

Data-Driven Control 1/2

A SPECIAL ISSUE SAMPLING FROM A VAST & DYNAMIC LANDSCAPE

Florian Dörfler

Data-centric and learning-based methods have pervaded all areas of science, engineering, technology, and society at large, including the field of control systems. Whatever is your take on these developments, they cannot be ignored. Recently, a gathering of around a hundred researchers from diverse backgrounds within the IEEE Control Systems Society came together to collaboratively brainstorm about a scientific roadmap for the future of our discipline. I quote from Section 4 of the *Control for Societal-Scale Challenges Roadmap 2030* report [1]:

“One of the major developments in control over the past decade — and one of the most important moving forward — is the interaction of machine learning and control systems.”

The topic of this double special issue, aptly named “data-driven control”, requires little further motivation. It encompasses a vast, diverse, and dynamic intellectual landscape, loosely identified with control engineering based on experimental data. It has become a prominent theme in the canvas of science, with numerous researchers contributing their own brushstrokes. In this editorial, I aim to frame this canvas, sketch some of the big ideas with broad strokes, and draw lines between different approaches.

Looking back: a blurry history

It is important to acknowledge that this canvas is not blank. Generations of researchers have been diligently working on adaptive control, system identification, and related learning-based topics; see [2]–[6] for excellent retrospectives. Their efforts have yielded not only scientific breakthroughs but also technological success stories, such as iterative learning control, auto-tuned regulators, adaptive predictive control approaches, and many others [7].

Whatever list of milestone results I provide will be judged by what I left out. Hence, an *incomplete* list of early research paradigms includes data-driven tuning (via iterative feedback [8], correlation [9], or virtual reference [10]); approaches blending identification and control (such as dual control [11], optimal controller identification [12], or identification for control [13]); data-driven optimization in feedback (e.g., using extremum seek-

ing [14] or real-time optimization [15]); data-driven takes on dynamic programming [16]–[18] nowadays broadly labeled as reinforcement learning; and many more. All of these sit square between model-free and model-based control, but the lines are blurred. I deliberately stopped this list about 15 years ago when real-time computation methods in the vein of predictive control and reinforcement learning led to further branching of the literature and many novel paradigms – some of which are contextualized in the articles of this double special issue.

Over the years, control engineers have embraced numerous concepts from the broad field of artificial intelligence (AI). Vice versa, many pivotal machine learning breakthroughs and entire subdisciplines are deeply rooted in our field, such as back-propagation [19] or reinforcement learning [20] – the authors of which are nowadays household names in AI. In these and other problem settings, both fields have repeatedly cross-fertilized another, and this process is accelerating in the internet age with readily available tutorial videos and actionable code. These success stories – and I apologize for not providing an exhaustive list – showcase the mutual influence between the two domains.

Delineating present versus past: deluge of data, computing power, and new problems & methods

Many of the aforementioned ideas came and went, some stuck, and some faded. Research often goes in circles, and ideas go through hype cycles – actually, repeatedly in the case of learning-based methods. However, at the time of writing, it appears that data-driven control has come to stay. To fully comprehend its current state and distinguish it from past hype cycles, it is crucial to contextualize data-driven control within the present time. This includes not only the widespread excitement surrounding big data and AI, but also the emergence of new applications and problem scenarios that are driving control research. For instance, the deployment of automation in unstructured environments or the utilization of complex sensing modalities, as in autonomous driving, are shaping the evolution of data-driven control.

Furthermore, recent technological advances have facilitated

the emergence of contemporary approaches that would have been inconceivable just a decade ago. We now have unprecedented access to vast amounts of data, owing to the widespread deployment of sensing and communication technology. Additionally, the availability of powerful computing resources has opened up new possibilities. Many approaches that were previously labeled as *brute force* or deemed impossible to implement in real-time are now efficiently and reliably implementable. Though, brute force does not reign yet in data-scarce applications or in online settings with few samples or severe real-time requirements.

Similarly, there have been significant advancements on the methodological front. Theory and deployment of optimization algorithms such as auto-differentiation, uncertainty quantification techniques such as distributional robustness, non-parametric regression methods based on reproducing kernels, advances in non-asymptotic and high-dimensional statistics, and many other innovations have brought a flurry of new ideas into our field and adjacent communities. These advancements have significantly shaped the landscape of data-driven control, enabling researchers to explore novel approaches and leverage the abundance of data and computational power available in the present era.

Sometimes, research paradigms from previous eras are also revisited. An example that I am intimately familiar with is the behavioral approach to systems theory [21] which provides a representation-free description of a system as a collection of trajectories. After receiving little attention for a decade, it is now gaining traction, as it is naturally suited for data driven control [22], [23]. Another approach that has been ahead of its time is regret-based analysis of learning-based optimal linear quadratic control [24], [25]. In both cases, today's researchers adopted a fresh perspective on these paradigms, leveraging our ever-advancing technology, contemporary methods for uncertainty quantification and robust design, demands from timely applications, and advanced performance specifications, e.g., safety, distributional uncertainty, finite sample complexity, and so on.

All articles in this double special issue exemplify the integration of timeless paradigms from our field with modern approaches and computational methods. For instance, the "model" in model predictive control does not need to be the familiar state-space representation. It can be a Gaussian process, a neural network, or a simple data matrix dating back to the heydays of subspace system identification or dynamic matrix control. Could this have been done decade(s) ago? Conceptually yes, but we now have advanced analysis methods, a broad interest in data-centric methods, practical business cases, sufficient amounts of data, and the (real-time) computing power to make it feasible.

Finally, I witness a cultural sea change within the academic control community which has traditionally favored theory over engineering solutions. The AI community has convinced many of us that there is an empirical side to our field that deserves more attention. And there is still theory to be developed: especially, via methods that were treated step-motherly so far, such as statistics.

Mapping the vast, diverse, & dynamic landscape

As the reader can discern by now, there are numerous exciting developments in the field of data-driven control, and it would be unfair to single out individual success stories while neglecting others. The field is ever growing and becoming more diverse. In the past, specialized workshops on system identification and adaptive control were the main venues for discussions on data-driven control. However, it has now become a mainstream topic at the largest control conferences, it has spread to events related to machine learning or applications, and the successful *Learning for Dynamics & Control Conference*¹ was inaugurated in 2019.

Data-driven control is a vast intellectual landscape to project our ideas upon, and the developments are fast paced. There are, however, few canonical problem settings, and the various research approaches are highly fragmented. To map this landscape, it is useful to categorize different approaches via *binary classifiers* (adopting AI terminology), such as online vs. offline data collection, batch vs. iterative implementations, certainty-equivalence vs. robust formulations, or parametric vs. non-parametric methods. The most significant divide is between direct vs. indirect methods, which refers to whether data is used directly for decision-making or to model the data generating process, which informs decision-making at a later stage. This divide is canonical in adaptive control [26], [27] and has been recognized across domains, as discussed in [28]. As technology and specifications become more complex, researchers often favor of end-to-end (direct) methods, whereas practitioners often prefer modular (indirect) solutions. This debate is raging in fields embracing AI and is unsettled even in the AI community [29].

Certainly, the classifiers mentioned above are not unambiguously defined, and the lines between the different approaches are blurred. I will not attempt to clear the fog, but I note that the articles in this double special issue broadly sample from this landscape and collectively form a diverse and rich collection that spans a substantial portion of the map in the field of data-driven control and its applications in different domains. The sidebar "The Articles of this Special Issue" provides an outline of the contents covered in this issue, which you are currently reading.

Synopsis: control no longer a hidden technology

What does the future hold for data-driven control? I leave it to the authors of the articles in this special issue to communicate their respective visions. Returning to my opening quote, only time will reveal the true significance of data-driven control as a long-term development within our field. However, in my opinion, data and learning have already had a significant impact on the scientific and public discourse, and technology will continue to drive this revolution in the foreseeable future. Algorithms, which are at the core of data-driven control, are no longer a *hidden technology* as once famously voiced by a towering figure in our field [30], but rather in the spotlight and gaining increasing attention. The continued advancements in technology, combined

¹See <https://14dc.seas.upenn.edu> for the latest edition.




FIGURE 1 Artist's illustration of data-driven control with the control algorithm taking a position in the spotlight rather than being merely a *hidden technology* – a famous term coined by K.J. Åström for the first time in [30] – serving an end product. Figure by Sonia Monti.

THE ARTICLES OF THIS SPECIAL ISSUE

The seven articles in this double special issue sample from the vast & dynamic scientific landscape of data-driven control and its applications in different domains. Collectively, they cover some of the main lines of thought. In what follows, I briefly summarize the four articles in the present special issue.

The article *“Kernel methods and Gaussian processes for system identification and control”* by Carè and co-workers surveys contemporary methods for system identification and learning-based control. The contents spans linear and nonlinear system identification in reproducing kernel spaces with a nice tutorial exposition of regularization methods and uncertainty quantification. Beyond system identification, the article also covers applications of the methodology to different learning-based control paradigms and experimental implementations.

The article *“Data-Driven Safety Filters”* by Wabersich and co-workers addresses the topic of safety in control systems. The authors introduce the concept of an ideal safety filter to enhance a controller with safety guarantees and present tutorials for three classes of safety filters: Hamilton-Jacobi reachability, control barrier functions, and predictive control techniques. The article covers applications to first-principle and data-driven mod-

els, a selection of model-based and learning-based controllers, and several illustrative case studies from the robotics domain.

The article *“Data-driven control based on the behavioral approach: From theory to applications in power systems”* by Markovsky and co-workers presents the behavioral approach to data-driven control. The cornerstone of this approach is that time series data, suitably assembled in a matrix, spans the set of all finite-length trajectories of an LTI system. In the deterministic LTI case, this result lends itself directly for optimal control. When deviating from this idealized setting, the authors robustify their methods by means of regularizations and apply them for control of power-electronics-dominated power systems.

The article *“Quasi-Stochastic Approximation—with Applications to Extremum Seeking Control”* by Lauand and Meyn surveys the theory of (quasi)stochastic approximation and discusses connections to extremum seeking control. These methods are concerned with a root finding problem based on random (respectively, deterministic) probing of the function of interest. The prime example and historical root is gradient-free optimization which relates to adaptive control and reinforcement learning. The paper gives a tutorial of the *ODE method* connecting discrete stochastic algorithms and their associated mean flows.

with the growing availability of data and the development of new learning approaches, are likely to shape the landscape of control systems research and applications in the years to come.

ACKNOWLEDGMENTS

I want to thank several colleagues that read drafts of this article and provided valuable feedback, opinions, and perspectives on the data-driven control landscape – among others, all authors of this double special issue, several of my colleagues at ETH Zürich and the NCCR Automation, as well as Anuradha Annaswamy, Manfred Morari, Alessandro Chiuso, Håkan Hjalmarsson, Travis Gibson, Ian Manchester, Andrea Iannelli, and Claudio De Persis.

AUTHOR INFORMATION

F. Dörfler (dorfler@ethz.ch) F. Dörfler is with the Automatic Control Laboratory, ETH Zürich, 8092 Zürich, Switzerland.

REFERENCES

- [1] “Control for societal-scale challenges roadmap 2030,” A. M. Annaswamy, K. H. Johansson, and G. J. Pappas, Eds., 2023. [Online]. Available: <https://ieeecs.org/control-societal-scale-challenges-roadmap-2030>
- [2] A. M. Annaswamy and A. L. Fradkov, “A historical perspective of adaptive control and learning,” *Annual Reviews in Control*, vol. 52, pp. 18–41, 2021.
- [3] M. Gevers, “A personal view of the development of system identification: A 30-year journey through an exciting field,” *IEEE Control Systems Magazine*, vol. 26, no. 6, pp. 93–105, 2006.
- [4] L. Ljung, “Perspectives on system identification,” *Annual Reviews in Control*, vol. 34, no. 1, pp. 1–12, 2010.
- [5] A. Fradkov and B. T. Polyak, “Adaptive and robust control in the ussr,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 1373–1378, 2020.
- [6] Z.-S. Hou and Z. Wang, “From model-based control to data-driven control: Survey, classification and perspective,” *Information Sciences*, vol. 235, pp. 3–35, 2013.
- [7] S. Yin, X. Li, H. Gao, and O. Kaynak, “Data-based techniques focused on modern industry: An overview,” *IEEE Transactions on Industrial Electronics*, vol. 62, no. 1, pp. 657–667, 2014.
- [8] H. Hjalmarsson, M. Gevers, S. Gunnarsson, and O. Lequin, “Iterative feedback tuning: theory and applications,” *IEEE Control Systems Magazine*, vol. 18, no. 4, pp. 26–41, 1998.
- [9] A. Karimi, L. Mišković, and D. Bonvin, “Iterative correlation-based controller tuning,” *International journal of adaptive control and signal processing*, vol. 18, no. 8, pp. 645–664, 2004.
- [10] M. C. Campi, A. Lecchini, and S. M. Savaresi, “Virtual reference feedback tuning: a direct method for the design of feedback controllers,” *Automatica*, vol. 38, no. 8, pp. 1337–1346, 2002.
- [11] B. Wittenmark, “Adaptive dual control methods: An overview,” *Adaptive Systems in Control and Signal Processing 1995*, pp. 67–72, 1995.
- [12] L. Camestrini, D. Eckhard, A. S. Bazanella, and M. Gevers, “Data-driven model reference control design by prediction error identification,” *Journal of the Franklin Institute*, vol. 354, no. 6, pp. 2628–2647, 2017.
- [13] H. Hjalmarsson, “From experiment design to closed-loop control,” *Automatica*, vol. 41, no. 3, pp. 393–438, 2005.
- [14] K. B. Ariyur and M. Krstic, *Real-time optimization by extremum-seeking control*. John Wiley & Sons, 2003.
- [15] B. Chachuat, B. Srinivasan, and D. Bonvin, “Adaptation strategies for real-time optimization,” *Computers & Chemical Engineering*, vol. 33, no. 10, pp. 1557–1567, 2009.
- [16] S. J. Bradtke, B. E. Ydstie, and A. G. Barto, “Adaptive linear quadratic control using policy iteration,” in *Proceedings of 1994 American Control Conference-ACC’94*, vol. 3. IEEE, 1994, pp. 3475–3479.
- [17] F. L. Lewis and D. Vrabie, “Reinforcement learning and adaptive dynamic programming for feedback control,” *IEEE Circuits and Systems Magazine*, vol. 9, no. 3, pp. 32–50, 2009.
- [18] D. Bertsekas and J. N. Tsitsiklis, *Neuro-dynamic programming*. Athena Scientific, 1996.
- [19] Y. LeCun, “A theoretical framework for back-propagation,” in *Proceedings of the 1988 connectionist models summer school*, vol. 1, 1988, pp. 21–28.
- [20] R. S. Sutton, A. G. Barto, and R. J. Williams, “Reinforcement learning is direct adaptive optimal control,” *IEEE Control Systems Magazine*, vol. 12, no. 2, pp. 19–22, 1992.
- [21] J. C. Willems, “The behavioral approach to open and interconnected systems,” *IEEE Control Systems Magazine*, vol. 27, no. 6, pp. 46–99, 2007.
- [22] I. Markovskiy and F. Dörfler, “Behavioral systems theory in data-driven analysis, signal processing, and control,” *Annual Reviews in Control*, vol. 52, pp. 42–64, 2021.
- [23] C. De Persis and P. Tesi, “Formulas for data-driven control: Stabilization, optimality, and robustness,” *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 909–924, 2019.
- [24] Y. Abbasi-Yadkori and C. Szepesvári, “Regret bounds for the adaptive control of linear quadratic systems,” in *Proceedings of the 24th Annual Conference on Learning Theory*. JMLR Workshop and Conference Proceedings, 2011, pp. 1–26.
- [25] “Asymptotically efficient adaptive control in stochastic regression models,” *Advances in Applied Mathematics*, vol. 7, no. 1, pp. 23–45, 1986.
- [26] K. J. Åström and B. Wittenmark, *Adaptive control*. Courier Corporation, 1995.
- [27] A. M. Annaswamy and K. S. Narendra, *Stable adaptive systems*. Prentice Hall, 1989.
- [28] M. C. Campi, A. Carè, and S. Garatti, “The scenario approach: A tool at the service of data-driven decision making,” *Annual Reviews in Control*, vol. 52, pp. 1–17, 2021.
- [29] T. Glasmachers, “Limits of end-to-end learning,” in *Asian conference on machine learning*. PMLR, 2017, pp. 17–32.
- [30] K. J. Åström, “Automatic control – the hidden technology,” in *Advances in Control: highlights of ECC’99*. Springer, 1999, pp. 1–28.