Chapter 8
Future developments

"The most fruitful areas for the growth of the sciences were those which had been neglected as a no-man's land between the various established fields"

- Norbert Wiener, *Cybernetics: Or Control and Communication in the Animal and the Machine, 1948*

The future of machine learning control is bright, bolstered by concrete success stories and fueled by steady advances in hardware and methodology. As pointed out by N. Wiener, there is particularly fertile ground between two well-developed fields [275], and it is likely that there will be many exciting new technological developments between the fields of machine learning and control theory. This field will remain vibrant as long as

(1) it addresses problems that are interesting, important and challenging,
(2) improving hardware facilitates increasingly ambitious demonstrations, and
(3) investments in fundamental research yield new paradigms and mathematical frameworks.

Machine learning methods are beginning to enter the control of complex systems with numerous demonstrations in closed-loop turbulence control. These machine learning methods already pervade technology and information infrastructure, with smartphones providing compelling examples. Potential applications in engineering and industry may have an equally transformative impact. Examples include drag reduction of cars, trucks, trains, ships, airplanes and virtually any ground-, air- and water-born transport vehicle, lift increase of airfoils, aerodynamic force control mitigating wind gusts, internal fluid transport in ventilation systems and oil pipelines, efficient operation of wind turbines, greener and stable combustion and efficient chemical processes, just to name a few. Other applications concern more broadly any fluid flows or networks with well defined inputs (actuators) and outputs (sensors), like water and electricity supply in civil engineering, drug dosage and interactions for medical treatments, or financial trading.

In this chapter, we outline promising methodological advances that are poised to transform machine learning control (Sec. 8.1). Section 8.2 is dedicated to promising enablers for systems with high-dimensional inputs and outputs. We also highlight a number of exciting future applications that we believe will be heavily impacted by MLC, resulting in innovative solutions to scientific, engineering, and technological challenges of the 21st century [43] (Sec. 8.3). These directions are not exhaustive, but merely offer a glimpse into the exciting future developments of machine learning in the control of complex nonlinear systems. This chapter concludes with an interview of Professor Belinda Batten (Sec. 8.5). Prof. Batten is an internally renowned scholar in reduced-order modeling and control and pushes the frontiers of renewable wind and water energy.

8.1 Methodological advances of MLC

One of the many benefits of MLC using genetic programming is its simplicity and generality, requiring little or no knowledge of the system being controlled. However, there are a number of targeted improvements that may expand the power and applicability of genetic programming for MLC to a wider range of problems with little added complexity.

Exploiting closeness of controllers: A major goal of future efforts is to reduce the training time required for MLC using GP. Genetic programming is an extremely flexible framework for building complex input–output functions, although these representations are not unique, and it is often unclear whether or not two GP controllers are similar in function. Improving this notion of *closeness* of controllers is critical to avoid redundant testing and to reduce training time. In addition, having an induced metric on the space of GP controllers may result in effective controller reduction, so that complicated, high-performance controllers can be reduced to their simplest form. One exciting development in this direction is related to embedding the space of GP controllers in a simple metric space using hash inputs. The function of a GP controller can be uniquely identified based on its response to a random input sequence, likely sampled on the attractor. These hash functions then define an embedding which is much lower dimensional and Euclidean; moreover, if the sampling is truly random and incoherent with respect to the *structure* of the controllers, then it is likely that the embedding will preserve the controller geometry. This is closely related to compressed sensing theory (see below) and the Johnson-Lindenstrauss theorem [147].

Multi-dimensional scaling: The progress of an evolutionary algorithm may be visualized in a two-dimensional plane with a suitable metric of the control laws and multi-dimensional scaling. Thus, close (distant) control laws in the metric space are close (distant) in the visualization plane. Each control law is associated with a value of a cost function. Thus, the visualization plane may indicate the topology of the cost function, e.g. the number of populated minima, maxima or saddle points. An example of such a visualization of an ensemble of control laws has been provided in Section 7.4.1. Multi-dimensional scaling may also be used to illustrate MIMO control laws. These visualizations may provide valuable information for on-line decisions during a control experiment, e.g. tuning the levels of exploration versus exploitation.

Library building and dictionary learning: Another straightforward modification to GP for MLC is the implementation of library building and dictionary learning. Throughout the GP training process, a tremendous range of control laws are explored, and typically, only information about the final controller performance and the GP function representation are used in future iterations. Although it is necessary to obtain a statistically averaged measure of performance, or fitness, for each controller, it is likely that data from individual control experiments will provide rich information about attractor control. In particular, there are transient events in turbulence control, and it is likely that controllers are locally exploring favorable regions of phase space. Understanding controller individual histories at a higher time resolution may provide higher performance, multi-resolution analysis and control.

Multiple operating conditions: In addition, library building may result in adaptive GP for MLC for systems with multiple operating conditions. In cases where there are slowly varying parameters that change the operating conditions of a system, such as temperature, altitude, chemical composition of feed stock, etc., there may be multiple optimal control strategies. Library building and dictionary learning provide a robust strategy to handle these cases with multiple attractors. First, the system parameters are characterized by comparing against a library of previously identified attractors; sparse identification algorithms are particularly promising, and will be discussed below [278, 34, 107, 46, 42]. After the parameters are roughly characterized, the controller jumps to the best GP controller that fits the situation; we refer to this as fast feedforward control. Subsequently, additional slow feedback control can be included to add robustness to un-modeled dynamics and disturbances. However, it is likely that the inherent feedback built into GP controllers will be sufficient for many situations.

Learning robustness from control theory: Finally, there is an opportunity to bridge the gap between the excellent performance obtained using MLC and the rigorous theoretical foundations provided by classical control theory. Investigating issues related to robust performance of MLC will continue to be important to provide theoretical guarantees. These guarantees are particularly important in many aerospace applications, such as boundary layer control on a wing or combustion control in a jet engine, since they allow for controller certification. Although MLC already provides impressive robustness to varying environmental conditions in practice, it may be possible to augment robustness using techniques from classical control theory. It will also be interesting to see how the introduction of filters and dynamic estimators into the MLC framework will impact the control of real engineering-relevant systems. Lastly, there is great interest in the community for *insight* gained from effective MLC.

Genetic programming is a powerful and flexible regression technique, but machine learning has many more techniques that may be used for MLC. There are a number of more general methodological advances that will likely lead to new innovations in machine learning control in the future. The fields of machine learning, control theory, and data science in general, are rapidly growing and steady progress is being made on many fronts. Additionally, innovations in sparse sensing and uncertainty quantification are resulting in improved sensor and actuator placement, which will undoubtedly impact MLC strategies. Here, we provide a high-level overview of some of these exciting recent techniques.

Machine learning applications: The field of machine learning is developing at an incredible pace, and these advances are impacting nearly every branch of the physical, biological, and engineering sciences [190, 92, 30, 194, 168]. This progress is fueled by a combination of factors, including concrete success stories on real-world problems, significant investment of resources by large corporations and governments, and a grass-roots enthusiasm in the domain sciences. Surrounding this movement is the promise of a better future enabled by data science.

Machine learning of the plant: In a sense, aspects of machine learning have already been in use in classical engineering design and control for decades in the form of system identification [150, 174]. The goal of machine learning, like system identification, is to train a model on data so that the model generalizes to new test data, ideally capturing phenomena that have not yet been observed. Recent techniques in machine learning range from the extremely simple and general to highly sophisticated methods. Genetic programming has been used to identify nonlinear dynamical systems from data [32, 239, 219], and sparse variants exist [274]. The sparse identification of nonlinear dynamics (SINDY) algorithm [44] uses sparse regression [265] for nonlinear system identification.

Neural network based control: Machine learning has already been used in a number of control schemes for decades. In fluid dynamics, neural networks are used to identify accurate input–output maps to describe phenomena such as the growth of structures in a boundary layer and reduce skin-friction drag [171] as an add-on to opposition control [66]. Other examples where neural networks have been used to model and control turbulence include [193, 189, 95]. Neural networks constitute a set of bio-inspired algorithms to model input–output behavior using a coupled network of individual computational units, or “neurons". These algorithms were widely celebrated in the 1990s, but fell out of mainstream use, when it became clear that they were unable to solve many challenging problems. For example, although neural networks may be tuned to approximate nearly any input–output function, they are prone to overfitting and local minima in optimization. However, recently, the advent of *deep* convolutional neural networks (CNNs) has brought these algorithms back into the forefront of research [78]. In particular, deep neural networks are being constructed to identify complex phrases to describe scenery in unstructured images [69] and process natural language [132]. Perhaps more famously, these networks have resulted in the Google *deep dream*.

Genetic algorithm based control: Genetic algorithms [137, 76, 122] have also been widely used to identify parameters of both models and controllers for complex systems, such as fluids. Unlike genetic programming, which identifies both the structure and parameters of a model, genetic algorithms only identifies parameters of a model with a fixed structure. However, both genetic algorithms and genetic programming are evolutionary algorithms, relying on advancing generations of individuals and evaluating their fitness. Genetic algorithms have been successful in many applications where the structure is known.

Clustering: In addition to algorithms that identify input–output maps, there is also wide use of machine learning algorithms for classification or clustering data. For example, many complex systems are non-stationary and may be characterized by multiple attractors in different parameter regimes. In these systems, it is often not necessary or even helpful to estimate the full high-dimensional state of the system, but instead it may be sufficient to characterize the coarse system behavior. In this case, classification algorithms [279] (K-means, K-nearest neighbors, LDA, QDA, SVM [252, 241, 258], Decision Trees, Random Forests, etc.) are particularly useful. These classification algorithms are generally categorized into *supervised* algorithms, where the training data is labeled, and *unsupervised*, where the relevant clusters are unknown ahead of time. Classification is especially important in control, as each classification may lead to a control decision. For example, in ultrafast laser systems with multiple operating regimes, sparse classification algorithms have been used to rapidly identify the operating conditions, after which the controller jumps to a pre-determined near-optimal control setting and a slower adaptive control algorithm adds stability and disturbance rejection [107, 42]. In fluids, classification of operating regimes is also being explored [34, 46, 14]

Reinforcement learning: In robotics, the combination of machine learning and control has been proceeding steadily for at least a decade. Techniques for autonomous learning and reinforcement learning [257] have been widely adopted to control robots, both autonomous and tethered. Iterative learning control is used for tracking control, for example of robot arm position [35]. Similar algorithms are also used in brain-machine-interface problems, such as to train prosthetic devices.

The field of machine learning, and data science more generally, relies on a number of key steps: 1) data scrubbing, 2) feature extraction, mining, or engineering, and 3) model building. Each stage is crucial. The argument that machine learning will replace highly skilled domain scientists and engineers is somewhat facetious, considering that these algorithms do not work without a notion of what is *important* in a problem. For example, determining which labels should be used for training and engineering relevant features to distinguish various aspects of the data remain highly creative and critical human tasks in many disciplines.

8.2 System-reduction techniques for MLC — Coping with high-dimensional input and output

This textbook contains examples multi-input multi-output (MIMO) control with few inputs and outputs. The learning of control laws can be expected to take longer as the number of inputs and outputs increases. This is particularly true for high-dimensional actuation like surface motion or high-dimensional sensing like real-time particle image velocimetry, like in Section 6.1, or image-based feedback. In this section, we outline approaches coping with such high-dimensional input and output.

System reduction: Dynamical systems and control strategies have long benefited from dimensionality reduction, relying on the fact that even complex high-dimensional systems typically exhibit low-dimensional patterns that are relevant for models and control [125, 138]. More generally, it has been observed for decades that nearly all natural data signals, such as audio, images, videos, laboratory measurements, etc., are inherently low-dimensional in an appropriate coordinate system, such as Fourier or Wavelets. In this new coordinate system, many of the coordinates of the data will be negligibly small, so that the original signal can be accurately represented by a sparse vector, containing mostly zeros, in the new basis. This inherent sparsity of natural signals is the foundation for data compression, such as MP3 for audio and JPEG for images. PCA provides a tailored basis for optimal low-rank representation of data.

Compressed sensing: A recent mathematical breakthrough, called compressed sensing [50, 84, 18, 266, 52], has upended the traditional paradigm of collecting and analyzing data. Instead of painstakingly measuring every variable in a system, only to compress and discard the majority of the negligible information in a transform coordinate system, it is now possible to measure significantly less data at the onset and infer the relevant terms in the sparse transform basis.

In the past, solving for these relevant terms from subsampled measurement data would amount to a brute-force combinatorial search. This class of non-polynomial-time, or NP-hard, problem does not scale favorably with the size of the data, and Moore's law of exponentially growing computer power does not scale fast enough to help. A major breakthrough in compressed sensing is an alternative algorithm based on convex optimization using the sparsity-promoting ℓ_1 norm: $\|\mathbf{x}\|_1 = \sum_{k=1}^{N} |x_k|$. It has been shown that under certain reasonable conditions, solving for the vector $\mathbf{x}$ with the smallest ℓ_1 norm will approximate the vector with fewest nonzero entries. The measure of nonzero entries is often called the ℓ_0 "norm", although it does not satisfy the properties of a norm; it is also equivalent to the Hamming distance of the vector from the zero vector, as in information theory. This convex optimization does scale favorably with Moore's law, providing a clear path to solve increasingly large problems in the future.

Another innovation surrounding compressed sensing is the establishment of clear conditions on when the theory will work. First, there must be sufficient measurements to estimate the nonzero transform coordinates; this minimum number of measurements depends on the original size of the vector, the expected sparsity of the vector in the transform basis, and the level of noise on the signal. Next, the samples must be sufficiently *incoherent* with respect to the vectors that define the sparsifying basis; in other words, measurements should be as orthogonal, as possible, to all of the basis directions. It was shown that random projection measurements of the state (i.e., taking the inner product of the state vector **x** with a vector of identical independently distributed Gaussian elements) provide nearly optimal compressed measurements, regardless of the sparsifying basis. This is truly profound, as this enables a *universal* sparse sampling strategy. However, random projections of the state are not necessarily physical measurements, since they still require having access to all of the information in **x**. A more useful engineering measurement is often the spatially localized point measurement, corresponding to a single physical sensor. Fortunately, single point measurements are incoherent with respect to the Fourier transform basis, so that many signals may be reconstructed from few point measurements.

Regardless of our interest in full signal reconstruction, there are huge implications of sparsity promoting techniques and compressed sensing in engineering. Many of the concepts are applicable more generally to machine learning, modeling and control, including measurement incoherence, working on compressed subspaces of the data, and the structure and sparsity of data in general. For example, if a categorical or a control decision is desired, as opposed to a full-state reconstruction, it is often possible to use randomly sampled measurements and still achieve high classification performance. It is also useful to consider the effect of subsampling or projecting data onto a low-dimensional measurement space. The Johnson-Lindenstrauss theorem [147] and the restricted isometry property (RIP) [51] provide powerful quantification of the distortion of high-dimensional inner products after the data is compressed. These concepts are closely related to unitarity, and they may be used to embed high-dimensional data, or functions, as in the case of genetic programming, in a low-dimensional metric space. Hash function embedding is already being used in fluid flow control [183].

Sparse sensor and actuator placement: The placement of sensors and actuators is a critically important stage in control design, as it affects nearly every downstream control decision. However, determining the optimal placement of sensors and actuators is an NP-hard problem, which does not have an elegant solution, but rather involves a brute-force combinatorial search. In particular, this strategy does not scale well to even moderately large problems of interest. To compound the difficulty, often sensors may be quite expensive, as in the case of putting tracers in the ocean or human vaccinators in a foreign country. Even if sensors are inexpensive, processing the vast streams of data from a sensor network may be intractable, especially for mobile applications which are power and computationally limited.

Recent advances in compressed sensing (see above) are poised to revolutionize sensor and actuator placement for control problems. From an engineering perspective, the ability to solve NP-hard optimization problems with convex optimization routines is transformative. Fortunately, the sensor placement problem may be thought of as a compressed sensing problem under certain conditions, as we are trying to find the few best locations to maximize a well-defined objective function. Sparse sensor optimization using compressed sensing techniques has already been implemented for categorical decision making in complex systems [39, 14]. Extending these methods to optimize sensors and actuators for a control objective is an important area of current research. Finally, extending the compressed sensing framework to dynamic point measurements, such as Lagrangian tracer elements, has huge potential for the fields of oceanographic and atmospheric sampling.

8.3 Future applications of MLC

MLC can be expected to address many applications from everyday life to industrial production. Many control problems have a well-defined cost function, a finite number of sensor signals (outputs) which monitor the state and a finite number of actuation commands (inputs) which shall improve the system performance. In short, we have a multiple-input multiple-output (MIMO) plant and search for a control logic which optimizes a cost function.

Most factory processes fall in the category of a MIMO control problem. Åström and Murray [222] provide an excellent introduction in the world of control problems. The key strength of MLC comes into play when the plant behavior departs from linear dynamics and eludes current methods of model-based control.

Feedback control of turbulence is a grand challenge problem for control design as the plant incorporates three key difficulties: (1) high dimension, (2) nonlinearity, and (3) potentially large time-delays (see Chapter 1). Turbulence increases the power required to pump gas and oil in pipe-lines and is a major source of drag of ground, airborne, and maritime transport. Wind turbulence increases maintenance costs of wind turbines via gusts and creates dangerous situations for ground- and airborne transport. In contrast, turbulence has desirable mixing properties in heat exchangers, combustion and chemical processes.

Effective turbulence control can contribute to a key challenge of modern society: renewable energy production and reduced energy consumption. The world's primary energy supply has more than doubled from $6,106\,Mtoe$[1] in 1973 to $13,371\,Mtoe$ in 2012 [3]. The consumption in 2012 was $1,555 \times 10^{17}\,Joule$ corresponding to an average rate of $17.75\,TW$ or $2.5\,kW$ per person on this planet. This corresponds to one powerful heater per person operating day and night. On the downside, the environmental cost of energy production is an increasing challenge. Part of this cost is immediately felt as smog in the city or as noise near streets, railroads and airports. Other costs are indirect long-term consequences: Coal and fuel consumption provided 60.4 % of the 2012 energy supply [3] and are particularly problematic because the CO_2 emissions affect our climate. Our very existence is threatened by man-

[1] Mtoe means million ton of oil or equivalent (11.63 TWh).

made global warming. The Nobel committee has appreciated this fact by awarding the 2007 Peace Nobel Price to the International Panel on Climate Change and Al Gore "for their efforts to build up and disseminate greater knowledge about man-made climate change, and to lay the foundations for the measures that are needed to counteract such change" [1].

On the demand side, 20% of the worldwide energy consumption is used for transport, mostly by oil [82]. This demand corresponds to a cube of oil with each side measuring nearly 14 km. In 2014, the USA used 28% of its energy consumption for transport with 92% provided by fuel. The economic price of world transport amounts to 10% of the world gross domestic product (GDP) in 2011 [82].

The relevance of sustainable traffic may be measured by the fact that the European Research Council supports over 11,000 research grants on this topic. In following, we describe several transport-related opportunities for MLC:

Drag reduction of trucks: About one third of the operating costs of trucking companies is for fuel. The average profit margin is 3–4% in the USA. Thus, a 5% reduction on fuel costs, for instance by active flow control, spells a substantial increase of the profit. Passive means, like vanes for the rear side are already for sale. Active control solutions have already been tested on the road and have proven to save 5% fuel already under real-world traffic conditions. Experiments with small-scale vehicles demonstrate the possibility of a 20% drag reduction [211, 19] at a fraction of the energy cost.

Drag reduction of cars: Personal cars are less subject to economic considerations as in the trucking industry. However, the European union has identified the development of sustainable transport as a cornerstone challenge for the next decades. To encourage efforts in that directions, European union legislation [202] has set mandatory emission reduction targets for the car industry. The new standards impose a reduction of 30% of greenhouse gas emissions and corresponding fuel consumption for new cars by 2020. Meanwhile, the transportation industry is a strategic economic sector in Europe accounting for 4.5% of the total employment (10 million people). However, the market for new vehicles in European countries is declining, whereas that of emerging countries is rapidly growing. To remain competitive in this market, European car manufacturers must develop disruptive technology concepts to produce safer and more sustainable vehicles.

Improving the aerodynamic performance of road vehicles by flow control can help fullfill these requirements, in particular at highway speeds. At highway speeds, overcoming aerodynamic drag represents over 65% of the total power expense [140, 184]. This explains significant aerodynamic improvements on automobiles during the last decades. Most existing flow control approaches to reduce vehicle drag on commercial road vehicles are passive [118]. They rely on aerodynamic shaping, such as well-rounded contours (especially at the front), rear slope angle, and add-on devices [70]. Such passive approaches, however, are restricted by design and practical considerations and cannot be 'turned off' when not needed or modified to adapt to changing conditions.

Recently, significant research has been focused on active flow control (AFC) solutions. Cattafesta & Shelpack [54] give an extensive overview of possible ac-

tuation mechanisms, whereas [65] present the most common AFC approaches on bluff bodies. Of the many available approaches, a large subset is considered as an academic exercise, such as rotary [81, 27], streamwise [58], and transverse [53, 31] oscillation of a bluff body. Of the practical and realistic AFC mechanisms, many were investigated on the Ahmed body [4], such as synthetic jets [207], pulsed pneumatic actuation [149, 225], microjects with steady blowing [10], and Coanda blowing [104, 116]. AFC on heavy truck vehicles were also investiged, such as synthetic jet actuators [96], and Coanda blowing [97, 98].
A drag reduction of 23% has been achieved on a truck model with steady Coanda blowing [211]. The saved towing power was 3 times the invested actuation energy. A better actuation efficiency of 7 has been achieved with high-frequency Coanda blowing for a blunt-edged Ahmed body by [20] resulting in a 18% drag reduction. Both studies were based on a working open-loop control. The implementation of MLC is the logical next step after a working AFC. Machine learning control has recently been shown to improve existing control of an Ahmed body and a small-scale Citroën model in OpenLab PPRIME/PSA, France.

Safety of cars and trucks under wind gusts: Side-wind stability is important for passenger comfort and safety of all ground vehicles. It is particularly critical and a safety issue for vehicles with large projected side area, such as trucks and buses. Accidents linked to crosswind have been reported by several governmental agencies [48]. Side-winds and side-gusts originate from different sources such as weather, surrounding traffic, or the topology of the terrain next to the road. Several numerical [126, 145] and experimental [118, 59] studies were conducted to understand the dynamics of such flows. Both steady [139, 15] and unsteady [59, 267] crosswind investigations have been researched. The findings of these aforementioned studies and others can be summarized as follows:

- The driving stability decreases with increasing speed.
- A reduction of the lateral projected area in the back will reduce the rear side force and thereby the yaw moment. However, this relation is not valid for well-rounded rear-end configurations (C- and D-pillars).
- A-pillar radiusing has a large influence on the yaw moment.
- The center of pressure of the aerodynamic forces has a large impact on the vehicle stability.
- Directional stability depends more on yaw moment than on the overall side force.
- For a wide range of vehicles, the yaw moment increases approximately linearly with the yaw angle up to 20°.

Despite the numerous studies on side-wind effects and driving stability, very few have tackled the issue using AFC. To our knowledge, only Englar [97, 98] and Pfeiffer et al. [211, 212] applied Coanda blowing on the 4 rear edges of a truck to reduce drag and to control the yaw moment. Whereas Englar [97, 98] only implemented AFC as an open-loop, Pfeiffer et al. [211, 212] successfully applied it as a closed-loop control. The latter were able to achieve a leading 23%

drag reduction and a complete authority over the yaw moment. However, further improvements on the previous results can be achieved:

- Distributed actuation comprising *all* locations with demonstrated effect on drag and yaw;
- Distributed sensing providing real-time flow information;
- Employing very general, nonlinear and robust control laws comprising periodic, multi-frequency, and adaptive in-time closed-loop control;
- Reduction in actuation power for both drag minimization and yaw control.

All these new possible improvements have two themes in common, safety and efficiency. MLC can optimize open-loop, adaptive and closed-loop control by using harmonic functions as additional arguments of the control laws.

High-lift configuration of passenger aircraft: Civil aircraft are highly optimized for cruise conditions. Typical lift-to-drag ratios are around 17. This means 1 Newtons of thrust lifts 17 Newton of weight at a cruise speed around 850 km/h. During landing the main goal is not a good lift-to-drag ratio but a steep descent at low velocity to reduce the noise signature on ground. This is achieved with a high-lift configuration through a reeled-out flap. Active control can prevent early separation and reduce the size of these flaps for a given lift. Thus the weight of the aircraft is reduced during cruise resulting in lower propulsion power and thus reduced fuel consumption. This opportunity is being pursued by the two major passenger aircraft producers, Airbus and Boeing, and is the subject of intense research.

The economic aspect of reduced fuel consumption may be appreciated from the following data: Average profit margins are between 1–2 %. Fuel is with a 28% contribution the most important operating cost of passenger airlines (2014, US airlines). Fuel may represent around 40% of the maximum take-off weight. The airline not only saves fuel costs with active flow control. More importantly, it may replace 80 kg fuel by a paying passenger. In a 100 passenger airplane one passenger amounts to 1% of the income and thus accounts for a significant change in the profit margin.

Nacelle of a passenger aircraft: During take-off, the engines produce about 6 times the thrust during cruise and can lift around 30% of the passenger aircraft. The large nacelle prevents the recirculation bubble from extending to the compressor, which would result in a dangerous loss of thrust. The large nacelle is only needed for about one minute of maximum thrust during take-off. Active flow control at the inlet may reduce size the nacelle and thus reduce weight. This is another opportunity to save fuel.

Cavity of aircraft and trains: Wheel cavities of aircraft are a major noise source during landing. Similarly, cavities between train wagons produce noise. This noise can be significantly mitigated by feedback flow control [231]. The cavity noise reduction is a subject of intense research.

Drag reduction of ships: All moving ships create a wave pattern on the water. These waves require energy to be created and cause additional drag on the ship. For small ships, wave drag may be 80% of the total drag! One may conceive active vanes at the front of the ship to mitigate part of the wave drag. For full-scale tankers, the wave drag reduces to 20% of the total drag but is still substantial. An expected 5% reduction of drag is considered as a threshold value for engaging in a new expensive ship design. The remaining portion of the drag is due to skin-friction.

Skin friction reduction: About 80% of the propulsion power of tankers and ocean liners are due to skin friction drag. Future ships may profit from drag reduction due to hydrophobic surfaces. Similarly, skin friction is the major drag source of passenger airplanes. Riblets have been shown to reduce skin friction drag by up to 11% and are actively pursued by aircraft manufacturers. Skin friction may also be mitigated with distributed local suction and blowing. Drag reductions of 25% have been reported in low Reynolds number simulations [66]. At larger Reynolds numbers the reduction decreases. The distributed suction/blowing acts on plus unit scale to fight sweeps and ejections. An experimental realization would require myriad of actuators and sensors and will not be feasible for a long time. A more realistic active control is oscillatory wall motion. Numerical simulations [153] with high frequency oscillation show a 40% drag reduction. In experiments at RWTH Aachen, a 5% drag reduction was obtained. Feedback control with wall motion has hardly been explored.

We briefly mention other opportunities of feedback turbulence control and hence MLC.

Rockets: Solid fuel rockets may develop acoustic instabilities in the interior body which may destroy the rocket. Evidently, this is an opportunity for closed-loop control.

Wind turbines: Wind gusts create sudden loads on the rotor hub and bearings of a wind turbine. These forces reduce the mean-time between failure and thus make the maintenance effort more costly. Feedback controlled flaps at the trailing edge of the airfoil can reduce the unequal loading. The development of suitable hardware solutions and corresponding control logic is an active area of research.

Water turbines: Many underwater flows, such as in a river, may produce a more steady loading. Preliminary results of the authors show that accelerating and decelerating a cross-flow turbine to excite unsteady fluid forces can increase the energy efficiency by 80 % [255].

Gas turbines: Lean combustion reduces CO_X and NO_X emission. However, in the lean limit, the danger of blow out or combustion instabilities increases. Closed-loop mitigation of combustion instabilities via the change of fuel supply or fluid mechanic actuators are a demonstrated opportunity towards greener energy production [114].

Combustion engines: The internal combustion engine is ubiquitous in many modes of transportation, providing propulsion for most motorbikes, cars and trucks and convert chemical energy through combustion in mechanical energy via a periodic piston operation. Despite the periodic piston movement, the flow exhibits large cycle-to-cycle variations. Some cycles are good for combustion, particularly if the fuel is well mixed in the piston by turbulent mixing. Some cycles are less efficient. An ongoing effort of all engine producers is to reduce these cycle-to-cycle variations and to stabilize a uniformly good mixing. Fuel injection or mechanical actuation may be elements of a feedback stabilization for a better operation.

Pipe-lines: The power needed to pump oil through pipe-lines can be significantly reduced by the addition of polymers. One polymer molecule per 1 million oil molecules may reduce drag by 40%. Closed-loop turbulence control can hardly compete with this efficiency. However, there are still large opportunites for closed-loop control for gas-liquid-solid mixtures close to an oil plant.

Air conditioning: The goal of air conditioning is increased comfort. Ventilation shall refresh the air at appropriate temperature without unpleasantly felt air streams. MLC may learn an optimal air stream management based on sensor information.
Air conditioning causes other problems in hospitals. Airborne viruses need to be neutralized at active surfaces in the ducts. Such surfaces cause additional pressure losses and require thus a larger ventilation power. The mixing problem may be formulated as follows: a fluid particle with a potential virus must be sufficiently close to the active surface at least once with near certainty. However, multiple encounters with the active surface would lead to unnecessary pressure losses. This is an exciting mixing problem for closed-loop turbulence control.

Pharmaceutical and chemical processing: The industrial production of food, e.g. chocolate, may happen in vessels with diameters up to 5 meters. Different constituents need to be uniformly distributed in these vessels. The mixing may be closed-loop controlled by the inlets and mechanical mixing and monitored by cameras. Numerous production processes require a good mixing in vessels or in pipes. This is another exciting closed-loop mixing problem for which MLC is predestined. The effect of mixing is highly deterministic yet hardly predictable. Mixing has largely eluded an intuitive understanding.

The previous examples have focused on turbulence control. The application of MLC requires only a finite number of actuators and sensors. In addition, the cost function needs to be strongly related to the control law. In other words, the controlled plant exhibits statistically stationary behavior. Many complex systems fall in this category. Examples include

Predictions for solar and wind energy: The operation of an electricity network shall satisfy the unsteady demand and prevent overproduction as the storage capacities are limited. Hence, the production of renewable energy needs to be predicted from weather and other data to identify potential supply-and-demand problems early in time. This is an estimation problem which can also be solved with genetic programming [219]. Here, the input is the weather and other data, the output the energy production and the cost function the difference between estimated and actual production.

Short-term trading at stock markets: Here, the plant is the stock market, the input the trading action of one player, the output the supply and demand curve, the control logic the automated trading, and the objective function the daily profit. Long-term trading is far more challenging for automated learning as the assumption of statistical stationarity is not satisfied. The Dow Jones index, for instance, has the long-term tendency to increase. Companies may be founded, new economies may appear, or old economies may disappear, etc.

Dog training: All dogs, in fact all animals, are trained with rewards (and punishment). Intriguingly, maximum animal performance is not obtained with a monotonic performance-reward curve. There exist many education concepts how the reward should depend on animal performance. MLC might yield novel concepts.

These examples show, pars pro toto, the many MLC opportunities for MIMO plants which are of large relevance in industry or everyday life. As a disclaimer, we note that there also exist many problems which are more suitable for a model-based control strategy. For instance, the effect of control surfaces on the motion of an aircraft is well represented by a linear model for the usual operating envelope. Control design based on these models reliably cover many operating conditions. Thus, the robustness of control laws may be estimated. Continuing to establish a deeper connection between MLC and traditional control analysis will be essential to ensure broad adoption of these methods, as discussed in Chapter 4.

The flow around the airplane is turbulent but the time-averaged effect of myriad of vortices yields a nearly linear relationship between motion of the control surface and aerodynamic force. Taking the time scale of the vortices as reference, the model-based operation of the control surface can be considered as slow adaptive control. Mathematically, the situation might be compared to statistical thermodynamics where myriad of gas molecule collisions, i.e. strongly nonlinear events, can still lead to a linear relation between pressure and temperature in a finite volume by statistical averaging.

8.4 Exercises

Exercise 8–1: Investigate each of the following machine learning clustering algorithms and think about a representative data set that works well with the method and a representative data set that does not work with the method: K-means, linear discriminant analysis (LDA), support vector machines (SVM), decision trees.

Exercise 8–2: Consider the following signal, which is obtained as the sum of two sine waves:

$$f(t) = \sin(114\pi\,\mathrm{Hz}\,t) + \sin(1042\pi\,\mathrm{Hz}\,t).$$

The Shannon-Nyquist sampling theorem [200, 247] indicates that one must measure at 1042 samples per second to accurately reconstruct this signal. However, because the signal is *sparse* in the Fourier domain (i.e., only two Fourier modes are active at 57 Hz and 521 Hz), we may reconstruct this signal from dramatically under-sampled measurements.

In this exercise, create a time-resolved signal $\mathbf{f}$ by sampling at 1042 samples per second for 10 seconds. Now sample this signal randomly for 10 seconds with an average sampling rate of 128 samples per second. You may think of the sampled vector $\mathbf{f}_s$ as the output after applying a measurement matrix $\mathbf{C}$ to the full time-resolved signal:

$$\mathbf{f}_s = \mathbf{C}\mathbf{f},$$

where $\mathbf{f}$ is a time-resolved vector of the signal: $\mathbf{f} = \left[f(\Delta t)\ f(2\Delta t)\ \cdots\ f(N\Delta t)\right]^T$, and $\mathbf{C}$ contains random rows of the $N \times N$ identify matrix.
Finally, we may solve for the active Fourier modes in the compressed sensing problem:

$$\mathbf{f}_s = \mathbf{C}\boldsymbol{\Psi}\hat{\mathbf{f}},$$

where $\boldsymbol{\Psi}$ is the inverse discrete Fourier transform. Use a convex optimization routine, such as `cvx` in Matlab®, to solve for the sparsest vector $\hat{\mathbf{f}}$ that solves the underdetermined system above.

Exercise 8–3: Consider the LQR stabilization task from Exercise 4–1. Here, we will define a metric on the space of controller functions in an attempt to improve the convergence time of genetic programming control. First, create one hundred random states $\mathbf{a}$ where each component a_1 and a_2 are Gaussian distributed about $a_1 = a_2 = 0$ with unit variance. Now, for each controller, $b = \mathbf{K}(\mathbf{a})$, evaluate the control law at these same one hundred random states, and store the value in a 100×1 vector $\mathbf{b}_{\text{hash}}$. It is now possible to use these vectors to construct a proxy distance between two controllers $b = \mathbf{K}_1(\mathbf{a})$ and $b = \mathbf{K}_2(\mathbf{a})$:

$$d(\mathbf{K}_1, \mathbf{K}_2) \triangleq \|\mathbf{b}_{\text{hash},1} - \mathbf{b}_{\text{hash},2}\|_2. \tag{8.1}$$

Using this induced metric, modify your genetic programming search strategy to improve the convergence rate by reducing the number of redundant controller functions tested. For instance, when mutating, you might check if a new individual is sufficiently similar to individuals from a previous generation, and impose additional distance criteria to improve exploration and exploitation.

Exercise 8–4: Pick a future application of MLC from Sec. 8.3 and explore the current control approaches being applied tho this system. What are the challenges with applying MLC to this problem?

8.5 Interview with Professor Belinda Batten

Belinda Batten is Professor of Mechanical Engineering at Oregon State University, OR, USA. Professor Batten is the Director of the Northwest National Marine Renewable Energy Center (NNMREC), a collaboration between Oregon State University, the University of Washington, and the University of Alaska Fairbanks. NNMREC supports wave, tidal, offshore wind, and in-river energy harvesting through research and testing. The consortium was established by the U.S. Department of Energy to facilitate the development of marine renewable energy technologies via research, education, and outreach.

Professor Batten is internationally renowned for her research in modeling and control of distributed parameter systems, especially for her development of computational algorithms for reduced-order controllers. Her current research projects include mathematical modeling and control of autonomous vehicles and wave energy devices. She has raised significant funding to support her research from DARPA, DoD, NSF, and DOE.

Professor Batten has been a Program Manager for dynamics and control at the Air Force Office of Scientific Research, after which she was elected member of the Scientific Advisory Board for the U.S. Air Force. She was also a professor of mathematics at Virginia Tech, and served as Department Head for Mechanical Engineering from 2003-2007 and as Head of the School of Mechanical, Industrial and Manufacturing Engineering from 2007-2011 at OSU. Her research was honored by national and international awards, for instance by the prestigious Alexander von Humboldt fellowship.

Authors: You are a leader in the field of marine renewable energy and have developed numerous innovative technologies and control solutions. What were the main trends in marine renewable energy, especially related to fluid modeling and control, in the past decade?

Prof. Batten: I will focus my comments on wave energy, as that is where my primary expertise lies. Prior to this past decade, there has been a good amount of research on developing models of wave energy converters, and some work on controlling them. The modeling work typically leveraged the great volume of literature on hydrodynamics of ocean structures. The work on control has typically been model-based optimal control, and often has ignored the “control costs" that arise through actuation. That is, the results on the amount of energy produced through active control were often highly theoretical.

In the last ten years, more work has been done in experimental validation of computational models. Computational codes for simulating wave energy converters continue to be developed. Recently, codes are being developed for design of marine energy arrays, both wave and tidal current energy converters.
Some researchers are working on high fidelity fluid-structure interaction models for wave and tidal energy converters. These codes are not currently suitable for control design and computation as they require millions of state variables. As computational capabilities continue to progress, these codes may be more useful in the control regime, but at this point, model reduction of some type is required to use them for purposes other than simulation, e.g., for control or optimization.

Authors: You were an early advocate of merging control theory with fluid mechanics. Do you see any key similarities or differences in the current efforts to bring machine learning into the fluid dynamics community?

Prof. Batten: One of the challenges in merging control theory with fluid dynamics in the early years was learning each others' language. Basic terms like "control" were used differently by control theorists and by fluid dynamicists. I remember giving a kick-off talk at a meeting of the two groups, laying the foundation for the discussion – including the typical nomenclature for mathematically describing a fluid system with a control. One of the fluid dynamicists responded, "you mean my entire life's work in actuator modeling is reduced to a B matrix?" So, part of getting the two groups to work together was facilitating a common understanding of what problems were viewed as easy, hard, unanswered, unanswerable, tractable, etc. Once these communities began to interact with each other, great progress was made, and a lot was learned.
I hope that bringing machine learning into the community is a little easier because some of the same people in the controls and fluid dynamics communities are involved, and collaborative work across disciplines starts with relationships. That said, I remember years ago when a seminar speaker talked about using neural networks to develop a model for the Navier-Stokes equation; several in the audience muttered about why anyone would throw out the physics.
And that illustrates the crux of the issue. There are some really hard problems for which developing high fidelity physics based models for control is not feasible—at least not with today's computational limitations. The fluid-structure interaction modeling of a wave energy converter in the ocean is one of these problems. So, to develop machine learning approaches that we can test and verify on smaller tractable problems can lead to solutions and insights on the more complex ones for which developing physical models for control is infeasible.

Authors: Could you comment on the societal impact of marine renewable energy, and how you see fluid mechanics and control helping to address the key challenges?

Prof. Batten: The reality is that the world cannot continue to rely on fossil fuels to meet energy needs. The supply of fossil fuels will eventually be exhausted, and in the meantime, the impact of fossil fuels on our environment is increasingly destructive. While people may debate about climate change, the acidification of our oceans is measureable, and the consequences of delaying the move to renewable energy are sobering.

In my opinion, replacing fossil fuels will require a portfolio of renewables that address particular needs of place. Marine renewables could be an important addition to this portfolio; this energy source is highly predictable, always present, and located within 50 miles of 50 percent the world's human population. Tides are predictable 100 years in the future; waves are predictable 84 hours out. Those horizons make it simple for a utility to manage electricity coming onto the grid, and one can control marine energy converter arrays to deliver what is needed. Unlike wind and solar, these energy sources are always present.

The key challenge for marine renewable energy is becoming cost competitive with other renewable energy sources. Control of marine energy converters has typically been directed toward maximizing energy extraction. Maximizing energy produced indeed contributes to a lower cost. However focusing solely on that objective ignores the reduction in life expectancy of a wave energy converter or increase in operations and maintenance costs due to fatigue or failure of converters under extreme events. It may be that for certain converters, control should also be used to minimize fatigue. At Oregon State University, my colleague Ted Brekken has been working with his students to develop life extending control to address this concept.

If computationally tractable models of fluid structure interactions of wave energy converters—especially under extreme wave conditions—could be developed, such models could be used to better predict the reliability and survivability of devices. Alternatively this topic could be fertile research ground for machine learning. Better understanding reliability and survivability of devices could be leveraged to understand how to lower the cost of energy of the wave energy converters, thus contributing to adoption of marine energy.

Authors: In other fields, such as astronomy and particle physics, there are well-funded grand-challenge problems that the community tackles in a coordinated effort. Is there a need for similar larger scale collaborative efforts in fluid flow control? If so, what do we need to do in order to get fluid dynamics to a similar position in terms of funding and support?

Prof. Batten: I think great progress has been made in fluid flow control since the early 2000s when I was challenging the communities to come together and work collaboratively on these problems. A lot has been learned about what control can do and what kinds of models and computational codes are useful. I think the big challenge problem for fluid dynamics related to marine energy is the fluid-structure interaction models, suitable for control design and optimization. A workshop that pulls community experts together to define other such problems might be an important next step.

Authors: As the director of the Northwest National Marine Renewable Energy Center (NNMREC), could you comment on some of the grand challenges in the marine renewable energy sector, and how they may be impacted by machine learning methods in the coming years?

Prof. Batten: If there is a fundamental grand challenge in marine renewable energy, it is how to lower the cost of energy of these technologies to make them competitive. While there are many subsystems within a marine energy array with highly technical underpinnings and each contributes to the overall cost, it is necessary to take a system level viewpoint of this problem. Applying machine learning to a system level performance of a marine energy array could have interesting outcomes.

Authors: This chapter considers future developments and extensions to machine learning control. If you were giving advice to a new student, what direction would you point them in for important research problems, and what tools would you want them to be equipped with?

Prof. Batten: I'll leave the advice about new research directions in machine learning to the experts in that area. Regarding the tools that I want students to know, I believe that it is important for students interested in controls, be it learning based methods or model based methods, to understand the value and limitations of the various approaches. When it comes down to it, any control approach is a tool, and it's important to know when to use a screw driver and when to use a hammer. There are times when you could use either, but one is probably better for the job than the other.

Authors: We look forward to your continued breakthroughs in control and renewable energy, and we thank you for this interview!