Journal of Guidance, Control, and Dynamics
Vol. 45, No. 12, December 2022

Facing Future Challenges in Feedback Control of Aerospace Systems Through Scientific Experimentation

Dennis S. Bernstein*
University of Michigan, Ann Arbor, Michigan 48109

https://doi.org/10.2514/1.G006785

I. Introduction

As the second quarter of the 21st century approaches, two schools of thought about control technology for aerospace vehicles are emerging. One school holds that the essential concepts and techniques of control are well established, and the principal challenge now is to take advantage of increasingly powerful digital technology to implement computationally demanding algorithms, such as reinforcement learning and model predictive control. Reaching this point has benefited from technological advances punctuated by costly and sometimes tragic failures, for example, X-15 Flight 3-65-97, Northwest Orient 710, United Airlines 232, USAir 427, Air France 447, US Airways 1549, Ariane 5 V88, Lion Air 610, NASA Helios, and NASA X-43A, each of which involved control technology, at least to some extent. Nevertheless, the lessons learned from these events have vastly improved the safety, reliability, and performance of subsequent flight vehicles.

The opposing school stresses that the challenges posed by future flight systems will require operation in regimes and envelopes outside of current experience. These regimes range from subsonic flight of flexible vehicles prone to aeroelastic instability [1] to hypersonic vehicles subject to severe thermal loads [2,3]. The extreme conditions of hypersonic flight, in particular, demand high-authority controllers that can learn rapidly under uncertain conditions. Another example is data-poor applications, such as turbulent combustion, for which extreme events are difficult to predict [4]. Meeting these needs will require technological innovation and new perspectives.

Motivated by the latter viewpoint, this article has three objectives. The first is to examine the widely held view that performance is the reason for feedback control, where performance may include the desire to modify the open-loop dynamics, for example, through stabilization. This article refines this view by demonstrating that, while the goal of *control* is performance, the reason for *feedback* control is to achieve the best possible performance in the face of uncertainty. In other words, *uncertainty is the underlying driver of feedback*. To support this perspective, we show that stabilization is inherently an uncertainty problem.

The second objective is to explore consequences of the fact that feedback control—by virtue of its ability to mitigate uncertainty—can mask physical variations. Consequently, the cause of the success or failure of a feedback controller may be hidden from the designer and operator.

Finally, to understand and expose these hidden effects, the last objective is to propose a call to action to establish a culture of scientifically meaningful feedback control experimentation. Unlike hardware demonstrations for education, and unlike flight testing for validation and verification, these experiments would be aimed at probing the ability of a feedback controller to provide robustness to physical variations that were not included in the control-oriented model, even to the point of probing the failure boundary. This objective is motivated by the fact that the failure of a feedback control system may have destructive consequences, and thus it is important to expose the potential weaknesses of a closed-loop system. Accordingly, this article is aimed at the challenges faced by feedback control technology for future aerospace systems.

The audience for this article includes all readers with an interest or stake in feedback control technology for aerospace applications. For readers with limited background in feedback control, this article provides a tutorial on the unique and perhaps surprising ability of feedback control to mitigate uncertainty using closed-loop action. For students and experts in control, some of this material, such as the view that stabilization is inherently an uncertainty problem and the observation that feedback exacerbates model incompleteness, is nonstandard. Finally, for researchers and instructors, the call to develop a culture of scientifically meaningful control experimentation can stimulate new and fruitful directions in the development and teaching of control concepts and methods.

The contents of this article are as follows. Sections II–IV present conceptual material that motivates later discussion. Sections V–VII demonstrate how uncertainty is the underlying driver of feedback. These sections are more technical than Secs. II–IV but are accessible to students of control. Section VIII discusses how the ability of feedback control to mask physical variations leads to the fundamental problem of feedback control, namely, the inability to definitively determine the reasons for the success or failure of a feedback controller. Section IX examines the rationale for feedback control experiments, which may be either physical or computational, and Section X discusses the implications of these experiments for research, education, and practice. Finally, Section XI summarizes the main points of the article and provides some concluding remarks.

II. Uncertainty Mitigation

Uncertainty constrains performance by impeding repeatability and predictability. Uncertainty is inherent in the fact that all models are approximate, all data are noisy, and the ability to observe the physical world is limited. Even if Newton's laws, the laws of thermodynamics, the Navier–Stokes equations, and Maxwell's equations are accepted as absolute ground truth, scientific and engineering experience show that physical phenomena can be exceedingly complex. In fact, it is all too easy to design and build circuits, structures, devices, and vehicles whose dynamics are too complex for us to reliably predict their behavior. In addition, simulation of chaotic, low-dimensional systems shows that unpredictability is not confined to high-dimensional dynamics. It is also important to stress that uncertainty need not be due to unknown science or chaotic dynamics. Rather, parameters may be uncertain due to limitations in testing and measurement, as well as unknown and unpredictable changes that occur during operation, such as icing and damage. The need to confront complex reality and limitations on knowledge of the physical world thus require uncertainty quantification [5] as well as uncertainty mitigation. This article focuses on the latter.

There are many ways to mitigate uncertainty. One approach is to improve the accuracy of models of the physical world; this is the goal of science. Another approach is redundancy, where, for example, multiple processors enhance reliability. Refined modeling and hardware redundancy are strategies for *open-loop uncertainty mitigation*.

Unfortunately, open-loop uncertainty mitigation may lead to over-engineering, which may be costly and sacrifice performance.

Received 3 March 2022; revision received 14 August 2022; accepted for publication 30 August 2022; published online Open Access 17 October 2022.

*Professor, Aerospace Engineering, 1320 Beal Avenue; dsbaero@umich.edu.

For example, the wings of a lightweight aircraft may undergo severe aeroelastic oscillations due to disturbances and aggressive maneuvers. An obvious fix is to make the wings more rigid, leading to a weight penalty. An alternative approach is to use feedback control to suppress structural oscillations. But why use *feedback* control rather than open-loop control?

Feedback control enables *closed-loop uncertainty mitigation*, which provides performance benefits in the presence of uncertainty that are not possible with open-loop uncertainty mitigation. As stated in [6, p. 3], "feedback allows a system to be insensitive both to external disturbances and to variations in its individual elements." Likewise, it is noted in [7] that "uncertainty and feedback have become inseparable in viable control applications." Furthermore, feedback control can be used to linearize the undesirable effect of nonlinear components [6, p. 3]. Perhaps most importantly, feedback control can be used for stabilization, a critical requirement in aerospace vehicles. As succinctly stated in [8],

> Feedback is used in control systems to change the dynamics of the system (usually to make the response stable and sufficiently fast), and to reduce the sensitivity of the system to signal uncertainty (disturbances) and model uncertainty.

By providing robustness to unknown physical variations, *feedback control aims to achieve the best possible performance in the face of uncertainty.* Although performance is the objective of *control*, uncertainty is the *raison d'être* for *feedback* control.

As obvious as it may be that feedback mitigates uncertainty; enhances robustness to variations in initial conditions, parameters, disturbances, dynamics, and nonlinearities; and makes stabilization possible, this ability has not always received the attention it deserves. For example, the history [9] of early control developments from ancient times to the 18th century provides no explicit recognition of uncertainty as the driver of feedback control, although uncertainty is implicit in the statement [9, p. 8] that "The purpose of a feedback control system is to carry out commands; the system maintains the controlled variable equal to the command signal in spite of external disturbances."

More recent developments in control technology are documented in [10], which extends the historical survey of feedback control given in [9] to 1930. Although [10] showcases the remarkable ingenuity of feedback control mechanisms along with early developments in control theory, uncertainty mitigation is not explicitly recognized as the driver of feedback. In the history of the next phase of control given in [11], however, robustness to unknown system variations is mentioned within the context of Black's negative feedback amplifier [11, p. 74]:

> Under this condition, the gain is dependent solely on the feedback circuit network, usually a passive network; changes in the behaviour of the active components in the forward path, for example because of the effect of variations in power supply voltage or ageing of the tube, have almost no effect on the output.

In other words, feedback control has the profound ability to precisely set the gain of an otherwise imprecise amplifier. Other than the above statements, there is no explicit recognition in [9–11] of uncertainty as the driver of feedback.

In contrast to the above sources, an eloquent description of the remarkable ability of feedback control to mitigate uncertainty is given in [12, pp. 57, 58] in relation to metal processing:

> The semi-miraculous aspect of feedback becomes apparent here: this correction takes place whatever the cause of the deviation. Whether the failure to reach the standard dimensions be due to a change in the malleability or the ductility of the metal, or whether it should stem from a gradual wearing of the rolls, or from a sudden reduction of the temperature, makes no odds! The regulator is unconcerned with causes; it will detect the deviation and correct it. The error may even arise from a factor whose influence has never been properly determined hitherto, or even from a factor whose very existence is unsuspected. Whatever the cause, the disturbance will be overcome none the less effectively.

Adding to this, Kelly [13, p. 121] notes that, "How the system finds agreement at any one moment is beyond human knowing, and more importantly, not worth knowing."

The ability of a feedback controller to mask physical variations, however, leads to the following fundamental problem: How can we determine the extent to which a feedback controller was or will be effective in controlling a system in view of the fact that the feedback controller suppresses unknown and perhaps unknowable variations? This question motivates consideration of the role of feedback control experiments in control research, control education, and control practice. The value of control experiments is described in [14, pp. 16, 17] as follows:

> Experimentation is becoming increasingly important. Perhaps the two primary goals of experimentation for control are model development or validation, and simulation of the closed-loop performance. It is important to emphasize that models suitable for open-loop simulation may not be appropriate for control studies. Therefore, the type of experiment needed by the control theorist can differ from the "standard" experiment. An obstacle here is the magnitude of some of the systems whose control is proposed, or the costliness of mistakes, as in the case of a space station or nuclear power system. Actual applications often allow no room for experimentation of the kind that could lead to genuine mathematical insight. Opportunities for small-scale experiments designed to test ideas on simple systems can be therefore of great benefit to theoretical progress.

In the update [15] to [14], a more forceful call for control experiments is made [15, p. 24]:

> Finally, experimentation on representative systems must be an integral part of the control community's approach. The continued growth of experiments, both in education and research, should be supported, and new experiments that reflect the new environment will need to be developed. These experiments are important for the insight into application domains that they bring, as well as the development of software and algorithms for applying new theory. But they also form the training ground for systems engineers, who learn about modeling, robustness, interconnection, and data analysis through their experiences on real systems.

In Sec. X, we will discuss the potential value of physical and computational experimentation in control research, education, and practice.

III. Observations on Modeling

The principal goal of science is to understand the physical world. Of course, "understand" is not a precise word, and there are many levels of understanding. For example, the model

$$f = \frac{GMm}{r^2} \tag{1}$$

shows how gravitational force depends on the masses and distance between a pair of idealized particles, as well as the physical constant G. For many, if not most, aerospace applications, this model provides excellent predictive capability. But there is much we do *not* understand. What is the nature of the gravitational field that produces a force that acts through space? What is mass? What determines G? It is easy to scratch the surface of this and many other models and raise fundamental questions. And, when we reach the next level of understanding, what new questions might we ask? As expressed in [16, p. 195],

> The principles need not pass any philosophical test or even be fully understood—thus, Newton considered himself to have explained the motions of the planets and the tides using his theory of gravity, although he offered no explanation of the causes of gravity itself.

In other words, the decisive test for scientific acceptance is predictive ability rather than conceptual plausibility.

We adopt the point of view that complex reality exists independently of how well we understand it and how well we can model it. In addition, we accept the fact that "understanding" reality is a process of uncovering an endless sequence of nested features. As stated in [17],

> Now it would be very remarkable if any system existing in the real world could be exactly represented by any simple model. However, cunningly chosen parsimonious models often do provide remarkably useful approximations. For example, the law $PV = RT$ (sic) relating pressure P, volume V and temperature T of an "ideal" gas via a constant R is not exactly true for any real gas, but it frequently provides a useful approximation and furthermore its structure is informative since it springs from a physical view of the behavior of gas molecules. For such a model there is no need to ask the question "Is the model true?". If "truth" is to be the "whole truth" the answer must be "No". The only question of interest is "Is the model illuminating and useful?".

More succinctly, Box and Draper [18, p. 74] state, "Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful." Even more succinctly, Skogestad and Postlethwaite [8, p. 24] state that "G is never an exact model." These observations express the *principle of ultimate unmodelability*. Nevertheless, despite ultimate unmodelability, science seeks the best possible understanding of complex reality—as it should.

The principle of ultimate unmodelability should not be construed as downplaying the value of models, whose usefulness in understanding complex reality and designing and operating aerospace vehicles cannot be overstated. Rather, our goal is to highlight the unique ability of feedback control to mitigate the limitations of modeling while exploring the implications of this ability.

IV. Models and Feedback Control

A. Terminology

A *model* of a physical system is a collection of pairs of inputs and outputs, where, to reflect initial conditions and uncertainty, multiple outputs may be associated with the same input. A model is *inaccurate* if, for at least one modeled input, none of the corresponding modeled outputs is identical to the output of the physical system. Furthermore, a model is *incomplete* if the set of modeled inputs does not include all possible inputs. A model is *approximate* if it is either inaccurate or incomplete; all models are inaccurate and incomplete. A model is *exact* if it is not approximate; there are no exact models, but we will use this word from time to time to facilitate discussion.

It is common practice to view a model as approximate due to erroneous parameter values or unmodeled physics. This view suggests, however, that the model can be made exact by providing the correct parameter values or including the missing dynamics. In reality, no model can be made exact.

A *representation* of a model is a prescription (such as an algorithm or computer code) for generating outputs from inputs. For convenience, a model will be identified with its representation rather than the collection of input–output pairs.

A *simulation model* is a model constructed to include as many details as desired. In contrast, a *control-oriented model* is typically a simpler model that is used as the basis for controller synthesis. A control-oriented model typically consists of two components, namely, 1) a nominal model and 2) an uncertainty model [19]. The term *model error* refers to discrepancies between the physical system and the control-oriented model. Ultimate unmodelability implies that the model error is unknowable and—most importantly—cannot be exactly represented by data or mathematical idealizations, such as idealized geometric objects and functions. Nevertheless, it is standard practice to use data and mathematics to characterize model error by constructing an uncertainty model as part of a control-oriented model. For the same nominal model, however, different modelers may have different uncertainty models, and thus different model errors, based on different modeling information and assumptions. Uncertainty models and model errors are thus *subjective entities that depend on the modeler.*

B. Uncertainty Models

As part of a control-oriented model, an uncertainty model must have two main attributes. First, it must be as nonconservative as possible relative to the possible physical variations in the physical system, and, second, it must be tractable for controller analysis and synthesis. For example, norm bounds provide rich uncertainty set that facilitates stability and performance guarantees through the small gain theorem and its variants. This approach is the basis for much of robust control.

But where do norm bounds come from? In practice, first-principles modeling and system identification are combined to determine a radius that encompasses the range of uncertainty. But all data are finite and noisy, and thus it is conceivable that the magnitude of the model error (viewing it for now as a mathematical object) might exceed the radius of the ball. If indeed a physical feedback-control experiment reveals unexpectedly poor performance, then the radius of the ball can be increased until acceptable results are (hopefully) obtained, suggesting that the model error (still viewed as a mathematical object) lies inside a larger ball. But is this a valid conclusion?

Let us optimistically assume for now—in violation of ultimate unmodelability—that the model error can be represented by a mathematical model. How can we be sure that the space of functions that we choose to work with includes the model error? Second, if this is in doubt, then how much testing is needed to ensure that a controller developed to accommodate the uncertainty model will work reliably on the physical system? Could the system possess a slow divergence that is revealed only after hours or weeks of operation, or could it have a destabilizing nonlinearity that is manifested only in response to rare disturbances? Third, open-loop system identification is performed under certain conditions, where either we or a computer code generates the test signals. Since the nominal model and the uncertainty model are reflections of the chosen input signals, it follows that these choices determine the incompleteness of the model.

Under feedback operation, however, the inputs to a physical system are generated by—itself. Since an exact model of the physical system is not available, we cannot know in advance what those signals will be and whether or not they are encompassed by the chosen control-oriented model. This phenomenon is articulated in [20] as the "*Modeling for control principle: The modeling and control problems are not separable and are necessarily iterative*." As an example of this principle [20],

> The Root Locus theory presumes a <u>fixed</u> system while the controller gain goes to infinity. But the fidelity of the system model *depends* upon the control gain. Hence, the same model of the system is not appropriate at *both* the vicinity of the open loop poles and the open loop zeros. [...] what the system "looks like" depends upon the controller being used.

A case in point is a rigid body with flexible appendages. As the controller bandwidth is increased, higher frequency modes become increasingly relevant to the closed-loop dynamics, and thus the control-oriented model must be chosen to accommodate properties of the controller. Along the same lines, in reference to experiments for system identification, Gevers [21] notes that "it pays to have experimental conditions that closely match the conditions in which the 'to-be-designed controller' will operate." A related phenomenon occurs when a linear controller is used for a nonlinear system whose unmodeled nonlinearities reduce the region of stabilization due to peaking [22]. These observations motivate the need for closed-loop system identification [23].

Modeling is also limited by the fact that mathematics consists of idealizations. Just as we cannot say that a physical line is straight and a physical plane is flat, we cannot say that a physical signal is in L_2 or is differentiable, or that a physical system resides in some space of operators. Finally—reinvoking ultimate unmodelability—the physical system is not a mathematical object, and thus the model error is not encompassed by any uncertainty model.

One way to potentially overcome these issues is to consider stochastic uncertainty models with probabilistic notions of stability and performance. Within this setting, however, the same issues arise; namely, the uncertainty model provides only rough guidance about the accuracy of a nominal model based on mathematical assumptions that are not satisfied in reality.

C. Robust and Adaptive Control

Linear-quadratic-Gaussian (LQG) control uses feedback to mitigate the effect of stochastic disturbances, whose statistics are known but whose time history is unknown. In addition, since LQG is an output-feedback controller, it has no knowledge of the initial state of the system. LQG thus uses feedback to mitigate two distinct sources of uncertainty. As shown in [24,25], however, LQG is sensitive to modeling errors in the system dynamics. By adopting a control-oriented model that characterizes the assumed system uncertainty, robust control seeks to overcome this shortcoming [26]. Within the context of robust control, however, model uncertainty is commonly viewed as an impediment to feedback control rather than its underlying driver, as in the case of Black's amplifier.

Since robust control accounts for the assumed uncertainty, it may sacrifice performance for the actual physical system. Adaptive control aims to overcome this tradeoff by incorporating data-driven learning. The challenge is then to learn during operation in the presence of sensor noise and unmodeled dynamics [27–29]. For systems that are open-loop unstable, an additional challenge is the need for sufficiently fast learning to avoid unacceptable transient response. In any event, learning requires data, and data require sensing; consequently, learning is a feedback process.

D. Feedback Control and Model Incompleteness

As discussed above, all models are approximate due to inaccuracy and incompleteness. Although feedback control mitigates the effect of inaccuracy, it can also exacerbate incompleteness. To illustrate how this can happen, consider the incompleteness of the standard resistor model $V = IR$. When the potential drop V across the resistor is 10 V, we have a reasonable model of the current I, but the behavior is harder to predict when the potential drop is 1000 V, in which case, as the temperature T of the resistor rises, the resistance is more accurately modeled as $R(T)$. The accuracy of a model of a physical system thus depends on the inputs to the physical system; this is the essence of incompleteness.

Of course, a model must be appropriate for its intended usage, which is directly related to the set of possible inputs. Since a feedback controller is driven by a physical system, the actual inputs to the physical system arise from the actual response of the actual physical system. Ultimate unmodelability thus implies that it is not possible to know a priori which inputs the physical system will be subjected to. Hence, the feedback controller may subject the system to input signals that are unknown when a system model is chosen and thus may be outside the assumed set of possible inputs. Therefore, feedback control compounds uncertainty by exacerbating model incompleteness.

V. Stabilization

A. Open-Loop Control of an Unstable System

In order to view stabilization from the perspective of uncertainty, we consider stabilization in terms of open-loop control. For simplicity, consider the scalar system

$$\dot{x}(t) = ax(t) + u(t) \tag{2}$$

where a is a real number, the initial state $x(0)$ is nonzero, and the goal is to bring $x(t)$ from the nonzero initial state $x(0)$ to zero. To do this, we view Eq. (2) as a *truth model*, that is, a true representation of reality.

If a is negative, then no control is needed, and $x(t) = e^{at}x(0)$ converges to zero. If, however, a is positive, then $x(t) = e^{at}x(0)$ diverges to $\pm\infty$, and thus the control $u(t)$ is needed. In this case, by choosing the open-loop control

$$u(t) = -2ae^{-at}x(0) \tag{3}$$

it follows that the solution to Eq. (2) is given by

$$x(t) = \underbrace{e^{at}x(0)}_{\text{free response}} - \underbrace{e^{at}x(0) + e^{-at}x(0)}_{\text{forced response}} \tag{4}$$

$$= e^{-at}x(0) \tag{5}$$

which converges to zero as $t \to \infty$.

Note that implementation of the open-loop control (3) requires knowledge of a and $x(0)$. In reality, however, no matter how deep our understanding of science is and no matter how advanced our engineering technology is, a and $x(0)$ cannot be known exactly, and thus the implemented control will be

$$u(t) = -2\hat{a}e^{-\hat{a}t}\hat{x}(0) \tag{6}$$

where $\hat{a}$ and $\hat{x}(0)$ are estimates of a and $x(0)$, respectively. Thus, Eq. (6) can be written as

$$u(t) = -2(a + \varepsilon)e^{-(a+\varepsilon)t}[x(0) + \eta] \tag{7}$$

where ε denotes the error in determining a and η denotes the error in determining $x(0)$. With the control (7), the solution of Eq. (2) is now given by

$$x(t) = \underbrace{e^{at}x(0)}_{\text{free response}} + \underbrace{\frac{2a + 2\varepsilon}{2a + \varepsilon}(e^{-(a+\varepsilon)t} - e^{at})[x(0) + \eta]}_{\text{forced response}} \tag{8}$$

$$= -e^{at}\left(\frac{\varepsilon}{2a+\varepsilon}x(0) + \frac{2a+2\varepsilon}{2a+\varepsilon}\eta\right) + \frac{2a+2\varepsilon}{2a+\varepsilon}e^{-(a+\varepsilon)t}[x(0)+\eta] \tag{9}$$

Hence, if either ε or η is nonzero, then, *no matter how small these numbers are*,

$$\lim_{t\to\infty}|x(t)| = \infty \tag{10}$$

and thus $x(t)$ diverges.

B. The Miracle of Feedback Stabilization

Consider now the feedback control

$$u(t) = -2ax(t) \tag{11}$$

which, with the truth model (2), yields the closed-loop dynamics

$$\dot{x}(t) = -ax(t) \tag{12}$$

whose solution $x(t) = e^{-at}x(0)$ converges to zero. As in the case of the open-loop control (7), however, the parameter a and the state $x(t)$ are not exactly known, and thus the implemented control is

$$u(t) = -2\hat{a}\hat{x}(t) \tag{13}$$

where $\hat{a}$ and $\hat{x}(t)$ are estimates of a and $x(t)$, respectively. Thus, Eq. (13) can be written as

$$u(t) = -2(a+\varepsilon)[x(t)+\eta(t)] \tag{14}$$

where ε denotes the error in determining a and $\eta(t)$ denotes noise in the measurement of $x(t)$. With the implemented feedback controller (14), the closed-loop dynamics are given by

$$\dot{x}(t) = -(a+2\varepsilon)x(t) - 2(a+\varepsilon)\eta(t) \tag{15}$$

whose solution is

$$x(t) = e^{-(a+2\varepsilon)t}x(0) - 2(a+\varepsilon)\int_0^t e^{-(a+2\varepsilon)(t-\tau)}\eta(\tau)\,\mathrm{d}\tau \tag{16}$$

Assuming that $2|\varepsilon| < a$ and that there exists $\eta_0 > 0$ such that, for all $t \geq 0$, $\eta(t) < \eta_0$, it follows from Eq. (16) that

$$\begin{aligned}|x(t)| &\leq e^{-(a+2\varepsilon)t}|x(0)| + 2(a+\varepsilon)\int_0^t e^{-(a+2\varepsilon)\tau}\,\mathrm{d}\tau\eta_0 \\ &= e^{-(a+2\varepsilon)t}|x(0)| + \frac{2(a+\varepsilon)\eta_0}{a+2\varepsilon}(1-e^{-(a+2\varepsilon)t})\end{aligned} \tag{17}$$

Hence,

$$\limsup_{t\to\infty}|x(t)| \leq \frac{2(a+\varepsilon)\eta_0}{a+2\varepsilon} \tag{18}$$

which implies that $x(t)$ is bounded with an asymptotic bound determined by ε and η_0.

The contrast between Eqs. (10) and (18) could not be more stark. For the open-loop control (7), the state diverges due to infinitesimally imperfect modeling information, whereas, for the closed-loop control (11), the state remains bounded with a bound that can be managed by reducing the level of the sensor noise. Consequently, feedback control has drastically mitigated the effect of uncertainty and enhanced robustness to physical variations.

VI. Enhancing Robustness to Physical Variations via Feedback Control

As another example of uncertainty mitigation, consider the servo loop in Fig. 1 with the system (2) as the truth model and the proportional controller $u(t) = ke(t)$. The command signal is the step function $r(t) = \bar{r}$, the error is $e(t) = r(t) - y(t)$, and a is assumed to be negative so that the uncontrolled system is asymptotically stable.

In the case where the feedback path is absent, the asymptotic output is

$$y_\infty(a) = \lim_{t\to\infty} y(t) = -\frac{k\bar{r}}{a} \tag{19}$$

Therefore, for the desired setpoint r_{des}, choosing k to be nonzero and setting $\bar{r} = -(ar_{\mathrm{des}}/k)$ yields $y_\infty(a) = r_{\mathrm{des}}$. Alternatively, in the case where $k > a$ and the loop is closed, the asymptotic output is

$$y_\infty(a) = \lim_{t\to\infty} y(t) = \frac{k\bar{r}}{k-a} \tag{20}$$

Therefore, setting $\bar{r} = ((k-a)/k)r_{\mathrm{des}}$ yields $y_\infty(a) = r_{\mathrm{des}}$. Note that, for both controllers, the commanded setpoint $\bar{r}$ depends on a model of the system, namely, the parameter a.

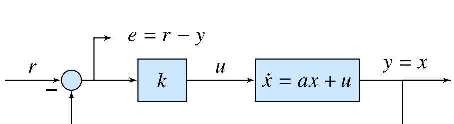

Fig. 1 Feedback stabilization of a scalar system.

To determine how a physical variation δa of a affects $e_\infty(a) = \bar{r} - y_\infty(a)$, note that

$$e_\infty(a+\delta a) \approx e_\infty(a) + \frac{de_\infty(a)}{da}\delta a \tag{21}$$

In the case of Eq. (19),

$$\frac{de_\infty(a)}{da} = -\frac{k\bar{r}}{a^2} \tag{22}$$

whereas, in the case of Eq. (20),

$$\frac{de_\infty(a)}{da} = \frac{k\bar{r}}{(k-a)^2} \tag{23}$$

Since $a < 0$, dividing the right-hand side of Eq. (23) by the right-hand side of Eq. (22) yields the parameter-sensitivity reduction ratio

$$\sigma(a) \triangleq -\frac{a^2}{(k-a)^2} = -\frac{1}{\left(1+\frac{k}{|a|}\right)^2} \tag{24}$$

which, for all $k > 0$, has magnitude less than 1. The asymptotic output is thus more robust to physical variations in a in the case of feedback control.

An alternative way to reach the same conclusion is to use the fact that the Laplace transform of the output is given in terms of the complementary sensitivity function $T(s) = kG(s)/(1+kG(s))$ by

$$\hat{y}(s) = T(s)\hat{r}(s) = \frac{kG(s)\bar{r}}{s+skG(s)} \tag{25}$$

and thus the Laplace transform of the error is given in terms of the sensitivity function $S(s) = 1/(1+kG(s))$ by

$$\hat{e}(s) = S(s)\hat{r}(s) = \frac{\bar{r}}{s+skG(s)} \tag{26}$$

Then, it can be shown [8, p. 22] that

$$S(s) = \frac{G(s)}{T(s)}\frac{dT(s)}{dG(s)} \tag{27}$$

which can be rewritten as

$$\frac{dT(s)}{T(s)} = S(s)\frac{dG(s)}{G(s)} \tag{28}$$

Setting $s = j\omega$ to obtain the frequency response, it follows from Eq. (28) that, at frequencies ω for which $|S(j\omega)| < 1$, the effect of variations of G on T is mitigated by S. Finally, noting that

$$S(j0) = \frac{1}{1+k\frac{1}{-a}} = \frac{-a}{k-a} \tag{29}$$

it follows from Eq. (24) that

$$\sigma(a) = -S(j0)^2 \tag{30}$$

which relates the DC gain of the sensitivity transfer function to the parameter-sensitivity reduction ratio.

VII. Uncertainty Reduction via Feedback Control

Although the feedback controller (11) enhances robustness to physical variations, it is interesting to ask whether it also *reduces* uncertainty. To investigate this question, consider, in place of Eq. (11), the feedback control

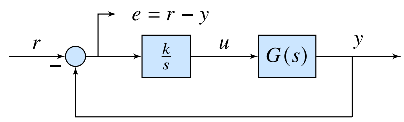

Fig. 2 Servo loop with integral controller.

$$u(t) = -kx(t) \tag{31}$$

and assume that this control law is applied to Eq. (2) with no knowledge of a. If $x(t)$ diverges, then we learn that $k < a$, whereas, if $x(t)$ remains bounded, then we learn that $k > a$. The qualitative response of the closed-loop system thus provides rudimentary knowledge about the parameter a, thus reducing uncertainty, at least within the context of the nominal model (2). Of course, this experiment may not be the safest way to learn.

As another example, consider the servo loop in Fig. 2 with the system $G(s)$ and the integral controller k/s, where k is chosen such that the closed-loop dynamics are asymptotically stable. This example mimics the fundamental discovery of Black in 1928 [30,31]. For the setpoint command $r(t) = \bar{r}$, the Laplace transform of the error is given by

$$\hat{e}(s) = S(s)\hat{r}(s) = \frac{\bar{r}}{s + kG(s)} \tag{32}$$

The final value theorem thus implies

$$e_\infty \triangleq \lim_{t\to\infty} e(t) = \lim_{s\to 0} s\hat{e}(s) = \lim_{s\to 0} \frac{s\bar{r}}{s + kG(s)} = 0 \tag{33}$$

as expected for integral control. Furthermore, the Laplace transform of the control input is

$$\hat{u}(s) = \frac{k\bar{r}}{s^2 + ksG(s)} \tag{34}$$

and thus, as a consequence of the internal model principle [8, p. 49], the final value theorem implies

$$u_\infty \triangleq \lim_{t\to\infty} u(t) = \lim_{s\to 0} s\hat{u}(s) = \lim_{s\to 0} \frac{k\bar{r}}{s + kG(s)} = \frac{\bar{r}}{G(0)} \tag{35}$$

Hence,

$$G(0) = \frac{\bar{r}}{u_\infty} \tag{36}$$

which, since u_∞ is known, shows that the asymptotic control input reveals the DC gain of the system. This example demonstrates that a controller that possesses memory is able to learn through feedback.

VIII. The Fundamental Problem of Feedback Control

A. Statement of the Problem and Control Engineering Practice

The fundamental problem of feedback control is the fact that *ultimate unmodelability makes it impossible to guarantee by mathematics, computation, and data analysis that a feedback controller will work reliably.* Nevertheless, most feedback control systems operate successfully in practice. There are multiple reasons for this success. First, few control systems are entirely model based. In practice, control systems are fine-tuned—tweaked—based on judgment and experience to work under changing, messy, unmodelable conditions. Tweaking helps to overcome the unmodelable dimensions of complex reality while allowing the user to satisfy performance criteria that are difficult to attain with textbook methods. Hence, despite the substantial development of control theory, feedback control is partly an art, where experience and judgment are relied on to accommodate unmodelable details.

Furthermore, feedback control systems are often engineered under stringent levels of modeled uncertainty with advantageous architectures. Sensors and actuators and their noise properties are carefully modeled; disturbance and command spectra are characterized; operational modes are monitored to facilitate gain scheduling, switching, and fault detection; and control architectures that merge sensors, actuators, processors, and data links are designed to minimize the complexity of the intervening dynamics by removing right-half-plane zeros, severe nonlinearities, and other effects that limit achievable performance and gain and phase margins [8]. In other words, domain knowledge and control-theoretic principles are used to design control architectures that are amenable to simple controllers that minimize sensitivity to modeled uncertainty.

B. Control Theory and Mathematical Idealizations

While all branches of science use mathematics to a greater or lesser extent, control theory has a special affinity for mathematics since its ideas transcend specific applications. In particular, the concepts and methods of control theory are often presented in a definition–theorem–proof format that facilitates abstraction of essential features and properties.

Mathematical abstractions are embraced for various reasons, such as efficiency, richness, beauty, simplicity, theoretical effectiveness, and—last but not least—practical utility. These abstractions, which are proposed, culled, and refined over long periods of time, are *idealizations*—intellectual constructions that facilitate the ability to think and reason about the physical world. Notions such as smoothness and convergence, for example, are bedrock ideas in mathematics but are subjective in practical applications, where all data are finite and noisy. Mathematical idealizations are used to construct models that approximate complex reality. How, then, can we know how close a given model is to something that cannot be known—especially since any attempt to quantify "closeness" depends on idealizations as well as perfect knowledge of complex reality? The difficulty is that the hypotheses of a theorem are mathematical idealizations, and thus it is not possible to determine whether mathematical closeness will capture the complex reality of a physical system. In other words, *no mathematical hypothesis involving complex reality can be verified.*

Engineers—especially control engineers—know that mathematical idealizations must be applied with discretion to physical systems. Margins provide confidence, but these are based on mathematical idealizations that account only for *modelable* uncertainty. On the other hand, feedback control enhances robustness to physical variations whether or not those variations can be modeled. So here is the challenge: Since *un*modelable uncertainty is *un*knowable, how can we use models to *guarantee* that a feedback controller will stabilize a physical system by sufficiently enhancing robustness to *un*modelable physical variations?

In feedback-control practice, uncertainty is mitigated by applying attention to detail, operational safeguards, and tweaking based on experience to deliver controllers that work reliably in plausible scenarios. But—and this is the key point—when we *are* successful, how can we know *why* a given controller worked? We may know a lot about some features of the system, but we cannot know everything about it—which is why feedback control was used in the first place. Ultimate unmodelability says that complex reality cannot be captured by any model, and thus the extent to which the success of a feedback controller is due to its ability to enhance robustness to physical variations is unknown.

IX. Toward Scientifically Meaningful Feedback Control Experiments

One approach to confronting the fundamental problem of feedback control is to perform feedback control experiments that reveal unmodeled effects and their potential impact on stability and performance. Although these experiments could be aimed at determining the ability of a controller to mitigate uncertainty, it is important to recognize that *uncertainty is a subjective quantity that depends on the state of knowledge of the experimentalist*. In other words, since different people have different knowledge and thus different levels of uncertainty, an experiment that depends on the knowledge state of the experimentalist violates separation between the experiment and the experimentalist. One solution to this quandary is to replace the subjective quantity of uncertainty mitigation with the objective quantity of robustness to physical variations. To explore how such experiments might be designed, performed, and analyzed, we will first consider physical experiments and then discuss how these can be supported by computational experiments.

A. Physical Feedback-Control Experiments

A scientific experiment begins with a hypothesis in the form of a theoretical prediction and ends with data that sheds light on the hypothesis—preferably, by either validating it (which is not possible) or invalidating it (which is the falsifiability paradigm of Popper), thereby discriminating between competing theories [16, pp. 18, 19]. In practice, experiments are often most valuable for uncovering unexpected effects leading to new theories and thus new predictions [18, pp. 7–10]. As stated in [32], however, "many experiments on robotics and applied control systems are often more like demonstrations than scientific experiments (which include hypotheses and controlled variables)." Within the context of uncertainty mitigation, a natural objective of a feedback control experiment is to *examine the validity of the hypothesis that the feedback controller is robust to a specified class of physical variations in the physical system being controlled.* A meaningful experiment would therefore need to encompass physical variations that can be realized and validated. Another objective is to uncover physical variations that can lead to performance degradation or failure. A relevant experiment would therefore need to probe the boundaries of system operation, even to the point of failure, which, for aerospace applications, typically corresponds to instability.

Since variations in a physical system are ultimately unmodelable, this experimental objective is not completely achievable. Nevertheless, a first step is to design a feedback-control experiment to provide data that can reveal hidden features to the greatest extent possible. A key point is to distinguish modeling information used to design and fine-tune a controller from modeling information used to analyze closed-loop response data. In effect, the idea is to use knowledge from an unrefined uncertainty model as the basis of controller synthesis, but use knowledge of a more refined uncertainty model to assess the robustness of the feedback controller to actual physical variations, at least to the extent that those variations are known. To do this, a system designed for a control experiment must be precisely and richly perturbable, which means that it can be varied in as many different reasonably well-known ways as possible in order to assess the resulting closed-loop response as a means for quantifying robustness to physical variations. In addition, it may be necessary to instrument the system not only with the sensors needed for feedback control but also with auxiliary sensors (such as external optical diagnostics that do not impact the system dynamics) that can provide data for exposing hidden details. Under this scenario, competing control-oriented models can be invoked and evaluated.

Since no physical system can be exactly modeled, a second objective of a physical control experiment is to discover potentially consequential effects that were captured by neither the unrefined nor the refined uncertainty models. This objective directly impacts control practice, as discussed below.

B. Computational Feedback-Control Experiments

Simulation can be used to perform a vast number of computational feedback-control experiments that can be informative in various ways. Aside from details such as machine constants, digital simulation (including stochastic simulation) is repeatable, which is crucial for computational experimentation. Furthermore, aside from variations in compilers and numerical issues such as roundoff, simulation provides omniscience, and thus the user can usually verify the code and the selected parameter values. This is not always the case, however, since some simulation codes are proprietary, extremely complex, or sparsely documented. Of course, the accuracy of simulation relative to complex reality depends on the accuracy of the chosen models and the errors introduced by spatial and temporal discretization.

Simulations can be used to probe a rich collection of system perturbations, including variations of linear dynamics, noise with specified statistical characteristics, and nonlinearities of endless variety, all of which are, of course, known to the computational experimentalist. The richness of these constructions can easily surpass the ability to vary the properties of a physical experiment. The computational environment is thus valuable for probing the ability of a feedback controller to enhance robustness to physical variations up to—and including—the failure boundary without destructive consequences. Simulation is limited, however, by the fact that a computer program includes only those effects that the modeler includes in the code. Ultimate unmodelability prevents us from simulating everything that can happen in complex reality, and thus simulation may not be able to discover physical variations that can adversely affect stability and performance.

Beyond probing the stability and performance robustness of a given controller, simulation is valuable for investigating anomalies observed in a physical control system. In particular, computational control experiments provide a diagnostic capability that can use experimental data to analyze the performance of a physical control system, especially when the closed-loop behavior is unexpected. A carefully designed, well-instrumented, and multidimensionally rich physical feedback-control experiment can yield a substantial amount of data for subsequent analysis. When a physical feedback-control experiment exhibits unexpected behavior, simulation can provide a what-if environment for deducing the cause of the anomaly including the ability to diagnose the reasons for failure.

The ease of varying the simulation model yields an embarrassment of riches, however. Simulation can be used to probe dozens of types of perturbations of varying magnitude. But what about simultaneous perturbations of two, three, or more different types? The number of possibilities is combinatorially large, and thus, when simulation is used to probe the performance of a physical feedback-control experiment, multiple effects can confound the ability to determine what actually occurred. This challenge is related to the distinction in experiment design between full factorial design and fractional factorial design [33].

X. Implications of Physical and Computational Control Experiments

A. Implications for Control Research

Experiments often reveal unexpected behavior that motivates theoretical investigations. Physical and computational feedback-control experiments can provide a serendipitous source of research questions and challenges that warrant theoretical consideration, especially relating to the mechanisms by which feedback control mitigates uncertainty and enhances robustness to physical variations.

B. Implications for Control Education

In control education, computational experiments can be used to assess the performance of a closed-loop system under various scenarios involving system perturbations, initial conditions (to determine the domain of attraction), commands, and disturbances. For physical control experiments, with the advent of low-cost processors, classroom experiments are commonplace. Although control laboratories help students to appreciate the value of modeling, identification, and controller synthesis, these "experiments," while educational, are essentially demonstrations. By investigating the extent to which a controller enhances robustness to physical variations, these demonstrations can be designed to be scientifically meaningful from an educational standpoint.

C. Implications for Control Practice

The principal impediment to experimentation in control practice is the cost and risk of performing experiments on full-scale vehicles and installations. The observation in [14, pp. 16, 17] quoted above that systems that are expensive or dangerous are not viable candidates for control experiments—a point made graphically in [34]—implies the need for scaled-down testbeds that capture the key physics and features of costly and dangerous systems. The challenge is then to determine whether or not crucial physical effects that may impact control of a full-scale system are also manifested in the testbed.

A control experiment is maximally informative when it is operated under rare or extreme conditions that reveal the failure boundaries. These experiments may include low-probability perturbations, aggressive maneuvers, and severe disturbances. The risks to a full-scale system under extreme experimental conditions are unacceptable, however, and thus tests on full-scale systems are confined to validation, verification, durability, fault-checking, and certification, where the goal is not to determine the ultimate limits to safe operation but rather to confirm that the system will operate reliably under the specifications and scenarios for which it was designed. In aerospace applications, this is done by test pilots. In practice, however, new and unexpected phenomena may arise as a novel system is scaled up and operates under real-world conditions. *This is especially the case when new technologies are developed and there is limited knowledge concerning failure modes and boundaries.* Although much can be learned when a full-scale system fails, the human and societal cost of such failures motivates and justifies the need for maximally informative control experiments at each stage of the development cycle [35].

Among the numerous X-planes developed over more than 75 years, the X-43A provides a case study of the implications of a full-scale engineering experiment. This program involved two phases, namely, a Mach 7 flight test followed by a Mach 10 flight test. For the first flight, the Mishap Investigation Board (MIB) investigated the cause of a failure related to the Hyper-X Launch Vehicle (HXLV) [36, p. 9]:

> The MIB report root cause specified that the "HXLV failed because the vehicle control system design was deficient for the trajectory flown due to inaccurate analytical models."

The response to this outcome included the following steps prior to return to flight (RTF) [36, p. 9]:

> The probability of occurrence and magnitude of impact were evaluated for each risk and mitigations for these were identified. Most of the RTF effort was spent on the risk reduction activities that were derived from these mitigations. Mitigations included additional testing and analysis for hardware and software, model and uncertainties evaluation, update and enhancement, and independent simulations and review where appropriate.

Lessons learned from the first flight led to modifications and subsequent success in the second phase flight test [36, p. 11].

XI. Conclusions

The goal of feedback control is to achieve the best possible performance through closed-loop uncertainty mitigation, recognizing that performance is the objective of *control*, but uncertainty is the *raison d'être* for *feedback* control. As such, feedback control is at odds with the main objective of science: Instead of trying to model complex reality to higher levels of accuracy, feedback control seeks the best possible performance despite unknown initial conditions, parameters, disturbances, dynamics, and nonlinearities as well as unmodeled and unmodelable physical variations. As a consequence of this objective, the fundamental problem of feedback control is the realization that it is not possible to definitively determine either the extent to which a feedback controller has succeeded or the reasons for which it has failed. This problem is not due to the shortcomings of a particular feedback control technique, but rather is due to the facts that 1) no control-oriented model can fully capture the complex reality of any physical system, and 2) feedback may mask the effect of physical variations.

To address future aerospace challenges, this article issued a call to action to establish a culture of scientifically meaningful feedback control experimentation, both physical and computational. In keeping with the underlying driver of feedback, a feedback control experiment can be aimed at determining the extent to which a feedback controller succeeds or fails to enhance robustness to physical variations. These experiments can be designed to probe the physical system in order to reveal physical variations that are masked by feedback or that contribute to degraded performance or instability. Theory and techniques for designing such experiments remain a challenge for future research and practice.

Acknowledgments

I wish to thank Danny Abramovitch, Randy Beard, Gary Bernstein, Karthik Duraisamy, James Forbes, Kunal Garg, Jesse Hoagg, Dan Inman, Aseem Islam, Harris McClamroch, Pete Seiler, Panagiotis Tsiotras, Elena Valcher, and the editor and reviewers for many helpful suggestions.

References

[1] Jonsson, E., Risoa, C., Luppa, C. A., Cesnik, C. E. S., Martins, J. R. R. A., and Epureanu, B. I., "Flutter and Post-Flutter Constraints in Aircraft Design Optimization," *Progress in Aerospace Sciences*, Vol. 109, 2019, Paper 100537.
https://doi.org/10.1016/j.paerosci.2019.04.001

[2] Morelli, E. A., "Flight-Test Experiment Design for Characterizing Stability and Control of Hypersonic Vehicles," *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 3, 2009, pp. 949–959.
https://doi.org/10.2514/1.37092

[3] Leyva, I. A., "The Relentless Pursuit of Hypersonic Flight," *Physics Today*, Vol. 70, No. 11, 2017, pp. 30–36.
https://doi.org/10.1063/PT.3.3762

[4] Hassanaly, M., and Raman, V., "Classification and Computation of Extreme Events in Turbulent Combustion," *Progress in Energy and Combustion Science*, Vol. 87, 2021, Paper 100955.
https://doi.org/10.1016/j.pecs.2021.100955

[5] Smith, R. C., *Uncertainty Quantification: Theory, Implementations, and Applications*, SIAM, Philadelphia, 2013.

[6] Astrom, K. J., and Murray, R. M., *Feedback Systems*, 2nd ed., Princeton Univ. Press, Princeton, NJ, 2021.

[7] Mitter, S., and Tannenbaum, A., "The Legacy of George Zames," *IEEE Transactions on Automatic Control*, Vol. 43, No. 5, 1998, pp. 590–595.
https://doi.org/10.1109/TAC.1998.668826.

[8] Skogestad, S., and Postlethwaite, I., *Multivariable Feedback Control*, 2nd ed., Wiley, Hoboken, NJ, 2005.

[9] Mayr, O., *The Origins of Feedback Control*, MIT Press, Cambridge, MA, 1970.

[10] Bennett, S., *A History of Control Engineering 1800-1930*, Inst. of Electrical and Electronics Engineers, New York, 1979.

[11] Bennett, S., *A History of Control Engineering 1930–1955*, Inst. of Electrical and Electronics Engineers, New York, 1993.

[12] de Latil, P., *Thinking by Machine: A Study of Cybernetics*, Houghton Mifflin, Boston, MA, 1956.

[13] Kelly, K., *Out of Control: The Rise of Neo-Biological Civilization*, Addison-Wesley, Reading, MA, 1994.

[14] Fleming, W. H. (ed.), *Future Directions in Control Theory: A Mathematical Perspective*, SIAM, Philadelphia, 1988.

[15] Murray, R. M. (ed.), *Control in an Information Rich World: Report of the Panel on Future Directions in Control, Dynamics, and Systems*, SIAM, Philadelphia, 2003.

[16] Strevens, M., *The Knowledge Machine: How Irrationality Created Modern Science*, Liveright Publishing Company, New York, 2020.

[17] Box, G. E. P., "Robustness in the Strategy of Scientific Model Building," *Robustness in Statistics*, Academic Press, New York, 1979, pp. 201–236.

[18] Box, G. E. P., and Draper, N. R., *Empirical Model-Building and Response Surfaces*, Wiley, Hoboken, NJ, 1987.

[19] Parker, J. T., Serrani, A., Yurkovich, S., Bolender, M. A., and Doman, D. B., "Control-Oriented Modeling of an Air-Breathing Hypersonic Vehicle," *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 3, 2007, pp. 856–869.
https://doi.org/10.2514/1.27830

[20] Skelton, R. E., "On the Structure of Modeling Errors and the Inseparability of the Modeling and Control Problems," *Model Error Concepts and Compensation*, Pergamon Press, Oxford, 1986, pp. 13–20.

[21] Gevers, M., "Identification for Control: From the Early Achievements to the Revival of Experiment Design," *European Journal of Control*, Vol. 11, No. 11, 2005, pp. 1–18.

[22] Kokotovic, P. V., "The Joy of Feedback: Nonlinear and Adaptive," *IEEE Control Systems Magazine*, Vol. 12, 1992, pp. 7–17. https://doi.org/10.1109/37.165507
[23] Forssell, U., and Ljung, L., "Closed-Loop Identification Revisited," *Automatica*, Vol. 35, 1999, pp. 1215–1241. https://doi.org/10.1016/S0005-1098(99)00022-9
[24] Doyle, J. C., "Guaranteed Margins for LQG Regulators," *IEEE Transactions on Automatic Control*, Vol. AC-23, No. 4, 1978, pp. 756–757.
[25] Nesline, F. W., and Zarchan, P., "Why Modern Controllers Can Go Unstable in Practice," *AIAA Journal of Guidance*, Vol. 7, No. 4, 1984, pp. 495–500. https://doi.org/10.2514/3.19884.
[26] Safonov, M. G., "Origins of Robust Control: Early History and Future Speculations," *Annual Reviews in Control*, Vol. 36, No. 2, 2012, pp. 173–181. https://doi.org/10.1016/j.arcontrol.2012.09.001
[27] Rohrs, C., Valavani, L., Athans, M., and Stein, G., "Robustness of Continuous-Time Adaptive Control Algorithms in the Presence of Unmodeled Dynamics," *IEEE Transactions on Automatic Control*, Vol. 30, No. 9, 1985, pp. 881–889. https://doi.org/10.1109/TAC.1985.1104070
[28] Anderson, B. D. O., "Failures of Adaptive Control Theory and Their Resolution," *Communications in Information & Systems*, Vol. 5, No. 1, 2005, pp. 1–20. https://doi.org/10.4310/CIS.2005.v5.n1.a1
[29] Anderson, B. D. O., and Dehghani, A., "Challenges of Adaptive Control—Past, Permanent and Future," *Annual Reviews in Control*, Vol. 32, No. 2, 2008, pp. 123–135. https://doi.org/10.1016/j.arcontrol.2008.06.001
[30] Black, H. S., "Inventing the Negative Feedback Amplifier," *IEEE Spectrum*, Vol. 14, 1977, pp. 55–60. https://doi.org/10.1109/MSPEC.1977.6501721
[31] Kline, R., "Harold Black and the Negative-Feedback Amplifier," *IEEE Control Systems Magazine*, Vol. 13, 1993, pp. 82–85. https://doi.org/10.1109/37.229565.
[32] How, J., "Why Do Experiments?" *IEEE Control Systems Magazine*, Vol. 39, 2019, pp. 4–6. https://doi.org/10.1109/MCS.2018.2888682
[33] Montgomery, D. C., *Design and Analysis of Experiments*, 10th ed., Wiley, Hoboken, NJ, 2020.
[34] Stein, G., "Respect the Unstable," *IEEE Control Systems Magazine*, Vol. 23, 2003, pp. 12–25. https://doi.org/10.1109/MCS.2003.1213600
[35] Robison, P., *Flying Blind: The 737 MAX Tragedy and the Fall of Boeing*, Doubleday, New York, 2021.
[36] Marshall, L. A., Corpening, G. P., and Sherrill, R., "A Chief Engineer's View of the NASA X-43A Scramjet Flight Test," *Proceedings of the AIAA/CIRA 13th International Space Planes and Hypersonics Systems and Technologies*, AIAA, Reston, VA, 2005, pp. 1–20; also AIAA Paper 2005-3332, 2005. https://doi.org/10.2514/6.2005-3332