DATA-DRIVEN ALGORITHMS FOR SIGNAL PROCESSING WITH TRIGONOMETRIC RATIONAL FUNCTIONS[*]

HEATHER WILBER[†], ANIL DAMLE[‡], AND ALEX TOWNSEND[§]

Abstract. Rational approximation schemes for reconstructing periodic signals from samples with poorly separated spectral content are described. These methods are automatic and adaptive, requiring no tuning or manual parameter selection. Collectively, they form a framework for fitting trigonometric rational models to data that is robust to various forms of corruption, including additive Gaussian noise, perturbed sampling grids, and missing data. Our approach combines a variant of Prony's method with a modified version of the AAA algorithm. Using representations in both frequency and time space, a collection of algorithms is described for adaptively computing with trigonometric rationals. This includes procedures for differentiation, filtering, convolution, and more. A new MATLAB software system based on these algorithms is introduced. Its effectiveness is illustrated with synthetic and practical examples drawn from applications including biomedical monitoring, acoustic denoising, and feature detection.

Key words. rational functions, signal processing, AAA algorithm, Prony's method

AMS subject classifications. 41A20, 94A12

1. Introduction. Recovering functions from noisy, incomplete, or corrupted samples is a ubiquitous task in signal and data processing [39]. Here, we recover underlying signals that contain impulses, shocks, or other algebraic singularities that can cause traditional Fourier-based methods to underperform or fail. Examples where these signals appear include sensor monitoring and event detection tasks in seismology and oceanography [14, 41], biomedical signal processing [23, 25], and time evolution along rays in nonsmooth media [13, 51]. We present a novel computing framework based on data-driven approximation with two complementary representations: (1) barycentric trigonometric rational approximations and (2) their Fourier transforms, which take the form of short sums of complex exponentials. Toggling between these two representations lets us overcome computational and data-related challenges.

Several families of functions, including rationals, wavelets, and radial basis functions, are well-suited for resolving sharp features in data and modeling phenomena with slow-decaying spectral content [55, 57]. However, methods that employ these functions often require the a priori selection of shape parameters [52], mother/father wavelets [17], initial pole configurations [29], or special rational basis functions [23]. One must carefully select parameters to avoid numerical instability and computational inefficiency. In contrast, we introduce flexible, data-driven, general-purpose software tools that can be applied without special knowledge about the locations or types of singularities in the signal. Our methods construct trigonometric rational representations of signals, and we develop a collection of algorithms for computing adaptively and efficiently with them. Our approach combines adaptations of two primary approximation methods, the AAA algorithm for rational approximation [45] and a Fourier inversion technique based on Prony's method [10]. Both of these methods

---

[†]Oden Institute for Computational and Engineering Sciences, University of Texas at Austin, Austin, TX , 78712-1229, United States (heather.wilber@oden.utexas.edu).

[‡]Department of Computer Science, Cornell University, Ithaca, NY 14853-4201, United States (damle@cornell.edu).

[§]Department of Mathematics, Cornell University, Ithaca, NY 14853-4201, United States (townsend@cornell.edu).

Barycentric form	Exponential sums
Differentiation (closed-form formula) [7]	Filtering and recompression [34]
Imputing missing data [45]	Pole symmetry preservation [10]
Stable evaluation [2, 36]	Robustness to noise [49]
Rootfinding [45], identifying extrema	convolution [49], cross-correlations

TABLE 1
Operations that are efficient and robust in the two representations. Having both representations and toggling between them allows us to compute a range of operations in signal processing.

can automatically construct near-optimal rational approximations to functions, and we develop a framework that unifies the representations they construct. The result is an automatic rational approximation method that enjoys two main advantages: (1) it is more robust to various forms of corruption than either method alone, and (2) taken collectively, the two representations are efficient for performing a range of post-processing operations (see Table 1).

Periodicity is a common assumption encountered in the practical analysis of signals, and the trigonometric rational functions we use arise naturally from this assumption. However, with some adjustments, our ideas can also be applied using rational functions in the non-periodic setting (see subsection 5.9). One might also consider settings where rational functions are fit to samples in frequency space, rather than signal space. Extensive literature on this topic is available from the digital filtering and optimal control communities [22, 46].

1.1. The approximation problem. Let $f\colon [0,1) \to \mathbb{R}$ be an unknown continuous periodic function of bounded variation, and suppose that for some integer N, we observe $2N+1$ noisy samples of f at $T \coloneqq \{x_j\}_{j=0}^{2N}$, i.e., $y_j = f(x_j) + s_j$ for $0 \le j \le 2N$. Here, s_j can be: (i) additive white Gaussian noise (i.i.d. normally distributed), (ii) popcorn noise (sparsely corrupted or arbitrarily large errors), or (iii) bounded deterministic errors. Throughout, we assume $\|f\|_\infty = 1$, where $\|\cdot\|_\infty$ is the infinity norm on $[0,1)$, and that the mean value of f over $[0,1)$ is 0. Our central approximation problem is to fit a special class of trigonometric rational functions (see section 2) to the $2N+1$ samples to construct r_m, a type $(m-1, m)$ trigonometric rational where m is selected adaptively so that the sampling error satisfies $\max_{x_j \in T} |f(x_j) - r_m(x_j)| \le \epsilon$, where $0 < \epsilon < 1$ is a tolerance parameter.

In practice, samples of f are often available under far from ideal circumstances. For example, T may consist of poorly distributed (i.e., not equally-spaced) points due to missing or corrupted data, N may be too small to adequately resolve features of interest, or frequencies of interest may be cut off or otherwise distorted during observations. We are interested in how well f can be recovered under such conditions.

1.2. Software. Once an approximant is constructed, we want to compute with it reliably. Answering the following questions has shaped our software's development:

1. *Which applications are of importance?* Our methods supply a general-purpose approach for working with noisy samples from periodic, univariate signals. They require no a priori knowledge about locations or types of singularities, and they are designed to be flexible enough for use in a range of applications that involve the detection and identification of events (e.g., ECG and geophone monitoring tasks, engineering and financial applications where change-point detection is important, and in the analysis of dynamical systems, such as periodic contagion modeling).

2. *What properties should our approximants have?* Like f, we want approximants

to be periodic, real-valued, and continuous on $[0, 1)$. We employ trigonometric rational functions, which are the periodic analogue of rational functions.

3. *What form of approximant should we use?* Trigonometric rationals can be expressed in many forms, but not all of these forms are numerically stable. For stable evaluation and rootfinding, we use the barycentric formula [6]. For efficient recompression and operations carried out in Fourier space, we represent the Fourier transforms of our trigonometric rational approximants using short sums of weighted complex exponentials (see Lemma 2.1).

4. *Which tools should we provide?* We offer a basic set of computational tools. This includes simple algebraic operations (addition, products), calculus-based operations (integration, differentiation), and tools for filtering, (de)convolving, and rootfinding/polefinding. Whenever possible, we automatically recompress representations as trigonometric rationals/exponential sums to maintain efficiency. One can combine these tools to perform more complicated tasks.

Accompanying this work is the open-source code REfit [59], which is written in MATLAB and uses two classes called `rfun` and `efun`. An `rfun` object stores a representation of f as a barycentric trigonometric rational function. An `efun` object stores a representation of the Fourier transform of f as a weighted sum of complex exponentials. After an `rfun` or `efun` object is constructed, it can be manipulated and analyzed through the operations implemented in the package (see Table 2). The commands are overloaded so that they can be applied to either type of object, and binary operators can be used between objects of different type.

Table 2
A selection of REfit commands.

command	Operation
`+, -, .*, ./`	basic arithmetic
`diff(·), cumsum(·)`	differentiation, indefinite integration
`conv(·)`	convolution
`corr(·,·)`	cross-correlation

1.2.1. Connections to other work. In addition to the AAA algorithm [45] and a large body of work on Prony's method [9, 10, 47, 49, 50], several recent developments in rational computing are related to our work. In [18], the authors also develop approximation methods that combine Prony's method with the AAA algorithm. Their work focuses on the recovery of parameters for extended models of exponential sums with polynomial coefficients[1] that can also include purely oscillatory terms, whereas our focus is on the development of general-purpose software and rational approximation tools. We consider the setting where exponential sums are applied in Fourier space (corresponding to trigonometric rationals in signal space), and in contrast, they use barycentric rational approximations in Fourier space to recover sums of exponentials in signal space. For this reason, they do not make use of trigonometric rational functions. While the models proposed here for the recovery of signals are different from those in [18], the developments in [18] are relevant to our setting. For example, the derivatives of a trigonometric rational function correspond

---

[1]This corresponds to using rational functions with poles of higher multiplicity. We find that it is sufficient for our purposes to use clusters of simple poles.

to exponential sums in Fourier space with polynomial coefficients. The parameters of such a model can be exactly recovered using the ideas in [18].

In [3], a trigonometric variant of AAA is introduced in the context of conformal mapping. It is similar to the pronyAAA method we introduce in subsection 2.3.1. However, a key difference is that it is not modified to account for any special connection to approximations in Fourier space. Other popular and effective rational approximation methods include the RKFIT algorithm (and software package) [5], the vector fitting method [29], and the Loewner framework of Mayo and Antoulas [38], which is closely connected to the AAA algorithm. These methods have not been explicitly adapted to the periodic setting, and the first two methods are not suited to our needs as they require and can be sensitive to a set of initialization parameters (e.g., guesses of the location of the poles). Recently, AAA has been combined with RKFIT to overcome this difficulty [20]. The development of our software package is largely inspired by Chebfun [19, 60], which is primarily designed for computing automatically and adaptively with smooth functions. In some respects, our software serves as an analogue to Chebfun for computing with non-smooth functions.

The rest of this paper is organized as follows: In section 2 we briefly review trigonometric rational functions and the barycentric form. We introduce the trigonometric variant of the AAA algorithm that we use to construct barycentric trigonometric interpolants (see subsection 2.3). In subsection 2.4, we apply the regularized Prony's method (RPM) to construct approximations in Fourier space. In section 3, we introduce stable Fourier and inverse Fourier transform methods for moving between representations in the time and frequency domains. In sections 4 and 5, we give examples and descriptions of the algorithms for computing with these representations.

2. Trigonometric rational functions and their Fourier transforms.
The trigonometric rationals are the periodic analogue of rational functions [35, 37]. A trigonometric rational of period 1 is the quotient of two trigonometric polynomials of period 1, i.e., a function of the form

$$(2.1) \qquad r(x) = \frac{p_\ell(x)}{q_m(x)} = \frac{\sum_{j=-\ell}^{\ell} a_j e^{2\pi i j x}}{\sum_{j=-m}^{m} b_j e^{2\pi i j x}}, \quad x \in [0, 1).$$

We call r a type (ℓ, m) trigonometric rational function. We follow the convention that unless stated explicitly, a type (ℓ, m) function is in reduced form, meaning that p_ℓ and q_m have no zeros in common. We restrict our interest to the family of period 1 trigonometric rationals r_m of type $(m-1, m)$ that are real-valued and continuous on $[0, 1)$. Furthermore, we assume that the roots of the denominator q_m are simple. Under these assumptions, if η_j is a root of q_m, then so is its complex conjugate $\bar{\eta}_j$, as well as $\eta_j \pm K$, where K is any integer. We say that a pole of r_m is any root, η_j, of q_m such that $0 \le Re(\eta_j) < 1$. In the same way, any root of p_{m-1} with a real part in the interval $[0, 1)$ is called a zero of r_m.

2.1. Why trigonometric rationals?.
Three key properties of this family of functions make them ideal for our setting. First, like standard rational functions, they are especially effective at resolving singularities. For example, for particular choices of f, such as $f(x) = |x - 1/2| - 1/4$, it is known that trigonometric rationals of type (m, m) can converge to f at a root-exponential rate with respect to m [55]. Second, these functions can be represented efficiently in Fourier space. We make extensive use of the following fact:
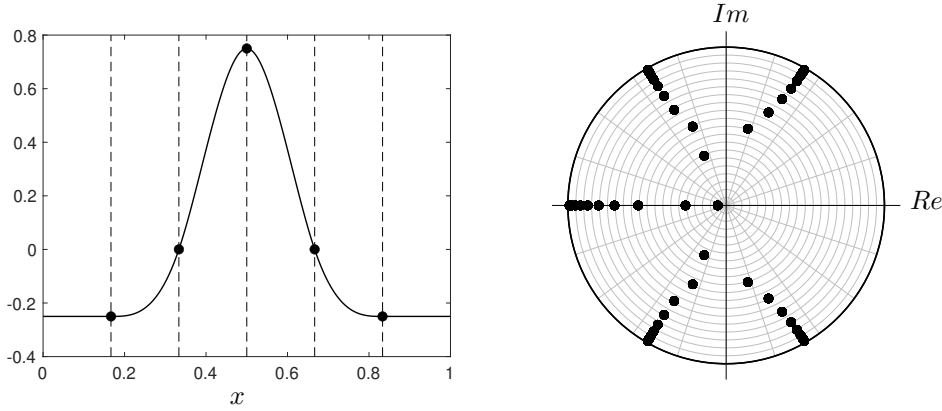
FIG. 1. *Left: A trigonometric rational function r_m (m = 44) approximating the shifted and scaled cubic B-spline on the interval $[0, 1)$ is constructed and plotted. The knots of the spline occur at the dotted vertical lines. Away from the singularities, the absolute approximation error is on the order of 10^{-11}. Approximate knot locations (black dots) are automatically computed using the poles of the rational $\tilde{r}_m(z) = r_m(x)$, where $z = e^{2\pi i x}$ (see (2.3)). Right: The poles of $\tilde{r}_m$ of magnitude ≤ 1 are plotted in the unit disk in the complex plane. They cluster toward the points $e^{2\pi i x_k}$ on the unit circle, where each x_k is a knot in the spline.*

LEMMA 2.1. *Let $r_m(x)$ be a type $(m-1, m)$ nonzero trigonometric rational function that is real-valued and continuous on $[0, 1)$ with exactly $2m$ simple poles. Let $\{\eta_j\}_{j=1}^m$ denote the collection of poles with $Im(\eta_j) > 0$. If the Fourier coefficients of r_m are given by $\{(\hat{r}_m)_k\}_{k=-\infty}^{\infty}$, then there exist ω_j such that*

$$(2.2) \qquad (\hat{r}_m)_k = R_m(k) := \begin{cases} \sum_{j=1}^m \omega_j e^{\alpha_j k}, & k \geq 0, \\ \sum_{j=1}^m \overline{\omega}_j e^{-\overline{\alpha}_j k}, & k < 0, \end{cases}$$

where $\alpha_j = 2\pi i \eta_j$. Here, $\overline{\omega}_j$ is the complex conjugate of ω_j.

Proof. See [53, Ch. 4]. □

Lemma 2.1 shows that the Fourier series of r_m can be represented as sum of m weighted decaying exponentials. We say that $\mathcal{F}(r_m) = R_m$ is the Fourier transform of r_m, and similarly, $\mathcal{F}^{-1}(R_m) = r_m$ is the inverse Fourier transform of R_m.

Third, trigonometric rationals can be used for feature detection. For example, in Figure 1, we plot the poles $z_j = e^{2\pi i \eta_j}$, $Im(\eta_j) > 0$, of the rational function

$$(2.3) \qquad \tilde{r}_m(z) = \frac{z^m p_{m-1}(z)}{z^m q_m(z)}, \qquad z = e^{2\pi i x},$$

where $r_m(x) = p_{m-1}(x)/q_m(x)$ is a trigonometric rational approximation to a shifted and scaled uniform cubic B-spline [16] f with knots at $x = \{1/6, \ldots, 5/6\}$. The knots are not easily identifiable in a plot of f, but the poles of $\tilde{r}_m$, which are chosen adaptively via Algorithm 2.3, cluster toward the singularities and reveal their locations. Using the five poles in $\{z_j\}_{j=1}^m$ with the largest magnitudes, estimates of the knot locations are given by $\tilde{x}_k = \arg(z_k)/2\pi$ for $1 \leq k \leq 5$. The radial coordinates of the poles in Figure 1 also encode information about f in the frequency domain. The Fourier coefficient $\hat{f}_k$ is well-approximated by $(\hat{r}_m)_k = R_m(k)$, where R_m is as in (2.2). The terms in R_m with $|z_j| \ll 1$ have negligible influence when k is small, so they capture aspects of the signal that are only observable at low frequencies.

2.2. Barycentric trigonometric rational functions. Large numerical errors can be incurred when evaluating trigonometric rationals that are numerically constructed using (2.1) directly [21, 55]. Another natural way to represent r_m is in a pole-residue format with respect to the pairs of poles $(z_j = e^{2\pi i \eta_j}, 1/\overline{z}_j)$ of the rational function $\tilde{r}_m$ in (2.3). Evaluation using this format is often reliable in practice, but it can potentially result in catastrophic cancellation if the evaluation point is too close to the poles. In the non-periodic setting, the AAA algorithm [45] safeguards against such instabilities by using barycentric rational interpolants with backward-stable evaluation on the interval of approximation [2, 36]. Type $(m-1, m-1)$ barycentric rational interpolants of a function f on $[0,1)$ have the form

$$(2.4) \qquad v_m^{\gamma,t} = \frac{\tilde{n}_m(x)}{\tilde{d}_m(x)} = \sum_{j=1}^{m} \frac{\gamma_j}{x - t_j} f(t_j) \bigg/ \sum_{j=1}^{m} \frac{\gamma_j}{x - t_j},$$

where for $1 \leq j \leq m$, $t_j \in [0,1)$. The trigonometric analogue of these interpolants is described in [35], and more recently, in [3, 37]. Type $(m-1, m)$ barycentric trigonometric rationals on $[0,1)$ take the following form:

$$(2.5) \qquad r_m^{\gamma,t}(x) = \frac{n_{m-1}(x)}{d_m(x)} = \frac{\sum_{j=1}^{2m} \gamma_j f_j \cot\left(\pi(x - t_j)\right)}{\sum_{j=1}^{2m} \gamma_j \cot\left(\pi(x - t_j)\right)}, \qquad \sum_{j=1}^{2m} \gamma_j f_j = 0.$$

The interpolating points $t = \{t_1, \ldots, t_{2m}\}$, always assumed to be distinct, are called barycentric nodes, $\gamma = \{\gamma_1, \ldots, \gamma_{2m}\}$ are called barycentric weights, and $f_j = f(t_j)$. It is easily shown that $r_m^{\gamma,t}(t_j) = f_j$ whenever each γ_j is nonzero [6], but it is not obvious from (2.5) that $r_m^{\gamma,t}$ is in fact a type $(m-1, m)$ trigonometric rational. This can be seen by using the trigonometric polynomial $\ell_t(x) = \prod_{j=1}^{2m} \sin(\pi(x - t_j))$, and rewriting $r_m^{\gamma,t}$ as $\ell_t n_{m-1}/\ell_t d_m$ (see [35]). The condition $\sum_{j=1}^{2m} \gamma_j f_j = 0$ enforces that the numerator is a trigonometric polynomial of degree $m-1$, rather than m. It is not clear from (2.5) where the poles of $r_m^{\gamma,t}$ lie, and in particular, whether they lie off $[0,1)$ (see subsection 2.3). Stability properties (and other advantages) of the barycentric form were popularized in the context of polynomial interpolation [8], where γ is always chosen so that $d_m(x) = 1$. Analyses can be found in [2, 36], with discussion on stability properties in the more general case, where $d_m(x) \neq 1$, in [21, Sec. 2.2–2.3].

2.3. Approximations in time. In the noiseless setting, the barycentric rational $r_m^{\gamma,t}$ can be directly constructed from samples of f using an analogue of the AAA algorithm (pronyAAA) that we introduce here. This supplies a fast, automated way to construct trigonometric rational representations of signals. A major advantage of AAA-type methods over other approximation schemes is that they can be blithely applied to samples from non-uniform grids and can even be used for recovering functions defined on disjoint sets of support [45]. In particular, these methods are robust to missing samples. In Figure 2, we use pronyAAA to recover a function from data that has been deleted in several corrupted regions.

2.3.1. Constructing barycentric trigonometric interpolants. Our construction algorithm is a straightforward extension of the AAA algorithm (see Algorithm 2.1), with differences being the use of trigonometric basis functions, the restriction that the numerator is always of degree $m-1$, and the requirement that the constructed interpolant has an even number of interpolating points.[2] These restrictions are so that the approximant corresponds to a length m exponential sum

---

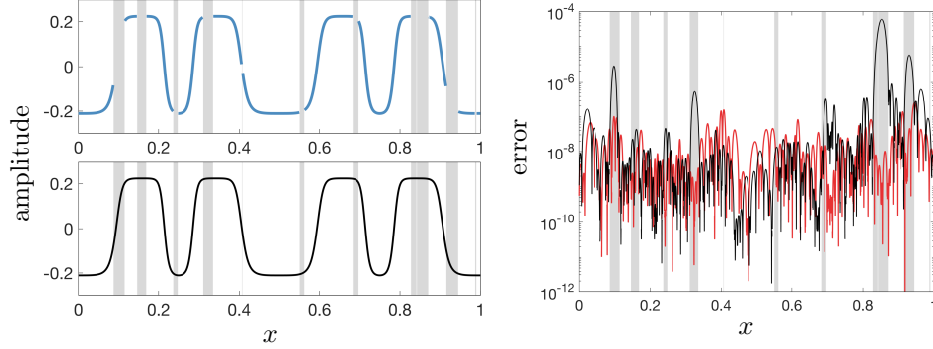[2]This is so that $r_m^{\gamma,t}$ may have an even number of simple poles.

Fig. 2. *Left: 3000 samples from a function f are taken on an equally spaced grid from $[0, 1)$. However, the data inside the grey regions are corrupted and thus deleted from the sample (a total of 609 observations are missing). The available data are plotted (blue, upper panel) and used to construct a type $(34, 35)$ trigonometric barycentric rational interpolant r_{35} via pronyAAA. The resulting approximant, evaluated on an equally spaced grid and plotted (black, lower panel), imputes the missing data in the grey regions. Right: The absolute error $|f(x) - r_{35}(x)|$ is plotted in black on a logarithmic scale against the original grid of sampled points. The absolute error when pronyAAA is applied to an uncorrupted sample of f with no missing data is also displayed (red).*

in Fourier space. As with AAA, pronyAAA employs a greedy residual minimization method to select interpolating points and iteratively build up an interpolant. We briefly discuss the process here, and refer to [45] for details.

Let $\{f_0, \ldots, f_{2N}\}$ be the samples of a function f we seek to approximate, where $T = \{x_j\}_{j=0}^{2N} \subset [0, 1)$ are the sample locations and $f(x_j) = f_j$. We first describe how the barycentric weights are updated at each iteration. Suppose that at the m^{th} iteration, the barycentric support points for $r_m^{\gamma, t}(x)$ in (2.5) are $t = \{t_1, \ldots, t_{2m}\} \subset T$. Let $\widetilde{T} = T \setminus \{t_1, \ldots, t_{2m}\}$. We must now choose $\{\gamma_j\}_{j=1}^{2m}$. Let $\gamma = (\gamma_1, \ldots, \gamma_{2m})^T$. As in (2.5), $r_m^{\gamma, t} = n_{m-1}/d_m$, where n_{m-1}, d_m are trigonometric polynomials. We select γ as the solution to the following constrained optimization problem:

$$(2.6) \quad \min_{\gamma \in \mathbb{C}} \sum_{x_j \in \widetilde{T}} (f(x_j) d_m(x_j) - n_{m-1}(x_j))^2, \quad \text{s.t.} \quad \sum_{j=1}^{2m} f(t_j) \gamma_j = 0, \quad \|\gamma\|_2 = 1.$$

This is similar to the constraint applied in the standard AAA algorithm (see step 2, line (iii), of Algorithm 2.1), except that it additionally requires n_{m-1} to be of degree $m-1$. The constraint is satisfied if we select a vector of the form $\gamma = Q\tilde{\gamma}$, $\|\tilde{\gamma}\|_2 = 1$, where Q is a $2m \times (2m-1)$ matrix with orthonormal columns that span the space orthogonal to the vector $(f(t_1), \ldots, f(t_{2m}))$. Since for each $x_j \in \widetilde{T}$,

$$(2.7) \quad f(x_j) d_m(x_j) - n_{m-1}(x_j) = \sum_{\ell=1}^{2m} (f(x_j) - f(t_\ell)) \gamma_\ell \cot(\pi x_j - \pi t_\ell),$$

we see that $\tilde{\gamma}$ is given by the last right singular vector of the matrix CQ, where $C \in \mathbb{R}^{|\widetilde{T}| \times 2m}$ has entries of the form $C_{j\ell} = (f(x_j) - f(t_\ell)) \cot(\pi x_j - \pi t_\ell)$, where each x_j is a unique member of $\widetilde{T}$.

Once γ is computed, we have constructed $r_m^{\gamma, t}$ and the iteration is complete. If the error $\max_{x_j \in \widetilde{T}} |f(x_j) - r_m^{\gamma, t}(x_j)|$ is not sufficiently small, we begin the next iteration by choosing two additional interpolating points (as we require an even number of interpolating points) from $\widetilde{T}$. It is not clear how to efficiently choose two

7

Algorithm 2.1 The standard AAA algorithm.

Input: tolerance parameter ϵ, sample locations $T = \{x_0, \ldots, x_{2N}\}$, samples $\{f(x_0), \ldots, f(x_{2N})\}$.

Output: barycentric support points t and weights γ defining $v_m^{\gamma,t}$ in (2.4).

1. Set $err = 1$, $m = 0$, $t = \{\}$, $\widetilde{T} = T$, $v_0^{\gamma,t} = 0$.

2. while $err > \epsilon$
  (i) Set $t_{m+1} = \text{argmax}_{x_j \in \widetilde{T}}\{|f(x_j) - v_m(x_j)|\}$.
  (ii) $t \leftarrow t \cup \{t_{m+1}\}$, $\widetilde{T} \leftarrow T \setminus t$.
  (iii) With $\tilde{n}_{m+1}$, $\tilde{d}_{m+1}$ as in (2.4), choose $\gamma = (\gamma_1, \ldots, \gamma_{m+1})^T$ such that it solves the following constrained optimization problem:
  $$\text{minimize } \sum_{j=1}^{|\widetilde{T}|} \left(f(x_j)\tilde{n}_{m+1}(x_j) - \tilde{d}_{m+1}(x_j) \right)^2 \text{ s.t. } \|\gamma\|_2 = 1.$$
  (iv) Set $err \leftarrow \max_{x_j \in \widetilde{T}} |f(x_j) - v_{m+1}^{\gamma,t}(x_j)|$.
  (v) $m \leftarrow m + 1$.
  end while
4. Compute poles and residues (see subsection 5.7 and [45, eq. 3.11]).
5. If there are spurious poles, apply cleanup (see subsection 2.3.2 and [45, sec. 4]).

---

points in a single step, so we introduce an interim step. The first point is chosen as $t_{2m+1} = \text{argmax}_{x_j \in \widetilde{T}}|f(x_j) - r_m^{\gamma,t}(x_j)|$, and then t and $\widetilde{T}$ are updated appropriately. To pick the second point, we construct an interim trigonometric interpolant $r_{m+1/2}^{\gamma,t}$ with an odd number of interpolating points, using the basis functions $\csc(\pi(x - t_j))$ instead of $\cot(\pi(x - t_j))$. The different basis functions ensure that $r_{m+1/2}^{\gamma,t}$ is a trigonometric rational: see [35]. The weights in $r_{m+1/2}^{\gamma,t}$ are updated by solving a problem similar to (2.6), though the constraint differs due to the different basis functions [35]. The second interpolating point is then chosen as $t_{2m+2} = \text{argmax}_{x_j \in \widetilde{T}}|f(x_j) - r_{m+1/2}^{\gamma,t}(x_j)|$. A different strategy for choosing the second point at each step is used in [3], but we find that such a choice adversely impacts convergence in our setting.[3]

2.3.2. Spurious poles. As with AAA, nothing in the pronyAAA algorithm directly controls where the poles of $r_m^{\gamma,t}$ occur. The advantage of this approach is that, unlike fixed-pole approximation methods, the pole locations are adaptively determined and automatically cluster in patterns that are highly effective for resolving singularities [56]. However, one drawback to this approach is that the conjugate-pair symmetry of the poles is not explicitly enforced. A more problematic issue is that so-called "spurious poles" may appear on the interval of approximation.

Poles that are undesirable or unnecessary are called spurious [6, 45]. They arise as artifacts related to numerical error, but can also appear within mathematically valid solutions to an interpolation problem constrained by (2.6). For example, a spurious pole may co-occur with a nearby zero that effectively cancels out its influence except within a small interval I, where $I \subset [x_j, x_{j+1}]$ for some neighboring sample locations $x_j < x_{j+1}$, so that for $x \notin I$ it holds that $|r_m^{\gamma,t}(x) - f(x)| < \epsilon$. This is a perfectly

---

[3]One might wonder if it is possible to choose the interpolation points so that the pole locations are restricted (or symmetries in pole locations are enforced). Such a challenge is of interest in rational-based numerical methods [15, 61]. The Floater–Hormann interpolants are one such example [7], but they have slower convergence rates than AAA-based interpolants.

Algorithm 2.2 The pronyAAA algorithm.

Input: tolerance parameter ϵ, sample locations $T = \{x_0, \ldots, x_{2N}\}$, samples $\{f(x_0), \ldots, f(x_{2N})\}$.
Output: barycentric support points t and weights γ defining $r_m^{\gamma,t}$ in (2.5).

1. Set $m = 1$, $t = \{t_1 = \mathrm{argmax}_{x_j \in T}|f(x_j)|, t_2 = \mathrm{argmax}_{x_j \in T \setminus t_1}|f(x_j)|\}$, $\widetilde{T} = T \setminus t$.
2. Solve for γ as in (2.7) to define $r_1^{\gamma,t}$. Set $err = \max_{x_j \in \widetilde{T}}|f(x_j) - r_1(x_j)|$.

3. While $err > \epsilon$
       for $\ell = 1, 2$
           (i) Set $t_{2m+\ell} = \mathrm{argmax}_{x_j \in \widetilde{T}}|f(x_j) - r_{m+(\ell-1)/2}^{\gamma,t}|$.
           (ii) $t \leftarrow \{t_1, \ldots, t_{2m+\ell}\}, \quad \widetilde{T} \leftarrow \widetilde{T} \setminus t_{2m+\ell}$.
           (iii) Update $\gamma = (\gamma_1, \ldots, \gamma_{2m+\ell})^T$ by solving a constrained
           optimization problem on $\widetilde{T}$ (see (2.7) and discussion).
       end for
       $err \leftarrow \max_{x_j \in \widetilde{T}}|f(x_j) - r_{m+1}^{\gamma,t}(x_j)|$.
       $m \leftarrow m + 1$.
   end while

4. Compute poles and residues (see subsection 5.7).
5. If there are spurious poles, apply cleanup routine (see subsection 2.3.2).

---

acceptable way to solve (2.6), but it leads to a solution where $\|r_m^{\gamma,t} - f\|_\infty$ is unbounded.

Spurious poles are eliminated in the standard AAA algorithm with an additional cleanup routine [45], which we adapt to our setting. Pairs of spurious poles are detected via their small residues or by the fact that they are real-valued. The barycentric nodes nearest to the spurious poles are eliminated. This forces the degree of the interpolant to decrease and requires the barycentric weights to be recomputed, which changes the number and location of the poles. When poles cannot be eliminated without destroying the approximation, it is usually because the function being interpolated is not well-approximated by low to moderate type $(m-1, m)$ trigonometric rationals. We observe this, for example, when pronyAAA is applied to samples perturbed by additive Gaussian noise. An example is described in subsection 4.2 that involves the reconstruction of ECG data from 645 samples. The direct application of pronyAAA to the noisy data is problematic. It results in 62 spurious poles appearing on $[0, 1)$ that create artificial peaks in the signal reconstruction. To overcome this, we need a method of rational approximation that is robust to Gaussian noise.

2.4. Approximations in Fourier space. The pronyAAA algorithm constructs trigonometric rational representations to signals directly from samples, but its exclusive use is inadequate in many practical settings. For example, one cannot apply it in the presence of Gaussian noise, and the barycentric form is not conducive to efficient recompression techniques (see subsection 5.1). These issues can often be remedied by representing f in Fourier space using the exponential sums in (2.2).

To construct the sums, we require the Fourier coefficients of f. We use the fast Fourier transform (FFT) to compute the coefficients $v = (\hat{f}_0, \ldots, \hat{f}_N)^T$ associated with samples $\{f(x_j)\}_{j=0}^{2N}$, where $x_j = j/(2N + 1)$. When f is not a bandlimited function (or has a bandlimit $> N$), this process introduces error into the Fourier coefficients. We describe the error with the following notion:

DEFINITION 2.2 (ϵ-resolution). *For $0 < \epsilon < 1$, the ϵ-resolution of f is the smallest non-negative integer N_ϵ such that*

$$\|f - f_{trunc}\|_\infty \leq \epsilon,$$

where f_{trunc} is the best $\mathcal{L}^\infty([0,1))$ projection of f onto the functions of bandlimit N_ϵ.

For bandlimited functions, N_0 is the bandlimit of f. The ϵ-resolution for non-bandlimited functions can be understood in relation to the smoothness of f and its region of analyticity in the complex plane [55]. In our setting, the assumption is that N_ϵ is large. However, when f is well-approximated by a type $(m-1, m)$ trigonometric rational, Lemma 2.1 indicates that $\mathcal{F}(f)$ can be represented with far fewer degrees of freedom via the exponential sum R_m. One way to find R_m is by fitting the nonlinear model $R_m(k) = \sum_{j=1}^m \omega_j e^{\alpha_j k}$ to the Fourier coefficients $\hat{f}_k$, $0 \leq k \leq N$. We emphasize that m is unknown in the general setting and must be determined adaptively.

2.4.1. Regularized Prony's method. To construct R_m, we follow an idea in [10] and use the regularized version of Prony's method (RPM) from [9]. Variants of this method go by many names across various disciplines, and we refer to [49] for an overview. The problem of finding R_m can be recast as a structured low rank approximation problem involving Hankel matrices. One can understand the connection using the following lemma, a version of which was first proven by Prony in 1795 [50].

LEMMA 2.3. *Let N be an even integer.[4] Let R_m be as in Lemma 2.1 and let H_{R_m} be an $(N/2 + 1) \times (N/2 + 1)$ Hankel matrix with entries $(H_{R_m})_{k\ell} = R_m(k + \ell)$, $0 \leq k, \ell \leq N/2$, where $N \geq 2m$. Let $\mathcal{P}_m(z) = \sum_{j=0}^m c_j z^j$ be a polynomial with roots $z_k = e^{\alpha_k}$, $1 \leq k \leq m$. Then, $\mathrm{rank}(H_{R_m}) = m$, and the null space of H_{R_m} is spanned by $\{c, Sc, \ldots, S^{N/2-m}c\}$, where $c = (c_0, \ldots, c_m, 0, \ldots, 0)^T$ and S is the forward shift matrix.*

Proof. See [48, Lem. 2.1]. □

In our case, we seek $R_m \approx \mathcal{F}^{-1}(f)$. Lemma 2.3 indicates that this is equivalent to finding a Hankel matrix H_{R_m} of rank m so that $H_{R_m} \approx H_v$, where $(H_v)_{jk} = \hat{f}_{j+k}$, $0 \leq j, k \leq N/2$. Moreover, it shows that the complex exponentials in (2.2) are the roots of a special polynomial, often referred to as Prony's polynomial [49], whose coefficients form a vector in the null space of H_{R_m}. The regularized Prony's method (RPM), described in pseudocode in Algorithm 2.3, finds a vector c of polynomial coefficients in the numerical null space of H_v, i.e., c such that $\|H_v c\|_2 \leq \epsilon$. Unlike in Lemma 2.3, H_v is not exactly of rank m, so the polynomial $\mathcal{P}(z)$ with monomial coefficients given by the entries of c generally has $N/2$ roots. Since we only want decaying exponentials in R_m, we keep only the m roots with modulus < 1. The exponents $\{\alpha_j\}_{j=1}^m$ for R_m are determined from these roots, and a least squares fit to a subset of Fourier coefficients of f supplies the weights $\{\omega_j\}_{j=1}^m$ (see Algorithm 2.3).

A qualitative error bound in terms of the singular values of H_v and more details about the RPM are given in [9]. In some settings, such as in the example in Figure 3, a good choice for the tolerance parameter ϵ may be unclear. In this case, we modify Algorithm 2.3 so that ϵ is chosen automatically by detecting gaps in the small singular values of H_v that indicate the presence of a numerical null space. Algorithm 2.3 naively implemented has an $\mathcal{O}(N^3)$ cost because it requires finding the singular value decomposition (SVD) of H_v. This is improved if one finds the SVD with an algorithm

---

[4]The statement and proof can be adjusted to account for odd N [9].

Algorithm 2.3 The regularized Prony method.

Input: tolerance parameter ϵ and Fourier coefficients $v = (\hat{f}_0, \ldots, \hat{f}_N)^T$.
Output: $\{(\omega_j, \alpha_j)\}_{j=1}^m$ defining R_m in (2.2), so that $|R_m(j) - \hat{f}_j| \approx \epsilon$.

1. Construct the Hankel matrix H_v, where $(H_v)_{j+k} = \hat{f}_{j+k}$ for $0 \leq j, k \leq N/2$.
2. Compute the SVD of H_v to find c, where $\|H_v c\|_2 \leq \epsilon$, $\|c\|_2 = 1$.
3. Set $\mathcal{P}(z) = \sum_{\ell=0}^{N/2} c_\ell z^\ell$.
4. Find the $m \leq N/2$ roots $\{z_j\}_{j=1}^m$ of $\mathcal{P}(z)$ with $|z_j| < 1$. Set $\alpha_j = \log z_j$.
5. Compute the least-squares solution to $V\omega = v$, where $V_{jk} = z_{k+1}^j$, $0 \leq j \leq N$, $0 \leq k \leq m-1$.
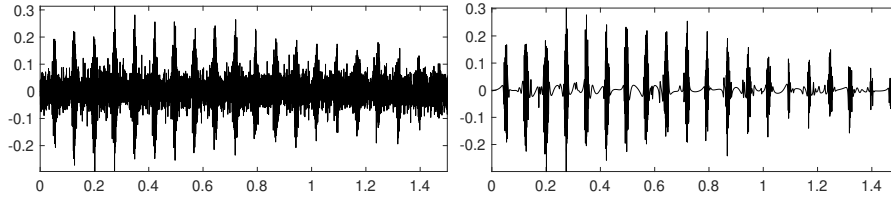

FIG. 3. *Left: A noisy recording of the trill portion of a Pacific blue whale's song. The sample consists of* 6001 *equally spaced observations recorded over* 1.5 *seconds [42]. Right: A type* (274, 275) *trigonometric rational approximation to the signal. The approximant is constructed by applying Algorithm 2.3 directly to the data with the tolerance parameter* $\epsilon = (2 \times 10^{-4})\|H_v\|_2$. *We additionally omit any terms in* R_m *with weights smaller than* ϵ. *The RPM filters out highly oscillatory noise that is not well-captured by trigonometric rationals, making it easier to identify the time-localized pulses in the trill. Once the approximant is constructed, one can toggle between a barycentric representation and the RPM-constructed representation as a sum of complex exponentials to perform various postprocessing tasks (see Table* 2*).*

that takes advantage of fast matrix-vector products for Hankel matrices (e.g., the randomized SVD [31] and Lanczos-based methods [26]).

2.4.2. The regularized Prony method as a filter. In practice, one expects that samples of f are corrupted by noise. The RPM has a natural interpretation as a type of filter. Rather than, for example, filtering out the high-frequency components of a signal, it separates a signal into the sum of two parts by splitting the Hankel matrix H_v into the sum $H_v = H_{R_m} + H_{\mathcal{N}}$. The first term encodes a sequence of coefficients that are well–approximated in Fourier space by a length m sum of exponentials (and thus correspond to a trigonometric rational). The second term encodes a sequence of coefficients that are not well approximated by such an expression. This is referred to as an annihilating filter in the literature on signals with so-called finite rates of innovation [58]. The example in Figure 3 displays noisy data collected by a hydrophone. The noise is not well represented by low to moderate degree trigonometric rationals, so this structured low rank approximation process filters it out.

3. Fourier and inverse Fourier transforms. The RPM and pronyAAA automatically construct compressed representations of f, but these representations are very different from one another. This section describes Fourier/inverse Fourier transforms that allow us to move between these representations. If R_m is a length m sum of exponentials, the existence of a trigonometric rational $r_m = \mathcal{F}^{-1}(R_m)$ is guaranteed by Lemma 2.1. However, the lemma does not reveal if or how r_m can be expressed in barycentric form. In the same way, given a trigonometric barycentric interpolant $r_m^{\gamma, t}$,

it is not clear from (2.5) how one can recover the sum of exponentials $R_m = \mathcal{F}(r_m^{\gamma,t})$. In fact, the exact recovery of one representation from the other is an ill-conditioned problem. With this in mind, we develop lossy but stable transform routines. In the REfit software [59], these transforms are accessed with the commands `ft` and `ift`.

3.1. The forward transform. Given $r_m^{\gamma,t}$ as in (2.5), we seek

$$R_m(k) = \mathcal{F}(r_m^{\gamma,t})(k) = \sum_{j=1}^{m} \omega_j e^{\alpha_j k}, \quad k \geq 0.$$

The parameters of interest can be expressed explicitly in terms of the poles and residues of $r_m^{-\gamma,t}$. For each j, $\alpha_j = 2\pi i \eta_j$ and $\omega_j = e^{-\eta_j} \mathrm{Res}(r_m^{\gamma,t}(z), e^{\eta_j})$, where $z = e^{2\pi i x}$ and $\{\eta_j\}_{j=1}^{m}$ are those poles of $r_m^{\gamma,t}$ with $Im(\eta_k) > 0$ [10]. However, using these formulas requires the accurate computation of the poles η_j, $1 \leq j \leq m$, and their residues. In general, this is an ill-conditioned problem involving extrapolation off the interval of approximation. Trigonometric rationals with different pole configurations can behave almost indistinguishably on $[0, 1)$, so it is often hopeless to recover the pole locations from observations on the interval. Known stability results depend on the poles of $r_m^{\gamma,t}$ being sufficiently well-separated from one another and the residues $\mathrm{Res}(r_m^{\gamma,t}(z), e^{\eta_j})$ being bounded well away from zero [Sec. 2][43]. In our setting we assume that $r_m^{\gamma,t}$ is an approximation to a function f that has algebraic singularities. Good resolution of these features is possible precisely because $r_m^{\gamma,t}$ has poles that cluster up near the singularities (see Figure 1). For this reason, we do not expect that one can compute the pole locations or residues with high accuracy. The exact recovery of the parameters of R_m from samples is a related problem that is also ill-conditioned [49, 54].

Instead of trying to recover R_m exactly, we apply a regularization that finds $\tilde{R}_{\tilde{m}} \approx R_m$, where $\tilde{m} \leq m$.[5] The poles of $r_m^{\gamma,t}$ can be approximately computed by solving a $(2m+1) \times (2m+1)$ generalized eigenvalue problem (see subsection 5.7). Suppose $\{\tilde{\eta}_k\}_{k=1}^{m}$ are the computed poles with $Im(\tilde{\eta}_k) > 0$. We set $\tilde{\alpha}_k = 2\pi i \tilde{\eta}_k$. Then, instead of computing $\{\omega_1, \ldots, \omega_m\}$ using the explicit formula, we find a vector of weights $\tilde{\omega} = (\tilde{\omega}_1, \ldots, \tilde{\omega}_m)^T$ by solving the overdetermined linear system $V_{\tilde{\alpha}} \tilde{\omega} = \hat{r}$, where $\hat{r} = [(\hat{r}_m^{\gamma,t})_0 \cdots (\hat{r}_m^{\gamma,t})_M]^T$ is a vector of Fourier coefficients of $r_m^{\gamma,t}$, and $(V_{\tilde{\alpha}})_{j,k} = e^{\tilde{\alpha}_{k+1} j}$, $0 \leq j \leq M$, $0 \leq k \leq m-1$. All of the Fourier coefficients of $\tilde{r}_{\tilde{m}} = \mathcal{F}(\tilde{R}_{\tilde{m}})$ are exactly produced by a length m sum of exponentials with exponents $\tilde{\alpha}_k = 2\pi i \tilde{\eta}_k$. (Some of the weights in the sum may vanish, and if so we reduce the length of the sum appropriately.) In infinite precision, we would only need $\tilde{m}$ Fourier coefficients of $\tilde{r}_{\tilde{m}}$ to solve for the weights $\tilde{\omega}$ (see Lemma 2.3). Instead, we must fit to the coefficients of the nearby rational $r_m^{\gamma,t}$, and so apply a modest level of oversampling. We then test the accuracy of $\tilde{R}_{\tilde{m}}$ against a randomized sample of the Fourier coefficients of $r_m^{\gamma,t}$, and systematically increase M as needed. It is typically sufficient to choose $M = 2m$.

Since finding $\{e^{\tilde{\alpha}_j}\}_{j=1}^{m}$ and solving $V_{\tilde{\alpha}} \tilde{\omega} = \hat{r}$ are each $\mathcal{O}(m^3)$ operations, the cost for computing $\tilde{R}_{\tilde{m}}$ is dominated by procuring an accurate sample $\hat{r}$. This is done first by evaluating $2N_\epsilon + 1$ samples of $r_m^{\gamma,t}$, where N_ϵ is the ϵ-resolution of $r_m^{\gamma,t}$, on an equally spaced grid, and then applying an FFT. By default, ϵ is taken to be near machine precision, and N_ϵ can be approximately found automatically using, for example, an adaptation of Chebfun's `chop` algorithm [1]. In total, computing $\tilde{R}_m$ from $r_m^{\gamma,t}$ requires $\mathcal{O}(N_\epsilon \log N_\epsilon + N_\epsilon m + m^3)$ operations.

---

[5]If one allows for some of the weights in R_m to be zero, then one can construct sums of exponentials where it is always true that $\tilde{m} = m$.

3.2. The inverse transform. We now assume $R_m(k) = \sum_{j=1}^{m} \omega_j e^{\alpha_j k}$ is given, where each α_j is distinct, $Re(\alpha_j) < 0$, $\omega_j \neq 0$, and $k \geq 0$. We seek an efficient representation for $r_m = \mathcal{F}^{-1}(R_m)$, which is defined to be

$$r_m(x) = \sum_{k=-\infty}^{-1} \overline{R_m(-k)} e^{2\pi i x k} + \sum_{k=0}^{\infty} R_m(k) e^{2\pi i x k}.$$

It is not always true that r_m is a type $(m-1, m)$ trignometric rational function. However, this result does hold under the additional assumption that $R_m(0) = 0$, or, equivalently, $\int_0^1 r_m(x) \mathrm{d}(x) = 0$. We take this to be the case, and our objective is to construct a barycentric interpolant $r_m^{\gamma,t}$ to r_m. We show in the next lemma that for any set of distinct points $t = \{t_1, \ldots, t_{2m}\} \subset [0,1)$, there is γ such that $r_m^{\gamma,t} = r_m$.

LEMMA 3.1. *Let r_m be a type $(m-1, m)$ trigonometric rational function with simple poles that is real-valued, continuous, and periodic on $[0,1)$. Let $t \subset [0,1)$ be a set of $2m$ distinct interpolating points. Then, there is a set of weights γ such that the trigonometric barycentric interpolant $r_m^{\gamma,t}$ recovers r_m exactly.*

Proof. Consider the denominator q_m in $r_m = p_{m-1}/q_m$, and assume that q_m has no shared zeros with p_{m-1}. Since q_m is a trigonometric polynomial, we can write it in barycentric form with respect to the interpolating points in t:

$$(3.1) \qquad q_m(x) = \ell_t(x) \sum_{j=1}^{2m} w_j q_m(t_j) \cot(\pi(x - t_j)), \qquad \ell_t(x) = \prod_{j=1}^{2m} \sin(\pi(x - t_j)),$$

where $w_j = 1/\prod_{k=1, j \neq k}^{2m} \sin(\pi(t_k - t_j))$ are the polynomial barycentric weights [8] associated with t. By setting $\gamma_j = q_m(t_j) w_j$ and $f_j = r_m(t_j)$, we have via (2.5) that there is $r_m^{\gamma,t}(x) = n(x) \ell_t(x)/q_m(x)$ for some function n, where for each j, $r_m^{\gamma,t}(t_j) = r_m(t_j)$. We must now show that $n(x)\ell_t(x) = p_{m-1}(x)$.

The barycentric trigonometric polynomial interpolant to p_{m-1} on t exists and is given by $p_{m-1}(x) = \ell_t(x) \sum_{j=1}^{2m} w_j p_{m-1}(t_j) \cot(\pi(x - t_j))$. Expanding this in exponential form, we have that $p_{m-1}(x) = c_m e^{2\pi i m x} + \ldots + c_{-m} e^{-2\pi i m x}$, where

$$c_m = \frac{1}{4i} \exp\left(-\pi i \sum_{j=1}^{2m} t_j\right) \sum_{j=1}^{2m} w_j p_{m-1}(t_j), \quad c_{-m} = -\exp\left(2\pi i \sum_{j=1}^{2m} t_j\right) c_m.$$

Since p_{m-1} is of degree $m-1$, $c_m = c_{-m} = 0$, so $\sum_{j=1}^{2m} w_j p_{m-1}(t_j) = 0$. This implies that $\sum_{j=1}^{2m} \gamma_j r_m(t_j) = 0$, as $r_m(t_j) = p_{m-1}(t_j)/q_m(t_j)$. Now it is clear that $n(x)\ell_t(x)$ is also a trigonometric polynomial of degree $m-1$. Since $n\ell_t$ interpolates p_{m-1} at $2m$ points, they must agree everywhere. □

The computation of γ as in Lemma 3.1 via polynomial barycentric weights and the evaluation of q_m is numerically unstable except in very special cases [55]. The stable computation of γ and subsequently, the error $\|r_m - r_m^{\gamma,t}\|_\infty$, depends strongly on the choice of t. Some of the more obvious methods for selecting the nodes perform poorly and lead to instabilities in the form of spurious poles. The following discussion is somewhat technical, but it introduces an effective heuristic for choosing a "good" set of barycentric nodes and then stably constructing $r_m^{\gamma,t} \approx r_m$.

3.2.1. A modified pronyAAA for rational recovery. A simple strategy for choosing nodes is to evaluate r_m on a fine enough grid and then apply m steps of

pronyAAA to construct the interpolant $r_m^{\gamma,t}$. This method does not usually exactly recover r_m (see the discussion of exact recovery in subsection 3.1), and it can be the case that the error $\|r_m^{\gamma,t} - r_m\|_\infty$ is unacceptably large. A few additional steps of pronyAAA may drive the error down, though this results in a trigonometric rational interpolant with more poles than r_m. However, a more pernicious problem with this approach is that demanding accuracy close to machine precision from AAA-based methods can result in spurious poles on the interval of approximation that cannot be eliminated without adversely impacting accuracy [45].

To avoid introducing spurious poles, we use the poles of r_m, which are known explicitly from R_m via Lemma 2.1. There is no hope of exactly preserving the poles. However, if m is fixed and $r_m^{\gamma,t}$ is constructed such that it approximately preserves the given poles, then it cannot also admit arbitrary spurious poles. This motivates a three-step procedure for constructing $r_m^{\gamma,t}$ that mixes a pole-preservation strategy involving a type $(m+K-1, m+K)$ trigonometric rational with a data-driven strategy:

(1) A candidate set $\tilde{t}$ of $2m + 2K$ barycentric nodes is chosen, where $K \geq 0$ is an oversampling parameter. Subsets of $\tilde{t}$ admit type $(m-1, m)$ barycentric trigonometric interpolants with poles close to those of r_m.

(2) The interpolant $r_{m+2K}^{\tilde{\gamma},\tilde{t}}$ is constructed via a pole-preserving least-squares fit to samples of r_m (discussed below), so that it has $2m$ poles close those of r_m.

(3) The pronyAAA cleanup procedure (see subsection 2.3.2) is applied to remove the $2K$ poles of $r_{m+2K}^{\tilde{\gamma},\tilde{t}}$ with the smallest residues. This selects t, a set of $2m$ barycentric nodes, from $\tilde{t}$. The barycentric weights $\{\gamma_1, \ldots, \gamma_{2m}\}$ are then computed via (2.6). Note that the poles of $r_m^{\gamma,t}$ must be recomputed.

A version of this method without oversampling (i.e., with $K = 0$, $\tilde{t} = t$, and $\tilde{\gamma} = \gamma$) is useful for motivating how the barycentric nodes in Step (1) are selected. In such a setting, Step (2) simplifies substantially and Step (3) is not needed. However, it is more stable to choose $K > 0$, and in practice, we usually take $K = 1$. We first describe the $K=0$ case, and then use it to explain the method for $K > 0$.

Case 1: $K = 0$. Suppose that T, the discretization of $[0,1)$ from which t is chosen, consists of points $x_0 < x_1 < \ldots < x_{2N}$. Let $P = \{\eta_1, \ldots, \eta_{2m}\}$ be the poles of r_m. Ideally, $r_m^{\gamma,t}$ can be constructed so that its poles are given by P. Noting that the poles of $r_m^{\gamma,t}$ are the zeros of the denominator polynomial $d_m(x) = \sum_{j=1}^{2m} \cot(\pi(x-t_j))$ in (2.5), we introduce the matrix $D_T \in \mathbb{C}^{(2m+1)\times(2N+1)}$:

$$(3.2) \qquad D_T = \begin{bmatrix} \ell_{1,0} & \cdots & \cdots & \ell_{1,2N} \\ \vdots & & & \vdots \\ \ell_{2m,0} & \cdots & \cdots & \ell_{2m,2N} \\ \hline r_m(x_0) & \cdots & \cdots & r_m(x_{2N}) \end{bmatrix}, \qquad \ell_{j,k} = \cot(\pi(\eta_j - x_k)).$$

Using D_T, we relate the selection of barycentric nodes to a column subset selection problem. Indexing from 0, denote by $(D_T)_k$ the kth column of D_T. The kth column is associated with the point x_k in T. The set of nodes $t = \{x_{k_1}, x_{k_2}, \ldots x_{k_{2m}}\}$ then corresponds to a set of columns that forms the submatrix $D_t = [(D_T)_{k_1}, \ldots, (D_T)_{k_{2m}}]$. From Lemma 3.1, there is $\gamma = (\gamma_1, \ldots, \gamma_{2m})^T$ such that $r_m^{\gamma,t} = r_m$, and γ is in the null space of D_t. The first $2m$ entries of $D_t\gamma$ are evaluations of d_m at its zeros. The last entry of $D_t\gamma$ is also zero, since numerator of $r_m^{\gamma,t}$ is of degree $m-1$ (see subsection 2.2). If γ can be computed from D_t accurately, then clearly t is an excellent set of interpolating points. However, the accuracy of this computation depends on properties of D_t. In particular, there are stable ways to compute γ if $2m-1$ of the columns of D_t form a well-conditioned matrix [28].

This suggests that we choose the points t by choosing a subset of columns from D_T that are close to orthogonal. Several kinds of rank-revealing algorithms can be applied to D_T to approximately solve this problem, including the column-pivoted QR (CPQR) algorithm. This constructs the factorization $D_T P = Q_T R$, where P is a permutation matrix, and the leading $\ell \leq \mathrm{rank}(D_T)$ columns of $D_T P$ have been greedily selected to minimize their linear dependence on one another [26, Sec. 6.4]. As a consequence of Lemma 3.1, any submatrix of D_T consisting of $2m$ or more columns is rank-deficient, so $\mathrm{rank}(D_T) \leq 2m-1$.[6] We choose $2m-1$ points in t by performing CPQR on D_T. In principle, the final point in the set t should be chosen so that the accuracy of the computed right singular vector in the nullspace of D_t in (3.2) is maximized. Instead, we choose the point associated with the column in the trailing $(|T|-2m+1)$ columns of $D_T P$ that has the smallest 2-norm. Though they are selected quite differently, these CPQR-selected barycentric nodes concentrate around singularities, just like the nodes selected by pronyAAA (see subsection 3.2.2 and Figure 4).

If one can compute the vectors in the null space of D_t accurately, then γ can be taken as the last right singular vector of D_t. However, this is rarely the case. The accurate recovery of γ from D_t can be problematic even with the best possible choice $t \subset [0,1)$. For this reason, we require a strategy that additionally incorporates a fit to samples. The simplest idea is to use the $2m$ points as selected above and then find the barycentric weights via (2.6). However, this strategy does not seem to eliminate spurious poles or reduce error as effectively as the procedure we describe below.

Case 2: $K \neq 0$. Typically, it is sufficient to take $K = 1$, though one can also choose a larger K. Construct an oversized candidate set $\tilde{t} = \{x_{k_1}, \ldots, x_{k_{2m+2K}}\}$ using the CPQR-based method from Case 1. Then $D_{\tilde{t}}$ is of size $(2m+1) \times (2m+2K)$ and has a numerically detectable null space. We compute the barycentric weights of the interpolant $r_{2m+2K}^{\tilde{t},\tilde{\gamma}}$ by requiring that $\tilde{\gamma} = \tilde{Q}\eta$, where the columns of $\tilde{Q}$ are orthogonal and approximately span the null space of $D_{\tilde{t}}$. We select η to minimize $\|C\tilde{Q}\eta\|_2$, with C constructed as in (2.7). Approximate poles and residues of $r_{2m+2K}^{\tilde{t},\tilde{\gamma}}$ can then be computed in $\mathcal{O}(m^3)$ operations (see subsection 5.7). It is almost always the case in practice that $2K$ poles of $r_{2m+2K}^{\tilde{t},\tilde{\gamma}}$ are negligible in that they have residues with tiny magnitudes. With this in mind, we sort the poles by the magnitude of their residues. As in the pronyAAA cleanup routine, for each of the $2K$ poles with the smallest residues, we eliminate the point in $\tilde{t}$ that is nearest to the pole. The remaining points in $\tilde{t}$ are taken as t, and the set of barycentric weights are found as in a standard step of pronyAAA, i.e., as the minimizer of (2.6).

This strategy first selects a set of interpolating points for which an interpolant with good properties (e.g., poles off $[0,1)$) is known to exist, and then fits the interpolant to samples of r_m. We remark that this is a heuristic. There is no guarantee that spurious poles are avoided or the original poles of r_m are preserved. It remains unclear why the solution in Step (3) often inherits the good pole properties associated with the initial solution in Step (2), and under what circumstances this inheritance can be assured. Nonetheless, we find that the method works extremely well in many cases where simply applying pronyAAA fails.

Implementational details. In practice, we start with $K = 0$. When γ can be recovered with high accuracy directly from D_t, we recover it and end the procedure. This can be checked by computing the singular values of D_t or by using estimates related to the CPQR routine [28]. When this isn't possible, we set $K = 1$ and enlarge our selection of candidate barycentric nodes, which requires no additional

---

[6]If we assume that r_m has a denominator of exactly degree m, then $\mathrm{rank}(D_T) = 2m-1$.
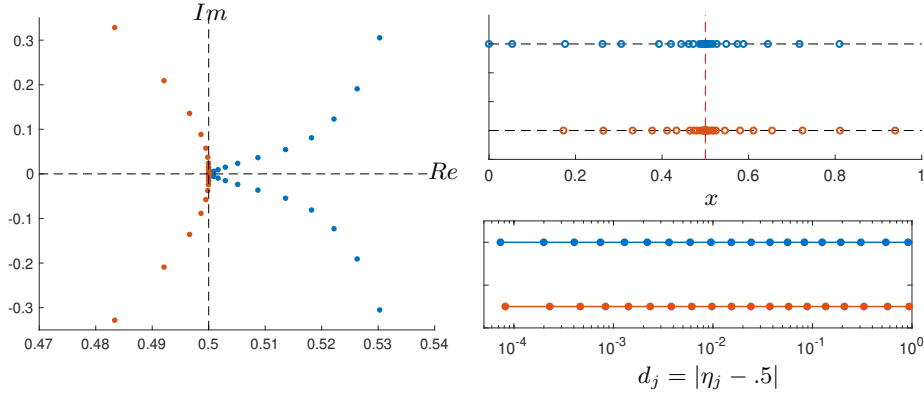
FIG. 4. *Left: poles of r_{bary} (blue) and r_{ift} (orange) are plotted in the complex plane. Here, r_{bary} is a type $(18, 19)$ barycentric interpolant to f constructed via pronyAAA, and $r_{ift} = \mathcal{F}^{-1}(R_{19}) \approx f$, where R_{19} is an exponential sum as in (2.2), and r_{ift} is constructed by applying the inverse Fourier transform from subsection 3.2.1 to R_{19}. The function f is given by $f(x) = |\sin(\pi(x - 1/2))| - \pi/2$, and has a singularity at $x = 1/2$. Right upper: The locations of the barycentric nodes for r_{bary} (blue) and r_{ift} (orange). Right lower: The distances $d_j = |\eta_j - .5|$ from the singularity, where each η_j is a pole with $\text{Im}(\eta_j) > 0$, are sorted by size and plotted on a logarithmic scale (shown in blue for r_{bary}, orange for r_{ift}). This shows that both the pronyAAA-constructed interpolant and the CPQR-based interpolant have poles that exhibit the tapered-type clustering associated with quasi-optimal rational approximation schemes [56].*

computation. Then, we move on to Steps (2) and (3). If the method fails and spurious poles are detected, we first try enlarging $\tilde{t}$ by setting $K = 2$ and trying again. If this fails, it can often be remedied by resampling r_m on a denser grid and starting over at Step (1). When resampling does not solve the issue, we instead construct a stable barycentric interpolant using pronyAAA by accepting a lower level of accuracy.

3.2.2. Example: Two types of barycentric interpolants. In Figure 4, we compare the properties of two types of rational approximations to the function $f(x) = |\sin(\pi(x - 1/2))| - \pi/2$. First, we apply pronyAAA to a set of 6000 samples of f taken on an equally-spaced grid T. This constructs r_{bary}, a type $(18, 19)$ trigonometric rational, where away from the singularity, $|f(x) - r_{bary}(x)| \approx 10^{-8}$. The locations of the barycentric nodes selected by pronyAAA are plotted (blue) in the upper right panel of Figure 4. In the left panel, a subset of the poles of r_{bary} are plotted (blue) in the complex plane. Both the nodes and the poles cluster up near the singularity $x = 1/2$. Shown in red in the same plots are the CPQR-selected barycentric nodes from subsection 3.2.1, and the poles of the barycentric trigonometric rational $r_{ift} = \mathcal{F}^{-1}(R_m)$, where $m = 19$. Here, R_m is an exponential sum constructed via the RPM using samples of f on T, and r_{ift} is constructed using the procedure in subsection 3.2.1. The nodes and poles of r_{ift} also cluster near $x = 1/2$, but in spatial patterns that are quite different from those of r_{bary}. A closer investigation of the pole clustering patterns (see Figure 4, lower right) reveals that in both cases, the sets of distances $d_1 \leq d_2 \leq \ldots \leq d_{19}$ from the singularity, where $d_j = |\eta_j - 1/2|$ and each η_j is a pole with $\text{Im}(\eta_j) > 0$, have the tapered-type spacing on a logarithmic scale that is associated with best (and near-best) convergence rates [56]. In addition to revealing that these functions have properties associated with quasi-optimal approximation power, the locations and clustering patterns of the nodes or poles can be used to identify singularities, extract features, or classify signals.
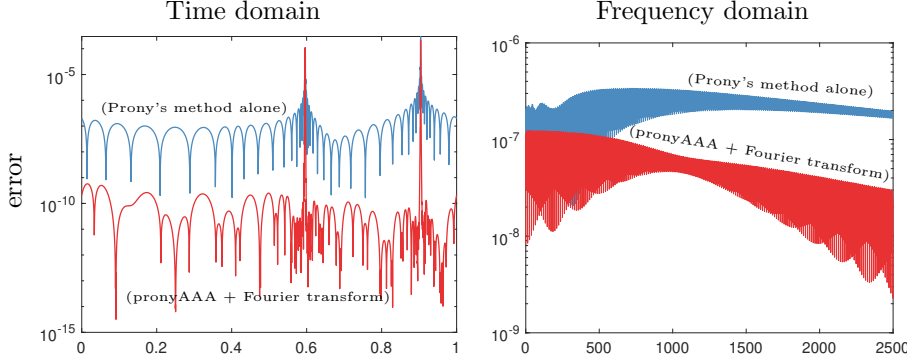
FIG. 5. *Left: The absolute error in approximating $f(x) = |\sin(2\pi x)/4 + \exp(\sin(2\pi x))|/4 + c$ with two different rational approximants, $\mathcal{F}(R_b)$ (blue) and $\mathcal{F}(R_r)$ (red), is plotted on a logarithmic scale at 3000 equally-spaced points: R_b is an exponential sum of length 14 that was adaptively constructed via the RPM (see Algorithm 2.3) from a sample consisting of only 1401 equally-spaced points. The tolerance parameter is set to $\epsilon = 10^{-10}$, but the coarseness of the sample limits the achievable accuracy of the representation. R_r (red) is constructed by first applying pronyAAA in signal space to construct a barycentric interpolant $r_r^{\gamma,t}$, and then using the Fourier transform function to compute $R_r = \mathcal{F}(r_r^{\gamma,t})$. Right: The absolute errors in Fourier space between accurately computed Fourier coefficients of f and the exponential sums R_b (blue) and R_r (red) are plotted on a logarithmic scale against the Fourier modes $0, 1, \ldots, 2500$.*

4. Signal reconstruction in time and frequency space.

With the Fourier and inverse Fourier transforms available, we can combine the advantages of pronyAAA and the RPM to overcome various issues, such as undersampling or noise. We illustrate this idea with two examples. Then in section 5, we describe a collection of algorithms for computing with trigonometric rational functions and exponential sums that exploits our ability to move stably between the representations.

4.1. An undersampled function.

In this example, we consider a function $f(x) = |\sin(2\pi x)/4 + \exp(\sin(2\pi x))|/4 + c$, which has Fourier coefficients that decay asymptotically like $\mathcal{O}(|k|^{-2})$, where k denotes the kth Fourier mode. Here, c is a normalization parameter ensuring that the mean value of f over $[0, 1)$ is zero. We suppose that f is sampled at 1401 equally spaced points, and that an exponential sum representing $\mathcal{F}(f)$ is desirable for downstream tasks. The direct application of Prony's method performs poorly because f is undersampled. We denote the constructed exponential sum as R_b. The error in the computation of the Fourier coefficients via the FFT is on the order of 10^{-6}, so we cannot expect accuracy much better than that. An alternative approach is to apply pronyAAA to construct the barycentric interpolant $r_r^{\gamma,t}$, and then compute $R_r = \mathcal{F}(r_r^{\gamma,t})$ as in subsection 3.1.

In Figure 5 (left), we use the pole-residue format to directly evaluate the values of the rationals associated with the two types of constructed exponential sums. There is a tiny band around the two singularities where the errors incurred by the two methods are approximately the same. Elsewhere, the accuracy achieved by first applying pronyAAA is nearly double that attained by Prony's method alone. The error in recovering the Fourier coefficients of f is diffuse and more accurate, especially in the extrapolation of the tail (see Figure 5, right). The exponential sum R_r, with only 13 terms, represents f with highly localized error behavior, and it is in a form efficient for storage, convolution, and other tasks (see section 5).

4.2. Reconstruction of an ECG signal.

Rational approximation methods are effective in many biomedical monitoring tasks, including the processing of elec-

17

trocardiogram (ECG) signals [23, 25]. In [23], rational functions constructed in the orthogonal rational Malmquist–Takenaka basis are used to reconstruct ECG signals and then classify them. The rationals perform with better overall compression properties and have several other advantages when compared to wavelets, splines, and other families of functions [40]. We do not expect to outperform such a highly specialized scheme with our approach. However, this example illustrates that our more general-purpose method effectively constructs a denoised representation of the signal.

In this example, we apply the RPM and fit a rational function directly to noisy ECG data taken from the PhysioNet MIT BIH arrhythmia database [44]. As in [23], the location of its poles can be used for classification and feature recognition tasks. Using the inverse Fourier transform function described in subsection 3.2.1, we can construct a barycentric trigonometric rational representation of the function, which is a convenient format for identifying local extrema (see section 5). This can all be done with three lines of code in REfit:

```
R = efun(data, 'tol', 1e-3);
r = ift(R);
extrema = [ min(r); max(r)];
```

If one tries to use pronyAAA directly, the result is a trigonometric rational with 200 poles, and the data set only contains 645 samples. Of these poles, 62 are spurious and lie on the interval of approximation. This happens because the pronyAAA algorithm does not distinguish between the signal and the noise, and it tries to induce a fit to noise by adding poles. A better approach is first to apply the RPM. Within the first two lines of the above code, several tasks are being executed: First, the exponential sum R_m (here, $m = 35$) stored in R is constructed via the RPM. The RPM automatically filters out additive noise on the sample with magnitudes approximately at or below the tolerance level $\epsilon = 10^{-3}$. Then, R is used to extrapolate high-frequency information that lies beyond the noise limitation (see Figure 6, left). This provides an enriched sample for selecting interpolating points and constructing the barycentric interpolant r. The construction of r in this way can be viewed as a form of super-resolution [12]. If one tries to construct a barycentric interpolant without enriching the sample by extrapolation in frequency space, then spurious poles appear that one cannot eliminate without destroying the accuracy of the approximation. This is because the signal is not well-resolved in the time domain at the original sample rate. Once r is available, one can then automatically and efficiently perform a variety of processing tasks, such as rootfinding and the detection of maxima and minima.

5. Algorithms for computing with rationals and exponential sums. This section gives an overview of several of the algorithms used in our software [59] to compute with trigonometric rationals and exponential sums. The Fourier and inverse Fourier transform functions are used to move between representations as needed. For operations on trigonometric rationals that return trigonometric rationals, we represent the resulting function using exponential sums or barycentric forms when possible.

5.1. Compression for suboptimal sums of exponentials. Exponential sums are closed under addition and multiplication, but a sum R_m resulting from the naive application of these and other operations is often suboptimal in the sense that a shorter sum $\tilde{R}_{\tilde{m}}$ exists, where $|R_m(j) - \tilde{R}_{\tilde{m}}(j)| < \epsilon$ for $j \geq 0$. A major advantage of the exponential sum format is that $\tilde{R}_{\tilde{m}}$ can be constructed with a cost that usually depends on m, rather than the ϵ-resolution parameter N_ϵ associated with R_m. Using
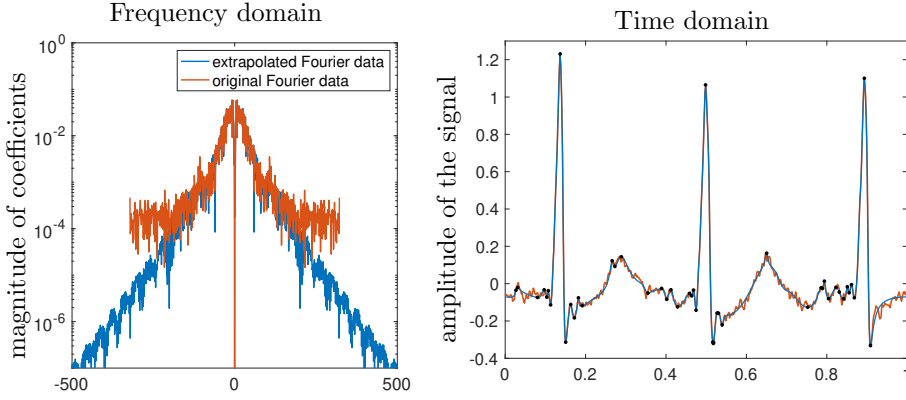
FIG. 6. *Left: The magnitude of the Fourier coefficients of the original signal (orange). The decay of the coefficients stagnates due to noise in the signal, and this pollutes the higher frequencies. Once an exponential sum representation R_m is constructed, we can extrapolate to higher frequencies by evaluating R_m, and thereby super–resolve the signal (blue). Right: A barycentric rational approximant (blue) in the time domain is computed using the extrapolated Fourier data. It is a denoised version of the original ECG signal (orange). Local extrema are identified (black dots) using the differentiation and rootfinding algorithms from section 5.*

AAK theory for finite rank Hankel operators, one can show (see [47, Thm. 3.2]) that there is a length $\tilde{m} \leq m$ approximation that satisfies

$$\|\mathcal{F}^{-1}(R_m) - \mathcal{F}^{-1}(\tilde{R}_{\tilde{m}})\|_{L_2} \leq 2\, \sigma_{\tilde{m}}(\Gamma_R),$$

where Γ_R is the infinite matrix with entries $(\Gamma_R)_{j+k} = R_m(j+k)$, $j,k \geq 0$. In [47], an $\mathcal{O}(m^3)$ algorithm for recovering $\tilde{R}_{\tilde{m}}$ directly from the parameters of R_m is developed; a closely-related approach using only properties of finite-dimensional Hankel matrices is described in [9]. This method is successfully employed in [33] within a scheme that uses rational approximations to solve Burger's equation. However, the implementation requires the judicious use of high-precision arithmetic, which we wish to avoid.

Instead, we note that when a length $\tilde{m} < m$ recurrence is approximately satisfied by the sequence $\{R_m(0), R_m(1), \ldots\}$, this fact is often captured well with a modified Prony's method that involves only a small sample of $2M + 1$ observations of R_m, where $M > m$. Specifically, we construct a small $(M+1) \times (m+1)$ rectangular Hankel matrix H with entries $H_{jk} = R_m(j+k)$. Then, we apply the RPM from Algorithm 2.3 on H to construct $\tilde{R}_{\tilde{m}}$. We check the error $|R_m(j) - \tilde{R}_{\tilde{m}}(j)|$ on a random sample of integers $0 \leq j \leq N$, where N is the ϵ-resolution parameter used in the original construction of R_m. When the error is too large, we increase M and try again. The cost to compute $\tilde{R}_{\tilde{m}}$ is $\mathcal{O}(Mm^2)$. In a worst-case scenario, M can grow as large as N. We observe experimentally that this approach is often very effective, but more work is needed to understand the conditions under which it is guaranteed that $M \ll N$.

5.2. Sums of trigonometric rationals. If S_ℓ and G_n are exponential sums, then $R_m = S_\ell + G_n$ can be constructed straightforwardly. However, R_m may be of suboptimal length. We apply the compression algorithm with $\epsilon \approx \epsilon_{mach}$ to R_m, where $m = \ell + n$, to find $\tilde{R}_{\tilde{m}}$. This "compression–plus" method is especially useful for tasks that involve repeated summations and require many recompressions. The compression–plus algorithm is applied in REfit when the '+' operator is used between efun objects. For summing trigonometric rationals s_ℓ and g_n represented as rfuns, we simply evaluate the sum and then apply pronyAAA to find
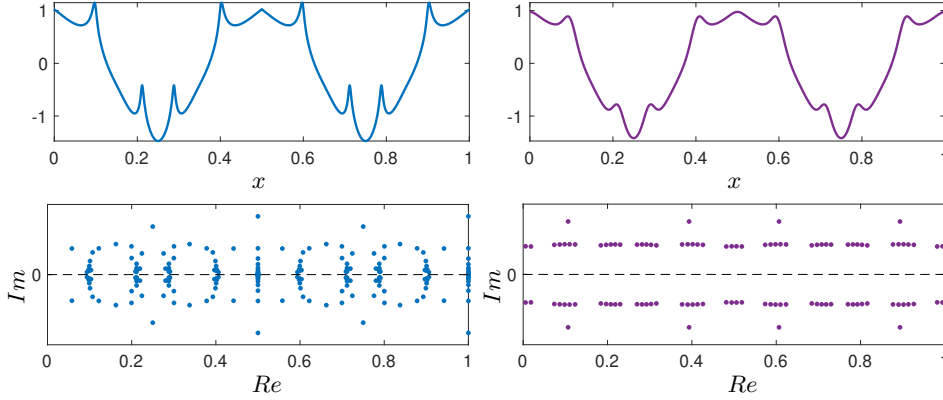
19

FIG. 7. *Left (top): an exponential sum S_ℓ representing a function f with singularities (inspired by Chebfun's "wild" function [19]) is constructed as an efun, and $s_\ell = \mathcal{F}^{-1}(S_\ell) \approx f$ is plotted in the time domain. Here, $\ell = 116$ and $\|s_\ell - f\|_\infty \approx 10^{-9}\|f\|_\infty$. Left (bottom): Locations of poles of s_ℓ occurring in a small strip of width .2 around the real line are plotted in the complex plane. Right (top): The convolution $r_m = s_\ell * g_n$ is computed ($m = 53$), stored as an efun, and plotted in the time domain. Here, g_n ($n = 13$) is a trigonometric rational approximation to a narrow Gaussian. Right (bottom): The poles of r_m still cluster near the singularities of f, but less tightly, reflecting the fact that $f * g_n$ is a smoothed version of f.*

$r_m = s_\ell + g_n$.[7] The rfun and efun objects can also be combined in various ways. The syntax `r = s + g` adds two rfuns and returns an rfun by default. The expression `[r, R] = s + g` automatically retrieves the efun `R = ft(r)` in addition to `r`. When an rfun and efun are summed together, both rfun and efun outputs are returned.

5.3. Convolutions of trigonometric rationals. The convolution of trigonometric rationals s_ℓ and g_n (denoted $s_\ell * g_n$) can be constructed in Fourier space by finding the exponential sum $R_{m_0} = (\mathcal{F}s_\ell)(\mathcal{F}g_n) = S_\ell G_n$. The product can be computed directly in a closed form, but this results in a large sum with $m_0 = \ell n$. To find a shorter sum $\tilde{R}_{\tilde{m}}$, we first find an upper bound m on $\tilde{m}$ by determining how many terms in R_{m_0} have a negligibly small influence. We also check the decay in the tail of the first m_0 Fourier coefficients given by $S_\ell G_n$ to determine if even fewer samples are needed. Then, we apply the compression algorithm using rectangular Hankel matrices of the form $H_{jk} = S_\ell(j + k)G_n(j + k)$. For efuns, this operation is accessed by typing `S.*G`. For rfuns, the command `conv(s,g)` uses the Fourier and inverse Fourier transform functions to apply the above scheme. Figure 7 displays an example where s_ℓ ($\ell = 116$) is an approximation to a function f with several singularities, and g_n ($n = 13$) is an approximation to the normalized Gaussian $g(x) = \frac{1}{\sqrt{2\pi}\sigma}e^{-x^2/2\sigma^2} - 1$, with $\sigma = .01$. Our convolution algorithm constructs $r_m \approx f * g$, a smoothed version of f. The method automatically chooses $m = 53$, about half the degrees of freedom required for approximating f with s_ℓ. Little accuracy is lost in this process: $\|R_m - S_\ell G_n\|_\infty \approx 10^{-11}\|S_\ell G_n\|_\infty$.

5.4. Products of trigonometric rationals. The product of two trigonometric rationals r_ℓ and s_n in the time domain is equivalent to their convolution in Fourier

---

[7]If this proves difficult due to spurious poles, we use the Fourier and inverse Fourier transforms to convert to efuns, perform the addition, and then convert back to an rfun.

space. If $R_\ell = \mathcal{F}(r_\ell)$ and $S_n = \mathcal{F}(s_n)$ are each sums of complex exponentials, then

$$(5.1) \quad \mathcal{F}(r_\ell s_n)(k) = (R_\ell * S_n)(k) = \sum_{j=-\infty}^{\infty} R_\ell(k-j)S_n(j) \approx \sum_{j=-N_\epsilon}^{N_\epsilon} R_\ell(k-j)S_n(j),$$

where N_ϵ is the ϵ-resolution parameter for S_n. The evaluation of (5.1) at M consecutive points requires a matrix-vector multiply with a Toeplitz matrix. Since $r_\ell s_m$ is a trigonometric rational with at most $2\ell+2n$ poles, we apply the compression algorithm to find $G_m \approx \mathcal{F}(r_\ell s_n)(k)$, with $m = \ell + n$. If R and S are efuns, this command is accessed by typing `conv(R,S)`. If r_ℓ and s_n are represented with rfuns r and s, respectively, then `r.*s` simply applies pronyAAA to the function $r_\ell(x)s_n(x)$ to construct a new rfun representing $r_\ell s_n$.[8]

5.5. Differentiation. The kth derivative of r_m, denoted $r_m^{(k)}$, is a trigonometric rational of type $(km-1, km)$. However, $r_m^{(k)}$ is fundamentally of a different form than the trigonometric rationals constructed via pronyAAA and the RPM. It has m conjugate pairs of poles, and each pole is of multiplicity k. It is possible to represent derivatives with trigonometric rationals with simple poles, but it isn't always sensible. By default, REfit returns a function handle for evaluating derivatives (or their Fourier transforms) whenever $\texttt{diff}(\cdot,\texttt{k})$ is applied to an rfun (or an efun). However, one can also use $\texttt{diff}(\cdot,\texttt{k},\texttt{'type'})$, to specify that an efun or rfun should be returned.

Differentiation in Fourier space. When r_m is represented by the complex exponential sum R_m in Fourier space, the Fourier coefficients of $r_m^{(k)}$ are given by $\mathcal{F}(r_m^{(k)})(j) = (2\pi i j)^k R_m(j)$. The command $\texttt{h=diff(R,k)}$ by default returns a handle for evaluating this function in Fourier space. If instead, for example, one types $\texttt{diff(R,k,'efun')}$, the RPM is applied to construct a representation of $\mathcal{F}(r_m^{(k)})$ as a sum of weighted complex exponentials (without polynomial coefficients).

Differentiation in the time domain. Derivatives of barycentric trignometric rational interpolants satisfy a recurrence relation and can be expressed in a simple closed form. To see this, consider the linearization of $r_m^{\gamma,t} = n_{m-1}/d_m$, which can be differentiated as $(r_m^{\gamma,t} d_m)' = (n_{m-1})'$. Plugging in the definitions from (2.5) results in the following formula, which holds everywhere on $[0, 1)$ except at the interpolating points:

$$(5.2) \qquad (r_m^{\gamma,t})'(x) = -\pi \frac{\sum_{j=1}^{2m} \gamma_j \csc^2(\pi x - \pi t_j)(f_j - r_m^{\gamma,t}(x))}{\sum_{j=1}^{2m} \gamma_j \cot(\pi x - \pi t_j)}.$$

To evaluate $(r_m^{\gamma,t})'$ at the interpolating points $t = (t_1, \ldots, t_{2m})^T$, we use the special differentiation matrices introduced in [4]. Explicit descriptions of recursive formulas for computing higher derivatives are also found in [4]. All of this is encoded within a function handle that is accessed in REfit by applying the command `diff` to an rfun.

5.6. Integration. The indefinite integral of a trigonometric rational r_m is not itself a trigonometric rational. In Fourier space, if $R_m = \mathcal{F}(r_m)$ is a sum of complex exponentials as in (2.2), then except at $k = 0$, the Fourier coefficients of $\mathcal{F}(g)$, where $g(y) = \int_0^y r_m(x)dx$, are given by $\hat{g}_k = R_m(k)/2\pi i k$. A function handle for evaluating $\mathcal{F}(g)$ is returned when `cumsum` is applied to an efun. One can also fit a new complex exponential to $\mathcal{F}(g)$ by typing $\texttt{cumsum}(\cdot,\texttt{'efun'})$, though it may not be an efficient representation.

---

[8]If this proves difficult due to spurious poles, we use the Fourier and inverse Fourier transforms to convert to efuns, perform the convolution in Fourier space, and then convert back to an rfun.

If `cumsum` is applied to an rfun, we supply a handle for g that applies Gauss-Legendre quadrature [30]. The stable evaluation property of the barycentric form is advantageous here, which is why we do not instead make use of the pole-residue form of the rational $r_m(z)$ in (2.3). To integrate r_m over a finite interval $[a, b] \subset [0, 1)$, the command `sum(·,a,b)` can be applied to an rfun or an efun.

5.7. Rootfinding and polefinding. The roots of the barycentric trigonometric rational $r_m^{\gamma,t}$ coincide with the generalized eigenvalues of a matrix pencil. Specifically, if $r_m^{\gamma,t}(\zeta_j) = 0$ and $\mu = e^{2\pi i \zeta_j}$, then there is nonzero y such that $Ey = \mu By$, where

$$(5.3) \quad E = \begin{bmatrix} e^{2\pi i x_1} & & & i\omega_1 e^{2\pi i x_1} \\ & \ddots & & \vdots \\ & & e^{2\pi i x_{2m}} & i\omega_{2m} e^{2\pi i x_{2m}} \\ \hline f_1 & \cdots & f_{2m} & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & & & i\omega_1 \\ & \ddots & & \vdots \\ & & 1 & i\omega_{2m} \\ \hline 0 & \cdots & 0 & 0 \end{bmatrix}.$$

The pencil (E, B) has at least two infinite eigenvalues and one eigenvalue at $\mu_0 = 0$ corresponding to $\zeta_0 = -\infty$ (this captures the asymptotic behavior of $r_m^{\gamma,t}$). Once the remaining $2m-2$ eigenvalues are found, the zeros of $r_m^{\gamma,t}$ are immediate. The command `roots` applied to rfuns or efuns applies this algorithm and returns real-valued roots. For efuns, this requires first converting to an rfun via the Fourier transform. The command `roots(·,'all')` additionally returns complex-valued roots.

The poles of $r_m^{\gamma,t}$ can be found in a similar way: the pencil $(\tilde{E}, B)$, where $\tilde{E}$ is identical to E except that each f_j in the last row is replaced by 1, has at least one infinite eigenvalue. If $\tilde{\mu}_j$ is one of the remaining finite eigenvalues, then $\eta_j = \log \tilde{\mu}_j / 2\pi i$ is a pole of $r_m^{\gamma,t}$. If it is important to preserve the pole symmetry, it is better to represent r_m with an exponential sum R_m.

Residues. We compute the residues of the barycentric interpolant $r_m^{\gamma,t}$ using the fact that $r_m^{\gamma,t} = n_{m-1}/d_m$, where n_{m-1} and d_m are trigonometric polynomials as in (2.5). Since the poles of $r_m^{\gamma,t}$ are simple, the residue at the pole η_j is given by $\text{Res}(r_m^{\gamma,t}, \eta_j) = n_{m-1}(\eta_j)/d_m'(\eta_j)$. The residues of the poles of $\mathcal{F}^{-1}(R_m)$, where R_m is an exponential sum, have a closed form formula (see subsection 3.1).

Minima and Maxima. The commands `min` or `max` a return the local minima (or maxima) attained by the represented trigonometric rational on the interval $[0, 1)$. The global minimum, for example, can be found by typing `min(min(·))`. To compute the extrema, we use a rfun and apply the differentiation formula in (5.2) to evaluate its derivative. We use this to construct a rfun representing $(r_m^{\gamma,t})'$, find its roots, and then test for concavity. For an application, see subsection 4.2.

5.8. Additional commands. The REfit package includes commands for data visualization and common tasks in signal processing, such as filtering and cross-correlation. Commands related to the pole–residue format of the rational $\tilde{r}_m(z)$ from (2.3) are also available, as this format is closely related to the notion of the z-transform [46] and is useful for interpretation and analysis.

5.9. Nonperiodic representations. It is natural to consider the extension of these ideas and algorithms to non-periodic signals. In signal space, a slight modification of the standard AAA algorithm can be used to construct type $(k-1, k)$ barycentric rational interpolants with basis functions as in (2.4). A type $(k-1, k)$ rational function with simple poles that have nonzero imaginary parts can be associated with a Laurent expansion, and the Laurent coefficients can be expressed in terms of two short sums of complex exponentials [53, Ch. 4]. The method in subsection 3.1 can be used to

find the parameters of the exponential sums, though one requires a scheme for computing the Laurent coefficients accurately from data samples. Modifications of the algorithms in section 5 can then be applied to compute with Laurent coefficients and the associated rational barycentric interpolants.

A general approach favored by many practitioners is to apply windowing strategies [24, 27, 32] to compute the Fourier transform of the non-periodic signal, or to use Fourier extension methods [11]. If the signal is first represented by a barycentric rational interpolant using AAA, then the pole or barycentric node locations can be used to adaptively select parameters for determining windows and designing quadrature schemes. We hope in future work to develop these observations into strategies for computing with time-frequency representations of general rational functions.

6. Conclusion. We have introduced a framework for signal reconstruction and automated computing that employs efficient representations in both time and frequency space. Our work integrates ideas from the harmonic analysis community involving exponential sums and Hankel operator theory [9, 47] with developments in adaptive barycentric rational interpolation [6, 35, 45]. An implementation of all of the described methods is publicly available in the REfit software package [59].

REFERENCES

[1] J. L. AURENTZ AND L. N. TREFETHEN, *Chopping a Chebyshev series*, acmtoms, 43 (2017), pp. 1–21.

[2] A. P. AUSTIN AND K. XU, *On the numerical stability of the second barycentric formula for trigonometric interpolation in shifted equispaced points*, IMA J. Numer. Anal., 37 (2017), pp. 1355–1374.

[3] P. J. BADDOO, *The AAAtrig algorithm for rational approximation of periodic functions*, SIAM J. Sci. Comput., 43 (2021), pp. A3372–A3392.

[4] R. BALTENSPERGER, *Some results on linear rational trigonometric interpolation*, Comp. Math. Appl., 43 (2002), pp. 737–746.

[5] M. BERLJAFA AND S. GÜTTEL, *The RKFIT algorithm for nonlinear rational approximation*, SIAM J. Sci. Comput., 39 (2017), pp. A2049–A2071.

[6] J. BERRUT, *Rational functions for guaranteed and experimentally well-conditioned global interpolation*, Comp. Math. Appl., 15 (1988), pp. 1–16.

[7] J.-P. BERRUT, R. BALTENSPERGER, AND H. D. MITTELMANN, *Recent developments in barycentric rational interpolation*, in Trends and Applications in Constructive Approximation, Springer, 2005, pp. 27–51.

[8] J.-P. BERRUT AND L. N. TREFETHEN, *Barycentric Lagrange interpolation*, SIAM Rev., 46 (2004), pp. 501–517.

[9] G. BEYLKIN AND L. MONZÓN, *On approximation of functions by exponential sums*, Appl. and Comp. Harmonic Analysis, 19 (2005), pp. 17–48.

[10] ——, *Nonlinear inversion of a band-limited Fourier transform*, Appl. and Comp. Harmonic Analysis, 27 (2009), pp. 351–366.

[11] O. P. BRUNO, *Fast, high-order, high-frequency integral methods for computational acoustics and electromagnetics*, in Topics in computational wave propagation, Springer, 2003, pp. 43–82.

[12] E. J. CANDÈS AND C. FERNANDEZ-GRANDA, *Towards a mathematical theory of super-resolution*, Comm. Pure Appl. Math., 67 (2014), pp. 906–956.

[13] V. ČERVENÝ, M. M. POPOV, AND I. PŠENČÍK, *Computation of wave fields in inhomogeneous media—Gaussian beam approach*, Geophys. J. Int., 70 (1982), pp. 109–128.

[14] A. CICHOWICZ, *An automatic s-phase picker*, Bulletin of the Seismological Society of America, 83 (1993), pp. 180–189.

[15] S. Costa and L. N. Trefethen, *AAA-least squares rational approximation and solution of Laplace problems*, arXiv preprint arXiv:2107.01574, (2021).

[16] C. De Boor, *On calculating with B-splines*, J. Approx. Theory, 6 (1972), pp. 50–62.

[17] L. Debnath, *Wavelets and signal processing*, Springer Science & Business Media, 2003.

[18] N. Derevianko and G. Plonka, *Exact reconstruction of extended exponential sums using rational approximation of their Fourier coefficients*, arXiv preprint arXiv:2103.07743, (2021).

[19] T. A. Driscoll, N. Hale, and L. N. Trefethen, eds., *Chebfun Guide*, Pafnuty Publications, Oxford, 2014.

[20] S. Elsworth and S. Güttel, *Conversions between barycentric, RKFUN, and Newton representations of rational interpolants*, Lin. Alg. Appl., 576 (2019), pp. 246–257.

[21] S.-I. Filip, Y. Nakatsukasa, L. N. Trefethen, and B. Beckermann, *Rational minimax approximation via adaptive barycentric representations*, SIAM J. Sci. Comput., 40 (2018), pp. A2427–A2455.

[22] B. Francis, *A course in H-infinity control theory*, Lecture notes in Control and Information Sciences, 88 (1987), p. R5.

[23] S. Fridli, L. Lócsi, and F. Schipp, *Rational function systems in ECG processing*, in Int. Conf. Comp. Aided Sys. Theory, Springer, 2011, pp. 88–95.

[24] D. Gabor, *Theory of communication. part 1: The analysis of information*, J. Inst. of Electrical Engineers-Part III: Radio and Comm. Engineering, 93 (1946), pp. 429–441.

[25] Z. Gilián, *ECG-based heart beat detection using rational functions*, in Conf. Dyadic Anal. Rel. Fields Appl., vol. 13, 2014.

[26] G. H. Golub and C. F. Van Loan, *Matrix Computations*, vol. 3, John Hopkins University Press, 2012.

[27] D. Gottlieb and C.-W. Shu, *On the Gibbs phenomenon and its resolution*, SIAM Rev., 39 (1997), pp. 644–668.

[28] M. Gu and S. C. Eisenstat, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM J. Sci. Comput., 17 (1996), pp. 848–869.

[29] B. Gustavsen and A. Semlyen, *Rational approximation of frequency domain responses by vector fitting*, IEEE Trans. Power Deliv., 14 (1999), pp. 1052–1061.

[30] N. Hale and A. Townsend, *Fast and accurate computation of Gauss–Legendre and Gauss–Jacobi quadrature nodes and weights*, SIAM J. Sci. Comput., 35 (2013), pp. A652–A674.

[31] N. Halko, P.-G. Martinsson, and J. A. Tropp, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.

[32] F. J. Harris, *On the use of windows for harmonic analysis with the discrete Fourier transform*, Proc. IEEE, 66 (1978), pp. 51–83.

[33] T. Haut and G. Beylkin, *Fast and accurate con-eigenvalue algorithm for optimal rational approximations*, SIAM J. Matrix Anal. Appl., 33 (2012), pp. 1101–1125.

[34] T. Haut, G. Beylkin, and L. Monzón, *Solving Burgers' equation using optimal rational approximations*, Appl. Comp. Harm. Appl., 34 (2013), pp. 83–95.

[35] P. Henrici, *Barycentric formulas for interpolating trigonometric polynomials and their conjugates*, Num. Math., 33 (1979), pp. 225–234.

[36] N. J. Higham, *The numerical stability of barycentric Lagrange interpolation*, IMA J. Numer. Anal., 24 (2004), pp. 547–556.

[37] M. Javed, *Algorithms for trigonometric polynomial and rational approximation*, PhD thesis, University of Oxford, 2016.

[38] D. Karachalios, I. V. Gosea, and A. C. Antoulas, *The Loewner framework for system identification and reduction*, in Model Reduction Handbook: Volume I: System-and Data-Driven Methods and Algorithms, De Gruyter, 2021, pp. 181–228.

[39] S. M. Kay, *Fundamentals of statistical signal processing*, Prentice Hall PTR, 1993.

[40] P. Kovács, S. Fridli, and F. Schipp, *Generalized rational variable projection with application in ECG compression*, IEEE Trans. Sig. Proc., 68 (2019), pp. 478–492.

[41] E. R. Martin, *A linear algorithm for ambient seismic noise double beamforming without explicit crosscorrelations*, Geophysics, 86 (2021), pp. F1–F8.

[42] Mathworks, *MATLAB Signal Processing Toolbox: (v R2020b)*, MathWorks, 2020.

[43] K. Miller, *Stabilized numerical analytic prolongation with poles*, SIAM J. Appl. Math., 18 (1970), pp. 346–363.

[44] G. Moody, R. Mark, and A. Goldberger, *Physionet: A research resource for studies of complex physiologic and biomedical signals*, in Computers in Cardiology 2000. Vol. 27 (Cat. 00CH37163), IEEE, 2000, pp. 179–182.

[45] Y. Nakatsukasa, O. Sète, and L. Trefethen, *The AAA algorithm for rational approximation*, SIAM J. Sci. Comput., 40 (2018), pp. A1494–A1522.

[46] A. V. Oppenheim, J. R. Buck, and R. W. Schafer, *Discrete-time signal processing. Vol. 2*, Upper Saddle River, NJ: Prentice Hall, 2001.

[47] V. Pototskaia and G. Plonka, *Application of the AAK theory and Prony-like methods for sparse approximation of exponential sums*, Proc. Appl. Math. Mech., 17 (2017), pp. 835–836.

[48] D. Potts and M. Tasche, *Nonlinear approximation by sums of nonincreasing exponentials*, Applicable Anal., 90 (2011), pp. 609–626.

[49] ———, *Parameter estimation for nonincreasing exponential sums by Prony-like methods*, Lin. Alg. Appl., 439 (2013), pp. 1024–1039.

[50] G. Prony, *Essai experimental et analytique*, J. de l'Ecole Polytechnique, 2 (1795).

[51] N. Ricker, *Transient waves in visco-elastic media*, vol. 10, Elsevier, 2012.

[52] S. Rippa, *An algorithm for selecting a good value for the parameter c in radial basis function interpolation*, Adv. in Comp. Math., 11 (1999), pp. 193–210.

[53] E. M. Stein and R. Shakarchi, *Complex analysis*, vol. 2, Princeton University Press, 2010.

[54] M. K. Transtrum, B. B. Machta, and J. P. Sethna, *Why are nonlinear fits to data so challenging?*, Phys. Rev. Letters, 104 (2010), p. 060201.

[55] L. N. Trefethen, *Approximation Theory and Approximation Practice*, SIAM, 2013.

[56] L. N. Trefethen, Y. Nakatsukasa, and J. Weideman, *Exponential node clustering at singularities for rational approximation, quadrature, and PDEs*, Numer. Math., 147 (2021), pp. 227–254.

[57] M. A. Unser and T. Blu, *Wavelets and radial basis functions: A unifying perspective*, in Wavelet Appl. Sig. Im. Proc. VIII, vol. 4119, Int. Soc. Opt. Phot., 2000, pp. 487–493.

[58] M. Vetterli, P. Marziliano, and T. Blu, *Sampling signals with finite rate of innovation*, IEEE Trans. Sig. Proc., 50 (2002), pp. 1417–1428.

[59] H. Wilber, *https://github.com/heatherw3521/refit*, 2021.

[60] G. B. Wright, M. Javed, H. Montanelli, and L. N. Trefethen, *Extension of Chebfun to periodic functions*, SIAM J. Sci. Comput., 37 (2015), pp. C554–C573.

[61] K. Xu and S. Jiang, *A bootstrap method for sum-of-poles approximations*, J. of Sci. Comp., 55 (2013), pp. 16–39.