathcal{U}.$$

The controller $K(s)$ is *invalidated* (i.e., its ability to meet the performance specification $\boldsymbol{T}_{spec}$ for all r is falsified) by the open-loop data u_{data}, y_{data}) if and only if the following condition fails to hold (cf. Fig. 2):

$$\boldsymbol{P}_{data} \cap \boldsymbol{K} \subset \boldsymbol{T}_{spec}. \tag{3}$$

The unfalsification condition (3) is simply a restatement of what it means mathematically for experimental data and a performance specification to be consistent with a particular controller. Yet, simple though it may be, it does not seem to have been incorporated in either classical or modern control problem formulations. It has some interesting and, I think, important implications:

- Condition (3) is nonconservative; i.e., it gives "if and only if" conditions on $K(s)$. It uses all the information in the data — and no more.

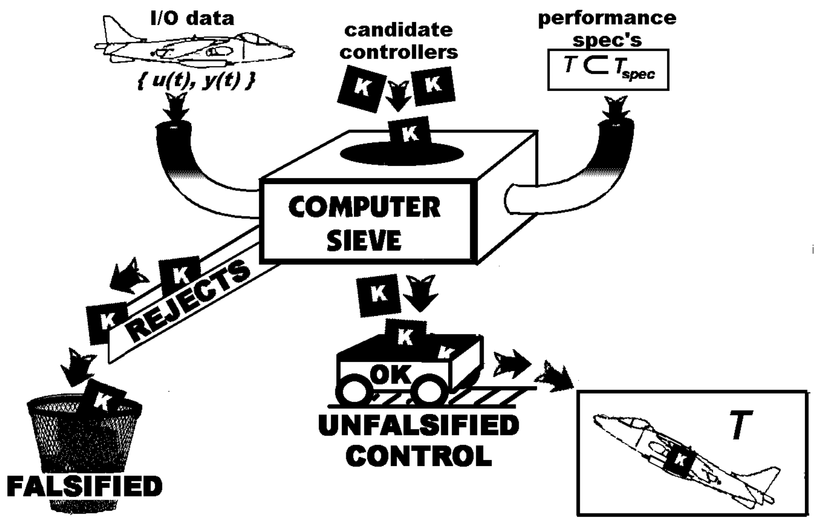

Figure 3: Computer "sieve" checks controllers for consistency with experimental data and performance specifications.

- Condition (3) provides a "plant-model free" test for controller invalidation. No plant model is needed to test its conditions. It depends only on the data, the controller and the specification.

- In (3), data (u, y) which invalidates a controller $K(s)$ need not have been generated with the controller K in the feedback loop; it may be open loop data or data generated by some other control law (which need not even be in $\mathbf{K}$).

- The unfalsification condition (3) accommodates noise and plant uncertainty. In the case that the system to be controlled is subject to random disturbances, sensor noise or uncertain nonlinear or dynamical "parasitics," then the performance criterion T_{spec} must be be chosen to reflect the fact that achievable performance will be correspondingly limited. But the terms P_{data} and K on the left-hand side of the falsification condition (3) are unaffected.

- The relation (3) provides the mathematical basis for constructing a computer "sieve" for screening controllers. This sieve non-conservatively identifies controllers which are consistent with both performance goals and past observations of plant input-output data. See Fig. 3.

Learning control is possible using the controller sieve provided by (3). If one has a class, say $\mathbf{K}$, of candidate controllers, then one may evaluate condition (3) for each candidate controller $K \in \mathbf{K}$ and thereby "identify" the subset $\mathbf{K}_{OK} \subset \mathbf{K}$ of unfalsified controllers, i.e., those controllers whose ability to meet the performance specification T_{spec} is not invalidated by the data u_{data}, y_{data}. With minor modifications described in [7, 8, 9, 10], this invalidation process may be carried out in real-time by recursively evaluating certain functions of the plant data u_{data}, y_{data}. The result is a practical "unfalsified adaptive control" theory which, in apparent contrast to previous theories, is focused squarely on what is knowable from the experimental data alone, independent of any assumptions on the plant.

Of course, if the theory is to lead to a non-empty set of unfalsified controllers, then it is essential that it begin with a sufficiently rich class of ideas for candidate controllers $\mathbf{K}$. As Linus Pauling put it [22],

"Just have lots of ideas and throw away the bad ones."

Conclusions

Because they fail to clearly distinguish information derived from new data from that which follows from previous knowledge, both classical and modern control problem formulations provide inadequate guidance when experimental results fail to corroborate a priori expectations based on plant models or old data. To address this concern, a non-traditional plant-model-free formulation of the control problem has been proposed in which the rather modest goal of "unfalsification" replaces traditional goals such as stability or optimality. Admittedly, the proposed goal deprives us of the ability to make grand claims about the plant's future behavior (e.g., asymptotic stability), but we console ourselves with the fact that such limits on our powers of clairvoyance would seem to be inherent in any unprejudiced scientific analysis of feedback control problems based solely on data. On the other hand, unfalsified control theory does give a precise mathematical characterization of how the available experimental data can be used to eliminate from consideration broad classes of controllers and, through this process of elimination, identify controllers which insofar as is possible to ascertain from data alone are at least not provably incapable of achieving traditional performance goals. The results complement traditional results based on a priori assumptions about the plant by providing a sharp characterization of how the set of suitable control laws shrinks when experimental data falsifies a priori expectations. And, when implemented in real time, the results lead directly to scientifically sound adaptive control algorithms.

The salient feature of our theory of learning control is that, for a given specification T_{spec} and given class of admissible controllers $\mathbf{K}$, plant models and other prejudicial assumptions about the plant or its uncertainty do not enter into the equations. Instead prejudices and models are relegated to the more appropriate role of influencing initial choices for the performance specification T_{spec} and for the class of candidate controllers $\mathbf{K}$. As Bertrand Russell said, *"The essential matter is an intimate association of hypothesis and observation."*

References

[1] Bertrand Russell. Nature and origin of scientific method. In British Broadcasting Corp., editor, *The Western Tradition*, pages 23–28, London, England, 1949. Vox Mundi.

[2] Aristotle. *Historia Animalium*, volume II. c. 300BC. In *The Works of Aristotle*, Vol. IV. Oxford, England: Clarendon, 1910. Trans. by D'Arcy W. Thompson under editorship of J. A. Smith and W. D. Ross.

[3] 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.

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

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

[6] Francis Bacon. *Novum Organum*. 1620. English translation by R. L. Ellis and J. Spedding (eds), in *The Philosophical Works of Francis Bacon*, Routledge, 1905.

[7] 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.

[8] M. G. Safonov and T. C. Tsao. The unfalsified control concept and learning. *IEEE Trans. on Automatic Control*, submitted 10/94, revised 7/95.

[9] T. F. Brozenec and M. G. Safonov. Controller invalidation. In *Proceedings IFAC World Congress*, San Francisco, CA, USA, July 19–23, 1996. Submitted June 1995.

[10] T. F. Brozenec and M. G. Safonov. Controller invalidation. *Automatica*, submitted 8/95.

[11] 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.

[12] R. Smith and J. Doyle. Model invalidation — a connection between robust control and identification. *IEEE Trans. on Automatic Control*, AC-37(7):942–952, July 1992.

[13] G. Zames. Functional analysis applied to nonlinear feedback systems. *IEEE Trans. on Circuit Theory*, CT-10:392–404, September 1963.

[14] I.W. Sandberg. On the L_2-boundedness of solutions of nonlinear functional equations. *Bell System Technical Journal*, 43(4):1581–1599, July 1964.

[15] I.W. Sandberg. A frequency-domain condition for the stability of feedback systems containing a single time-varying nonlinear element. *Bell System Technical Journal*, 43(4):1601–1608, July 1964.

[16] 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*, AC-15(2):228–238, April 1966.

[17] G. Zames. On the input–output stability of time-varying nonlinear feedback systems — Part II: Conditions involving circles in the frequency plane and sector nonlinearities. *IEEE Trans. on Automatic Control*, AC-15(3):465–467, July 1966.

[18] H. H. Rosenbrock. *State Space and Multivariable Theory.* NY:Wiley, 1970.

[19] M. G. Safonov. *Robustness and Stability Aspects of Stochastic Multivariable Feedback System Design.* PhD thesis, Dept. Elec. Eng., MIT, 1977. Supervised by Michael Athans.

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

[21] J. C. Willems. Paradigms and puzzles in the theory of dynamical systems. *IEEE Trans. on Automatic Control*, AC-36(3):259–294, March 1991.

[22] J. Angier (exective producer). Linus Pauling: Crusading scientist. *Transcript of broadcast of NOVA*, 417, 1977.

[23] R. Virchow. *Der Mensch (On Man).* Berlin, 1849. In *Disease, Life and Man Selected Essays of Rudolf Virchow* (trans. L. J. Rather), Stanford University Press, Stanford, CA, pp. 67–70, 1958.