Future Controls Courses

As discussed in [1] and [2], the key future application domains for control systems are very broad, and the systems are likely to be very complex. In addition, for control research to have a large impact on those problems, there will likely be less emphasis placed on mathematical theory applied to abstract models and, instead, more emphasis on dealing with the realities (nonlinearity, noise, and modeling errors/uncertainty) of the systems under control. The system complexity could also require a larger focus on the design to meet performance goals, using techniques such as online optimization, rather than proving stability.

The complexity of the systems and the decision-making processes involved suggest the need for a fresh look at the future of controls curriculum, although it is not discussed much in [1]. Future interest is on large-scale systems, for which models might not be available and/or easily constructed. Furthermore, many of the systems of interest involve connected, distributed subsystems with numerous subcomponents, such as the network in communication systems; the sensors in perception-based systems; and complex, possibly poorly known, time-varying dynamics. This educational approach should also provide a more broadly applicable perspective of control systems as being a sequential decision-making process that involves feedback and decision making under uncertainty, with objectives that focus on performance, stability, resilience, as well as cost.

These goals differ significantly from the traditional focus on graduate-level controls that dwell on linear state-space models, with lengthy discussions of observability, controllability, zeroes, and model representations; control design and weight selection for linear quadratic regulator (LQR) control; estimator design for systems impacted by Gaussian noise; and optimal control design using linear-quadratic Gaussian (LQG) control. There is a lot of useful mathematical training in learning and interpreting these results, but the emphasis in the coursework likely leaves the impression that the results are more broadly applicable than they are in real life. Real-world control problems are complex and "messy," and as they are predicated on the assumption of an accurate linear model, these techniques typically cannot handle that complexity. Although nonlinear and/or adaptive control courses are typically offered beyond these traditional state-space courses, the solutions usually do not scale up to the problem complexity of interest.

To address these future research needs, controls courses should focus more on online numerical solutions of optimal control techniques. Techniques of interest include online optimization-based approaches such as differential dynamic programming [3], iterative LQR [4], and iterative LQG (iLQG) [5], which provide approximate solutions using online iterative reoptimization. These approaches typically consist of three steps: compute derivatives along a given nominal trajectory, a backward pass to update the value function estimates, and then a forward pass to compute a new estimate of the best state trajectory. These techniques require a significant amount of computation (iLQG is comparable to a Gauss–Newton Hessian approximation), but those calculation are tractable for many systems of interest, given the central processing units and graphics processing units currently available.

While strong statements can be made about these algorithms for linear systems with quadratic costs, they can also be applied to more complex cases (nonlinear dynamics and nonquadratic costs), although convergence guarantees are lost. Educational objectives here would be to ensure that the students are aware of what is lost in these more general cases and the significance of what is lost. It is of interest to note that similar material was covered in an optimal control course I took in the late 1980s based on [6], but, given the limited computational resources available at that time, the approach seemed unrealistic for the mechanical/aerospace systems of interest. Then, over time, that material was replaced with the mu analysis and synthesis that formed the main components of the advanced control material covered in the 1990s [7]. However, it is interesting to note that, at least at the Massachusetts Institute of Technology, neither components is currently covered in depth in the course material.

Further topics of interest include optimization-based control using model predictive control (MPC) [8]–[11]. For example, [5] discusses how MPC retains the benefits of optimal control but avoids the curse of dimensionality of dynamic programming, handles constraints very easily, and although model based, can be easily integrated with a model-learning algorithm. While commonplace in the chemical process industry, with the substantial recent increase in computation power available for embedded systems, MPC has been more widely adopted for the online planning of robotic systems. Educational goals would be to point out the issues of picking a

Digital Object Identifier 10.1109/MCS.2017.2786418
Date of publication: 16 March 2018

good cost function for the optimal control, understanding modeling/abstraction approaches, creating good cost to go, deciding on the plan horizon length and time step (Ts), and computation and parallel implementations.

With models available, these MPC techniques integrate the data available through feedback of the current state. However, many of the problems of interest have nonlinear dynamics that are difficult to model accurately. Thus, an interesting educational challenge is to provide the students with the skills necessary to determine how best to utilize the data beyond just using it for feedback, such as for learning models and/or control policies.

A recently developed course at the University of California, Berkeley, by S. Levine [12] is an exciting exemplar of a new class designed to achieve these objectives. It provides a good mix of model-based control and direct data-to-decisions approaches. Reinforcement learning (RL) is discussed for model and policy learning. The benefits of deep RL techniques are also discussed (avoids the need for manual crafting—the features in the representation are self-learned from the data), which are particularly useful for perception-based feedback loops. The course touches on optimization-based control but has few details of these techniques as a control system. The ideal course would provide a more balanced perspective on the implications of closing the loop on the system. This could be accomplished by drawing on the known analysis of the model-based online optimization for linear quadratic problems and then highlighting the limits of that analysis for more complex systems, as would typically occur in a controls class.

Controls courses must provide a solid foundation, but they should also provide students with the tools necessary to work with models if they exist, create (if possible) models in real time from the available data, and/or can yield good control policies if modeling not possible. It is likely that a deep fusion of model- and data-based control techniques will be needed to solve the complex problems envisaged for the future, and the controls curriculum must evolve and expand to reflect that change. This change will likely require a tighter integration of the material covered in optimization-based control and reinforcement/deep learning classes, with a combined teaching effort between the two communities.

REFERENCES

[1] F. Lamnabhi-Lagarrigue, A. Annaswamy, S. Engell, A. Isaksson, P. Khargonekar, R. M. Murray, H. Nijmeijer, T. Samad, D. Tilbury, and P. Van den Hof, "Systems & control for the future of humanity, research agenda: Current and future roles, impact and grand challenges," *Annu. Rev. Control*, vol. 43, no. 1, pp. 1–64, 2017.

[2] J. P. How, "Future directions in control [from the editor]," *IEEE Control Syst.*, vol. 37, no. 5, pp. 3–4, Oct. 2017.

[3] D. H. Jacobson, "Second-order and second-variation methods for determining optimal control: A comparative study using differential dynamic programming," *Int. J. Control*, vol. 7, no. 2, pp. 175–196, 1968.

[4] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proc. American Control Conf.*, 2005, pp. 300–306.

[5] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2012, pp. 4906–4913.

[6] P. Dyer and S. R. McReynolds, *The Computation and Theory of Optimal Control*. New York: Academic, 1970.

[7] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2nd ed. New York: Wiley, 2005.

[8] J. A. Rossiter, *Model-Based Predictive Control: A Practical Approach*. Boca Raton, FL: CRC, 2003.

[9] E. F. Camacho and C. B. Alba, *Model Predictive Control*. New York: Springer, 2007.

[10] J. M. Maciejowski, *Predictive Control: with Constraints*. Englewood Cliffs, NJ: Prentice Hall, 2002.

[11] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2017.

[12] S. Levine. (2017). CS 294: Deep reinforcement learning. [Online]. Available: http://rll.berkeley.edu/deeprlcourse/

Jonathan P. How

Intelligence to Information

Reading the work of Ralph Hartley, Shannon said, was "an important influence on my life." Not simply on his research or his studies: Shannon spent much of his life working with the conceptual tools that Hartley built, and for the better part of his life, much of his public identity—"Claude Shannon, Father of Information Theory" —was bound up in having been the one who extended Hartley's ideas far beyond what Hartley, or anyone, could have imagined. Aside from George Boole, that obscure logician, no one shaped Shannon's thought more. In the 1939 letter in which Shannon first laid out the study of communications that he would complete nine years later, he used Nyquist's "intelligence." By the time the work was finished, he used Hartley's crisper term: "information."

—Jimmy Soni and Rob Goodman, *A Mind at Play: How Claude Shannon Invented the Information Age*, Simon and Schuster, 2017, p. 130.