

Model-Free Optimal Control: A Critical Analysis

Vijaysekhar Chellaboina^(✉)

Mahindra École Centrale, Hyderabad, India
vijay@mechyd.ac.in

Abstract. In this note, we present a critical analysis of machine learning techniques for applications involving optimal (feedback) control. Specifically, we will focus on the question of using reinforcement learning and other similar techniques in providing provably stable optimal controllers.

1 Introduction

The traditional control methodology is heavily based on the assumption that dynamical systems are represented in terms of clearly defined mathematical models [3, 14]. The two corner stones of control theory are stability and performance. While performance is a standard requirement in any design process, stability is inherently connected with the dynamic nature of dynamical systems. The definitions of stability are intricately related to the type of model used. Such models are either derived from first principles or through a rigorous system identification techniques based on experimental data [11, 20].

An ideal control methodology would be a model-free approach for deriving optimal controllers (which may be model-free themselves) just based on the input-output data of the system. At this time, machine learning techniques such as reinforcement learning come to close to providing a model-free approach for optimal control [24].

In this note, we provide a brief overview of both optimal control and machine learning techniques. In the discussion of optimal control, the need for the concept of the state and the state-space approach are considered in relation to stability of a system. In contrast, the overview of machine learning shows the difficulty in identifying either state or stability of systems consisting of machine learning *blocks*.

2 Optimal Feedback Control

In this section, we present a basic description of the problem of optimal control starting with a general description of dynamical systems, stability, performance, optimal, robust, and adaptive control. For more details on these well established topics one may refer to many of the standard references such as [2, 3, 6, 14, 22].

2.1 Dynamical Systems and Control

A *dynamical system* is a system that changes with time. Specifically, the complete description of a dynamical system consists of *input* and *output signals* (which are functions of time) and the relation between the input signals and output signals. Let \mathcal{U} denote the set of input signals $u : \mathbb{T} \rightarrow \mathbb{R}$ and \mathcal{Y} denote the set of output signals $y : \mathbb{T} \rightarrow \mathbb{R}$. Let the relation between u and y be given by a mapping $\mathcal{G} : \mathcal{U} \rightarrow \mathcal{Y}$ so that $y = \mathcal{G}u$, $u \in \mathcal{U}$ so that the description of the dynamical system is complete by specifying \mathcal{G} , \mathcal{U} , and \mathcal{Y} . If $\mathbb{T} = \mathbb{R}$ then the system is a *continuous-time* system and if \mathbb{T} is a finite set then it is a *discrete-time* system. In this section, we focus only on the continuous-time version. Analogues for discrete-time version can be easily developed and are well documented in the literature.

In the case where \mathcal{G} is a linear mapping, the description may be given in terms of the Laplace transforms of the input and output signals, that is, $Y(s) = G(s)U(s)$, where $G(s)$ is called the *transfer function* of the system. A system described by a transfer function may be equivalently represented in its *state-space* form given in terms of ordinary differential equations

$$\dot{x}(t) = Ax(t) + Bu(t), \quad (1)$$

$$y(t) = Cx(t) + Du(t), \quad (2)$$

where $x(t)$ is the *state* vector and A , B , C , D are *system matrices* are such that $G(s) = C(sI - A)^{-1}B + D$. Such a state-space description in terms of a (vector) ordinary differential equation is always possible if the transfer function is *real rational* and *proper*, that is, $G(s)$ is described in terms of a ratio of real polynomials with the numerator order less than or equal to that of the denominator. Even in the more general case, under mild technical assumptions, it is always possible to represent a linear system in a state-space form involving infinite-dimensional states. The most general description of a dynamical system is in terms of a state-function which maps initial state and inputs to the state at a future time [6, 7, 15].

In practice, the output signals are typically signals that can be measured using an instrument or a function of such signals. The input signals are typically divided into two categories (i) control inputs and (ii) disturbance inputs or noise. A *control problem* can then be stated as *determination of appropriate control input signals so that specific output signals follow a desirable pattern*. An *optimal control problem* is *determination of appropriate optimal control input signals so that specific output signals follow a desirable pattern and maximize a chosen performance*. It is understood that the performance will be in terms of only the input and the output signals.

2.2 Stability and State-Space Models

In control theory, every control system is designed for (i) stability and (ii) performance. The need for performance is clear and has been included in the control problem statement above. However, neither the definition nor the need for

stability is obvious. The difficulty starts with the fact that there are numerous definitions of stability. The definitions of stability can be broadly classified into two categories (i) bounded-input, bounded-output (BIBO) stability and (ii) equilibrium-state stability.

A *dynamical system* is *BIBO stable* if for every bounded input, the output remains bounded. It should be noted that input and output signals are functions of time and hence there is no unique way of measuring the size (in terms of a norm) of these signals and hence the same system (represented by, say, \mathcal{G}) may be BIBO stable with respect to certain choice of input-output norm pairs but not stable with respect to other choices. The most standard choice (though not necessarily most natural) for these norms is the Euclidean norm or the L_2 -norm. The *classical control theory* provides stability results for transfer functions has the interpretation of BIBO stability with the L_2 -norm.

As described above, systems can also be represented in a state-space form. And for the state-space model, one may identify special states called equilibria. A state is called an *equilibrium* if the system (in the absence of inputs) starts in an equilibrium state then it remains there. An equilibrium is said to be *stable* if the system starts close to an equilibrium then it remains close to the equilibrium and approaches the state asymptotically. If the system has only one equilibrium and it is stable (as per the definition above) then such a system may be called as a *stable system*. The definition of equilibrium stability given here is known as *asymptotic stability* in the literature. There are many more equally interesting forms of stability (all connected to one another) but will not be discussed here.

The reader may be wondering if there is a relation between these two broad categories of stability. The connection between the two categories is most interesting. *A linear system is BIBO stable if and only if it is asymptotically stable.* The case of nonlinear systems is slightly more complicated. There are classical results proving that equilibrium-state stability implies BIBO stability. Indeed most of the results in the literature on BIBO stability of nonlinear systems rely on equilibrium-state stability.

It should be noted that the concept of equilibrium-state stability is inherently connected to the definition of a state. Though neither a state-space description nor the concept of equilibrium-state stability is necessary for defining stability of a dynamical system but it is interesting to note that most results on stability are given for state-space models. The above discussion clearly shows that state-space description and the corresponding stability notions makes it convenient for designing a control system. However, it is not clear if these notions are absolutely necessary.

The popularity of state-space approach has a strong underlying reason. Most physical systems can be modeled using first principles (such as Newton's laws of motion) and state-space descriptions are very natural. In such case, the state of the system is identified with physical variables. It is therefore important that the control systems for physical systems are designed such that the entire state remains bounded. Hence, nonlinear system control is almost synonymous with control for equilibrium-state stability and state-boundedness (see [14] and

references therein). However, this is not the case if the system is not physical (for example, economic or socio-economic systems).

Are the concepts of state-boundedness and equilibrium-state stability essential for designing control systems?

2.3 Optimal Robust Control

In the previous section, we described the importance given to the concept of stability in systems theory. Here we focus on the issue of optimality and incomplete knowledge of system models. The most important goal in designing control systems is to design for best performance while satisfying all physical constraints. This is the subject of optimal control. Design of an optimal control for a given system requires the complete knowledge of the system. For example, in the linear case, the so called LQR or Kalman filters are optimal controllers (for specific performance criteria) and require the exact knowledge of the system matrices [2].

In practice, it is almost impossible to have the exact knowledge of system and hence every model has parameters or functions that are uncertain. Hence, the stability results will have to be extended for models where part of the model is unknown or uncertain. Such stability results form the concept of *robust stability* [8, 14, 27]. For example, if a system satisfies an input-output property known as *passivity* [14] and there is no other information available about the system one can design a passive controller to make the overall system stable. More generally, the concept of dissipativity can be used developed stability results for systems with different classes of model uncertainty [14]. The dissipativity-based results are applicable for systems with and without an explicit state-space characterization. These provide methods for designing a stabilizing controller for a set of (uncertain) systems and the performance obtained will be the worst-case performance. Hence, the optimal control methodology based on robust stability concepts can only be used to design controllers maximizing the worst-case performance. If the uncertainty set is *large* then the optimal performance will be poor.

2.4 Adaptive Control

An alternate method to control uncertain systems is to use the idea of *adaptive control* [4, 17, 22]. The main idea of adaptive control is to adapt the parameters of the controllers so that the performance of the closed-loop system is optimal at each and every operating point. In the 1950's and the 1960's, NASA had an extensive research-airplane program to test the adaptive control methodology and it was finally shut down due to a fatal accident [18]. The analysis of the accident revealed that the failure is due to overall stability issues (as opposed to stability at every operating point). This led to the development of stable or provably correct adaptive systems [22] based on rigorous Lyapunov approach applied to state-space models with parametric uncertainty. Stable adaptive control, as compared to robust control, provides a framework for stabilization as

well as optimal performance. See [10] for a recent analysis of the NASA X-15 program and the lessons learned from it.

Stability takes priority over performance

3 Machine Learning for Control

In this section, we discuss the relevance of machine learning techniques in the context of feedback control with specific focus on the issues of stability, performance, and uncertainty. First, we present a brief overview of machine learning techniques and their specific role in control systems. For an introduction to machine learning techniques and its allied topics see [1, 5, 13, 21, 24, 26] and numerous references within.

3.1 Overview of Machine Learning

The ever increasing ability to manipulate and to compute with large sets of data makes it possible to implement many of the machine learning techniques for a variety of applications. *Machine learning* may simply be defined as extracting *information* from *data* using computing machines. Extracting information from data is as old as science and every fundamental laws of nature is an example of such extraction. Hence, every model discovery is such an example. If a computational tool is utilized in extracting such information or a model from data then the methodology is dubbed as *artificial intelligence* or more modestly *machine learning*.

Machine learning techniques process large data (of input and output of a given system) to essentially provide a *black-box* model of the system, which may then be used in predicting output for an input that is not part of the original data. The black box may contain any number of models available including *neural networks*, *support vector machines*, *decision trees*, and a variety of other underlying models. These models are derived using error minimization techniques including *back propagation algorithms* and *reinforcement learning algorithms*. Since these techniques are primarily based on the paradigm of processing large amount of data the models thus obtained may also be significantly high dimensional, essentially rendering them useless in terms of rigorous mathematical analysis. Hence, the abstraction of data into one of the machine learning models is dubbed as *model-free* approach. Here, the model-free approach also refers to the fact that a model is not developed from the first principles (laws of nature) but only from the available data (input as well as output).

Machine learning provides a model-free data-based approach to predict outputs for unknown inputs

3.2 Model-Free Control

The literature on machine learning applications to control systems is a lot sparser compared to that of other applications. See, for example, [9, 19, 23, 25]. With exception of neural network based control (see for example, [12, 16]) most of the machine learning based control do not focus on the proof of stability. The primary focus of such literature is on performance maximization or optimizing parameters of a stabilizing controller. The available literature on machine learning applications to control systems can be broadly classified as follows:

- (i) Neural-network or similar model-based optimal control where the weights/parameters of the model are adapted for stabilization and optimal performance. In this case, the models are invariably state-space based and the stability proof is in terms of equilibrium-state using Lyapunov-like approaches. This is simply a large scale version of model-based state-space approach to control.
- (ii) Reinforcement-learning or a similar technique is used to develop optimal controllers based only on (large amount of) input-output data with appropriate control problem stated in terms of input and output signals. In most of these cases, (BIBO) stability is an inherent quality of the system or no proof of stability is considered. This is truly a model-free approach to control, that is, no model is derived from first principles or the available large-scale model is not useful for formal analysis.

Note that the above model-free approach is ideal for systems which are inherently bounded (for example, finite-state machines). In this case, the reinforcement learning and other techniques can be used to extract maximum performance from the system. However, in the case of systems where the boundedness is not inherent to the system then it is not clear that the model-free approach is sound. Though there have been multiple demonstrations of model-free approach to control on a variety of problems, only time will tell if it is indeed a safe approach in every operating condition. As the NASA X-15 program taught us that performance does not imply stability. The following two questions (and their derivatives) remain unanswered at this time:

- (i) Is there a provably-correct stable machine learning control (that is different from adaptive control)?
- (ii) In the case of model-free approach, what are the definitions of state or stability? More, fundamentally is there a need for such paradigms?

4 Conclusion

In this note, we considered the question of stability in a model-free approach to control. Specifically, we first present an overview of traditional control concepts with specific focus on the issue of stability and its related concepts such as state and state-space models. This is followed by a very high level introduction to machine learning techniques for control and discussed the difficulty as well as need for the concept of stability in such a model-free approach.

References

1. Alpaydin, E.: Introduction to Machine Learning. MIT Press, Cambridge (2014)
2. Antsaklis, P.J., Michel, A.N.: Linear Systems, vol. 1. Birkhäuser, Boston (2006)
3. Åström, K.J., Murray, R.M.: Feedback Systems: An Introduction for Scientists and Engineers. Princeton University Press, Princeton (2010)
4. Åström, K.J., Wittenmark, B.: Adaptive Control. Courier Corporation, Mineola (2013)
5. Bertsekas, D.P., Tsitsiklis, J.N.: Neuro-Dynamic Programming, 1st edn. Athena Scientific, Belmont (1996)
6. Bhatia, N.P., Szegő, G.P.: Dynamical Systems: Stability Theory and Applications, vol. 35. Springer, Heidelberg (2006). <https://doi.org/10.1007/BFb0080630>
7. Chellaboina, V., Bhat, S.P., Haddad, W.M.: An invariance principle for nonlinear hybrid and impulsive dynamical systems. *Nonlinear Anal. Theory Methods Appl.* **53**(3), 527–550 (2003)
8. Dullerud, G.E., Paganini, F.: A Course in Robust Control Theory: A Convex Approach, vol. 36. Springer, New York (2013). <https://doi.org/10.1007/978-1-4757-3290-0>
9. Duriez, T., Brunton, S.L., Noack, B.R.: Machine Learning Control-Taming Nonlinear Dynamics and Turbulence. Springer, Heidelberg (2017). <https://doi.org/10.1007/978-3-319-40624-4>
10. Dydek, Z.T., Annaswamy, A.M., Lavretsky, E.: Adaptive control and the NASA X-15-3 flight revisited. *IEEE Control Syst.* **30**(3), 32–48 (2010)
11. Fogel, D.B.: System Identification Through Simulated Evolution: A Machine Learning Approach to Modeling. Ginn Press, Needham Heights (1991)
12. Ge, S.S., Hang, C.C., Lee, T.H., Zhang, T.: Stable Adaptive Neural Network Control, vol. 13. Springer, New York (2013). <https://doi.org/10.1007/978-1-4757-6577-9>
13. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge (2016)
14. Haddad, W.M., Chellaboina, V.: Nonlinear Dynamical Systems and Control: A Lyapunov-Based Approach. Princeton University Press, Princeton (2008)
15. Haddad, W.M., Chellaboina, V., Nersisov, S.G.: Impulsive and Hybrid Dynamical Systems. Princeton Series in Applied Mathematics (2006)
16. Hovakimyan, N., Cao, C.: \mathcal{L}_1 Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation. SIAM, Philadelphia (2010)
17. Ioannou, P.A., Sun, J.: Robust Adaptive Control, vol. 1. PTR Prentice-Hall, Upper Saddle River (1996)
18. Jenkins, D.R.: Hypersonics before the shuttle: A concise history of the X-15 research airplane (2000)
19. Lewis, F.L., Liu, D.: Reinforcement Learning and Approximate Dynamic Programming for Feedback Control, vol. 17. Wiley, New York (2013)
20. Ljung, L.: System identification. In: Procházka, A., Uhlíř, J., Rayner, P.W.J., Kingsbury, N.G. (eds.) Signal Analysis and Prediction. Applied and Numerical Harmonic Analysis, pp. 163–173. Springer, Heidelberg (1998). https://doi.org/10.1007/978-1-4612-1768-8_11
21. Michalski, R.S., Carbonell, J.G., Mitchell, T.M.: Machine Learning: An Artificial Intelligence Approach. Springer, Heidelberg (2013)
22. Narendra, K.S., Annaswamy, A.M.: Stable Adaptive Systems. Courier Corporation, Mineola (2012)

23. Ng, A.Y., Coates, A., Diel, M., Ganapathi, V., Schulte, J., Tse, B., Berger, E., Liang, E.: Autonomous inverted helicopter flight via reinforcement learning. In: Ang, M.H., Khatib, O. (eds.) *Experimental Robotics IX*. Springer Tracts in Advanced Robotics, vol. 21, pp. 363–372. Springer, Heidelberg (2006). https://doi.org/10.1007/11552246_35
24. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*, vol. 1. MIT press, Cambridge (1998)
25. Sutton, R.S., Barto, A.G., Williams, R.J.: Reinforcement learning is direct adaptive optimal control. *IEEE Control Syst.* **12**(2), 19–22 (1992)
26. Witten, I.H., Frank, E., Hall, M.A., Pal, C.J.: *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, San Francisco (2016)
27. Zhou, K., Doyle, J.C.: *Essentials of Robust Control*, vol. 104. Prentice hall, Upper Saddle River (1998)