Abstract

Due to their relevance in systems analysis and controller design, we consider the problem of determining system-theoretic input-output properties of linear time-invariant (LTI) systems. While, in practice, the input-output behavior of dynamic systems is often undisclosed and obtaining a suitable mathematical model via first principles can be a cumbersome task, data of the system in form of input-output trajectories is often and increasingly available. Therefore, we present different methods to determine system-theoretic input-output properties directly from data without deriving or identifying a mathematical model first. In particular, we study iterative methods, where data is actively sampled by performing experiments on the unknown system, as well as approaches based on available (offline) data. Considering these offline approaches, we first develop necessary and sufficient conditions requiring only one input-output trajectory to certify system properties on the basis of Willems' fundamental lemma and then introduce robust approaches providing guaranteed bounds on the respective input-output property in the case of noisy data.

Iterative methods. We generalize state-of-the-art iterative sampling schemes determining the $\mathcal{L}_2$-gain of an LTI system to a more general framework in Chapter 3, which includes the consideration of passivity properties and conic relations. Such iterative schemes require active sampling of input-output data by performing experiments or simulations on the unknown system with iteratively updated input signals. We first show how the respective input-output system properties can be reformulated in terms of optimization problems and present sampling strategies based on gradient dynamical systems and saddle point flows to find their optimizers. These sampling strategies are based on the fact that the gradients of the optimization problems can be evaluated from only input-output data samples even though the input-output behavior of the system is unknown. This leads us to evolution equations, whose convergence properties are then discussed in continuous time and discrete time. Multiple numerical examples show the potential and applicability of the introduced approaches.

Offline methods. Since the requirement of iterative experiments can be limiting, we introduce a necessary and sufficient condition for LTI systems to verify input-output system properties from only one input-output trajectory on the basis of Willems' fundamental lemma in Chapter 4. More specifically, we consider the general classes of dissipativity properties and integral quadratic constraints (IQCs). For important classes of such input-output system properties, we provide convex optimization problems in form of semidefinite programs (SDPs) to retrieve the optimal, i.e., the tightest, system property description that is satisfied by the unknown system. Furthermore, we provide insights and results on the difference between finite and infinite horizon IQCs and finally illustrate the effectiveness of the proposed scheme in a variety of simulation studies including noisy measurements and a high dimensional system.

Robust methods. While the previous iterative and offline approaches for data-driven dissipativity analysis guarantee the dissipativity condition only over a finite-time horizon and provide no quantitative guarantees on robustness in the presence of noise, we provide a framework to verify dissipativity properties from noisy data with desirable guarantees in Chapter 5. We first consider the case of input-state measurements, where we provide nonconservative and computationally attractive dissipativity conditions in the presence of unknown but bounded process noise. We then extend this approach to input-output data in the noise-free as well as in the noisy case. Finally, we apply the proposed approach to real-world data of a two-tank water system and illustrate its applicability and advantages compared to established methods based on system identification.

The main goal of this thesis is to introduce a framework to determine input-output system properties directly from data without deriving or identifying a mathematical model first. To this end, we develop different methods based on various control-theoretic results and insights, which all have their individual advantages, limitations, and their corresponding application cases.

Chapter 1

Introduction

1.1 Motivation

Most state-of-the-art systems analysis and controller design methods are based on an accurate mathematical model that adequately describes the system at hand. While the available literature on model-based systems analysis and controller design is quite elaborate and comes with rigorous guarantees for stability, performance, and robustness, acquiring a suitable and accurate mathematical model of the plant via first principles can be a time-consuming task that depends on expert knowledge. At the same time, data are becoming ubiquitous, cheap, and increasingly available in engineering applications in form of input-output trajectories of dynamic systems from experiments or simulations. These data potentially capture sufficient information about the otherwise unknown system needed for systems analysis and controller design. Therefore, there has recently been an increasing interest in the field of data-driven methods for systems analysis and control that systematically extract and directly exploit the information captured in the data.

This increasing interest in data-driven methods led to a significantly growing area of data-driven control approaches. In the survey paper [40] on data-driven controller design techniques, some rather established approaches to data-driven control are summarized such as virtual reference feedback tuning [17], iterative feedback tuning [39], and unfalsified control [45]. Since the appearance of this survey in 2013, many more approaches in the field of data-driven control have emerged and the literature is expanding rapidly. These recent approaches include stabilizing, optimal or robust control for linear time-invariant (LTI) systems on the basis of data-driven closed-loop parametrizations [75, 102, AK3], data-driven predictive control exploiting Willems' fundamental lemma [8, 20, 111], linear quadratic regulator (LQR) tuning via Gaussian process optimization [50, 66], data-driven control of networks and multi-agent settings [4, 24, AK24], mentioning only a few examples.

With the increasing availability of and interest in data, also the well-established field of system identification or data-driven modeling has seen a renewed interest. While the general area of system identification has a long history and provides a well-established and extensive amount of literature (see, e.g., [47] and references therein), there are still open questions. Even for LTI systems, finding nonasymptotic guarantees on a model from noisy data is yet a challenging problem [97] and part of ongoing research. Some promising results in this direction have been presented in [56], where the authors identify a model with guaranteed bounds on the uncertainty from data that are perturbed by stochastic noise, which can then be accounted for in robust controller design approaches [21]. Two further results in this direction include [74, 91] providing nonasymptotic system identification guarantees. However, all of these approaches require rather strong assumptions on the data and the noise. In [21], for example, all but the last state transition in each trajectory are discarded, in [91] state measurements corrupted by Gaussian independent and identically distributed (i.i.d.) noise are required, and assumptions in [74] include Gaussian i.i.d. noise and input, known system order n, and zero initial condition. System identification approaches in the case of deterministic noise typically rely on set membership estimation, where the key challenge is in providing both nonconservative but computationally tractable error bounds [61].

One complementary approach to the direct controller design from data or to identifying a full mathematical model of the system, respectively, is to learn and analyze certain system-theoretic input-output properties from data. Such input-output properties not only provide valuable insights into the system, but can also be leveraged to design a controller. In fact, properties such as dissipativity properties or integral quadratic constraints (IQCs) allow for the direct application of well-known feedback theorems for controller design, as shown for example in [22, 23, 99, 113]. These controllers come with desirable guarantees on stability and robustness, which is still an open problem for many data-driven control approaches. Furthermore, by first determining input-output properties, the controller structure is not a priori determined or parametrized in contrast to many data-driven control and data-driven tuning approaches. At the same time, learning such system properties from data can retain many of the desired advantages of data-driven methods compared to system identification approaches. They can be simple to apply without requiring expert knowledge, avoid the computational load of identifying a full mathematical model, and skip a possibly unnatural fit to a parametric system model potentially introducing additional error. One specifically valuable application area of input-output system properties includes cooperative control and multi-agent settings, where dissipativity can be used for compositional certification of stability, performance, and safety [2]. Besides the application of controller design via well-known feedback theorems, applications of input-output system properties from data also include controller validation [36] as well as fault detection and mitigation [112].

Therefore, we aim to provide a framework to determine system-theoretic input-output properties such as dissipativity properties and IQCs directly from data with guarantees. We strive to introduce methods that are easy to apply, computationally attractive and do not require expert knowledge. In this thesis, we focus on approaches to determine input-output properties of LTI systems, while we will discuss potential extensions to nonlinear systems in the individual summaries as well as in the conclusions of this thesis.

1.2 Related work

Due to the relevance of input-output properties in systems analysis and control and the well-established literature on dissipativity-based controller design, there has been a considerable number of approaches to determine such properties from data. Thus, the following literature review summarizes existing results for determining input-output properties from data as well as points out related methods and approaches.

Very generally, the literature on data-driven verification of input-output system properties can be roughly categorized into three conceptually different setups, as presented in the following.

Input-output properties of nonlinear systems from big data

Firstly, there exist some interesting ideas for determining input-output system properties for rather general classes of nonlinear systems from large amounts of input-output data tuples that are stored and available for analysis. In [62], the authors derive overestimates on the $\mathcal{L}_2$-gain, the shortage of passivity, and a cone containing all input-output samples based on finite, but densely sampled, input-output data. This idea was further analyzed in [89] and extended to nonlinearity measures in [54]. All these methods are based on Lipschitz assumptions on the nonlinear system, exploited in a similar fashion also in nonlinear data-driven control approaches [68, 69]. More general dissipation inequalities were considered in [AK20], where the ordering of the supply rates via the S-procedure allows to infer system properties from only a finite amount of input-output data. Another approach to determine passivity properties via Gaussian process optimization was introduced in [AK23], which additionally allows for active sampling schemes (i.e., choosing an input and performing an experiment) to increase data efficiency. This idea was extended to a deterministic setting in [53] via successive graph approximation to infer nonlinearity measures.

However, to receive quantitative bounds on certain dissipation inequalities, all these approaches require immense amounts of input-output trajectories, which limits their applicability. While we focus only on LTI systems in this thesis, we discuss extensions for certain classes of nonlinear systems, which potentially reduce the conservatism or the required number of data samples significantly compared to the existing literature on data-driven analysis of nonlinear systems introduced above.

Input-output properties of LTI systems from iterative sampling schemes

The second category of data-driven systems analysis methods is based on input signal optimization and iterative experiments for LTI systems, closely related to the general idea of iterative learning control (ILC), see, e.g., [14, 67]. The basic idea for iterative data-driven systems analysis has already been introduced in 2005 [38, Section 12.2], where the author applies a so-called power iteration method to receive an $\mathcal{L}_2$-gain estimate of the otherwise unknown system. The power iteration method is based on the well-known power method to determine the dominant eigenvector and eigenvalue pair of a matrix (see, e.g., [31]). It can be shown that this power iteration method, based on iterative experiments, is guaranteed to converge to the gain of the system [106], which has been further analyzed in terms of asymptotic statistical properties in [105]. While the method proposed in [106] requires two experiments per iteration, it is shown in [82] how this can be reduced to one experiment per iteration. In [72], this sampling scheme is extended to multiple-input multiple-output (MIMO) systems and applied to robust active vibration isolation.

Another similar approach is related to game theory, where the input domain is discretized at specific frequencies and the inputs over the iterations are adapted such that the excitation level is increased in those regions of the frequency domain where the peak is expected to be located [79]. Along the same lines, the multi-armed bandit approach is adopted to $\mathcal{L}_2$-gain estimation in [64] with the goal to maximize the probability of choosing the arm drawing samples, i.e., frequencies, with the highest amplitude gain in the output. This approach has been extended in [63] and improved in [65] by combining power iterations and Weighted Thompson Sampling.

However, all these iterative approaches for LTI systems [63, 64, 65, 72, 79, 82, 105, 106] only consider estimating the $\mathcal{L}_2$-gain. In Chapter 3 in this thesis, we generalize the idea of input signal optimization and iterative experiments for LTI systems to infer the $\mathcal{L}_2$-gain to a more general framework, which includes the consideration of passivity properties and conic relations.

Input-output properties of LTI systems from one offline trajectory

Finally, the third category of methods to verify dissipativity properties from data includes offline computational approaches from one input-state or input-output trajectory of LTI systems. In this direction, Willems' fundamental lemma [108] offers a fruitful basis by proving that the full behavior of an LTI system can be described by a Hankel matrix containing one previously measured input-output trajectory, given that the input component is persistently exciting. This result, as discussed in [9, 104], provides an equivalent description of LTI systems based only on data, which hence allows for systems analysis and control with desirable guarantees from input-output trajectories. Contributions to data-driven approaches based on Willems' fundamental lemma range from data-driven simulation and output-matching [52] to data-driven controller design [109, AK27] and data-driven model predictive control [8, 20, 111].

Considering data-driven systems analysis, Willems' fundamental lemma allows to optimize over the input signals not via iterative experiments but offline from only one input-output trajectory since it provides a data-based characterization of all trajectories of an unknown LTI system. The idea of determining dissipativity from a one-shot trajectory on the basis of the behavioral framework via Willems' fundamental lemma was introduced in [57]. However, their approach results in a nonconvex indefinite quadratic program, which is generally very hard to solve. Therefore, we introduce in Chapter 4 of this thesis how dissipativity properties, and more generally IQCs, can be certified by verifying positive semidefiniteness of a single data-dependent matrix.

One limitation to all of the introduced data-driven dissipativity results is, however, that no quantitative guarantees can be provided in the case of noisy data. To approach this, we make use of another fruitful line of work related to Willems' fundamental lemma, which is based on the idea of finding a parameterization of the closed loop from input-state trajectories of an otherwise unknown system [75]. This idea becomes especially interesting when it comes to noise-corrupted trajectories, where a suitable parameterization can represent all possible closed loops that are consistent with the data [AK3], hence allowing for robust data-driven state-feedback design from noise-corrupted data. This idea has been further extended and improved, e.g., in [10, 102]. In Chapter 5 of this thesis, we utilize this line of work to derive rigorous deterministic bounds on dissipativity properties even from noise-corrupted data.

1.3 Contribution and outline of this thesis

In the following, we detail the outline of this thesis and clarify the contributions.

Chapter 2: Background

In Chapter 2, we present the basic problem setup underlying this thesis. In particular, we define and discuss the input-output system properties of interest. Furthermore, we provide equivalent representations from different viewpoints of such system properties, which will be exploited throughout the thesis.

Chapter 3: Iterative schemes to determine input-output properties

The thrust of Chapter 3 is to present a systematic approach to iteratively determine certain dissipativity properties from input-output samples, where the input-output map remains undisclosed. To this end, we introduce different sampling strategies that can be summarized by (i) formulating the system property of interest as an optimization problem, and (ii) iteratively performing experiments that update the input along the gradient of the resulting optimization problem. In particular, we use multiple input-output trajectories from iterative experiments to investigate the $\mathcal{L}_2$-gain, the shortage of passivity and conic relations. We start in Section 3.1 with discrete-time LTI systems with a thorough analysis of continuous-time optimization as well as the implications for the iterative scheme (discrete-time optimization), where advanced sampling schemes can improve the convergence rate. In Section 3.2, we generalize the framework presented in Section 3.1 by, firstly, showing how the introduced iterative approach to determine input-output system properties can be extended to continuous-time LTI systems in Section 3.2.1. Secondly, we extend the framework to MIMO systems in Section 3.2.2 and additionally provide results on the robustness of the presented framework to measurement noise in Section 3.2.3. Finally, we apply the introduced approaches to different simulation examples in Section 3.3, including an oscillating system and a high dimensional system, and we end with a short summary in Section 3.4.

The results of Chapter 3 have been previously presented in [AK8, AK19, AK21, AK22].

Chapter 4: Offline approaches to determine input-output properties

In Chapter 4, we provide a computationally simple approach to certify input-output system properties from offline data. More precisely, we introduce necessary and sufficient conditions for a discrete-time LTI system to satisfy dissipativity properties, or more generally IQCs, given only one input-output data trajectory in Section 4.1. These conditions can be verified by simply checking whether a single data-dependent matrix is positive semidefinite, with the only requirements being i) a persistently exciting (but otherwise arbitrary) input signal and ii) knowledge of an upper bound on the lag of the system. This theory is based on Willems' fundamental lemma [108], a result originally developed in the context of behavioral systems theory, which provides a data-based characterization of all trajectories of an unknown LTI system. On the basis of the developed simple condition to infer IQCs from data, we extend the approach to noisy measurements, where we provide a very promising heuristic relaxation. Additionally, we characterize optimal system properties and provide semidefinite programs (SDPs) to find such IQCs in Section 4.2, which can provide more informative and tighter descriptions of the unknown system allowing for, e.g., less conservative robust controller designs. While in most of Chapter 4 the IQC property is only considered over a finite-time horizon, we infer bounds on the respective system property over the infinite-time horizon in Section 4.3. We conclude this chapter with simulation studies on a high dimensional numerical example in Section 4.4, demonstrating the potential of the introduced approach, together with a short summary in Section 4.5.

The results of Chapter 4 have been previously presented in [AK4, AK16].

Chapter 5: Bounds on input-output properties via a robust viewpoint

In Chapter 5, we develop guarantees on dissipativity properties from trajectories corrupted by bounded process noise. To this end, we first introduce an equivalent dissipativity characterization on the basis of one input-state trajectory in Section 5.1. In contrast to many other approaches to determine input-output properties from data (e.g., iterative schemes [106, AK8] and methods based on Willems' fundamental lemma [57, AK16]), we exploit the state-space definition of dissipativity in this chapter, which can be verified by taking a difference viewpoint, i.e., looking at the difference at two time points. This yields the advantage that guarantees on system properties over the infinite horizon can be obtained. More importantly, this allows to introduce a rigorous treatment of noisy data in Section 5.2, where we provide a computationally attractive and nonconservative robust verification framework for dissipativity properties from input-state data corrupted by bounded process noise. To this end, we first characterize all system matrices that are consistent with the data and an a priori known noise bound and then verify the dissipativity property of interest for all systems in this set. As the requirement of state measurements can be restrictive in practice, we extend the results to input-output trajectories in the noise-free case in Section 5.3, followed by a consideration of noise-corrupted input-output trajectories in Section 5.4. This finally yields a method to guarantee dissipativity properties over the infinite horizon from one noise-corrupted input-output trajectory of finite length via SDPs. Lastly, we apply the introduced approach to real-world data of a two-tank water system in Section 5.5 and conclude with a short summary in Section 5.6.

The results of Chapter 5 have been previously presented in [AK5, AK6].

Chapter 6: Conclusions

In this chapter, we summarize the contributions of this thesis, contrast the different approaches to data-driven inference of input-output system properties, and provide some directions for future research.