Chapter 1
Introduction

"I think it's very important to have a feedback loop, where you're constantly thinking about what you've done and how you could be doing it better."

- Elon Musk

1.1 Feedback in engineering and living systems

Feedback processes are critical aspects of most living and engineering systems. *Feedback* occurs when the output of a system influences the input of the same system. *Feedback control* is a process of creating such a feedback loop to modify the behavior of a dynamical system through actuation that is informed by measurements of the system.

The very existence of humans and other endothermic animals is based on a robust feedback control: They maintain their body temperature within narrow limits despite a large range of environmental conditions and disturbances. This temperature regulation is performed with temperature monitoring and control actions, such as increasing metabolism or sweating. Similarly, air conditioning also keeps a room temperature in a narrow interval by heating or cooling via a ventilating air stream.

The world around us is actively shaped by feedback processes, from the meandering path of a river to the gene regulation that occurs inside every cell in our body. A child's education may be considered a feedback control task, where parental and societal feedback guide the child's actions towards a desired goal, such as socially acceptable behavior and the child becoming a productive member of society. The order achieved in a modern society is the result of a balance of interests regulated through active policing and the rule of laws, which are in turn shaped by a collective sense of justice and civil rights. Financial markets and portfolio management are also feedback processes based on a control logic of buying and selling stocks to reach an optimal growth or profit at a given risk over a certain time horizon. In fact, currency inflation is actively manipulated by changing interest rates and issuing bonds. Our very thoughts and actions are intimately related to a massively parallel feedback architecture in our brain and nervous system, whereby external stimuli are collected and assimilated, decisions are made, and control actions are executed, resulting in our interaction with the world. Finally, the earth's climate and temperature are maintained through a delicate balance of forcing from sources including solar irradiance, greenhouse gases, vegetation, aerosols and cloud formation, many of which are coupled through feedback.

The feedback control of fluid systems is an immensely important challenge with profound implications for technologies in energy, security, transportation, medicine, and many other endeavors. Flow control is an academically exciting research field undergoing rapid progress — comprising many disciplines, including theoretical, numerical and experimental fluid mechanics, control theory, reduced-order modeling, nonlinear dynamics and machine learning techniques. Flow control has applications of epic proportion, such as drag reductions of cars, trucks, trains, ships and submarines, lift increase of airplanes, noise reduction of ground or airborne transport vehicles, combustion efficiency and NO_X reduction, cardiac monitoring and intervention, optimization of pharmaceutical and chemical processes and weather control. The flows found in most engineering applications are turbulent, introducing the complexities of high-dimensionality, multi-scale structures, strong nonlinearities and frequency crosstalk as additional challenges.

Feedback turbulence control shares a significant overlap with the other feedback systems described above, in the sense that

- the control goal can be defined in mathematical terms;
- the control actions are also in a well-defined set;
- the unforced system has its own internal *chaotic* nonlinear dynamics, where neighboring states may rapidly diverge to different behaviors within the prediction horizon;
- the full state is only partially accessible by limited sensors;
- there is an underlying evolution equation (i.e., the Navier-Stokes equation) which provides a high-fidelity description of the system, but may not be useful for control decisions in a real-life experiment.

The last three properties are a generic consequence of high-dimensional nonlinear dynamics. However, unlike many of the systems described above, turbulence control is more benign, as the system quickly forgets its past treatment and the control experiments tend to be more reproducible. In other words, the unforced and forced systems have a statistical stationarity, i.e. statistical quantities like mean values and variances are well defined. Regardless, feedback turbulence control is significantly more complex than most academic control theory tasks, such as stabilization of an inverted pendulum. Hence, improving feedback control architectures that work for turbulence control may have significant impact in other complex systems.

Nature offers compelling examples of feedback flow control that may provide inspiration for engineering efforts. For example, eagles are expert flyers, capable of rising on thermals or landing gently on a rock or tree despite strong wind gust perturbations and other challenging weather conditions. These maneuvers require active feedback control by sensing the current position and velocity and dynamically adjusting the control actions involving the motion of wings and feathers. An eagle's flight is robust to significant uncertainty in the environment and flight conditions, including harsh weather and significant changes to its own body, including mass, geometry, and wing shape. It is unlikely that eagles, or other flying animals, such as birds, bats, or insects, are operating based on a high-fidelity model of the underlying Navier-Stokes equations that govern fluid flow. Instead, it is more likely that these animals have adapted and learned how to sense and modify dominant coherent structures in the fluid that are most responsible for generating forces relevant for flight. Airplanes similarly move on prescribed trajectories at predetermined speeds under varying wind and weather conditions by adjusting their control surfaces, such as flaps and ailerons, and engine thrust. However, there is still a tremendous opportunity to improve engineering flight performance using bio-inspired techniques.

This book outlines the use machine learning to design control laws, partially inspired by how animals learn control in new environments. This *machine learning control* (MLC) provides a powerful new framework to control complex dynamical systems that are currently beyond the capability of existing methods in control.

1.2 Benefits of feedback control

Figure 1.1 illustrates a general feedback control system. The physical system, also called the *plant*, is depicted in the blue box. The system is monitored by sensors **s** and manipulated by actuators **b** through a control logic depicted in the yellow box. Moreover, the plant is subjected to sensor noise and exogenous disturbances **w** and the control shall be optimized with respect to a cost function J.

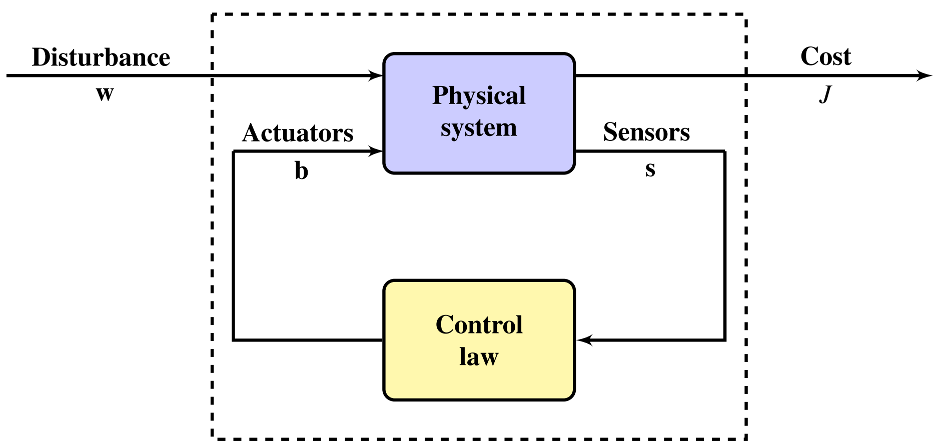

Fig. 1.1 General optimization framework for feedback control. The behavior of the physical system is modified by actuators (inputs, **b**) through a control law informed by sensor measurements of the system (outputs, **s**). The control logic is designed to shape the closed-loop response from the exogenous disturbances **w** to a high-level objective encoded by the cost function J.

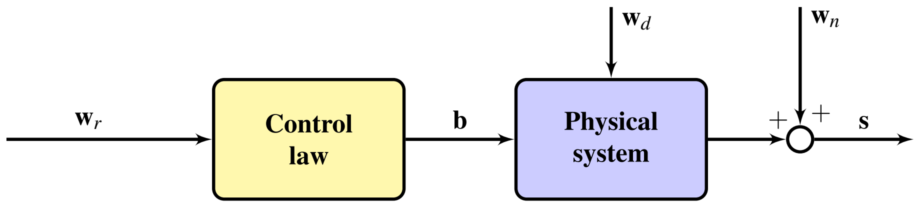

Fig. 1.2 Open-loop control diagram. A reference signal $\mathbf{w}_r$ is fed directly into an open-loop controller which specifies a pre-determined actuation signal $\mathbf{b}$. External disturbances $(\mathbf{w}_d)$ and sensor noise $(\mathbf{w}_n)$, as well as un-modeled system dynamics and uncertainty, degrade the overall performance.

One example of an optimization tasks is drag reduction. A physically meaningful optimization problem penalizes the actuation. A well-posed drag reduction problem requests a minimization of the power required to overcome drag J_{drag} plus the invested actuation power J_{act}, i.e. the net gain $J = J_{\text{drag}} + J_{\text{act}}$. Other examples include lift increase, mixing increase and noise reduction. To keep an airplane on a desired trajectory, the thrust and lift need to be kept at a well-defined level. Thus, the control task becomes a *reference tracking problem*, in which a reference force — or other quantity — is commanded. In this case, the cost function penalizes the deviation from the desired state and the invested actuation level.

In the case of reference tracking, it is natural to first consider the open-loop control architecture shown in Fig. 1.2. In this strategy, the actuation signal $\mathbf{b}$ is chosen based on knowledge of the system to produce the desired output that matches the commanded reference signal. This is how many toasters work, where the heating element is turned on for a fixed amount of time depending on the desired setting. However, open-loop control is fundamentally incapable of stabilizing an unstable system, such as an inverted pendulum, as the plant model would have to be known perfectly without any uncertainty or disturbances. Open-loop control is also incapable of adjusting the actuation signal to compensate for disturbances to the system.

Instead of making control decisions purely based on the desired reference, as in open-loop control, it is possible to *close the loop* by feeding back sensor measurements of the system output so that the controller knows whether or not it is achieving the desired goal. This *closed-loop feedback control* diagram is shown in Fig. 1.3. Sensor-based feedback provides a solution to the issues that occur with open-loop control. It is often possible to stabilize an unstable system with the aid of sensor feedback, whereas it is never possible to stabilize an unstable system in open-loop. In addition, closed-loop control is able to compensate for external disturbances and model uncertainties, both of which should be measured in the sensor output.

Summarizing, feedback control is, for instance, necessary for the following tasks:

- Optimize a state or output with respect to a given cost function;

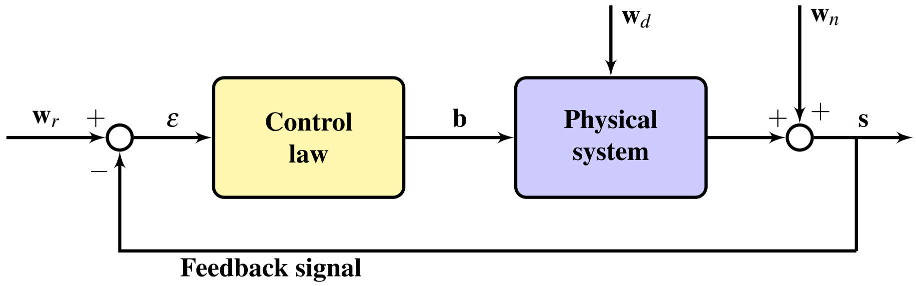

Fig. 1.3 Closed-loop feedback control diagram. The sensor signal $\mathbf{s}$ is fed back and subtracted from the reference signal $\mathbf{w}_r$. The resulting error ε is used by the controller to specify the actuation signal $\mathbf{b}$. Feedback is generally able to stabilize unstable plant dynamics while effectively rejecting disturbances $\mathbf{w}_d$ and attenuating noise $\mathbf{w}_n$.

- Stabilize an unstable system;
- Attenuate sensor noise;
- Compensate for exogenous disturbances and model uncertainty.

Mathematical formulation of feedback control task

There is a powerful theory of feedback control based on dynamical systems. In this framework, the plant is modeled by an input–output system:

$$\frac{d}{dt}\mathbf{a} = \mathbf{F}(\mathbf{a}, \mathbf{b}, \mathbf{w}_d), \tag{1.1a}$$

$$\mathbf{s} = \mathbf{G}(\mathbf{a}, \mathbf{b}, \mathbf{w}_n), \tag{1.1b}$$

consisting of a coupled system of possibly nonlinear differential equations in a state variable $\mathbf{a} \in \mathbb{R}^{N_a}$, where N_a is the dimension of the state. The actuation input is given by the vector $\mathbf{b} \in \mathbb{R}^{N_b}$ and this input directly affects the state dynamics in Eq. (1.1a), along with exogenous disturbances $\mathbf{w}_d$. The sensor measurements are given by the output vector $\mathbf{s} \in \mathbb{R}^{N_s}$, and these measurements may be nonlinear functions of the state $\mathbf{a}$, the control $\mathbf{b}$ and noise $\mathbf{w}_n$.

The control task is generally to construct a controller

$$\mathbf{b} = \mathbf{K}(\mathbf{s}, \mathbf{w}_r), \tag{1.2}$$

so that the closed-loop system has desirable properties in terms of stability, attenuation of noise, rejection of disturbances, and good reference tracking characteristics. The commanded reference signal is $\mathbf{w}_r$. These factors are encoded in the cost function J, which is generally a function of the sensor output, the actuation input, and the various external signals $\mathbf{w}_r$, $\mathbf{w}_d$, and $\mathbf{w}_n$.

With a well-designed sensor-based feedback control law, it is often possible to obtain a closed-loop system that performs optimally with respect to the chosen cost function and is robust to model uncertainty, external disturbances, and sensor noise. In fact, most modern control problems are posed in terms of optimization via cost minimization. The perspective taken in this book is that machine learning provides a powerful new set of techniques to obtain high-performance control laws even for extremely complicated systems with non-convex cost functions.

1.3 Challenges of feedback control

Most textbooks start with simple feedback control problems. An airplane, for instance, may need to keep a certain ground speed. The airplane has a steady-state map (model) indicating the required thrust (actuation) under ambient flow condition and for an average airplane. Thus, the right thrust may be commanded in an open-loop manner based on the model, as illustrated in Fig. 1.2.

Yet, each airplane has its own steady-state map and an aging process (model uncertainty). Moreover, the wind (exogenous disturbance) may change the ground velocity. Model uncertainty and disturbances require a feedback element: The ground speed needs to be measured (tachometer) and the thrust needs to be adjusted. If the ground speed is too low (high), the thrust needs to be increased (decreased). The general feedback scheme is illustrated in Fig. 1.3.

Evidently, the control design is simple. There is a single state variable a (speed) which is sensed s (tachometer) and acted upon b (thrust) in a highly predictable manner and with negligible time delay. We refer to the excellent textbook of Åström & Murray [222] for the applicable control design.

The stabilization of steady solutions to the equations for laminar or transitional flows requires more refined methods.Navier-Stokes equations A sufficiently detailed discretized of the Navier-Stokes equation results in a system with a high-dimensional state, making it computationally expensive to design and implement controllers. In addition, time-scales may be very small in real-world fluid applications, such as flow over a wing or in a combustor, making controllers very sensitive to time delays; these time-delays may be due to sensor and actuator hardware or the computational overhead of enacting a control law. Sensor and actuator placement is also a challenge in high-dimensional fluid systems, with competing goals of decreasing time delays and increasing downstream prediction. Finally, many fluid systems are characterized by strongly non-normal linearized dynamics, meaning that the linearized Navier-Stokes equations have nearly parallel eigenvectors resulting in large transient growth of these modes in response to excitation [67, 262].

Despite inherent nonlinearity, stabilizing a steady state brings the system closer to the equilibrium solution where linearization is increasingly valid. Thus, the fluid dynamics literature contains a rich set of success stories based on linear control methods. Examples include the stabilization of the cylinder wake [226, 115, 218, 65], of the cavity flow [231], of the boundary layer [172, 11], and of the channel flow [29], just to name a few [43]. The linear quadratic regulator (LQR) and linear quadratic Gaussian (LQG) are among the most widely used methods for control based on computational fluid mechanics. The model-based control of experimental plants requires reduced-order models for computationally tractable on-line decisions. For details, we refer to excellent reviews on the applications of linear control theory in fluids mechanics [231, 161, 13, 250]. The associated reduced-order modeling efforts are summarized in these reviews and elaborated in [197, 138].

Optimization of turbulent flows tends to be much more complex. In addition to the challenges outlined above, the system is strongly nonlinear and is sufficiently far from a fixed point or limit cycle that linearization is not typically useful. The nonlinearity manifests in frequency crosstalk, where actuation at a given frequency may excite or suppress entirely different frequencies. Fully turbulent dynamics are typically chaotic and evolve on a high-dimensional attractor, with the dimension of the attractor generally increasing with the turbulence intensity. These are mathematical issues in turbulence control, but there are also more practical engineering issues. These include the cost of implementing a controller (i.e., actuator and sensor hardware, modifying existing designs, etc.), computational requirements to meet exceedingly short time scales imposed by fast dynamics and small length scales, and achieving the required control authority to meaningfully modify the flow.

As a motivating example, let us assume we want to minimize the aerodynamic drag of a car with, say, 32 blowing actuators, distributed over all four trailing edges and the same number of pressure sensors distributed over the car. A control logic for driving the actuators based on the sensor readings shall help to minimize the effective propolsion power required to overcome drag. This highlights the significant challenges associated with in-time control:

- High-dimensional state;
- Strong nonlinearity;
- Time delays.

A direct numerical simulation of a suitably discretized Navier-Stokes equation has not been performed for wind-tunnel conditions. Even a simplifying large eddy simulation requires at minimum tens of millions of grid points and still has a narrow low-frequency bandwidth for actuation. Secondly, the turbulent flow does not respond linearly to the actuation, so that there is no superposition principle for actuation effects. The changes to the flow caused by two actuators acting simultaneously is not given by the sum of the responses of the two actuators acting alone. Moreover, actuating at twice the actuation amplitude does not necessarily lead to twice the output. The trend may even be reversed. Thirdly, the effect of actuation is generally not measured immediately. It may take hundreds or thousands of characteristic time scales to arrive at the converged actuated state [21, 205]. We refer to our review article on closed-loop turbulence control [43] for in-depth coverage of employed methods.

1.4 Feedback turbulence control is a grand challenge problem

A high-dimensionsional state space and nonlinear dynamics do not necessarily imply unpredictable features. One liter of an ideal gas, for instance, contains $O(10^{24})$ molecules that move and collide according to Newton's laws. Elastic collisions signify strongly nonlinear dynamics, and indeed, the numerical simulation of Newton's laws at macro-scale based on molecular dynamics will remain intractable for decades to come. Yet, statistical averages are well described as an analytically computable maximum entropy state. This is the statistical foundation of classical thermodynamics. In contrast, the search for similar closures of turbulence has eluded any comparable success [198]. One reason is the ubiqituous Kolmogorov turbulence cascade. This cascade connects large-scale energy carrying anisotropic coherent structures with nearly isotropic small-scale dissipative structures over many orders of magnitudes in scale [106]. The multi-scale physics of turbulence has eluded all mathematical simplifications. Feynman has concluded that 'Turbulence is the most important unsolved problem of classical physics.' In other words: a grand challenge problem.

Turbulence control can be considered an even harder problem compared to finding statistical estimates of the unforced state. The control problem seeks to design a small $O(\varepsilon)$ actuation that brings about a large change in the flow. Many approaches would require a particularly accurate control-oriented closure. The necessary control mechanism might be pictured as a Maxwellian demon who changes the statistical properties of the system by clever actions. Control theory methods often focus on stabilization of equilibria or trajectories. Turbulence, however, is too far from any fixed point or meaningful trajectory for the applicability of linearized methods. In the words of Andrzej Banaszuk (1999):

> 'The control theory of turbulence still needs to be invented.'

1.5 Nature teaches us the control design

In the previous section, a generic control strategy for turbulence has been described as a grand challenge problem. Yet, an eagle can land on a rock performing impressive flight maneuvers without a PhD in fluid mechanics or control theory. Nature has found another way of control design: learning by trial and error.

It is next to impossible to *predict* the effect of a control policy in a system such as turbulence where we scarcely understand the unforced dynamics. However, it may be comparatively easy to *test* the effectiveness of a control policy in an experiment. It is then possible to evolve the control policy by systematic testing, exploiting good control policies and exploring alternative ones. Following these principles, Rechenberg [223] and Schwefel [242] have pioneered *evolutionary strategies* in design problems of fluid mechanics more than 50 years ago at TU Berlin, Germany.

In the last 5 decades, biologically inspired optimization methods have become increasingly powerful. Fleming & Purshouse [103] summarize:

> 'The evolutionary computing (EC) field has its origins in four landmark evolutionary approaches: evolutionary programming (EP) (Fogel, Owens, & Walsh, 1966), evolution strategies (ES) (Schwefel, 1965; Rechenberg, 1973), genetic algorithms (GA) (Holland, 1975), and genetic programming (GP) (Koza, 1992).'

EP, GA and GP can be considered regression techniques to find input–output maps that minimize a cost function. Control design can also be considered a regression task: find the mapping from sensor signals to actuation commands which optimizes the goal function. Not surprisingly, evolutionary computing is increasingly used for complex control tasks. For example, EP is used for programming robot missions [272]. GA are used to find optimal parameters of linear control laws [90, 23]. And since almost two decades, GP has been employed to optimize nonlinear control laws [91]. Arguably GP is one of the most powerful regression techniques as it leads to analytical control laws of almost arbitrary form. All evolutionary methods are part of the rapidly evolving field of *machine learning*. There are many other machine learning techniques to discover input–output maps, such as decision trees, support vector machines (SVM), and neural networks, to name only a few [279]. In fact, the first example of feedback turbulence control with machine learning methods has employed a neural network [171]. In the remainder of this book, we refer to *machine learning control* as a strategy using any of the aforementioned data-driven regression techniques to discover effective control laws.

1.6 Outline of the book

The outline of the book is as follows. Chapter 2 describes the method of machine learning control (MLC) in detail. In Chapter 3, linear control theory is presented to build intuition and describe the most common control framework. This theoretical foundation is not required to understand or implement MLC, but it does motivate the role of feedback and highlight the importance of dynamic estimation. In Chapter 4, MLC is benchmarked against known optimal control design of linear systems without and with noise. We show that MLC is capable of reproducing the optimal linear control but outperforms these methods even for weak nonlinearities. In Chapter 5 we illustrate MLC for a low-dimensional system with frequency crosstalk. A large class of fluid flows are described by such a system. We show that the linearized system is uncontrollable while MLC discovers the enabling nonlinearity for stabilization. In Chapter 6, we highlight promising results from MLC applied in real-world feedback turbulence control experiments. Chapter 7 provides a summary of best practices, tactics and strategies for implementing MLC in practice. Chapter 8 presents concluding remarks with an outlook of future developments of MLC.

1.7 Exercises

Exercise 1–1: Name two examples of feedback control systems in everyday life. Define the inputs and outputs of the system, the underlying system state and dynamics, and describe the objective function. Describe the uncertainties in the system and the types of noise and disturbances that are likely experienced.

Exercise 1–2: Consider the following plant model:

$$s = b.$$

(a) Design an open-loop controller $b = K(w_r)$ to track a reference value w_r.
(b) Now, imagine that the plant model is actually $s = 2b$. How much error is there in the open-loop controller from above if we command a value $w_r = 10$?
(c) Instead of open-loop control, implement the following closed-loop controller: $b = 10(w_r - s)$. What is the error in the closed-loop system for the same command $w_r = 10$?

Exercise 1–3: Choose a major industry, such as transportation, energy, healthcare, etc., and describe an opportunity that could be enabled by closed-loop control of a turbulent fluid. Estimate the rough order of magnitude impact this would have in terms of efficiency, cost, pollution, lives saved, etc. Now, hypothesize why these innovations are not commonplace in these industries?