Comparison between methods of analytical continuation for bosonic functions

J. Schött,[1] E. G. C. P. van Loon,[2] I. L. M. Locht,[1,2] M. I. Katsnelson,[2] and I. Di Marco[1]

[1]*Department 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, Netherlands*

In this paper we perform a critical assessment of different known methods for the analytical continuation of bosonic functions, namely, the maximum entropy method, the non-negative least-squares method, the non-negative Tikhonov method, the Padé approximant method, and a stochastic sampling method. Four functions of different shape are investigated, corresponding to four physically relevant scenarios. They include a simple two-pole model function; two flavors of the tight-binding model on a square lattice, i.e., a single-orbital metallic system and a two-orbital insulating system; and the Hubbard dimer. The effect of numerical noise in the input data on the analytical continuation is discussed in detail. Overall, the stochastic method by A. S. Mishchenko *et al.* [Phys. Rev. B **62**, 6317 (2000)] is shown to be the most reliable tool for input data whose numerical precision is not known. For high-precision input data, this approach is slightly outperformed by the Padé approximant method, which combines a good-resolution power with a good numerical stability. Although none of the methods retrieves all features in the spectra in the presence of noise, our analysis provides a useful guideline for obtaining reliable information of the spectral function in cases of practical interest.

I. INTRODUCTION

Strongly correlated materials exhibit a wide range of exotic physical phenomena, ranging from magnetism to superconductivity [1]. This rich physics is interesting for the scientific community since it has a high potential for future technological applications. In the last few decades several theories have been developed to describe the electronic structure of strongly correlated materials with a good accuracy, e.g., the combination of density-functional theory (DFT) and dynamical mean-field theory (DMFT) [2–5]. Practical calculations for systems at finite temperature are usually performed using the Green's function formalism for complex energies [6,7]. In this approach determining physical observables requires an analytical continuation from complex energies (Matsubara frequencies) to real energies, as depicted in Fig. 1.

Traditionally, the analytical continuation of one-particle Green's functions has received the most attention since it is related to the one-particle spectral function. The latter not only is of fundamental importance but can also be probed (more or less) directly via various types of photoemission experiments. Over the last few years, however, dynamical *two-particle* quantities have become more and more important [8–16]. Significant examples include the electron-energy-loss spectrum (EELS) [9,13] and dynamical susceptibilities [15,17,18]. Physically, the two-particle spectrum contains several interesting features, such as low-energy dispersive Goldstone modes, Landau damping of collective modes, and sharp transitions between isolated energy levels. Even in a weakly interacting Fermi liquid, the two-particle spectrum presents some of these interesting features [7]. From a more technical point of view, two-particle quantities also play an important role in the treatment of nonlocal interaction in DMFT-based approaches such as extended DMFT (EDMFT) [19–23] and its extension dual boson [24] by describing the feedback of collective excitations on the one-particle spectrum.

Similar to the one-particle Green's function, two-particle quantities are usually obtained on the Matsubara axis. However, whereas the one-particle Green's function is determined on the *fermionic* Matsubara frequencies, the two-particle quantities are calculated on the *bosonic* Matsubara frequencies. Some physical information, such as the occurrence of charge-order transitions [11,25,26], can be obtained directly from computational data on the Matsubara axis at zero frequency. However, the full dynamical susceptibilities, as well as related observables such as the EELS and the plasmon spectrum, require an analytical continuation to real energies. It is easy to understand why finding a reliable method for the analytical continuation of two-particle functions has evolved from a niche problem to a necessity in computational many-body physics.

The usual methods of analytical continuation of fermionic functions can, with the appropriate kernel and symmetry relations, also be used for bosonic functions [27], as well as for more exotic and sophisticated cases [28,29]. It was previously shown that the maximum entropy method (MEM) is able to perform the analytical continuation of the optical conductivity with good accuracy [30]. More recently, in the context of EDMFT, Huang *et al.* [12] focused on the retarded interaction, emphasizing that a modified version of MEM [31] leads to a good analytical continuation. While these studies provide a very interesting and useful analysis, they do not cover the applicability of various methods to functions of different shapes and characters. So far, a more complete study of the analytical continuation problem for bosonic functions is lacking. In this work we intend to provide such an analysis by considering five different methods of analytical continuation in four physically relevant scenarios. In addition to the MEM [27,30,32–37], we will also test the non-negative least-squares (NNLS) method [38,39], the non-negative Tikhonov (NNT) method [39,40], the Padé approximant (Padé) method [41,42], and Mishchenko's stochastic sampling method [43,44]. Furthermore, we will investigate the effect of numerical noise on the quality of the continuation in connection to the usage of quantum Monte Carlo (QMC) techniques (see, e.g., Ref. [45]).

This paper is structured as follows: in Sec. II we review bosonic Green's functions and their analytical and symmetry
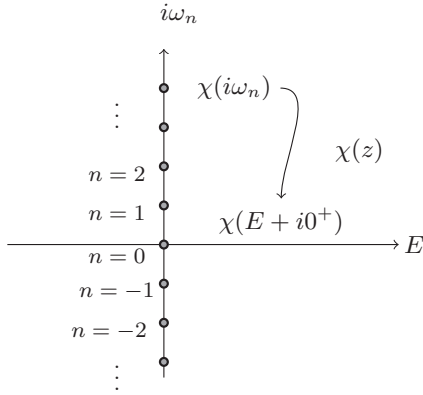
FIG. 1. Schematic illustration of the analytic continuation of a bosonic response function $\chi(z)$. The values $\chi(i\omega_n)$ at the Matsubara frequencies $i\omega_n$ are used to reconstruct the function $\chi(E + i\delta^+)$ just above the axis of real energies.

properties. In Sec. III a brief introduction to the various continuation schemes is given. Then, in Sec. IV, we describe the different testing cases and highlight the reasons why they are interesting. In Sec. V, we show and analyze the results of the tests. Finally, in Sec. VI we present the conclusions of our study, including a comparative summary of the performance of the continuation methods investigated.

II. BOSONIC GREEN'S FUNCTIONS

The analytical continuation consists of obtaining a function $\chi(z)$ in the whole complex plane from a finite set of values. The focus in this work is on functions with bosonic symmetry that are known on a finite number of bosonic Matsubara frequencies $i\omega_n = 2n\pi T i$, where T is the temperature and n is an integer number. From this partial knowledge, we intend to extract the values of the function in the entire complex plane, especially just above the real axis, for $z = E + i\delta$.

The analytical continuation problem can be formulated in terms of the spectral function $\rho(E)$ by means of the Hilbert transform

$$\chi(i\omega_n) = \int_{-\infty}^{\infty} dE \frac{1}{i\omega_n - E} \, \rho(E). \tag{1}$$

This equation is an example of the Fredholm equation of the first kind. This inversion problem, as well as the methods used to solve it, actually appear across many different disciplines. We focus on bosonic functions, with the symmetry that its spectral function is an odd function, i.e., $\rho(E) = -\rho(-E)$. Examples of odd bosonic functions are the density-density correlation function $\langle \hat{n}\hat{n} \rangle$ and the spin-spin correlation function $\langle \hat{S}_z \hat{S}_z \rangle$. Correlation functions of noncommuting observables, such as $\langle \hat{S}_x \hat{S}_y \rangle$, are, in general, not odd. Our study does not include this type of function. The oddness allows us to simplify Eq. (1) to

$$\chi(i\omega_n) = \int_0^{\infty} dE \frac{-2E}{\omega_n^2 + E^2} \, \rho(E). \tag{2}$$

From this equation we find that for large Matsubara frequencies $\chi(i\omega_n)$ is asymptotically proportional to ω_n^{-2}. Without the odd symmetry of ρ, one would obtain an asymptotic behavior

as ω_n^{-1}, like for the one-particle Green's function. Furthermore, $\chi(i\omega_n)$ is purely real, and $\chi(i\omega_n) = \chi(-i\omega_n)$. It follows from the analytical properties of $\chi(z)$ that $\chi(i\omega_n)$ decreases monotonously as a function of $i\omega_n$ [46]. The correct symmetry can be enforced in the analytical continuation procedure either explicitly by using Eq. (2) or implicitly by including both positive and negative Matsubara frequencies as input points. This issue is studied in more detail in Appendix A.

III. CONTINUATION ALGORITHMS

The analytical continuations are performed by means of five different methods. We use our own in-house implementations for all algorithms except for Mishchenko's method. For Mishchenko's method we use an message passing interface-parallelized version of the original code presented in Ref. [43]. Here, a brief introduction to these algorithms is given.

A. NNLS method

Discretizing the integral in Eq. (2) makes it possible to reformulate the problem as a system of linear equations

$$\chi(i\omega_n) = \sum_j f_j \underbrace{\frac{-2E_j}{\omega_n^2 + E_j^2}}_{K_{n,j}} \rho_j, \tag{3}$$

where f_j is a quadrature weight and $K_{n,j}$ is the matrix to invert. However, this approach would not work in practice since obtaining $\rho(E)$, once $\chi(i\omega_n)$ is given, is an ill-posed problem. The NNLS method solves Eq. (3) in a least-squares sense and stabilizes the solution by enforcing the known symmetry property $\rho(E)E \geqslant 0$. This means finding a non-negative solution to a least-squares problem:

$$\min_{\rho \geqslant 0} \| \chi - K \rho \|^2. \tag{4}$$

Here, χ and ρ are vectors containing the values $\chi(i\omega_n)$ and ρ_j on the discrete points of Eq. (3). The NNLS problem can be solved iteratively as described extensively in Ref. [38].

B. NNT method

For ill-posed problems, regularizations are commonly used. One of the most famous regularization methods is the Tikhonov method [40]. Applying the Tikhonov method to the problem of analytical continuation defines the NNT method as

$$\min_{\rho \geqslant 0} \| \chi - K \rho \|^2 + \alpha \| \rho \|^2, \tag{5}$$

where α is a weight parameter that can be determined by the L-curve method [40], which selects α corresponding to the smallest value of $\ln(\| \chi - K \rho_\alpha \|^2) + \ln(\| \rho_\alpha \|^2)$. Here, ρ_α denotes the solution of Eq. (5) for a fixed α.

C. MEM

Another famous regularization method is the maximum entropy method [27,30,32–37]. This method maximizes the Neumann entropy for the spectral function. Formulated as a minimization problem, the equations become

$$\min_{\rho \geqslant 0} \| \chi - K \rho \|^2 + \alpha \| S[\rho] \|, \tag{6}$$

where the entropy

$$S[\rho] = \int dE \rho(E) \ln\left(\frac{\rho(E)}{m(E)}\right)$$

$$\approx \sum_j f_j \rho_j \ln\left(\frac{\rho_j}{m_j}\right) \qquad (7)$$

contains a default model $m(E)$, which can incorporate *a priori* knowledge about the spectral function.

D. Padé approximant method

The Padé approximant method [41,42] is based on a fitting procedure. This method starts with a rational polynomial ansatz for $\chi(z)$ with unknown coefficients and fits this ansatz to the known values at the Matsubara axis $\chi(i\omega_n)$. In this way one can reconstruct the function in the entire complex plane. The Padé method can be formulated as a matrix problem [41], which improves its numerical accuracy. As recently proposed [47], one can enhance the stability of the continuations (especially in the presence of Matsubara noise) by taking an average of several Padé approximants, obtained by varying the number of Padé coefficients and Matsubara points in the fitting procedure.

E. Mishchenko's stochastic sampling method

Finally, we consider the stochastic sampling method proposed by Mishchenko *et al.* [43,44]. In this method the spectral function is approximated by a set c of rectangles R, i.e.,

$$\rho_c(\omega) = \sum_{R \in c} R_{\{h,b,m\}}(\omega).$$

The rectangles R are defined by their height h, width b, and center position m. A stochastic algorithm updates the rectangles randomly and accepts the changes according to a Metropolis algorithm based on the difference between the known function $\chi(i\omega_n)$ and the corresponding function obtained from $\rho_c(\omega)$. Combining this update scheme with a deterministic minimization of the Matsubara difference further improves the performance of the method. An average over different independent Monte Carlo chains is taken to minimize the influence of noise. Since the algorithm uses a condition for the sum of the rectangle weights, we scale ρ by introducing $\tilde{\rho}(E) = \frac{-2}{E\chi_0}\rho(E)$. Instead of finding ρ by solving Eq. (2), we calculate $\tilde{\rho}$, which has to obey

$$\tilde{\chi}(i\omega_n) = \frac{\chi(i\omega_n)}{\chi_0} = \int_0^\infty dE \frac{E^2}{\omega_n^2 + E^2} \tilde{\rho}(E). \qquad (8)$$

For $n = 0$ we have the desired normalization condition $1 = \int_0^\infty dE \, \tilde{\rho}(E)$. Once $\tilde{\rho}$ is found, we scale back to obtain ρ. This rescaling can, in principle, be used for the other methods as well, with the exception of Padé where the continuation is not done by solving Eq. (2).

IV. MODEL DESCRIPTIONS

To investigate the quality of the analytical continuation provided by the different methods we focus on test functions that are known in the entire complex plane, including the real-energy axis. The functions are first evaluated at the Matsubara energies and then analytically continued to the real axis. In an ideal situation, the solution obtained by analytical continuation coincides with the original exact values. However, in practice, some differences arise due to the limited precision and amount of input data.

The continuation becomes more difficult when we add numerical noise to the Matsubara data, which models what happens in real calculations such as by, e.g., QMC methods. There, the noise typically scales inversely with the square root of the computational time. As a practical example, the numerical noise in dual boson calculations [25] is often in the 0.1% range. The numerical noise is modeled by adding relative Gaussian noise to each Matsubara point, $\chi(i\omega_n)(1 + \epsilon)$, where ϵ is sampled from a Gaussian distribution with zero mean and a standard deviation σ. We use the same σ for all frequencies. In realistic QMC simulations, the noise level will usually depend on the Matsubara frequency $i\omega_n$ in a complicated way. This issue will be investigated more in detail in Sec. V D for the Hubbard dimer.

It is important to individuate test functions that are realistic and cover bosonic functions of different character as one would encounter in practice. Naturally, the exact real-energy solutions of many of the most interesting systems are not known, so we need to find realistic but tractable cases. Here, we consider four scenarios: (i) a two-pole model, (ii) the momentum-resolved Lindhard function of the two-dimensional doped tight-binding model, (iii) that of a band insulator with two identical bands separated by a gap, and (iv) the Hubbard dimer. These four functions, which are explicitly described in Sec. V, cover sharp features, a sharp low-energy zero-sound mode (that gets broadened by Landau damping at larger momenta), high-frequency modes, features far away from zero energy, and a combination of sharp features with a broad base. The Hubbard dimer is chosen to model realistic QMC noise.

All these functions are evaluated for Matsubara frequencies corresponding to an inverse temperature $\beta = 50$ and at a distance above the real axis $\delta = 0.05$. For all methods we used 100 Matsubara frequencies, except for Padé, where a spectral average is done by varying the number of Matsubara points between 50 to 100 in steps of 4. For the Hubbard dimer, instead, we used 51 Matsubara frequencies. It is important to use a sufficiently large number of Matsubara frequencies. We study the dependence of the quality of the continuation on the number of Matsubara frequencies in Sec. V A. We also investigate the effect of including a few negative frequencies in Appendix A.

V. RESULTS

In this section we present a comparison between the different methods of analytical continuation when applied to four selected test functions.

A. Two-pole model

The first function we address is a simple model function whose analytical structure is

$$\chi(z) = \frac{a_1}{z^2 - E_1^2} + \frac{a_2}{z^2 - E_2^2}, \qquad (9)$$
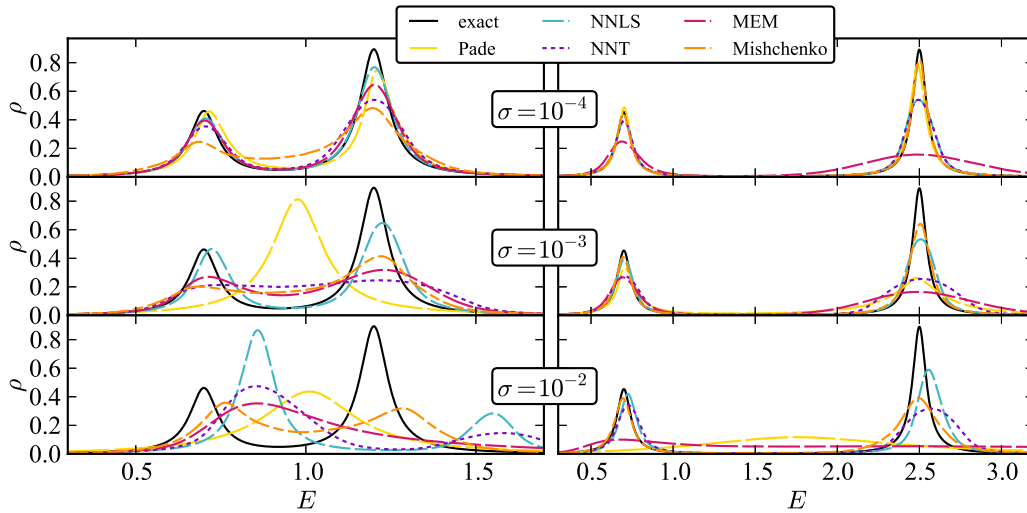
FIG. 2. Spectra for the two-pole test function. Left: $a_1 = 0.1$, $a_2 = 0.335663$, $E_1 = 0.7$, and $E_2 = 1.2$; right: $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$, and $E_2 = 2.5$, as defined in Eq. (9).

where the parameters have the following values: $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$, and $E_2 = 2.5$. For $\delta \rightarrow 0^+$, the exact spectrum has two distinct peaks at energies E_1 and E_2. Therefore this function has similarities to the Gaussian two-peak structure used in Ref. [12].

To understand how the position of the high-energy peak influences the quality of the attained spectra, we also perform tests with parameters $E_2 = 1.2$ and $a_2 = 0.335663$. This choice preserves the height of both peaks in the two chosen setups. The two spectra, corresponding to $E_2 = 1.2$ and $E_2 = 2.5$, present different degrees of difficulty. For $E_2 = 1.2$, the two peaks in the spectrum are close to each other, and therefore it is difficult to resolve them as two separate peaks. For $E_2 = 2.5$, on the other hand, the second peak is far away from the first peak but also far away from zero energy (the complex $i\omega_n$ axis). This makes it difficult to resolve the distant peak.

In Fig. 2 the two exact spectra are reported and compared to the analytical continuation obtained through the five methods for three different noise levels. It is known that the Padé method is much more sensitive to the precision of the input data than the other methods. The various noise levels are therefore chosen to illustrate the critical input precision where the Padé method starts to fail. These noise levels are realistic for QMC simulations [45], as also discussed in Sec. V D.

As the top panels of Fig. 2 show, for relatively low noise ($\sigma = 10^{-4}$), all methods capture the two peaks and their positions correctly for both sets of parameters. The Padé method performs the best if we also consider the height of the peaks. Conversely, the MEM has the most difficulties to resolve the distant peak at $E_2 = 2.5$. This picture changes drastically for intermediate noise levels ($\sigma = 10^{-3}$), shown in the middle panels of Fig. 2. For $E_2 = 1.2$ (left panel), the Padé method is not able to reproduce the two peaks and finds only the one at low energy. The other methods also gain spectral weight in between the peaks but are still able to resolve both peaks. For $E_2 = 2.5$ (right panel), results are significantly better for all methods. Interestingly, the MEM performs worse than the Padé method. In situations where

the input data are characterized by high noise ($\sigma = 10^{-2}$), depicted in the bottom panels of Fig. 2, both the MEM and the Padé method fail to find distinct peaks and merge them into one single broad peak. The other methods find distinct peaks but locate them at the wrong energies, which is especially evident for $E_2 = 1.2$ (left panel). In any case, these methods are still able to describe the basic physics correctly, i.e., the presence of two well-defined peaks. In general, Mishchenko's method seems to offer the best results across various levels of noise. Even for the highest noise ($\sigma = 10^{-2}$), Mishchenko's method gives an acceptable description of the exact spectra. Nevertheless, one has to notice, if two peaks are close to each other, this approach tends to increase the spectral weight in between the peaks. As a matter of fact, this problem is not an intrinsic property of Mishchenko's algorithm and is related to an insufficient number of local updates in the sampling. The dependence of Mishchenko's results with respect to the number of local updates is investigated in Appendix B. For the two-pole model, for small noise levels, Mishchenko's method can achieve results as good as those of the Padé method, although at the cost of a much bigger computational effort.

Finally, from the comparison of the left and right panels of Fig. 2, one can conclude that partially overlapping peaks seem to be more difficult to describe than having one peak close to and one peak far from zero energy. This is not so trivial since, usually, analytical continuation methods are supposed to perform well to describe (even complex) structures close to zero energy.

Number of Matsubara frequencies

It is also useful to illustrate how the quality of the analytical continuation depends on the number of input points, e.g., the number of Matsubara frequencies. The quality of the continuation can be measured through the difference between the obtained $\rho(E)$ and the exact function on the real axis $\rho_{\text{exact}}(E)$ in the following way:

$$\text{Error} = \frac{\int_0^5 dE |\rho(E) - \rho_{\text{exact}}(E)|}{\int_0^5 dE |\rho_{\text{exact}}(E)|}, \quad (10)$$
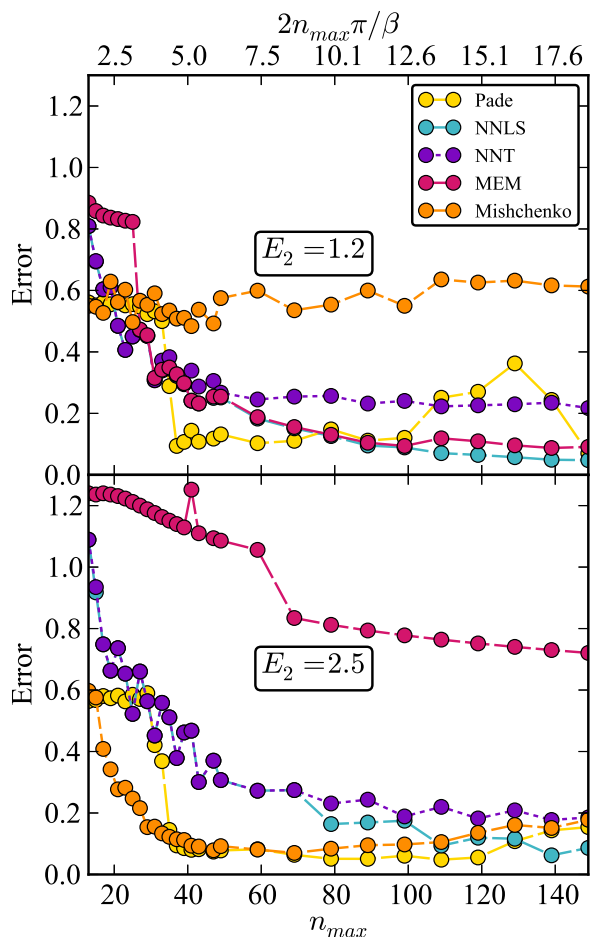
FIG. 3. Integrated real-axis error for the two-pole model as a function of the highest Matsubara point index $n_{\max}$ used. Top: $a_1 = 0.1$, $a_2 = 0.335663$, $E_1 = 0.7$ and $E_2 = 1.2$; bottom: $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$, and $E_2 = 2.5$, as defined in Eq. (9). Matsubara noise level $\sigma = 10^{-4}$ is used.

where the upper bound of the integrals is chosen to be larger than the actual extension of the function. The error measure given by Eq. (10) for the two-pole model with noise $\sigma = 10^{-4}$ is plotted in Fig. 3 with respect to different numbers of the maximum Matsubara index $n_{\max}$. Note that for the Padé method the notion of $n_{\max}$ is a bit more complicated to define since an average is done over several continuations with varying numbers of input points and Padé coefficients [47]. To make optimal usage of the averaging procedure, both the number of input points and Padé coefficients take values between $n_{\max} - 47$ and $n_{\max} + 1$ in steps of 4 under the constraint of not having more coefficients than Matsubara points. If $n_{\max} - 47 < 4$, then 4 is used as the lower boundary instead.

In Fig. 3, the general trend is that the quality of the continuation initially improves with an increasing number of Matsubara points. At some point, the error stabilizes. When the typical energy scale on the real axis (E_2) increases, this stabilization point moves to higher Matsubara frequencies. However, there are small variations to this general picture for the different methods.

The error between the Padé continuation and the exact function quickly drops at $n_{\max} \approx 30$ and is very small until $n_{\max} \approx 100$ (or 120 for $E_2 = 2.5$). When taking into account more Matsubara points, the error increases. We attribute this behavior mainly to the number of physical continuations in the averaging procedure. For small $n_{\max}$, the total number of continuations is limited. For large $n_{\max}$, the number of physical continuations decreases significantly. For this particular test, at least 30 to 40 physical continuations were needed to optimally use the power of the averaging procedure. The continuations obtained with NNT and NNLS, on the other hand, continue improving when additional Matsubara frequencies are taken into account. However, the improvement becomes smaller and smaller. The MEM error quickly drops at $2n_{\max}\pi/\beta \approx 3E_2$ and decreases only slightly for higher $n_{\max}$. The error of the Mishchenko continuation stabilizes at a relatively small number of Matsubara frequencies. For this method, we have used the same number of local updates (15 000) for all continuations. The behavior for more local updates is discussed in Appendix B. It is conceivable that adding additional frequencies does give extra information but also requires a longer run time of the stochastic sampling.

B. Doped tight-binding model

For the second test, we use a metallic system, namely, the tight-binding model on a square lattice with nearest-neighbor hopping. The density-density correlation function depends on the Matsubara frequency and on the momentum **q**. The random-phase approximation [7,48] tells us that the susceptibility in the presence of a weak interaction $V(\mathbf{q})$ can be obtained from the susceptibility χ of the noninteracting system as

$$\chi(i\omega_n, \mathbf{q})^{-1} = \chi(i\omega_n, \mathbf{q})^{-1} - V(\mathbf{q}). \tag{11}$$

An example is the Hubbard model, with $V(\mathbf{q}) = U$. Since the interaction $V(\mathbf{q})$ is independent of the Matsubara frequency, for the purpose of investigating the analytical continuation, it is sufficient to study the noninteracting density-density susceptibility, which is given by the Lindhard function

$$\chi(i\omega_n, \mathbf{q}) = \sum_{\mathbf{p}} \frac{n_F(\epsilon_{\mathbf{p}}) - n_F(\epsilon_{\mathbf{p+q}})}{\epsilon_{\mathbf{p}} - \epsilon_{\mathbf{p+q}} + i\omega_n}, \tag{12}$$

where the sum denotes an average over the Brillouin zone, n_F is the Fermi distribution, and the dispersion relation is

$$\epsilon_{\mathbf{p}} = -2t[\cos p_x + \cos p_y]. \tag{13}$$

The energy scale of this model is given by the half bandwidth $D = 4t$, where t is the hopping parameter. Therefore, we fix the energy unit to $D = 4t = 1$. The chemical potential is set to $\mu = -0.5$, which corresponds approximately to a filling of 0.185 per spin flavor.

In Fig. 4 the analytical continuation of the momentum-resolved Lindhard function for all methods is reported. The panels in the first three columns correspond to data for various levels of noise σ. The panels in the last column illustrate the exact spectrum on the real axis, repeated in each row for an easier comparison.

At small wave vectors (close to Γ), the susceptibility exhibits a single low-energy mode, the zero-sound mode. This
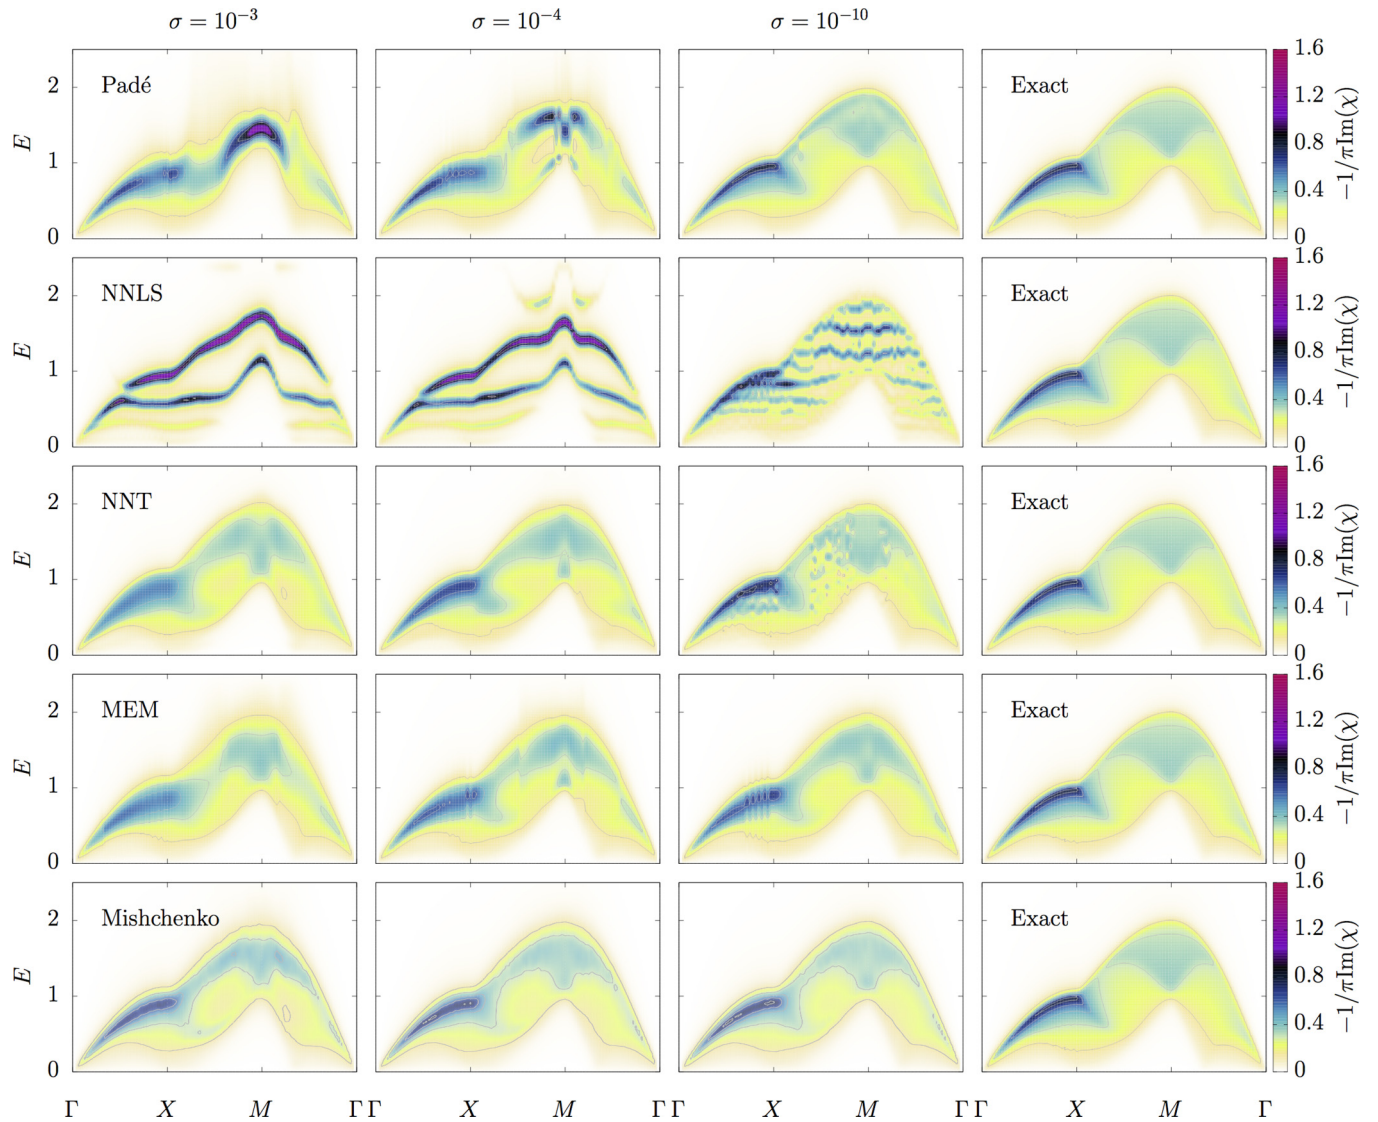
FIG. 4. Spectrum of the Lindhard function $(-\frac{1}{\pi}\text{Im}[\chi(\omega + i\delta)])$ of the two-dimensional doped tight-binding model along the Brillouin zone path $\Gamma \to X \to M \to \Gamma$, obtained using Padé, NNLS, NNT, MEM, and Mishchenko's methods. The noise level varies among the values $\sigma = 10^{-3}$, $\sigma = 10^{-4}$, and $\sigma = 10^{-10}$. All panels on the right-hand side show the exact spectrum for an easier comparison. The color map denotes the intensity of the spectra, and contour lines are added at $-1/\pi\,\text{Im}(\chi) = 0, 0.15, 0.30, 0.45, 0.60, 0.75$. For the NNLS plots not all contour lines are displayed (see main text). The spectra for the M point separately are also reported in Fig. 5.

mode consists of low-energy excitations close to the Fermi surface. Farther away from Γ, Landau damping broadens this zero-sound mode. All methods correctly capture the zero-sound mode. However, NNLS shows an unphysical splitting of this mode into two branches roughly halfway from $\Gamma \to X$. The tendency of NNLS to introduce spurious peaks is even more pronounced when going from $X \to M$, and in fact, it is not even possible to illustrate all contour lines in Fig. 4 without compromising its readability. The complex spectral evolution in the path $X \to M$ seems, indeed, rather difficult to describe for all methods, especially for high noise. At the M point the exact spectrum acquires a seemingly simpler structure, consisting of a relatively flat bump between $E = 1$ and $E = 2$. This point is analyzed in more detail in Fig. 5, which shows that it is not so easy to capture the broad flat feature instead of two separate peaks. Similar to the previous test function,

Padé correctly reproduces the spectrum at low noise but gives a single sharp peak instead of a broad flat mode when σ is increased. NNLS performs very poorly here, with a spectrum of a few sharp peaks even at low noise, $\sigma = 10^{-10}$. NNT does a lot better compared to NNLS; the regularization flattens the spectrum, although an unphysical local minimum remains at $\sigma = 10^{-4}$. Similarly, MEM also captures the general shape of the spectrum, but again, an unphysical local minimum forms between the main peaks. Mishchenko's method is the only one that correctly reproduces the very flat spectrum between $E = 1$ and $E = 2$ even for high noise. This feature is consistent with the behavior observed in the top left panel of Fig. 2.

In general Mishchenko's method performs the best across various levels of noise, while NNLS leads to the poorest continuation for this function. As expected, the quality of the continuation obtained with Padé worsens quickly when
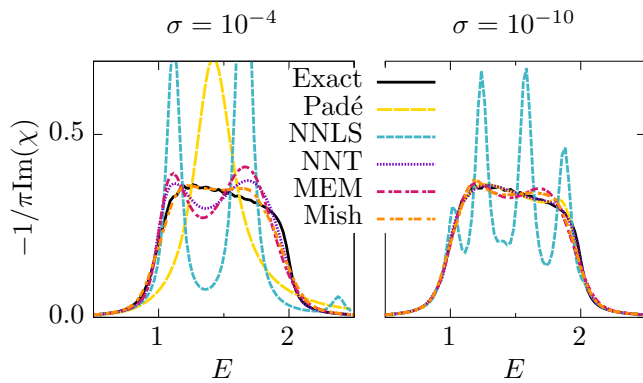
FIG. 5. Spectrum at the M point for the Lindhard function of the doped tight-binding model reported in Fig. 4. Data for two noise levels, $\sigma = 10^{-4}$ and $\sigma = 10^{-10}$, are shown.

σ is increased. Interestingly, NNT seems to perform slightly better for higher noise. For very low noise, the features are sharper, and some spurious wiggles occur. These properties influence strongly the applicability of NNT and Padé to realistic problems where the exact function is unknown. The MEM is rather robust to external noise, but for high-precision data its quality seems inferior to both Mishchenko's method and Padé.

C. Band gap model

The final model function is useful to investigate realistic features that are not close to zero energy. We again use a tight-binding model on a square lattice with nearest-neighbor hopping. Different from the previous test, we now consider two nondegenerate bands. Again, we take $4t = 1$ for the half bandwidth. The two bands are shifted with respect to each other by an energy $E_{\text{shift}} = 3$. Since the half bandwidth of both bands is 1, there is an indirect gap of $E_{\text{gap}} = 1$. In Eq. (12), this would correspond to two bands with shifted dispersions $\epsilon_{\mathbf{p}}$ and $\epsilon'_{\mathbf{p}} = \epsilon_{\mathbf{p}} + 3$. We consider both excitations within one band as well as between the two bands. The chemical potential $\mu = 1.7$ lies in the gap between the two bands. This model could correspond to the presence of a magnetic field that splits the dispersion of the electrons with spins up and down. The electronic states with spin down have energy $E_{\downarrow} \in [-1, 1]$; that is, they are all below the Fermi energy. The electronic states with spin up, on the other hand, have energy $E_{\uparrow} \in [2, 4]$; that is, they are all above the Fermi energy. Excitations between the two bands have energies between $E = 1$ and $E = 5$.

In Fig. 6 the analytical continuation of the momentum-resolved Lindhard function for all methods is reported. As above, the panels in the first three columns correspond to data for various levels of noise σ, while the panels in the last column illustrate the exact spectrum on the real axis. At the Γ point, the exact spectrum shows the presence of a single peak at $E = 3$. This essential feature is captured by all methods and for all levels of noise, although the MEM and the NNT method exhibit the usual tendency to broaden the peak. Going from Γ to X, the single peak splits into two distinct peaks, with some spectral weight remaining in the intermediate region. To illustrate this, Fig. 7(a) shows the spectrum half way between

Γ and X. At $\sigma = 10^{-4}$ MEM, NNT, and Mishchenko's method all capture the widening of the peak as we move away from the Γ point. However, only Mishchenko's method captures the gradual splitting into two peaks at this point. MEM and NNT find the splitting closer to X. For smaller noise, i.e., for $\sigma = 10^{-10}$, the MEM is also capable of resolving the two-peak structure at $\frac{1}{2}(\Gamma \to X)$. NNLS, NNT, and Padé fail to offer an acceptable description of the spectrum at this point since they also introduce one or more spurious peaks in between the two physical peaks.

At the M point, the exact spectrum has a sharp peak at $E = 3$ with a very broad base. All methods, except for NNLS, give the correct width but are not very accurate on the symmetric character of the spectrum. This improves for lower noise levels, but the sharp peak in the middle stays slightly broadened, as can be seen in Fig. 7(b). These data show that Mishchenko's method is the only one to offer a fully satisfactory description of the spectrum for $\sigma = 10^{-4}$. It is also worth noticing that although the Padé method does not fail too badly at M, the general trend for $\sigma = 10^{-4}$ observed in Fig. 6 is rather negative and violates symmetry constraints. For even smaller noise, $\sigma = 10^{-10}$, all methods but NNLS give a very good description of the spectrum at the M point.

In summary, for this model function, Mishchenko's method is stable with respect to input noise and clearly performs best of all the methods for $\sigma = 10^{-3}$. However, great care has to be taken concerning the number of updates in the sampling chain, as is investigated in Appendix B. The Padé method performs quite well for high-precision data but fails badly when noise is present. Finally, for intermediate levels of noise, MEM and NNT also lead to reasonable results, although Mishchenko's method is still superior.

D. Hubbard dimer

Until now, we have shown only tests where a Gaussian relative noise was added to the exact Matsubara data. In practical applications, the size of the noise can vary strongly for different Matsubara frequencies, and also the noisy data can be correlated between Matsubara frequencies. To study how this more realistic type of noise affects the continuation procedure, here, we will apply a QMC solver to an exactly solvable system: the Hubbard dimer.

The Hubbard dimer consists of two sites. These sites can be occupied by electrons with spin either up or down. Electrons can move to the other site, with an associated matrix element $-t$, and two electrons with opposite spin on the same site undergo a Coulomb repulsion U, as described with the following Hamiltonian:

$$H = -t \sum_{\substack{\sigma = \uparrow, \downarrow \\ i = 1,2 \\ j \neq i}} c_{i\sigma}^{\dagger} c_{j\sigma} + U \sum_{i=1,2} n_{i\uparrow} n_{i\downarrow}, \quad (14)$$

where i and j are site indices, σ is the spin index, and $c^{(\dagger)}$ and n are the annihilation (creation) and number operators. From this Hamiltonian, one can see that the total number of states of the model is only 16. We consider this model in the grand-canonical ensemble. The exact expressions for the dynamic susceptibility are given in Ref. [49]. For our model function,
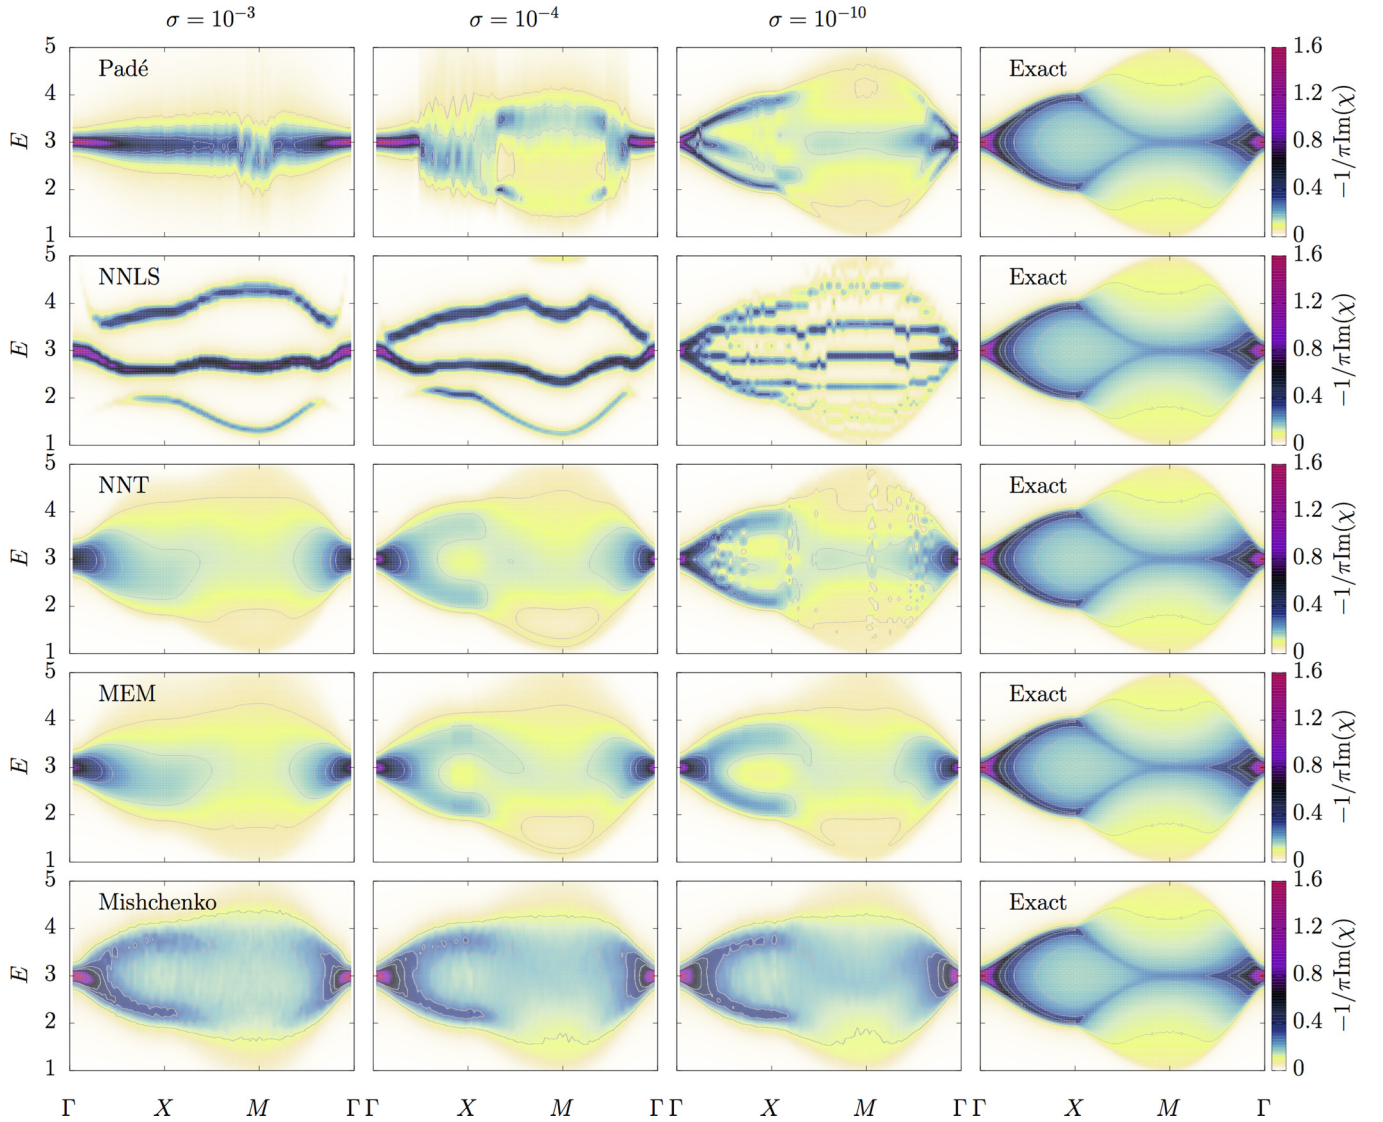
FIG. 6. Spectrum of the Lindhard function for the band gap model ($-\frac{1}{\pi}\text{Im}[\chi(\omega + i\delta)]$) along the Brillouin zone path $\Gamma \rightarrow X \rightarrow M \rightarrow \Gamma$, obtained using Padé, NNLS, NNT, MEM, and Mishchenko's method. The noise level varies among the values $\sigma = 10^{-3}$, $\sigma = 10^{-4}$, and $\sigma = 10^{-10}$. All right panels show the exact spectrum. The color map denotes the intensity of the spectra, and contour lines are added at $-1/\pi\text{Im}(\chi) = 0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6$. The spectra at the M point and at the $\frac{1}{2}(\Gamma \rightarrow X)$ point are reported separately in Fig. 7.

we use $t = 0.25$, $U = 1$, $\beta = 50$ at half filling. We consider the equal-spin susceptibility $4\langle n_{i,\uparrow}n_{i,\uparrow}\rangle$, which is the sum of the local charge and S_z susceptibilities, as found in Ref. [49]. This susceptibility contains two distinct energy scales, i.e., the energy scale of the charge fluctuations and that of the spin fluctuations.

We also solve the Hubbard dimer using determinant QMC [50] in the open-source implementation QUEST [51]. The QMC procedure and the measurements are performed in imaginary time τ using a discrete uniform grid of 1000 τ points. We perform a Fourier transform to Matsubara frequencies using Simpson's rule for integration. This simulation is repeated 300 times, using different random seeds.

As input data for the analytical continuations, we have used both the exact data [49] with added relative Gaussian noise and the average of the QMC simulations that contains realistic QMC noise. The frequency dependence of the standard

deviation of the average QMC data normalized by the average values is shown in Fig. 8(a). The Gaussian noise level added to the exact data is also displayed for comparison. The correlation of the QMC data between different Matsubara frequencies can be seen from the correlation matrix, shown in Fig. 8(b). A purely diagonal matrix means the data at different Matsubara frequencies are not correlated. From Fig. 8(b), however, one clearly sees that our QMC data are rather correlated. To take into account the frequency dependence of the standard deviation and the correlation between the different Matsubara frequencies, the Cholesky decomposition of the inverse covariance matrix [39,44(a)] is used. This requires, however, more Monte Carlo measurements than the considered Matsubara frequencies. It is not always possible to satisfy this requirement. Therefore it is also convenient to perform the analytical continuation for a covariance matrix that is set to identity. This corresponds to assuming no correlation between
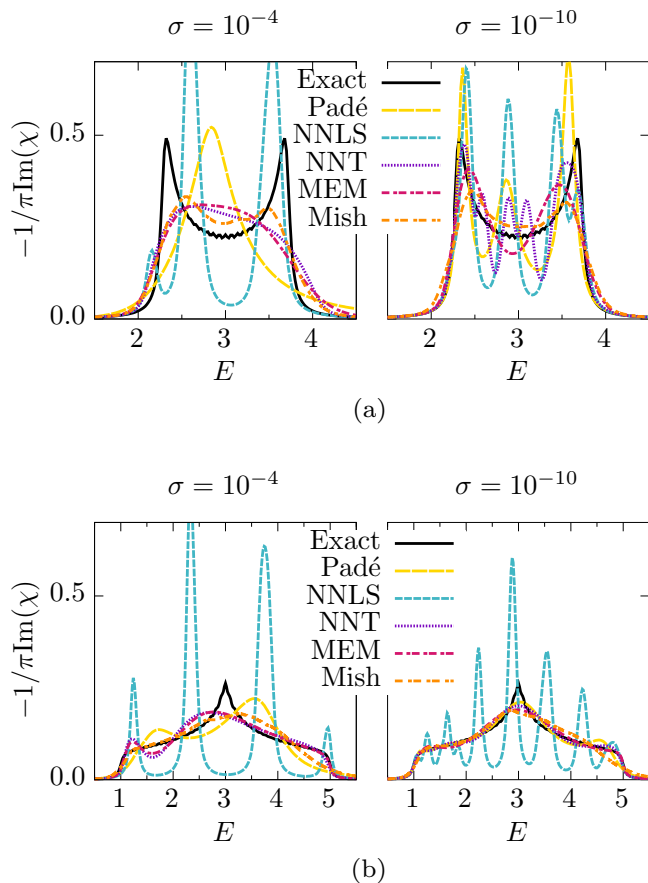
FIG. 7. Spectrum at (a) the $\frac{1}{2}(\Gamma \rightarrow X)$ and (b) M points of the Lindhard function for the band gap model in Fig. 6 for two noise levels. $\sigma = 10^{-4}$ and $\sigma = 10^{-10}$.
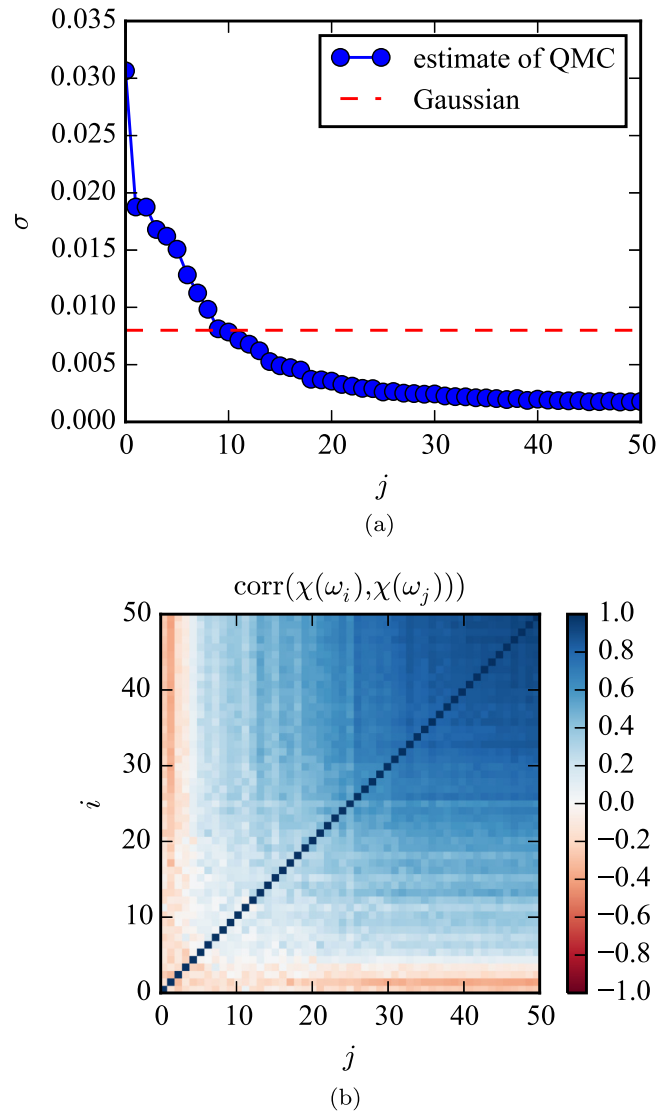

FIG. 8. The noise characteristics of the input data for the analytical continuations for the Hubbard dimer. (a) The standard deviation of the average QMC data between different simulations normalized by the average values (blue) compared to the relative Gaussian noise that is chosen to have $\sigma = 8 \times 10^{-3}$ (red). (b) Correlation matrix (normalized covariance matrix) of the QMC data. The color coding denotes the amount of correlation between two Matsubara frequencies i, j.

data at different Matsubara frequencies and the same standard deviation at all Matsubara frequencies.

Implementing the corrections due to the covariance matrix is straightforward for NNLS, NNT, and MEM. In Mishchenko's method, it is a bit difficult to combine this correction with the original relative deviation measure [43] and the expression in Eq. (8). Therefore we instead use a nonrelative least-squares deviation measure together with the expression in Eq. (2) for this test model. Concerning the Padé method, instead, implementing a correction due to the covariance matrix is even more complex. A procedure similar to what is done for the other methods leads to an impracticable nonlinear system of equations, at least in Beach's formulation [41,47]. Therefore we have not considered the Padé method for this analysis. We plan to address this issue more in detail in future works.

In Fig. 9 we present the spectra obtained for the two sets of input data, i.e., the exact data with Gaussian noise and the QMC data. Let us first analyze the continuation of the QMC data. Without including the covariances, all methods that are able to resolve the high-energy peak locate it above the exact solution. Taking into account the covariance matrix shifts the high-energy peak down in energy, closer to the exact position. The low-energy peak is slightly broadened by considering the covariances, but these differences do not

seem very significant. We can then focus on the continuation of the input data with Gaussian noise. The Gaussian noise considered corresponds to $\sigma = 8 \times 10^{-3}$, which is slightly bigger than the relative σ for the QMC data averaged over the Matsubara points. For this model noise, the obtained spectra are very similar to those obtained for the QMC input data when the covariances are considered. Focusing on the various methods for analytical continuation, we confirm the features found in the previous sections. The NNLS method and Mishchenko's method are the only methods that resolve both the low-energy and high-energy peaks. For the NNT and MEM methods, the high-energy peak appears only as a very broad shoulder. The covariance matrix increases the accuracy
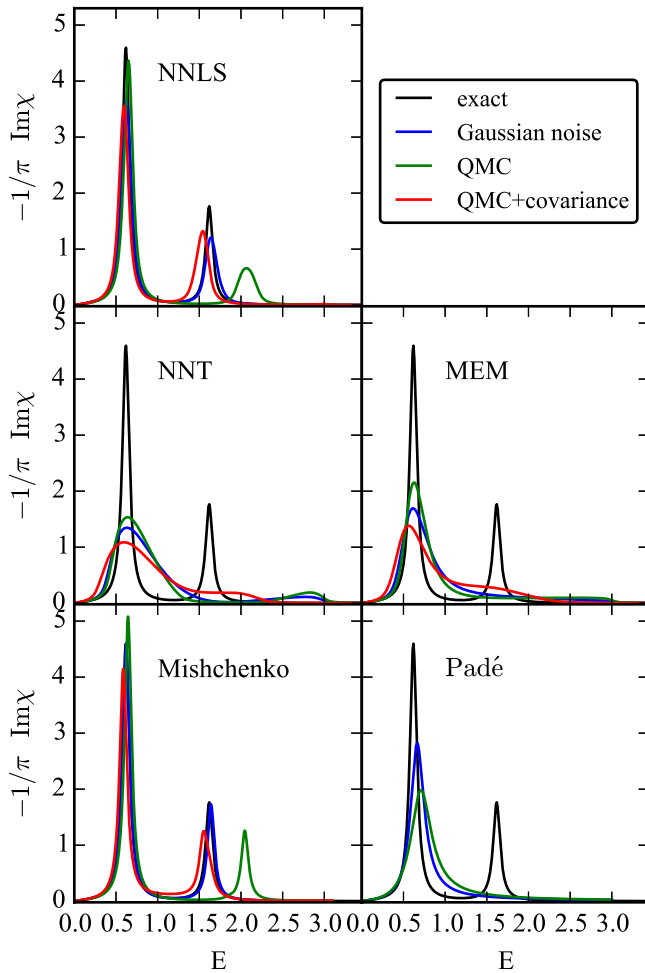
FIG. 9. Spectra for the Hubbard dimer obtained with different methods of analytic continuation. The blue spectra are obtained from the exact data on the Matsubara axis with added relative Gaussian noise $\sigma = 8 \times 10^{-3}$. The red and green spectra are obtained from QMC data, where for the red spectra we used the Cholesky decomposition of the inverse covariance matrix to consider the frequency-dependent standard deviation and the correlation between different input data points [39,44(a)]. The exact solution is given in black.

of the continuation and gives higher spectral weight around the correct position of the high-energy peak. The high-energy peak is not captured using Gaussian noise data or QMC data for the Padé method. It is not trivial to treat covariances within the Padé method. Including the covariance matrix in a way similar to what is done for the other methods leads to an impracticable nonlinear system of equations, at least in Beach's formulation [41,47]. Therefore we have not considered it for this analysis. Nevertheless, we plan to address this issue in more detail in future works.

The previous analysis shows that considering a more realistic noise directly from QMC simulations does not lead to different conclusions than working with a model Gaussian noise. The positive and negative features of the different methods for analytical continuation seem not to depend on the different inputs, although the obtained spectra show some

differences, as expected. Overall, this more realistic test case reinforces the conclusions drawn in the previous sections.

E. Computational time

As a final note, we want to make a brief comment about the computational effort required to perform the analytical continuation of these functions. Momentum-resolved calculations, such as the Lindhard function spectra for the tight-binding model shown above, typically involve linear dimensions of approximately 100 sites. Even considering a minimal effort, analytical continuation is required on a path through the high-symmetry points of the lattice, involving about a hundred analytical continuations. Additionally, the analytical continuation should be done several times with different continuation parameters to ascertain its stability. Therefore, the computational time required to perform a single continuation becomes a factor of practical importance.

It is not simple to compare the fundamental computational effort required by the five methods investigated here because the employed codes have different degrees of optimization and parallelization. Nevertheless, we can still provide some estimates. The fastest methods are undoubtedly NNLS and NNT, which require less than a second per continuation. The MEM and Padé are slower and usually require from several to a hundred seconds per continuation, depending on the parameters used. Mishchenko's method is by far the most demanding method among those tested here. To obtain properly converged results one has to perform the stochastic sampling for a time ranging from a few hours to weeks. The effective time can be significantly reduced by using a parallelized code, like we do in the present work. However, parallelization is easy to implement only over global updates in the Monte Carlo chain, while local updates still have to be handled serially. This limits the scalability of the code, and therefore the reduction in computational time, to about a hundred CPUs. Even in this case, we must notice that it is not always possible to dedicate such a large number of resources to the analytical continuation problem. This sets a limit on the applicability of Mishchenko's method in the current research. We are currently exploring an alternative formulation of Mishchenko's method based on a graphical processing unit (GPU) [52].

VI. CONCLUSIONS

We have presented results for the analytical continuation of bosonic functions using five different methods and four realistic test cases, corresponding to different scenarios of physical relevance for strongly correlated systems. An overview of the major features emphasized in the previous section is given in Table I. The first conclusion of our work is that, at the moment, for low input accuracy, none of the existing methods is good enough to be able to continue all different types of two-particle quantities from the Matsubara frequencies to real energies. Nevertheless, a combination of several methods can give enough information to reconstruct the physical picture behind the analytical continuation problem. Table I can offer useful guidelines to perform this type of analysis.

TABLE I. Summary of the test performance.

	σ	Padé	NNLS	NNT	MEM	Mishchenko
Two very close poles	10^{-2}	One peak, position wrong	Broadened, positions wrong	Broadened, positions wrong	One peak, position wrong	Broadened, positions wrong
	10^{-4}	Excellent	Good	Good	Good	Good
Two close poles	8×10^{-3}	One peak	Very good	Broadened, position wrong	One peak	Excellent
	QMC		Good	Broadened	Broadened	Good
Two separated poles	10^{-2}	One peak	Good	Good, broadened	One peak	Good, broadened
	10^{-4}	Excellent	Good	Good	Good, broadened	Excellent
One sharp feature	10^{-3}	Good	Spurious features	Broadened	Broadened	Good
	10^{-10}	Excellent	Spurious features	Broadened, unstable	Broadened	Good
Broad, flat feature	10^{-3}	Too sharp	Two sharp peaks	Good, unphysical minima	Good, unphysical minima	Good
	10^{-10}	Good	Spurious features	Good	Good	Good
Two peaks in a broad plateau	10^{-3}	One peak	Wrong positions plus additional peak	One broad peak	One broad peak	Good, broadened
	10^{-10}	Additional peak	Spurious peaks	Spurious peaks	Good, broadened	Good, broadened
One peak in broad plateau	10^{-3}	One peak	Spurious peaks	Plateau	Broad peak plus plateau	Chopped triangle
	10^{-10}	Good	Spurious peaks	Good	Good	Good

Conversely, for high-precision input data, our results are encouraging. Among all methods, the Padé averaging scheme [47] is the one that performs best when the numerical noise is absent. For complicated spectra, such as that of the band gap model, even Padé cannot completely resolve the spectrum for $\sigma = 10^{-10}$, and it is indeed not even clear that it performs better than Mishchenko's method due to small violations of the symmetry characterizing the exact spectrum. In general, Mishchenko's method seems to be the most robust approach across various levels of noise. As a matter of fact, its results do not seem very affected by the different noise levels for all the tested functions. The accuracy of Mishchenko's method is also good, although care is needed when choosing the numerical parameters of the stochastic sampling. For optimal parameter values, the accuracy of Mishchenko's method is definitely good.

Overall, a single sharp peak close to zero energy is resolved by all methods even for a relatively high noise. The MEM, NNT method, and Mishchenko's method are able to resolve broad features. However, these methods also have the tendency to smear several pronounced peaks, sometimes into a single broad feature. In Mishchenko's method, this problem is related to the number of local updates performed in the sampling. As shown in Appendix B, increasing the number of local updates leads to more peaked spectra, but those are not always in better agreement with the exact results. In general, the tendency to smear pronounced features into broad peaks and, at the same time, to create spurious peaks makes it difficult to know if a calculated broad spectral feature really is a broad feature or a combination of several peaks. As an example of this, the MEM continuation of the two-pole function for $E_2 = 1.2$ at $\sigma = 10^{-4}$ (Fig. 2) looks similar to the MEM continuation of the two-dimensional doped tight-binding model at the M point for the same noise level (Fig. 5), although the exact functions are very different.

The very broad modes that appear, e.g., around the M point of the doped two-dimensional tight binding model, are difficult to resolve correctly. Functions with two distinct poles, on the other hand, are much easier. NNLS, in particular, performs well for the two-pole model, whereas it performs very poorly for more complicated spectra. This shows the importance of varied, realistic test functions for assessing the potential of each continuation method.

As the results for the Hubbard dimer show, considering realistic data directly from QMC calculations does not seem to affect the previous conclusions. NNT and MEM give smeared-out features, but NNLS and Mishchenko's method perform very well. The Padé method captures only the low-energy peak, and this is not surprising, considering the high level of noise. In general, for both types of noise investigated, the performance of the analytical continuation methods is in line with that obtained for the two-pole model. This is also expected since the exact spectra of the Hubbard dimer and the two-pole model are rather similar. Including a correction for the covariance between the input data from QMC does not change these conclusions but does lead to an improvement of the quality of the analytical continuation.

Finally, we compare our work to previous studies. Two previous works [12,27] found that Padé generally performs poorer than the other methods. Huang *et al.* [12] found that MEM outperforms Padé for a spectrum generated from a combination of two Gaussian distributions. In particular, their results show that the Padé method sometimes merges two peaks, just as we have seen in the two-pole spectra. In Ref. [27], the authors test several methods of analytical continuation on the optical conductivity for relatively high noise levels ($\sigma = 10^{-2}$ to 10^{-3}). They find that all methods (Padé, singular-value decomposition, sampling, and MEM) perform similarly but that the Padé method generally gives slightly less accurate results and sometimes finds unphysical continuations. However, in both studies, the averaging scheme for Padé [47] was not used. This averaging scheme solves the problem of unphysical continuations and should improve the

overall performance of the Padé method. We would also like to stress two more fundamental conclusions of this study. First, we want to emphasize the importance of varied, realistic test functions in the assessment of continuation algorithms. The dependence of the performance of each method on the type of function investigated is a clear indication that no analysis can be meaningful without exploring a proper number of test cases. Second, we want to stress that our findings fully support what was already claimed in Ref. [27], namely, although no universal tool is currently available, better insight into the continuation problem can be obtained by using a combined approach of several methods.

APPENDIX A: IMPOSING MIRROR SYMMETRY

The MEM, the NNT, and the NNLS method conserve the symmetry of the spectra $\rho(E) = -\rho(-E)$ since they use the symmetric definition of the Hilbert transform, Eq. (2). The Padé method does not have such symmetry imposed. For fermionic one-particle Green's functions and self-energies with zero imaginary part at the origin of the complex plane, it is shown [47] that the continuation improves by taking a few negative Matsubara frequencies into account to impose the mirror symmetry. Such a procedure, with a single negative Matsubara frequency, has also been used for the optical conductivity [27]. We have tried the same scheme for the Padé continuation of the bosonic functions in this paper, but in most cases imposing mirror symmetry actually gave a worse continuation. In particular, when the exact spectrum has a broad, rather flat spectrum, such as at the M point of the doped tight-binding model, imposing the symmetry can lead to spurious sharp features. Broad modes in the spectrum are related to branch cuts in the complex function $\chi(z)$. The Padé approximant uses a finite number of poles to simulate $\chi(z)$. When the mirror symmetry is imposed on the Padé approximant, the poles move closer to the real axis, which results in sharper spectral features. Hence broad features are harder to accurately simulate by imposing mirror symmetry.

APPENDIX B: UPDATES FOR MISHCHENKO'S SAMPLING METHOD

The stochastic sampling method by Mishchenko *et al.* [43] consists of many iterative steps. In every step, the configuration
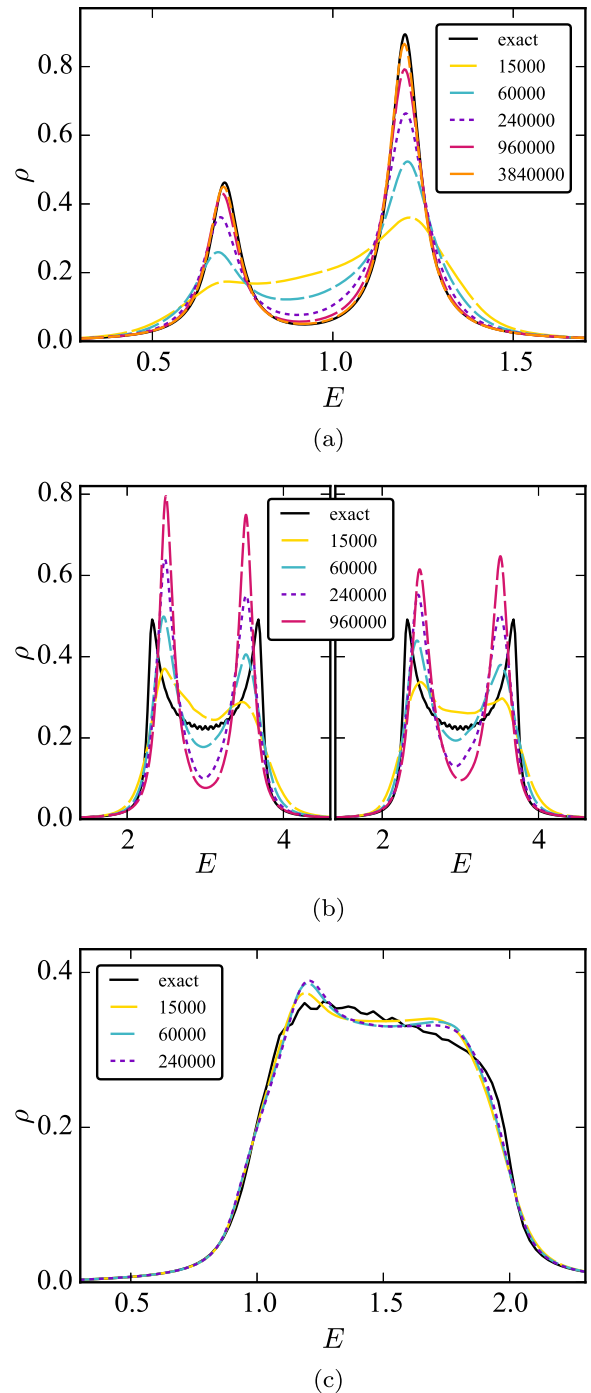

FIG. 10. The dependence on the number of iterations (see legend) performed with the Mishchenko method. (a) The two-pole function in the presence of input noise with $\sigma = 10^{-4}$. (b) The band gap model at the k point $\frac{1}{2}(\Gamma \to X)$ for two different noise levels of $\sigma = 10^{-3}$ (left) and $\sigma = 10^{-4}$ (right). (c) The doped tight-binding model at the M point and Matsubara noise level $\sigma = 10^{-4}$.

of the rectangles is altered. In the main text, results with 15 000 iterations are presented, which is a suitable setup to obtain a reasonable output in a decent computational time (30 CPU hours using 128 CPUs). However, it is important to check how the spectral functions evolve when varying the number of local updates. In Fig. 10 we increase that number by a factor
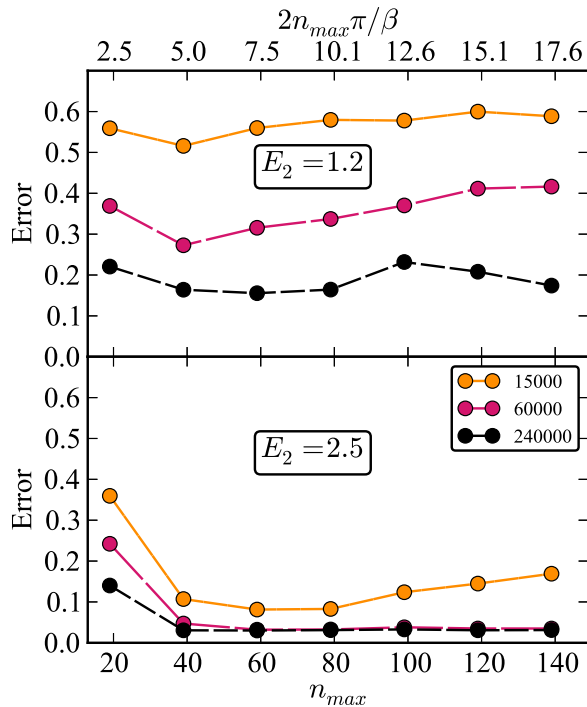
FIG. 11. Integrated real-axis error for the two-pole model as a function of the highest Matsubara point index n_{max} used in Mishchenko's method. Different numbers of local iterations are investigated and are indicated by the legend. Top: $a_1 = 0.1$, $a_2 = 0.335663$, $E_1 = 0.7$, and $E_2 = 1.2$; bottom: $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$, and $E_2 = 2.5$, as defined in Eq. (9). Matsubara noise level $\sigma = 10^{-4}$ is used.

the number of iterations improves the resulting spectrum [see Fig. 10(a)]. However, for a more complicated spectrum, such as the band gap model at the k point $\frac{1}{2}(\Gamma \rightarrow X)$ [see Fig. 7(a)], increasing the number of iterations does not lead to an improvement [see Fig. 10(b)]. This issue is particularly pronounced when the input noise is big. The tendency to produce too sharp features resembles the overfitting occurring in the NNLS method. For the doped tight-binding model at the M point (see the left panel in Fig. 5) the spectra converge already for 15 000 iterations [see Fig. 10(c)]. In summary, the number of iterations giving the best spectral function is highly dependent on the function under investigation and the noise level. With an increasing number of iterations, the computational cost increases as well as the risk of overfitting to the input Matsubara data. Hence care should be taken when choosing the number of iterations. A possible solution could be to stop the local updates when the deviation of the backwards-transformed output data from the input data has a magnitude which is comparable to the standard deviation of the input data.

Further, it is interesting to repeat the analysis of the accuracy of the continuation with respect to the number of Matsubara points (Fig. 3) while varying the number of local updates. These results are illustrated in Fig. 11. While increasing the number of local updates leads to a better continuation, the general trend reported in Fig. 11 is very similar to Fig. 3. Increasing the number of Matsubara frequencies, the quality of the continuation reaches saturation very quickly and with a threshold that does not seem to depend on the number of local updates. *A posteriori*, this shows that the approach used in Fig. 3, i.e., keeping fixed the number of local updates while increasing the number of Matsubara points, is justifiable.

of 4 several times and show the resulting spectral function. For the two-pole model (see Fig. 2) it is clear that increasing

[1] V. Anisimov and Y. Izyumov, *Electronic Structure of Strongly Correlated Materials*, edited by M. Cardona, Springer Series in Solid-State Sciences (Springer, Berlin, 2010).

[2] W. Metzner and D. Vollhardt, Phys. Rev. Lett. **62**, 324 (1989).

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

[4] A. I. Lichtenstein and M. I. Katsnelson, Phys. Rev. B **57**, 6884 (1998).

[5] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Rev. Mod. Phys. **78**, 865 (2006).

[6] J. W. Negele and H. Orland, *Quantum Many-Particle Systems* (Addison-Wesley, Boston, 1988).

[7] G. D. Mahan, *Many Particle Physics*, 3rd ed. (Plenum, New York, 2000).

[8] L. Boehnke and F. Lechermann, Phys. Rev. B **85**, 115128 (2012).

[9] T. Ayral, P. Werner, and S. Biermann, Phys. Rev. Lett. **109**, 226401 (2012).

[10] P. Hansmann, T. Ayral, L. Vaugier, P. Werner, and S. Biermann, Phys. Rev. Lett. **110**, 166401 (2013).

[11] T. Ayral, S. Biermann, and P. Werner, Phys. Rev. B **87**, 125149 (2013).

[12] L. Huang, T. Ayral, S. Biermann, and P. Werner, Phys. Rev. B **90**, 195114 (2014).

[13] E. G. C. P. van Loon, H. Hafermann, A. I. Lichtenstein, A. N. Rubtsov, and M. I. Katsnelson, Phys. Rev. Lett. **113**, 246407 (2014).

[14] H. Hafermann, E. G. C. P. van Loon, M. I. Katsnelson, A. I. Lichtenstein, and O. Parcollet, Phys. Rev. B **90**, 235105 (2014).

[15] A. Galler, C. Taranto, M. Wallerberger, M. Kaltak, G. Kresse, G. Sangiovanni, A. Toschi, and K. Held, Phys. Rev. B **92**, 205132 (2015).

[16] A. S. Mishchenko, N. Nagaosa, G. De Filippis, A. de Candia, and V. Cataudella, Phys. Rev. Lett. **114**, 146401 (2015).

[17] A. A. Katanin, A. I. Poteryaev, A. V. Efremov, A. O. Shorikov, S. L. Skornyakov, M. A. Korotin, and V. I. Anisimov, Phys. Rev. B **81**, 045117 (2010).

[18] L. Boehnke, H. Hafermann, M. Ferrero, F. Lechermann, and O. Parcollet, Phys. Rev. B **84**, 075145 (2011).

[19] Q. Si and J. L. Smith, Phys. Rev. Lett. **77**, 3391 (1996).

[20] H. Kajueter, Ph.D. thesis, Rutgers University, 1996.

[21] J. L. Smith and Q. Si, Phys. Rev. B **61**, 5184 (2000).

[22] R. Chitra and G. Kotliar, Phys. Rev. Lett. **84**, 3678 (2000).

[23] R. Chitra and G. Kotliar, Phys. Rev. B **63**, 115110 (2001).

[24] A. N. Rubtsov, M. I. Katsnelson, and A. I. Lichtenstein, Ann. Phys. (N.Y.) **327**, 1320 (2012).

[25] E. G. C. P. van Loon, A. I. Lichtenstein, M. I. Katsnelson, O. Parcollet, and H. Hafermann, Phys. Rev. B **90**, 235135 (2014).

[26] E. G. C. P. van Loon, M. I. Katsnelson, and M. Lemeshko, Phys. Rev. B **92**, 081106 (2015).

[27] O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, Phys. Rev. B **82**, 165125 (2010).

[28] H. Matsui, A. S. Mishchenko, and T. Hasegawa, Phys. Rev. Lett. **104**, 056602 (2010).

[29] A. S. Mishchenko, H. Matsui, and T. Hasegawa, Phys. Rev. B **85**, 085211 (2012).

[30] O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, Phys. Rev. B **81**, 155107 (2010).

[31] M. Casula, A. Rubtsov, and S. Biermann, Phys. Rev. B **85**, 035115 (2012).

[32] R. N. Silver, D. S. Sivia, and J. E. Gubernatis, Phys. Rev. B **41**, 2380 (1990).

[33] M. Jarrell and J. Gubernatis, Phys. Rep. **269**, 133 (1996).

[34] R. Bryan, Eur. Biophys. J. **18**, 165 (1990).

[35] J. E. Gubernatis, M. Jarrell, R. N. Silver, and D. S. Sivia, Phys. Rev. B **44**, 6011 (1991).

[36] S. Fuchs, T. Pruschke, and M. Jarrell, Phys. Rev. E **81**, 056701 (2010).

[37] A. Dirks, P. Werner, M. Jarrell, and T. Pruschke, Phys. Rev. E **82**, 026701 (2010).

[38] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (Society for Industrial and Applied Mathematics, Philadelphia, 1995).

[39] K. Ghanem, MS thesis, RWTH Aachen University, 2013.

[40] P. C. Hansen, in *Computational Inverse Problems in Electro-cardiology*, edited by P. Johnston, Advances in Computational Bioengineering (WIT Press, Southampton, 2000), pp. 119–142.

[41] K. S. D. Beach, R. J. Gooding, and F. Marsiglio, Phys. Rev. B **61**, 5147 (2000).

[42] H. J. Vidberg and J. W. Serene, J. Low Temp. Phys. **29**, 179 (1977).

[43] A. S. Mishchenko, N. V. Prokof'ev, A. Sakamoto, and B. V. Svistunov, Phys. Rev. B **62**, 6317 (2000).

[44] M. Jarrell, in *Correlated Electrons: From Models to Materials*, edited by E. Pavarini, E. Koch, F. Anders, and M. Jarrell (Forschungszentrum Julich GmbH, Julich, 2012), pp. 365–394; A. S. Mishchenko, *ibid.*, pp. 395–422.

[45] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).

[46] L. D. Landau and E. M. Lifshitz, *Statistical Physics* (Pergamon Press, Oxford, 1980).

[47] J. Schött, I. L. M. Locht, E. Lundin, O. Grånäs, O. Eriksson, and I. Di Marco, Phys. Rev. B **93**, 075104 (2016).

[48] D. Pines and P. Nozières, *The Theory of Quantum Liquids: Normal Fermi Liquids* (Benjamin, Philadelphia, 1966).

[49] A. Avella, F. Mancini, and T. Saikawa, Eur. Phys. J. B **36**, 445 (2003).

[50] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, Phys. Rev. D **24**, 2278 (1981).

[51] A. Tomas, C.-C. Chang, Z.-J. Bai, and R. Scalettar, Quantum Electron Simulation Toolbox QUEST 1.3.0, http://quest.ucdavis.edu/.

[52] J. Nordström, J. Schött, I. L. M. Locht, and I. Di Marco, SoftwareX **5**, 178 (2016).