

Analytic continuation by averaging Padé approximants

Johan Schött,¹ Inka L. M. Loch,^{1,2} Elin Lundin,¹ Oscar Grånäs,^{1,3} Olle Eriksson,¹ and Igor Di Marco¹

¹*Dept. of Physics and Astronomy, Uppsala University, Box 516, SE-75120 Uppsala, Sweden*

²*Institute of Molecules and Materials, Radboud University of Nijmegen, Heyendaalseweg 135, 6525 AJ Nijmegen, The Netherlands*

³*School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*
(Dated: November 13, 2015)

The ill-posed analytic continuation problem for Green's functions and self-energies is investigated by revisiting the Padé approximants technique. We propose to remedy to the well-known problems of the Padé approximants by performing an average of several continuations, obtained by varying the number of fitted input points and Padé coefficients independently. The suggested approach is then applied to several test cases, including Sm and Pr atomic self-energies, the Green's functions of the Hubbard model for a Bethe lattice and of the Haldane model for a nano-ribbon, as well as two special test functions. The sensitivity to numerical noise and the dependence on the precision of the numerical libraries are analysed in detail. The present approach is compared to a number of other techniques, i.e. the non-negative least-square method, the non-negative Tikhonov method and the maximum entropy method, and is shown to perform well for the chosen test cases. This conclusion holds even when the noise on the input data is increased to reach values typical for quantum Monte Carlo simulations. The ability of the algorithm to resolve fine structures is finally illustrated for two relevant test functions.

PACS numbers: 71.10.Fd, 71.15.Dx, 02.70.Hm, 05.30.d

I. INTRODUCTION

Strongly correlated materials are currently of great interest due to that they exhibit a plethora of exotic effects which may be important for technological applications^{1,2}. One of the crucial problems behind the investigation of strongly correlated materials is determining their electronic structure. To this aim several computational methods have been developed in the last twenty years, including the dynamical mean-field theory (DMFT)^{3,4} or the GW approach⁵. These methods are often based on the Green's function formalism, and are not always applied directly to real energies. In fact, for technical reasons, it may be more convenient to work with complex energies and then obtain the Green's functions (and related observables) for real energies by means of analytic continuation, as illustrated in Fig. 1.




FIG. 1. Typical example of analytic continuation of the Green's function in the complex plane. The Matsubara frequencies on the imaginary axis are labelled following the convention explained in section II.

The analytic continuation of functions whose formulas are known is a rather simple task in complex analysis. In numerical problems, instead, difficulties arise. Although there still exists a unique continuation if a function is known at infinitely many discrete points on the imaginary axis and with infinite precision⁶, these conditions are not fulfilled in standard computations. The problem depicted in Fig. 1 is therefore ill-posed, especially if there is no a priori knowledge of the function structure on the real axis. Owing to these difficulties, several methods were proposed to perform the analytic continuation of Green's functions and self-energies. The most celebrated approaches include the maximum entropy method (MEM)⁷⁻¹⁴, the Singular Value Decomposition (SVD)^{14,15} and the closely related non-negative Tikhonov (NNT) method¹⁶, the non-negative least-square (NNLS) method¹⁷, stochastic regularization methods¹⁸ and sampling methods¹⁹⁻²². These methods contributed greatly to advance our understanding of the analytic continuation, but in practical terms they still suffer of major drawbacks such as requiring prior information on the function to continue, smearing high energy states or involving great computational efforts.

Another well-known technique for analytic continuation is the Padé approximant method. It consists in parametrising the numerical function by means of a ratio of two polynomials, or equivalently by a terminating continued fraction. There are different schemes for finding the Padé approximant. One is Thiele's reciprocal difference method²³, which is a numerically fast method for attaining the function values at selected points in the complex plane by recursion. It was first applied to condensed matter physics by Vidberg and Serene²⁴ to address the Eliashberg equations. A second approach is to determine the polynomial coefficients explicitly, giv-

ing the function everywhere in the complex plane. In the scheme proposed by Beach *et al.*²⁵, the Padé coefficients are calculated through a standard matrix problem, which makes it possible to use widely available routines for linear algebra.

In general, all the schemes based on the Padé approximant method have the advantage of not requiring any prior information on the function that one intends to continue. However, such an unbiased continuation often results in artefacts and spurious features. In the worst cases the continued function may violate important physical constraints, e.g. leading to a positive imaginary part of the Green's function. These problems are more severe for Padé schemes based on recursive algorithms, due to the propagation of errors. Increasing the precision of the input data as well as that of the numerical routines employed in the algorithms removes spurious features from the continuation²⁵. Unfortunately, obtaining input data with a precision higher than double precision (about 64 bits) is not so straightforward for problems of current scientific interest. Conversely, one must often work with Green's functions and self-energies plagued by numerical noise, as arising from quantum Monte-Carlo (QMC) methods²⁶. Hence, the Padé approximant method has achieved less attention than the other techniques outlined above.

The errors in the Padé approximant method are due to the presence of approximate zero-pole pairs in the polynomials²⁵. These pairs arise from including too many points in the Padé fitting and should ideally cancel each other when taking the ratio between polynomials. The limited precision prevents this from happening, leading to a sort of overfitting problem. Beach *et al.* proposed a procedure to identify the ideal number of points to include in the algorithm but this approach is hard to apply unless Matsubara data of very high precision is available²⁵. More recently, in Refs. 27 and 28, the authors suggested some heuristic procedures to identify and eliminate the zero-pole pairs from the polynomials. Despite representing an improvement with respect to the original formulation, all these schemes still suffer of a certain degree of arbitrariness and require an ad-hoc analysis, case by case.

In the present work, we show that the overfitting problem can be greatly reduced by reformulating the Beach algorithm to use fewer Padé coefficients than input points, which has been suggested already in Ref. 29. This implies transforming the problem into a least square (LS) minimisation. We build on this approach to tackle the sensitivity of the Padé method by constructing a set of analytic continuations where the number of coefficients and input points are varied independently. These continuations can then be used to perform a weighted average in which the spurious zero-pole pairs have a strongly reduced contribution. We show that this procedure stabilises the Padé method and removes unphysical structures. The advantages of our approach are shown to be especially important for low input precision data. Even

for numerical noise comparable to standard QMC simulations our Padé scheme provides good analytic continuations. A further improvement is obtained by exploiting the symmetries of the Green's function, i.e. by considering a small number of negative Matsubara frequencies. Finally, we benchmark our approach on various test-cases, including the atomic self-energies of Sm and Pr, the Green's functions of the Hubbard model for a Bethe lattice and of the Haldane model for a nano-ribbon, as well as some simple test functions. These results are compared with exact solutions as well as with analytic continuations obtained via MEM, NNT and NNLS. The proposed Padé method is shown to often perform better than the other methods over a wide range of noise magnitudes.

The paper is organised as follows. Section I is an introduction, section II is dedicated to a short review of the relevant properties of Green's functions and self-energies which we intend to continue. In section III, our Padé scheme is described in detail. The results are illustrated in section IV, together with an description of step-by-step improvements, tests of the accuracy and stability with respect to Matsubara noise and precision of the numerical routines. A wider range of tests as well as a comparison with some other methods for analytic continuation are presented in section V. Finally the resolving power of our algorithm is analysed in section VI, which is followed by the conclusions.

II. GREEN'S FUNCTION FORMALISM

The one-particle Green's function is a key quantity in the solution of models exhibiting strong correlations. Here we are going to briefly summarise the main properties of the Green's function, which will be then used in section III for explaining and improving the Padé approximant method. In finite temperature Matsubara formalism, the one-electron Green's function of a time-independent Hamiltonian is written as

$$G(\tau) = -\langle T[c(\tau)c^\dagger(0)] \rangle \quad (1)$$

where τ is the imaginary time, T is the time ordering super-operator, $c(\tau)$ and $c(\tau)^\dagger$ are respectively annihilation and creation operators in Heisenberg representation. For simplicity we consider here models with a single orbital, so that no subscripts are needed. The generalisation to multi-orbital systems is straightforward. The expectation value in Eq. (1) denotes the thermal average over the grand canonical ensemble, and $G(\tau)$ is defined on the interval $\tau \in (-\beta, \beta]$, where β is the inverse temperature. A fermionic Green's function is anti-periodic due to the anti-commutation relation $\{c, c^\dagger\} = 1$ and the trace invariance under cyclic permutations. By periodically repeating $G(\tau)$, its Fourier representation is

$$G(\tau) = \frac{1}{\beta} \sum_{n=-\infty}^{\infty} G_n e^{-i\omega_n \tau} \quad (2)$$

with Fourier coefficients

$$G_n = \int_0^\beta d\tau e^{i\omega_n \tau} G(\tau) \quad (3)$$

and fermionic Matsubara frequencies $\omega_n = (2n - 1)\pi/\beta$. A unique spectral function $\rho(\omega) \in \mathbb{R}^+$ exists for a complete set of G_n -values, related to G_n through the Hilbert transform

$$G(z) = \int_{-\infty}^{\infty} d\omega \frac{1}{z - \omega} \rho(\omega) \quad (4)$$

by setting $z = i\omega_n$. Time reversal symmetry is ensured by $G(z^*) = G(z)^*$ and follows from Eq. (4), as is shown in Ref. 25. Due to causality, the Green's function is analytic in the whole complex plane except along the real axis where its imaginary part has a discontinuity. The spectral function in Eq. (4) can be expressed as

$$\rho(\omega) = -\frac{1}{\pi} \text{Im}[G(\omega + i\delta^+)] \quad \text{with } \delta \rightarrow 0^+, \quad (5)$$

and is $1/(2\pi)$ times the magnitude of the discontinuity in $\text{Im}[G(z)]$ on the real axis. The parameter δ in Eq. (5) is the distance above the real axis. The spectral function obeys the sum rule $s = \int_{-\infty}^{\infty} d\omega \rho(\omega) = 1$, thus the Green's function has an asymptote $G(z) \rightarrow s/z$ for $|z| \rightarrow \infty$, according to Eq. (4).

Determining the spectral function $\rho(\omega)$ from a finite set of values G_n with a finite precision is the goal of this paper. Sometimes one can divide the Green's function into a non-interacting Green's function and a correction part due to interactions, expressed by a self-energy function $\Sigma(z)$. Often one can determine the analytical continuation of the non-interacting Green's function exactly, and therefore it may be more convenient to perform the analytic continuation of the self-energy $\Sigma(z)$, instead of the full Green's function $G(z)$. The self-energy has the following analytical form:

$$\Sigma(z) = \Sigma_0 + \Sigma_G(z) \quad (6)$$

with a static part $\Sigma_0 \in \mathbb{R}$ and a dynamic part $\Sigma_G(z) \in \mathbb{C}$. The latter has the same analytic properties as $G(z)$ with the exception of the normalisation value, $s \in \mathbb{R}^+$, which is not necessarily equal to one²⁵. Once the self-energy has been analytically continued to the real axis, then the spectral function may be obtained via the Dyson equation and Eq. (5). In general the procedure outlined above leads to smaller errors than performing directly the analytic continuation of $G(z)$ ^{22,30,31}.

III. PADÉ APPROXIMANTS

A Padé approximant can be expressed as a $[k/r]$ rational polynomial:

$$P_{k,r}(z) = \frac{\sum_{i=1}^{k+1} a_i z^{i-1}}{\sum_{i=1}^r b_i z^{i-1} + z^r}, \quad (7)$$

and has in general complex coefficients a_i and b_i , and an asymptote $a_{k+1} z^{k-r}$ for large $|z|$. Since $G(z)$ and $\Sigma_G(z)$ have asymptotes s/z , they are suitably fitted using a $[(r-1)/r]$ Padé approximant, i.e.

$$P_r(z) = \frac{\sum_{i=1}^r a_i z^{i-1}}{\sum_{i=1}^r b_i z^{i-1} + z^r}. \quad (8)$$

The coefficient $a_r = s \in \mathbb{R}^+$ is significant as it determines the asymptotic behaviour. The number of coefficients are in total $N = 2r$ and can be found as described below.

A. Plain Padé

In Beach's algorithm²⁵, the N coefficients for the Padé approximant $P_r(z)$ of a complex function $f(z)$ are found by selecting N points $z_i \in \mathbb{C}$ where $f(z)$ is known and requiring $P_r(z_i) = f(z_i)$. This results in inverting a N by N matrix. For sake of simplicity, let us call this approach Plain Padé.

B. Least Square (LS) Padé

Using the same number of unknown coefficients as fitting points may lead to overfitting and can give unstable continuations with spurious or even unphysical spectra, e.g. negative intensities. Hence, it may be more useful to have the number of input points M and the number of coefficients N independent under the condition $N \leq M$, as was also suggested by the authors in Ref. 29. Requiring $P_r(z_i) = f(z_i)$ for M points, with $N \leq M$, yields a LS matrix equation

$$K\mathbf{v} = \mathbf{y}, \quad (9)$$

where

$$\mathbf{v} = \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} = \begin{bmatrix} a_1 \\ \vdots \\ a_r \\ b_1 \\ \vdots \\ b_r \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} z_1^r f(z_1) \\ z_2^r f(z_2) \\ \vdots \\ z_M^r f(z_M) \end{bmatrix} \quad (10)$$

and

$$K = \begin{bmatrix} 1 & \cdots & z_1^{r-1} & -f(z_1) & \cdots & -f(z_1)z_1^{r-1} \\ 1 & \cdots & z_2^{r-1} & -f(z_2) & \cdots & -f(z_2)z_2^{r-1} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \cdots & z_M^{r-1} & -f(z_M) & \cdots & -f(z_M)z_M^{r-1} \end{bmatrix}. \quad (11)$$

C. Average Padé approximants

As mentioned in the introduction, for input data of high precision, the performance of Plain Padé is excellent. When the precision is reduced, however, the quality

of the continuation quickly degrades, and shows a strong dependence on the number N of coefficients and input points. A small N may lead to an approximant with not enough poles to reproduce the structure of the function $f(z)$. A big N increases the risk of spurious poles coming from imperfect zero-pole pairs, leading to errors and unphysical features. Due to that the zero-pole pairs arise from numerical noise and have therefore a random distribution²⁷, we suggest to perform the analytic continuation for several N and then average the spectra in contrast to just picking a single N -value. Moreover, for the LS Padé, a similar argument can be formulated for the number of (independent) fitting points M . An average continuation is now defined by the interval of coefficients N and the interval of input points M , as well as by the way the input points are picked. We use subsequent Matsubara points starting at ω_{n_0} (see Fig.1 for explanation of n_0), even though other distributions are possible, like e.g. a logarithmic one. We can now identify a single continuation with the values of $\{n_0, N, M\}$, which we label with the configuration subscript c . The Padé coefficients of c are labeled as \mathbf{v}_c and its spectrum as $\rho_c(\omega) = -1/\pi\text{Im}[P^{(c)}(\omega + i\delta)]$. We now have to decide how to vary the independent coefficients $\{n_0, N, M\}$. For simplicity we choose to use the following sets:

$$n_0 \in \{n_{0_{\min}}:n_{0_{\text{step}}}:n_{0_{\max}}\} \quad (12)$$

$$M \in \{M_{\min}:M_{\text{step}}:M_{\max}\} \quad (13)$$

$$N \in \{N_{\min}:N_{\text{step}}:\min(N_{\max}, M)\}, \quad (14)$$

where for example $n_{0_{\min}}:n_{0_{\text{step}}}:n_{0_{\max}}$ indicates $n_{0_{\min}}, n_{0_{\min}} + n_{0_{\text{step}}}, \dots, n_{0_{\max}}$. Once we have computed the analytic continuations for all the parameters defined above, we can take a weighted average by summing over all configurations c , with weight w_c ,

$$\rho(\omega) = \frac{1}{\sum_c w_c} \sum_c w_c \rho_c(\omega). \quad (15)$$

We have also considered averaging the Padé approximant coefficients, instead of the spectra, for configurations with the same N and then average the spectra from different N . However this turned out to give worse results and will not be discussed further in this article.

The next step is to specify the averaging weight distribution w_c . A necessary condition we apply to all considered distributions is the absence of unphysical configurations. A configuration c is considered unphysical if $\rho_c(\omega) < 0$ for some ω . With this premise, we analyse three possible choices of distributions:

1. Average diagonal Padé

The distribution w_c^D includes only physical configurations with $N = M$, i.e. $w_c^D = \delta_{N,M}$ for physical configurations and $w_c^D = 0$ for unphysical configurations. The superscript D refers to use of only diagonal configurations for which $N = M$.

2. Average LS Padé

The distribution w_c^{LS} includes all physical LS configurations (including those where $N = M$), hence $w_c^{LS} = 0$ for unphysical configurations and $w_c^{LS} = 1$ otherwise.

3. Average similar LS Padé

The distribution w_c^S includes a subset of all physical continuations. Spurious structures may arise due to the finite precision of the ill-conditioned matrix K which needs to be inverted. Being randomly distributed, these spurious structures can be eliminated by favouring spectra similar to each other. To this aim we define a deviation-value between a physical spectral function from all the other physical ones:

$$\Delta_c = \sum_{c' \neq c} w_{c'}^{LS} \int_{-\infty}^{\infty} d\omega |\rho_c(\omega) - \rho_{c'}(\omega)|. \quad (16)$$

In order to filter out the best configurations, we introduce two criteria. Two sets of configurations will be generated, one for each fulfilled criterion, and we pick the intersection. One criterion is to only include configurations with Δ_c smaller than 30% of the average deviation. The latter is given by

$$\bar{\Delta} = \frac{1}{\sum_c w_c^{LS}} \sum_c w_c^{LS} \Delta_c. \quad (17)$$

The other criterion is to restrict the number of configurations to be included to 51% of all the physical ones. This is done by picking the ones with lowest Δ_c .

With the averaging weights, one can also calculate the energy resolved variation from the mean, with big variation for a particular energy indicating instability and possibly unresolved sharp features.

D. Enforcing symmetry

To improve the resulting continuation, it may be advantageous to exploit the symmetry $f(z^*) = f(z)^*$ by including negative Matsubara points, hence $n_0 < 1$, in the fitting. If a symmetric mesh is considered with equally many negative as positive Matsubara points ($n_0 = -M/2 + 1$), spurious poles appear on the real axis, which results in a poor spectrum¹⁴. The more negative n_0 gets, the more spurious peaks appear. By averaging over several continuations many of the peaks will be suppressed due to their random positions. This makes it possible to have a more negative n_0 -value compared to a non-averaging approach. By using a span of different n_0 , the averaging procedure determines how many continuations are taken into account for each n_0 . Assigning higher weights to continuations fulfilling the symmetry criterion $\text{Im}[a_r] = 0$, in the spirit of Ref. 25, sometimes




FIG. 2. (Color online) Illustration of the different improvements done on Plain Padé for a Sm atom, with relative noise of magnitude $\sigma = 10^{-6}$ on the Matsubara points. Description of the four obtained spectra (from left to right panel): Plain Padé, using $N = M = 70, n_0 = 1$, and characterised by unphysical features; Padé average with w_c^D and $n_0 = 1$; Padé average with w_c^{LS} and $n_0 = 1$; Padé average with w_c^S and with mirror symmetry imposed by using $n_0 \in \{-5:1:0\}$.

improves on the accuracy even further, but this strategy has not been analysed in the present article.

IV. RESULTS AND DISCUSSION

A. Averaging methods

In this study, we performed analytic continuation for a variety of functions. An interesting test for illustrating the averaging methods described in section III C is offered by the self-energy of a Sm atom. This function is interesting for two reasons. First of all, it is a rather demanding test for a Padé approximant method, due to the presence of several narrow peaks at close distance from each other, which are associated with the atomic multiplets (see black lines in Fig. 2). Second, this Sm self-energy represents a realistic test case which may arise in state-of-the-art computational problems. The present function has been obtained by means of the electronic structure code RSPt³², when addressing a cluster of seven Sm atoms in the Hubbard I approximation^{30,33}. The Hubbard I approximation can be used to calculate the Green's function on the Matsubara axis. However, it can also be applied directly on the real axis which makes it possible to obtain a reference term to judge the quality of the analytic continuation. Further computational details about this test-case are described in Appendix A. The spectral functions are evaluated at a distance $\delta = 0.01$ Ry above the real axis. For the self-energy Σ , we estimate Σ_0 with $\tilde{\Sigma}_0$ just before doing the continuations, by fitting the asymptote of the real part to $\tilde{\Sigma}_0 + c/\omega_n^2$. This gives us approximately Σ_G , which is suitably fitted by a $P[(r-1)/r]$ Padé approximant described in section III. More sophisticated fittings of the asymptotic tail³⁴ have been tested, and lead to similar results. The dependence

of the Padé performance on the quality of the asymptotic fitting is analysed separately in Appendix B.

The Sm self-energy is continued by using the different averaging methods. The number of coefficients and input points to use in the average are set to $M \in \{50:4:98\}$ and $N \in \{50:4:\min(98, M)\}$. The distributions of averaging weights and input points are instead the following: 1) Plain Padé, $w_c = \delta_{M,70}\delta_{N,M}$ and $n_0 = 1$; 2) Average diagonal Padé, w_c^D and $n_0 = 1$; 3) Average LS Padé, w_c^{LS} and $n_0 = 1$; 4) Average similar LS Padé with mirror symmetry, w_c^S and $n_0 \in \{-5:1:0\}$. The results of the analytic continuations for these four different methods are shown in Fig. 2. The Matsubara self-energy has relative noise of magnitude $\sigma = 10^{-6}$ on every Matsubara point. The dependence of the results on the noise level magnitude will be discussed in the next subsection, here we instead focus on the different averaging techniques. The leftmost panel of Fig. 2 clearly illustrates that for this noise magnitude Plain Padé leads to an unphysical self-energy, whose spectrum becomes negative at about -0.4 Ry. When averages are considered, this pathology is cured. Already with Average diagonal Padé, i.e. by including only physical configurations with $N = M$, the self-energy is acceptable, as illustrated in the second panel (from the left) of Fig. 2. In the third panel, one can see the effect of averaging over more configurations with $N \leq M$. High energy peaks are now captured and the only significant deficiency is that the two peaks just below the Fermi level (zero energy) are merged into one. The distributions of configurations contributing to the averages for the first three cases are visualised in Fig. 3. This figure shows a few physical continuations but it shall be noted that this is related to the particular test function as well as to the magnitude and representation of the external noise. For a generic function the number of physical continuations is typically larger. Moreover, the number of physical continuations can be increased by decreasing the steps M_{step}




FIG. 3. (Color online) Illustration of configurations in (N, M) space with $n_0 = 1$ for the self-energy of a Sm atom, with relative noise of magnitude $\sigma = 10^{-6}$ on Matsubara points. A black cross denotes an unphysical, $\rho_c < 0$, continuation. A circular dot denotes a physical continuation, $\rho_c \geq 0$, and the colour denotes the deviation Δ_c .

and N_{step} . As a general consideration, it is not surprising that the Plain Padé often runs into problems since so many continuations are not even physical. A final improvement to the spectral function is obtained when using the Average similar LS Padé together with symmetry constraints. The rightmost figure of Fig. 2 shows that this method gives an analytic continuation that is very close to the exact result. The two-peak structure at -0.15 Ry is finally resolved. The only noticeable differences are a certain tendency to broaden peaks which are close to each other. This broadening is not surprising and originates directly from the averaging. Not all the continuations contributing to the average are able to reproduce those peaks, and their inclusion results into a broadening.

B. Stability with respect to the numerical noise

The overview presented above is illustrative of the improvements due to the averaging in the Padé scheme. It is important to analyse how these results depend on the numerical precision used in the computation. It is well-known that Beach's algorithm performs well for high precision problems²⁵. A fundamental question here is if only the precision of the input Matsubara data matters or if the internal numerics of Padé also play a role, and to what extent.

We show the influence of the Matsubara precision by illustrating how the analytic continuation of the Sm self-energy presented in the previous subsection changes by adding random noise of variable magnitude. For each Matsubara point, both the real and imaginary parts of the self-energy are multiplied with a factor $(1+\epsilon)$, where ϵ is a normally distributed random variable with zero mean and standard deviation value σ . Although the quality of a continuation is better judged by a close inspection of the resulting function with respect to the exact result, this procedure is difficult to follow for an assessment of a plethora of tests. To avoid this problem we follow Beach's strategy and devise a measure of the error performed on the spectral function:

$$E = \frac{\int_{-\infty}^{\infty} d\omega |\rho_{\text{exact}}(\omega) - \rho(\omega)|}{\int_{-\infty}^{\infty} d\omega |\rho_{\text{exact}}(\omega)|}. \quad (18)$$

The integration bounds are of course limited in computations and were chosen as -1 Ry and $+1$ Ry, unless differently specified. To check how the error changes depending on noise representation, both mean $\langle E \rangle$ and $\sigma_E = \langle (E - \langle E \rangle)^2 \rangle$ are calculated using 10 different seeds for random number generation.

In the top panel of Fig. 4 we report $\langle E \rangle$ and σ_E for the four different methods. Average diagonal Padé, Average LS Padé, and Average similar LS Padé are considered by using the same setup as above but keeping $n_0 = 1$ for all of them. For Plain Padé we use $N = M = 70$, which is representative of a typical calculation. The circles in Fig. 4 indicate the errors $\langle E \rangle$ associated to the curves reported in Fig. 2. This may be useful to understand how the errors quantify different agreements with the exact spectral function. The data reported in the top panel of Fig. 4 illustrates how all the continuations using averaging leads to improved results over a wide range of noise magnitudes. Among the different averages, the schemes based on the LS minimisation perform better, but the improvements are minor. In the bottom panel of Fig. 4 we report analogous data, but using $n_0 \in \{-5:1:0\}$ for all methods, except for the Plain Padé where $n_0 = -2$ is used. Three points become evident. First of all, introducing this small amount of negative frequencies leads to a great improvement of the continuations for noise of all magnitudes smaller than $\sigma = 10^{-4}$. Second, the improvements due to this symmetrisation are much larger for the Padé schemes based on averages, which again emphasises the limitations of the original Padé formulation. Third, the curves for all different Padé schemes based on averages exhibit some sort of plateaux when reducing the amount of noise. These are most likely due to the resolving power of the algorithms and will be analysed separately in section VI.

It is also instructing to look at the number of configurations included in w_c^D , w_c^{LS} and w_c^S . These, resolved for each n_0 , are reported in Table I for a single noise representation of magnitude $\sigma = 10^{-6}$. Notice that in our representation, $n_0 = 0$ corresponds to including the first negative Matsubara frequency, $n_0 = -1$ corresponds to




FIG. 4. (Color online) Average error $\langle E \rangle$ versus standard deviation σ of the numerical noise. The error deviation σ_E (see main text) is represented as error bars. The circles in the plots indicate those continuations corresponding to the spectra reported in Fig. 2. Top panel: analytic continuations performed with $n_0 = 1$. Bottom panel: analytic continuations performed with $n_0 \in \{-5:1:0\}$, except for Plain Padé in red where $n_0 = -2$ is used.

including also the second one, and so on. Interestingly, the three averaging methods give no solutions for $n_0 = 0$, although this is not a general feature of the method but specifically related to this function and noise. By using a lower n_0 , the number of physical configurations increase. In principle, however, there is a non-trivial relation between physical configurations and n_0 . For instance, we can see from Table I that the averaging for w_c^D picks up mostly configurations with $n_0 = -3$ and $n_0 = -5$.

Similar tests as those reported in Fig. 2 have been performed for several other functions (not shown) and one can in general conclude that w_c^S performs better than w_c^{LS} which performs better than w_c^D . This hierarchy is also seen in Fig. 4, although differences are smaller than compared to Plain Padé. Finally notice, especially in the bottom figure, of Fig. 4 the big σ_E for Plain Padé compared to the other schemes.

C. Precision of the numerical routines

As mentioned above, an important question concerns the role of internal numerical precision of the Padé approximant method. For Plain Padé, i.e. Beach's original algorithm, one can estimate the needed precision for the inversion of the matrix K by con-

n_0	-5	-4	-3	-2	-1	0	1
$\#w_c^D > 0$	13	8	11	3	1	0	2
$\#w_c^{LS} > 0$	91	75	81	44	3	0	8
$\#w_c^S > 0$	70	48	29	3	0	0	5

TABLE I. Number of configurations contributing in averages, resolved for each n_0 . Note that w_c^S is determined by comparing separate sets of configurations with each other. In this case, all the configurations arising from $n_0 \in \{-5:1:0\}$ are compared with each other. Configurations with $n_0 = 1$ are instead compared separately with each other. For each n_0 , 91 configurations are in total considered and 13 configurations on the diagonal. The numbers in the rightmost column can be counted from Fig. 3. The analysis is for a self-energy of a Sm atom, with $\sigma = 10^{-6}$ noise added on the Matsubara points.

sidering the ratio between the biggest and smallest elements, which is approximately given by $\xi = (w_{n_{\max}})^r = ((2(M + n_0 - 1) - 1)\pi/\beta)^r$. This yields^{25,29} about $2 \log_2(\xi)$ binary digits, which in our setup varies between 73 and 256 depending on which N, M and n_0 are used. From this one might expect 32 bits (single precision) and 64 bits (double precision) to perform poorly, and perhaps one should resort to 512 bits. To clarify these questions we perform several tests using the method which offered the best results for the Sm self-energy, i.e. Average similar LS Padé with symmetry. We interface our code to a multiple precision arithmetic library MPACK³⁵, using the standard routine *CGESV*. We explore four different precisions of this routine: 64, 128, 256 and 1024 bits. The results of the tests are shown in Fig. 5. It is clear that standard double precision is not enough to solve the numerical problems arising in Padé, unless a very big noise is present. Luckily, quadruple precision offers already good results. Increasing the precision even more, one obtains that 256 and 1024 bits give the same error, which means that no further improvements can be obtained. It is important to stress that these results depend on the choice of test function. In particular, a faster convergence to the exact continuation can be obtained for a more “curvy” function, without the delta-like peaks which are typical for atomic self-energies calculated with the Hubbard I approximation. Nevertheless, Fig. 5 illustrates how significant improvements to the analytic continuation can be obtained by using an increased precision for the LS-routines, even for big Matsubara noise. We therefore use a modified LAPACK³⁶ routine *ZGELS* with a precision of 128 bits for the LS minimisation. In our tests LAPACK performed better than MPACK, and 128 bits LAPACK routines result into errors comparable to 256 bits MPACK. This most likely arises from using different algorithms for the LS problem. In light of these results the LAPACK routines with 128 bits precision were used for all the data presented in the rest of the paper, as well as for those presented in the previous sections.




FIG. 5. (Color online) In the bottom panel, the real axis error $\langle E \rangle$ as a function of input noise σ is reported for various precisions of the numerical routines (MPACK library). σ_E is represented by the error bars. The circles indicate continuations whose spectra are plotted in the top panel, using the same lines as in the bottom panel, and shifted with an offset for a better visualization. Exact results are also reported (lowest curve, in grey).

V. COMPARISON WITH OTHER CONTINUATION METHODS

In this section we compare the most accurate Padé method, as described in the previous section, with other existing methods for analytical continuation. The latter include our in-house implementations of NNLS¹⁷, NNT¹⁶ and MEM⁷⁻¹⁴. A schematic overview of these approaches is given in appendix C. The test functions presented in the following were chosen to cover different types of analytic structures, from the multiplet structures of a Sm atom to the semi-circular density of states of the non-interacting infinite-dimensional Bethe lattice. Both Green's functions and self-energies are considered and the exact function is also known on the real axis to evaluate the accuracy and stability of the analytic continuation.

A. Sm atom in a cluster

The first test case is represented by the function used in the previous sections, i.e. the atomic-like self-energy of one of the sites of a cluster of seven Sm atoms calculated in the Hubbard I approximation, as described in Appendix A. The error as a function of noise for several

methods is presented on the bottom panel of Fig. 6a. We note that the MEM gives a smaller error than NNLS and NNT, for most noise levels. Overall these three methods show a similar behaviour when varying the noise. The averaging Padé scheme gives the most accurate results, except when $\sigma = 10^{-4}$. Even in this case, nevertheless, the difference between Padé and the best performing method is small. Finally, we note that none of the tested methods converge to the exact spectrum when decreasing noise up to the native input precision, and one can still observe a plateau in $\langle E \rangle$, as we emphasise in the section above. This confirms that this self-energy, with spiky features, is a difficult case for analytic continuation. The biggest contribution to the real-axis error for the two spectra in the top panel of Fig. 6a originates from the underestimation of the peak heights for the three peaks closest to the Fermi energy, which is clearly seen from the middle panel of Fig. 6a.

B. Pr atom in the bulk

The second test-case is given by the self-energy of a Pr atom, as obtained from LDA+DMFT simulations of bulk fcc Pr in the Hubbard I approximation (see also appendix A). The error as a function of noise for several methods is illustrated on the bottom panel of Fig. 6b. As for Sm, we again see that the Padé scheme outperforms the other methods for the majority of noise levels. For noise where σ is between 10^{-5} and 10^{-8} the MEM shows a comparable agreement with the exact results. For big noise, it is very surprising that Padé performs better than the MEM. However, other tested self-energies based on the Hubbard I approximation, that are not shown in this article, also show the same behaviour. A significant difference between the left and right bottom panels of Fig. 6 is that for Pr no error plateaux is observed for the Padé scheme, when decreasing noise. As explained in section VI, this is probably related to the analytic structure of the test function, i.e. the distribution of delta-like peaks within the considered energy range. From the upper and middle panel of Fig. 6b, one can see that the advantage of the averaging Padé scheme with respect to other methods is that, for comparable errors, the features it resolves are better captured, such as the height and width of the main peak. This gives a better agreement with the exact continuation, even if more peaks are actually detected by other methods than by Padé.

C. Hubbard model on the Bethe lattice

An interesting test function is the non-interacting Green's function of the Hubbard model on the Bethe-lattice with infinite nearest neighbors³. This function has a smooth semi-circular shape, which makes it much different from the other functions considered so far. The error as a function of noise for $\beta = 100 \text{ Ry}^{-1}$ is shown




FIG. 6. (Color online) Comparison of the Padé scheme with three other well-known methods for analytic continuation. In the bottom panels, the real axis error $\langle E \rangle$ as a function of the Matsubara noise σ is shown. The errors bars represent σ_E . The circles indicate continuations whose spectra are plotted in the top panels. Spectra shifted with an offset for a better visualisation. Exact results are also reported (lowest curves, in grey). In the middle panels the energy-resolved errors of the spectra in the top panels with respect to the exact result are reported.

in Fig. 7 for several methods. For the Padé scheme, we see a stable monotonic decrease of $\langle E \rangle$ as σ is decreased. Even for big noise the error remains relatively small as a few spurious oscillations arise in the spectral function. The performance of the other methods varies greatly. The curve for the MEM shows no improvements for σ smaller than 10^{-8} , due to tiny oscillations around the exact function. We stress here that the MEM is known to have problems in describing sharp band edges far from the Fermi energy. The spectral function for NNT is also characterised by small spurious oscillations (see top panel of Fig. 7), but they disappear with decreasing σ . The NNLS completely fails (at least in our implementation) in describing the present spectral function as many spurious peaks arise on the contour of the semi-circle.

Another interesting observation from Fig. 7 is that $n_0 = 1$ actually performs better than $n_0 \in \{-5:1:0\}$. The latter gives more spurious features, which makes it less precise. We have also obtained this result for some

other test functions. This trend we believe can be traced back to the presence of a discontinuity of the function when going from positive to negative Matsubara frequencies. If the function goes to zero as $i\omega$ goes to zero, like for the atomic self-energies of Sm and Pr, then adding a small number of negative frequencies improves the result significantly. However, if the function goes to a finite value, like for the metallic Green's function of the Bethe lattice, then trying to describe the discontinuity in the imaginary part at the real energy axis is detrimental to Padé. This qualitative difference is not a problem for the Padé method, since the asymptotic behaviour for small frequencies is known from the input Matsubara data, and one can decide beforehand whether to include negative frequencies or not.

D. Haldane model for a nano-ribbon

We also investigate a non-interacting Green's function of an edge-atom of a nano-ribbon, having a honeycomb lattice for the Haldane model³⁷. The spectral function contains many Van Hove singularities which makes it a very challenging test for analytic continuation. We use $\beta = 100$ in units of mRy and a Hamiltonian hopping parameter ratio $\lambda/t = 0.1$. In Fig. 8 real axis errors $\langle E \rangle$ are shown for NNT, MEM and Padé. We see that Padé, NNT and MEM reproduce the exact low-energy spectrum successfully. The overall distribution of spectral weight at higher energy is also captured. Resolving the high energy peaks require data of higher precision than those used here. NNLS creates a plethora of spurious peaks for all noise levels and its real axis error is therefore not shown. This problem is related to the complicated structure of the exact function but also reflects the general tendency of NNLS to resolve a lot of features while leading to substantial errors on the peak positions. This behaviour is different from the present Padé scheme. If the latter resolves the peak structure of a function the




FIG. 7. (Color online) Comparing the Padé scheme with three other well-known methods for the analytic continuation of the non-interacting Green's function of the Hubbard model on a Bethe-lattice with infinite nearest neighbors³. In the bottom panel, the real axis error $\langle E \rangle$ as a function of the Matsubara noise σ is shown. The error bars represent σ_E . The circles indicate continuations whose spectra are plotted in the top panel, using the same lines as in the bottom panel, and shifted with an offset for a better visualisation. Exact spectrum is also reported (lowest curve, in grey).




FIG. 8. (Color online) Comparing the Padé scheme with three other well-known methods for the analytic continuation of the non-interacting Green's function of a nano-ribbon. In the bottom panel, the real axis error $\langle E \rangle$ as a function of the Matsubara noise σ is shown. The error bars represent σ_E . The circles indicate continuations whose spectra are plotted in the top panel, using the same lines as in the bottom panel. Exact spectrum is also reported (in grey).

positions are usually correct.

In the top panel of Fig. 8, we see that the Padé scheme reproduces the exact function very well at low energy and also the shoulders at ± 0.8 mRy are captured. These are completely missed by the other methods. Overall, we again observe that the Padé scheme leads to the best results among the tested methods. In light of the discussion at the end of section V C, we use $n_0 = 1$ for this system, due to its metallic character.

VI. ALGORITHM RESOLVING POWER

From our previous analysis, it may be expected that decreasing the noise on the input Matsubara data should lead to continuous improvement of the analytical continuation obtained with Padé. Nevertheless, the data reported on the bottom panel of Fig. 4 as well as in Fig. 5 illustrate how below a certain noise no further decrease of $\langle E \rangle$ is observed up to the native precision of the Matsubara data. This issue is related to the nature of the analytical continuation problem. The amount of information stored in the input data is not enough to resolve high

energy features, even if one possesses the perfect continuation scheme. In this context, the Sm atom presented in the previous section is a particularly hard test, since there are several high energy satellites, like the double-peak at about -0.6 Ry seen in Fig. 2. The plateaux in $\langle E \rangle$ can then be explained by the fact that the Padé approximant method first manages to resolve all structures close to the Fermi level, but then saturates until the data precision is high enough to resolve also the high energy structures. This unfortunately does not happen in our test, see e.g. Fig. 5, as the native Matsubara data is limited to double precision. A similar problem occurs for the Green's function of the nano-ribbon, where none of the tested methods could resolve the series of high energy peaks. This discussion motivates us to analyse the ability of Padé (and other methods) to resolve fine spectral structures, by using a simple test function. The latter is a Green's function which consists of two poles on the real axis, i.e.:

$$G(z) = \frac{1/2}{z - (\omega_0 + \Delta\omega/2)} + \frac{1/2}{z - (\omega_0 - \Delta\omega/2)}. \quad (19)$$

The spectrum is centered at ω_0 , while the two distinct peaks are located at a distance $\Delta\omega$ from each other. This Green's function is first generated at Matsubara frequencies for an inverse temperature $\beta = 100 \text{ mRy}^{-1}$. From the Matsubara data G is continued to $\delta = 0.1 \text{ mRy}$ above the real energy axis, using Padé settings w_c^S and $n_0 \in \{-5:1:0\}$. Since we are treating a Green's function, no static part is subtracted from the asymptote. We can then identify the maximal noise magnitude σ_{\max} for which the two peaks can still be resolved. Success in resolving the two peaks is defined as the Padé spectrum having two maxima, in a majority of 10 noise representations. In addition, the position of each maximum has to be within a distance δ from the position of the exact peak. The maximum allowed σ_{\max} for $\omega_0 \in \{1, 2, 4, 6, 8\}$ is reported in the left panels of Fig. 9, for two different distances $\Delta\omega$. If the two-peaks could not be resolved for any noise level, no point is plotted. As expected the accuracy of Padé decreases when ω_0 increases, and also when the peaks are closer to each other. It is also interesting to analyse how the algorithm performs when other spectral features are present. To this aim we add a pole centered at zero energy and containing half of the total spectral weight, see the right panels of Fig. 9. The resolving power of the algorithm is slightly decreased, but the scaling with respect to ω_0 seems to be the same as for the original two-peak function.

Before comparing to the other methods, let us discuss the scaling of the curves reported in Fig. 9. The monotonic decrease of σ_{\max} as ω_0 increases is expected and can be understood by comparing $G(z)$ with a single pole function: $1/(z - \omega_0)$ at $z = i\omega_n$. Expanding the difference of them in $\Delta\omega$ one gets $\Delta\omega^2/(i\omega_n - \omega_0)^3$, where it is clear the two-peak structure should be hard to resolve for a big ω_0 or a small $\Delta\omega$. From Fig. 9 we conclude it is less than exponentially hard for Padé to resolve a two




FIG. 9. (Color online) Left panel: highest possible noise magnitude to resolve the two-peak structure in Eq. 19 as a function of the center position ω_0 . Results for the Padé scheme as well as NNLS, NNT and MEM are reported. Right panel: same as left panel but for a function with an additional peak at zero energy.

peak structure as a function of ω_0 . Overall, the data reported in Fig. 9 offer a good quantitative measure of the resolving power of the Padé scheme. When considering pairs of peaks at different distance from zero energy one can explain the real-axis error saturation as a function of Matsubara precision observed in the previous sections for the Padé algorithm.

The resolving power of the other methods shown in Fig. 9 also suffers from the ill-posed nature of the analytic continuation. In addition, the saturation observed when decreasing Matsubara noise also comes from real-axis discretisation and regularisation constraints. In an overall comparison of the methods, the performance in resolving the two-peak structure is somehow the opposite to the performance for smooth curves, see e.g. subsection V C. For the two-peak test function NNLS performs best, MEM worst, and NNT and the Padé scheme are in the middle. This confirms what is outlined in the previous section, i.e. that NNLS resolves many more peaks than the other methods, although at the price of a significant loss of accuracy. Notice also how quickly the performance of MEM degrades with respect to the increase of ω_0 . Smearing of peaks by MEM is a well known issue¹⁹ and is here seen when comparing with the other analytic continuation methods. Finally, notice that only the Padé scheme was able to resolve the peaks centered at $\omega_0 = 6 \text{ mRy}$ and separated by $\Delta\omega = 0.2 \text{ mRy}$.

VII. CONCLUSIONS

We have developed a new Padé scheme to perform the analytic continuation of Green's functions and self-energies. We have first decoupled the number of input Matsubara points and the number of Padé coefficients in Beach's algorithm. To improve the stability and the accuracy we propose to average several physical LS Padé approximants as well as to use numerical routines of higher precision than the Matsubara data has. The averaging is shown to remove spurious features, arising due to the presence of imperfect zero-pole pairs and high numerical precision routines are shown to be crucial for resolving spectrum features, even for noisy Matsubara data. Enforcing mirror symmetry further improves the accuracy, for test functions with small spectral weight at zero energy. The joint usage of averaging physical LS continuations and mirror symmetry is shown to perform better than other well-known methods for analytic continuation, such as NNLS, NNT and MEM, for a variety of test functions. Even for noisy Matsubara data, the Padé scheme is shown to give a better agreement with exact results than other methods. The problem of unphysical continuations, which is considered to be typical of Padé, is here cured for all tested functions. The performance of our algorithm in resolving close peaks at high energies is also high, and inferior only to NNLS, among the investigated methods. But for functions with smooth and several peaky features, the Padé scheme gives higher accuracy than NNLS.

An accurate Padé scheme for analytic continuation presents several advantages with respect to other methods. The Padé method assumes an ansatz for the shape of the exact function and does not depend on any further assumption, such as model functions, normalisation, regularisation or asymptote of the Matsubara data. This makes it tolerant to systematic errors. Moreover, spectral methods, such as MEM, need to discretise the real axis and sometimes the attained spectrum is very dependent on how this is done. This is not a problem for the Padé method, where the function becomes known in the whole complex plane and can just be evaluated on the real energy window of interest. For the spectral methods, the energy window has to be chosen with greater care, since one wants a window which is small enough to have a dense mesh but big enough to ensure the whole spectral weight is included. A too small energy window does not only exclude spectral features but also reduces spectral accuracy inside, since left-out features have to be compensated.

The developed Padé scheme can find several applications in computational condensed matter physics. In LDA+DMFT it can be used to extract spectral functions when calculations are performed on the Matsubara axis. Moreover, increasing applications of the exact diagonalisation solver³⁸⁻⁴¹ demands for the development of new techniques to fit the hybridisation function more accurately. One way to do this is to fit the function in

the whole complex plane instead of using a selected axis (real or imaginary). This may be useful for codes working on the Matsubara axis, for which the exact diagonalisation fitting may be non-trivial^{38,40}. Finally, several implementations of LDA+DMFT include analytic continuation in their computational schemes, such as e.g. the KKR/EMTO implementations of Refs. 27, 42, and 43. A more controlled formulation of Padé like the one presented here may lead to an improvement of the stability and the accuracy of those codes.

VIII. ACKNOWLEDGMENTS

This work was sponsored by the Swedish Research Council (VR), the Swedish strategic research programme eSSSENCE and the Knut and Alice Wallenberg foundation (KAW). The computations were performed on resources provided by the Swedish National Infrastructure for Computing (SNIC). We thank Andreas Östlin and Liviu Chioncel for insightful discussions, Maurits Haverkort for suggesting the section about resolving power, Andrea Droghetti for discussion and providing the input data for the Haldane test function and Erik van Loon for providing test functions and discussions.

Appendix A: Atomic self-energies for Sm and Pr

The self-energies of the Sm and Pr atoms used in the main text were calculated using DMFT within the Hubbard I approximation. In particular we used the RSPt code, where DMFT is interfaced with Full-Potential Linear Muffin-Tin Orbital (FP-LMTO) method³²⁻³⁴. The aim of this appendix is to illustrate the basic concepts behind the calculations and also to present the computational details needed to reproduce our test functions. For an extensive explanation of methods and code, we redirect the reader to the aforementioned references.

In the Hubbard I approximation the many-body problem is reduced to an atomic problem, whose Hamiltonian is:

$$\hat{H} = \sum_{i,j} \hat{H}_{i,j}^{\text{at}} \hat{c}_i^\dagger \hat{c}_j + \frac{1}{2} \sum_{i,j,k,l} U_{i,j,k,l} \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_l \hat{c}_k. \quad (\text{A1})$$

Here the indices i, j, k, l run over the correlated $4f$ orbitals, and the operators \hat{c} and \hat{c}^\dagger are respectively the annihilation and creation operators for those orbitals. The single particle atomic Hamiltonian matrix $\hat{H}_{i,j}^{\text{at}}$ and the Coulomb repulsion tensor $U_{i,j,k,l}$ fully determine the problem to solve. We assume the orbitals i to be atomic-like, whose angular part corresponds to complex spherical harmonics $Y_{l,m}$ for $l = 3$. In this representation, the Coulomb repulsion term is further parametrised using the Slater integrals F^0, F^2, F^4 and F^6 , as explained in Ref. 4. For simplicity let us consider the chemical potential as being at zero energy.

Once the Hamiltonian in Eq. A1 is given, we can solve the atomic problem by direct diagonalisation, which leads to eigenvalues E_ν and eigenvectors $|\nu\rangle$. The Lehmann representation then gives the atomic interacting Green's function:

$$G_{i,j}(z) = \frac{1}{Z} \sum_{\nu_1, \nu_2} \frac{\langle \nu_1 | c_i | \nu_2 \rangle \langle \nu_2 | c_j^\dagger | \nu_1 \rangle}{z + (E_{\nu_1} - E_{\nu_2})} (e^{-\beta E_{\nu_1}} + e^{-\beta E_{\nu_2}}). \quad (\text{A2})$$

From here, the self-energy can be obtained via inverse Dyson equation

$$\Sigma_{i,j}(z) = [\mathcal{G}_0^{-1}(z)]_{i,j} - [G^{-1}(z)]_{i,j}, \quad (\text{A3})$$

where the \mathcal{G}_0 is the non-interacting atomic Green's function. In this article, we focus on the trace of the self-energy, i.e.:

$$\Sigma(z) = \text{Tr}[\Sigma_{i,j}(z)] = \sum_i \Sigma_{i,i}(z). \quad (\text{A4})$$

The input points z are chosen to be Matsubara frequencies for $\beta = 100 \text{ Ry}^{-1}$. The analytic continuation to the real energy axis is evaluated for a distance $\delta = 0.01 \text{ Ry}$.

The local Hamiltonians for a Sm atom in a cluster and a Pr atom in the bulk on the basis of the 14 correlated $4f$ orbitals are respectively

$$\text{Sm: } \hat{H}_{i,j}^{\text{at}} = \begin{pmatrix} -2.187 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -2.193 & 0 & 0 & 0 & 0 & 0 & 0.015 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -2.200 & 0 & 0 & 0 & 0 & 0 & 0.019 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -2.206 & 0 & 0 & 0 & 0 & 0 & 0.021 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -2.212 & 0 & 0 & 0 & 0 & 0 & 0.021 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -2.217 & 0 & 0 & 0 & 0 & 0 & 0.019 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -2.223 & 0 & 0 & 0 & 0 & 0 & 0.015 & 0 \\ 0 & 0.015 & 0 & 0 & 0 & 0 & 0 & -2.223 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.019 & 0 & 0 & 0 & 0 & 0 & -2.217 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.021 & 0 & 0 & 0 & 0 & 0 & -2.212 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.021 & 0 & 0 & 0 & 0 & 0 & -2.206 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.019 & 0 & 0 & 0 & 0 & 0 & -2.200 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.015 & 0 & 0 & 0 & 0 & -2.193 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -2.187 \end{pmatrix}$$

$$\text{Pr: } \hat{H}_{i,j}^{\text{at}} = \begin{pmatrix} -0.678 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -0.683 & 0 & 0 & 0 & 0 & 0 & 0.010 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -0.686 & 0 & 0 & 0 & 0 & 0 & 0.013 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.691 & 0 & 0 & 0 & 0 & 0 & 0.014 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -0.694 & 0 & 0 & 0 & 0 & 0 & 0.014 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -0.699 & 0 & 0 & 0 & 0 & 0 & 0.013 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.702 & 0 & 0 & 0 & 0 & 0 & 0.010 & 0 \\ 0 & 0.010 & 0 & 0 & 0 & 0 & 0 & -0.702 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.013 & 0 & 0 & 0 & 0 & 0 & -0.699 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.014 & 0 & 0 & 0 & 0 & 0 & -0.694 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.014 & 0 & 0 & 0 & 0 & 0 & -0.691 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.013 & 0 & 0 & 0 & 0 & 0 & -0.686 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.010 & 0 & 0 & 0 & 0 & 0 & -0.683 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.678 \end{pmatrix}$$

The orbitals are ordered from $m = -3$ to $m = 3$ for minority spin and then again for majority spin, for both rows and columns. The Slater parameters for determining the Coulomb tensor are instead

$$\begin{aligned} F^0 &= 0.441 \text{ Ry} & F^2 &= 0.876 \text{ Ry} \\ F^4 &= 0.585 \text{ Ry} & F^6 &= 0.433 \text{ Ry} \end{aligned}$$

for Sm and

$$\begin{aligned} F^0 &= 0.514 \text{ Ry} & F^2 &= 0.745 \text{ Ry} \\ F^4 &= 0.488 \text{ Ry} & F^6 &= 0.360 \text{ Ry} \end{aligned}$$

for Pr. As a final remark, please consider that these two tests should be intended as realistic examples but not as physical solutions of given problems.

Appendix B: Dependence on the static self-energy asymptote

In this appendix we will investigate how sensitive the different continuation methods are to a small static part in the self-energy, which might arise from imprecise

asymptotic fitting. The self-energy can be decomposed as $\Sigma = \Sigma_0 + \Sigma_G$, where $\Sigma_G(z) \rightarrow 0$ as $|z| \rightarrow \infty$. All the considered continuation methods work with Σ_G and assume that it has the correct asymptote. However, since we estimate Σ_0 by fitting a constant $\tilde{\Sigma}_0$ to the noisy data, the function to continue may have a small static part $\Delta\Sigma_0 = \Sigma_0 - \tilde{\Sigma}_0$. Therefore, it is interesting to test how sensitive the continuation methods are to a residual static part. We perform this test for the self-energy of a Sm atom, which is described extensively in the main text. We first subtract the static part Σ_0 from the raw data, then add a small static part $\Delta\Sigma_0$ and perform the analytic continuation. Finally we calculate the real-axis error, by comparing to the exact self-energy, which is reported in Fig. 10. It is clear that the Padé algorithm is much less sensitive to a finite $\Delta\Sigma_0$ than the other methods. This is typical of the method and not only of our particular Padé scheme. This tolerance to errors on the asymptote is one of the reasons why the Padé scheme works so well also for big noise magnitudes compared with the other methods. Note that Green's functions do not have a finite static part, but this test may be relevant for them too, since systematic errors in numer-




FIG. 10. (Color online) Real-axis error as a function of the residual static self-energy $\Delta\Sigma_0$. Both positive and negative $\Delta\Sigma_0$ -values are considered, hence the two lines per method. The test is for the self-energy of a Sm atom with double precision Matsubara data.

ical routines (e.g. Fourier transforms) may affect their asymptotic behaviour.

Appendix C: Overview of other methods for analytic continuation

In this appendix we provide a brief description of the other methods for analytic continuation that are used in section V. For a more detailed description of all techniques mentioned here, we redirect the reader to appropriate references. Note that in this implementation we use standard double precision (64 bits). In a few cases some continuations improved by increasing the precision to quadruple (128 bits) but we have not investigated this issue in detail, since it is too far from the scope of this work.

1. Non-Negative Least Square (NNLS)

Instead of fitting to input Matsubara data f_n using a Padé approximant, one can use the Hilbert transform, Eq. (4), to calculate the unknown $\rho(\omega)$. By discretising

the real axis, the integral in Eq. (4) becomes

$$f_n = \sum_{j=1}^N \frac{w_j}{i\omega_n - \omega_j} \rho_j, \quad (\text{C1})$$

where w_j is an integration weight and $\rho_j = \rho(\omega_j)$. The previous equation can be conveniently written as matrix equation of the form $\mathbf{f} = K\boldsymbol{\rho}$. By choosing M Matsubara points, such that $M \geq 2N$, and by respecting the non-negativity of the spectral function, the task becomes solving the following NNLS problem

$$\min_{\boldsymbol{\rho} \geq 0} |\mathbf{f} - K\boldsymbol{\rho}|^2. \quad (\text{C2})$$

A procedure to solve NNLS problems iteratively is described extensively in Ref. 17.

2. Non-Negative Tikhonov (NNT)

Ill-posed problems, such as analytic continuation, are often tackled by introducing regularisations. The Tikhonov regularisation⁴⁴ is based on adding a minimisation condition on ρ , which results into transforming Eq. (C2) into

$$\min_{\boldsymbol{\rho} \geq 0} |\mathbf{f} - K\boldsymbol{\rho}|^2 + |\alpha\boldsymbol{\rho}|^2 = \min_{\boldsymbol{\rho} \geq 0} |\tilde{\mathbf{f}} - \tilde{K}_\alpha\boldsymbol{\rho}|^2. \quad (\text{C3})$$

For a fixed α this is simply a NNLS problem and a suitable α can be determined by e.g. the L-curve method⁴⁴.

3. Maximum Entropy Method (MEM)

The MEM is the most widespread method for analytic continuation⁷⁻¹⁴. Its success is due to the robustness to Matsubara noise. The lack of information on the noisy Matsubara data is counteracted by additional information provided through a model function \mathbf{m} . In the MEM the spectral function is regularised by maximising the entropy. It is a Bayesian approach with a normal probability distribution as a likelihood function and a priori distribution set by the Shannon entropy. In the simplified case where the Matsubara data is assumed to be uncorrelated and with equal precision, the equation to solve can be formulated as:

$$\min_{\boldsymbol{\rho} \geq 0} |\mathbf{f} - K\boldsymbol{\rho}|^2 - \alpha S[\boldsymbol{\rho}], \quad (\text{C4})$$

where S is the entropy which is maximised for $\boldsymbol{\rho} = \mathbf{m}$.

¹ J. Bednorz and K. Müller, Zeitschrift für Physik B Condensed Matter **64**, 189 (1986).

² G. Jonker and J. V. Santen, Physica **16**, 337 (1950).

³ A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996).

- ⁴ G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, *Rev. Mod. Phys.* **78**, 865 (2006).
- ⁵ F. Aryasetiawan and O. Gunnarsson, *Reports on Progress in Physics* **61**, 237 (1998).
- ⁶ C. Titchmarsh, *The Theory of Functions*, Fundamental Theories of Physics (Oxford University Press, London, England, 1939).
- ⁷ R. N. Silver, D. S. Sivia, and J. E. Gubernatis, *Phys. Rev. B* **41**, 2380 (1990).
- ⁸ M. Jarrell and J. Gubernatis, *Physics Reports* **269**, 133 (1996).
- ⁹ R. Bryan, *European Biophysics Journal* **18**, 165 (1990).
- ¹⁰ J. E. Gubernatis, M. Jarrell, R. N. Silver, and D. S. Sivia, *Phys. Rev. B* **44**, 6011 (1991).
- ¹¹ S. Fuchs, T. Pruschke, and M. Jarrell, *Phys. Rev. E* **81**, 056701 (2010).
- ¹² A. Dirks, P. Werner, M. Jarrell, and T. Pruschke, *Phys. Rev. E* **82**, 026701 (2010).
- ¹³ O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, *Phys. Rev. B* **81**, 155107 (2010).
- ¹⁴ O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, *Phys. Rev. B* **82**, 165125 (2010).
- ¹⁵ C. E. Creffield, E. G. Klepfish, E. R. Pike, and S. Sarkar, *Phys. Rev. Lett.* **75**, 517 (1995).
- ¹⁶ V. S. A. Tikhonov, A. Goncharksky and A. Yagola, *Numerical Methods for the Solution of Ill-Posed Problems* (Springer-Science+Business Media, B.V., Moscow, 1995).
- ¹⁷ C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (Society for Industrial and Applied Mathematics Philadelphia, Philadelphia, 1995).
- ¹⁸ I. S. Krivenko and A. N. Rubtsov, eprint arXiv:cond-mat/0612233 (2006), cond-mat/0612233.
- ¹⁹ K. Vafayi and O. Gunnarsson, *Phys. Rev. B* **76**, 035115 (2007).
- ²⁰ A. W. Sandvik, *Phys. Rev. B* **57**, 10287 (1998).
- ²¹ A. S. Mishchenko, N. V. Prokof'ev, A. Sakamoto, and B. V. Svistunov, *Phys. Rev. B* **62**, 6317 (2000).
- ²² P. Staar, B. Ydens, A. Kozhevnikov, J.-P. Locquet, and T. Schulthess, *Phys. Rev. B* **89**, 245114 (2014).
- ²³ G. A. Baker, *Essentials of Padé Approximants* (Academic Press, New York, 1975).
- ²⁴ H. J. Vidberg and J. W. Serene, *J. Low Temp. Phys.* **29**, 179 (1977).
- ²⁵ K. S. D. Beach, R. J. Gooding, and F. Marsiglio, *Phys. Rev. B* **61**, 5147 (2000).
- ²⁶ E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, *Rev. Mod. Phys.* **83**, 349 (2011).
- ²⁷ A. Östlin, L. Chioncel, and L. Vitos, *Phys. Rev. B* **86**, 235107 (2012).
- ²⁸ D. Sokolovskii, E. Akhmatkaya, and S. K. Sen, *Computer Physics Communications* **182**, 448 (2011).
- ²⁹ Ž. Osolin and R. Žitko, *Phys. Rev. B* **87**, 245135 (2013).
- ³⁰ A. Grechnev, I. Di Marco, M. I. Katsnelson, A. I. Lichtenstein, J. Wills, and O. Eriksson, *Phys. Rev. B* **76**, 035107 (2007).
- ³¹ X. Wang, E. Gull, L. de' Medici, M. Capone, and A. J. Millis, *Phys. Rev. B* **80**, 045101 (2009).
- ³² J. Wills, A. Delin, and O. Eriksson, *Full-Potential Electronic Structure Method* (Springer, New York, 2000).
- ³³ P. Thunström, I. Di Marco, A. Grechnev, S. Lebègue, M. I. Katsnelson, A. Svane, and O. Eriksson, *Phys. Rev. B* **79**, 165104 (2009).
- ³⁴ O. Grånäs, I. Di Marco, P. Thunström, L. Nordström, O. Eriksson, T. Björkman, and J. Wills, *Computational Materials Science* **55**, 295 (2012).
- ³⁵ M. Nakata, *The MPACK (MBLAS/MLAPACK): A Multiple Precision Arithmetic Version of BLAS and LAPACK*, <http://mplapack.sourceforge.net>, 0th ed.
- ³⁶ E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999).
- ³⁷ F. D. M. Haldane, *Phys. Rev. Lett.* **61**, 2015 (1988).
- ³⁸ P. Thunström, I. Di Marco, and O. Eriksson, *Phys. Rev. Lett.* **109**, 186401 (2012).
- ³⁹ Y. Lu, M. Höppner, O. Gunnarsson, and M. W. Haverkort, *Phys. Rev. B* **90**, 085102 (2014).
- ⁴⁰ J. Kolorenč, A. I. Poteryaev, and A. I. Lichtenstein, *Phys. Rev. B* **85**, 235136 (2012).
- ⁴¹ M. Granath and H. U. R. Strand, *Phys. Rev. B* **86**, 115111 (2012).
- ⁴² L. Chioncel, L. Vitos, I. A. Abrikosov, J. Kollár, M. I. Katsnelson, and A. I. Lichtenstein, *Phys. Rev. B* **67**, 235106 (2003).
- ⁴³ J. Minár, L. Chioncel, A. Perlov, H. Ebert, M. I. Katsnelson, and A. I. Lichtenstein, *Phys. Rev. B* **72**, 045125 (2005).
- ⁴⁴ P. C. Hansen, in *Computational Inverse Problems in Electrocardiology*, ed. P. Johnston, *Advances in Computational Bioengineering* (WIT Press, 2000) pp. 119–142.