

Plug-and-play adaptive surrogate modeling of parametric nonlinear dynamics in frequency domain

Phillip Huwiler* Davide Pradovera[†] Jürg Schiffmann*

February 8, 2024

Abstract

We present an algorithm for constructing efficient surrogate frequency-domain models of (nonlinear) parametric dynamical systems in a *non-intrusive* way. To capture the dependence of the underlying system on frequency and parameters, our proposed approach combines rational approximation and smooth interpolation. In the approximation effort, locally adapted sparse grids are applied to effectively explore the parameter domain even if the number of parameters is modest or high. Adaptivity is also employed to build rational approximations that efficiently capture the frequency dependence of the problem. These two features enable our method to build surrogate models that achieve a user-prescribed approximation accuracy, without wasting too many resources in “oversampling” the frequency and parameter domains.

Thanks to its non-intrusiveness, our proposed method, as opposed to projection-based techniques for model order reduction, can be applied regardless of the complexity of the underlying physical model. Notably, our algorithm for adaptive sampling can be used even when prior knowledge of the problem structure is not available.

To showcase the effectiveness of our approach, we apply it in the study of an aerodynamic bearing. Our method allows us to build surrogate models that adequately identify the bearing’s behavior with respect to both design and operational parameters, while still achieving significant speedups.

Keywords: model order reduction, gas bearing, frequency domain, nonlinear dynamics, high-dimensional approximation.

1 Introduction

Frequency-domain models of (nonlinear) dynamical systems are often used by engineers in domains such as control systems, fluid mechanics, acoustics, optics, electronics and communications. A good understanding of nonlinear effects can be crucial for stability analysis,

*Laboratory for Applied Mechanical Design, Department of Mechanical Engineering, EPFL, 2000 Neuchâtel, Switzerland. Emails: phillip.huwiler@epfl.ch, jurg.schiffmann@epfl.ch.

[†]Department of Mathematics, University of Vienna, 1090 Vienna, Austria. Email: davide.pradovera@univie.ac.at.

performance evaluation, and mitigation of undesirable phenomena. Specifically, with the objective of resolving complex behaviors in physical or data-driven models, high-fidelity *parametric* models are developed. The parameters are embedded within the model as “tuning knobs” of the interesting inputs to the system.

Simultaneously, engineers are often not only unwilling to accept the exploding computational cost that accompanies an up-scaling in modeling complexity, but some applications also necessitate significant reductions in computing time and memory. This is particularly relevant in the so-called multi-query setting, where calculations have to be repeated for many different values of frequency and parameters, e.g., in the context of design optimization or of real-time monitoring and control. Model order reduction (MOR) addresses this by replacing the high-fidelity model with a cheap-to-evaluate surrogate one, which (hopefully) has a small approximation error. Evaluation efficiency is not the only concern when considering modeling alternatives: surrogate costs can come in the form of variable (so-called “*online*”) costs on a per-evaluation basis and overhead (so-called “*offline*”) costs for building the surrogate model. Examples of the latter include (re)-training and/or development effort associated with modifications in the reference model, which could be the case for techniques exploiting specific knowledge about the inner workings of the reference model.

We wish for the reader not to lose sight of the general applicability of our algorithm to frequency-dependent, nonlinear, and parametric dynamical systems, whose abstract mathematical formulation is outlined in Section 2. Yet, we are also practitioners and see value in showcasing our approach on a tractable engineering application. For this, we narrow-in on herringbone grooved journal bearings (HGJBs), a type of non-contact bearings where the pressure distribution in a thin film of gas separates a rotating shaft from a static bearing part. A reference HGJB model is summarized in section Section 4. Several numerical experiments with the HGJB are then performed in Section 5. We remind the reader that the details about the reference model are merely given to establish a context about the functioning and performance of the algorithm. They are not needed for its application. In this sense, our proposed algorithm can be thought of as *plug-and-play* in an approximation setting for *black-box* functions.

1.1 Model reduction background and contributions

By now, a vast body of works in MOR is available, and many different approaches have been developed for building surrogate models. In the specific case of nonlinear systems, some additional difficulties arise, especially when employing the most popular class of MOR techniques, namely, projection-based approaches such as proper orthogonal decomposition (POD) or reduced basis (RB) methods (see, e.g., [13] for an overview): due to the presence of nonlinear terms, it is difficult to produce a surrogate model that is both efficient and accurate. Strategies to address this issue have been studied under the umbrella term “hyperreduction”, with the most known approach being the discrete empirical interpolation method (DEIM) [6]. There are also two notable weaknesses of DEIM: (i) it is very problem-dependent and often requires considerable “user supervision” (limiting the extent to which it can be applied in an automated way), and (ii) it is *intrusive* (i.e., it requires knowledge of the nonlinearities and access to the high-dimensional system “state”, e.g., the pressure or displacement field that solves the nonlinear equation, cf. Section 2). This second item is a serious issue in

many practical settings, specifically for severe nonlinearities and whenever one cannot easily access the high-fidelity state, e.g., due to closed-source code. Even “less” intrusive methods like “non-intrusive RB” [11] or physics-informed neural networks [14, 27] are affected by this problem. Note that further difficulties arise for intrusive *parametric* MOR methods since dedicated techniques contingent on the nature of parameter dependencies are needed to adapt the bases for given parameter values, see, e.g., [8, 36].

In this work, we address this issue by developing a *non-intrusive* MOR approach for building surrogate models of quantities of interest (QoIs) arising in frequency-domain analyses of nonlinear dynamics. Our method builds the surrogate model from samples of the QoIs only, without requiring any access to –nor knowledge of– the underlying nonlinear model. This allows us to treat the high-fidelity model as a “black box”.

We build a surrogate model by interpolation. More precisely, a rational interpolation framework is employed with respect to frequency, whereas smooth interpolation (e.g., by piecewise polynomials or radial basis functions) is applied to approximate the dependence on parameters. The frequency dependence is modeled via the (less regular) class of rational functions as a way to address any potential lack of smoothness of the QoIs with respect to frequency. In this context, rational functions have been often used as a general-purpose approximation class [12, 19, 24, 29, 33], often yielding (quasi-)optimal approximation quality, regardless of whether such “non-smooth effects” are actually present or not. This being said, we refrain from using rational interpolation to model the joint dependence of the QoIs on frequency and parameters, because of the intrinsic instabilities and computational difficulties associated to multivariate rational approximation, which are ultimately due to the *curse of dimension*.

Lastly, we consider one additional (and, in our view, crucial) issue in MOR: the selection of the sample points. Given the task of building a surrogate model, one inevitably faces the questions of how many and where (in frequency-and-parameter space) to take the expensive high-fidelity samples that the surrogate will be eventually based on. Answering these questions well is crucial to guarantee accuracy and efficiency of the surrogate model, but also to limit the cost of the so-called “offline phase”, i.e., the assembly of the surrogate. Sampling is a particularly tricky issue in non-intrusive methods: the usual residual-based greedy strategies [13] cannot be employed, due to the absence of an error estimator. In practice, this hinders an efficient identification of points where the approximation error is large. In our method, we employ a fully non-intrusive approach for adaptive sampling, based on a predictor-corrector strategy: a modestly sized test set is used to assess the quality of the current surrogate, and also to determine where to add new sample points in case of insufficient accuracy.

In terms of novelty, an MOR approach similar to ours was already considered in [2] for generic parametric problems. In that work, a “fully smooth” approximation framework was used, relying on, e.g., piecewise polynomials or radial bases for building an interpolatory surrogate model with respect to both frequency and parameters. In this paper, we modify the method to account for the frequency component in the high-fidelity model, effectively specializing the approximation strategy from [2] to frequency-domain modeling. From a practical viewpoint, this requires combining the rational-in-frequency and smooth-in-parameters approximations. The rationale behind our proposed hybrid method is threefold: (i) rational function are better suited to address any potential lack of smoothness in the QoIs; (ii) ra-

tional functions can often achieve spectral approximation accuracy, much higher than more “local” function classes like piecewise polynomials; (iii) univariate rational approximation can be endowed with very efficient adaptive sampling strategies [25, 26], in contrast to the more costly approaches for high-dimensional adaptivity, e.g., based on locally adapted sparse grids [2]. Some numerical tests on the last two above-mentioned items are carried out in Section 5.1.

A hybrid rational/smooth approximation method was also proposed in [21], where, however, the focus was on tracking how singularities with respect to frequency (the *resonating frequencies* of the target system) evolve as functions of the parameter. In the present setting, such resonant behavior is generally absent (and, if present, is not too relevant for the sake of approximation accuracy), cf. Section 5, thus making the method from [21] misplaced. See Remark 2 for further discussion on this.

2 Frequency-domain problem framework

We start in the time domain and consider a time-evolving physical system, possibly subject to some input. Such a system may be modeled as a *parametric* time-invariant nonlinear ordinary differential equation system (ODES)

$$\frac{d\mathbf{x}}{dt}(t; \mathbf{p}) = \boldsymbol{\phi}(\mathbf{x}(t; \mathbf{p}); \mathbf{p}), \quad (1)$$

endowed with suitable initial conditions. The *state* $\mathbf{x} \in \mathbb{R}^n$ is a collection of n *degrees of freedom*, which characterize the instantaneous configuration (e.g., positions, velocities, temperatures, pressures, voltages, etc.) of the system. The right-hand side $\boldsymbol{\phi}$ drives the evolution of the state in response to both internal effects and external forces, including the system’s dependencies on a set of input parameters \mathbf{p} having a practical significance to the application. Typical examples of parameters are system design (e.g., geometry and material properties) and operational variables (e.g., external forces and boundary conditions). We note that one may be interested in more general systems, e.g., including higher-order time derivatives. Still, one can often cast such problems as first-order ODES like Eq. (1), e.g., by augmentation. While, for convenience, we proceed assuming a first-order structure, it is of no consequence to the applicability of our algorithm.

Simulating the time-evolution of the state \mathbf{x} , e.g., by discretizing Eq. (1) with time-stepping schemes, can be demanding, due to (i) the large size n of the state, (ii) the fine time-resolution needed to follow the (nonlinear) behavior of $\boldsymbol{\phi}$, and (iii) the potentially long simulation time to reach a steady state or periodic limit, if damping is low. These issues can, at least in part, be counteracted by moving to the frequency domain. To this aim, direct transformation of $\boldsymbol{\phi}$ to the frequency domain can be done by, e.g., Fourier or Laplace transforms (see, e.g., [22]). Notable for their frequent usage in solving engineering problems are a subset of perturbation-based methods that introduce small oscillatory excitations in $\boldsymbol{\phi}$ [20]. These perturbation techniques often leverage assumptions about the form of \mathbf{x} , e.g., (quasi-) periodicity, to aid the simplification. Finally, we note that, sometimes, the problem is cast in the frequency domain not just out of convenience, but because frequency-domain quantities are the ultimate objective of the analysis.

In a first-order perturbation method, one assumes that the state is a small *harmonic* perturbation of a stationary state \mathbf{x}_0 , namely,

$$\mathbf{x}(t; \mathbf{p}) \approx \mathbf{x}_0(\mathbf{p}) + \operatorname{Re}(\mathbf{x}_1(\omega; \mathbf{p})e^{i\omega t}). \quad (2)$$

The amplitude \mathbf{x}_1 of the perturbation term is complex to account for possible phase shifts. Note that, while \mathbf{x}_1 depends on the perturbation frequency $\omega \in \mathbb{R}$, \mathbf{x}_0 does not. Further harmonic perturbation terms with different frequencies may be added in the expansion of \mathbf{x} , giving rise to higher-order perturbation methods.

This expansion is then plugged into Eq. (1). Under the assumption of a small perturbation \mathbf{x}_1 , a first-order Taylor expansion of ϕ is then performed, resulting in a problem of the form

$$\begin{cases} \phi(\mathbf{x}_0(\mathbf{p}); \mathbf{p}) = 0, \\ \psi(\mathbf{x}_1(\omega; \mathbf{p}), \omega; \mathbf{x}_0(\mathbf{p}), \mathbf{p}) = 0. \end{cases} \quad (3)$$

The first equation characterizes stationary equilibrium states, whereas the second equation is frequency-dependent and identifies the perturbation amplitude *in function of the equilibrium* \mathbf{x}_0 . Additional terms in higher-order perturbation methods would be identifiable through further equations, related to higher-order terms in a Taylor expansion of ϕ .

In most applications, the target of the system's analysis is not the state \mathbf{x} , but rather a set of *quantities of interest* (QoIs), functions of \mathbf{x} whose choice is driven by the specific application. For instance, the QoIs may characterize the (averaged) boundary values of the system state (such as position and pressure values for a fluid-structure interaction [3]) at an interface or interconnection point with other systems. Other examples of QoIs are point evaluations of the state, to be compared with real-time sensor measurements in the context of, e.g., feedback control or system-health monitoring [1].

Generally, QoIs depend on the whole \mathbf{x} , i.e., in the context of Eq. (2), on both equilibrium state \mathbf{x}_0 and perturbation \mathbf{x}_1 . For this reason, QoIs are, in general, functions of both perturbation frequency ω and parameters \mathbf{p} , and their computation requires solving both equations in Eq. (3). We denote generic QoIs by

$$\xi : \mathbb{R}^{1+n_p} \ni (\omega, \mathbf{p}) \mapsto \xi(\omega, \mathbf{p}) \in \mathbb{C}^{n_\xi}. \quad (4)$$

Multiple QoIs are allowed, by letting $\xi(\omega, \mathbf{p})$ be vector-valued ($n_\xi > 1$).

This being said, for our later discussion, it is important to remark that, in some cases, QoIs may depend on the equilibrium state \mathbf{x}_0 and parameters \mathbf{p} only. For instance, in some applications, e.g., predator-prey or oscillator models, the equilibrium \mathbf{x}_0 is a QoI in itself. In a similar way, it is sometimes of interest to compute time-averaged quantities (over infinite time intervals), for which the oscillating component of Eq. (2) disappears. See Section 4 for further examples. The corresponding QoIs are functions of \mathbf{p} , but not of ω . Moreover, we note that evaluating these QoIs requires solving only the first equation in Eq. (3), and not the second one. We will denote such frequency-independent QoIs by

$$\zeta : \mathbb{R}^{n_p} \ni \mathbf{p} \mapsto \zeta(\mathbf{p}) \in \mathbb{C}^{n_\zeta}. \quad (5)$$

Again, $\zeta(\mathbf{p})$ is allowed to be vector-valued ($n_\zeta > 1$).

Remark 1 Above and in the rest of this work, we assume that we are able to compute QoIs at arbitrary values of ω and \mathbf{p} , by (i) expensively solving the frequency-domain problem in Eq. (3) and then (ii) post-processing the computed \mathbf{x} . In some situations, instead of solving a problem in the frequency domain, it is possible to obtain frequency-domain information by (discrete) Fourier transform of time-domain snapshots. For instance, this might be relevant in applications where a frequency-domain model is out of reach, e.g., because it is too complex to obtain or to simulate.

3 Approximation setup

Consider the abstract setup introduced in the previous section. We are interested in finding an approximation for the two following functions: the generic QoIs $\xi : \mathbb{R}^{1+n_p} \rightarrow \mathbb{C}^{n_\xi}$ and the frequency-independent QoIs $\zeta : \mathbb{R}^{n_p} \rightarrow \mathbb{C}^{n_\zeta}$, cf. Eqs. (4) and (5). (To keep our presentation as general as possible, we even allow n_ξ or n_ζ to be 0, in case only one type of QoIs is relevant.) We wish the respective approximations of ξ and ζ , which we denote by $\tilde{\xi}$ and $\tilde{\zeta}$ respectively, to be accurate for frequencies $\omega \in \Omega \subset \mathbb{R}$ and for parameters $\mathbf{p} \in \mathcal{P} \subset \mathbb{R}^{n_p}$, with Ω and \mathcal{P} being the (application-dependent) frequency and parameter ranges, respectively.

Due to its additional dependence on ω , the approximation of ξ requires special care. This behooves us to use a non-standard framework, taking inspiration from the “marginalized” approximation strategy described in [21]. The main idea is to treat ω and \mathbf{p} separately. First, one selects T distinct *collocation* parameter points $\mathbf{p}_1, \dots, \mathbf{p}_T \in \mathcal{P}$. At each collocation parameter \mathbf{p}_j , a *local* surrogate for ξ is built: by fixing $\mathbf{p} = \mathbf{p}_j$, we are left with the rather easier task of approximating a function of ω only. This step is carried out by *minimal rational interpolation* (MRI), cf. Section 3.1. We denote the corresponding approximation by $\tilde{\xi}_j : \Omega \rightarrow \mathbb{C}^{n_\xi}$, which should satisfy $\tilde{\xi}_j(\omega) \approx \xi(\omega, \mathbf{p}_j)$ for all $\omega \in \Omega$.

Then, the *global* surrogate is constructed by combining the local surrogates, ultimately resulting in an approximation of the form

$$\tilde{\xi}(\omega, \mathbf{p}) = \sum_{j=1}^T \tilde{\xi}_j(\omega) \varphi_j(\mathbf{p}). \quad (6)$$

In the expression above, for all j , $\varphi_j : \mathcal{P} \rightarrow \mathbb{R}$ is a weight function, responsible for combining the different local surrogates. Note that, in order to recover the local surrogates at all collocation points, i.e., to have

$$\tilde{\xi}(\omega, \mathbf{p}_j) = \tilde{\xi}_j(\omega) \approx \xi(\omega, \mathbf{p}_j) \text{ for all } \omega \in \Omega \text{ and all } j, \quad (7)$$

the weight functions should satisfy a Lagrange property: $\varphi_j(\mathbf{p}_j) = 1$ and $\varphi_j(\mathbf{p}_k) = 0$ if $j \neq k$. For instance, multivariate piecewise-linear “hat” functions and Lagrange polynomials are candidate families of weight functions.

Remark 2 As already mentioned, our approach differs from the one in [21] because our focus is on ξ , not on its poles and residues. More explicitly, we see from Eq. (6) that the

global surrogate $\tilde{\xi}$ is obtained by superposition of the local surrogates $\tilde{\xi}_j$. Instead, [21] suggests building $\tilde{\xi}$ through a superposition of poles and residues of $\tilde{\xi}_j$, namely,

$$\tilde{\xi}_j(\omega) = \sum_i \frac{\tilde{r}_j^i}{\omega - \tilde{\lambda}_j^i} \quad \rightsquigarrow \quad \tilde{\xi}(\omega, \mathbf{p}) = \sum_i \frac{\sum_{j=1}^T \tilde{r}_j^i \varphi_j(\mathbf{p})}{\omega - \sum_{j=1}^T \tilde{\lambda}_j^i \varphi_j(\mathbf{p})}.$$

In [21], this idea is justified by the approximation target having poles in the frequency range Ω , thus making an accurate identification of such poles a priority.

In our setting, we assume the approximation target ξ to be sufficiently smooth over Ω . Accordingly, the poles $\tilde{\lambda}_j^i$ lose importance and meaning. In effect, tracking the evolution of the poles with respect to \mathbf{p} would result in spurious and numerically unstable effects, whereas Eq. (6) is accurate and stable.

Constructing an approximation for ζ is usually easier, since ω is absent. In this work, for simplicity, we employ a structure that is rather similar to that of $\tilde{\xi}$ in Eq. (6), namely,

$$\tilde{\zeta}(\mathbf{p}) = \sum_{j=1}^T \zeta(\mathbf{p}_j) \varphi_j(\mathbf{p}). \quad (8)$$

Note that, while $\tilde{\xi}$ is obtained as a \mathbf{p} -superposition of *approximations* $\tilde{\xi}_j$, $\tilde{\zeta}$ is built as a \mathbf{p} -superposition of *exact* values $\zeta(\mathbf{p}_j)$. Also, we remark that, since the weight functions φ_j are assumed to satisfy a Lagrange property, we recover exact interpolation: $\tilde{\zeta}(\mathbf{p}_j) = \zeta(\mathbf{p}_j)$ for all j . We note that a rather similar MOR approach for building $\tilde{\zeta}$ is the one described in [2], where radial basis functions are used as weight functions for interpolation.

In general, it makes sense to choose the number T of collocation points differently for $\tilde{\xi}$ and for $\tilde{\zeta}$, e.g., if ξ depends on \mathbf{p} in a more (or less) “complicated” way than ζ does. This might even correspond to completely disjoint sets of collocation points \mathbf{p}_j for ξ and ζ , and to different weight functions φ_j . In presenting our approach, we mostly ignore this possibility to avoid an overly heavy notation. At the same time, we wish to mention here that our ultimate algorithm of choice, see Sections 3.2 and 3.3, usually employs different values of T for ξ and ζ .

Remark 3 Before proceeding, we wish to discuss the cost of evaluating the surrogates $\tilde{\xi}$ and $\tilde{\zeta}$ at new values of $(\omega \text{ and}) \mathbf{p}$. In the case of $\tilde{\xi}$, we see from Eq. (6) that we only need to evaluate T weight functions φ_j , which is very inexpensive, and T local surrogates $\tilde{\xi}_j$, which, as we will show, are also very cheap to evaluate. In particular, their evaluation cost is independent of the cost of solving the high-fidelity problem.

On the other hand, the evaluation of $\tilde{\zeta}$ requires evaluating the T weight functions, but also obtaining the T values of ζ at the collocation points. This second group of operations can be significantly sped up by employing an offline-online decomposition. More explicitly, since the T values $\{\zeta(\mathbf{p}_j)\}_{j=1}^T$ are independent of the new parameter \mathbf{p} at which $\tilde{\zeta}$ is to be evaluated, we can (expensively) compute them once and then store them. Afterwards, when $\tilde{\zeta}$ is to be evaluated, it suffices to cheaply access the values $\zeta(\mathbf{p}_j)$ from memory.

3.1 Minimal rational interpolation

Here we consider the problem of building a local surrogate $\tilde{\xi}_j : \Omega \rightarrow \mathbb{C}^{n_\xi}$ at some collocation point \mathbf{p}_j : the target is $\tilde{\xi}_j(\omega) \approx \xi(\omega, \mathbf{p}_j)$ for all $\omega \in \Omega$. For this, we employ the MRI method introduced in [24], which we summarize here in its barycentric formulation.

Assume that S_j *support* frequencies $\omega_{j,1}, \dots, \omega_{j,S_j} \in \Omega$ are fixed. Note the presence of the index j , a reminder that different collocation points may have different support frequencies. In fact, we will soon outline how such support frequencies can be *adaptively* selected, in a process that may yield different results at each collocation point \mathbf{p}_j .

Once the costly samples $\{\xi(\omega_{j,i}, \mathbf{p}_j)\}_{i=1}^{S_j}$ have been computed, we define the local surrogate as

$$\tilde{\xi}_j(\omega) = \sum_{i=1}^{S_j} \xi(\omega_{j,i}, \mathbf{p}_j) \underbrace{\frac{q_{j,i}}{\omega - \omega_{j,i}}}_{\phi_{j,i}(\omega)} \bigg/ \sum_{k=1}^{S_j} \frac{q_{j,k}}{\omega - \omega_{j,k}}. \quad (9)$$

The *rational* weight functions $\phi_{j,i}$ are defined in *barycentric* form, with $q_{j,1}, \dots, q_{j,S_j} \in \mathbb{C}$ being coefficients that will be defined in the next paragraph. Note that this choice of weight functions satisfies the Lagrange property for any (nonzero) $q_{j,i}$, so that, in particular, we always have exact reconstruction of the samples: $\tilde{\xi}_j(\omega_{j,i}) = \xi(\omega_{j,i}, \mathbf{p}_j)$ for all i .

The barycentric coefficients $q_{j,i}$ are computed by solving the minimization problem

$$\begin{aligned} \min_{q_{j,1}, \dots, q_{j,S_j}} \quad & \left\| \sum_{i=1}^{S_j} q_{j,i} \xi(\omega_{j,i}, \mathbf{p}_j) \right\| \\ \text{subject to} \quad & \sum_{i=1}^{S_j} |q_{j,i}|^2 = 1 \end{aligned} \quad (10)$$

where $\|\cdot\|$ denotes the Euclidean norm. We refer to [4] for more details on why the barycentric coefficients are chosen as the solution of this problem and on how to easily compute them via singular value decomposition. Alternative approaches for finding the weights $q_{j,i}$ could be based, e.g., on the Loewner framework or the adaptive Antoulas–Anderson algorithm [19].

The number S_j and the locations of the support frequencies $\omega_{j,i}$ are crucial for the approximation properties and efficiency of the MRI surrogate: too few samples yield an inaccurate approximation, whereas too many make the training of the local surrogate overly expensive, and can even result in numerical instabilities [26]. A practical and effective way of choosing the support frequencies is the *greedy* MRI (gMRI) method, introduced in [25].

In gMRI, the set of support frequencies is incrementally built, starting from the two endpoints of the interval Ω . Each new support frequency ω_{j,S_j+1} is selected based on the current surrogate $\tilde{\xi}_j$, specifically, based on the denominator of the weight functions $\phi_{j,i}$, cf. Eq. (9). The termination condition is based on a user-defined tolerance tol_ω , which determines the desired accuracy in frequency of the local surrogate.

3.2 Adaptive parameter sampling strategy

In this section, we restrict our attention to the approximation of ξ . More details on our strategy for building $\tilde{\xi}$ follow in Section 3.3. Also, for simplicity, from here onward we assume that the parameter range $\mathcal{P} \subset \mathbb{R}^{n_p}$ is a hypercube. This is the most common

situation in practice, since a certain range is usually prescribed independently for each of the n_p parameter components $\mathbf{p} = (p_1, \dots, p_{n_p})$.

Building a good and efficient approximation $\tilde{\xi}$ requires a careful choice of the T collocation points \mathbf{p}_j . To this aim, we take inspiration from the sampling approach described in [2, 21], adaptively selecting the collocation points by exploring \mathcal{P} via a *locally adapted sparse grid* (LASG) structure.

LASGs are able to explore hypercubes even in high dimension at a relatively quick pace. Some comparable methods are locally refined tensor grids and quasi-random (e.g., Halton points or Latin hypercube) sampling. See, e.g., [30]. However, tensor grids suffer greatly from the so-called “curse of dimension” and, as a consequence, are much more inefficient than LASG even for modest n_p (e.g., > 2). On the other hand, quasi-random sampling is very competitive even in the high-dimensional case, but does not allow for localized adaptive refinements. We refer to [2, 17, 21] for more details on LASGs.

In our sampling strategy, we rely on a non-intrusive predictor-corrector approach to “grow” the set of collocation points using a tree-like structure. Given a certain *training set* (i.e., the set of collocation points \mathbf{p}_j where $\tilde{\xi}_j$ is available), a disjoint *test set* is defined. For instance, the test set can be defined in a procedural way, as the “discrete LASG-neighborhood” of the training set. See [21, Section 3.4] for more details.

A *prediction* step is carried out at each test point \mathbf{p}_\star : the current surrogate model (which is based on training points only) is evaluated at $\mathbf{p} = \mathbf{p}_\star$. The discrepancy between surrogate and high-fidelity model at \mathbf{p}_\star is then quantified through the *prediction error* (see Eq. (11) below). If the prediction error is larger than a user-defined tolerance tol_ξ , the test point \mathbf{p}_\star is flagged for *correction*. This means that \mathbf{p}_\star is added to the training points, thus deflating the prediction error at \mathbf{p}_\star by the interpolation property in Eq. (7). The test set is then automatically refined around \mathbf{p}_\star for the next iteration. An example of the overall procedure is shown in Fig. 1 for $n_p = 2$. We can indeed observe local refinements of the LASG near points that are flagged for correction. The algorithm terminates when the prediction error is smaller than the tolerance tol_ξ at all test points.

Remark 4 *Within a prediction-correction iteration, each test parameter \mathbf{p}_\star is explored independently. This means that the above strategy is amenable to a parallel implementation.*

Above, by “prediction error” at some parameter point $\mathbf{p} = \mathbf{p}_\star$, we mean the following quantity:

$$e_{\star, \xi} = \sqrt{\frac{1}{S} \sum_{i=1}^S \frac{\left\| \tilde{\xi}(\omega_i^{\text{test}}, \mathbf{p}_\star) - \xi(\omega_i^{\text{test}}, \mathbf{p}_\star) \right\|^2}{\left\| \xi(\omega_i^{\text{test}}, \mathbf{p}_\star) \right\|^2}}. \quad (11)$$

Essentially, we look at the relative error between approximated and exact QoIs, taking the average (in the root-mean-square sense) over a set of user-defined “test frequencies” $\omega_1^{\text{test}}, \dots, \omega_S^{\text{test}} \in \Omega$.

Computing the prediction error by Eq. (11) might be rather expensive if S is modest or large, since S high-fidelity samples of ξ must be computed. To alleviate this issue, we replace the exact QoI at $(\omega_i^{\text{test}}, \mathbf{p}_\star)$ with a proxy, namely, an *ad hoc* approximation $\tilde{\xi}_\star(\omega_i^{\text{test}})$




Figure 1: Example of evolution of the LASG in the adaptive algorithm for 2 parameters $\mathbf{p} = (p_1, p_2)$. The components p_1 and p_2 are on the x- and y-axes, respectively. The black squares are the current training set (the collocation points). The red circles are the test set (the points where the prediction error is computed). Points flagged for correction at each step are circled in blue.

computed by gMRI at \mathbf{p}_\star . If $\tilde{\xi}_\star$ is a sufficiently good approximation of $\xi(\cdot, \mathbf{p}_\star)$ (notably, if it is closer to $\xi(\cdot, \mathbf{p}_\star)$ than $\tilde{\xi}(\cdot, \mathbf{p}_\star)$ is), the resulting approximated prediction error

$$\bar{e}_{\star, \xi} = \sqrt{\frac{1}{S} \sum_{i=1}^S \frac{\left\| \tilde{\xi}(\omega_i^{\text{test}}, \mathbf{p}_\star) - \tilde{\xi}_\star(\omega_i^{\text{test}}) \right\|^2}{\left\| \tilde{\xi}_\star(\omega_i^{\text{test}}) \right\|^2}} \quad (12)$$

can be expected to be close to the (expensive) exact one in Eq. (11).

It is useful to note that, if a test point \mathbf{p}_\star is flagged for correction, then the gMRI surrogate $\tilde{\xi}_\star$ that we propose to use in Eq. (12) has to be computed anyway, since it becomes one of the local surrogates in the expansion of the new $\tilde{\xi}$, cf. Eq. (6). With this observation, one can avoid repeated computations in the training phase.

Remark 5 *Since $\tilde{\xi}$ is built based on the local surrogates $\tilde{\xi}_j \approx \xi(\cdot, \mathbf{p}_j)$, it is crucial that such local surrogates be accurate enough. Otherwise, the prediction error might stagnate at an excessive value, resulting in a non-terminating \mathbf{p} -sampling loop. Importantly, if the local surrogates are computed by gMRI, we must set the tolerance tol_ω to a sufficiently smaller value than tol_ξ .*

3.3 Post-enrichment technique for frequency-independent QoIs

As already mentioned in Section 3, the approximation of ζ is fundamentally simpler than that of ξ . Indeed, no frequency modeling is necessary, and the approximation $\tilde{\zeta}$ consists of a linear combination of values of ζ at some collocation points, cf. Eq. (8). Still, computing the

values $\zeta(\mathbf{p}_j)$ for $j = 1, \dots, T$ is a rather expensive task. We can speed it up thanks to the following observation: if $\xi(\omega, \mathbf{p}_j)$ has been computed at some ω , this means that the high-fidelity problem in Eq. (3) has been solved at the corresponding frequency and parameters. In particular, the steady state $\mathbf{x}_0(\mathbf{p}_j)$ has already been computed. As such, since evaluating ζ only requires post-processing of \mathbf{x}_0 , the computation of $\zeta(\mathbf{p}_j)$ is very cheap, usually much less expensive than solving the high-fidelity problem from scratch.

Since our target is to obtain an approximation of *both* ξ and ζ , this allows us to employ the following strategy:

1. Build an approximation $\tilde{\xi}$, e.g., using the adaptive strategy from Section 3.2. This requires evaluating ξ at all the adaptively selected collocation points \mathbf{p}_j for $j = 1, \dots, T$, at the corresponding support frequencies $\omega_{j,1}, \dots, \omega_{j,S_j} \in \Omega$.
2. At every collocation point \mathbf{p}_j , compute also $\zeta(\mathbf{p}_j)$. As mentioned above, depending on the implementation, this step might be very cheap, if $\zeta(\mathbf{p}_j)$ can be obtained by post-processing of the already available steady state $\mathbf{x}_0(\mathbf{p}_j)$.
3. Check if the resulting $\tilde{\zeta}$, i.e., the surrogate obtained by combining the samples of ζ at the collocation points of $\tilde{\xi}$ (according to Eq. (8)), is accurate enough.
4. If $\tilde{\zeta}$ is not accurate enough, add more collocation points. (Note that such additional points affect only $\tilde{\zeta}$, and not $\tilde{\xi}$!)

The addition of new collocation points for $\tilde{\zeta}$ can be carried out by using the LASG approach from Section 3.2. However, in this case, since ω is absent, the prediction error, cf. Eq. (11), can be more easily computed as the relative approximation error

$$e_{\star, \zeta} = \left\| \tilde{\zeta}(\mathbf{p}_\star) - \zeta(\mathbf{p}_\star) \right\| / \left\| \zeta(\mathbf{p}_\star) \right\|. \quad (13)$$

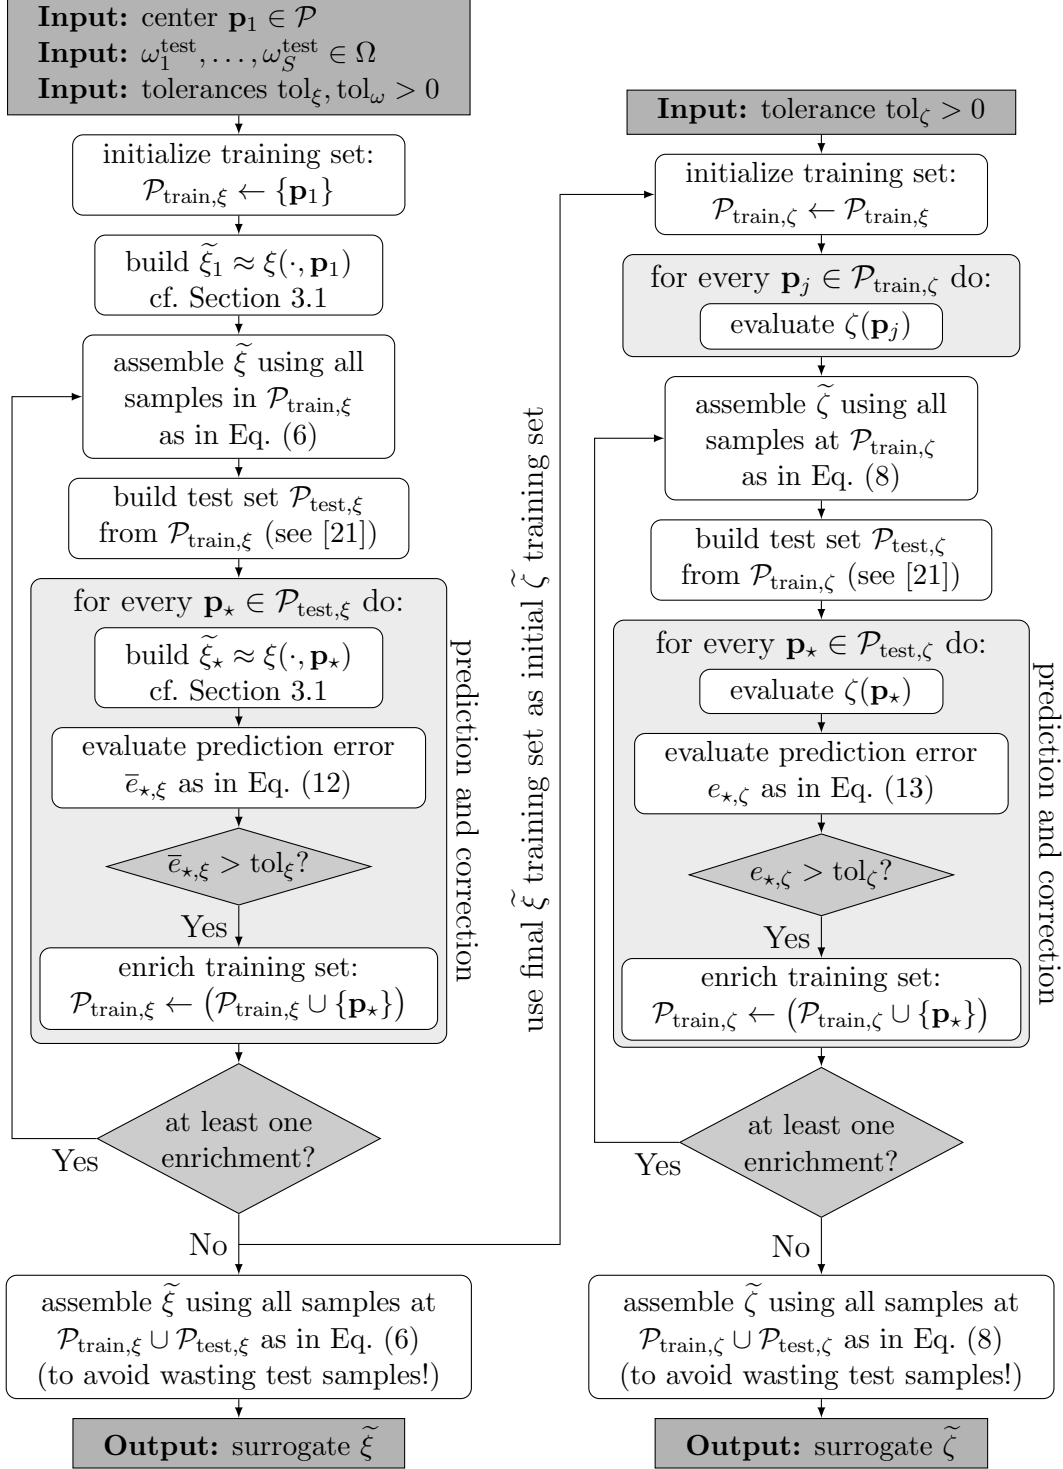
Note that this error measure greatly penalizes inaccuracies at locations where ζ is small, since the denominator above approaches zero there. In some applications, one might consider applying a “regularized” relative error, e.g.,

$$e_{\star, \zeta, \delta} = \left\| \tilde{\zeta}(\mathbf{p}_\star) - \zeta(\mathbf{p}_\star) \right\| / \max \{ \left\| \zeta(\mathbf{p}_\star) \right\|, \delta \},$$

for some (small) regularizing parameter $\delta > 0$. Similar considerations are relevant also for the error estimators used for ξ , namely, Eqs. (11) and (12).

Note that a complementary strategy exists, where the surrogate $\tilde{\zeta}$ is built first, and then $\tilde{\xi}$ is obtained by adding more collocation points. However, in our numerical testing, approximating $\tilde{\zeta}$ usually required more collocation points than $\tilde{\xi}$. This explains why we present the first- ξ -then- $\tilde{\zeta}$ method only. Of course, depending on the application, the complementary approach may be preferable.

We summarize the overall procedure in Fig. 2. In the diagram, we can see that, as described in Section 3.2, the training set grows from the center of \mathcal{P} into a LASG, by progressively adding any test point \mathbf{p}_\star with an excessive prediction error. Once the training

Figure 2: Adaptive parametric approximation of ξ , and of ζ by post-enrichment.

phase of $\tilde{\xi}$ ends, the final training set $\mathcal{P}_{\text{train},\xi}$ is selected as the initial training set $\mathcal{P}_{\text{train},\zeta}$ for $\tilde{\zeta}$, which is then subject to adaptive refinements.

We note that, while training both $\tilde{\xi}$ and $\tilde{\zeta}$, test points \mathbf{p}_\star with sufficiently small prediction errors might remain in the test set for several iterations of the predictor-corrector scheme. In such cases, the testing phase in Fig. 2 requires building the local surrogate $\tilde{\xi}_\star$ (or evaluating $\tilde{\zeta}(\mathbf{p}_\star)$) over and over in the prediction-correction phase. These expensive operations can be avoided by storing any such $\tilde{\xi}_\star$ (and $\tilde{\zeta}(\mathbf{p}_\star)$). The memory requirements for this are extremely low, thus resulting in a net time save with basically no extra costs.

In Fig. 2, we also employ the following trick to improve the accuracy of the surrogates. For both $\tilde{\xi}$ and $\tilde{\zeta}$, let us put ourselves at the end of the outer adaptivity loop, i.e., when the LASG reaches convergence. The prediction error, as encoded by Eqs. (12) and (13), has just been computed at all test points. This means that both the “test surrogates” $\tilde{\xi}_\star \approx \xi(\cdot, \mathbf{p}_\star)$ and the “test evaluations” $\tilde{\zeta}(\mathbf{p}_\star)$ are available. Such quantities can be included in the global surrogates $\tilde{\xi}$ and $\tilde{\zeta}$ by simply appending the test sets $\mathcal{P}_{\text{test}}$ to the training sets $\mathcal{P}_{\text{train}}$. This improves the accuracy of the two approximations with virtually no computational effort, since the added “expensive” values are already available in memory from the prediction-correction phase.

Before proceeding further, we also note that we provide an open-source MATLAB[®] version of the algorithm at https://github.com/pradovera/rational_lasg_matlab.

4 Frequency-domain bearing model

At this point, we are ready to explore the example of our approximation applied to the modeling of an aerodynamic bearing, specifically a *herringbone grooved journal bearing* (HGJB). Our focus in this and the next section is to establish a reference model in the framework of Section 2, referring readers to [9] for a broader overview of the modeling of grooved bearings.

A schematic representation of an HGJB is shown in Fig. 3 (top left). It consists of a segment of a larger rotor, rotating with angular speed ω_{rot} inside a bushing. A tight clearance between the parts (typically, on the order of microns) is occupied by the lubricating gas. The HGJB’s distinguishing feature is the V-shaped groove (and ridge) pattern. This acts as a viscous pump, self-pressurizing the gas film through its own rotation, by forcing fluid toward the axial center from both ends, where ambient conditions exist. As with other aerodynamic bearings, the relative motion of two bearing parts (bushing and rotor) in HGJBs generates a pressure distribution in a thin gas film between their surfaces, thus acting as a cushion to resist contact between the bearing parts. Such bearings are characterized by a low-friction and oil-free operation, enabling small-scale high-speed turbo-machinery that is well suited for distributed energy conversion technology, e.g., fuel cells [32] or waste-heat recovery [28], needed for a sustainable energy transition. Notably, compared to other aerodynamic bearings, HGJBs have the advantage of higher stiffness and load capacity. In applications, e.g., in the transportation sector, gas-bearing-supported rotors also see significant external loading, including shock and vibration. Predicting bearing behavior over a range of design and operational parameters is challenging, especially in the contexts of (near-)real-time stability or response prediction during machine operation.



ambient conditions:

 μ_a^* (dynamic viscosity) P_a^* (pressure)

boundary conditions:

— ambient periodic

For our purposes, the design and operation of the HGJB can be completely described in nondimensional terms. The translating rotor has a trajectory inside the bearing, whose position, depicted in the section view, is given in bushing-fixed Cartesian (x, y) -coordinates with the origin at bearing's center (Fig. 3). With the characteristic length in this frame is taken as the concentric ridge clearance $h_{r,nom}$, the magnitude of the rotor position vector cannot exceed unity, which would correspond to the (undesired) contact between the rotor and the bushing. We describe any bearing fluid film location using bushing-fixed cylindrical (θ, z) -coordinates, with z aligned with the rotor's rotation vector. The grooves are designed with a symmetric V-shape in the $\theta - z$ plane, parametrized by the angle of the groove p_β , the width ratio of the groove p_α , groove-ridge clearance ratio p_H , and the length of the bearing p_L , as depicted. It is worth mentioning that an eccentric rotor has θ -dependent local film thicknesses $h_r(\theta)$ and $h_g(\theta)$. Finally, the compressibility number p_Λ describes the combined effects of rotor speed and the fluid properties at a given operating point, and is of utmost importance for both design and operational purposes. In Fig. 3, we refer to a few characteristic dimensional quantities, for which we use the superscript “ \star ”, to ground our nondimensional description to a physical bearing.

4.1 Time-domain model

At present, modeling of generic aerodynamic bearings requires solving a suitable adaptation of the Reynolds equation, the classical governing equation describing the pressure distribution within the fluid film. More specifically, in the case of HGJBs, the groove-ridge interfaces in the θ -direction lead to discontinuities in the local film thickness $h(\theta)$ (see Fig. 3). This makes it necessary to employ a very fine spatial (and temporal) meshing in numerical implementations, resulting in high simulation costs.

As an alternative, we employ the narrow groove theory (NGT), pioneered in the 1950s by Whipple [34, 35]. The NGT allows a convenient modification of the Reynolds equation for bearings with periodic grooving features in the θ -direction. Two assumptions lead to the NGT result [7]: (i) locally incompressible fluid, giving a triangular pressure profile across a groove-ridge pair, and (ii) an infinite number of grooves, which transforms the saw-toothed pressure profile into a smooth one. The NGT equation describes the smooth pressure distribution with a second-order nonlinear partial differential equation given by

$$2p_\Lambda \partial_t (a_5 P) = \partial_\theta (P (a_1 \partial_\theta P + a_2 \partial_z P)) + \partial_z (P (a_2 \partial_\theta P + a_3 \partial_z P)) + a_s (\sin(p_\beta) \partial_\theta (a_4 P) - \cos(p_\beta) \partial_z (a_4 P)) - p_\Lambda \partial_\theta (a_5 P), \quad (14)$$

where P is the pressure field, while the coefficients $a_d(\theta)$ with $d = 1, \dots, 5$ and a_s are defined in Appendix A, along with the nondimensionalization scheme used. Refer to Fig. 3 for definitions of parameters \mathbf{p} . As boundary conditions, we impose ambient pressure at the two axial edges of the bearing.

After spatial discretization of the fluid film and numerical approximation (typically, by finite-difference or -element methods) of the spatial partial derivatives, we can transform Eq. (14) into the form of the first-order ODE system in Eq. (1). We refer to [15] for more details on the discretization. In our simulations below, we discretize the fluid film on a 65×65 rectangular grid in the (θ, z) -plane using a centered-differences scheme.

In the interest of studying the rotordynamics, the time-domain ODE system can be augmented with the equations of motion for the relevant inertial components. Specifically, in a fluid-structure-interaction fashion, our rotor is subject to the pressure of the fluid film, so that the total force is obtained by integration of P , to be added to any external forces (e.g., due to gravity).

4.2 Frequency-domain model

To obtain a frequency-domain version of the bearing model, we follow [16, 10]. Consider again Fig. 3 with the rotor in an arbitrary equilibrium position (x_0, y_0) due to a static loading (e.g., inertial forces, gravity, etc.). The rotor is then perturbed to yield simple harmonic motion (with frequency ω) around the equilibrium position. Independent perturbations in the x and y directions are allowed, with magnitudes Δ_x and Δ_y respectively, which are complex to account for phase shifts. Specifically, the local film thickness takes the time-harmonic form

$$h_r(\theta, t) \approx 1 - (x_0 + \operatorname{Re}(\Delta_x e^{i\omega t})) \cos(\theta) - (y_0 + \operatorname{Re}(\Delta_y e^{i\omega t})) \sin(\theta). \quad (15)$$

As a consequence of this expansion, it is assumed that the total pressure solution takes the form

$$P(\theta, z, t) \approx P_0(\theta, z) + \operatorname{Re}((\Delta_x P_{1,x}(\theta, z) + \Delta_y P_{1,y}(\theta, z)) e^{i\omega t}). \quad (16)$$

Note the presence of two independent perturbation terms $P_{1,x}$ and $P_{1,y}$, whose sum (weighted by Δ_x and Δ_y) gives the total first-order pressure perturbation. As the notation suggests, they correspond to perturbations of h_r in the x and y directions, respectively.

After injecting Eqs. (15) and (16) into Eq. (14) and keeping only terms up to order $\mathcal{O}(\Delta_x + \Delta_y)$, we can split the resulting equation into three separate equations by collecting and balancing coefficients of Δ_x and Δ_y (see [10] for the detailed equations):

- A nonlinear steady-state equation, with no analytical solution, is given by 0th-order terms. Its solution is the equilibrium point P_0 in Eq. (16), whose discrete version is indicated by \mathbf{x}_0 in Eq. (3). To find it, we use the Newton method after spatial discretization, and this accounts for a significant portion of the computation time.
- Two ω -dependent linear equations contain first-order perturbation terms, corresponding to Δ_x and Δ_y , respectively. Their solutions are the ω -dependent first-order perturbation fields $P_{1,x}$ and $P_{1,y}$. These equations are linear and thus, after spatial discretization, can be solved more directly by an off-the-shelf linear system solver. Still, it is required to inject the equilibrium \mathbf{x}_0 , which must therefore be precomputed.

Note that varying the parameters \mathbf{p} will generally affect all the above-mentioned equations. As such, for each parameter value of interest, it is necessary to assemble and solve the nonlinear equation, and then assemble and solve the linear equations at each excitation frequency ω of interest.

Remark 6 *The frequency-domain formulation of the bearing described above is not merely a matter of modeling convenience. The typical axisymmetric bearing geometry and (quasi-)periodic nature of system excitations (e.g., rotor unbalance, external vibrations, self-excited whirl) contribute to strongly frequency-dependent responses. This observation is supported by experimental studies for the HGJB, e.g., in [18, 5].*

Until now, we have not yet specified how to quantify the “bearing behavior” that we wish to model with respect to the parameters \mathbf{p} . We do this through specific QoIs, which are used to assess the bearing’s static and dynamic performance. These QoIs are extracted from the fluid film pressure solution via a post-processing step. For assessing the bearing’s dynamic performance near its static equilibrium, we model the rotor as a 2D nonlinear harmonic oscillator with complex direct and cross-coupled impedances computed by integration (practically, by numerical quadrature) of the perturbed pressure field as

$$\mathbf{Z} = \begin{bmatrix} Z_{xx} & Z_{xy} \\ Z_{yx} & Z_{yy} \end{bmatrix} = - \int_{-p_L}^{p_L} \int_0^{2\pi} \begin{bmatrix} P_{1,x} \cos \theta & P_{1,x} \sin \theta \\ P_{1,y} \cos \theta & P_{1,y} \sin \theta \end{bmatrix} d\theta dz \quad (17)$$

Decomposing entries of \mathbf{Z} into real and imaginary parts ($\mathbf{Z} = \mathbf{K} + i\omega\mathbf{C}$) leads us directly to eight real coefficients, interpreted as the dynamic stiffness and damping of the bearing. The evolution of the impedance matrix with respect to frequency ω and design/operating parameters \mathbf{p} can be directly used to obtain critical information needed in harmonic-response calculations and in stability analyses. When considered in the context of more complex rotordynamic systems, e.g., with multiple bearings and/or system degrees-of-freedom, these

bearing dynamic coefficients can also be integrated into the respective impedance matrices of the larger rotordynamic system for more sophisticated performance analyses.

Moreover, when the system is at the static equilibrium at (x_0, y_0) , the force transmitted to the rotor due to the static pressure distribution exactly offsets the external static load applied. Having solved the equilibrium equation for the static pressure, we can obtain the entries of the static force vector by integration (in practice, quadrature) of the rise in the static pressure as

$$\mathbf{f} = \begin{bmatrix} f_x \\ f_y \end{bmatrix} = - \int_{-p_L}^{p_L} \int_0^{2\pi} (P_0 - 1) \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix} d\theta dz. \quad (18)$$

The load that the bearing can support at equilibrium provides practitioners with useful information on the bearing's static performance. Specifically, a large magnitude of the static force \mathbf{f} means high bearing load capacity, which is generally desirable to keep the rotor operating safely.

A high-level process summary for obtaining the QoIs from the reference model is included in Fig. 4.

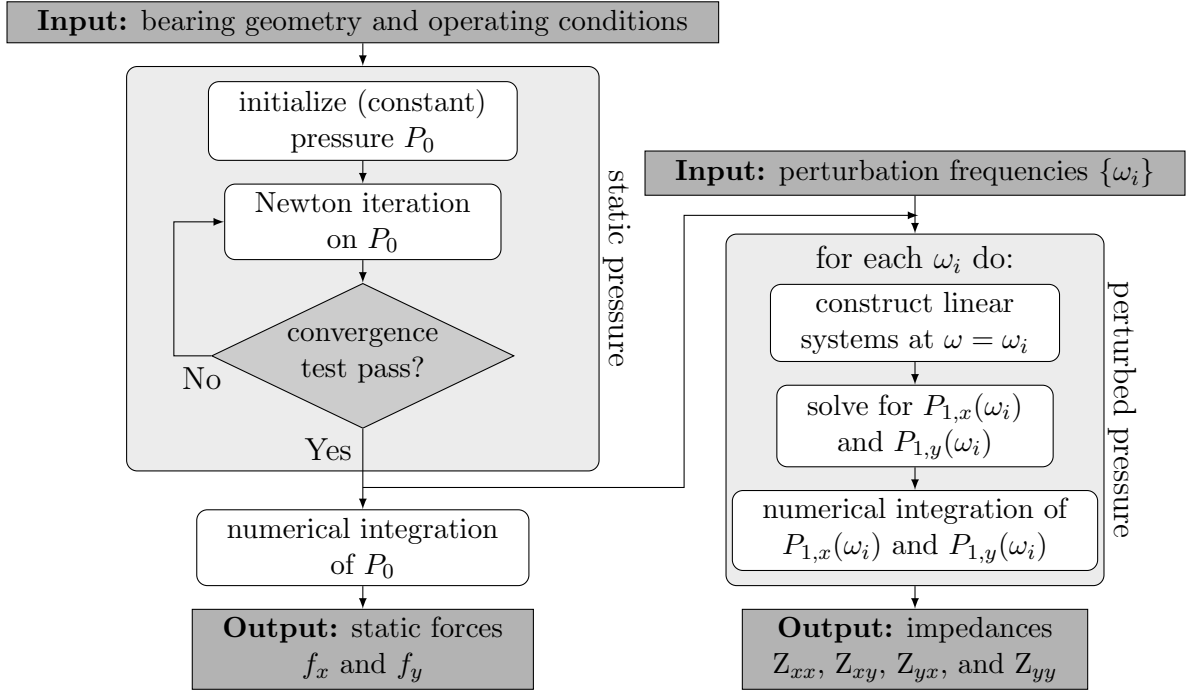


Figure 4: Eccentric NGT HGJB reference model numerical procedure.

Remark 7 Above, to obtain a frequency-domain formulation, we have perturbed the NGT equations around a generic equilibrium rotor position. In the literature, it is common to find perturbation methods that act around the bearing's concentric position only. Such models, which we refer as “concentric NGT” (cNGT) models, trade accuracy to gain computational efficiency. Specifically, cNGT models predict the bearing stiffness and damping properties (cf. Eq. (17)) as functions of perturbation frequency ω and parameters \mathbf{p} without consideration of static eccentricity effects. Such a concentric analysis generally yields a pessimistic

stability and load capacity estimation. An efficient implementation of cNGT for an HGJB was provided in [31], and the rudimentary timings of this model are compared with our surrogate in Section 5.

5 Numerical results

In many applications (e.g., in optimal design or real-time control), one needs efficient and accurate prediction tools for bearing behavior in response to a variety of sources of excitation. However, the nonlinear frequency-domain model described in the previous section is often too expensive to be directly applied in such multi-query settings. We propose to replace the high-fidelity and high-cost model with a surrogate one, obtained with our proposed adaptive strategy from Section 3. In the following sections, some experiments are carried out in this direction, for several choices of parameters \mathbf{p} . Our numerical tests were carried out in MATLAB[®] R2022b on a desktop computer with a 6-core 3.50 GHz Intel[®] processor.

In all cases, the QoIs are the impedance parameters and the static forces, cf. Eqs. (17) and (18). More precisely, we seek approximations for the frequency-and-parameter-dependent

$$\xi := [Z_{xx}, Z_{xy}, Z_{yx}, Z_{yy}]^\top : \mathbb{R}^{1+n_p} \rightarrow \mathbb{C}^4$$

and for the parameter-dependent

$$\zeta := \mathbf{f} = [f_x, f_y]^\top : \mathbb{R}^{n_p} \rightarrow \mathbb{R}^2.$$

While we test several input variable sets, perturbation frequency and eccentricity are common inputs in all our simulations, because they are the very essence of the eccentric perturbation frequency-domain model we wish to approximate, cf. Section 4. In our formulation, ω represents the nondimensional ratio between the (dimensional) excitation frequency and the rotor speed (see Appendix A). Considering the consequential nature of an unbalanced rotor, it is of practical relevance to ensure that the ω domain of interest includes the rotor frequency, i.e., in nondimensional terms, $\omega = 1$. Concerning eccentricity, for convenience we consider equilibrium eccentricities only in the right direction, i.e., $x_0 \geq 0$ and $y_0 = 0$ in Eq. (15). This can be done without loss of generality due to the bearing’s symmetry.

5.1 (1+1)D numerical results

For our first numerical experiment, we consider a frequency $\omega \in \Omega = [10^{-5}, 1.701]$ and a single parameter $\mathbf{p} := x_0 \in \mathcal{P} = [0, 0.95]$, the eccentricity in the x direction. The remaining parameters of the bearing model are set to their nominal values, which are reported in Table 1. Since Ω contains several orders of magnitude, we space the $S = 1001$ test frequencies $\omega_1^{\text{test}}, \dots, \omega_S^{\text{test}}$ logarithmically over Ω . We set $\text{tol}_Z = \text{tol}_f = 10^{-3}$ and $\text{tol}_\omega = \text{tol}_Z/3$, and we prescribe piecewise-linear “hat basis functions” for \mathbf{p} -interpolation (φ_j in Eqs. (6) and (8)).

Our algorithm samples the parameter space $\Omega \times \mathcal{P}$ adaptively. The resulting sample points are shown in Fig. 5 (left). With respect to frequency, we can observe that, except for the left endpoint $\omega = 10^{-5}$ (which is always sampled by gMRI), all the sampled frequencies are clustered at large frequencies, in the interval $[0.1, 1.701]$. Also, we note that the number and locations of the sampled frequencies depend on the value of x_0 .

Table 1: Nominal values and ranges of bearing parameters in numerical simulations. See Section 4 for a more detailed description of the parameters and of their significance in the HGJB model.

parameter	nominal value	range	
		min	max
ω	—	10^{-5}	1.701
x_0	0	0	0.95
p_Λ	25	1	40
p_α	0.65	0.3	0.7
p_β	19°	-50°	-10°
p_L	1.0	0.7	2.0
p_H	3.2	1.0	4.0

In Fig. 5 (right), we can see the \mathbf{p} -sample points that are added at each greedy iteration. For the first 5 iterations, we can observe uniform refinements, whereas, at the final iteration, localized refinements are performed around small values of x_0 . This means that the approximation of \mathbf{Z} is particularly difficult there, a fact that was unknown *a priori*.




Figure 5: On the left: (ω, x_0) sample points. On the right: x_0 sample points as the \mathbf{Z} predictor-corrector loop advances. The LASG tree structure is superimposed as dotted lines.

To validate our approximation, we evaluate $\tilde{\mathbf{Z}}$ on a uniform grid of 100×50 (ω, x_0) -points. There, we also expensively compute the exact \mathbf{Z} . The approximated $\tilde{\mathbf{Z}}$ is shown in Fig. 6, where we focus on two stiffness entries ($\mathbf{K}_{xx} = \text{Re}(\mathbf{Z}_{xx})$ and $\mathbf{K}_{xy} = \text{Re}(\mathbf{Z}_{xy})$) and on two damping entries ($\mathbf{C}_{xx} = \text{Im}(\mathbf{Z}_{xx})/\omega$ and $\mathbf{C}_{xy} = \text{Im}(\mathbf{Z}_{xy})/\omega$). We show the absolute error $|\tilde{\mathbf{Z}} - \mathbf{Z}|$, which is quite small in all cases, namely, below 10^{-3} .

As a final test to investigate the behavior of the employed LASGs, we look at the prediction error indicators e_\star in Eqs. (11) and (13) over the parameter space \mathcal{P} . In Fig. 7 (left), we evaluate e_\star on a grid of new test points, which do not fall within the set of support parameters. We see that the error values are uniformly below the prescribed tolerance. Moreover, the gap between tolerance and actual error is rather small: for most values of x_0 , the error is only around an order of magnitude smaller than the tolerance. This is actually a very

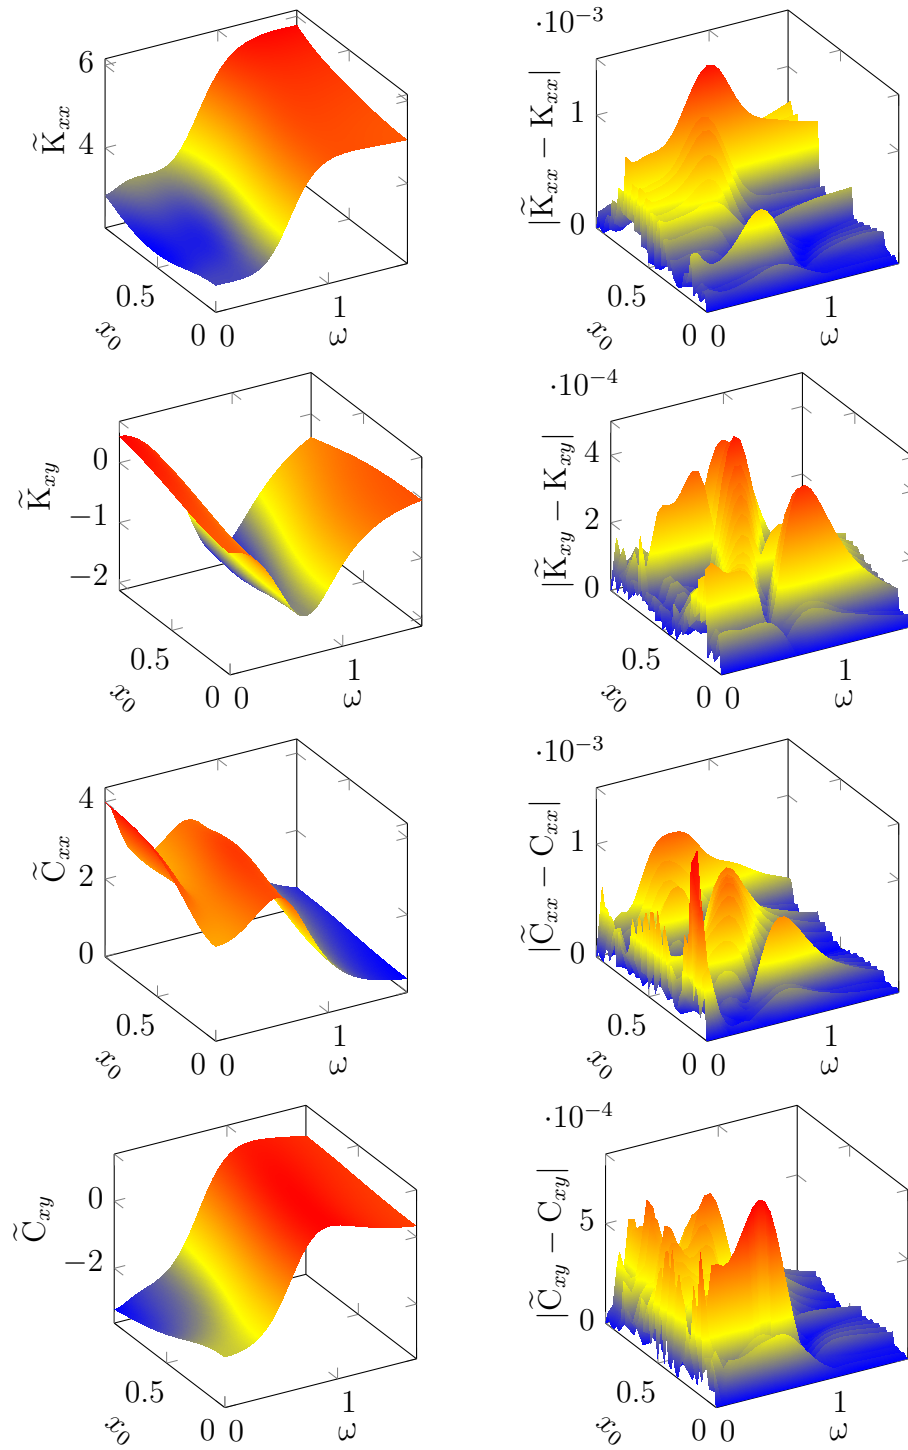


Figure 6: Predicted values and errors for some stiffness and damping entries.

beneficial property. Indeed, if the gap between tolerance and prediction error were larger, it would be an indication that our algorithm takes too many samples in the training phase, resulting in less efficient training, and also in more costly evaluation of the surrogate model.




Figure 7: Validation errors for impedance and static forces. From left to right, the approximations \tilde{Z} and \tilde{f} are obtained with piecewise-linear, polynomial, and spline \mathbf{p} -interpolation. The tolerances are $\text{tol}_Z = \text{tol}_f = 10^{-3}$, as denoted by the horizontal dotted line. On top of both plots, the locations of x_0 sample points are shown as crosses, with dots denoting points reserved for \tilde{f} .

Additionally, we recall that our algorithm is “aware” of such indicators only for values of x_0 in the training and test set. Moreover, the impedance prediction error is available only in “surrogate” form, cf. Eqs. (11) and (12). Nonetheless, the prescribed tolerance seems nicely attained over the whole parameter range.

5.1.1 Comparing more (and less) global approximation strategies

We now want to discuss some of the consequences of our choice of interpolation strategy over the parameter x_0 . We recall that, in order to obtain our results above, we have employed piecewise-linear hat functions. Such basis is local, in the sense that each local value, i.e., \tilde{Z}_j for \tilde{Z} and $\mathbf{f}(\mathbf{p}_j)$ for \tilde{f} , affects the approximation only around the collocation point \mathbf{p}_j . This leads to some advantages: it can be used to approximate general non-smooth (but continuous) functions and it is cheap to evaluate. However, it also has one main disadvantage: we are building an approximation that is locally only linear with respect to x_0 , which may not be too accurate.

Under smoothness assumptions, global bases can be employed to gain accuracy. To show this empirically, we repeat our experiments by replacing hat functions with global Lagrange polynomials. The training phase of our algorithm ends much earlier: only 17 samples of x_0 are needed to build the approximation, compared to the 37 needed with the previous setup. Despite the lower number of sample points, the approximation error (quantified by $e_{\star,Z}$ and $e_{\star,f}$) is fairly similar to the previous case for most values of x_0 , as can be seen in Fig. 7 (middle).

However, we can clearly notice some Runge oscillations near both endpoints of \mathcal{P} , where the tolerance seems not to be attained. These numerical instabilities are intrinsically related to polynomial interpolation over scattered points. Specifically, (locally refined subsets of)

uniformly spaced points may have extremely large Lebesgue constants [23]. As such, we could expect our algorithm, endowed with global polynomial \mathbf{p} -interpolation, to fail whenever the collocation points are unfavorably placed. Indeed, in such case, the numerical instabilities might cause the tolerance to never be attained, by introducing oscillations near the endpoints of the parameter domain.

With splines [23], one can sometimes achieve higher accuracy than with local bases (like hat functions) and better stability than with global bases (like polynomials). To test this, we repeat our simulation, using cubic splines to interpolate over \mathbf{p} . The corresponding algorithm reaches convergence after exploring 17 \mathbf{p} -points, which is about half as many as with hat functions. The corresponding accuracy (again, quantified by $e_{\star,z}$ and $e_{\star,f}$) is shown in Fig. 7 (right). It is uniformly below the tolerance, without showing any unstable behavior. This being said, splines cannot be easily generalized to higher dimensions ($n_p \geq 2$) if the sampling set is unstructured. Due to adaptivity, this makes splines unavailable in our method as soon as more than one parameter is present. In effect, this means that piecewise-linear hat functions remain the natural choice for a robust interpolation framework over general LASGs.

5.1.2 Validation of a gMRI+LASG approximation

As a final test involving the (1+1)D model, we compare our method to the state-of-the-art non-intrusive adaptive algorithm from [2]. There, LASGs are used to adaptively sample over *all* parameters, i.e., in our framework, over frequency and parameters alike. Effectively, this means that the rational approximation routine via gMRI is skipped, at the cost of increasing the number of parameters by 1. In comparing the two approaches, we refer to our method as “hybrid”, since it combines rational and hat functions, whereas we dub the reference state-of-the-art method “joint”.

To make the comparison more fair, in the joint method, we treat frequency on a linear scale, as opposed to the logarithmic one used for the hybrid approach. (Using a logarithmic scale would exacerbate the sampling issues that we are about to describe.) Moreover, since the joint approach does not require a “frequency tolerance”, we use the larger “tol” value as tolerance for both frequency and eccentricity.




Figure 8: (ω, x_0) points sampled by the “joint” method.

The joint greedy method requires 8 iterations to converge, as opposed to the 6 needed by our hybrid approach. In turn, this corresponds to a set of 622 collocation (ω, x_0) -points (at 61 distinct eccentricity values), many more than the 287 (ω, x_0) -points (at 37 distinct eccentricity values) used by the hybrid approach. We show the collocation points of the joint method in Fig. 8, which should be visually compared to Fig. 5 (left), containing the samples of the hybrid approach. We can conclude that the joint method is less effective in achieving the target tolerance, when compared to our hybrid approach. In our view, there are two main reasons for this:

- When used to approximate the frequency dependence of the QoIs, hat functions, due to their locality and lower degree, converge more slowly than rational functions. From this point of view, our proposed hybrid method improves the joint one by increasing the approximation efficiency in one of the parameters, namely, frequency. Notably, in contrast to our observations concerning polynomials in Section 5.1.1, the chosen higher-order approximation strategy in frequency is numerically stable, due to the beneficial properties of the barycentric expansion, cf. Eq. (9).
- The greedy MRI method performs adaptive sampling more effectively than LASGs. In particular, both gMRI and LASGs use a predictor-corrector approach, evaluating the current approximation error over a set of test parameter points. However, the size of the test set used in gMRI is smaller, consisting of a *single* point at each iteration. This enables a quicker adaptivity routine, which is more parsimonious in terms of total number of samples taken. (Of course, the main limitation of gMRI is its being limited to the approximation of *univariate* functions.)

5.2 (1+2)D numerical results

We consider a slightly more complex setup with operational relevance. The compressibility number p_Λ is added as a parameter, on top of frequency ω and static eccentricity x_0 : $\mathbf{p} := (x_0, p_\Lambda)$. The effect of a static loading on the bearing is captured by x_0 , while p_Λ envelopes the combined effects of the specific fluid conditions and rotor speed at a given operating point. See Table 1 for the ranges of ω , x_0 , and p_Λ , as well as for the nominal values of the remaining bearing parameters.

In this section, we specifically wish to investigate the effect of the tolerances (tol_z , tol_f , and tol_ω) on the behavior of the algorithm, as well as on the quality of the resulting surrogate model. For simplicity, we fix the relationships $\text{tol}_z = \text{tol}_f =: \text{tol}$ and $\text{tol}_\omega = \text{tol}/3$, so that we only have one free tolerance value to vary. Each value of tol requires an independent training, and yields a different surrogate model. We train 6 surrogates, corresponding to tolerance levels between 10^{-3} and $2.7 \cdot 10^{-2}$. For reference, in Fig. 9 we display the sets of training \mathbf{p} -points that are explored for some values of tol . We can observe anisotropic refinements near small values of p_Λ and mid-to-large values of x_0 .

We assess the accuracy of our approximations by evaluating the surrogates and the exact model on a 30×30 uniform grid of (x_0, p_Λ) -points. The frequency-independence of \mathbf{f} affords us the convenience of plotting the results. In Fig. 10, we visualize the exact f_y , the y -component of \mathbf{f} , and the evolution of the corresponding surrogate absolute error for selected values of tol . We can observe that the error peaks, in absolute terms, are driven down across




Figure 9: Adaptively chosen training \mathbf{p} -points for three values of tol , which are reported on top of the plots. The white dots are added by the \mathbf{f} predictor-corrector loop.




Figure 10: On the left: f_y computed with the high-fidelity training model. On the right: absolute errors $|\tilde{f}_y - f_y|$ for surrogates obtained at different tol levels. The same z -scale and color scale are used for all plots on the right-hand side.

the grid as tol is decreased: adding collocation points improves the surrogate approximation in these critical areas of the domain.

Further, we look at the trade-offs between training cost, online evaluation cost, and desired accuracy of a surrogate model for our current application. In Fig. 11, we show two different kinds of scaling with respect to tol .

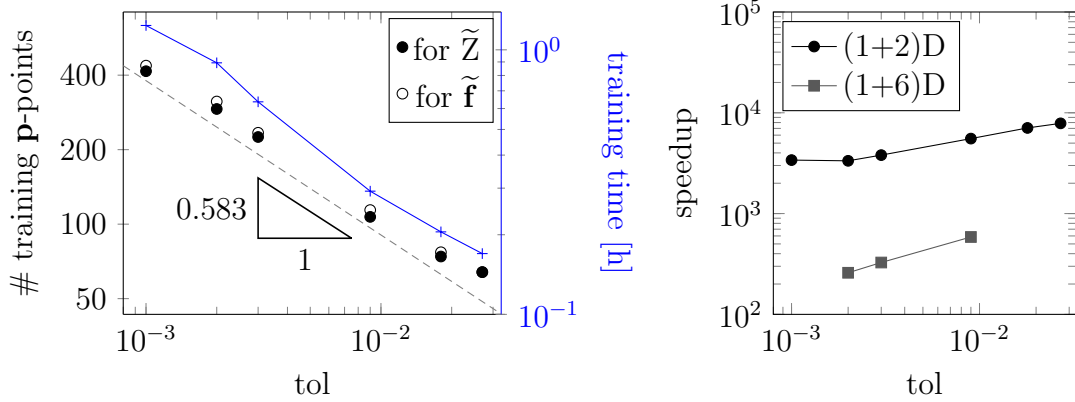


Figure 11: Timing results for different values of tol . The left and right plots refer to the offline and online phases, respectively. Left plot: the black dots represent the sizes of the training sets, for $\tilde{\mathbf{Z}}$ and $\tilde{\mathbf{f}}$ (full and empty dots, resp.), with the scale on the left axis. The wall-clock time for the training phase is shown as a full blue line, with the scale on the right axis. A trend line is also included, showing the average slope of the training cost. Right plot: online speedup for the (1+2)D and (1+6)D surrogates.

Offline cost. In Fig. 11 (left), we show how the number of explored training \mathbf{p} -points increases as tol decreases. In the same plot, we also show that the training time is approximately proportional to the size of the training set. This confirms the fact that computing high-fidelity samples is the bottleneck of the training phase. A trend line is also shown, investigating how the training cost scales with respect to the target accuracy tol . The observed rate ($\sim \text{tol}^{-0.583}$) is particularly interesting if compared to the theoretical approximation rate by hat functions in 1D, namely, $\sim \text{tol}^{-0.5}$, and in 2D over a tensor grid (as opposed to a sparse grid), namely, $\sim \text{tol}^{-1}$. This comparison reminds us of the fact that sparse grids can often recover approximation rates similar to those in just 1 dimension, being only weakly affected by the “curse of dimension”.

Online cost. In Fig. 11 (right), we display the average “online speedup” of the surrogate model. To assess the online speedup, each of the trained surrogates and the high fidelity training model are evaluated at 300 sample points chosen by a Latin hypercube sampling in \mathbf{p} -space. Then the “speedup” is defined as the factor (averaged over the 300 sample points) by which evaluating the surrogates $\tilde{\mathbf{Z}}$ and $\tilde{\mathbf{f}}$ is less expensive than computing these values with the high-fidelity training model. Since \mathbf{Z} also depends on ω , for fairness to the high-fidelity models, we compute the impedance speedup by comparing evaluations of $\tilde{\mathbf{Z}}$ and \mathbf{Z} at 100 different values of ω , for each fixed \mathbf{p} .

We observe that the speedup decreases as we go to smaller values of tol : the surrogate model becomes more complex there, since more collocation points are required. Still, even in the worst case, the speedup is always on the order of 10^3 .

As a final note on the online costs, we also compare our surrogate models with the simplified cNGT bearing model referred to in Remark 7, which neglects eccentricity dependence. Despite its simplicity (one of the parameters is completely ignored), the cNGT model

evaluation is more than 55 times slower than all our surrogates.

5.3 (1+6)D numerical results

In this section, we present results of our most complex surrogate, a complete parametric model of the classical HGJB. Both design and operational parameters are allowed to vary: $\mathbf{p} := (x_0, p_\Lambda, p_\alpha, p_\beta, p_L, p_H)$. The parameter ranges for each of the (1+6)-parameters are given in Table 1. We train three surrogates, differing in the tolerance level $\text{tol} \in \{2 \cdot 10^{-3}, 3 \cdot 10^{-3}, 9 \cdot 10^{-3}\}$.

For the purpose of validating the surrogates against the high-fidelity model, we sample the approximation errors over a set of 300 points generated by Latin hypercube sampling. At each of the sample points, the metrics $\bar{e}_{\star, \xi}$ and e_ζ (see Section 3.2) are calculated, in order to quantify the errors between surrogate and high-fidelity values of Z and \mathbf{f} . Fig. 12 shows the corresponding empirical probability density functions. We observe that, as tol decreases, the error distributions are shifted accordingly.




Figure 12: Histograms of testing errors, measured at 300 new values of \mathbf{p} . Impedance and static forces errors are shown in the top and bottom rows, respectively. Two different values of tol are considered in the two columns. The dashed vertical lines are placed at tol . The proportion of samples that do *not* satisfy the tolerance is indicated in the top right of each plot.

For Z , only 13% and 5% of the errors in our random sample on the training domain exceed the respective tol values. Even when the tolerance is not attained, the errors remain approximately within a $(2 \cdot \text{tol})$ -level. This behavior is to be expected, since our way of enforcing the tolerance cannot prevent larger-than-desired errors to arise far from the test set used in the predictor-corrector step. On the other hand, \mathbf{f} behaves better, with all error samples falling below tol .

Table 2: Summary of computational costs for training (1+6)D surrogates.

		value of tol		
		$2 \cdot 10^{-3}$	$3 \cdot 10^{-3}$	$9 \cdot 10^{-3}$
building \tilde{Z}	#training p -points	69745	50043	19584
	training time [h]	248	184	68
post-refinement for $\tilde{\mathbf{f}}$	#training p -points	29900	22690	9407
	training time [h]	102	80	31

Further, the tol level qualitatively falls within or just above the upper tail of the distributions. The close proximity of each distribution to its “tol” value speaks to the effectiveness of greedy sampling. Indeed, in the interest of achieving both a quick training phase and a favorable speedup factor, we want the algorithm to use as few training points as necessary. Error distributions much lower than their training tol level would have suggested that fewer collocation points could have been used to get the desired model accuracy.

In Fig. 11 (right), we also show the online speedup factors for the (1+6)D surrogates. As in the previous example, the speedups are computed based on the 300 **p**-points used in post-processing. We see that, despite four additional parameters, our algorithm still maintains a speedup on the order of 10^2 . Even in the worst case, there is still a significant speedup ($>5x$) compared to the simplified cNGT model, which neglects the effect of static eccentricity (cf. Remark 7).

As a final point, we revisit training costs. The training times and the required number of training points are reported in Table 2. As in the (1+2)D case, our surrogates maintain a proportionality between training time and the number of training points required. The training set size is observed to scale as $\sim \text{tol}^{-0.839}$. This is a further justification for the use of sparse grids as a means to combat the “curse of dimension” in our application. Compared to theoretical 6D tensor grid scaling ($\sim \text{tol}^{-3}$) our observed rate again is more similar to theoretical 1D scaling ($\sim \text{tol}^{-0.5}$).

6 Conclusions

In this work, we have proposed a strategy for building surrogate models of QoIs arising from generic parametric frequency-domain problems, as exemplified by the costly nonlinear systems resulting from the study of gas-lubricated bearing dynamics under harmonic perturbation. Non-intrusiveness is paramount in our approach: we do not make any assumptions on the structure of the target high-fidelity problem, so that our method can be applied even if the underlying model is accessible only via “samples” obtained through closed-source code. Via locally adapted sparse grids, we can deal even with a modest number of parameters.

Numerical results show, by example, that with our method one can build accurate surrogate models, with significant speedups compared to their reference high-fidelity model. The obtained surrogates have the potential to enable (near-)real-time behavior characterizations of (nonlinear) dynamical systems where they would otherwise be unfeasible, due to the high cost of the underlying model. We look to engineers and practitioners in fields like real-time control and design optimization to leverage potential gains in model efficiency into

advances in performance, efficiency, and reliability of real-world applications. Notably, tasks such as anomaly detection and stability analysis (even with feedback from measurements), or optimization of a rich physical model become attainable “online”.

The price to pay for a reliable and cheap surrogate model is a lengthy training phase, which however can be performed only once and “offline”, i.e., before the surrogate is deployed. The training phase can be especially costly when the number of parameters is large, due to the necessity to (non-intrusively) explore high-dimensional parameter domains with sparse grids. In our numerical tests, adaptivity allowed us to tackle problems depending on up to 7 parameters, which would have otherwise been intractable, e.g., with tensorized sampling grids. Still, in modest-to-high-dimensional settings, adaptivity can only partially reduce the unavoidable training burden.

It is in this context that we envision the most interesting further research directions, specifically in the study of more effective (adaptive) sampling strategies in high-dimension. Also, since we are sampling with the objective of interpolating, there is also the question of the interpolation basis: in our experiments, we have employed piecewise-linear hat functions, because they are intrinsically hierarchical and naturally suited for approximation over LASGs. Still, their weakness is the lack of accuracy. Alternatives like polynomials, splines, or radial basis functions are often unstable or unwieldy when it comes to high-dimensional *unstructured* interpolation. On the other hand, approximation methods rooted in “machine learning”, like Gaussian process regression or neural networks may be more promising. We feel that, in this direction, there is much room for research, both in theory and in numerical practice.

A Reference model details

The nondimensional formulation of the NGT equation given by Eq. (14) and the frequency-domain formulation in Section 4.2 can be derived using the scalings

$$\begin{aligned} P &= P^*/P_a^*, & \omega &= \omega^*/\omega_{rot}^*, & t &= t^*\omega_{rot}^*, \\ h_r &= h_r^*/h_{r,nom}^*, & \text{and} & & h_g &= h_g^*/h_{r,nom}^*, \end{aligned} \quad (19)$$

where quantities with a “ \star ” as superscript are dimensional, while quantities without superscript are not. We refer to Fig. 3 for context.

For brevity, we introduced coefficients in Eq. (14), which can be expanded as

$$\begin{aligned} a_1 &= \frac{b_1 + b_2 \sin^2 p_\beta}{b_3}, & a_2 &= \frac{b_2 \sin p_\beta \cos p_\beta}{b_3}, & a_3 &= \frac{b_1 + b_2 \cos^2 p_\beta}{b_3} \\ a_4 &= \frac{h_g^3 - h_r^3}{b_3}, & a_5 &= p_\alpha h_g + (1 - p_\alpha) h_r, & \text{and} \\ a_s &= -p_\Lambda p_\alpha (1 - p_\alpha) (p_H - 1) \sin p_\beta, \end{aligned} \quad (20)$$

using the quantities

$$b_1 = h_r^3 h_g^3, \quad b_2 = (h_g^3 - h_r^3)^2 p_\alpha (1 - p_\alpha), \quad b_3 = (1 - p_\alpha) h_g^3 + p_\alpha h_r^3. \quad (21)$$

We note that a_1, \dots, a_5 are space-dependent coefficients and thus they have a perturbation form obtained by injection of Eq. (15) in the development of the frequency-domain model in Section 4.2.

Acknowledgments

The authors extend their heartfelt thanks to Dimitri Goutaudier for his invaluable feedback on a preliminary version of this work. His contributions significantly improved the quality of this manuscript.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] Adil Raheed, O. San, and T. Kvamsdal. Digital Twin: Values, Challenges and Enablers From a Modeling Perspective. *IEEE Access*, 8:21980–22012, 2020.
- [2] F. Alsayyari, Z. Perkó, D. Lathouwers, and J. L. Kloosterman. A nonintrusive reduced order modelling approach using Proper Orthogonal Decomposition and locally adaptive sparse grids. *Journal of Computational Physics*, 399, 2019.
- [3] R. L. Bisplinghoff, H. Ashley, and R. L. Halfman. *Aeroelasticity*. Courier Corporation, June 2013.
- [4] F. Bonizzoni, D. Pradovera, and M. Ruggeri. Rational-approximation-based model order reduction of Helmholtz frequency response problems with adaptive finite element snapshots. *Mathematics in Engineering*, 5(4), 2023.
- [5] V. Castelli and J. H. Vohr. Performance characteristics of herringbone- grooved journal bearings operating at high eccentricity ratios and with misalignment. Technical report, NASA, Mar. 1967.
- [6] S. Chaturantabut and D. C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, 2010.
- [7] V. N. Constantinescu and V. Castelli. On the Local Compressibility Effect in Spiral-Groove Bearings. *Journal of Lubrication Technology*, 91(1):79–86, Jan. 1969.
- [8] D. Goutaudier, F. Nobile, and J. Schiffmann. A new method to interpolate POD reduced bases–Application to the parametric model order reduction of a gas bearings supported rotor. *International Journal for Numerical Methods in Engineering*, 124(18):4141–4170, 2023.
- [9] L. Gu, E. Guenat, and J. Schiffmann. A Review of Grooved Dynamic Gas Bearings. *Applied Mechanics Reviews*, 72(1), 2020.

- [10] E. Guenat and J. Schiffmann. Effects of humid air on aerodynamic journal bearings. *Tribology International*, 127:333–340, Nov. 2018. Reporter: Tribology International.
- [11] M. Guo and J. S. Hesthaven. Reduced order modeling for nonlinear structural analysis using Gaussian process regression. *Computer Methods in Applied Mechanics and Engineering*, 341:807–826, 2018.
- [12] M. A. Hassini and M. Arghir. A Simplified Nonlinear Transient Analysis Method for Gas Bearings. *Journal of Tribology*, 134(1), Jan. 2012. Number: 1 Reporter: Journal of Tribology.
- [13] J. S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. SpringerBriefs in Mathematics. Springer International Publishing, Cham, 2016.
- [14] J. S. Hesthaven and S. Ubbiali. Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics*, 363:55–78, June 2018.
- [15] W. Liu, P. Bättig, P. H. Wagner, and J. Schiffmann. Nonlinear study on a rigid rotor supported by herringbone grooved gas bearings: Theory and validation. *Mechanical Systems and Signal Processing*, 146:106983, Jan. 2021.
- [16] J. W. Lund. Calculation of Stiffness and Damping Properties of Gas Bearings. *Journal of Lubrication Technology*, 90(4):793–803, Oct. 1968. Number: 4 Reporter: Journal of Lubrication Technology.
- [17] X. Ma and N. Zabaras. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *Journal of Computational Physics*, 228(8):3084–3113, 2009.
- [18] S. B. Malanoski. Experiments on an Ultrastable Gas Journal Bearing. *Journal of Lubrication Technology*, 89(4):433–438, Oct. 1967.
- [19] Y. Nakatsukasa, O. Sète, and L. N. Trefethen. The AAA algorithm for rational approximation. *SIAM Journal on Scientific Computing*, 40(3):A1494–A1522, 2018.
- [20] A. H. Nayfeh. *Introduction to Perturbation Techniques*. John Wiley & Sons, Apr. 2011.
- [21] F. Nobile and D. Pradovera. Non-intrusive double-greedy parametric model reduction by interpolation of frequency-domain rational surrogates. *ESAIM: Mathematical Modelling and Numerical Analysis*, 55(5):1895–1920, 2021.
- [22] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab. *Signals & systems*. Prentice-Hall signal processing series. Prentice Hall, Upper Saddle River, New Jersey, second ed. edition, 1997.
- [23] G. M. Phillips. *Interpolation and approximation by polynomials*, volume 41 of *CMS books in mathematics*. Springer, New York, 2004.

- [24] D. Pradovera. Interpolatory rational model order reduction of parametric problems lacking uniform inf-sup stability. *SIAM Journal on Numerical Analysis*, 58(4):2265–2293, 2020.
- [25] D. Pradovera and F. Nobile. Frequency-Domain Non-intrusive Greedy Model Order Reduction Based on Minimal Rational Approximation. *Scientific Computing in Electrical Engineering*, 36:159–167, 2021.
- [26] D. Pradovera and F. Nobile. A technique for non-intrusive greedy piecewise-rational model reduction of frequency response problems over wide frequency bands. *Journal of Mathematics in Industry*, 12, 2022.
- [27] S. A. Renganathan, R. Maulik, and V. Rao. Machine learning for nonintrusive model order reduction of the parametric inviscid transonic flow past an airfoil. *Physics of Fluids*, 32(4):047110, Apr. 2020.
- [28] K. Rosset, O. Pajot, and J. Schiffmann. Experimental Investigation of a Small-Scale Organic Rankine Cycle Turbo-Generator Supported on Gas-Lubricated Bearings. *Journal of Engineering for Gas Turbines and Power*, 143(5), Mar. 2021.
- [29] E. Saff and R. Varga. *Padé and rational approximation: theory and applications*. Academic Press Inc, 1977.
- [30] E. Saliby and F. Pacheco. An empirical evaluation of sampling methods in risk analysis simulation: Quasi-Monte Carlo, descriptive sampling, and Latin Hypercube Sampling. In *Winter Simulation Conference Proceedings*, volume 2, pages 1606–1610, 2002.
- [31] J. Schiffmann. *Integrated design, optimization and experimental investigation of a direct driven turbocompressor for domestic heat pumps*. PhD thesis, EPFL, 2008.
- [32] P. H. Wagner, Z. Wuillemin, D. Constantin, S. Diethelm, J. Van herle, and J. Schiffmann. Experimental characterization of a solid oxide fuel cell coupled to a steam-driven micro anode off-gas recirculation fan. *Applied Energy*, 262:114219, Mar. 2020.
- [33] J. L. Walsh. *Interpolation and Approximation by Rational Functions in the Complex Domain*. American Mathematical Society, 1960.
- [34] R. T. P. Whipple. Theory of the spiral grooved thrust bearing with liquid or gas lubricant. Technical Report, Great Britain Atomic Energy Research Establishment, 1951.
- [35] R. T. P. Whipple. The inclined groove bearing. Technical Report AERE-T/R-622(Rev.), United Kingdom Atomic Energy Authority. Research Group. Atomic Energy Research Establishment, Harwell, Berks, England, Oct. 1958.
- [36] R. Zimmermann. A Locally Parametrized Reduced-Order Model for the Linear Frequency Domain Approach to Time-Accurate Computational Fluid Dynamics. *SIAM Journal on Scientific Computing*, 36(3):B508–B537, Jan. 2014. Publisher: Society for Industrial and Applied Mathematics.