

Two-step greedy algorithm for reduced order quadratures

Harbir Antil, Scott E. Field, Frank Herrmann, Ricardo H. Nochetto,
Manuel Tiglio

the date of receipt and acceptance should be inserted later

Abstract We present an algorithm to generate application-specific, global *reduced order quadratures* (ROQ) for multiple fast evaluations of weighted inner products between parameterized functions. If a reduced basis (RB) or any other projection-based model reduction technique is applied, the dimensionality of integrands is reduced dramatically; however, the cost of evaluating the reduced integrals still scales as the size of the original problem. In contrast, using discrete empirical interpolation (DEIM) points as ROQ nodes leads to a computational cost which depends linearly on the dimension of the reduced space. Generation of a reduced basis via a greedy procedure requires a training set, which for products of functions can be very large. Since this direct approach can be impractical in many applications, we propose instead a two-step greedy targeted towards approximation of such products. We present numerical experiments demonstrating the accuracy and the efficiency of the two-step approach. The presented ROQ are expected to display very fast convergence whenever there is regularity with respect to parameter variation. We find that for the particular application here considered, one driven by gravitational wave physics, the two-step approach speeds up the offline computations to build the ROQ by more than two orders of magnitude. Furthermore, the resulting ROQ rule converges exponentially with the number of nodes, and a factor of ~ 50 savings, without loss of accuracy, is observed in evaluations of inner products when ROQ are used as a downsampling strategy for equidistant samples using the trapezoidal rule. While the primary focus of this paper is on quadrature rules for inner products of parameterized functions, our method can be easily adapted to integrations of single parameterized functions, and some examples of this type are also considered.

1 Introduction

Many application areas deal with parameterized problems. Throughout this paper, we consider these to be a set of functions

$$\mathcal{F}_W := \{h_\mu : \Omega \rightarrow \mathbb{C} \mid \mu \in \mathcal{P}, h_\mu \in \mathcal{C}(\Omega)\},$$

Harbir Antil
Department of Mathematical Sciences, George Mason University, Fairfax, VA 22030, USA. Tel.: +703-993-2390
Fax: +703-993-1491
E-mail: hantil@gmu.edu

Scott E. Field
Department of Physics, Joint Space Science Institute, Maryland Center for Fundamental Physics, University of Maryland, College Park, MD 20742, USA.

Frank Herrmann
Center for Scientific Computation and Mathematical Modeling, Department of Physics, Maryland Center for Fundamental Physics, University of Maryland, College Park, MD 20742, USA.

Ricardo H. Nochetto
Department of Mathematics, and Institute of Physical Science and Technology, University of Maryland, College Park, MD 20742, USA.

Manuel Tiglio
Center for Scientific Computation and Mathematical Modeling, Department of Physics, Joint Space Science Institute, Maryland Center for Fundamental Physics, University of Maryland, College Park, MD 20742, USA.

where Ω , \mathcal{P} denote the physical and parameter domains, respectively, \mathbb{C} is the set of complex numbers, $\mathcal{C}(\Omega)$ is the set of continuous functions on compact domain $\Omega \subset \mathbb{R}$ and $\mathcal{F}_W \subset \mathcal{H}_W$ denotes a compact subset of the Hilbert space $\mathcal{H}_W := L^2_W(\Omega)$. In general, both Ω and \mathcal{P} are multi-dimensional and can be irregular domains. Here we assume \mathcal{P} to be compact in \mathbb{R}^N and take the scalar product and norm between two functions $f, g \in \mathcal{H}_W$ to be

$$\langle f, g \rangle_{L^2_W} := \int_{\Omega} f^*(x)g(x)W(x)dx, \quad \|f\|_{L^2_W}^2 := \langle f, f \rangle_{L^2_W}, \quad (1)$$

with $0 < W \in \mathcal{C}(\Omega)$ some weight function and f^* denoting complex conjugation. We absorb the weight above into each integrand function through¹

$$f(x) \rightarrow W^{\frac{1}{2}}(x)f(x),$$

with $x \in \Omega$, and define the set of functions

$$\mathcal{F} = \{h_{\mu}\sqrt{W} \in \mathcal{C}(\Omega) \mid h_{\mu} \in \mathcal{F}_W\},$$

which is now a compact subset of the Hilbert space $\mathcal{H} := L^2(\Omega)$. This allows us, without loss of generality, to use the standard L^2 inner product and the corresponding norm for elements of \mathcal{H} , i.e.,

$$\langle f, g \rangle = \int_{\Omega} f^*(x)g(x)dx, \quad \|f\|^2 := \langle f, f \rangle. \quad (2)$$

Furthermore, the discrete L^2 inner product and the corresponding norm are here denoted by

$$\langle f, g \rangle_{\mathbf{d}} = \sum_{k=1}^M \omega_k f^*(x_k)g(x_k), \quad \|f\|_{\mathbf{d}}^2 = \langle f, f \rangle_{\mathbf{d}}, \quad (3)$$

where $\{x_k, \omega_k\}_{k=1}^M$ are arbitrary quadrature points and weights. Additionally, we often denote discrete objects and vectors through bold notation, for example, $\mathbf{f} = (f(x_1), \dots, f(x_M))^T$ and $\boldsymbol{\omega} = (\omega_1, \dots, \omega_M)^T$. More details about the notation used throughout this paper are given in Appendix A. We assume exactness between continuous and discrete inner products within machine precision.

There are several challenges associated with parameterized problems and Reduced Order Modeling (ROM). One of them, when the set of functions \mathcal{F} is not known a priori but requires expensive numerical simulations, is how to select “on the fly” which parameters to solve for in a nearly optimal way, and a compact –application-specific spectral–representation for every element in \mathcal{F} . Reduced Basis (RB) [1] is a leading candidate for such problems. We refer to Section 3.1 for details about it; for the purposes of this paper it suffices to describe in that section the algorithm and some of its key features. Another often desired aspect of any approach to parameterized problems is, once a reduced basis has been built, the ability to accelerate the computation of a particular quantity of interest. This can be fast online prediction of new solutions by solving a reduced problem [1] or, as in the case of this paper, fast online evaluations of inner products between elements of \mathcal{F} . Additional benefits may be gained from an interpolation-based approximation. The leading candidates for this are the Empirical Interpolation Method (EIM) [2,3] and its discrete counterpart the Discrete Empirical Interpolation Method (DEIM) [4,5]. Again, for the purposes of this paper it suffices to summarize the algorithm in Section 3.2 and some of its properties.

The specific goal of this paper is to build application-specific global quadratures for **inner products of functions** (such as those appearing in convolutions, matched filtering, Markov Chain Monte Carlo simulations, etc.) combining both ROM and the DEIM, which for brevity we refer to as *reduced order quadratures* (ROQ). Here we explicitly consider RB for ROM, but the approach remains the same for other choices of bases generated using, for example, Proper Orthogonal or Singular Value Decompositions (POD, SVD) [6].

The construction of ROQ is somewhat similar in spirit to that one of Gaussian quadratures: the function to be integrated is approximated by a truncated expansion in a reduced basis, this expansion is replaced by interpolation (here DEIM) at special nodal points and, finally, the interpolant is integrated to compute effective quadrature weights for any scalar product of functions in the set of interest \mathcal{F} . For many applications of interest, there are several advantages of ROQ with respect to Gaussian quadratures, though. First, regularity with respect to parameter variation, not with respect to the integration variable, is exploited to provide very fast (in many cases of interest, exponentially fast) convergence.

¹ For generic target applications we envision these functions not to be polynomials, and the weights not to be standard ones. In particular, our proposed reduced order quadrature construction does not rely on classical orthogonal polynomial theory for which the precise form of the weight W is essential.

For example, for the gravitational wave test cases considered in this paper, ROQ prove to be more efficient than Gaussian quadratures (see Fig. 8). In addition, the method is adaptive by nature, and the quadrature nodes are hierarchical. Our approach also gives, when acquiring data from experiments or observations, a nearly optimal way of downsampling it for matched filtering [7, 8] or other purposes, while at the same time preserving the accuracy in computing integrals with the full data set. All these benefits of ROQ come at the cost of a potentially expensive offline procedure to pre-assemble both the basis and the ROQ rule. For many applications where multiple fast evaluations are required, especially if they are needed online or in real time, this tradeoff is indeed desirable.

To the best of our knowledge, leveraging the advantages of the Empirical Interpolation Method for fast numerical integration was first suggested by Maday et al [2] (see also Refs. [9, 10]) and further investigations were carried out by Aanonsen [10]. However, our approach differs from previous ones in several respects. First, the empirical interpolant coefficients are usually found by carrying out a potentially costly (compared to a competitive quadrature rule) matrix-vector multiplication. Here we absorb this cost into an *offline* computation of the reduced order quadrature weights and explicitly write the ROQ rule as a vector-vector product (see Eq. (10) in Section 2). Additionally, the RB-DEIM model reduction used here allows for a natural factorization of the parameter and physical dependences. Finally and most important, our main objective here is fast computation of **inner products** of the form $\langle h_{\mu_i}, h_{\mu_j} \rangle$. While this is closely related to the integration of single functions in \mathcal{F} , for large problems there are practical and serious obstacles towards a reduced basis representation of products of functions, denoted here as

$$\tilde{\mathcal{F}} = \{h_{\mu_i}^* h_{\mu_j} \in \mathcal{C}(\Omega) \mid h_{\mu_i}, h_{\mu_j} \in \mathcal{F}\}, \quad (4)$$

in terms of sizes. Most notably very large training spaces. To address this we suggest a simple two-step approach targeted towards such products which reuses the algorithms necessary for reducing the underlying lower order space \mathcal{F} , with dramatic savings in the offline stage. Without this two-step procedure and/or adaptive sampling of the training space, building ROQ would be simply not possible in many realistic application problems driving our work, even if carried out offline and using large supercomputers.

This paper is organized as follows. Section 2 provides an overview of the main ideas to be developed throughout the paper as well as brief preliminary discussions on quadrature rules, RB-DEIM, and reduced order quadratures. Section 3 collects necessary results and algorithms required for developing ROQ. In particular, in Section 3.1 we describe a RB-greedy model reduction algorithm to generate a set of basis vectors which are subsequently interpolated at a set of points chosen by the DEIM algorithm as summarized in Section 3.2. The main difficulties in constructing a reduced basis for $\tilde{\mathcal{F}}$, as well as our approach for overcoming them, are considered in Section 3.3. Fast integrations using the reduced model and ROQ are constructed in Section 4.1, along with some error estimates in Section 4.2. Straightforward extensions of the ROQ algorithm for resampled functions are discussed in Section 4.3. Two numerical experiments are documented in Section 5. To compare with well known results in one dimension, our first experiment considers polynomials on the standard interval $[-1, 1]$ as basis, and integration of single functions. The second example draws from a non-trivial problem in gravitational wave physics and showcases the potential savings of ROQ to compute scalar products in large problems. In Section 6 we summarize the results of this paper, and comment on their possible extensions.

2 Overview and main ideas

The main goal of this paper is to efficiently compute approximations to integrals such as those in Eq. (2). We refer to the exact, continuum one as I_c ,

$$I_c(i, j) := \langle h_{\mu_i}, h_{\mu_j} \rangle, \quad (5)$$

where $\{\mu_i, \mu_j\}$ are any two values in the parameter space \mathcal{P} . Evaluation of this integral may be costly and for some applications must be repeatedly computed, perhaps in real-time; in such situations a Gaussian quadrature rule might be either limiting or unavailable (for example, if the integrands are not smooth enough or their value is not known at Gaussian nodes). Problems which depend on a high dimensional parameter or physical spaces constitute such cases. In addition, if the target functions $\{h_{\mu_j}\}$ come from acquired data, each function h_{μ_j} is usually sampled at equally spaced or scattered points and the integrations are usually carried out using an extended trapezoidal rule, with slow convergence in general.

Next we discuss different approximations to continuum inner products I_c such as those in Eq. (5).

Integration by quadrature: Let us recall a typical quadrature rule, since the proposed ROQ of this paper follow a somewhat similar pattern. With $\{x_k, \omega_k\}_{k=1}^M$ denoting an arbitrary set of quadrature points and weights, respectively,

we can approximate the integral (5) as

$$I_c(i, j) \approx I_d(i, j) := \langle h_{\mu_i}, h_{\mu_j} \rangle_d = \sum_{k=1}^M \omega_k h_{\mu_i}^*(x_k) h_{\mu_j}(x_k). \quad (6)$$

Many integration rules can be written in the form of Eq. (6), including the extended trapezoidal rule, Gaussian quadratures, and integration using domain decomposition. Among these many choices an optimal strategy should allow for a given target integration accuracy with the smallest number of operations. If the number of quadrature points M is not too large, one can evaluate (6) with standard methods provided not too many integrals are required. However, if M is large (perhaps the functions are sharply peaked, have different length scales and/or many cycles) and/or many evaluations of I_c are needed, computing the approximations I_d can become very costly.

Integration in a reduced space: A preliminary reduced order approximation to $I_d(i, j)$ would involve using an expansion in basis vectors $\{e_\ell\}_{\ell=1}^n$. Without loss of generality we can assume that the basis is orthonormal with respect to the discrete ² inner product $\langle \cdot, \cdot \rangle_d$. Then any function $h_\mu \in \mathcal{F}$ can be approximated by

$$h_\mu \approx \mathcal{P}_n h_\mu := \sum_{\ell=1}^n \langle e_\ell, h_\mu \rangle_d e_\ell,$$

where we have introduced the orthogonal projection operator \mathcal{P}_n , and an approximation to (6) would be

$$I_d(i, j) \approx \langle \mathcal{P}_n h_{\mu_i}, \mathcal{P}_n h_{\mu_j} \rangle_d = \sum_{k=1}^n \langle h_{\mu_i}, e_k \rangle_d \langle e_k, h_{\mu_j} \rangle_d. \quad (7)$$

Now, computing the approximation (7) requires knowledge of the projection coefficients $\{\langle h_{\mu_i}, e_k \rangle_d\}$. Modulo the latter, performing an integration via Eq. (7) is of improved computational cost (n multiplications and $n - 1$ additions) whenever the number of basis elements is smaller than the number of quadrature points, $n < M$. A greedy construction of a *reduced basis* (see Sec. 3.1), as opposed to a standard basis choice such as Jacobi polynomials, involves the use of a problem-dependent *training space*. If the functions to be integrated, namely h_{μ_i} and h_{μ_j} , are members of this training space, the projection coefficients on the right hand side of Eq. (7) have been precomputed offline while building the basis (see Section 3.1). The more interesting case in practice is that one in which h_{μ_i} and h_{μ_j} were not members of the training space, or one is unable to store the set of projection coefficients.

Integration with reduced order quadratures (ROQ): We advocate an application-specific quadrature scheme for efficiently evaluating the discrete integral in (6), with essentially no loss of accuracy, with a computational cost that depends only on the reduced basis space dimension, and without requiring projection coefficients. This new approach combines the flexibility of quadrature rules with powerful dimensionality reduction. We seek to approximate integrands of the form $g_{ij} = h_{\mu_i}^* h_{\mu_j}$, which is where the main challenge lies; namely, in integrating products of target functions and not individual ones. Throughout this paper tildes will refer to approximations or quantities related to $\tilde{\mathcal{F}}$, defined in Eq. (4). We are therefore seeking to approximate integrals of

$$g_{ij} \approx \tilde{\mathcal{P}}_m g_{ij} := \sum_{\ell=1}^m \langle \tilde{e}_\ell, g_{ij} \rangle_d \tilde{e}_\ell,$$

with $g_{ij} \in \tilde{\mathcal{F}}$, and a truncated expansion is carried with a set of basis vectors $\{\tilde{e}_\ell\}_{\ell=1}^m$ approximating elements in $\tilde{\mathcal{F}}$. As before, we are faced with computing or approximating the projection coefficients $\langle \tilde{e}_\ell, g_{ij} \rangle_d$, the cost of which can in principle be expensive. Discrete empirical interpolation (DEIM) offers an attractive alternative: consider these m basis $\{\tilde{e}_\ell\}_{\ell=1}^m$, let $\{\tilde{p}_\ell\}_{\ell=1}^m \subset \{x_k\}_{k=1}^M$ be a set of points (the generation of which will be described later), and let the DEIM interpolant of g_{ij} be

$$\tilde{\mathcal{I}}_m[g_{ij}] := \sum_{\ell=1}^m \tilde{c}_\ell \tilde{e}_\ell,$$

² In carrying out the RB-greedy algorithm we resolve the integrals within machine precision, so for all practical purposes the discrete and continuum scalar products agree with each other.

where the interpolation coefficients $\{\tilde{c}_\ell\}_{\ell=1}^m$ are so that $\tilde{\mathcal{I}}_m[g_{ij}](\tilde{p}_k) = g_{ij}(\tilde{p}_k)$. This means that we enforce nodal exactness of $\tilde{\mathcal{I}}_m$ and $\{\tilde{c}_\ell\}_{\ell=1}^m$ solves the linear system

$$\sum_{\ell=1}^m \tilde{c}_\ell \tilde{e}_\ell(\tilde{p}_k) = g_{ij}(\tilde{p}_k), \quad k = 1, \dots, m. \quad (8)$$

If the interpolant is accurate then $g_{ij} \approx \tilde{\mathcal{I}}_m[g_{ij}]$ and, substituting this into Eq.(6), our approximation to the quadrature rule $\{x_k, \omega_k\}_{k=1}^M$ becomes

$$I_d(i, j) \approx \sum_{k=1}^M \omega_k \tilde{\mathcal{I}}_m[g_{ij}](x_k) = \sum_{k=1}^M \omega_k \sum_{\ell=1}^m \tilde{c}_\ell \tilde{e}_\ell(x_k) = \sum_{\ell=1}^m \left[\sum_{k=1}^M \omega_k \tilde{e}_\ell(x_k) \right] \tilde{c}_\ell. \quad (9)$$

Written in this way, the bracketed expression in the last term of Eq. (9) is seen to be the integration of the basis and acts as an effective quadrature weighting. Later on (Algorithm 4 and, in particular, Eq.(26)) we show that this expression can be rewritten in a more recognizable form

$$I_d(i, j) \approx I_{\text{ROQ}}(i, j) := \sum_{\ell=1}^m \omega_\ell^{\text{ROQ}} h_{\mu_i}^*(\tilde{p}_\ell) h_{\mu_j}(\tilde{p}_\ell) = \sum_{\ell=1}^m \omega_\ell^{\text{ROQ}} g_{ij}(\tilde{p}_\ell). \quad (10)$$

Eq. (10) is our proposed Reduced Order Quadrature, which has an online evaluation cost of $\mathcal{O}(m)$. The interpolation points $\{\tilde{p}_\ell\}_{\ell=1}^m \subset \{x_k\}_{k=1}^M$ are outputs of the DEIM algorithm. Both these points and the weights ω_ℓ^{ROQ} only depend on the underlying quadrature rule given by $\{x_k, \omega_k\}_{k=1}^M$ and on the choice of reduced basis vectors \tilde{e}_ℓ ; they are precomputed off-line and do not depend on the integrand g_{ij} . Note that the ROQ rule (10) achieves our efficiency requirements whenever $m < M$. In many cases one can expect exponential convergence of $I_{\text{ROQ}} \rightarrow I_d$, and for $I_d \approx I_c$ we expect this convergence rate to be inherited by the limit to the continuum $I_{\text{ROQ}} \rightarrow I_c$ with respect to m . Furthermore, even though this might be problem-dependent, in our numerical experiments (see Section 5.2) we have found that $m \sim 2n \ll M$ (the equality $m = 2n$ exactly holds for polynomial bases). For the case $m = M$, when all degrees of freedom have been exhausted, consistency requires both quadrature rules $\{\tilde{p}_\ell, \omega_\ell^{\text{ROQ}}\}_{\ell=1}^m$ and $\{x_k, \omega_k\}_{k=1}^M$ to be identical. Indeed, in this case the quadratures points trivially coincide and one can show $\omega_\ell^{\text{ROQ}} = \omega$ (see Corollary 1).

In Sec. 3 we propose and analyze an algorithm for constructing the product basis $\{\tilde{e}_\ell\}_{\ell=1}^m$ as well as identifying the nearly optimal interpolation points $\{\tilde{p}_\ell\}_{\ell=1}^m$. A straightforward and largely self contained blueprint to generate the quadrature rule $\{\tilde{p}_\ell, \omega_\ell^{\text{ROQ}}\}_{\ell=1}^m$ is given in Algorithm 4.

3 Reduced basis and the discrete empirical interpolation method

The ROQ rule proposed here requires a compact and accurate interpolation representation for all integrands in the space of interest. It has been shown that the EIM [2, 3] or its discrete counterpart, the DEIM [4, 5], offers an attractive interpolation approach for reducing the complexity of a problem without sacrificing accuracy. In turn, the approach relies on knowledge of a problem-specific (ordered) basis such that any function of interest can be accurately represented by its L^2 projection onto the space spanned by these basis vectors. There are many ways to generate such a basis set, such as Proper Orthogonal/Singular Value Decomposition [6] and the Reduced Basis (RB)-greedy algorithm [1]. Here we favor a RB-greedy approach, since it is able to handle very large problems, such as those that we are interested in, but our proposal is directly applicable to other basis choices.

In this section we describe both the RB-greedy and DEIM algorithms. Sec. 3.1 introduces the RB-greedy approach to approximate the set of functions \mathcal{F} or $\tilde{\mathcal{F}}$, and the DEIM algorithm of Sec. 3.2 provides empirical interpolation points for these basis functions. Instead of directly approximating $\tilde{\mathcal{F}}$, in Sec. 3.3 we propose an alternative, the two-step RB-greedy approach for large sets of products of functions. Without loss of generality we assume all functions to be normalized with respect to the discrete inner product $\langle \cdot, \cdot \rangle_d$.

3.1 Greedy construction of a reduced basis

The RB approach, combined with a greedy algorithm to generate the basis, provides a way to construct an application-specific expansion, where the basis elements are members of the space under consideration itself. Here the sets of interest

are \mathcal{F} and $\tilde{\mathcal{F}}$. For definiteness in notation we will describe the approach for \mathcal{F} , since an approach for $\tilde{\mathcal{F}}$ or any other space is identical; we will refer to this approach applied to $\tilde{\mathcal{F}}$ as a *direct* or one-step RB-greedy. Later in Sec. 3.3, we mention some of its drawbacks and our proposal to overcome them, via a two-step RB-greedy.

A greedy approach is highly efficient in practice. It yields a nested basis set that is hierarchically constructed. As is customary for spectral expansions, having an orthonormal basis $\{e_\ell\}_{\ell=1}^n$ simplifies computing the orthogonal projection $\mathcal{P}_n h_\mu$ with respect to $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ of a function $h_\mu \in \mathcal{F}$ onto the span F_n of the basis,

$$F_n := \text{span}\{e_1, \dots, e_n\},$$

namely,

$$\mathcal{P}_n h_\mu = \sum_{\ell=1}^n \langle e_\ell, h_\mu \rangle_{\mathcal{H}} e_\ell.$$

We recall (see, for example, [11]) that the Kolmogorov n -width of \mathcal{F} in \mathcal{H}

$$d_n(\mathcal{F}; \mathcal{H}) = \inf_{\dim X_n \leq n} \sup_{h_\mu \in \mathcal{F}} \inf_{f \in X_n} \|h_\mu - f\|, \quad (11)$$

measures the error of the best n -dimensional subspace $X_n \subset \mathcal{H}$ approximating \mathcal{F} , and since \mathcal{H} is a Hilbert space,

$$\inf_{f \in X_n} \|h_\mu - f\| = \|h_\mu - \mathcal{P}_n h_\mu\|.$$

Computation of the n -width d_n , or any basis achieving it, is in most practical applications not possible (however, see [11] for some simple cases where it can be done). Nevertheless, obtaining a convergence rate for the n -width provides valuable information towards understanding the approximability of a parameterized space by an algorithm such as the greedy.

We summarize a greedy strategy to build a reduced basis $\{e_\ell\}_{\ell=1}^n$, with an approximation error

$$\sigma_n(\mathcal{F}; \mathcal{H}) := \sup_{h_\mu \in \mathcal{F}} \|h_\mu - \mathcal{P}_n h_\mu\|, \quad (12)$$

which is nearly optimal with respect to the Kolmogorov n -width defined in Eq. (11). Then we will state some of the available convergence rates for the *greedy error* $\sigma_n(\mathcal{F}; \mathcal{H})$ based on [12, 13], which make precise the quoted near optimality, under the assumption that up to machine precision

$$\|h_\mu - \mathcal{P}_n h_\mu\| \simeq \|h_\mu - \mathcal{P}_n h_\mu\|_{\mathcal{H}}.$$

While the greedy error (12) has been defined over \mathcal{F} , in practice one discretely samples the continuum using a *training set* $\mathcal{T}_K := \{\mu_i\}_{i=1}^K \subset \mathcal{P}$ of size K and a *training space* $\mathcal{F}_{\text{train}}$ of associated normalized functions $\{h_{\mu_i}\}_{i=1}^K$. Typically, if there is redundancy, the number of reduced basis needed to represent the training space is much smaller than the number of samples: $n \ll K$.

Next let $\epsilon > 0$ be a user-specified error tolerance. Its role is to guarantee that the approximation error ensured by the RB-Greedy algorithm is strictly bounded as

$$\sigma_n(\mathcal{F}_{\text{train}}; \mathcal{H}) := \sup_{h_\mu \in \mathcal{F}_{\text{train}}} \|h_\mu - \mathcal{P}_n h_\mu\|_{\mathcal{H}} \leq \epsilon. \quad (13)$$

To a-priori ensure that the training space $\mathcal{F}_{\text{train}}$ is a faithful approximation of the continuum \mathcal{F} is, in general, difficult. Construction and adaptive management of the training space $\mathcal{F}_{\text{train}}$ is an area of active research; see, for example, [14, 9, 15, 16]. A good choice is dictated by the problem, further details for our application are given in the numerical results Section 5.2. To check that the reduced basis space F_n is a faithful approximation of \mathcal{F} we typically do convergence tests with respect to the number of samples K in \mathcal{T}_K , and Monte Carlo reconstruction studies of the continuum; see for example [17, 18, 19] and Section 5.2. For any given tolerance $\epsilon > 0$, Algorithm 1 ensures the strict bound (13) over the entire training space. However, a-posteriori validation has allowed us to establish a more impressive bound in all of our applications in gravitational wave physics so far [17, 18, 19],

$$\|h_\mu - \mathcal{P}_n h_\mu\|_{\mathcal{H}} \leq \sigma_n(\mathcal{F}; \mathcal{H}) \lesssim \epsilon \quad \forall \mu \in \mathcal{P}. \quad (14)$$

Such observation is obviously problem-dependent, and in general it should read $\sigma_n(\mathcal{F}; \mathcal{H}) \leq \tilde{\epsilon}$ where in principle $\tilde{\epsilon} \geq \epsilon$.

For any $\epsilon > 0$ and training set \mathcal{T}_K the greedy algorithm to build a reduced basis is as follows:

Algorithm 1 (RB Greedy Algorithm)

$[\{e_\ell\}_{\ell=1}^n, \{\mu_\ell\}_{\ell=1}^n] = \text{RB-Greedy}(\epsilon, \mathcal{T}_K)$

- (1) Set $n = 0$ and define $\sigma_0(\mathcal{F}_{\text{train}}; \mathcal{H}) := 1$
- (2) Choose an arbitrary $\mu_1 \in \mathcal{T}_K$ and set $e_1 := h_{\mu_1}$ **Comment:** $\|h_{\mu_1}\|_{\mathcal{H}} = 1$
- (3) do, while $\sigma_n(\mathcal{F}_{\text{train}}; \mathcal{H}) \geq \epsilon$
 - (a) $n = n + 1$
 - (b) $\sigma_n(h_\mu) := \|h_\mu - \mathcal{P}_n h_\mu\|_{\mathcal{H}}$ for all $\mu \in \mathcal{T}_K$
 - (c) $\sigma_n(\mathcal{F}_{\text{train}}; \mathcal{H}) = \sup_{\mu \in \mathcal{T}_K} \{\sigma_n(h_\mu)\}$
 - (d) $\mu_{n+1} := \text{argsup}_{\mu \in \mathcal{T}_K} \{\sigma_n(h_\mu)\}$ (greedy sweep)
 - (e) $e_{n+1} := h_{\mu_{n+1}} - \mathcal{P}_n h_{\mu_{n+1}}$ (Gram-Schmidt Orthogonalization)
 - (f) $e_{n+1} := e_{n+1} / \|e_{n+1}\|_{\mathcal{H}}$ (normalization)

The above algorithm returns a nested, hierarchical set of n greedy points $\{\mu_\ell\}_{\ell=1}^n$ and orthonormal reduced basis $\{e_\ell\}_{\ell=1}^n$ which are nearly optimal with respect to the n -width (11), in the following sense. Recall from [13, Corollary 3.3] that if the n -width defined in (11) decays exponentially,

$$d_n(\mathcal{F}; \mathcal{H}) \leq C e^{-c_0 n^\alpha}, \quad (15)$$

where C , c_0 , and α are positive constants, then the greedy error in (12) decays as

$$\sigma_n(\mathcal{F}; \mathcal{H}) \leq \sqrt{2C} e^{-c_1 n^\alpha}, \quad (16)$$

with $c_1 := 2^{-1-2\alpha} c_0$. Similar statements can be made if the Kolmogorov n -width decays with a polynomial order i.e., if the n -width in (11) decays as

$$d_n(\mathcal{F}; \mathcal{H}) \leq C_a n^{-\beta}, \quad (17)$$

where C_a , β are positive constants, then

$$\sigma_n(\mathcal{F}; \mathcal{H}) \leq 2^{5\beta+1} C_a n^{-\beta}. \quad (18)$$

In Section 5.2 we summarize the observed exponential decay of the greedy error for a particular family of smoothly parameterized gravitational wave solutions; see [20, 18, 17, 19] for further studies.

We have also found that for any given representation error and any parameter domain \mathcal{T}_K we can reconstruct any member function in \mathcal{F} , not necessarily in the training space used to build the basis, with a *finite number of basis* elements. That is, as the density of \mathcal{T}_K increases, a finite asymptotic number of basis is needed to represent it; see, for example, [18], and [21] for a similar result using RB and SVD respectively.

Remark 1 While it is obvious that $d_n(\mathcal{F}; \mathcal{H}) \leq \sigma_n(\mathcal{F}; \mathcal{H})$, the decay rate of the Kolmogorov n -width gives an estimate for the greedy error decay according to (16) (exponential) or (18) (algebraic). A direct comparison obtained in [13, Corollary 3.3] reads

$$\sigma_n(\mathcal{F}; \mathcal{H}) \leq \sqrt{2 d_{n/2}(\mathcal{F}; \mathcal{H})},$$

for any n -width decay rate. If $d_{n/2}(\mathcal{F}; \mathcal{H})$ decays as in (15), then one gets (16). Similarly, if $d_{n/2}(\mathcal{F}; \mathcal{H})$ decays as (17) then, one gets (18).

We now turn to some immediate applications of the reduced space expressed with matrix-vector notation, which will prove convenient throughout the remainder of this paper. Let

$$\mathbf{V} = [\mathbf{e}_1, \dots, \mathbf{e}_n] \in \mathbb{C}^{M \times n},$$

where the ℓ^{th} column of the matrix \mathbf{V} is $\mathbf{e}_\ell = (e_\ell(x_1), \dots, e_\ell(x_M))^T \in \mathbb{C}^M$ i.e., it corresponds to the ℓ^{th} reduced basis e_ℓ sampled at the set of quadrature points $\{x_i\}_{i=1}^M$ used in the greedy algorithm³. In this matrix-vector notation the expansion coefficients $c_\ell(\mu) = \langle e_\ell, h_\mu \rangle_{\mathcal{H}}$ for $\mathcal{P}_n h_\mu = \sum_{\ell=1}^n c_\ell(\mu) e_\ell$ now read

$$\mathbf{c} = \mathbf{V}^\dagger(\omega \circ \mathbf{h}_\mu),$$

³ For certain applications it may be desirable to evaluate the functions at a different set of points $\{y_i\}_{i=1}^{M'}$, see Section 4.3 for details.

where $\mathbf{c} = (c_1, \dots, c_n)^T \in \mathbb{C}^n$, \mathbf{V}^\dagger denotes the conjugate transpose of \mathbf{V} , and \circ denotes the standard Hadamard (Schur) or pointwise product between the vectors (in this case $\boldsymbol{\omega}$ and \mathbf{h}_μ) components. Therefore the reduced discrete integral (7) becomes

$$\begin{aligned} I_{\mathbf{d}}(i, j) &\approx \langle \mathcal{P}_n h_{\mu_i}, \mathcal{P}_n h_{\mu_j} \rangle_{\mathbf{d}} = \sum_{k=1}^n \langle h_{\mu_i}, e_k \rangle_{\mathbf{d}} \langle e_k, h_{\mu_j} \rangle_{\mathbf{d}} \\ &= \left(\mathbf{V}^\dagger (\boldsymbol{\omega} \circ \mathbf{h}_{\mu_i}) \right)^T \left(\mathbf{V}^\dagger (\boldsymbol{\omega} \circ \mathbf{h}_{\mu_j}) \right). \end{aligned} \quad (19)$$

Here we have reduced the overall *dimensionality* of the problem to n , but two issues remain. First, the number of $\{h_\mu(x_k)\}_{k=1}^M$ functional evaluations depends on M . Second, the dominant operation count for the approximation is seen to be of order $\mathcal{O}(Mn)$. We will return to these issues shortly.

3.2 The discrete empirical interpolation method (DEIM)

Recall from Sec. 3.1 that an approximation by projection, such as $\mathcal{P}_n h_\mu$, requires knowledge of projection coefficients, the computation or approximation of which relies on potentially expensive quadrature rules. The main idea behind the DEIM is to replace a potentially expensive approximation of projection coefficients by a relatively inexpensive interpolation, without sacrificing accuracy. The DEIM algorithm identifies a set of basis-specific interpolation points through a greedy selection criteria. When applied to several physical dimensions, irregular shaped domains, and rather generic parameterized spaces, the DEIM [4, 5] and EIM [2, 10, 9] points have shown remarkable robustness and efficiency.

Suppose we have a set of basis vectors. In particular, the algorithm applies to the reduced basis for the space F_n or the space \tilde{F}_m to be constructed in Sec. 3.3. The algorithm selects the DEIM points in physical space Ω and builds an associated interpolation matrix \mathbf{P} . This matrix \mathbf{P} interpolates the columns of the matrix \mathbf{V} introduced in Sec. 3.1.

Algorithm 2 (Selection of DEIM Points)

$[\mathbf{P}, \{p_i\}_{i=1}^n] = \text{DEIM}(\mathbf{V}, \{x_k\}_{k=1}^M)$ **Comment:** the column vectors of \mathbf{V} must be linearly independent

- (1) $j = \arg\max |\mathbf{e}_1|$ **Comment:** here $\arg\max$ takes a vector and returns the *index* of its largest entry
- (2) Set $\mathbf{U} = [\mathbf{e}_1]$, $\mathbf{P} = [\hat{\mathbf{e}}_j]$, $p_1 = x_j$ **Comment:** $\hat{\mathbf{e}}_j$ is a unit column vector with a single unit entry at index j
- (3) for $i = 2, \dots, n$ do
 - (4) Solve $(\mathbf{P}^T \mathbf{U}) \mathbf{c} = \mathbf{P}^T \mathbf{e}_i$ for \mathbf{c}
 - (5) $\mathbf{r} = \mathbf{e}_i - \mathbf{U} \mathbf{c}$
 - (6) $j = \arg\max |\mathbf{r}|$
 - (7) Set $\mathbf{U} = [\mathbf{U} \ \mathbf{r}]$, $\mathbf{P} = [\mathbf{P} \ \hat{\mathbf{e}}_j]$, $p_i = x_j$

The DEIM algorithm described above is nearly identical to the one given in Ref. [4] with the exception of Step 7 which is $\mathbf{U} = [\mathbf{U} \ \mathbf{e}_i]$ in that reference. In Appendix C we show these algorithms to be equivalent and, furthermore, we show a relationship between DEIM and Gauss-elimination with partial pivoting. Owing to the lower triangular form of $\mathbf{P}^T \mathbf{U}$, in Appendix B we show that Algorithm 2 has a reduced computational cost.

Remark 2 The columns of \mathbf{P} are unit vectors with a single unit entry at the location of the empirical interpolation points and zero elsewhere. For example, the first column of \mathbf{P} has a unit entry at the location $\arg\max |\mathbf{e}_1|$. Furthermore, $\mathbf{P}^T \mathbf{h}_\mu$ extracts an n -subvector which is exactly equivalent to evaluating h_μ at n empirical interpolation points $\{p_i\}_{i=1}^n$.

With empirical interpolation points $\{p_i\}_{i=1}^n$ we are now in a position to present an explicit and practical expression for the interpolant. Consider the approximation

$$(h_\mu(x_1), \dots, h_\mu(x_M))^T = \mathbf{h}_\mu \approx \sum_{i=1}^n c_i(\mu) \mathbf{e}_i = \mathbf{V} \mathbf{c}(\mu),$$

where $\mathbf{c}(\mu)$ are unknowns to be found via interpolation. Multiply both sides of this expression by \mathbf{P}^T to deduce that the interpolation step described in Eq. (8) is equivalent to

$$\mathbf{P}^T \mathbf{h}_\mu = \mathbf{P}^T \mathbf{V} \mathbf{c}(\mu),$$

which can be solved by (note that $\mathbf{P}^T \mathbf{V}$ is invertible when the columns of \mathbf{V} are linearly independent [4])

$$\mathbf{c}(\mu) = (\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T \mathbf{h}_\mu.$$

Hence, an explicit form for the DEIM approximation is

$$\mathbf{h}_\mu \approx \mathcal{I}_n[\mathbf{h}_\mu] := \mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T \mathbf{h}_\mu, \quad (20)$$

which is indeed an interpolant:

$$\mathbf{P}^T \mathcal{I}_n[\mathbf{h}_\mu] = \mathbf{P}^T (\mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T \mathbf{h}_\mu) = \mathbf{P}^T \mathbf{h}_\mu.$$

That is, the interpolant agrees with the original function at the interpolating points, as it should.

Substituting $\mathcal{I}_n[\mathbf{h}_\mu]$ from Eq. (20) into Eq. (6) gives

$$I_d(i, j) \approx \langle \mathcal{I}_n[\mathbf{h}_{\mu_i}], \mathcal{I}_n[\mathbf{h}_{\mu_j}] \rangle_d \approx \left((\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T \mathbf{h}_{\mu_i} \right)^T \left((\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T \mathbf{h}_{\mu_j} \right). \quad (21)$$

As we only have to evaluate $\mathbf{P}^T \mathbf{h}_\mu \in \mathbb{C}^n$, with $n < M$ in many applications, we have reduced the complexity (that is, the number of required functional evaluations) to $\mathcal{O}(n)$ (compare with (19)). However, due to the two matrix-vector multiplications an evaluation cost of $\mathcal{O}(n^2)$ has appeared. For many (if not all) interesting problems one should expect $\mathcal{O}(n^2) > \mathcal{O}(M)$. Such unacceptable scaling arises when approximating the individual functions h_μ instead of the products (integrands) $h_{\mu_i}^* h_{\mu_j}$ directly. We consider the generation of a reduced basis for products of functions next.

3.3 Two-step greedy approach for the approximation of product of target functions

$$\begin{array}{ccc} & \nearrow \text{Path \#1} \searrow & \\ \tilde{\mathcal{F}} & \xrightarrow[\text{Path \#2 (i)}]{F_{n^2}} & \tilde{F}_m \end{array} \quad \xrightarrow[\text{Path \#2 (ii)}]{F_m} \quad (22)$$

As discussed at the end of Secs. 3.1 and 3.2, we seek to approximate all possible integrands $g_{ij} = h_{\mu_i}^* h_{\mu_j} \in \tilde{\mathcal{F}}$, where $\tilde{\mathcal{F}}$ is a subset of the Hilbert space $\mathcal{H} = L^2(\Omega)$. A natural approach, which we refer to as a *direct* one, is to build a reduced basis for $\tilde{\mathcal{F}}$ through a greedy algorithm as described in Algorithm 1 and is marked as *Path #1* in (22). In that case the results of Secs. 3.1 and 3.2 are directly applicable, with \mathcal{F} and F_n replaced by $\tilde{\mathcal{F}}$ and \tilde{F}_m respectively. However, if a faithful approximation of \mathcal{F} requires a training set \mathcal{T}_K of size $\mathcal{O}(K)$ then an appropriate training set for a faithful representation of $\tilde{\mathcal{F}}$ would in principle be of size $\mathcal{O}(K^2)$. Since for large problems $\mathcal{O}(K)$ can already be computationally challenging, this direct approach might be unfeasible in some applications, even when building the basis is an offline and parallelizable calculation. In Section 5.2 we argue that this direct approach can be impractical. Advanced algorithms for sampling or building the training space, perhaps adaptively [14, 9, 15, 16], might overcome this issue, though, and allow for a direct approach to such large problems.

Here we advocate what we call the *two-step greedy approach* via *Path #2* in (22). The basic idea is to take advantage of the observation that all integrands can be accurately approximated by

$$h_{\mu_i}^* h_{\mu_j} \approx (\mathcal{P}_n h_{\mu_i})^* \mathcal{P}_n h_{\mu_j} = \left(\sum_{k=1}^n \langle e_k, h_{\mu_i} \rangle_d e_k \right)^* \left(\sum_{l=1}^n \langle e_l, h_{\mu_j} \rangle_d e_l \right),$$

which involves a sum of n^2 terms featuring products of functions purely of the form $e_k^* e_l$. We then carry out a second layer of dimensional reduction – a second greedy for the products $\{e_k^* e_l\}_{k,l=1}^n$ selected by the first greedy. The training space for this second greedy is of size $\mathcal{O}(n^2)$, as opposed to $\mathcal{O}(K^2)$ for the direct approach (with, usually, $n^2 \ll K^2$).

In more detail, our two-step greedy works as follows. In the first step (*Path #2 (i)* in (22)) we generate, through a first greedy as in Algorithm 1, a reduced basis – whose span we denote by F_n – approximating \mathcal{F} . Next one might construct a set F_{n^2} , consisting of all n^2 products $e_k^* e_l / \|e_k^* e_l\|_d$. Another option, and indeed the one used in our numerical experiments, is to use a training set $\mathcal{T}_n^2 := \{(\mu_i, \mu_j)\}_{i,j=1}^n$ and associated normalized products $h_{\mu_i}^* h_{\mu_j} / \|h_{\mu_i}^* h_{\mu_j}\|_d$ to define F_{n^2} , where $\{\mu_i\}_{i=1}^n$ are the greedy points identified by the first greedy algorithm when approximating \mathcal{F} by F_n . In other words, we do not necessarily use an orthonormal reduced basis for F_n (see Remark 3, below). In either case, we approximate F_{n^2} through a second RB-greedy (*Path 2 (ii)* in (22)), carried out again as in Algorithm (1) but with F_{n^2}

as the training space. The result is an orthonormal set of m reduced basis $\{\tilde{e}_i\}_{i=1}^m$ such that $\tilde{F}_m := \text{span}\{\tilde{e}_1, \dots, \tilde{e}_m\}$ accurately approximates F_{n^2} , and in turn the full set $\tilde{\mathcal{F}}$.

At the end of the day, either using a direct approach (*Path #1*) or our proposed two-step one (*Path #2*), given any integrand $g_{ij} = h_{\mu_i}^* h_{\mu_j} \in \tilde{\mathcal{F}}$ there is an accurate reduced order approximation $\tilde{\mathcal{P}}_m g_{ij}$ of the form

$$g_{ij} \approx \tilde{\mathcal{P}}_m g_{ij} := \sum_{k=1}^m \langle \tilde{e}_k, g_{ij} \rangle_{\mathbf{d}} \tilde{e}_k. \quad (23)$$

For definiteness, we choose the tolerance for the second greedy step to be equal to the first one (ϵ in Eq.(13)). To better explain the algorithm it will be useful to introduce new notation for elements of \mathcal{T}_n^2 : define $\tilde{\mu}_{k(i,j)} := (\mu_i, \mu_j)$ to be a re-indexing of a two dimensional array $\{(\mu_i, \mu_j)\}_{i,j=1}^n$ into a one dimensional array $\{\tilde{\mu}_k\}_{k=1}^{n^2}$, where k is a one-to-one function which takes two integers i and j and returns a unique integer $k(i, j)$.

Remark 3 The first greedy returns both the greedy points $\{\mu_i\}_{i=1}^n$, an orthonormal basis $\{e_i\}_{i=1}^n$, and the *greedy basis functions* $\{h_{\mu_i}\}_{i=1}^n$, serving as a non-orthonormal basis set for the same space F_n . As discussed above, from either $\{h_{\mu_i}\}_{i=1}^n$ or $\{e_i\}_{i=1}^n$ one can build the set F_{n^2} which is approximated below by Algorithm 3. In our numerical experiments we choose $\{h_{\mu_i}\}_{i=1}^n$ as basis, partly because we need not store the orthonormal basis vectors $\{e_i\}_{i=1}^n$ but only the greedy points $\{\mu_i\}_{i=1}^n$. Clearly the implementation and interpretation of Algorithm 3 will depend on this choice. For example, when using $\{h_{\mu_i}\}_{i=1}^n$ it makes sense to discuss the selected greedy parameter points $\tilde{\mu}_k$ while for $\{e_i\}_{i=1}^n$ no such interpretation holds; instead one should view Step 5d as returning a column index when F_{n^2} is thought of as an $M \times n^2$ matrix.

To summarize, the second step of our proposed two-step greedy algorithm to build \tilde{F}_m is

Algorithm 3 (Two-step RB-Greedy approximation for \tilde{F}_m)

$[\{\tilde{e}_\ell\}_{\ell=1}^m, \{\tilde{\mu}_\ell\}_{\ell=1}^m] = \text{Two-step RB-Greedy}(\epsilon, \{\mu_i\}_{i=1}^n, \{e_i\}_{i=1}^n)$

Comment: $\{e_i\}_{i=1}^n$ are either orthonormal or greedy basis functions (cf. remark 3)

- (1) Let $\mathcal{T}_n^2 = \{(\mu_i, \mu_j)\}_{i,j=1}^n$
- (2) Let $F_{n^2} = \{e_i^* e_j / \|e_i^* e_j\|_{\mathbf{d}}\}_{i,j=1}^n$
- (3) Set $m = 0$ and define $\sigma_0(F_{n^2}; \mathcal{H}) := 1$
- (4) Choose an arbitrary $g \in F_{n^2}$ and set $\tilde{e}_1 := g$ **Comment:** $\|g\|_{\mathbf{d}} = 1$
- (5) do, while $\sigma_m(F_{n^2}; \mathcal{H}) \geq \epsilon$
 - (a) $m = m + 1$
 - (b) $\sigma_m(g) := \|g - \tilde{\mathcal{P}}_m g\|_{\mathbf{d}}$ for all $g \in F_{n^2}$
 - (c) $\sigma_m(F_{n^2}; \mathcal{H}) = \sup_{g \in F_{n^2}} \{\sigma_m(g)\}$
 - (d) $\tilde{\mu}_{m+1} := \text{argsup}_{g \in F_{n^2}} \{\sigma_m(g)\}$ (greedy sweep)
 - (e) $\tilde{e}_{m+1} := g_{\tilde{\mu}_{m+1}} - \tilde{\mathcal{P}}_m g_{\tilde{\mu}_{m+1}}$ (Gram-Schmidt)
 - (f) $\tilde{e}_{m+1} := \tilde{e}_{m+1} / \|\tilde{e}_{m+1}\|_{\mathbf{d}}$ (normalization)

Paralleling the discussion in Sec. 3.1 and Sec. 3.2, a discrete form of the approximation (23) is accomplished by defining the matrix

$$\tilde{\mathbf{V}} = [\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_m] \in \mathbb{C}^{M \times m},$$

where the ℓ^{th} column is $\tilde{\mathbf{e}}_\ell = (\tilde{e}_\ell(x_1), \dots, \tilde{e}_\ell(x_M))^T$. Then in matrix-vector notation $\tilde{\mathcal{P}}_m \mathbf{g}$ is

$$\tilde{\mathcal{P}}_m \mathbf{g} = \tilde{\mathbf{V}} \tilde{\mathbf{V}}^\dagger (\omega \circ \mathbf{g}),$$

where we have dropped the indices i, j for brevity. Continuing, we build the interpolation matrix $\tilde{\mathbf{P}}$ from Alg. (2) with $\tilde{\mathbf{V}}$ as input. The DEIM interpolant is found by mimicking the steps in Sec. 3.2, which produces $\tilde{\mathcal{I}}_m[\mathbf{g}]$ satisfying

$$\mathbf{g} \approx \tilde{\mathcal{I}}_m[\mathbf{g}] := \tilde{\mathbf{V}} (\tilde{\mathbf{P}}^T \tilde{\mathbf{V}})^{-1} \tilde{\mathbf{P}}^T \mathbf{g}. \quad (24)$$

A discrete integral rule stemming from the above interpolant leads to an efficient and simple reduced order quadrature rule (10) and is the subject of the next Section.

4 Reduced order quadratures (ROQ)

In classical theory of Gaussian quadratures one seeks to maximize the exactness for the integration of polynomials. Here the goal is to empirically maximize the accuracy of inner products between elements of \mathcal{F} . Notice that an approximation of (discrete) inner products may be carried out in either a reduced space $F_n \approx \mathcal{F}$ or $\tilde{F}_m \approx \tilde{\mathcal{F}}$. In Sec. 3.2 we noted that working in F_n leads to a pessimistic $\mathcal{O}(n^2)$ cost for inner products (see the discussion below Eq. (21)). Thus, for applications where one needs fast evaluations of many scalar products the approximation space \tilde{F}_m is preferable. For the construction of \tilde{F}_m we refer to (22).

Section 4.1 focuses on a ROQ stemming from \tilde{F}_m , although only very minimal changes are needed to adapt Algorithm 4 for obvious variations. Indeed, one such variation is considered in our first numerical experiment. In Section 4.3 we consider extensions of the ROQ construction to situations where one might only be able to evaluate integrands at a set of points $\{y_i\}_{i=1}^{M'}$ which do not correspond to the nodes used in the quadrature rule $\{x_k, \omega_k\}_{k=1}^M$. In Section 4.2.1 we present the DEIM interpolation error estimates, followed by some ROQ error estimates in Section 4.2.2.

4.1 ROQ Algorithm

Suppose we are given a set of functions \mathcal{F} and an arbitrary quadrature rule $\{x_k, \omega_k\}_{k=1}^M$ for the discrete inner product (6). The following algorithm generates the ROQ nodes and weights.

Algorithm 4 (Construction of Reduced Order Quadratures)

$$\left[\{\tilde{p}_\ell, \omega_\ell^{\text{ROQ}}\}_{\ell=1}^m \right] = \text{ROQ}(\epsilon, \mathcal{T}_K, \{x_k, \omega_k\}_{k=1}^M)$$

(1) **Approximation of $\tilde{\mathcal{F}}$:** We have two options to approximate $\tilde{\mathcal{F}}$ by \tilde{F}_m , (cf. (22)).

Path #1 Consider the training set $\mathcal{T}_K^2 = \{(\mu_i, \mu_j)\}_{i,j=1}^K$ and apply the Greedy Algorithm 1, with obvious modifications, to generate $\tilde{F}_m = \text{span}\{\tilde{e}_\ell\}_{\ell=1}^m$ and greedy points $\{\tilde{\mu}_\ell\}_{\ell=1}^m \subset \mathcal{T}_K^2$ such that any product $h_{\mu_i}^* h_{\mu_j} \in \tilde{\mathcal{F}}$ can be approximated by its projection $\tilde{\mathcal{P}}_m$ as:

$$\left\| h_{\mu_i}^* h_{\mu_j} - \tilde{\mathcal{P}}_m(h_{\mu_i}^* h_{\mu_j}) \right\|_{\text{d}} \leq \epsilon, \quad \forall (\mu_i, \mu_j) \in \mathcal{T}_K^2.$$

Note: This direct approach can be expensive; we therefore advocate the following two-step procedure.

Path #2 (i) **Approximation of \mathcal{F} :** Consider the training set $\mathcal{T}_K = \{\mu_i\}_{i=1}^K$ and apply the Greedy Algorithm 1 to generate $F_n = \text{span}\{e_\ell\}_{\ell=1}^n$ and greedy points $\mathcal{T}_n = \{\mu_\ell\}_{\ell=1}^n \subset \mathcal{T}_K$ such that any $h_\mu \in \mathcal{F}$ can be approximated by its projection \mathcal{P}_n as:

$$\left\| h_\mu - \mathcal{P}_n h_\mu \right\|_{\text{d}} \leq \epsilon, \quad \forall \mu \in \mathcal{T}_K.$$

Training set \mathcal{T}_n^2 : Define the training set $\mathcal{T}_n^2 = \mathcal{T}_n \times \mathcal{T}_n = \{(\mu_i, \mu_j)\}_{i,j=1}^n$, and the set

$$F_{n^2} = \{e_i^* e_j / \|e_i^* e_j\|_{\text{d}}\}_{i,j=1}^n.$$

Another choice for F_{n^2} is $\{h_{\mu_i}^* h_{\mu_j} / \|h_{\mu_i}^* h_{\mu_j}\|_{\text{d}}\}_{i,j=1}^n$, we refer to Remark 3 for justification.

(ii) Apply the Greedy Algorithm 3 to generate $\tilde{F}_m = \text{span}\{\tilde{e}_\ell\}_{\ell=1}^m$ and $\{\tilde{\mu}_\ell\}_{\ell=1}^m \subset \mathcal{T}_n^2 \subset \mathcal{T}_K^2$ such that any product $e_i^* e_j \in F_{n^2}$ can be approximated by its projection $\tilde{\mathcal{P}}_m$ as:

$$\left\| e_i^* e_j - \tilde{\mathcal{P}}_m(e_i^* e_j) \right\|_{\text{d}} \leq \epsilon.$$

(2) Define the projection matrix

$$\tilde{\mathbf{V}} = [\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_m] \in \mathbb{C}^{M \times m},$$

where, for example, the ℓ^{th} column is $\tilde{\mathbf{e}}_\ell = (\tilde{e}_\ell(x_1), \dots, \tilde{e}_\ell(x_M))^T$.

(3) **Generation of empirical interpolation points:** Apply the DEIM Algorithm 2 with $\tilde{\mathbf{V}}$ as an input to compute the DEIM points $\{\tilde{p}_\ell\}_{\ell=1}^m \subset \{x_i\}_{i=1}^M$ and the interpolation matrix $\tilde{\mathbf{P}}$.

(4) Compute the ROQ weights by

$$\left(\omega^{\text{ROQ}}\right)^T := \omega^T \tilde{\mathbf{V}} (\tilde{\mathbf{P}}^T \tilde{\mathbf{V}})^{-1}. \quad (25)$$

Step 4 deserves a bit more explanation. With the interpolation matrix $\tilde{\mathbf{P}}$ from Step 3 we first build an expression for the interpolant of some product $g \in \tilde{\mathcal{F}}$ using Eq. (24). The reduced order weights are then found by substituting the DEIM interpolant into Eq. (6) to produce the desired approximation with $\mathbf{g} = \mathbf{h}_{\mu_i}^* \circ \mathbf{h}_{\mu_j}$,

$$I_d(i, j) = \omega^T \mathbf{g} \approx \omega^T \tilde{\mathcal{I}}_m[\mathbf{g}] = \left[\omega^T \tilde{\mathbf{V}} (\tilde{\mathbf{P}}^T \tilde{\mathbf{V}})^{-1} \right] \tilde{\mathbf{P}}^T \mathbf{g} = \sum_{i=1}^m \omega_i^{\text{ROQ}} g(\tilde{p}_i). \quad (26)$$

Notice that the term $\tilde{\mathbf{P}}^T \mathbf{g}$ is just evaluation of g at the DEIM points $\{\tilde{p}_i\}_{i=1}^m$ (which need not be ordered, but can be if desired). Furthermore, the weights are explicitly parameter independent and for less accuracy we might consider using the first $m' \leq m$ functional evaluations of g

$$\omega^T \tilde{\mathcal{I}}_{m'}[\mathbf{g}] \approx \sum_{i=1}^{m'} \alpha_i^{\text{ROQ}} g(\tilde{p}_i),$$

which, as this is equivalent to carrying out the computation in an m' -dimensional subspace of \tilde{F}_m , is also a valid integration rule. In general the weights α_i^{ROQ} associated with this m' point quadrature rule will not be a subset of the weights ω_i^{ROQ} computed for the m point rule. Notice that the hierarchical nature of the empirical interpolation set allows one to build application-specific nested quadratures of arbitrary order and depth (i.e. more than one embedded method). For example, suppose $m = 20$; one might combine the 10-point, 15-point, and 20-point ROQ rules to yield a nested quadrature scheme.

Next we state some results which highlight the importance of using an accurate quadrature rule $\{x_k, \omega_k\}_{k=1}^M$, these will be verified in Section 5.1.

Theorem 1 *Integration of the basis functions $\{\tilde{e}_\ell\}_{\ell=1}^m$ with either of the quadrature rules $\{x_i, \omega_i\}_{i=1}^M$ or $\{\tilde{p}_\ell, \omega_\ell^{\text{ROQ}}\}_{\ell=1}^m$ yield identical results.*

Proof The integration of basis $\{\tilde{e}_i\}_{i=1}^m$ using ROQ (cf. (25), (26) in vector notation) is

$$I_{\text{ROQ}} = \left(\omega^{\text{ROQ}}\right)^T \tilde{\mathbf{P}}^T \tilde{\mathbf{e}}_i = \omega^T \tilde{\mathbf{V}} (\tilde{\mathbf{P}}^T \tilde{\mathbf{V}})^{-1} \tilde{\mathbf{P}}^T \tilde{\mathbf{e}}_i, \quad i = 1, \dots, m.$$

Recalling $\tilde{\mathbf{V}} = [\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_m]$, integration of all the basis $\{\tilde{e}_i\}_{i=1}^m$ can be written as:

$$\left(\omega^{\text{ROQ}}\right)^T \tilde{\mathbf{P}}^T \tilde{\mathbf{V}} = \omega^T \tilde{\mathbf{V}},$$

where $\omega^T \tilde{\mathbf{V}}$ is precisely the integration of all the basis vectors $\{\tilde{e}_i\}_{i=1}^m$ with the quadrature rule $\{x_k, \omega_k\}_{k=1}^M$.

Corollary 1 (Consistency of ROQ rule) *When $m = M$ the quadrature rules $\{x_i, \omega_i\}_{i=1}^M$ and $\{\tilde{p}_\ell, \omega_\ell^{\text{ROQ}}\}_{\ell=1}^m$ are identical up-to ordering of the quadrature points.*

Proof Using Theorem 1, we have

$$\left(\omega^{\text{ROQ}}\right)^T \tilde{\mathbf{P}}^T \tilde{\mathbf{V}} = \omega^T \tilde{\mathbf{V}},$$

apply $\tilde{\mathbf{V}}^{-1}$ on the right to deduce $\left(\omega^{\text{ROQ}}\right)^T \tilde{\mathbf{P}}^T = \omega^T$.

Remark 4 If the quadrature rule $\{x_i, \omega_i\}_{i=1}^M$ integrates the basis functions $\{\tilde{e}_\ell\}_{\ell=1}^m$ exactly then ROQ built using Algorithm 4 also integrates those basis functions exactly.

Remark 5 As a special case consider $\tilde{\mathcal{F}}$ to be a space of polynomials of degree $m - 1$ with weight $W = 1$ and suppose the basis is specified by an ordered set of Legendre polynomials. Then the m -point ROQ rule will exactly integrate any polynomial of at most degree $m - 1$ provided the quadrature rule $\{x_k, \omega_k\}_{k=1}^M$ is at least $m - 1$ exact.

4.2 Convergence Estimates

4.2.1 DEIM Convergence

We shall first state a DEIM interpolation error estimate with respect to the discrete L^∞ -norm. Then we will recall the DEIM error estimate with respect to the discrete L^2 -norm from [4] and, finally, combine this estimate with those for the RB-greedy basis construction from [13].

Theorem 2 (discrete empirical interpolation method) *Let the set of reduced basis $\{e_\ell\}_{\ell=1}^n$ be orthonormal with respect to the discrete inner product defined in (3) and $h_\mu^{\text{opt}} \in F_n$ be the optimal approximation of h_μ with respect to the discrete L^∞ -norm. Then for every $\mu \in \mathcal{P}$ and x_k , with $k = 1, \dots, M$,*

$$\max_{1 \leq k \leq M} |h_\mu(x_k) - \mathcal{I}_n[h_\mu](x_k)| \leq (1 + \Lambda_{n,\infty}) \max_{1 \leq k \leq M} |h_\mu(x_k) - h_\mu^{\text{opt}}(x_k)|, \quad (27)$$

where $\Lambda_{n,\infty} = \|\mathcal{I}_n\|_\infty = \|\mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T\|_\infty$ denotes the Lebesgue constant.

In the case of the discrete L^2 -norm, $h_\mu^{\text{opt}} = \mathcal{P}_n h_\mu$ and we have for all $\mu \in \mathcal{P}$

$$\|h_\mu - \mathcal{I}_n[h_\mu]\|_{\mathbf{d}} \leq \Lambda_{n,2} \|h_\mu - \mathcal{P}_n h_\mu\|_{\mathbf{d}}, \quad (28)$$

where $\Lambda_{n,2} = \|\mathcal{I}_n\|_2 = \|\mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T\|_2$ ⁴.

Proof The proof of (27) and (28) are straightforward and are omitted. Furthermore, (28) follows from [4, Lemma 3.2], but we state it for completeness. Since $\mathcal{I}_n[h_\mu^{\text{opt}}](x_k) = h_\mu^{\text{opt}}(x_k)$ for $k = 1, \dots, M$, we get

$$\begin{aligned} \|h_\mu - \mathcal{I}_n[h_\mu]\|_{\mathbf{d}} &= \|(1 - \mathcal{I}_n)[h_\mu - h_\mu^{\text{opt}}]\|_{\mathbf{d}} \\ &\leq \|\mathcal{I}_n\|_2 \|h_\mu - h_\mu^{\text{opt}}\|_{\mathbf{d}}, \end{aligned}$$

where the last equality follows from the fact that $\|1 - \mathcal{I}_n\|_2 = \|\mathcal{I}_n\|_2$ (see, for example [22, 23]).

Remark 6 Notice that the DEIM Algorithm 2 seeks to minimize the interpolation error as measured by a discrete L^∞ -norm (cf. Step 6, Algorithm 2), for which an error bound of the type (27) is closely related. The RB-greedy Algorithm 1, however, exactly minimizes the term $\|h_\mu - \mathcal{P}_n h_\mu\|_{\mathbf{d}}$ on the right hand side of the error bound (28) and is therefore computable. Observing that $\Lambda_{n,2} \leq \|\mathbf{V}\|_2 \|(\mathbf{P}^T \mathbf{V})^{-1}\|_2$ (because $\|\mathbf{P}^T\|_2 = 1$), one might consider directly minimizing the error $\|h_\mu - \mathcal{I}_n[h_\mu]\|_{\mathbf{d}}$ by minimizing the norm $\|(\mathbf{P}^T \mathbf{V})^{-1}\|_2$. Indeed, whenever $\mathbf{V}^\dagger \mathbf{V} = \mathbf{1}_{n \times n}$ we have exactly $\Lambda_{n,2} = \|(\mathbf{P}^T \mathbf{V})^{-1}\|_2$ which is more efficient to compute than $\|\mathbf{V}(\mathbf{P}^T \mathbf{V})^{-1} \mathbf{P}^T\|_2$. Throughout the numerical experiments section 5, in addition to the error bound (28) we monitor

$$\|h_\mu - \mathcal{I}_n[h_\mu]\|_{\mathbf{d}} \leq \|\mathbf{V}\|_2 \|(\mathbf{P}^T \mathbf{V})^{-1}\|_2 \|h_\mu - \mathcal{P}_n h_\mu\|_{\mathbf{d}}, \quad (29)$$

which we have written here for future reference.

Corollary 2 (RB-Greedy-DEIM error bounds) *The following practical RB-greedy-DEIM estimate holds*

$$\|h_\mu - \mathcal{I}_n[h_\mu]\|_{\mathbf{d}} \lesssim \Lambda_{n,2} \epsilon, \quad \forall h_\mu \in \mathcal{F}.$$

Furthermore, if the Kolmogorov n -width decays exponentially (with order α) as in (15), then the error (28) decays with the same order,

$$\|h_\mu - \mathcal{I}_n[h_\mu]\|_{\mathbf{d}} \lesssim \sqrt{2C} \Lambda_{n,2} e^{-c_1 n^\alpha},$$

whereas if the Kolmogorov n -width decays algebraically with order β as in (17), then (28) again decays with the same order β

$$\|h_\mu - \mathcal{I}_n[h_\mu]\|_{\mathbf{d}} \lesssim 2^{5\beta+1} C_a \Lambda_{n,2} n^{-\beta},$$

where $C_a > 0$ and $\beta > 0$ are some constants.

Proof Consider (14) and (28) and apply (16) and (18).

⁴ See Appendix A for the definition of the 2-norm $\|\mathbf{A}\|_2$ of an operator \mathbf{A} .

4.2.2 ROQ Convergence

Next we present the error estimate for the case in which the reduced basis $\{\tilde{e}_\ell\}_{\ell=1}^m$ for $\tilde{\mathcal{F}}$ are generated through a direct approach via *Path #1* in (22).

Theorem 3 (approximation error ROQ (direct approach)) *Let $\epsilon > 0$ be an error tolerance and $\tilde{\mathcal{P}}_m : \tilde{\mathcal{F}} \rightarrow \tilde{F}_m$ be the projection operator defined in Algorithm 4 (Path #1), with accuracy $\epsilon > 0$ in the discrete L^2 -norm $\|\cdot\|_{\mathbf{d}}$. Furthermore, let $\tilde{\mathcal{I}}_m : \tilde{\mathcal{F}} \rightarrow \tilde{F}_m$ be a DEIM operator built using Step (3) (RB-greedy-DEIM) of Algorithm 4. Given two arbitrary functions $h_{\mu_i}, h_{\mu_j} \in \mathcal{F}$, the following estimate holds*

$$|I_{\mathbf{d}}(i, j) - I_{\text{ROQ}}(i, j)| \lesssim \epsilon \left(\|1\|_{\mathbf{d}} \|\tilde{\mathcal{I}}_m\|_2 \right) \|h_{\mu_i}^* h_{\mu_j}\|_{\mathbf{d}}.$$

Proof We recall that given $h_{\mu_i}, h_{\mu_j} \in \mathcal{F}$, with $\mu_i, \mu_j \in \mathcal{P}$, we have

$$I_{\mathbf{d}}(i, j) := \langle h_{\mu_i}, h_{\mu_j} \rangle_{\mathbf{d}} = \sum_{k=1}^M \omega_k h_{\mu_i}^*(x_k) h_{\mu_j}(x_k).$$

On the other hand, we have the following expression for the ROQ integral

$$I_{\text{ROQ}}(i, j) = \sum_{k=1}^M \omega_k \tilde{\mathcal{I}}_m[h_{\mu_i}^* h_{\mu_j}](x_k).$$

An application of Cauchy-Schwarz implies

$$\begin{aligned} |I_{\mathbf{d}}(i, j) - I_{\text{ROQ}}(i, j)| &\leq \sum_{k=1}^M \left| \omega_k \left(h_{\mu_i}^*(x_k) h_{\mu_j}(x_k) - \tilde{\mathcal{I}}_m[h_{\mu_i}^* h_{\mu_j}](x_k) \right) \right| \\ &\leq \|1\|_{\mathbf{d}} \left\| h_{\mu_i}^* h_{\mu_j} - \tilde{\mathcal{I}}_m[h_{\mu_i}^* h_{\mu_j}] \right\|_{\mathbf{d}}. \end{aligned}$$

Next the fact that $\tilde{\mathcal{I}}_m \left[\tilde{\mathcal{P}}_m[h_{\mu_i}^* h_{\mu_j}] \right] = \tilde{\mathcal{P}}_m[h_{\mu_i}^* h_{\mu_j}]$, implies

$$h_{\mu_i}^* h_{\mu_j} - \tilde{\mathcal{I}}_m[h_{\mu_i}^* h_{\mu_j}] = (1 - \tilde{\mathcal{I}}_m) \left[h_{\mu_i}^* h_{\mu_j} - \tilde{\mathcal{P}}_m[h_{\mu_i}^* h_{\mu_j}] \right]$$

Finally using $\|1 - \tilde{\mathcal{I}}_m\|_2 = \|\tilde{\mathcal{I}}_m\|_2$ (cf. [22, 23]), we get

$$\begin{aligned} |I_{\mathbf{d}}(i, j) - I_{\text{ROQ}}(i, j)| &\leq \|1\|_{\mathbf{d}} \|\tilde{\mathcal{I}}_m\|_2 \left\| h_{\mu_i}^* h_{\mu_j} - \tilde{\mathcal{P}}_m[h_{\mu_i}^* h_{\mu_j}] \right\|_{\mathbf{d}} \\ &\leq \|1\|_{\mathbf{d}} \|\tilde{\mathcal{I}}_m\|_2 \sigma_m \left(\tilde{\mathcal{F}}; \mathcal{H} \right) \|h_{\mu_i}^* h_{\mu_j}\|_{\mathbf{d}}. \end{aligned}$$

Then the greedy estimate *Path #1* in Algorithm 4 gives the required estimate.

Our numerical examples indicate the ROQ built from a two-step greedy (using *Path #2*) has an error bound as in Theorem 3, but proving such bound is beyond the scope of this paper.

Remark 7 (FLOP-Count) Let K be the number of parameters in the training space to approximate \mathcal{F} by F_n . If \hat{m} and m denote the number of reduced basis for direct (*Path #1*) and two-step greedy (*Path #2*) respectively, then numerically we observe that $m \approx \hat{m}$. The cost of ROQ Algorithm 4 using a direct and a two-step greedy is $\mathcal{O}(K^2 M \hat{m})$ and $\mathcal{O}(n^2 M m + K M n)$, respectively. This implies that for $K \gg n$ (as is observed in our numerical experiments), the two-step greedy is highly efficient. For the details of FLOP counts we refer to Appendix B.

4.3 ROQ at new quadrature points

We now consider the scenario in which the integrands are available at a set of points which is different from the ones used to build the reduced basis. A typical example of this would be the case in which one is allowed to choose the sampling points and quadrature rule to build the basis, for example in RB-greedy algorithm, but not the points at which the data is available for computing the ROQ.

We continue denoting the set of points and quadrature rule used to build the basis as $\{x_k, \omega_k\}_{k=1}^M$, and we denote those for which the online data are available as $\{y_i, \tau_i\}_{i=1}^{M'}$. It is usually cheaper to generate $\tilde{F}_m = \text{span}\{\tilde{e}_\ell\}_{\ell=1}^m$ and $\{\tilde{\mu}_\ell\}_{\ell=1}^m$ using $\{x_k, \omega_k\}_{k=1}^M$ (for example using Step 1, 2 of Algorithm 4) than $\{y_i, \tau_i\}_{i=1}^{M'}$, if we can choose the former. Therefore, it is not unreasonable to assume, for simplicity, that $M' > M$. Our goal here is to generate an ROQ with respect to $\{y_i, \tau_i\}_{i=1}^{M'}$ under this scenario.

Algorithm 5 (Construction of ROQ for $\{y_i, \tau_i\}_{i=1}^{M'}$)
 $[\{\tilde{p}_\ell, \tau_\ell^{\text{ROQ}}\}_{\ell=1}^m] = \text{ROQ-NEW}(\{\tilde{\mu}_\ell\}_{\ell=1}^m, \{y_i, \tau_i\}_{i=1}^{M'})$

- (1) Orthogonalize the (ordered) product of functions $\{g_{\tilde{\mu}_\ell}\}_{\ell=1}^m$ using the quadrature rule $\{y_i, \tau_i\}_{i=1}^{M'}$ and let the resulting vectors form the columns of a matrix $\tilde{\mathbf{V}}$. That is,

$$\tilde{\mathbf{V}} = [\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_m] \in \mathbb{C}^{M' \times m},$$

where, for example, the ℓ^{th} column is $\tilde{\mathbf{e}}_\ell = (\tilde{e}_\ell(y_1), \dots, \tilde{e}_\ell(y_{M'}))^T$.

- (2) Apply Steps 3 and 4 from Algorithm 4 to generate the ROQ points $\{\tilde{p}_\ell\}_{\ell=1}^m \subset \{y_i\}_{i=1}^{M'}$ and the interpolation matrix $\tilde{\mathbf{P}}$. Then the ROQ weights are given by

$$(\tau^{\text{ROQ}})^T := \tau^T \tilde{\mathbf{V}} (\tilde{\mathbf{P}}^T \tilde{\mathbf{V}})^{-1}.$$

Accuracy and conditioning of the resulting ROQ rule are guaranteed provided (i) we are able to carry out orthogonalization in Step 1 up to a certain tolerance and (ii) the new quadrature rule is accurate enough such that $\|h_\mu - \mathcal{P}_n h_\mu\|_\tau \lesssim \epsilon$.

5 Numerical examples

Here we discuss two sets of numerical experiments. As a first example we compare ROQ using orthogonal polynomials on the interval $[-1, 1]$ as basis (that is, the basis is not built through a greedy approach) with the well known case of Gaussian quadratures and, in particular, integration of Runge's function. This study reveals several important features: i) ROQ is well conditioned, ii) ROQ and Gaussian quadrature rules have a similar point and weight distribution, iii) the importance of an accurate quadrature rule $\{x_i, \omega_i\}_{i=1}^M$ (cf. Theorem 1).

Next we turn to a challenging example drawn from gravitational wave physics, namely the calculation of overlaps between "chirp" gravitational waves for compact binary coalescences. Details on implementation, error bounds, and a discussion of computational cost are provided, as well as the physical motivation of the problem. In particular, we find that ROQs exhibit exponential asymptotic convergence in the calculation of overlaps, with a profile of the error function in terms of quadrature points qualitatively similar to the greedy error curve of the training space in terms of the number of reduced basis elements. We emphasize that this exponential convergence in computing overlaps is, as expected, present even when the function is sampled at equally spaced quadrature points. In particular, this means that reduced order quadratures can be used as an efficient down-sampling strategy when function evaluations from experimental observations are given at (for example) fixed intervals. We also show the staggering offline savings of our two-step greedy approach compared to a direct one.

5.1 Comparison with Gauss-Legendre quadratures

In order to highlight the essential features of our approach, in this subsection we use Legendre polynomials $\{P_\ell(x)\}_{\ell=1}^m$ defined on $x \in [-1, 1]$ as an orthonormal basis (with respect to $W(x) = 1$) instead of a reduced basis generated in Step 1 Algorithm 4. Steps 3 and 4 of Algorithm 4 require specifying an underlying quadrature rule, which in this subsection will be either an M point Gauss-Legendre or extended trapezoidal rule. The resulting ROQ point and weight distribution (see Subsection 5.1.1), conditioning (see Subsection 5.1.2) and application to Runge's function (see Subsection 5.1.3) are considered next.

5.1.1 Point and weight distribution of ROQ

To build the ROQ points and weights an extended trapezoidal rule is used and all $m = 24$ Legendre polynomials are sampled at $M = 1,000$ equidistant points. Fig. 1(a) compares the distribution of points and weights for the Gauss-Legendre quadrature (red dots) and ROQ (blue crosses). Both weight and point distributions are similar, and in particular the points are not equally spaced but instead cluster towards the end points at $x = \pm 1$. One of the ROQ weights is equal to -0.00496089441576999 at 0.775775775775776 . Negative weights can be problematic in quadrature rules due to roundoff errors leading to poorly conditioned formula [24]. We turn to this issue next.

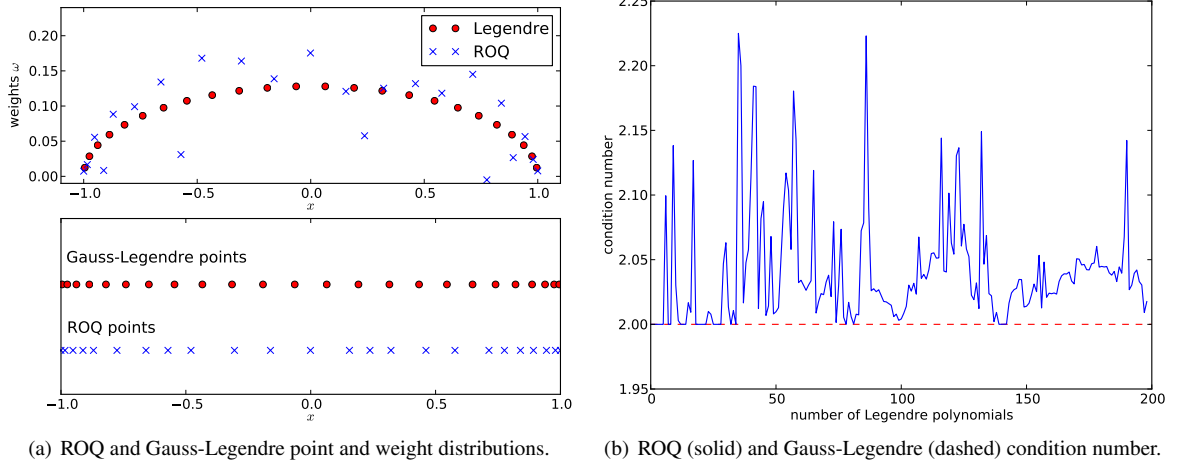


Fig. 1 The top left scatter plot shows the weight ω_k and point distributions for Gauss-Legendre and ROQ using $m = 24$ Legendre polynomials as a basis (see Subsection 5.1.1 for more details), whereas the bottom left one shows just the point locations. The right plot gives the condition number $\sum_{k=1}^m |\omega_k|$ for an m -point Gauss-Legendre and ROQ rule, with $m \in [2, 200]$ (see Subsection 5.1.2 for more details).

5.1.2 Conditioning of ROQ

Consider the absolute condition number $\sum_k |\omega_k|$ of a quadrature formula which, on the integration domain $x \in [-1, 1]$, is exactly 2 if all the weights ω_k are positive as can be seen by integration of unity. If some weights are negative then the condition number is necessarily larger than 2. A classic example of ill-conditioning is the Newton-Cotes rule, where $\sum_k |\omega_k| > 2$ and grows without bound when using more than 8 nodes [24]. Fig. 1(b) shows the absolute condition number $\sum_k |\omega_k|$ for a Gauss-Legendre (dashed red line) and reduced order (solid blue line) quadrature. The dashed red line shows the optimal value of 2, which is achieved by the Gauss-Legendre quadrature case. The condition number for ROQ remains below 2.25 for the first 200 Legendre polynomials (see Subsection 5.1.1) and has no noticeable growth trend. This plot was generated using $M = 1,000$ equally spaced points on $x \in [-1, 1]$; this resolution is sufficient for demonstrating well conditioning, but the resulting ROQ integration rule would be of low accuracy. We have considered up to 100,000 points, always obtaining behavior similar to that one of Fig. 1(b).

5.1.3 Integration of Runge's function

Many general features of ROQ can be ascertained by integration of Runge's function

$$\int_{-1}^1 \frac{1}{(1+x^2)} dx = 2 \tan^{-1}(1),$$

which is a classic example of oscillatory errors due to high-order interpolation. In the following experiment, integration of Runge's function is carried out with three different ROQ rules built from a basis consisting of m Legendre polynomials. In each case the Legendre polynomials are evaluated at $M_1 = 10,000$ equally spaced and $M_2 = 400$ Gauss-Legendre points, and the integration errors are computed as $|2 \tan^{-1}(1) - I_{\text{ROQ}}|$.

First, integration of Runge's function is carried out using an m -point ROQ-trapezoidal rule which is built from an M_1 -point extended trapezoidal rule and the m DEIM point selections. The resulting integration error is given by the dash-dot black line in Fig. 2. Notice that the ROQ-trapezoidal error curve flattens out around 10^{-9} . Beyond this value the 10,000 point trapezoidal rule is unable to accurately integrate the Legendre polynomials of degree higher than ~ 20 , thereby limiting the ROQ accuracy (see Theorem 1).

Next we show how accuracy may be restored in this particular case. Notice that the factor $\omega^T \tilde{\mathbf{V}}$ in Eq. (25) for the ROQ weights is precisely the numerical integration of m Legendre polynomials. Furthermore, we know that the first Legendre polynomial integrates to $\sqrt{2}$ while all others integrate to zero. Hence, in this particular case we may integrate them "by hand" by setting $\omega^T \tilde{\mathbf{V}} = [\sqrt{2}, 0, \dots, 0]$, while continuing to generate interpolation points by way of the DEIM algorithm. The resulting ROQ rule, denoted in the figure as ROQ-equidistant, is then used to integrate Runge's function with an error given by the solid blue line in Fig. 2. This highlights the importance of accurate integration of the basis functions which, in a general setting where the basis vectors are selected by a greedy algorithm, should be carried out with a good underlying quadrature rule.

Lastly, an ROQ rule is constructed from a $M_2 = 400$ point Gauss-Legendre rule (dashed red line of Fig. 2) and shows excellent agreement with the ROQ-equidistant case. Of particular noteworthiness, all three ROQ rules clearly show exponential convergence with the number of quadrature points and no apparent issues stemming from Runge's phenomena. Notice that the dashed red and solid blue lines are nearly indistinguishable, confirming the expectation that if the basis functions are sampled densely enough (a statement which certainly depends on m) and integrated accurately enough by the underlying quadrature rule, the resulting rules perform similarly.




Fig. 2 Convergence of the integral $\int_{-1}^1 (1+x^2)^{-1} dx$ to $2 \tan^{-1}(1)$ using three different ROQ rules whose generation (using Legendre polynomials as basis) is described in Sec. 5.1.3. Integration error from a ROQ-trapezoidal constructed with a $M_1 = 10,000$ point extended trapezoidal rule is given by the dash-dot black line. Due to the extended trapezoidal rule's inability to accurately integrate higher order Legendre polynomials this curve flattens out around 10^{-9} . If we *exactly* integrate these basis functions by setting $\omega^T \tilde{\mathbf{V}} = [\sqrt{2}, 0, \dots, 0]$ in Eq. (25) while continuing to use the ROQ-trapezoidal's quadrature points the accuracy is restored as seen from the ROQ-equidistant error curve (solid blue). Integration error from an ROQ constructed with a $M_2 = 400$ point Gauss-Legendre is given by the dashed red line.

5.2 Gravitational waves

The inspiral and merger of neutron stars (NSs), black holes (BHs) or mixed pairs, known generically as compact binary coalescences, is believed to be one of the main sources of gravitational waves to be detected by the upcoming generation of earth-based observatories [25, 26, 27, 28, 29, 30]. A passing wave distorts the length between any two points separated by a non-zero spatial distance. These distortions are measured by gravitational wave detectors as an effective strain $h_\mu(x)$, where x denotes either time or frequency. In this problem the parameter μ depends on quantities such as the mass and spin of the compact objects, or detector orientation. In principle h_μ is obtained by numerically solving Einstein's

equations, which are a set of parameterized quasilinear hyperbolic-elliptic equations [31, 32, 33, 34]. However, due to the cost of these simulations and the properties of the target system, simplified models are often used.

Here we focus on testing our proposed ROQ on the post-Newtonian (PN) approximation [35], under which the gravitational waves in the leading order Stationary Phase Approximation (SPA) are given by [36, 37, 38],

$$h_{\mathcal{M}_c}(f) = \mathcal{A}f^{-7/6} \cdot \exp\left(i\left\{-\frac{\pi}{4} + \frac{3}{128}\left(\pi \cdot \frac{G}{c^3} \cdot f \cdot \mathcal{M}_c\right)^{-5/3}\right\}\right), \quad (30)$$

where f denotes the frequency, G Newton's gravitational constant [39], c the speed of light, and \mathcal{A} an overall amplitude that depends on quantities such as the distance to the compact binary coalescence. We will refer to Eq. (30) as a gravitational *waveform*, two representative examples are shown in Fig. 3. The single parameter μ in this model is the *chirp mass* \mathcal{M}_c , which in turn is a function of the compact objects' masses $\{m_i\}_{i=1}^2$, $\mathcal{M}_c := (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$ [38]. All quantities here use MKS units.




Fig. 3 Gravitational waveforms h_A (left) and h_B (right) correspond to the smallest ($A = 2.611651689888372M_\odot$) and largest ($B = 26.11651689888372M_\odot$) chirp masses of the training set \mathcal{T}_K defined in (32). Waveforms are complex valued, but for clarity we do not show the imaginary part of h_A .

There are various reasons one might need to repeatedly compute inner products between waveforms. One setting where this occurs is in gravitational wave searches using matched filtering [40, 41, 42, 38]. Given data s recorded by a detector, a matched filtering search is carried out by computing all possible inner products of the data s against waveforms drawn from a (usually large) catalog. It can be shown [43, 44] that the optimal matched filtering strategy requires one to compute inner products with a weight given by the reciprocal of the detector's power spectral density $S(f) = W^{-1}$. As a representative example we use $S(f)$ from the initial Laser Interferometer Gravitational Wave Observatory (LIGO) [45], which can be modeled by the curve [46]

$$S(y) = 9 \times 10^{-46} \left[(4.49y)^{-56} + 0.16y^{-4.52} + 0.52 + 0.32 \cdot y^2 \right], \quad y = \frac{f}{150\text{Hz}}. \quad (31)$$

To ensure that no signals are missed and that the parameters are correctly estimated many inner products are needed for each segment of data [18, 47, 48, 49]. These significant computational costs motivate a need for faster ways of performing these integrals, often in real-time [50, 51, 52].

In the following series of experiments we consider gravitational waves of the form given by Eq. (30), generated by inspiralling binary black hole (BBH) systems with mass components in the interval $m_i \in [3, 30]M_\odot$, where $1M_\odot = 1.98892 \times 10^{30} \text{Kg}$ is the astrophysical unit of mass known as a *solar mass*. Therefore the parameter domain under consideration is $\mathcal{P} = [A, B]$, where $A = 2.611651689888372M_\odot$ and $B = 26.11651689888372M_\odot$. The physical domain frequency interval is given by $\Omega = [40, 366.3383434841933] \text{Hz}$. Here the lower value is set by the Earth's seismic noise which decreases the detector's sensitivity. The upper bound is set, roughly speaking, by the limitations of the PN approximation which breaks down at frequencies corresponding to the *Innermost Stable Circular Orbit* (ISCO).

While this is a mass dependent frequency, for simplicity we choose the upper limit as the maximum over the mass range here considered.

Our primary goal in the forthcoming sections is to describe the construction and performance of an ROQ rule for computing inner products between gravitational waveforms given by Eq. (30). These steps are conveniently summarized in Algorithm 4. The first step, an approximation for $\tilde{\mathcal{F}} \approx \tilde{F}_m$, can be carried out either by means of a direct greedy (Path #1) or two-step greedy (Path #2). Due to a *significantly* reduced offline cost we advocate the two-step greedy approach, which in turn requires an intermediate approximation $\mathcal{F} \approx F_n$. Results for this intermediate approximation are considered in Section 5.2.1. We contemporaneously remark on the performance of the DEIM algorithm when applied to single functions to facilitate a better understanding of this algorithm in a simpler setting. In Sections 5.2.2 and 5.2.3 we complete the second half of the two-step greedy approximation, identify the interpolation points and finally generate an ROQ rule. In this section we also answer the fundamental question: *Do the direct and two-step greedy approaches give comparable results and is the former much more expensive than the latter?* As we will see, we have strong numerical evidence that, as predicted by our estimates, the two-step greedy gives enormous computational savings without sacrificing the accuracy or compactness of the reduced basis.

5.2.1 RB-greedy and DEIM for single functions

Here we consider experiments focused on approximations of the set of functions $\mathcal{F} = \{h_{\mathcal{M}_c}(f) : \mathcal{M}_c \in \mathcal{P}\}$; with $h_{\mathcal{M}_c}$ as in Eq. (30) and the range $\mathcal{P} = [A, B]$ for the chirp mass \mathcal{M}_c —serving here as the parameter μ —as described above. We find the DEIM points and numerically confirm the DEIM error bounds from Sec. 4.2.2. This subsection is also precursor to Sections 5.2.2 and 5.2.3, where ROQs are constructed for waveform inner products.

Reduced Basis: We begin by building a reduced basis space F_n approximating \mathcal{F} by using the RB-greedy approach. Since the number of cycles is proportional to an inverse power of \mathcal{M}_c , we have found it advantageous to build the training space with a logarithmic spacing between the samples, thereby clustering more points at low values of \mathcal{M}_c . More precisely, a training set of size K is here given by

$$\mathcal{T}_K = \left\{ A \left(\frac{B}{A} \right)^{\frac{i}{K-1}} \mid i = 0, \dots, K-1 \right\}, \quad (32)$$

and an associated training space of normalized waveforms

$$\mathcal{F}_{\text{train}} = \{h_{\mathcal{M}_c}(f) \mid \mathcal{M}_c \in \mathcal{T}_K\}.$$

Through numerical experiments we have found that for this problem the number of basis for any given greedy error saturates with at most $K = 3,000$ training space elements, with—for example—178 RB elements $\{e_i\}_{i=1}^{178}$ needed to achieve a tolerance⁵ of $\epsilon^2 = 10^{-12}$. Therefore, the results shown below use 3,000 training space samples. Fig. 4(a) shows the distribution of points selected by the greedy algorithm. They cluster at low \mathcal{M}_c , which corresponds to more lower-frequency oscillations in Eq. (30). Such clustering is expected, since the number of cycles in a frequency range $f \in [f_{\min}, f_{\max}]$ is given to lowest post-Newtonian approximation order by (graphically seen from Fig 3)

$$\mathcal{N}_{\text{cycles}} = \mathcal{N}(f_{\min}) - \mathcal{N}(f_{\max}),$$

where (see Eq. (4.23), and Eq. (5.247) from Ref. [44])

$$\mathcal{N}(f) = 1/(32\pi^{8/3}) \left(G\mathcal{M}_c/c^3 \right)^{-5/3} f^{-5/3}.$$

The truly interesting aspect is that the distribution of points is nearly optimal in a precise sense, as described in Section 3.1, and is automatically given by the RB-greedy algorithm.

The solid black lines labeled “projection error” in both plots of Figure 5 show the square of the greedy error $\sigma_n(\mathcal{F}_{\text{train}}; \mathcal{H})$, defined in Eq. (12), over the training space $\mathcal{F}_{\text{train}}$, as a function of the number of RB elements. After a slow decay, the error decays with a very fast exponential falloff, a feature that we have found in this family of inspiral gravitational waveforms and generalizations thereof [18, 19].

⁵ Our numerical experiments are carried out with double precision arithmetic. This translates into a double precision computation of the quantity $\|h_\mu - \mathcal{P}_n h_\mu\|_d^2$ found in step 3b of Algorithm 1 (RB-Greedy Algorithm), and hence an accuracy of about 10^{-7} in the computation of its square root. These observations motivate a choice for the greedy error tolerance to be $\epsilon = 10^{-6}$. Refs. [53, 54] address this issue in greater detail and propose alternatives for improving error computations.




Fig. 4 The left figure (a) shows the parameters \mathcal{M}_c selected by greedy Algorithm 1. The right figure (b) shows the frequency points selected by DEIM Algorithm 2 out of 20,000 equidistant frequency sample points. We see from the histograms that both the chirp masses selected and the frequency interpolating points cluster at small values. This is intuitively expected, because smaller chirp masses correspond to a greater number of waveform cycles in Eq. (30). The point here is that the RB-greedy and DEIM algorithms automatically select both the parameter and physical relevant points in a nearly optimal way, without “human intervention” or “selection by eye”.

Discrete Empirical Interpolation: Since the reduced basis set $\{e_i\}_{i=1}^{178}$ generated above is application-based (as opposed to, say, Jacobi polynomials), an appropriate set of interpolation points is *in principle* not known, and that is where the DEIM algorithm of Section 3.2 enters the game. We now generate a hierarchical set of DEIM points for the reduced basis built (these results are not used in the construction of ROQ for inner products), choosing the maximum number of interpolating points which equals the dimension of the RB space. That is, using $n = 1, 2, 3, \dots, 178$ reduced basis elements, we sequentially computed the equivalent number of DEIM points. We recall that the RB-greedy approach selects points in parameter space, in this case the chirp mass, while the DEIM generates interpolating points in physical space (here frequency), and that both methods are hierarchical. That is, a seed choice for the first parameter value defines the first RB element, and from it the first interpolating point can be computed. Next, the greedy method chooses a second parameter value and the second RB element. By applying the DEIM to this new set (of so far two basis elements), the second interpolating point can be computed. And so on. Equivalently (the output is exactly the same), one can generate the whole RB first, up to the desired representation error tolerance, and then generate all of the interpolation points. Put differently, if there are n RB elements, up to n DEIM points can be generated, and generating some $n_1 < n$ DEIM only requires the first n_1 RB elements. This discussion, though perhaps trivial, might be helpful to keep in mind when later discussing the results from Fig. 5.

As input to the DEIM algorithm we must provide the RB functions sampled on some set of frequency points. Two cases are here considered: i) the basis vectors are sampled at 1,701 Gauss-Legendre points, which was the integration rule used to build the reduced basis in the Subsection 5.2.1, ii) the basis vectors are sampled at 20,000 equidistant points⁶. Figure 4(b) depicts the frequency distribution of selected points for the equidistant sampling case, which is of particular practical importance when one seeks to downsample experimental data. The overall structure for the Gauss-Legendre points case was found to be essentially identical, and is therefore not shown. This, we emphasize once again, is important because ROQ allow downsampling data without loss of accuracy, while producing a fast converging quadrature rule, regardless of the sampling.

Next, we randomly pick 10,000 waveforms, not necessarily in the training space, and represent each of them as a DEIM interpolant (that is, using Eq. (24)). Both sampling at Gauss-Legendre and equidistant points are considered. For each waveform the DEIM error is calculated in the discrete L^2 norm as $\|h_\mu - \mathcal{I}_n[h_\mu]\|_a^2$. Fig. 5 compares the largest DEIM error over all 10,000 waveforms (solid blue line) with the greedy error (solid black line). Notice that the black line lies strictly below the blue line; for waveforms in the training space this is guaranteed by the optimality of the L^2 orthogonal projection $\mathcal{P}_n h_\mu$. In turn, the DEIM interpolation error is bounded, again as expected, from above by the a-priori theoretical error estimates given in by Eq. (28) (red line) and Eq. (29) (purple line). In principle these two error estimates strictly hold only for waveforms in the training space, but our numerical results show that they evidently

⁶ See the discussion in Section 4.3 related to the subtleties involved when sampling the RB functions at a set different from those used to compute the underlying quadratures to build the RB itself.

continue to hold for waveforms between elements of the training space when the latter is sufficiently dense. The red line provides a sharper bound on the error but it is also more expensive to calculate (cf. Remark 6). Finally, notice that, at least for elements in the training space, the asymptotic convergence rate of the DEIM is predicted to be bounded by the RB-greedy representation error (with the Lebesgue constant as a proportionality constant), see Corollary 2. From the figure we can see that the similar asymptotic convergence rates for the DEIM and reduced basis representation continue to hold for waveforms not necessarily in the training space.




Fig. 5 Maximum projection and interpolation errors, as well as the predicted error bounds, for single waveforms as a function of the number of reduced basis (equal to the number of DEIM points) used. The left plot refers to ROQ built from frequency points sampled at Gauss-Legendre nodes, while the right one shows ROQ built from equidistant points, whence the error bounds given by Eq. (29) (error bound 1, magenta) and Eq. (28) (error bound 2, red) are identical up to numerical accuracy (cf. remark 6). In both figures “projection error” plots the square of the greedy error $\sigma_n(F_{\text{train}}; \mathcal{H})$ and “interpolation error” plots a maximum error $\|h_{\mathcal{M}_c} - \mathcal{I}_n[h_{\mathcal{M}_c}]\|_{\mathcal{A}}^2$ taken over 10,000 randomly drawn values of \mathcal{M}_c not necessarily in the training space. For both cases the maximum interpolation error is nearly identical, as expected.

5.2.2 Direct and two-step greedy: results and comparison

In the previous subsection we constructed an approximation space $F_n \approx \mathcal{F}$. It was found that 178 basis elements are needed to represent any member of a 3000 member training set given by (32) with an accuracy better than $\epsilon^2 = 10^{-12}$. We continue following Algorithm (4) with an aim towards approximating $\tilde{\mathcal{F}}$. Results for both the two-step and direct greedy algorithms are given, followed by a short discussion.

Two-step greedy results: As discussed in Sec. 3.3 a second RB-greedy is used to build a space \tilde{F}_m which approximates $\tilde{\mathcal{F}}$. Following the prescription outlined there, we begin by building a training set $\mathcal{T}_n^2 = \{(\mu_i, \mu_j)\}_{i,j=1}^n$ and a training space F_{n^2} of associated normalized products $\{h_{\mu_i}^* h_{\mu_j} / \|h_{\mu_i}^* h_{\mu_j}\|_{\mathcal{A}}\}_{i,j=1}^n$, where $\{\mu_i\}_{i=1}^n$ are the greedy points identified upon the first greedy approximation $\mathcal{F} \approx F_n$ carried out in Sec. 5.2.1. Clearly our definition of F_{n^2} is not the only choice and, in particular, one may build a training space from normalized products of the orthogonalized basis $\{e_i\}_{i=1}^{178}$ (cf. Remark 3).

Through numerical experiments we have found that 339 reduced basis elements are needed to achieve a tolerance of $\epsilon^2 = 10^{-12}$ for approximation of F_{n^2} . Fig. 6 (left) shows the distribution of points selected by the greedy algorithm (the two-step greedy results are denoted by black circles). As expected they cluster towards lower values of \mathcal{M}_c (see Sec. 5.2.1). The solid blue line labeled “projection error, Two-step” in Fig. 6 (right) plots the square of the greedy error $\sigma_m(F_{n^2}; \mathcal{H})$, defined in Eq. (12), over the training space of F_{n^2} , as a function of the number of reduced basis elements. Furthermore, through Monte Carlo sampling of the continuum we find any waveform to be accurately represented (see “interpolation error” of Fig. 7, right) by these same 339 basis elements.

Direct greedy results: One may consider direct approximation of $\tilde{\mathcal{F}}$ through a single RB-greedy (Path # 1 in Algorithm 4). For most problems this will be prohibitively expensive even as an offline computation. Nevertheless, we provide details on it here mainly for comparison with the two-step approach (Path # 2 in Algorithm 4). We take our training set to be a Cartesian product $\mathcal{T}_K^2 = \{(\mu_i, \mu_j)\}_{i,j=1}^K$ of the $K = 3000$ element training set given by (32) and a training space F_{K^2} to be the associated normalized products $h_{\mu_i}^* h_{\mu_j} / \|h_{\mu_i}^* h_{\mu_j}\|_{\mathcal{A}}$. To achieve an approximation tolerance




Fig. 6 Selected chirp mass \mathcal{M}_c values (left) and square of the greedy error (right) for both the direct and two-step greedy approaches. Each approach seeks to compress a training space which consists of normalized products of waveforms. However, these training spaces are constructed in very different ways. The two-step training space is built from a Cartesian product of greedy points identified while building an RB space for waveforms (as opposed to their products) and thus has 178^2 members. A direct greedy uses a training space from the Cartesian product of the \mathcal{T}_K defined in Eq. (32) and thus has 3000^2 members. Remarkably, despite these differences, we find that the exact number (339) of basis vectors are needed to achieve an error tolerance of $\epsilon^2 = 10^{-12}$ for both the two-step and direct greedy. The left figure shows the point distribution is slightly different with more clustering at lower masses for the two-step process (black circles) as compared with the direct one (red squares).

of $\epsilon^2 = 10^{-12}$ we have found that 339 RB elements were needed. Fig. 6 (left) shows the distribution of points selected by the greedy algorithm (the direct greedy results are denoted by red squares). The dashed black line labeled “projection error, Direct” in Figure 6 (right) shows the square of the greedy error $\sigma_m(\mathcal{F}_{K^2}; \mathcal{H})$, defined in Eq. (12), over the training space \mathcal{F}_{K^2} , as a function of the number of RB elements.

Discussion and comparison of two-step and direct greedy: We first notice that when building F_{n^2} from products of waveforms (as opposed to products of orthonormal basis vectors (cf. Remark 3)) one has $F_{n^2} \subset F_{K^2}$. In light of this observation one may view the two-step greedy as a smarter choice of sampling $\tilde{\mathcal{F}}$. Furthermore, to better assist with our comparison, in the numerical experiments carried out above we have initialized both the two-step and direct greedy algorithms with the same seed.

Remarkably, despite the differences in the two approaches, we find the greedy error curves are nearly identical and in both cases *exactly* 339 basis vectors are needed to achieve an error tolerance of $\epsilon^2 = 10^{-12}$ (the greedy errors are defined with respect to different training spaces, however). In both cases after a slowly decaying region we observe very fast exponential convergence of the form $Ce^{-c_0 n^\alpha}$. For the two-step $\sigma_m(F_{n^2}; \mathcal{H})$ can be fitted by

$$C = 4.19 \times 10^{-3}, c_0 = 0.981, \text{ and } \alpha = 0.923,$$

while for the direct $\sigma_m(\mathcal{F}_{K^2}; \mathcal{H})$ can be fitted by

$$\hat{C} = 3.98 \times 10^{-3}, \hat{c}_0 = 1.07 \text{ and } \hat{\alpha} = 0.875.$$

The number m of basis found to be needed so that \tilde{F}_m represents $\tilde{\mathcal{F}}$ within machine precision, $m = 339$, is remarkably close to twice that one needed so that F_n represents \mathcal{F} with the same accuracy, $n = 178$. If we were dealing with polynomials, it would be exactly $m = 2n$. From a practical perspective, with the two-step greedy approach we remove significant redundancy amongst elements of F_{n^2} where the dimensionality of the space to represent products of waveforms is compressed from n^2 to $\sim 2n$, with a compression ratio of about 90 for this problem.

Notice the savings in the offline stage when building the reduced basis space \tilde{F}_m using our two-step greedy approach, compared to a direct one. For the problem here considered, which is not particularly large in terms of the number of physical parameters (typically an 8 dimensional parameter space for a faithful description of compact binary coalescences [19]), we needed 3,000 training space points to build F_n . Using a direct approach to build \tilde{F}_m one needs 9×10^6 training space points, compared to the 178^2 needed in the two-step greedy, *with offline savings when building \tilde{F}_m of ~ 284* . At the end of the day, both basis are able to accurately represent any waveform product and hence perfectly well suited for an ROQ construction. Clearly the two-step is preferable when considering these costs; we refer to Remark 7.

Remark 8 The savings for larger problems, for example when using a lower cutoff frequency of 10Hz, as estimated for the upcoming generation of earth-based detectors, or including spins in the modeling of each compact objects would be considerably larger (even in the absence of precession). For such cases we have typically used $\sim 10^6$ elements in the training space in order to build a high accuracy RB, with less than $\sim 2,000$ RB elements needed to represent F_n [17, 18, 19]. For those cases a direct RB-greedy construction of \tilde{F}_m would in principle require $\sim 10^{12}$ training space elements – our two-step greedy would then save a cost of $\sim 10^5$.

5.2.3 DEIM and ROQ for overlap/inner-product integration

In the previous subsection we constructed the approximation space $\tilde{F}_m \approx \tilde{\mathcal{F}}$ using a direct and two-step greedy algorithms. While both approaches identify an accurate and compact reduce basis set $\{\tilde{e}_i\}_{i=1}^m$ approximating \mathcal{F} , the two-step has a much smaller computational cost (see Remark 7). We now continue with Algorithm (4). The results shown are for a reduced basis generated from the two-step greedy.

Discrete Empirical Interpolation: At this point we have a reduced basis set $\{\tilde{e}_i\}_{i=1}^{339}$ approximating $\tilde{\mathcal{F}}$, which for this problem can be viewed as continuous functions of frequency f (they are known in closed form for any values of f). We now generate the corresponding set of interpolation points using the DEIM algorithm of Section 3.2. Results for Gauss-Legendre sampling (case 1) are shown in Figure 7, equidistant sampling is qualitatively similar. A distribution of selected DEIM points is depicted in the leftmost plot. Notice that the overall structure is very similar to the single waveform case shown in Figure 4(b).

For 20,000 randomly sampled normalized waveform products $g \in \tilde{\mathcal{F}}$, not necessarily in the training space, the DEIM interpolant is evaluated using Eq. (24). The DEIM interpolation error is calculated in the L^2 norm as $\|g - \tilde{\mathcal{I}}_m[g]\|_d^2$. Fig. 7 (right) compares the largest DEIM interpolation error over all 20,000 products (solid blue line) with the square of the greedy error $\sigma_m(F_{n2}; \mathcal{H})$ (solid black line). Notice that the DEIM interpolation error is bounded from above by the a-priori error estimates Eq. (28) (red line) and Eq. (29) (purple line). These two error estimates strictly hold for waveforms in the training space, and evidently continue to hold for waveforms outside of the training space when the latter is sufficiently dense. The red line provides a sharper bound on the error, but its also more expensive to compute (cf. Remark 6).




Fig. 7 The left figure shows the distribution of empirical interpolation points selected by the DEIM algorithm with the basis vectors sampled at 1,701 Gauss-Legendre points. The right figure shows maximum projection and interpolation error (and error bounds) for products of waveforms as a function of the number of reduced basis used; notice that the profile is similar to that one of Fig. 5. On the right figure “projection error” plots the square of the greedy error $\sigma_m(F_{n2}; \mathcal{H})$ and “interpolation error” plots a maximum error $\|g - \tilde{\mathcal{I}}_m[g]\|_d^2$ taken over 20,000 randomly drawn normalized products $g \in \tilde{\mathcal{F}}$ not necessarily in the training space. A- priori interpolation error bounds given by Eq. (28) (error bound 2) and Eq. (29) (error bound 1) are also shown. DEIM results for equidistant sampling are qualitatively similar to the Gauss-Legendre case.

Reduced Order Quadratures: Having assembled a set of basis vectors and empirical interpolation points for products, we compute the reduced order quadrature weights ω_i^{ROQ} with Eq. (25) to complete the ROQ rule. Like the previous

experiments of this section, we randomly sample 20,000 inner products (with normalized integrands) to compute, and monitor the maximum error from this computation. Figure 8 compares accuracy versus computational degrees of freedom (number of quadrature nodes) for a variety of integration schemes. The black and blue lines denote trapezoidal and Gauss-Legendre quadratures, respectively. Two cases are considered within reduced order quadratures. The ROQ leading to the red line stems from a reduced basis space $\{\tilde{e}_i\}_{i=1}^{339}$ with Gauss-Legendre nodes and underlying quadrature rule. The magenta one, in turn, uses equally spaced samples and a trapezoidal underlying rule. From Figure 8, one can see that ROQ have a factor of ~ 2 of savings when compared to Gauss-Legendre points for a maximum error below $10^{-2} - 10^{-1}$, while the savings compared to the extended trapezoidal rule at high resolutions are greater than 50. This latter comparison is the one of interest for many data analysis applications where samples are often given on an equidistant grid spacing (downsampling).




Fig. 8 We randomly draw 20,000 normalized pair products $g \in \tilde{\mathcal{F}}$ and for each compute a very accurate inner product which we take to be its exact value I_c . Next, for each product we compute an inner product using a i) Gauss-Legendre quadrature (blue line), ii) trapezoidal rule (black line), iii) ROQ built from Gauss-Legendre quadratures (red line), and iv) ROQ built from the trapezoidal rule (pink line). In each of the four cases we monitor the maximum errors $|I_c - I_d|$ (for the Gauss-Legendre and trapezoidal rules) and $|I_c - I_{ROQ}|$ (for the ROQ accelerated Gauss-Legendre and trapezoidal rules). Once the underlying product is well resolved by the empirical interpolant we see very fast exponential convergence of $I_{ROQ} \rightarrow I_c$ (compare with “interpolation error” plotted in the right panel of Figure 7). Notice that both ROQ rules perform similarly, which is to be expected whenever the underlying integration scheme is able to accurately integrate the reduced basis (cf. the discussion following Algorithm 4 and Fig. 2). After an error of about 10^{-1} both ROQ rules outperform their discrete counterparts and in some cases significantly so.

6 Final remarks

Numerical integration is a well studied topic (see Refs. [55, 24] for excellent introductions). One may wish to identify when Reduced Order Quadratures (ROQ) are likely to be a competitive option over alternative, more standard integration methods. First and foremost, the set of functions to be integrated should be well approximated by a small set of basis vectors. Although we have here focused on bases obtained through the Reduced Basis-greedy approach, ROQ apply unchanged to any other basis sets; for example, those obtained through a Proper Orthogonal/Singular Value decomposition. Second, since the cost of building the reduced basis either by a RB-greedy or POD/SVD is potentially large, it should be done *offline*. For larger problems, however, a direct reduced order modeling approach might be unfeasible, even when building the basis is an offline and parallelizable calculation. To overcome this we proposed a two-step greedy targeted towards approximation of products of functions (such as those appearing in weighted inner products) which, for the problem considered here, accelerated offline ROQ computations by two orders of magnitude. Finally, and crucially, one should determine which *online* costs to reduce. Benefits of ROQ are greatest when multiple fast evaluations of weighted inner products (or perhaps simply weighted integrals) are required between parametrized functions. Are the functional evaluations themselves costly? Quadrature rules which are designed for high-order integration of polynomials might lead to significantly more functional evaluations compared to an ROQ rule. Can one sample the function at arbitrary

points or is the data sampling specified? In our numerical experiments with gravitational waves the savings were moderate (a factor of ~ 2) when we were able to dictate the sampling location whereas they were much greater (a factor of ~ 50) when the function was sampled at equally spaced points and integrated with an extended trapezoidal method. This suggests the possibility of using ROQ as an efficient downsampling criteria when doing matched filtering, for example (see [52]).

There are many potential applications of ROQ which deserve further consideration. As pointed out in Ref. [2], empirical interpolation is remarkably versatile in its handling of multi-dimensional problems on irregular domains, and a natural application of ROQ would be to such problems. Discontinuous functions (in physical space) with smooth variation with respect to parameters might admit a reduced basis approximating to the full space with fast convergence, and if so it might be possible to generate fast converging ROQ for a particular family of discontinuous functions. Furthermore, as pointed out in Sec. 4, the hierarchical construction of ROQ node locations allows for very natural application-specific nested quadrature rules of arbitrary order and depth.

Acknowledgments

This work has been supported in part by NSF Grants DMS-0807811, PHY0801213, PHY1005632, and DMS-1109325 to the University of Maryland, and PHY1125915 to the University of California at Santa Barbara. We thank Priscilla Canizares for comments on the manuscript and Chad Galley for useful discussions. MT and SEF thank the Kavli Institute for Theoretical Physics, University of California at Santa Barbara, where this work was completed, for its hospitality. H. A. and M. T. thank Tryst DC, where parts of this work were done, for its hospitality.

A Notation

List of symbols:

- \mathcal{F} : parameterized space of target functions.
- μ : parameter (in general, a multi-dimensional one).
- μ_i, μ_j : numerical values of the parameter μ .
- $\tilde{\mathcal{F}}$: parameterized space of product of target functions.
- $\tilde{\mu}_k$: in sets arising from products, such as \mathcal{T}_n^2 , $\tilde{\mu}_{k(i,j)} = (\mu_i, \mu_j)$ is a re-indexing of a two dimensional array $\{(\mu_i, \mu_j)\}_{i,j=1}^n$ into a one dimensional array $\{\tilde{\mu}_k\}_{k=1}^{n^2}$, where k is a one-to-one function which takes two integers i and j and returns a unique integer $k(i, j)$.
- K : number of training space elements sampling \mathcal{F} .
- K^2 : number of training space elements in principle needed to *directly* sample $\tilde{\mathcal{F}}$.
- n^2 : number of training space elements used in our two-step greedy approach to sample $\tilde{\mathcal{F}}$. In general $n^2 \ll K^2$.
- M : number of samples or quadrature points where functions to be integrated are evaluated.
- n, m : number of DEIM points (and reduced basis) to approximate elements in \mathcal{F} and $\tilde{\mathcal{F}}$ receptively.
- Discrete L^∞ norm of a g evaluated at M quadrature points is $\|g\|_\infty = \max_{1 \leq i \leq M} |g(x_k)|$.
- The spectral 2-norm of a matrix $\mathbf{Q} \in \mathbb{C}^{M \times n}$ is defined by

$$\|\mathbf{Q}\|_2 := \max_{\mathbf{u} \neq 0} \frac{\|\mathbf{Q}\mathbf{u}\|_d}{\|\mathbf{u}\|_d} = \left(\lambda_{\max}(\mathbf{Q}^\dagger \mathbf{Q}) \right)^{1/2},$$

with † denoting conjugate transposition. The matrix ∞ -norm is given by

$$\|\mathbf{Q}\|_\infty = \max_{1 \leq k \leq M} \sum_{\ell=1}^n |q_{k\ell}|,$$

respectively, where $q_{k\ell}$ denotes the (k, ℓ) entry of \mathbf{Q} .

B FLOP Count

In this appendix we provide the asymptotic FLOP count for a variety of algorithms discussed through the paper. These numbers are used for comparing the relative costs of different algorithms.

B.1 DEIM FLOP Count

In Algorithm 2 as $\mathbf{P}^T \mathbf{U}$ is a lower triangular matrix therefore (after m iterations) Step (4) costs $\mathcal{O}(m^3)$. Step (5) costs $\mathcal{O}(Mm)$ subtractions and $\mathcal{O}(Mm^2)$ matrix vector multiplications, where latter is the dominant one. The improvement in our implementation of Algorithm 2 as compared to [4] is that the our implementation scales as $\mathcal{O}(m^3)$ as compared to $\mathcal{O}(m^4)$ in [4]. Therefore the total DEIM cost is:

$$\mathcal{O}(Mm^2 + m^3).$$

B.2 ROQ FLOP Count

The dominant cost (assuming $M > m$) of computing reduced order quadrature weights $(\omega^{\text{ROQ}})^T = \omega^T \tilde{\mathbf{V}} (\tilde{\mathbf{P}}^T \tilde{\mathbf{V}})^{-1}$ in (25) is $\mathcal{O}(Mm^2)$. Here we used the fact that the operation $\tilde{\mathbf{P}}^T \tilde{\mathbf{V}}$ is equivalent to selecting m rows corresponding to DEIM points $\{\tilde{p}_\ell\}_{\ell=1}^m$.

Then the overall cost to compute ROQ using one-step (Algorithm 1 [Path #1] with modified *Gram-Schmidt* in greedy is:

$$\mathcal{O}(K^2 M \hat{m} + M \hat{m}^2)$$

and two-step (Algorithm 1 [Path #2] is:

$$\mathcal{O}(n^2 M m + K M n + M m^2).$$

Here m, \hat{m} denote the number of reduced basis used in approximation of $\tilde{\mathcal{F}}$ via two-step and direct approach respectively and n is the number of reduced basis used to approximate \mathcal{F} .

C DEIM and LU Decomposition with Partial Pivoting

The original DEIM described in [4] has as Step 7

$$\mathbf{U} = [\mathbf{U} \quad \mathbf{e}_i], \quad i = 2, \dots, n.$$

Since the columns of \mathbf{U} are linearly independent, $\mathbf{P}^T \mathbf{U}$ is always invertible, but in this form the matrix $\mathbf{P}^T \mathbf{U}$ can be dense. Next we show that with a slight modification to the original DEIM, we get Algorithm 2 and at every iteration $i = 2, \dots, n$ with $\mathbf{r}_i := \mathbf{r}$,

$$\mathbf{U} := [\mathbf{U} \quad \mathbf{r}_i],$$

gives the same result. One of the advantages of looking at DEIM in this format is that the matrix $\mathbf{P}^T \mathbf{U}$ is lower triangular. Next we prove that we get the same result using both formats.

Proposition 1 In [4], Step 7 of Algorithm 2 was presented as $\mathbf{U} = [\mathbf{U} \quad \mathbf{e}_i], i = 2, \dots, n$, we can replace it by

$$\mathbf{U} = [\mathbf{U} \quad \mathbf{r}_i], \quad i = 2, \dots, n,$$

where $\mathbf{r}_i := \mathbf{r}$, at every iteration, and \mathbf{r} is as given in Step 5 of Algorithm 2.

Proof Step 4 of Algorithm 2, at $i = n$, gives the coefficient matrix

$$\mathbf{P}^T \mathbf{U} = \begin{pmatrix} \mathbf{e}_1(p_1) & \mathbf{e}_2(p_1) & \cdots & \mathbf{e}_{n-1}(p_1) \\ \mathbf{e}_1(p_2) & \mathbf{e}_2(p_2) & \cdots & \mathbf{e}_{n-1}(p_2) \\ \mathbf{e}_1(p_3) & \mathbf{e}_2(p_3) & \cdots & \mathbf{e}_{n-1}(p_3) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{e}_1(p_{n-1}) & \mathbf{e}_2(p_{n-1}) & \cdots & \mathbf{e}_{n-1}(p_{n-1}) \end{pmatrix}.$$

Next we write the row reduced echelon form using forward Gaussian elimination for $(\mathbf{P}^T \mathbf{U})^T$, where the first step implies

$$\begin{pmatrix} \mathbf{e}_1(p_1) & 0 & \cdots & 0 \\ \mathbf{e}_1(p_2) & \mathbf{r}_2(p_2) & \cdots & \mathbf{e}_{n-1}(p_2) - \frac{\mathbf{e}_{n-1}(p_1)}{\mathbf{e}_1(p_1)} \mathbf{e}_1(p_2) \\ \mathbf{e}_1(p_3) & \mathbf{r}_2(p_3) & \cdots & \mathbf{e}_{n-1}(p_3) - \frac{\mathbf{e}_{n-1}(p_1)}{\mathbf{e}_1(p_1)} \mathbf{e}_1(p_3) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{e}_1(p_{n-1}) & \mathbf{r}_2(p_{n-1}) & \cdots & \mathbf{e}_{n-1}(p_{n-1}) - \frac{\mathbf{e}_{n-1}(p_1)}{\mathbf{e}_1(p_1)} \mathbf{e}_1(p_{n-1}) \end{pmatrix}.$$

The final row reduced echelon form for $(\mathbf{P}^T \mathbf{U})^T$ is

$$\begin{pmatrix} \mathbf{e}_1(p_1) & 0 & \cdots & 0 \\ \mathbf{e}_1(p_2) & \mathbf{r}_2(p_2) & \cdots & 0 \\ \mathbf{e}_1(p_3) & \mathbf{r}_2(p_3) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{e}_1(p_{n-1}) & \mathbf{r}_2(p_{n-1}) & \cdots & \mathbf{r}_n(p_{n-1}) \end{pmatrix}. \quad (33)$$

Hence the proposition.

We note that advantage of using \mathbf{r} in Step 7 of Algorithm 2, is that then in Step 4, $(\mathbf{P}^T \mathbf{U})$ is a lower triangular matrix, therefore the system can be solved for \mathbf{c} with $\mathcal{O}(n^2)$ operations, as compared to using \mathbf{e}_i , which gives a dense matrix $(\mathbf{P}^T \mathbf{U})$, solving for \mathbf{c} requires $\mathcal{O}(n^3)$ operations.

Remark 9 Algorithm 2 has a flavor of Gauss elimination with row pivoting. We showed the elimination part in Proposition 1. Next we will show that the DEIM points $\{p_1, \dots, p_n\}$ are in fact the pivots, similar to Gauss elimination. Our goal is to arrive at (33) row echelon form of $(\mathbf{P}^T \mathbf{U})^T$ with with partial pivoting. Consider the matrix

$$\begin{pmatrix} \mathbf{e}_1(p_1) & \mathbf{e}_2(p_1) & \cdots & \mathbf{e}_{n-1}(p_1) \\ \mathbf{e}_1(p_2) & \mathbf{e}_2(p_2) & \cdots & \mathbf{e}_{n-1}(p_2) \\ \mathbf{e}_1(p_3) & \mathbf{e}_2(p_3) & \cdots & \mathbf{e}_{n-1}(p_3) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{e}_1(p_{n-1}) & \mathbf{e}_2(p_{n-1}) & \cdots & \mathbf{e}_{n-1}(p_{n-1}) \end{pmatrix}$$

Let the location of the first pivot, i.e., the maximum of the absolute value of the first column be p_1 (this is the same as Step 1 in Algorithm 2). Apply the first step of forward elimination as in the proof of Proposition 1. Next compute the second pivot p_2 , which is the maximum of the absolute value of the so-obtained second column. Following in this manner we arrive at the matrix in (33).

References

1. A. Patera, G. Rozza, Reduced Basis Approximation and A Posteriori Error Estimation for Parametrized Partial Differential Equations, to appear in (tentative rubric) MIT Pappalardo Graduate Monographs in Mechanical Engineering, http://augustine.mit.edu/methodology/methodology_bookPartI.htm.
2. Y. Maday, N. C. Nguyen, A. T. Patera, S. H. Pau, A general multipurpose interpolation procedure: the magic points, *Communications on Pure and Applied Analysis* 8 (2009) 383–404.
3. M. Barrault, Y. Maday, N. C. Nguyen, A. T. Patera, An empirical interpolation method: application to efficient reduced-basis discretization of partial differential equations, *Comptes Rendus Mathématique* 339 (2004) 667–672. doi:10.1016/j.crma.2004.08.006.
4. S. Chaturantabut, D. C. Sorensen, Nonlinear model reduction via discrete empirical interpolation, *SIAM Journal on Scientific Computing* 32 (5) (2010) 2737–2764. doi:10.1137/090766498. URL <http://link.aip.org/link/?SCE/32/2737/1>
5. S. Chaturantabut, D. Sorensen, Discrete empirical interpolation for nonlinear model reduction, in: *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, 2009, pp. 4316–4321. doi:10.1109/CDC.2009.5400045.
6. M. Hinze, S. Volkwein, Proper orthogonal decomposition surrogate models for nonlinear dynamical systems: error estimates and sub-optimal control, in: *Dimension reduction of large-scale systems*, Vol. 45 of *Lect. Notes Comput. Sci. Eng.*, Springer, Berlin, 2005, pp. 261–306.
7. D. Manolakis, V. Ingle, S. Kogon, *Statistical and adaptive signal processing: spectral estimation, signal modeling, adaptive filtering, and array processing*, Artech House signal processing library, Artech House, 2000. URL <http://books.google.com/books?id=3RQfAQAATAAJ>
8. L. A. Wainstein, V. D. Zubakov, *Extraction of signals from noise*, Prentice-Hall, London, 1962.
9. J. L. Eftang, B. Stamm, Parameter multi-domain hp empirical interpolation, submitted.
10. T. O. Aanonsen, Empirical interpolation with application to reduced basis approximations, Ph.D. thesis, Norwegian University of Science and Technology (2009). URL <http://ntnu.diva-portal.org/smash/record.jsf?pid=diva2:348771>
11. A. Pinkus, *N-widths in approximation theory*, Springer, Amsterdam, 1985.
12. P. Binev, A. Cohen, W. Dahmen, R. A. DeVore, G. Petrova, P. Wojtaszczyk, Convergence rates for greedy algorithms in reduced basis methods., *SIAM J. Math. Analysis* 43 (3) (2011) 1457–1472. URL <http://dblp.uni-trier.de/db/journals/siamma/siamma43.html#BinevCDDPW11>
13. R. DeVore, G. Petrova, P. Wojtaszczyk, Greedy algorithms for reduced bases in banach spaces, *Arxiv preprint arXiv:1204.2290*.
14. J. L. Eftang, A. T. Patera, E. M. Ronquist, An hp certified reduced basis method for parametrized elliptic partial differential equations, *SIAM J. Sci. Comput.* 32 (2010) 3170–3200.
15. T. Bui-Thanh, Model-constrained optimization methods for reduction of parameterized large-scale systems, PhD thesis, Massachusetts Institute of Technology.
16. B. Haasdonk, M. Dohlmann, M. Oehlberger, A training set and multiple bases generation approach for parametrized model reduction based on adaptive grids in parameter space, *Tech. Report*, SRC SimTech.
17. S. Caudill, S. E. Field, C. R. Galley, F. Herrmann, M. Tiglio, Reduced Basis representations of multi-mode black hole ringdown gravitational waves, *Class.Quant.Grav.* 29 (2012) 095016. arXiv:1109.5642.
18. S. E. Field, C. R. Galley, F. Herrmann, J. S. Hesthaven, E. Ochsner, M. Tiglio, Reduced basis catalogs for gravitational wave templates, *Phys. Rev.Lett.* 106 (2011) 221102. arXiv:1101.3765, doi:10.1103/PhysRevLett.106.221102.
19. F. Herrmann, S. E. Field, C. R. Galley, E. Ochsner, M. Tiglio, Towards beating the curse of dimensionality for gravitational waves using Reduced BasisarXiv:1205.6009.
20. K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, et al., Singular value decomposition applied to compact binary coalescence gravitational-wave signals, *Phys. Rev. D* 82 (2010) 044025. arXiv:1005.0012, doi:10.1103/PhysRevD.82.044025.
21. K. Cannon, C. Hanna, D. Keppel, Efficiently enclosing the compact binary parameter space by singular-value decompositionarXiv:1101.4939.
22. J. Xu, L. Zikatanov, Some observations on babuška and brezzi theories, *Numer. Math.* 94 (1) (2003) 195–202. doi:10.1007/s002110100308. URL <http://dx.doi.org/10.1007/s002110100308>
23. D. B. Szyld, The many proofs of an identity on the norm of oblique projections, *Numer. Algorithms* 42 (3-4) (2006) 309–323. doi:10.1007/s11075-006-9046-2. URL <http://dx.doi.org/10.1007/s11075-006-9046-2>

24. F. S. Alfio Quarteroni, Riccardo Sacco, Numerical Mathematics, Springer, Berlin, 2010.
25. B. C. Barish, R. Weiss, LIGO and the detection of gravitational waves, *Phys. Today* 52N10 (1999) 44–50.
26. S. J. Waldmann, Status of LIGO at the start of the fifth science run, *Class. Quantum Grav.* 23 (2006) S653–S660.
27. S. Hild, The status of GEO 600, *Class. Quantum Grav.* 23 (2006) S643–S651.
28. F. Acernese, et al., The Virgo status, *Class. Quantum Grav.* 23 (2006) S635–S642.
29. B. Abbott, et al., LIGO: The Laser Interferometer Gravitational-Wave Observatory [arXiv:0711.3041](https://arxiv.org/abs/0711.3041).
30. J. Abadie, et al., Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors, *Class. Quantum Grav.* 27 (2010) 173001. [arXiv:1003.2480](https://arxiv.org/abs/1003.2480), doi:10.1088/0264-9381/27/17/173001.
31. J. Centrella, J. G. Baker, B. J. Kelly, J. R. van Meter, Black-hole binaries, gravitational waves, and numerical relativity, *Rev. Mod. Phys.* 82 (2010) 3069. [arXiv:1010.5260](https://arxiv.org/abs/1010.5260), doi:10.1103/RevModPhys.82.3069.
32. F. Pretorius, Binary Black Hole Coalescence [arXiv:0710.1338](https://arxiv.org/abs/0710.1338).
33. T. W. Baumgarte, S. L. Shapiro, Numerical relativity and compact binaries, *Physics Reports* 376 (2) (2003) 41–131. [arXiv:gr-qc/0211028](https://arxiv.org/abs/gr-qc/0211028).
34. O. Sarbach, M. Tiglio, Continuum and Discrete Initial-Boundary-Value Problems and Einstein's Field Equations To appear in *Living Reviews in Relativity*, 191 pages. [arXiv:1203.6443](https://arxiv.org/abs/1203.6443).
35. L. Blanchet, Gravitational radiation from post-newtonian sources and inspiralling compact binaries, *Living Reviews in Relativity* 9 (4). URL <http://www.livingreviews.org/lrr-2006-4>
36. L. Blanchet, T. Damour, B. R. Iyer, Gravitational waves from inspiralling compact binaries: Energy loss and wave form to second post-Newtonian order, *Phys. Rev. D* 51 (1995) 5360. [arXiv:gr-qc/9501029](https://arxiv.org/abs/gr-qc/9501029), doi:10.1103/PhysRevD.51.5360.
37. C. M. Will, A. G. Wiseman, Gravitational radiation from compact binary systems: gravitational waveforms and energy loss to second post-Newtonian order, *Phys. Rev. D* 54 (1996) 4813–4848. [arXiv:gr-qc/9608012](https://arxiv.org/abs/gr-qc/9608012).
38. B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, J. D. Creighton, FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries [arXiv:gr-qc/0509116](https://arxiv.org/abs/gr-qc/0509116).
39. L.-C. Tu, Q. Li, Q.-L. Wang, C.-G. Shao, S.-Q. Yang, et al., New determination of the gravitational constant G with time-of-swing method, *Phys. Rev. D* 82 (2010) 022001. doi:10.1103/PhysRevD.82.022001.
40. B. J. Owen, Search templates for gravitational waves from inspiraling binaries: Choice of template spacing, *Phys. Rev. D* 53 (1996) 6749.
41. B. J. Owen, B. S. Sathyaprakash, Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement, *Phys. Rev. D* 60 (1999) 022002. [arXiv:gr-qc/9903108](https://arxiv.org/abs/gr-qc/9903108).
42. S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, B. S. Sathyaprakash, A template bank to search for gravitational waves from inspiralling compact binaries: I. Physical models, *Classical and Quantum Gravity* 23 (2006) 5477–5504. [arXiv:gr-qc/0604037](https://arxiv.org/abs/gr-qc/0604037), doi:10.1088/0264-9381/23/18/002.
43. L. S. Finn, Detection, measurement and gravitational radiation, *Phys. Rev. D* 46 (1992) 5236–5249. [arXiv:gr-qc/9209010](https://arxiv.org/abs/gr-qc/9209010), doi:10.1103/PhysRevD.46.5236.
44. M. Maggiore, Gravitational Waves Volume 1: Theory and Experiments, Oxford University Press, Oxford, 2008.
45. B. Abbott, et al., LIGO: The Laser Interferometer Gravitational-Wave Observatory, *Rept. Prog. Phys.* 72 (2009) 076901. [arXiv:0711.3041](https://arxiv.org/abs/0711.3041), doi:10.1088/0034-4885/72/7/076901.
46. P. Ajith, S. Bose, Estimating the parameters of non-spinning binary black holes using ground-based gravitational-wave detectors: Statistical errors, *Phys. Rev. D* 79 (2009) 084032. [arXiv:0901.4936](https://arxiv.org/abs/0901.4936), doi:10.1103/PhysRevD.79.084032.
47. B. Abbott, et al., Einstein@home search for periodic gravitational waves in ligo s4 data, *Phys. Rev. D* 79 (2009) 022001. doi:10.1103/PhysRevD.79.022001. URL <http://link.aps.org/doi/10.1103/PhysRevD.79.022001>
48. P. R. Brady, T. Creighton, Searching for periodic sources with LIGO. 2. Hierarchical searches, *Phys. Rev. D* 61 (2000) 082001. [arXiv:gr-qc/9812014](https://arxiv.org/abs/gr-qc/9812014), doi:10.1103/PhysRevD.61.082001.
49. M. Vallisneri, Beyond Fisher: exact sampling distributions of the maximum-likelihood estimator in gravitational-wave parameter estimation, *Phys. Rev. Lett.* 107 (2011) 191104. [arXiv:1108.1158](https://arxiv.org/abs/1108.1158), doi:10.1103/PhysRevLett.107.191104.
50. J. Kanner, et al., LOOC UP: Locating and observing optical counterparts to gravitational wave bursts, *Class. Quantum Grav.* 25 (2008) 184034. [arXiv:0803.0312](https://arxiv.org/abs/0803.0312), doi:10.1088/0264-9381/25/18/184034.
51. P. S. Shawhan, LIGO Scientific Collaboration, Virgo Collaboration, LOOC-UP: Seeking Optical Counterparts to Gravitational-Wave Signal Candidates, in: *Bulletin of the American Astronomical Society*, Vol. 42 of Bulletin of the American Astronomical Society, 2010, pp. 230–+.
52. K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, N. Fotopoulos, et al., Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence, *Astrophys. J.* 748 (2012) 136. [arXiv:1107.2665](https://arxiv.org/abs/1107.2665).
53. F. Casenave, Accurate a posteriori error evaluation in the reduced basis method, *Comptes Rendus Mathématique* 350 (9) (2012) 539–542. doi:10.1016/j.crma.2012.05.012. URL <http://www.sciencedirect.com/science/article/pii/S1631073X12001483>
54. C. Canuto, T. Tonn, K. Urban, A posteriori error analysis of the reduced basis method for nonaffine parametrized nonlinear pdes, *SIAM J. Numer. Anal.* 47 (3) (2009) 2001–2022. doi:10.1137/080724812. URL <http://dx.doi.org/10.1137/080724812>
55. W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, Numerical Recipes, 2nd Edition, Cambridge University Press, New York, 1992.