

Data-Driven Control 2/2

HOT TAKE: WHY NOT GO WITH MODELS?

Florian Dörfler

A recurring question that all authors of this special issue encounter is “*why not go with models?*” Two terms need to be clarified: in this context, a *model* is understood as a parametric system representation often endowed with an interpretable structure, e.g., a state-space representation with a readily discernible $F = m \cdot a$ equation. Further, the term *data-driven control*, as we employ it in this special issue, is not just about using data from a black box to inform decision-making. Researchers are exploring different paradigms, among others model-based control design, where the model and uncertainty estimates are learnt from data using contemporary system identification and uncertainty quantification techniques. In classical adaptive control terminology [1], [2], this two-stage approach is referred to as *indirect*. In contrast, *direct* data-driven control by-passes models in the decision making; see Figure 1 for a graphical illustration of the two paradigms. Hence, the more precise question should be “*When to embrace direct or indirect data-driven control?*” I will delve into the expected “*it depends*” answer in this editorial.

The typical answers one encounters often reason with the shortcomings of system identification, take the vantage point of a specific application, or quote the widely recognized success of reinforcement learning in computer and board games.

Let me put good automatic control practice into use, abstract the problem, and see more clearly through the lens of mathematics. Since the minimization of a fitting criterion is the prevalent

formulation in system identification, I will take an optimization perspective. Allow me to leave the uncertainty aspect aside for now. In this idealized setup, the indirect paradigm – first identify a model within a pre-specified class, then perform model-based control – is abstracted as a nested optimization problem:

indirect data-driven control	
minimize	control cost(u, y)
subject to	trajectory (u, y) compatible with the model
where	model \in argmin fitting criterion (u^d, y^d)
	subject to model belongs to a certain class

The indirect approach can be first and foremost described as being *modular* with two well separated levels. In comparison, the direct approach is more lean and seeks a decision compatible with the data. Abstractly, it is an *end-to-end* monolithic problem:

direct data-driven control	
minimize	control cost(u, y)
subject to	trajectory (u, y) compatible with data (u^d, y^d)

While the direct versus indirect classification is simple and useful, the line between both paradigms is often blurred. Let me now dive into the relative merits of both formulations.

On models: The distinguishing feature of both approaches is whether to use a model for decision making. Compared to raw data, models are tidied-up representations, i.e., compressed, de-noised, and typically also approximate. These features are most obvious in subspace system identification, where they are achieved by singular value thresholding of data matrices [3]. Crucially, models are interpretable, often structured, and physically intuitive, e.g., think of the class of Port-Hamiltonian systems [4]. Further, models are amenable to powerful control design methods (e.g., semidefinite programming for optimal control in state space [5]), and they are obviously useful beyond control, e.g., for analysis, simulation, or system design. Hence,




FIGURE 1 Direct and indirect paths from input/output data (u^d, y^d) to deciding upon control and output trajectories u and y . In practice, the design is often iterative (i.e., the diagram contains loops) and involves further processing steps, such as filtering of the raw data.

there will always be a place for models and indirect design.

However, there are indeed arguments to be made against models. Think about systems with complex physics (e.g., soft robotics), complex disturbances (e.g., wind farms), complex sensing modalities (e.g., perception-based control), or operating in complex environments (e.g., autonomous cars). In such instances, even if first-principle models were available, they might be too complex to be useful for control design, and one may argue for the direct approach. This brings to light that “*the main issue in modeling from data is approximation*” [6]. To take it a step further, rather than modeling the possibly complex data-generating mechanism, the data *should* be used to directly inform the decision-making, as argued in [7]–[9]. Of course, the “should” is debated. More concretely, the ultimate objective is the control policy, and often it might be easier to learn a control policy than learning a model. This catchy statement has often been voiced recently [10], and some historic examples include the widely deployed Ziegler-Nichols PID tuning from a single system response [11] or finite-time optimal control design based on few Markov parameters [12] or step responses [13]. Further, while persistence of excitation is a necessary identifiability condition in the indirect approach, unstable systems can sometimes be directly stabilized but not identified with limited data [14].

I close this point with a famous quote from [9] – “*When solving a given problem, try to avoid solving a more general problem as an intermediate step*” – and a simple disarming example: Consider asymptotic rejection of a constant disturbance. Indirect approaches require estimates of the model and the disturbance, whereas a direct approach is simply integral control.

Lack of separation principle: In general, the indirect approach might be suboptimal since there is no separation principle between identification and control, i.e., the best model fitting the data may not be the best model for the ultimate control task. For example, in classical frequency domain control, the DC gain and closed-loop bandwidth are often the two key control specifications, which should again inform at which frequencies to mainly

identify the system. This lack of a separation principle stimulated many approaches blending the inner identification and outer control objectives, e.g., dual control [15], [16], identification for control [17], [18], or approaches blending the identification and control objectives [19]–[22]. However, in some settings, indirect design via system identification is optimal, as discussed next.

On system identification: The inner system identification problem serves two crucial roles. First, system identification filters the data and reduces variance. If a parameterization of the data-generation mechanism is known, an indirect approach based on maximum-likelihood estimation and certainty-equivalence design is optimal in mean-square error [18] or regret metrics (for online LQR settings) [23]. Second, one can also interpret system identification as a projection of the data on a specific model class which is either a priori known or obtained in pre-processing. Indeed, this innermost model selection step is often hard but allows to inject prior knowledge, structure, and any sort of physical bias, such as stability, dissipativity, or positivity, see e.g., approaches on kernel-based system identification [24]–[26].

Model selection is evidently absent in the direct approach which makes it harder to include side information on the plant. Conversely, no bias error can be incurred due to an erroneously selected model class or inconsistent parameter estimates. While this is not a universal statement (e.g., some non-iterative direct model-reference approaches may also induce an undesired bias), it has been illustrated for specific settings in [27]–[30]. In a nutshell, direct approaches with imperfect learning can be more robust. Especially, in adaptive control settings it is argued that “*the indirect approach [aiming at model matching] is motivated by optimality and the direct approach [aiming at output matching] is motivated by stability*” [31]. Though, obviously no universal conclusions can be drawn for all methods, especially since the direct versus indirect distinction is blurry to begin with.

In my opinion, the robustness and bias-variance trade-offs of both approaches are still to be fully explored and quantified.

On a last note, system identification in practice is often an art, cumbersome, and may require re-adaptation. Further, it is argued in [18] and confirmed by the author’s experience that the engineering cost due to modeling, identification, and (re)commissioning is significant. The industrial report [32] bluntly concludes that “*about the only place the cost of dynamic modeling is ever warranted is during MPC implementation.*”

On uncertainty: Let me quote [33] – “*The most outstanding point of [direct] approaches is that the twinborn problem of unmodeled dynamics and robustness in traditional [model-based] theory do not exist under [the direct] framework*” – and illustrate this insight with an example: Consider a batch of noisy data, fed through non-convex prediction error identification, and yielding a nominal model of certain order together with a stochastic parametric uncertainty estimate accounting for both noise as well as unmodeled (e.g., higher-order) dynamics. In the indirect approach, these are then later used for optimal \mathcal{H}_∞ -control subject to an unstructured worst-case uncertainty. Observe the *twinborn* problem of the control design having to be robust due

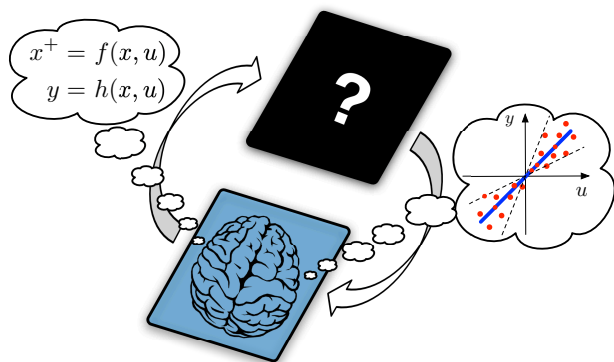


FIGURE 2 A control engineer confronted with a stylized black box model may question whether to adopt a direct data-driven design (right) or an indirect approach based on a mechanistic model (left).

to the model being only approximate. Further, in the indirect approach, uncertainty on the data needs to be propagated through the inner system identification problem, and the identification uncertainty estimates may be incompatible with the uncertainty quantification preferred for control design.

In comparison, direct approaches incorporate uncertainty on the data directly in the control design [34]–[36], and they do so in a transparent way without propagating it through system identification or approximating uncertainty estimates.

Complexity & implementation: There are many other arguments to be brought up. For instance, there is generally no winning approach in terms of “complexity”, i.e., there are instances of both methods that are data-efficient (or wasteful), that are analytically (in)tractable, that have (non)convex problem formulations, that can be carried out partially or entirely offline (respectively, online), that are easier (respectively, harder) to code or debug, and that require more or less human supervision.

While on the topic of implementation, the indirect approach is obviously modular with well-understood subtasks, which makes it a reliable and interpretable building block in a layered architecture. This is in stark contrast with the monolithic nature of the direct approach. Researchers in academia, in industry, and in application domains are arguing about the relative merits. For instance, in robotics this dichotomy is often referred to as *end-to-end* versus a layered *autonomy stack*. Proponents argue in favor of either approach when it comes to implementations, suitability for complex specifications, and all the previously listed topics.

This debate is far from settled, it epitomizes the grand challenge of architecture selection [37], and the infamous *middle road* might be the best path forward, as shown in many robotics implementations blending both paradigms, e.g., [38], [39].

Synopsis: I hope to have illuminated the unsatisfying “*it depends*” answer to my opening question from sufficiently many angles, clarified a few of the paradigms and research gaps, and gave food for thought for when to implement either approach. “The Articles of this Special Issue” each have take their own take on this question and develop a rich data-driven control theory.

Let me conclude with a few personal thoughts. Science and engineering based on traditional (i.e., first-principle and parametric) models have brought us far – literally to the moon and back. However, there are notoriously challenging problems with the traditional approach, and the impact of contemporary data-driven methods will be judged (among others) by whether they will overcome these challenges. For instance, can they take a stab at hard system classes (nonlinear, infinite dimensional, etc.), can they be implemented online with streaming data, and are there applications with a true business case for these methods?

ACKNOWLEDGMENT

I want to thank several colleagues that read drafts of this article and provided valuable feedback – among others, all authors of this double special issue, several of my colleagues at ETH Zürich and the NCCR Automation, as well as Andrea Iannelli, Alexandre Sanfelici Bazanella, Claudio De Persis, Anuradha

Annaswamy, Ian R. Manchester, Alessandro Chiuso, Håkan Hjalmarsson, and Alireza Karimi. I was hard pressed to include all opinions and different perspectives on this “hot take” topic.

AUTHOR INFORMATION

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

REFERENCES

- [1] K. J. Åström and B. Wittenmark, *Adaptive control*. Courier Corporation, 1995.
- [2] A. M. Annaswamy and K. S. Narendra, *Stable adaptive systems*. Prentice Hall, 1989.
- [3] P. Van Overschee and B. De Moor, *Subspace identification for linear systems: Theory, Implementation, Applications*. Springer Science & Business Media, 2012.
- [4] A. Van Der Schaft, D. Jeltsema *et al.*, “Port-hamiltonian systems theory: An introductory overview,” *Foundations and Trends® in Systems and Control*, vol. 1, no. 2-3, pp. 173–378, 2014.
- [5] C. Scherer and S. Weiland, “Linear matrix inequalities in control,” *Lecture Notes, Dutch Institute for Systems and Control, Delft, The Netherlands*, vol. 3, no. 2, 2000.
- [6] J. C. Willems, “In control, almost from the beginning until the day after tomorrow,” *European Journal of Control*, vol. 13, no. 1, p. 71, 2007.
- [7] L. Breiman, “Statistical modeling: The two cultures,” *Statistical science*, vol. 16, no. 3, pp. 199–231, 2001.
- [8] 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.
- [9] V. Vapnik, *The nature of statistical learning theory*. Springer science & business media, 1999.
- [10] B. Recht, “A tour of reinforcement learning: The view from continuous control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, pp. 253–279, 2019.
- [11] J. G. Ziegler and N. B. Nichols, “Optimum settings for automatic controllers,” *Transactions of the American society of mechanical engineers*, vol. 64, no. 8, pp. 759–765, 1942.
- [12] G. Shi and R. E. Skelton, “Markov data-based LQG control,” *J. Dyn. Sys., Meas., Control*, vol. 122, no. 3, pp. 551–559, 2000.
- [13] C. R. Cutler and B. L. Ramaker, “Dynamic matrix control – a computer control algorithm,” in *joint automatic control conference*, no. 17, 1980, p. 72.
- [14] H. J. Van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, “Data informativity: a new perspective on data-driven analysis and control,” *IEEE Transactions on Automatic Control*, vol. 65, no. 11, pp. 4753–4768, 2020.
- [15] A. Feldbaum, “Dual control theory problems,” *IFAC Proceedings Volumes*, vol. 1, no. 2, pp. 541–550, 1963.
- [16] B. Wittenmark, “Adaptive dual control methods: An overview,” *Adaptive Systems in Control and Signal Processing 1995*, pp. 67–72, 1995.
- [17] M. Gevers, “Identification for control: From the early achievements to the revival of experiment design,” *European Journal of Control*, vol. 11, pp. 1–18, 2005.
- [18] H. Hjalmarsson, “From experiment design to closed-loop control,” *Automatica*, vol. 41, no. 3, pp. 393–438, 2005.
- [19] M. C. Campi and P. Kumar, “Adaptive linear quadratic gaussian control: the cost-biased approach revisited,” *SIAM Journal on Control and Optimization*, vol. 36, no. 6, pp. 1890–1907, 1998.
- [20] S. Formentin and A. Chiuso, “Control-oriented regularization for linear system identification,” *Automatica*, vol. 127, p. 109539, 2021.
- [21] F. Dörfler, P. Tesi, and C. De Persis, “On the certainty-equivalence approach to direct data-driven LQR design,” *IEEE Transactions on Automatic Control*, 2023.
- [22] 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.
- [23] M. Simchowitz and D. Foster, “Naive exploration is optimal for online LQR,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 8937–8948.
- [24] M. Khosravi, “Side-information in linear and nonlinear system identification,”

THE ARTICLES OF THIS SPECIAL ISSUE

While all articles in this double special issue touch upon the direct versus indirect design questions from different perspectives, they also address numerous further challenges and all have their own unique foci. Specifically, in the present special issue, the articles focus on the informativity of data for different analysis and synthesis tasks, sample complexity estimates for identification and control, as well as automated tuning of hyperparameters in data-driven control design.

The article *The informativity approach to data-driven analysis and control* by van Waarde and co-workers puts informativity of the data as the central concept. The authors are particularly interested in the setting when the data are not sufficiently informative to (uniquely) identify a system. Nevertheless, analysis tasks or control design can be carried out for all systems sets that are unfalsified by the data. For noise-free data the approach leads to a robust theory for affine sets of systems, and for noisy data the methods draw inspiration from classical robust control.

The article *Statistical Learning Theory for Control: A finite sample perspective* by Tsiamis and co-workers surveys recent advances in statistical learning theory relevant to control and system identification. The article focuses on non-asymptotic sample complexity estimates for linear system identification and linear quadratic regulation leveraging modern tools from high-dimensional statistics and learning theory. The tutorial-style and self-contained presentation makes the topic accessible to control engineers. The article also outlines fruitful future directions.

Finally, the article *AutoDDC: hyperparameter tuning for direct data-driven control* by Breschi and Formentin addresses the challenges of automatically optimizing hyperparameters, i.e., degrees of freedom whose tuning requires significant human expertise. The authors focus on direct data-driven control methods to showcase the potential and the limitations of automatic hyperparameter tuning, namely, virtual reference feedback tuning and data-driven predictive control. They illustrate their strategies with applications from the automotive domain.

Ph.D. dissertation, ETH Zurich, 2021, available at <https://www.research-collection.ethz.ch>.

- [25] G. Pillonetto, T. Chen, A. Chiuso, G. De Nicolao, and L. Ljung, *Regularized system identification: Learning dynamic models from data*. Springer Nature, 2022.
- [26] H. J. van Waarde and R. Sepulchre, “Kernel-based models for system analysis,” *IEEE Transactions on Automatic Control*, 2022.
- [27] A. Becker, P. Kumar, and C.-Z. Wei, “Adaptive control with the stochastic approximation algorithm: Geometry and convergence,” *IEEE Transactions on Automatic Control*, vol. 30, no. 4, pp. 330–338, 1985.
- [28] V. Krishnan and F. Pasqualetti, “On direct vs indirect data-driven predictive control,” in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 736–741.
- [29] S. Formentin, K. Van Heusden, and A. Karimi, “A comparison of model-based and data-driven controller tuning,” *International Journal of Adaptive Control and Signal Processing*, vol. 28, no. 10, pp. 882–897, 2014.
- [30] F. Dörfler, J. Coulson, and I. Markovskiy, “Bridging direct and indirect data-driven control formulations via regularizations and relaxations,” *IEEE Transactions on Automatic Control*, vol. 68, no. 2, pp. 883–897, 2022.
- [31] A. M. Annaswamy, “Adaptive control and intersections with reinforcement learning,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 6, 2023.
- [32] L. Desborough and R. Miller, “Increasing customer value of industrial control performance monitoring-honeywell’s experience,” in *AIChE symposium series*, no. 326. New York: American Institute of Chemical Engineers; 1998, 2002, pp. 169–189.
- [33] 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.
- [34] A. Bisoffi, C. De Persis, and P. Tesi, “Data-driven control via Petersen’s lemma,” *Automatica*, vol. 145, p. 110537, 2022.
- [35] H. J. van Waarde, M. K. Camlibel, and M. Mesbahi, “From noisy data to feedback controllers: Nonconservative design via a matrix S-lemma,” *IEEE Transactions on Automatic Control*, vol. 67, no. 1, pp. 162–175, 2020.
- [36] J. Berberich, C. W. Scherer, and F. Allgöwer, “Combining prior knowledge and data for robust controller design,” *IEEE Transactions on Automatic Control*, 2022.
- [37] “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>
- [38] M. P. Polverini, S. Formentin, L. Merzagora, and P. Rocco, “Mixed data-driven and model-based robot implicit force control: A hierarchical approach,” *IEEE Transactions on Control Systems Technology*, vol. 28, no. 4, pp. 1258–1271,

2019.

- [39] R. T. Fawcett, K. Afsari, A. D. Ames, and K. A. Hamed, “Toward a data-driven template model for quadrupedal locomotion,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7636–7643, 2022.