Frequency-domain non-intrusive greedy Model Order Reduction based on minimal rational approximation

Davide Pradovera and Fabio Nobile

Abstract We present a technique for Model Order Reduction (MOR) of frequency-domain problems relying on rational interpolation of vector-valued functions. The selection of the sample points is carried out adaptively according to a greedy procedure. We describe several options for the choice of *a posteriori* error indicators, which are used to drive the greedy algorithm and define its termination condition. Namely, we illustrate a tradeoff between each indicator's accuracy and its "intrusiveness", i.e. how much information on the underlying high-fidelity model needs to be available. We test numerically the effectiveness of this technique in solving a non-Hermitian eigenproblem and a microwave frequency response analysis.

1 Introduction

Consider the function $\mathbf{u} : \mathbb{C} \ni \mu \mapsto \mathbf{u}(\mu) \in \mathbb{C}^n$ implicitly defined as the solution of the linear parametric problem with a single parameter

$$\mathrm{A}(\mu)\mathbf{u}(\mu) = \mathbf{f}(\mu), \tag{1}$$

with A and $\mathbf{f}$ smooth functions taking values in $\mathbb{C}^{n\times n}$ and $\mathbb{C}^n$, respectively. To be more specific, we consider here parametric problems (1) arising from spatial discretization (e.g., by FEM [3]) of frequency domain problems, with the parameter μ representing the frequency. For most such problems, $\mathrm{A}(\mu)$ depends at most quadratically on μ:

$$\mathrm{A}(\mu) = \mathrm{A}_0 + \mu\mathrm{A}_1 + \mu^2\mathrm{A}_2,$$

whereas $\mathbf{f}$ is usually of the form $\mathbf{f}(\mu) = \theta_0(\mu)\mathbf{f}_0$, with $\theta_0 : \mathbb{C} \to \mathbb{C}$.

Davide Pradovera, Fabio Nobile
CSQI, Institute of Mathematics, EPFL, 1015 Lausanne, Switzerland, e-mail: davide.pradovera@epfl.ch, fabio.nobile@epfl.ch

In applications, it is often computationally unfeasible to solve (1) as many times as needed for a frequency response analysis. In recent years, this issue has been solved through MOR, whose main purpose is the construction of a surrogate $\widetilde{\mathbf{u}}(\mu)$ for $\mathbf{u}(\mu)$, much cheaper to evaluate at any given μ than solving (1), and with good approximation properties.

2 Available MOR strategies

A plethora of MOR techniques have been employed to compute surrogates for frequency response problems; some notable ones include:

- projective techniques, e.g. the Reduced Basis (RB) and multi-moment-matching methods [2], are extremely powerful, but require knowledge of, and access to, the specific structure of A and $\mathbf{f}$;
- strategies based on rational approximation, e.g. the Löwner framework [4] or the Vector Fitting (VF) algorithm, are *non-intrusive*, i.e. they rely only on evaluations of $\mathbf{u}$ at few frequencies, which we will refer to as *snapshots* or *samples*; in particular, there is no need for any information on (nor access to) the specific structure of A and $\mathbf{f}$ in (1); the price to pay for this additional flexibility is a reduced accuracy of the method for a given number of snapshots.

More recently, the minimal rational interpolation (MRI) technique was proposed [7], trying to achieve non-intrusiveness and optimal snapshot management at the same time. We summarize here a practical scheme for MRI:

1. fix a set of sample points $\mu_1,\dots,\mu_S \in \mathbb{C}$, and a polynomial basis $\{\psi_i\}_{i=0}^{S-1} \subset \mathbb{P}^{S-1}(\mathbb{C})$ (e.g., one could choose monomials, or Chebyshev polynomials); also, let $\{\ell_j\}_{j=1}^{S} \subset \mathbb{P}^{S-1}(\mathbb{C})$ be the Lagrangian basis associated to the sample points;
2. build the Vandermonde matrix $\mathrm{V} \in \mathbb{C}^{S\times S}$ and the diagonal weight matrix D:
$$(\mathrm{V})_{ij} = \psi_j(\mu_i) \quad \text{and} \quad \mathrm{D} = \mathrm{diag}\left(\left[\frac{\mathrm{d}^{S-1}\ell_1}{\mathrm{d}\mu^{S-1}},\dots,\frac{\mathrm{d}^{S-1}\ell_S}{\mathrm{d}\mu^{S-1}}\right]\right) \in \mathbb{C}^{S\times S};$$
3. compute the snapshots $\mathbf{u}(\mu_1),\dots,\mathbf{u}(\mu_S)$ and assemble a QR decomposition of the snapshot matrix
$$\left[\mathbf{u}(\mu_1)\middle|\mathbf{u}(\mu_2)\middle|\cdots\middle|\mathbf{u}(\mu_S)\right] = \mathrm{WR}, \quad \text{with } \mathrm{W} \in \mathbb{C}^{n\times S}, \mathrm{R} \in \mathbb{C}^{S\times S}; \tag{2}$$
4. compute a minimal eigenvector $\mathbf{q} \in \mathbb{C}^S$ of the positive semidefinite (Gramian) matrix $(\mathrm{RDV})^H \mathrm{RDV}$, and define the surrogate denominator as
$$Q \in \mathbb{P}^{S-1}(\mathbb{C}), \quad Q(\mu) = \sum_{i=0}^{S-1} (\mathbf{q})_i \psi_i(\mu);$$
5. define the *reduced* minimal rational approximation $\mathring{\mathbf{u}}$ as

$$\mathbb{C} \ni \mu \;\mapsto\; \mathring{\mathbf{u}}(\mu) = \frac{\mathrm{R\ diag}\left([Q(\mu_1),\ldots,Q(\mu_S)]\right)}{Q(\mu)} \sum_{i=0}^{S-1} \left(V^{-\top}\right)_{:i} \psi_i(\mu) \in \mathbb{C}^S,$$

where $(\mathrm{A})_{:i}$ denotes the i-th column of matrix A; then the *full* minimal rational approximation $\widetilde{\mathbf{u}} \approx \mathbf{u}$ can be found as $\widetilde{\mathbf{u}}(\mu) = \mathrm{W}\mathring{\mathbf{u}}(\mu)$.

2.1 Greedy approach

A common feature of all the techniques cited above is that a "sufficiently large" number of samples is needed to guarantee the accuracy of the surrogate model; in the particular case of frequency-domain problems, there exist lower bounds [3] for the number of samples required to achieve reasonable accuracy. Unfortunately, such number depends on the unknown spectral properties of A, and on the approximability of $\mathbf{f}$. For RB and MRI, one can identify adaptively the correct number of samples by relying on the so-called *greedy* algorithm, which can be summarized as follows:

1. Initialize a set $V = \{\mathbf{u}_1,\ldots,\mathbf{u}_{S_0}\}$ with some preliminary snapshots at $\mu_1,\ldots,\mu_{S_0}$.
2. Build a surrogate model (e.g., by MRI) based on V.
3. Choose a measure $r(\mu)$ of the discrepancy between exact and surrogate solution, and find its maximal point $\widehat{\mu}$: $r = r(\mu) \leq r(\widehat{\mu})$ for all μ.
4. If $r(\widehat{\mu})$ is smaller than a prescribed tolerance, terminate.
5. Compute a snapshot at $\widehat{\mu}$, add it to V, and go to 2.

The main difficulty in setting up the greedy algorithm is choosing a good r. Given the presence of resonances, it is standard [3] to use as *a posteriori* estimator the residual of (1), namely, given some suitable norm $\|\cdot\|_\star$,

$$r(\mu) = \|\mathrm{A}(\mu)\widetilde{\mathbf{u}}(\mu) - \mathbf{f}(\mu)\|_\star. \tag{3}$$

2.2 A posteriori indicators

In an intrusive framework, an efficient way to compute (3) has been known in the RB literature for quite a while, see e.g. [2], assuming $\mathbf{f}(\mu)$ to depend affinely on μ, i.e.

$$\mathbf{f}(\mu) = \sum_{i=0}^{N_\mathbf{f}-1} \theta_i(\mu)\mathbf{f}_i,$$

with $\mathbf{f}_i \in \mathbb{C}^n$ and $\theta_i : \mathbb{C} \to \mathbb{C}$ for all i. Then, as long as the matrices A_i, the vectors $\mathbf{f}_i$, the weights $\theta_i(\mu)$, and the reduced surrogate solution $\mathring{\mathbf{u}}(\mu)$ are available, we can evaluate the residual at μ as

$$r(\mu)^2 = \sum_{i,j=0}^{N_{\mathbf{f}}-1} \overline{\theta_i(\mu)}\theta_j(\mu)\langle \mathbf{f}_j, \mathbf{f}_i\rangle_\star + \mathring{\mathbf{u}}(\mu)^* \left(\sum_{i,j=0}^{2} \overline{\mu}^i \mu^j \langle \mathrm{A}_j\mathrm{W}, \mathrm{A}_i\mathrm{W}\rangle_\star \right) \mathring{\mathbf{u}}(\mu) - 2\mathrm{Re}\left(\left(\sum_{i=0}^{N_{\mathbf{f}}-1} \sum_{j=0}^{2} \overline{\theta_i(\mu)}\mu^j \langle \mathrm{A}_j\mathrm{W}, \mathbf{f}_i\rangle_\star \right) \mathring{\mathbf{u}}(\mu) \right) \quad \text{(I)}$$

in $O((S+N_{\mathbf{f}})^2)$ operations. This idea can be employed in MRI as well, at the cost of making the procedure intrusive. However, we propose here some alternatives.

In [7] it was observed that, if $\widetilde{\mathbf{u}}$ is the MRI of $\mathbf{u}$ with samples at $\{\mu_j\}_{j=1}^S$ and denominator Q, and both $\mathrm{A}(\mu)$ and $\mathbf{f}(\mu)$ depend at most linearly on μ (i.e. $\mathrm{A}_2 = 0$ and $\mathbf{f}(\mu) = \mathbf{f}_0 + \mu\mathbf{f}_1$) or μ^2 (i.e. $\mathrm{A}_1 = 0$ and $\mathbf{f}(\mu) = \mathbf{f}_0 + \mu^2\mathbf{f}_2$), then

$$r(\mu) = \frac{c}{|Q(\mu)|} \prod_{j=1}^{S} |\mu - \mu_j|, \tag{4}$$

with $c = c(\mu_1, \dots, \mu_S, \mathrm{A}, \mathbf{f})$ independent of μ. In particular, since the location of the maximum of r does not depend on c, see (4), $\widehat{\mu}$ can be found even without knowing c. In order to determine the value of c non-intrusively, it is enough to compute r using (3) at a single new point μ' (in practice, we take $\mu' = \widehat{\mu}$):

$$r(\mu) = r(\mu') \left| \frac{Q(\mu')}{Q(\mu)} \right| \prod_{j=1}^{S} \left| \frac{\mu - \mu_j}{\mu' - \mu_j} \right|. \tag{R}$$

In certain situations, however, a direct evaluation of the quantity $r(\mu')$ within each greedy iteration might be impossible (e.g. if the solver used to evaluate $\mathbf{u}$ is a black-box that does not allow residual evaluation) or too computationally expensive. In order to have a viable greedy loop, we need to design an alternative termination condition in step 4. In this case, we propose to employ some heuristic indicator based on snapshot collinearity [3]: more explicitly, let W be the Q-factor in the QR factorization of the current snapshot matrix (2), and assume that the sample at $\widehat{\mu}$ has been computed; we opt to terminate the greedy algorithm if

$$\|\mathbf{u}(\widehat{\mu}) - \mathrm{W}\mathrm{W}^*\mathbf{u}(\widehat{\mu})\| < \mathrm{tol}\, \|\mathbf{u}(\widehat{\mu})\|. \tag{C}$$

For this last indicator, as long as the greedy iterations continue, the extra snapshot does not go wasted, since it is precisely the one which gets added to V in step 5: on the whole, this procedure computes only one "extra" snapshot, at the final greedy iteration, with respect to the two previous versions of the algorithm. Actually, one can adjust the greedy algorithm so as to employ even the extra snapshot in the final surrogate model: it suffices to build an updated MRI using all the samples, including the last one, once the termination condition has been satisfied.

We remark that the last two strategies rely on (4), which is valid only under some strong assumptions (linear dependence on the parameter) on A and $\mathbf{f}$. However, (4) can still be used for general parametric problems (1), and will give a reasonable estimation of the residual as long as $\frac{\mathrm{d}^2}{\mathrm{d}\mu^2}\mathrm{A}$ and $\frac{\mathrm{d}^2}{\mathrm{d}\mu^2}\mathbf{f}$ are small.

3 Numerical examples

Here, through two practical examples, we showcase the usefulness of the greedy MRI procedure, as well as the effectiveness of the three termination strategies based on (I), (R), and (C).

3.1 An eigenproblem in magneto-hydrodynamics

Take the generalized eigenproblem from [5]: $\mathrm{K}\mathbf{v} = \mu\mathrm{M}\mathbf{v}$ in $\mathbb{C}^n$, with $n = 4800$; it stems from modal analysis of a FE discretization of a dissipative problem in MHD. Here, we restrict our interest to the part of the spectrum with positive imaginary part: the eigenvalues are located on 3 so-called Alfvén branches around the branch point $\mu_b \approx -0.082 + 0.613\mathrm{i}$. Our aim is to approximate the number and location of eigenvalues around $\mu_0 = -0.175 + 0.5\mathrm{i}$; more precisely, we focus on the disk $D = \{\mu \in \mathbb{C} : |\mu - \mu_0| \leq 0.175\}$.

In order to cast this problem in the form (1), let $\mathbf{f} \in \mathbb{C}^n$ be a (normal Gaussian) random vector: we define the non-homogeneous problem

$$\text{find } \mathbf{u} : \mathbb{C} \to \mathbb{C}^n \quad \text{s.t.} \quad (\mathrm{K} - \mu\mathrm{M})\mathbf{u}(\mu) = \mathbf{f},$$

and build a surrogate for $\mathbf{u}$ using MRI. Then, our estimates for the eigenvalues will be the roots of the MRI denominator Q.

As a first MOR method, we apply greedy MRI: in particular, we employ indicator (I) with relative tolerance 10^{-2}, and the initial snapshots are at the $S_0 = 30$ shifted roots of unity[1] $\{\mu_0 - 0.175\mathrm{e}^{2\mathrm{i}\pi j/S_0}\}_{j=1}^{S_0}$.

Additionally, we consider a non-greedy approach as a reference: we build an MRI starting from $S = 45$ samples at shifted roots of unity $\{\mu_0 - 0.175\mathrm{e}^{2\mathrm{i}\pi j/S}\}_{j=1}^{S}$. The number of samples is chosen so that, overall, the two methods employ exactly the same number of snapshots: the only difference is *where* the samples are taken.

The results are shown in Fig. 1. At the beginning of the greedy procedure (which corresponds to a standard MRI with $S = 30$), the spectrum is approximated quite poorly; this is correctly identified by the a posteriori indicator, which shows that the prescribed tolerance is not satisfied over a large portion of D. After 15 iterations of the greedy procedure, the residual is globally below the tolerance, and the algorithm ends. We can verify that all the eigenvalues in D are well captured.

In the standard approach with $S = 45$, most of the eigenvalues are well identified, but the quality of the approximation deteriorates around μ_b. In particular, among the two surrogates obtained with 45 snapshots, the greedy one is clearly superior. However, the improved accuracy of the greedy approach is accompanied by some risks:

[1] The shifted roots of unity are chosen as sample points because they allow for very stable and efficient interpolation schemes, relying on Fast Fourier Transform. We refer to [1] for a more detailed discussion of their properties.

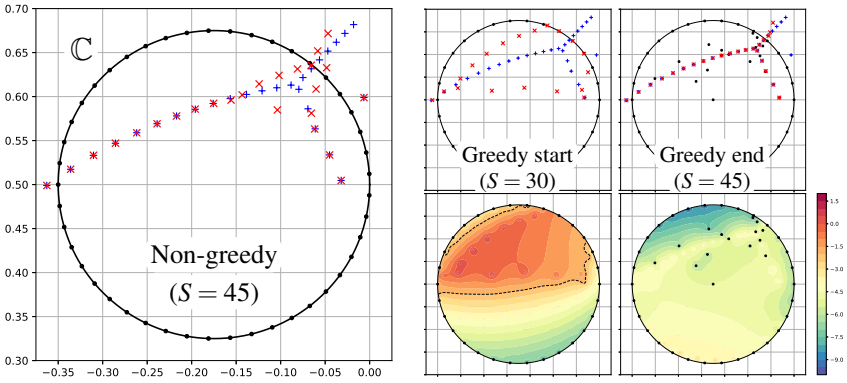

Fig. 1 Results of standard (left) and greedy (right) MOR. The exact and approximate eigenvalues are pluses and crosses, respectively, whereas the sample points are full dots. The contour plots show the logarithm of the greedy residual indicator; the dashed line represents the locus $\{\mu : r(\mu) = \text{tol}\}$, i.e. the boundary of the set where the prescribed tolerance is not satisfied.

- The locations of the greedy snapshots are close to the exact eigenvalues, since, according to the residual indicator (4), sampling there yields "the most information" for the surrogate model. However, sampling close to an eigenvalue may require solving numerically an ill-conditioned or even singular linear system.
- While the Vandermonde matrix for samples at the roots of unity has optimal condition number, adding new sample points at arbitrary locations is guaranteed to hinder the well-conditioning of the interpolation problem. Indeed, the residual indicator at the end of the greedy iterations shows a slightly unstable behavior near the bottom of D.

3.2 Frequency response of a waveguide diplexer

We consider a frequency response problem involving the FE discretization ($n = 90258$) of a waveguide diplexer [3], for frequencies μ between 9.5 and 11 GHz. We are interested in approximating the scattering parameters

$$\mathrm{S} : \mathbb{C} \ni \mu \mapsto \mathrm{I} - 2\left(\mathrm{I} + \mathrm{i}\mu\sqrt{\frac{1-(\mu_c/\mu_0)^2}{1-(\mu_c/\mu)^2}}\mathrm{F}^*\underbrace{(\mathrm{K}-\mu^2\mathrm{M})^{-1}\mathrm{F}}_{\mathrm{U}(\mu)\in\mathbb{C}^{90258\times 3}}\right)^{-1} \in \mathbb{C}^{3\times 3}, \qquad (5)$$

where we set $\mu_c = 6.56$ GHz, $\mu_0 = 10$ GHz, and the state matrix $\mathrm{U}(\mu)$ has one column for each port of the waveguide.

We build a surrogate for U using greedy RB and MRI, employing indicators (I) and (R) with relative tolerance 10^{-2}, and (C) with $\text{tol} = 10^{-4}$. The reduced tolerance for (C) can be justified by the considerably different nature of the indicator. To

obtain an approximation of S, we just replace the exact state with the surrogate one in (5).

Table 1 Timing results of greedy MOR (average over 3 simulations with the same parameters for each method). All simulations were carried out on a single node of the Fidis cluster at EPFL [6].

Method	No. of snapshots	Average time per iteration		
		state solve	indicator	surrogate update
RB+(I)	24	97.1 s	15.0 s	3.2 s
MRI+(I)	23		4.31 s	2.9 s
MRI+(R)	23		1.04 s	
MRI+(C)	22(+1)		1.08 s	

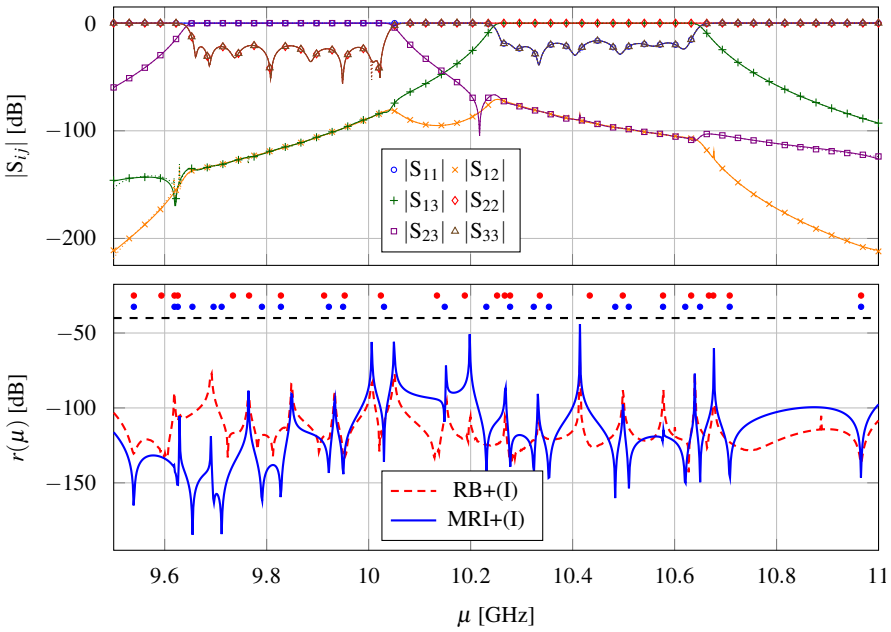

Fig. 2 Results of greedy MOR. On top the surrogate scattering parameters: the points are measurements from the original problem (5), whereas full and dotted lines are the surrogates obtained with MRI+(I) and MRI+(C), respectively. On the bottom the relative residual at the end of the greedy iterations for RB+(I) and MRI+(I); the points indicate the snapshot positions.

The results are summarized in Table 1 and visually depicted in Fig. 2. We remark that, by construction, the snapshot "history" of MRI is independent of the indicator, i.e. the parameter value $\widehat{\mu}$ which is selected at a given iteration is the same: the only effect of the choice of the indicator, besides timing, is the number of greedy iterations which are carried out before termination. In this regard, we observe that MRI+(I) and MRI+(R) yield exactly the same indicator, whereas MRI+(C) termi-

nates one snapshot sooner[2], causing some slight instability in the approximations of the scattering parameters for low frequencies, noticeable mostly in S_{13}.

A comparison of the surrogate S obtained by MRI+(I) and RB+(I) shows that the two methods yield very similar approximations, and reconstruct well the exact values. In fact, the approximated scattering parameters for RB are not included in Fig. 2, as they are almost indistinguishable from those obtained with MRI+(I). However, RB requires one more snapshot, and the locations of the snapshots (and the residual profiles) for RB and MRI are quite different.

In terms of computing time, the efficiency of MRI seems quite remarkable, particularly for the two "least intrusive" indicators: the overhead time needed for the evaluation of indicators (R) and (C) is just a fraction of the time required for computation of (I) in RB.

4 Conclusions

We have presented several a posteriori indicators which can be employed in the greedy MRI algorithm, characterized by different degrees of intrusiveness and applicability. Good approximation properties, as well as a substantial speed-up in residual computation with respect to classical methods, have been observed in two numerical examples.

Acknowledgements This work has been funded by the Swiss National Science Foundation through FNS Research Project number 182236.

References

1. Austin, A.P., Kravanja, P., Trefethen, L.N.: Numerical Algorithms Based on Analytic Function Values at Roots of Unity. SIAM J. Numer. Anal., **52**, 1795–1821 (2014)
2. Benner, P., Gugercin, S., Willcox, K.: A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems. SIAM Rev., **57**, 483–531 (2015)
3. De La Rubia, V. and Mrozowski, M.: A compact basis for reliable fast frequency sweep via the reduced-basis method. IEEE Trans. Microw. Theory Tech., **66**, 4367–4382 (2018)
4. Ionita, A.C., Antoulas, A.C.: Data-Driven Parametrized Model Reduction in the Loewner Framework. SIAM J. Sci. Comput., **36**, A984–A1007 (2014)
5. Kooper, M.N., Van Der Vorst, H.A., Poedts, S., Goedbloed, J.P.: Application of the implicitly updated Arnoldi method with a complex shift-and-invert strategy in MHD. J. Comput. Phys., **118**, 320–328 (1995)
6. SCITAS: EPFL Fidis cluster webpage. Online document (2020). URL http://www.epfl.ch/research/facilities/scitas/hardware/fidis. Cited Feb 11, 2020
7. Pradovera, D.: Interpolatory rational model order reduction of parametric problems lacking uniform inf-sup stability. Accepted for publication in SINUM (2020)

[2] Here we are discarding the final extra snapshot used to check the termination condition (C). If it had been included, we would have recovered the same surrogate model as MRI+(I)/(R).