PAPER • OPEN ACCESS

Probabilistic behavioral distance and tuning—reducing and aggregating complex systems

View the article online for updates and enhancements.

Journal of Physics: Complexity

PAPER

Probabilistic behavioral distance and tuning—reducing and aggregating complex systems

Frank Hellmann[1,*] , Ekaterina Zolotarevskaia[1,*], Jürgen Kurths[1] and Jörg Raisch[2]

1 Potsdam Institute for Climate Impact Research, Potsdam, Germany
2 TU Berlin, Berlin, Germany
* Authors to whom any correspondence should be addressed.

E-mail: hellmann@pik-potsdam.de and zolotarevskaia@pik-potsdam.de

Abstract

Given two dynamical systems, we quantify how similar they are with respect to their interaction with the outside world. We focus on the case where simpler systems act as a specification for a more complex one. Combining a behavioral and probabilistic perspective we define several useful notions of the distance of a system to a specification. We show that these distances can be used to tune a complex system. We demonstrate that our approach can successfully make non-linear networked systems behave like much smaller networks, allowing us to aggregate large sub-networks into one or two effective nodes. Finally, we discuss similarities and differences between our approach and H_∞ model reduction.

1. Introduction

Given two dynamical systems with a well-defined interaction with the outside world, we want to quantify how similar these systems are from the point of view of the outside world. For example for the case of power grids, a distribution grid is typically coupled to the transmission grid at a single point of common coupling. If we want to specify how the distribution grid should behave, then as far as the transmission grid is concerned, all we care about is the behavior at the point of common coupling.

In order to make such a similarity measure useful in real world systems, we need to ensure that it can be efficiently approximated numerically. This way it can be used to tune a complex system so that it typically appears like one of a class of simple systems that we specify.

To study this problem it is natural to take the behavioral approach to the theory of dynamical systems. The hallmark of this approach is that it focuses on the possible trajectories that a system can exhibit in interactions with other systems [1, 2]. In the context of control theory these could be for example the inputs and outputs of a closed loop system. This perspective allows one to speak about specification and system behavior on the same level. The basic notion is: A dynamical system is a set of possible trajectories, a specification is a set of permitted trajectories and a dynamical system satisfies a specification if every possible trajectory is permissible.

This paper introduces several notions of distance between a system and a specification. These distance concepts are probabilistic, meaning they can be usefully approximated. Obtaining meaningful and usable notions of the distance of system and specification enables us to tune complex systems to satisfy specifications at least approximately. The context we have in mind is tuning the control of a subsystem of a larger network to present a unified aggregate behavior towards the 'rest of the world'. In this setting the links towards the larger network act as inputs/outputs. Another application would be the hierarchical decomposition of a larger control task, where the specification of lower levels in the hierarchy serve as systems for the higher level [3].

Two different distance notions are introduced, one suited for optimizing systems towards a specification, and one for probabilistically validating their compliance. We show that standard tools from differential

programming and non-linear optimization can be used successfully to tune complex systems by minimizing the sampled approximation of the distance by tuning several complex non-linear dynamical networks.

The method developed here is designed with applications to power grids in mind. In parallel to this paper we present a software stack based on the capabilities of the Julia language [4, 5], that allows tuning dynamical properties in power systems [6–8]. In this context a behavioral approach is particularly natural. Typically the regulatory specification of how the system should behave is not given in terms of a precise set of differential equations, but rather by describing general properties of the trajectories. For example the system limits for the rate of the change of frequency (RoCoF) at which disconnections of major generators will occur, and thus cascading blackouts become likely, is defined by the European transmission system operators by the conditions that the moving average of the RoCoF stays within: $\pm 2\,\mathrm{Hz\,s}^{-1}$ for a 500 ms window, $\pm 1.5\,\mathrm{Hz\,s}^{-1}$ for a 1000 ms window and $\pm 1.25\,\mathrm{Hz\,s}^{-1}$ for a 2000 ms window [9]. Any frequency trajectory that stays inside this curve is considered acceptable. At the same time, the demands, perturbations and faults that the power grid experiences are varied and random, necessitating a probabilistic approach. This has long been standard in the analysis of static properties, e.g. [10, 11], but is increasingly also used to systematically understand dynamic aspects of the system [12–15].

2. Systems and specifications

The highly abstract definitions of behavioral dynamical systems theory are hard to work with directly. The setting of this paper is to consider behaviors given by parametrized input-output differential equations. Fix some time interval $T = [0, t_{\mathrm{final}}]$, state space X, input space I and output space O. We denote the space of functions from from T to X by X^T. Denote the trajectories of the internal states as $x \in X^T$, the input states $i \in I^T$, and the output states $o \in O^T$. *We will always use x, i and o to refer to the function, and $x(t) \in X$ to refer to a concrete value.* The dynamical system is then given by specifying the dynamics f, the output function g, and the initial conditions x_0, all of which can depend on parameters $p \in \mathcal{P}$. Note that for notational simplicity we have chosen to include the initial conditions in the set of parameters. If the choice of initial conditions is free in X this means that the space of parameters factorizes as $\mathcal{P} = \tilde{\mathcal{P}} \times X$ for some $\tilde{\mathcal{P}}$. Below we will also consider systems with fixed initial conditions. Finally the inputs are restricted to some set $\mathcal{B}_i$, the equations then are:

$$\dot{x}(t) = f(x(t), i(t), p) \tag{1}$$
$$o(t) = g(x(t), p)$$
$$x(0) = x_0(p)$$
$$p \in \mathcal{P} \text{ and } i \in \mathcal{B}_i \subset I^T.$$

We always assume that these equations can be integrated for the time period T. The set of possible trajectories of this set of equations is parametrized by $\mathcal{P} \times \mathcal{B}_i$.

In what follows it will be important to distinguish between a system with parameters, a system without parameters and a specification. Even though all behaviors we consider will be of the input-output (IO) form (1), we will denote specifications using $z \in Z^T$ for internal state, and q for parameters, and we will denote the complete system of the specification as $\mathcal{C}$:
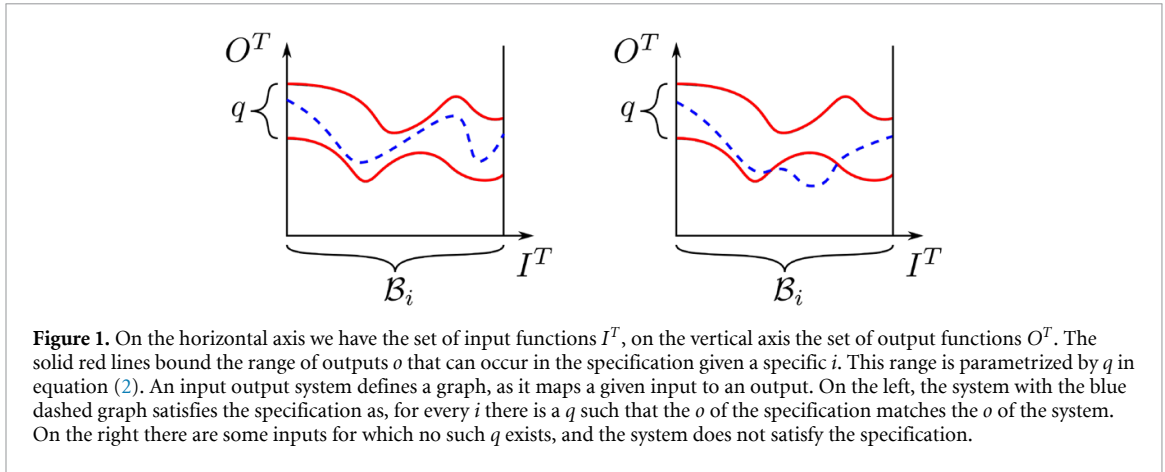
$$\mathcal{C} : \dot{z}(t) = f^{\mathcal{C}}(z(t), i(t), q) \tag{2}$$
$$o^{\mathcal{C}}(t) = g^{\mathcal{C}}(z(t), q)$$
$$z(0) = z_0(p)$$
$$\text{and } q \in \mathcal{Q}.$$

A parametrized system is denoted by $\mathcal{S}$:

$$\mathcal{S} : \dot{x}(t) = f^{\mathcal{S}}(x(t), i(t), p) \tag{3}$$
$$o^{\mathcal{S}}(t) = g^{\mathcal{S}}(x(t), p)$$
$$x(0) = x_0(p)$$
$$\text{and } p \in \mathcal{P}$$

Finally, given a system with no parameter freedom (e.g. $|\mathcal{P}| = 1$), we call the system unparametrized and denote it as $\overline{\mathcal{S}}$. The behavior of the system is then determined entirely by the set of inputs $\mathcal{B}_i$. We denote the unparametrized system obtained by fixing the parameter of a system $\mathcal{S}$ at some $p \in \mathcal{P}$ by $\mathcal{S}|p$.

Using these notions of behavior, we can now see explicitly what it means in our case for a system to satisfy a specification. Take a specification $\mathcal{C}$ parametrized by $\mathcal{Q}$. Take an unparametrized system $\overline{\mathcal{S}}$. The output o is

Figure 1. On the horizontal axis we have the set of input functions I^T, on the vertical axis the set of output functions O^T. The solid red lines bound the range of outputs o that can occur in the specification given a specific i. This range is parametrized by q in equation (2). An input output system defines a graph, as it maps a given input to an output. On the left, the system with the blue dashed graph satisfies the specification as, for every i there is a q such that the o of the specification matches the o of the system. On the right there are some inputs for which no such q exists, and the system does not satisfy the specification.

completely determined by the input i. $\overline{\mathcal{S}}$ satisfies the specification if for every input i there is a $q[i] \in \mathcal{Q}$, such that the specification system $\mathcal{C}$ will produce the output o. This is illustrated in figure 1.

Remark. Note that it is not necessary that there is one set of parameters q that matches the system behavior for all inputs for $\overline{\mathcal{S}}$ to satisfy the specification. In particular we do not require that the two ODEs transfer functions match. The behavioral condition is strictly weaker.

Given a system with parameters, the tuning problem we want to address is to find a p such that $\mathcal{S}|_p$ satisfies a specification. To do so we will now introduce several notions of distance between system and specification.

3. Distance to the specification

In practice it might often be impossible but also unnecessary to satisfy the specification exactly. This can mean both, that it is acceptable to fail for some inputs, or that the outputs are not exactly but only approximately the same. Our goal will be to get the system to be close to the specification. To formalize this idea we introduce a notion of the distance of a system to the specification.

The first important ingredient for this is a distance on the set of output functions, $\Delta(o_1, o_2)$. In what follows we will always take the square of the L2-norm:

$$\Delta(o_1, o_2) = \int_T \|o_1(t) - o_2(t)\|^2 \mathrm{d}t \tag{4}$$

As noted above, the outputs are functions of the inputs and the parameters. We will write $o[i, p]$ or $o[i]$ if no parameters are present. At fixed input i and parameter p, $o[i, p]$ is a function of time. Now given an unparametrized system $\overline{\mathcal{S}}$ and a specification $\mathcal{C}$ of the form (3) and (2) and a given input i we can consider

$$\min_q \Delta(o^{\overline{\mathcal{S}}}[i], o^{\mathcal{C}}[i, q]), \tag{5}$$

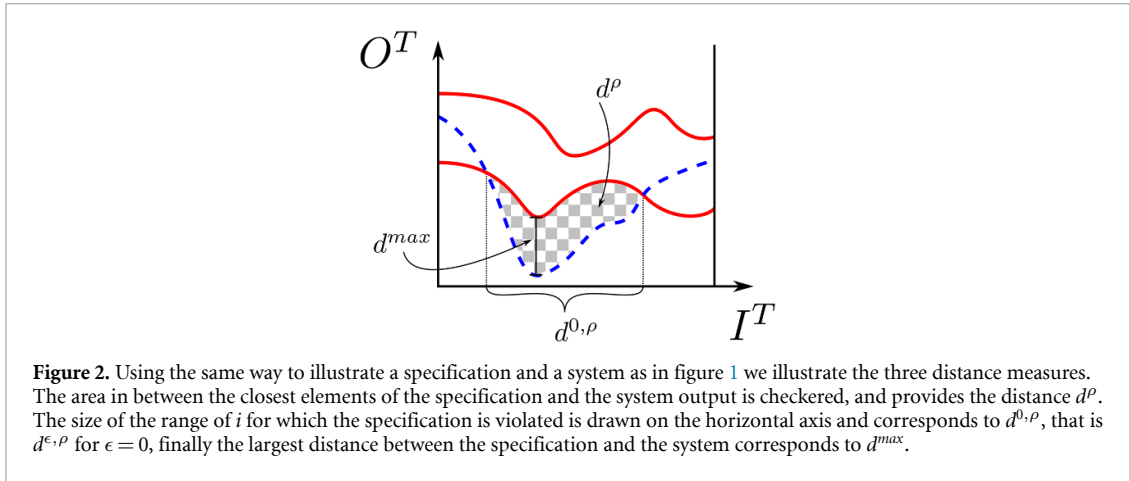as a distance of the system to the specification at input i.

In order to understand how much the outputs typically diverge we need to provide information on what inputs the system typically encounters. This can be formalized by providing a probability measure ρ on $\mathcal{B}_i$. The main distance measure we will investigate in what follows is then given by taking the average distance of the system to the specification in the sense of (5):

$$d^\rho(\overline{\mathcal{S}}, \mathcal{C}) = \int_{\mathcal{B}_i} \min_q \Delta\left(o^{\overline{\mathcal{S}}}[i], o^{\mathcal{C}}[i, q]\right) \rho \tag{6}$$

$$= E_\rho \left[\min_q \Delta\left(o^{\overline{\mathcal{S}}}[i], o^{\mathcal{C}}[i, q]\right) \right],$$

where E_ρ denotes the expectation value.

Note that $d^\rho(\overline{\mathcal{S}}, \mathcal{C}) = 0$ implies that the system satisfies the spec for all i up to a set of measure zero in ρ. Importantly, if the expectation is small we also know that the output of the system is close to an output allowable according to the specification for most inputs.

Figure 2. Using the same way to illustrate a specification and a system as in figure 1 we illustrate the three distance measures. The area in between the closest elements of the specification and the system output is checkered, and provides the distance d^ρ. The size of the range of i for which the specification is violated is drawn on the horizontal axis and corresponds to $d^{0,\rho}$, that is $d^{\epsilon,\rho}$ for $\epsilon = 0$, finally the largest distance between the specification and the system corresponds to d^{max}.

A second complementary distance measure that will be useful for validating the performance of a system in a more rigorous way is the fraction of inputs for which the distance of the system to the specification is larger than a margin ϵ. Let $\Theta_\epsilon : \mathbb{R} \to \{0,1\}$ be the shifted Heaviside step function with $\Theta_\epsilon(z) = 1$ if $z > \epsilon$ and 0 otherwise. Then we have:

$$
\begin{aligned}
d^{\epsilon,\rho}(\overline{\mathcal{S}},\mathcal{C}) &= \int_{\mathcal{B}_i} \Theta_\epsilon\left(\min_q \Delta\left(o^{\overline{\mathcal{S}}}[i], o^{\mathcal{C}}[i,q]\right)\right)\rho \\
&= E_\rho\left[\Theta_\epsilon\left(\min_q \Delta\left(o^{\overline{\mathcal{S}}}[i], o^{\mathcal{C}}[i,q]\right)\right)\right].
\end{aligned}
\tag{7}
$$

As ρ is a probability measure, $d^{\epsilon,\rho}$ will vary between 0 and 1. Note that $d^{\epsilon,\rho} = 0$ does not guarantee that the system satisfies the specification. It only implies that almost all inputs produce an output that differs from the specification by less than or equal to ϵ. Only for $\epsilon = 0$ do we guarantee that the specification is exactly satisfied for almost all i. Note that these two distances are genuinely complementary and do not coincide.

Finally we note that there is a natural distance that does not depend on a probability measure, but just on some norm on $\mathcal{B}_i$, namely the maximum distance to the specification given a fixed norm of the input:

$$
d^{\max}(\overline{\mathcal{S}},\mathcal{C}) = \max_{i:\|i\|=1} \min_q \Delta(o^{\overline{\mathcal{S}}}[i], o^{\mathcal{C}}[i,q]).
\tag{8}
$$

This notion of distance closely resembles an operator norm. Whereas d^ρ and $d^{\epsilon,\rho}$ capture the typical performance of the system, $d^{\max}$ is concerned with the worst case performance only. We will return to this in section 8 where $d^{\max}$ will allow us to connect and contrast the perspective taken here to conventional notions of the H_∞ operator norm and model reduction.

The three distances defined in (6)–(8) are illustrated in figure 2.

4. Sampling based approximations

The distances introduced above require a probability distribution on a set of inputs. Probability distributions on spaces of trajectories in time are called stochastic processes. The distances we introduced are thus given by expectation values of stochastic processes, and that in turn means they can be approximated by evaluating them on a sample $B_i \subset \mathcal{B}_i$, a set of realizations of the process.

Note that, while the various kinds of stochastic differential equations are the most familiar and best studied class of stochastic processes, they are not necessarily the best suited for the differential equation setting we study here. Instead we will make use of smooth random functions or random ODEs with smooth solutions [16]. We will give an example of this below.

Now given a sample B_i of realizations of the stochastic process ρ, with cardinality $|B_i|$, we can introduce the estimators

$$
\hat{d}^\rho = \frac{1}{|B_i|} \sum_{i \in B_i} \min_p \Delta(o^{\overline{\mathcal{S}}}[i], o^{\mathcal{C}}[i,q]),
\tag{9}
$$

for (6) and

$$\hat{d}^{\epsilon,\rho} = \frac{1}{|B_i|} \sum_{i \in B_i} \Theta_\epsilon(\min_q \Delta(o^{\overline{S}}[i], o^{\mathcal{C}}[i,q])), \tag{10}$$

for (7). It is challenging to understand the quality of the first estimator rigorously without further information on the summand. However, the summand for the second estimator is either 0 or 1. This means we can interpret this sample as a Bernoulli trial with randomly drawn i and 'success' defined as $\min_q \Delta(o^{\overline{S}}[i], o^{\mathcal{C}}[i,q]) \leqslant \epsilon$. The naive estimator $\hat{d}^{\epsilon,\rho}$ then approximates the true probability of the Bernoulli trial $d^{\epsilon,\rho}$. We can use standard methods [17] to obtain a center point correction and confidence interval for this estimator. For example, the 'add two successes and failures' 95% confidence interval is given in terms of $\tilde{n} = |B_i| + 4$ and

$$\tilde{d} = \frac{|B_i|\hat{d}^{\epsilon,\rho} + 2}{\tilde{n}},$$

by:

$$d^{\epsilon,\rho} \approx \tilde{d} \pm 2\sqrt{\frac{\tilde{d}(1-\tilde{d})}{\tilde{n}}}. \tag{11}$$

In practice it will typically not be possible to evaluate the minima in (10) exactly. Note however, that this leads to an overestimation of the distance. Thus some successes will be incorrectly counted as failures, and our estimator will tend to be conservative. The confidence interval is most accurately thought of as the confidence interval for an estimator for the probability that an upper bound on the distance $\min_q \Delta$ is below ϵ.

While it is easier to provide rigorous statements on the quality of the estimator for $d^{\epsilon,\rho}$, the estimate for d^ρ has the advantage of not requiring an appropriate choice of ε and of varying smoothly as the fit between system and specification changes. This makes it better suited as a basis for tuning the system using optimization techniques.

5. Tuning the system

Given a parametrized system of the form (3), we can formulate an optimization problem to find the set of parameters for which the system is the closest to the specification. Recall that d^ρ, as defined in (6), is the expectation value of a minimum. Optimizing it is thus a non-linear 2-stage stochastic programming problem. By using the approximation $\hat{d}^\rho$ we can explicitly give a discretized extensive form.

Recall that we denote as $\mathcal{S}|p$ the unparametrized system obtained by setting the parameters of $\mathcal{S}$ to p. Thus we want to find:

$$p_{tuned} = \arg\min_p \hat{d}^\rho(\mathcal{S}|p, \mathcal{C}). \tag{12}$$

In the expectation value underlying d^ρ there is a minimization for each input. To make this explicit, we will denote the parameter q of the specification $\mathcal{C}$ for a given input i as q_i. The inner optimizations are independent and we can exchange the order of the sum and minimization:

$$p_{tuned}$$
$$= \arg\min_p \frac{1}{|B_i|} \sum_{i \in B_i} \min_{q_i} \Delta(o^{\mathcal{S}}[i,p], o^{\mathcal{C}}[i,q_i])$$
$$= \arg\min_p \min_{\{q_i\}} \frac{1}{|B_i|} \sum_{i \in B_i} \Delta(o^{\mathcal{S}}[i,p], o^{\mathcal{C}}[i,q_i]). \tag{13}$$

This is as a joint parameter optimization in p and the set $\{q_i\}$ of a large differential equation with a trajectory based target function. This type of optimizations can be implemented in a straightforward manner using DiffEqFlux [5] in the Julia language [18]. Our implementation is available at https://github.com/PIK-ICoNe/ProBeTune.jl. Thanks to the ability to differentiate through ODE solvers a wide range of optimizers are available to perform this optimization. Crucially, the distance measure is designed in such a way that we can perform a joint optimization rather than having to perform an optimization of an optimum, a much harder problem.

6. Demonstration for a non-linear network

We will demonstrate the distances and their tuning, by considering two paradigmatic examples of complex non-linear dynamical networks connected at one node to the outside world. We will tune the complex network of identical dynamical nodes to react to outside inputs like a simple network of nodes of the same type, thus demonstrating that probabilistic behavioral tuning can aggregate complex networks. The two systems we will consider are a diffusively coupled network with tunable non-linear stabilizing forces and a system of Kuramoto oscillators with inertia. The coupling to the outside world is always at node 1 and we observe the reaction of the state at node 1 relative to the input.

6.1. Diffusive non-linear network

Consider the networked system $\mathcal{A}^N$ with N nodes denoted by $n = 1 \ldots N$:

$$
\begin{aligned}
\dot{x}_n &= -x_n - p_n x_n^3 + \sum_{m=1}^{N} A_{nm}(x_n - x_m) + \delta_{n1}(x_n - i), \\
o(t) &= i(t) - x_1(t) \\
x_n(0) &= 0 \\
p_n &\in \mathbb{R}^+
\end{aligned}
\tag{14}
$$

$$
\mathcal{B}_i = \left\{ i : i(t) = \operatorname{Re} \sum_{l=0}^{L} a_l e^{j(2\pi lt + \theta_l)} \right\},
\tag{15}
$$

for some fixed graph with adjacency matrix A_{nm}. The trajectories of this system are always bounded and therefore smooth solutions always exist. We can specify a probability distribution on $\mathcal{B}_i$ by specifying a probability for the parameters $a_l \in \mathbb{R}$ and $\theta_l \in [0, 2\pi)$.

Our goal now is to tune a full system with N nodes to behave, as far as the input-output relationship is concerned, like a two node system $\mathcal{A}^2$. In the concrete example we choose the Barabasi-Albert model [19] with $N = 10$ to generate A_{nm}, as such a scale-free network has a rich irregular structure.

Thus we have $\mathcal{S} = \mathcal{A}^{10}$ and $\mathcal{C} = \mathcal{A}^2$. In (13) this implies that we will jointly optimize over $(\mathbb{R}^+)^{10 + 2|B_i|}$. That is, the system $\mathcal{A}^N$ with N parameters and one copy of the system $\mathcal{A}^2$ for each element of the sample $B_i \subset \mathcal{B}_i$, with each such copy of the system being parametrized by two parameters.

6.2. Kuramoto oscillators system

The other system we consider is a network of Kuramoto oscillators with inertia with a tunable network. Thus while above we have a fixed adjacency matrix A_{ij}, here the network itself is part of the parameters and thus denoted p_{ij}. The system $\mathcal{K}^N$ is defined as:

$$
\begin{aligned}
\ddot{\phi}_n &= \Omega_n - p_n \dot{\phi}_n - K \sum_{m=1, m \neq n}^{N} p_{nm} \sin(\phi_n - \phi_m) + \delta_{n1} i, \\
&\quad n = 1 \ldots N, \\
o(t) &= i(t) - \phi_1(t) \\
\phi_n(0) &= 0 \\
p_n &\in \mathbb{R}^+ \\
p_{nm} &\in \mathbb{R}^+ \\
&1 < n, m < N \text{ and } m \neq n
\end{aligned}
\tag{16}
$$

$$
\mathcal{B}_i = \left\{ i : i(t) = \operatorname{Re} \sum_{l=0}^{L} a_l e^{j(2\pi lt + \theta_l)} \right\},
\tag{17}
$$

where K is the coupling constant, and Ω_n is the intrinsic node frequency drawn from a Gaussian distribution centered at 0 with width 1.

We have $\mathcal{S} = \mathcal{K}^{10}$ and $\mathcal{C} = \mathcal{K}^1$. Thus the specification is a single oscillator.

Table 1. Sequence of optimization steps in the tuning process of the diffusive non-linear system.

Tuning pipeline step	$\hat{d}^\rho$
Estimate $\hat{d}^\rho$ with 10 samples. This provides a q_i for each element of the sample.	8.9445
Tuning $\hat{d}^\rho$ using q_i from the previous step as initial parameters, 50 steps of ADAM(0.01) and 200 steps of AMSGrad(0.01)	1.0411
Estimating $\hat{d}^\rho$ on the original sample after tuning	0.882
Resampling the system, sample of the same size (10). Estimating $\hat{d}^\rho$ for the new sample.	1.356
Resampling the system, sample of size 100. Estimating $\hat{d}^\rho$ for the new sample to find initial values of the parameters.	1.072
10 repetitions of the following: 50 iterations of ADAM(0.01) and 200 iterations of AMSGrad(0.01)	0.199
Resampling the system, new sample of size 100. Estimating $\hat{d}^\rho$ for the new sample.	0.2053

7. Numerical results

We will now show the tuning for these systems. In both cases we begin by taking the system and the specification and determining the d^ρ of the untuned system at randomly guessed initial parameters p, and then tune the system to improve d^ρ. We will also study the question of whether the optimization is overfitting to the sample. Overfitting here means that for a small number of samples and a large number of parameters, it can occur that the parameters fit the specific sample, rather than the underlying distribution. To rule this out we resample after tuning the system, evaluating the final d^ρ achieved on a sample different from the one used to optimize the parameters.

Specific details of tuning the two example systems are presented below. Tuning is done using the Algorithms ADAM, AMSGrad and BFGS as implemented in the package DiffEqFlux.jl. By trying different combinations of optimizers and number of iterations we found that BFGS is best suited for low dimensional parameter spaces, as occur in the estimation of $\hat{d}^\rho$. The tuning is a high dimensional optimization problem and works well with gradient descent methods, such as ADAM and AMSGrad.

Step-by-step description of the tuning algorithm and the code can be found in the project repository at https://github.com/PIK-ICoNe/ProBeTune.jl. The version on which the results of this paper are based is published in [20]. The tuning was performed on a laptop with i7-8665U CPU. The typical runtime of the tuning is of the order of 10 min.
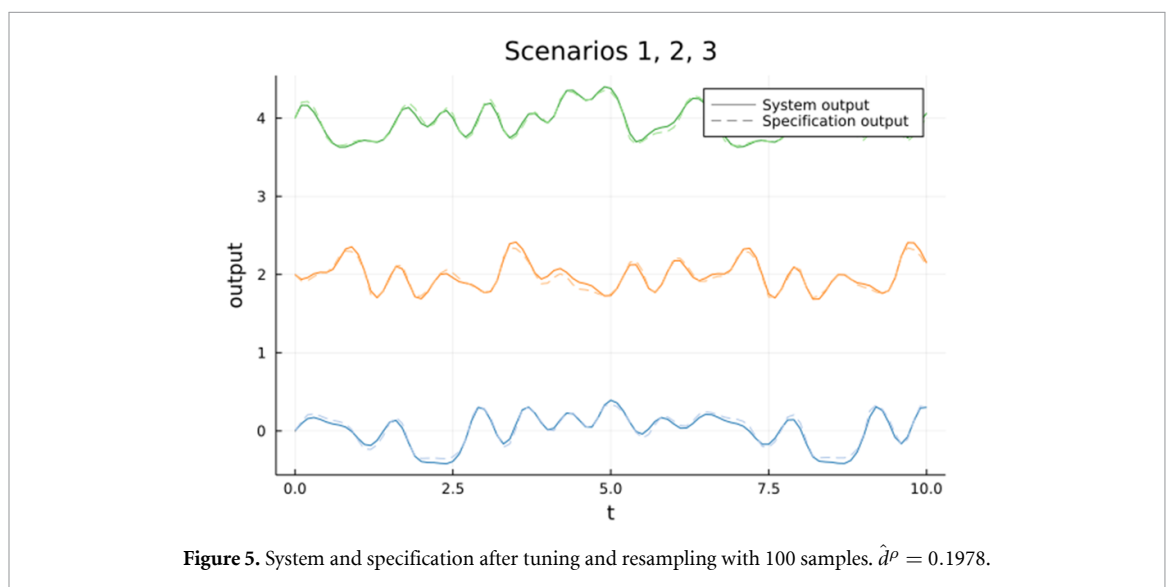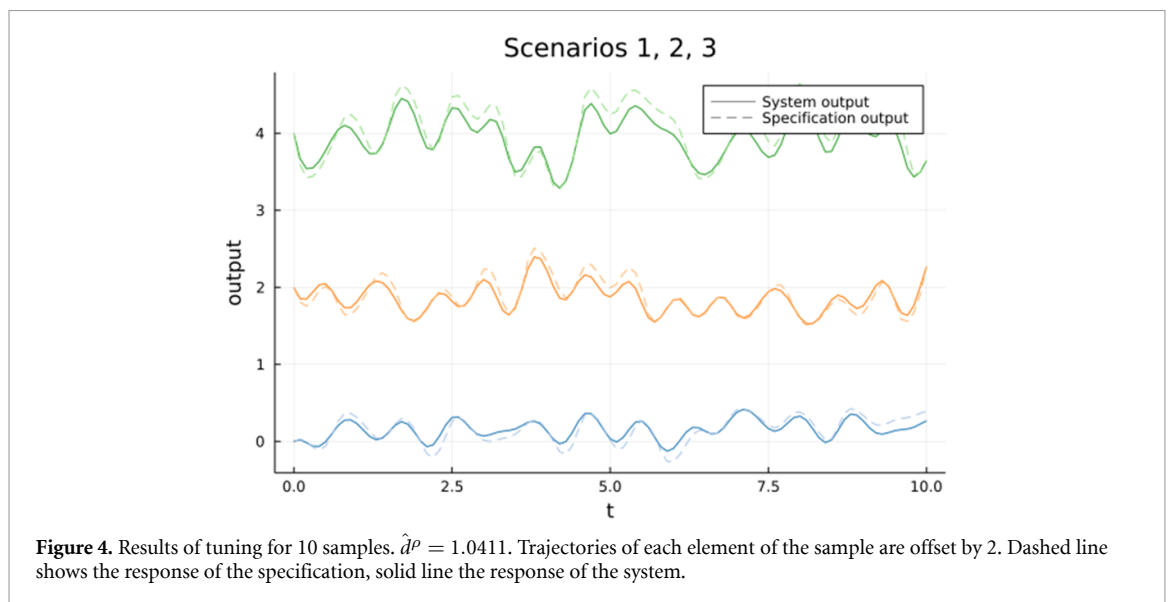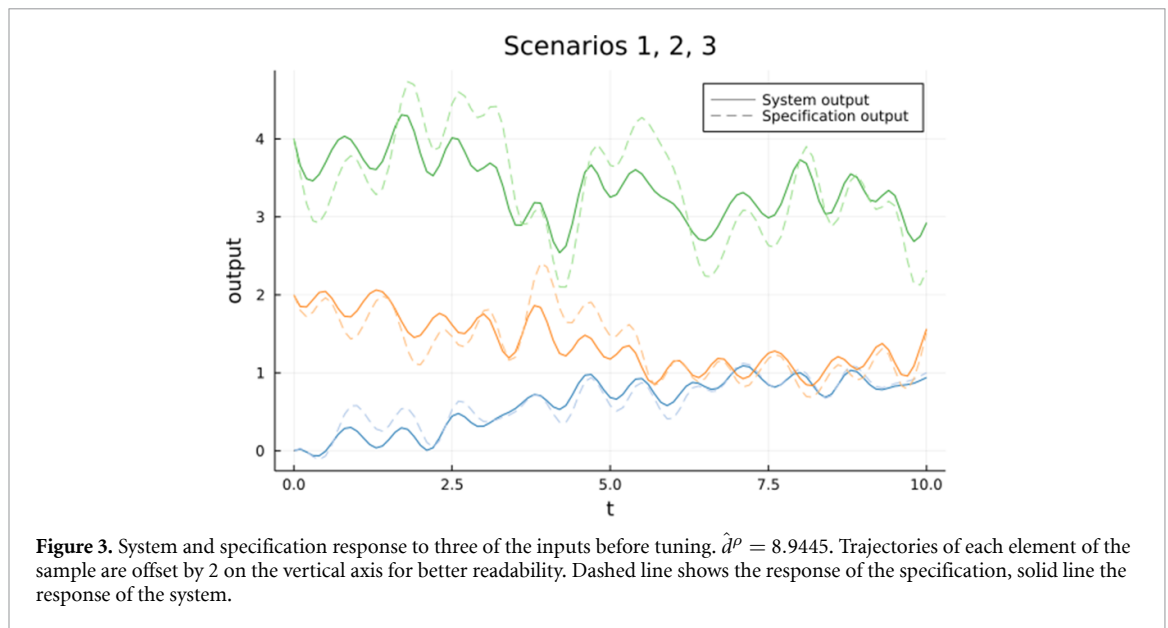
7.1. Diffusive non-linear network

Table 1 shows the schedule of estimating behavioral distance, resampling and tuning. We begin by estimating $\hat{d}^\rho$ for our initial guess of p with an input sample of size ten, i.e. $|\mathcal{B}_i| = 10$, obtaining $\hat{d}^\rho \sim 8.9$. The output of system and specification for three samples is shown in figure 3. Then we tune the system using this sample, and evaluate the distance again. This way we obtain a significant reduction in distance to the specification to $\hat{d}^\rho \sim 1.0$. Resampling shows that this included only moderate overfitting, the resulting outputs are shown in figure 4. This might seem surprising as we only have 10 samples but should be viewed in the context that each sample consists of a timeseries and contains considerable information.

Further tuning the system using a sample of 100 inputs shows a much further reduction of $\hat{d}^\rho$ to ~ 0.2 that persists after resampling. The outputs of the system and specification are now barely distinguishable, as shown in figure 5. Overall the L_2 distance of the difference of the output signals is a factor 45 smaller than in the untuned system.

To obtain a rigorous statement about the systems performance, we use the second distance measure (10). Rather than fixing a single ε, we can plot $\hat{d}^{\epsilon,\rho}$ and its 95% confidence interval, as a function of ε, similar to [21]. This is shown in figure 6. We thus can state that with 95% confidence, more than 80% of inputs produce a response that deviates less than 0.1 from the specification in the square L2 norm chosen.

7.2. Kuramoto oscillators system

The Kuramoto oscillators of equation (16) are considerably more complex, exhibiting oscillating behavior. The ability to make them behave like a single oscillator will crucially depend on whether they can maintain synchrony. This in turn is controlled by the spread of the intrinsic frequencies s. We will see that indeed a synchronous system is easier to tune to the specification using our method. To account for the fact that we
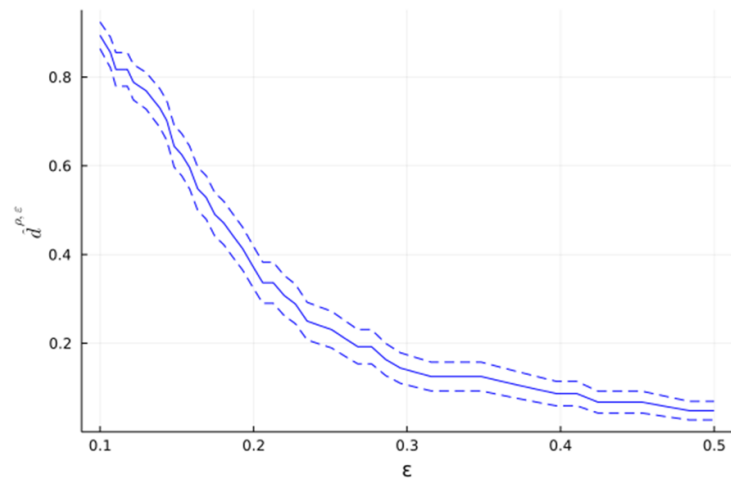
Figure 3. System and specification response to three of the inputs before tuning. $\hat{d}^{\rho} = 8.9445$. Trajectories of each element of the sample are offset by 2 on the vertical axis for better readability. Dashed line shows the response of the specification, solid line the response of the system.

Figure 4. Results of tuning for 10 samples. $\hat{d}^{\rho} = 1.0411$. Trajectories of each element of the sample are offset by 2. Dashed line shows the response of the specification, solid line the response of the system.

Figure 5. System and specification after tuning and resampling with 100 samples. $\hat{d}^{\rho} = 0.1978$.

Figure 6. $\hat{d}^{\epsilon, \rho}$ with 95% confidence interval.

Table 2. Sequence of optimization steps in the tuning process of the second order Kuramoto oscillators system.

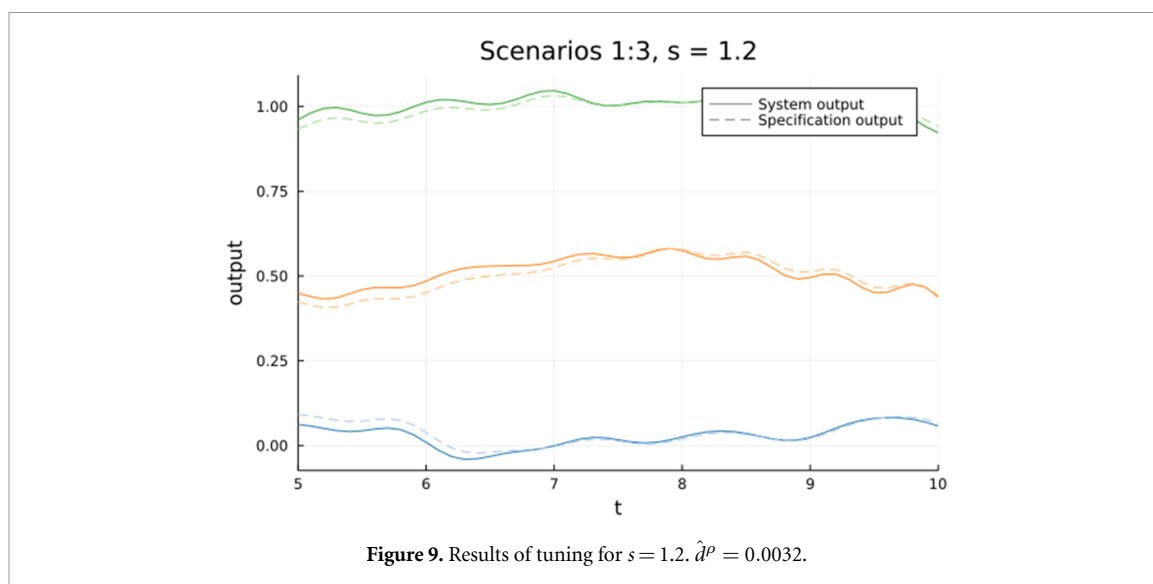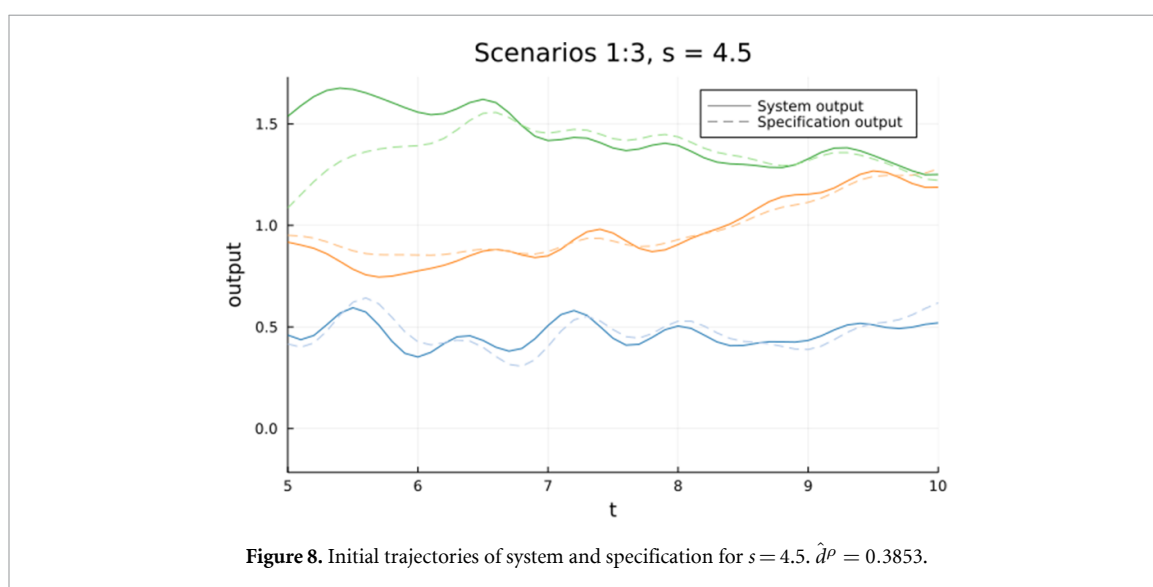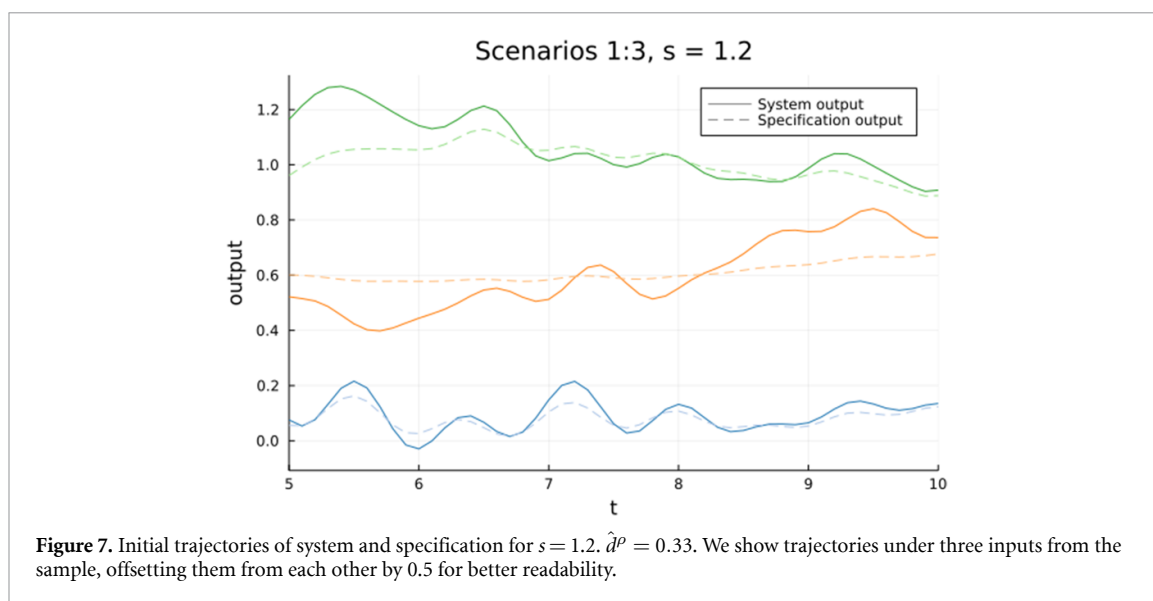| Tuning pipeline step	$\hat{d}^{\rho}$ for $s = 1.2$	$\hat{d}^{\rho}$ for $s = 4.5$
Estimate $\hat{d}^{\rho}$ with 10 samples. This provides a q_i for each element of the sample.	0.33	0.3853
Tuning $\hat{d}^{\rho}$, 100 steps of ADAM(0.01)	0.0389	0.1136
Tuning $\hat{d}^{\rho}$ further using BFGS	0.0122	0.0885
Resampling the system, sample of the same size (10). Estimating $\hat{d}^{\rho}$ for the new sample.	0.0093	0.113
Tuning $\hat{d}^{\rho}$ using BFGS	0.0066	—
Resampling the system, estimating $\hat{d}^{\rho}$ for the new sample.	0.0053	—
Tuning $\hat{d}^{\rho}$ further using BFGS	0.0032	—
Resampling the system, estimating $\hat{d}^{\rho}$ for the new sample.	0.0037	—

are not tuning initial conditions to a fixed point of the system we only considered the deviation of the system and specification after initial transients have subsided.
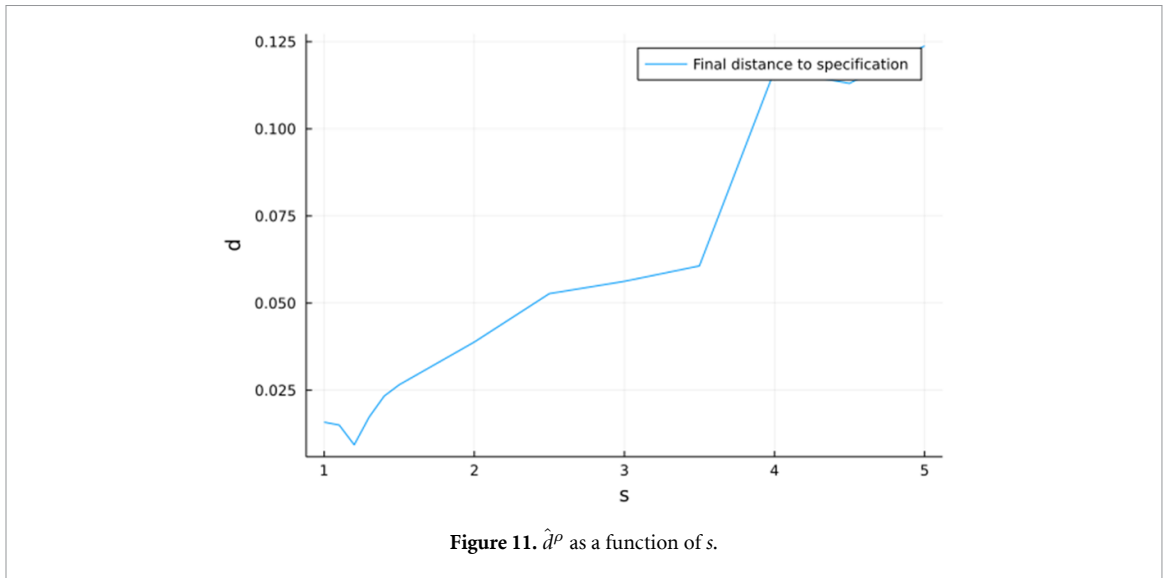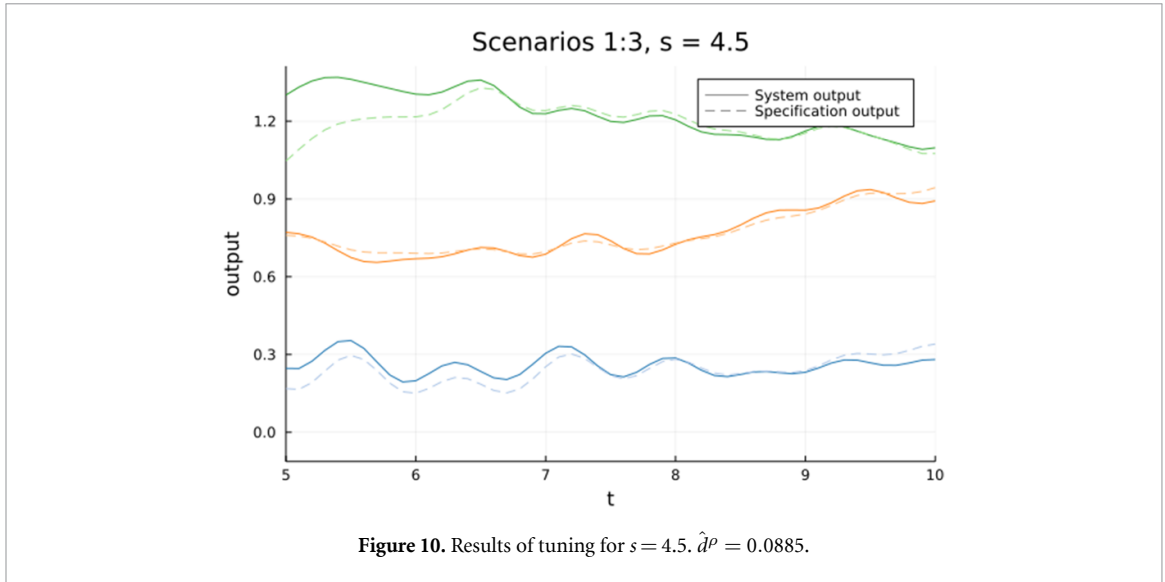
We used a fixed set of Ω_n with $\langle \omega_n \rangle = 0$ and $\omega_1 = 0$ and otherwise drawn from a Gaussian distribution of variance 1, scaling them by a factor s. At $s = 1$ the system exhibits robust synchrony, but at $s = 5$ the system does not fully synchronize. We tuned the system with s ranging from 1 to 5 with the same optimization parameters and input functions. We will show in more detail results for $s = 1.2$ and $s = 4.5$. Table 2 provides the detailed results of two optimization schedules. For $s = 4.5$ the further tuning using the BFGS optimizer failed.

Figures 7 and 8 show the untuned output behavior with $\hat{d}^{\rho}$ of ∼0.33 and ∼0.38 respectively.

Then we tune the systems using BFGS and ADAM algorithms. With the same schedule the respective values could be tuned down to ∼0.01 and ∼0.11. This shows that the synchronous behavior can be tuned much more easily in our case. Further tuning the system with $s = 1.2$ we could achieve another factor three improvement of the distance, leading to an overall reduction in the square L2 norm of the outputs of a factor of 100, as compared to a factor 3 for the system with wider frequency spread. Resulting output trajectories are shown in figures 9 and 10.

To compare the effect of system synchronicity on the tuning result, we show the final value of $\hat{d}^{\rho}$ for all studied systems from s = 1 to 5 in figure 11. For all systems we used the first four steps of the tuning pipeline of table 2 with identical parameters.

Figure 7. Initial trajectories of system and specification for $s = 1.2$. $\hat{d}^\rho = 0.33$. We show trajectories under three inputs from the sample, offsetting them from each other by 0.5 for better readability.

Figure 8. Initial trajectories of system and specification for $s = 4.5$. $\hat{d}^\rho = 0.3853$.

Figure 9. Results of tuning for $s = 1.2$. $\hat{d}^\rho = 0.0032$.

Figure 10. Results of tuning for $s = 4.5$. $\hat{d}^\rho = 0.0885$.

Figure 11. $\hat{d}^\rho$ as a function of s.

8. Relation to H_∞ model reduction

Above we defined several notions of distance between parametrized differential equations with the same inputs and outputs. These are in many ways comparable to operator norms of the difference of transfer operators of such systems. Such operator norms have been used extensively in control theory [22]. If the specification has no parameters and we deal with linear systems, our d^ρ is just the $H2$-norm of the difference of the systems transfer functions [23].

To more thoroughly illuminate the similarities and differences to such approaches we will discuss the relationship of our tuning to H_∞ model reduction in more detail. H_∞ model reduction relies on parametrized classes of transfer operators and optimizes a distance in the space of dynamical systems, and thus closely resembles some aspects of our approach here.

To begin with we first consider the non-probabilistic d^{max} distance introduced above as this can be related explicitly to the H_∞ operator norm, and its tuning to model reduction. To see this let us consider the case where both specification and system are given in terms of parametrized linear transfer operators in Laplace space $T(s)$, such that $o[i](s) = T(s)i(s)$.

The goal of H_∞ model reduction is the following. Given a system with the transfer operator T, find a reduced system T_{red} out of some class of systems, such that the difference in the induced operator norm is small. Parametrizing the reduced systems with q, and writing $T_{red}[q]$, we want to minimize the norm of the difference between the original and lower-order system (see e.g. [24] for an introduction):

$$\min_q \|T - T_{red}[q]\|_\infty = \min_q \max_{i:\|i\|_2=1} \|(T - T_{red}[q])i\|_2 \qquad (18)$$

We can show that the $d^{\max}$ distance of behaviors introduced in (8) is bounded by the H_∞ norm of the optimal reduction among the T_{red}. Take the space of possible reduced models to be the specification $\mathcal{C}$ and the system $\overline{\mathcal{S}}$ given by the transfer operators $T^{\mathcal{C}}$ and $T^{\overline{\mathcal{S}}}$. Then we have:

$$
\begin{aligned}
d^{\max}(\overline{\mathcal{S}}, \mathcal{C}) &= \max_{i:\|i\|_2=1} \min_q \|o^{\overline{\mathcal{S}}}[i] - o^{\mathcal{C}}[i,q]\|_2 \\
&= \max_{i:\|i\|_2=1} \min_q \|(T^{\overline{\mathcal{S}}} - T^{\mathcal{C}}[q])i\|_2.
\end{aligned}
\tag{19}
$$

This differs from the norm of the optimal model reduction by the order of the min and max. By the min-max inequality [25] we then have

$$
\begin{aligned}
d^{\max}(\overline{\mathcal{S}}, \mathcal{C}) &\leqslant \min_q \max_{i:\|i\|_2=1} \|(T^{\overline{\mathcal{S}}} - T^{\mathcal{C}}[q])i\|_2 \\
&= \min_q \|T^{\overline{\mathcal{S}}} - T^{\mathcal{C}}[q]\|_\infty.
\end{aligned}
\tag{20}
$$

Thus $d^{\max}$ is a lower bound of the quality of the optimal H_∞ model reduction of $T^{\overline{\mathcal{S}}}$ in the class $T^{\mathcal{C}}[q]$.

This comparison shows both the structural similarities and differences between standard model reduction, and our behavioral approach. Model reduction asks to have one particular simple system that behaves like the full system. Fulfilling behavioral specifications, in the setting of this paper, requires that given an input, such a system exists, but not that these simple systems are the same across all inputs. Requiring that we always have the same system is more difficult, thus the H_∞ distance to the best reduction in the class under consideration is bounded by our $d^{\max}$ to the class of reductions from below.

The tuning problem we consider is to tune the complex system $\overline{\mathcal{S}}$ towards better reducibility. However, both the d^{max} distance and the H_∞ norm require solving min-max problems that are difficult for non-linear and non-convex systems, and that might not be easy to approximate. Further, they are given by the behavior of the system given the most challenging input. Depending on the purpose of the tuning, a focus on the worst case might not be appropriate. The probabilistic distances we introduce above instead focus on the typical performance. Ignoring rare or weak failures of the specification is what allows these distances to be well approximated by sampling typical inputs. We trade a hard to track min-max problem for an easily approximated probabilistic estimation, at the price of having to provide a meaningful probability distribution on the set of inputs.

9. Discussion

In this paper we show how to combine probabilistic and behavioral concepts to provide novel distance measures that quantify how well a system conforms to a specification. Further, we demonstrate that these distance measures are well suited to tuning a complex system to a specification. Thus they enable us to aggregate a complex network into a vastly simpler specification. While they are probabilistic, we can give mathematically precise confidence intervals for the performance of the tuned system.

We demonstrate that the method can be efficiently used by implementing it in Julia, which has excellent library support for the type of optimizations required here. We use this implementation to successfully tune a diffusive non-linear networked system with 10 nodes to behave as a 2-node system, by jointly optimizing the system and one hundred copies of the specification over a sample of likely inputs, where each copy of specification corresponds to one possible input.

We also tune a system of 10 Kuramoto oscillators, making them behave in terms of input-output behavior as a single oscillator. This example is important for potential application of these ideas to power grids, as in that case the system is oscillating. The conceptual setup does not change, but the performance of the tuning depends on the system synchronicity. We have explored the impact of the intrinsic frequencies spread on the tuning result and find that synchronous systems are easier to tune to represent as a single oscillator, while for non-synchronous systems the same quality cannot be achieved.

The quality of tuning is evaluated using probabilistic notions. We count distances below and above a certain threshold ε as successes and failures respectively, interpreting it as a Bernoulli trial to provide a rigorous confidence interval. This leads to a stability curve similar to those underlying the approach in [21].

Finally we also discuss the relationship of this approach to model reduction and of our distances to H_∞ norms in this context. While the use cases are different, there are considerable structural similarities.

The method introduced here is in principle well suited to establishing novel control hierarchies in complex multi-modal systems. The motivating example being future renewable power grids. In this context, the challenge is to optimize parameters of an energy cell, where the input and output characterize the power

flow at the grid connection point. This sub-network should then be tuned to a specification that ensures that the continental scale system made of these cells is stable.

However, many open questions remain in order to employ this approach in such a realistic context, and the approach introduced here raises many new questions. Most importantly, in order to realize the potential of novel control hierarchies based on probabilistically satisfied specifications, we need to understand how to safely compose such specifications in a way that the guaranteed probabilistic properties are preserved.

We also focused here on specifications provided by parametrized differential equations. As noted in the introduction, the specifications for power grids usually given in terms of direct properties of the trajectory though. In this context it might be possible to explicitly solve for the specification compliant output that has the least distance to the system output. This would bypass the spec parameter optimization, and all our remaining concepts would carry through the same way.

Applying the method to non-linear systems relies on the ability of optimization algorithms to perform efficient searches in the parameter space. While the ability to differentiate through ODE solvers means that a wide variety of solvers are available for this task in Julia, the systems explored so far do not allow a comprehensive picture of their performance characteristics. Finally hand tailored optimization algorithms for this problem are also an intriguing possibility.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: https://doi.org/10.5281/zenodo.5759486 [20].

Acknowledgments

ORCID iD

Frank Hellmann ⓘ https://orcid.org/0000-0001-5635-4949

References

[1] Willems J C 1989 *Dynamics Reported* (Berlin: Springer) pp 171–269
[2] Willems J C and Polderman J W 1997 *Introduction to Mathematical Systems Theory: A Behavioral Approach* vol 26 (New York: Springer Science & Business Media)
[3] Moor T, Raisch J and Davoren J M 2003 *IFAC Proc. Vol.* **36** 349
[4] Rackauckas C and Nie Q 2017 *J. Open Res. Softw.* **5** 15
[5] Rackauckas C, Innes M, Ma Y, Bettencourt J White L and Dixit V 2019 arXiv:1902.02376
[6] Büttner A, Würfel H, Plietzsch A, Lindner M and Hellmann F 2022 *2022 Open Source Modelling and Simulation of Energy Systems (OSMSES)* (IEEE) pp 1–6
[7] Lindner M, Lincoln L, Drauschke F, Koulen J M, Würfel H, Plietzsch A and Hellmann F 2021 *Chaos* **31** 063133
[8] Plietzsch A, Kogler R, Auer S, Merino J, Gil-de Muro A, Liße J, Vogel C and Hellmann F 2021 arXiv:2101.02103
[9] ENTSO-E 2018 Rate of change of frequency (rocof) withstand capability (available at: www.entsoe.eu/network_codes/cnc/cnc-igds/)
[10] Borkowska B 1974 *IEEE Trans. Power Appar. Syst.* **PAS-93** 752
[11] Anders G J 1989 *Probability Concepts in Electric Power Systems* (New York: Wiley)
[12] Menck P J, Heitzig J, Kurths J and Schellnhuber H J 2014 *Nat. Commun.* **5** 1
[13] Hellmann F, Schultz P, Grabow C, Heitzig J and Kurths J 2016 *Sci. Rep.* **6** 1
[14] Hellmann F, Schultz P, Jaros P, Levchenko R, Kapitaniak T, Kurths J and Maistrenko Y 2020 *Nat. Commun.* **11** 1
[15] Liemann S, Strenge L, Schultz P, Hinners H, Porst J, Sarstedt M and Hellmann F 2020 *IEEE Madrid PowerTech 2021*
[16] Filip S, Javeed A and Trefethen L N 2019 *SIAM Rev.* **61** 185
[17] Agresti A and Coull B A 1998 *Am. Stat.* **52** 119
[18] Bezanson J, Edelman A, Karpinski S and Shah V B 2017 *SIAM Rev.* **59** 65–98
[19] Barabási A-L and Albert R 1999 *Science* **286** 509
[20] Zolotarevskaia E, Hellmann F, Würfel H and Büttner A 2021 PIK-ICoNe/ProBeTune.jl: Version 1.0 (v1.0) *Zenodo* https://doi.org/10.5281/zenodo.5759487
[21] Schultz P, Hellmann F, Webster K N and Kurths J 2018 *Chaos* **28** 043102
[22] Khalil I S, Doyle J and Glover K 1996 *Robust and Optimal Control* (Hoboken, NJ: Prentice Hall)
[23] Dragan V, Morozan T and Stoica A 2004 *Automatica* **40** 1103
[24] Chen B M 2013 *Robust and H-∞ Control* (New York: Springer Science & Business Media)
[25] Boyd S, Boyd S P and Vandenberghe L 2004 *Convex Optimization* (Cambridge: Cambridge University Press)