

Extending the average spectrum method: Grid point sampling and density averaging

Khaldoon Ghanem^{1,2} and Erik Koch^{1,3}

¹Jülich Supercomputer Centre, Forschungszentrum Jülich, 52425 Jülich, Germany

²Max Planck Institut für Festkörperforschung, 70569 Stuttgart, Germany

³JARA High-Performance Computing, 52425 Jülich, Germany



(Received 3 April 2020; revised 14 June 2020; accepted 29 June 2020; published 6 July 2020)

Analytic continuation of imaginary time or frequency data to the real axis is a crucial step in extracting dynamical properties from quantum Monte Carlo simulations. The average spectrum method provides an elegant solution by integrating over all nonnegative spectra weighted by how well they fit the data. In a recent paper, we found that discretizing the functional integral, as in Feynman's path-integrals, does not have a well-defined continuum limit. Instead, the limit depends on the discretization grid whose choice may strongly bias the results. In this paper, we demonstrate that sampling the grid points, instead of keeping them fixed, also changes the functional integral limit and rather helps to overcome the bias considerably. We provide an efficient algorithm for doing the sampling and show how the density of the grid points acts now as a default model with a significantly reduced biasing effect. The remaining bias depends mainly on the width of the grid density, so we go one step further and average also over densities of different widths. For a certain class of densities, including Gaussian and exponential ones, this width averaging can be done analytically, eliminating the need to specify this parameter without introducing any computational overhead.

DOI: [10.1103/PhysRevB.102.035114](https://doi.org/10.1103/PhysRevB.102.035114)

I. INTRODUCTION

Quantum Monte Carlo (QMC) simulations have become an indispensable tool for studying quantum many-body systems. They often compute Green's or correlation functions on the imaginary-time axis or Matsubara frequencies, which then need to be analytically continued to the real axis to extract dynamical information about the system of interest. One important example of analytic continuation is obtaining the spectral function $A(\omega)$ at real frequencies from finite-temperature Green's function values $\mathcal{G}(\tau)$ at imaginary times $\tau \in [0, \beta]$, where the two functions are related by

$$\mathcal{G}(\tau) = -\frac{1}{2\pi} \int d\omega \frac{e^{-\tau\omega}}{1 + e^{-\beta\omega}} A(\omega), \quad (1)$$

and $\beta = 1/T$ is the inverse temperature. Another example is determining the optical conductivity spectrum $\sigma(\omega)$ from the current-current correlation function $\Pi(i\omega_m)$ evaluated at bosonic Matsubara frequencies $\omega_m = 2m\pi/\beta$. The relation between the two reads

$$\Pi(i\omega_m) = \frac{2}{\pi} \int_0^\infty d\omega \frac{\omega^2}{\omega_m^2 + \omega^2} \sigma(\omega). \quad (2)$$

In general, the analytic continuation problem reduces to solving an integral equation. The difficulty is that, in the

presence of noise, this is an inherently ill-posed problem. When evaluating the data on the imaginary axis, oscillations and sharp features in the spectrum get smoothed and noise gets damped due to the integration. This makes the inverse problem of reconstructing the details of the spectrum extremely challenging. Without regularization, small noise on the data leads to catastrophically large errors on the best data-fitting spectrum.

There are different approaches to tackle this problem including the maximum entropy method [1–5], the average spectrum method [6–13], Pade approximation techniques [14–17], and some recent machine-learning based approaches [18–21]. The most commonly used approach is the maximum entropy method (MaxEnt), which is rooted in Bayesian inference. It tries to find a spectrum by balancing the fit to the data and the entropy relative to some default model. This entropy term acts as a regularization that penalizes deviations from the featureless default model and thus suppresses rapid oscillations that otherwise would dominate the solution.

The average spectrum method (ASM) is an alternative Bayesian approach with the following premise: Since the data are not exact, all spectra that fit the data equally well, up to the noise level on that data, should be considered equally. As a result, ASM integrates over all admissible spectra weighted by how well they fit the data. The average spectrum method makes no assumptions about the smoothness of the spectrum and any regularization comes from averaging only, which is expected to smooth out details not supported by the data: Larger noise leads to more smoothing.

We continue here our work published in Ref. [13], which in the following will be referred to as ASM1. In ASM1, we showed that a naive discretization of the functional integral involved in the average spectrum method does not produce

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

unique results. The results are biased by the discretization grid on which the spectrum is represented. We constructed the grids explicitly by mapping the real-frequency axis to the unit interval using a density function and then discretizing this interval using a regular grid. We found that the density function of the grid points plays the role of a default model while the number of grid points acts as a regularization parameter. We proposed a practical recipe for choosing a reliable grid by comparing the results of different grid densities and choosing the one with the least dependence on the number of grid points.

In the present paper, we generalize the average spectrum method by releasing the grid points and sampling their position from a prior grid density. Although, we still need to specify a grid density and a number of points, we show that the bias is significantly reduced and we observe that it depends mainly on the width rather than the shape of its density. The proper width can be chosen according to the same type of recipe used earlier for the fixed grid. We can, however, go further and extend the method to sample over a whole class of grid densities of variable widths. The method is then able to automatically relocate the grid points and concentrate them into the important region of the real-frequency axis. Test cases show that this width-sampling method gives good results resolving the features of the spectrum without the need for fine-tuning the grid.

The structure of the paper is as follows. We first review the average spectrum method and reformulate it in a way that facilitates subsequent developments. We then introduce released-grid ASM and explain an efficient algorithm for sampling the grid points. We examine the effect of this sampling on the grid bias using a problematic test case discussed earlier in ASM1. To illustrate the remaining width-dependence, we then consider another test case and use it to motivate the width-sampling ASM derived afterward. Finally, we discuss the role of the grid size as a regularization parameter and provide a simple recipe for choosing a good number of grid points. We end with a summary and a discussion.

II. AVERAGE SPECTRUM METHOD

A. Background

Mathematically, the analytic continuation problem can be formulated as a Fredholm integral equation of the first kind

$$g(y) = \int dx K(y, x) f(x), \quad (3)$$

where the left-hand side $g(y)$ represents QMC data, while the integral kernel $K(y, x)$ is a continuous function known analytically. The goal is estimating the spectrum $f(x)$, an integrable nonnegative function.

QMC provides noisy and incomplete samples of the data evaluated at a finite number of y -coordinates. Using the central limit theorem, one can assume that the exact data vector \mathbf{g} has a Gaussian distribution with mean equal to the sample mean $\bar{\mathbf{g}}$ and covariance matrix equal to the sample covariance matrix \mathbf{C} . Then the likelihood of a spectrum f being the exact one is proportional to

$$\exp\left(-\frac{1}{2}(\bar{\mathbf{g}} - \mathbf{g}[f])^\dagger \mathbf{C}^{-1}(\bar{\mathbf{g}} - \mathbf{g}[f])\right) =: e^{-\chi^2[f]/2}, \quad (4)$$

where $\mathbf{g}[f]$ is the data corresponding to the spectrum f and $\chi^2[f]$ is its fit to the measured data. Maximizing the likelihood (or equivalently minimizing the fit χ^2) gives a spectrum dominated by diverging amplitudes that are extremely sensitive to the noise of the data. This ill-posedness constitutes the primary difficulty of the analytic continuation problem.

The average spectrum method uses the likelihood alongside our prior knowledge about the nonnegativity of the spectrum and computes a weighted average over all nonnegative spectra as an estimate of the true spectrum

$$f_{\text{ASM}}(x) \propto \int_{f(x) \geq 0} \mathcal{D}f e^{-\chi^2[f]/2} f(x). \quad (5)$$

Other exact prior knowledge, like sum rules, can be easily incorporated by restricting the averaging to the spectra satisfying them.

Despite the apparent elegance of this functional integral formulation, we found in ASM1 that it is not a well-defined expression because the result depends on how the spectrum $f(x)$ is discretized. In ASM1 the discretization was specified by a grid density function $\rho(x)$ and the total number of grid points N . Using $\rho(x)$, we mapped the domain of the variable x into a unit interval of an auxiliary variable $z(x) := \int_{x_{\min}}^x dx' \rho(x')$, discretized it uniformly with N points, and mapped the points back to x . We showed for such grids that the density $\rho(x)$ plays the role of a default mode, while N plays the role of a regularization parameter.

To demonstrate this role of the grid density $\rho(x)$ and simultaneously introduce some notation, we repeat the argument of appealing to symmetry when the data contains no information other than normalization to unity. In that case, the spectral integrals over different grid intervals $\mathcal{I}_i = [\bar{x}_i, \bar{x}_{i+1}]$ should be equal

$$\bar{f}_i := \int_{\bar{x}_i}^{\bar{x}_{i+1}} dx f(x) = \frac{1}{N}. \quad (6)$$

The widths of these intervals are inversely proportional to the grid density

$$\Delta x_i \approx \frac{1}{N\rho(x_i)}, \quad (7)$$

with $x_i \in \mathcal{I}_i$ being the grid point representing the i th interval. Therefore, in the absence of data, the estimated mean values of the spectrum are equal to the grid density

$$f(x_i) = f_i := \frac{\bar{f}_i}{\Delta x_i} \approx \rho(x_i), \quad (8)$$

justifying calling it a default model.

B. Formalism

To spell out the dependence of the average spectrum method on the discretization grid explicitly, let us parametrize the spectrum by its grid points \mathbf{x} and its integrals over the grid intervals \mathbf{f} , defined in (6). Together they are enough to determine the data (up to a discretization error) without knowing the details of the spectrum inside the grid intervals. This follows from the first mean-value theorem for integrals, which states that for any nonnegative integrable function $f(x)$,

there is a specific point $x_i^* \in [\bar{x}_i, \bar{x}_{i+1}]$ such that

$$\int_{\bar{x}_i}^{\bar{x}_{i+1}} dx K(y, x) f(x) = K(y, x_i^*) \int_{\bar{x}_i}^{\bar{x}_{i+1}} dx f(x). \quad (9)$$

Using this, we can approximate the data as follows:

$$\mathbf{g}[\bar{\mathbf{f}}, \mathbf{x}] \approx \sum_i K(\mathbf{y}; x_i) \bar{f}_i =: \mathbf{K}[\mathbf{x}] \bar{\mathbf{f}}. \quad (10)$$

The approximation comes from using the grid points x_i instead of the unknown optimal points x_i^* , which depend on both the spectrum $f(x)$ and the kernel $K(x, y)$. The approximation error is proportional to the difference between the maximum and minimum values of the kernel $K(x, y)$ inside each interval. Since the kernel is a continuous function of x , this error gets smaller, the smaller the intervals are and the smoother the kernel is. Using a fine-enough grid, the error becomes so small that it is negligible in comparison to the noise on the data.

Contrary to the data, the spectrum cannot be determined entirely from the spectral integrals $\bar{\mathbf{f}}$ and some assumptions about its behavior inside the grid intervals are needed. We may also need to specify how to go from grid points \mathbf{x} to the intervals edges $\bar{\mathbf{x}}$. To stay as general as possible, we encode whatever assumptions we have in the object $f(\bar{\mathbf{f}}, \mathbf{x}; x)$, which maps a set of grid points \mathbf{x} and integrals $\bar{\mathbf{f}}$ into a nonnegative integrable function of the continuous variable x . A simple choice is using delta functions located at the grid points

$$f(\mathbf{x}, \bar{\mathbf{f}}; x) = \sum_i \bar{f}_i \delta(x - x_i), \quad (11)$$

another is a piecewise constant function

$$f(\mathbf{x}, \bar{\mathbf{f}}; x) = \sum_i \frac{\bar{f}_i}{\bar{x}_{i+1} - \bar{x}_i} (\Theta(x - \bar{x}_i) - \Theta(x - \bar{x}_{i+1})). \quad (12)$$

The advantage of the above general formulation is that we do not need to know the exact form of $f(\bar{\mathbf{f}}, \mathbf{x}; x)$ when sampling the spectral functions. It becomes relevant only when averaging them, which will be discussed in Sec. III B.

The functional integral of the average spectrum can now be written as a multidimensional integral over parameterized functions $f(\bar{\mathbf{f}}, \mathbf{x}; x)$ given the mean data vector $\bar{\mathbf{g}}$, the covariance matrix \mathbf{C} and the grid points \mathbf{x} . The averaging is carried over all spectral integrals $\bar{\mathbf{f}}$ on that grid, weighted by their fit

$$f_{\text{ASM}}(\mathbf{x}; x) \propto \int_0^\infty d\bar{\mathbf{f}} e^{-\chi^2[\bar{\mathbf{f}}, \mathbf{x}]/2} f(\bar{\mathbf{f}}, \mathbf{x}; x), \quad (13)$$

where the dependence of the fit χ^2 on the grid points \mathbf{x} is through the discretized kernel matrix \mathbf{K} evaluated at these points

$$\chi^2[\bar{\mathbf{f}}, \mathbf{x}] := (\bar{\mathbf{g}} - \mathbf{K}[\mathbf{x}] \bar{\mathbf{f}})^\dagger \mathbf{C}^{-1} (\bar{\mathbf{g}} - \mathbf{K}[\mathbf{x}] \bar{\mathbf{f}}). \quad (14)$$

III. RELEASING GRID POINTS

The form of Eq. (13) spells out the grid dependence of ASM explicitly and is suggestive of a straightforward extension of the method. Instead of having the grid points fixed at regular intervals based on some density function $\rho(x)$, let us sample them freely from this distribution and average the

results

$$f_{\text{ASM}}(\rho, N; x) \propto \int d\mathbf{x} \prod_{i=1}^N \rho(x_i) \int_0^\infty d\bar{\mathbf{f}} e^{-\frac{1}{2} \chi^2[\bar{\mathbf{f}}, \mathbf{x}]} f(\bar{\mathbf{f}}, \mathbf{x}; x). \quad (15)$$

Although we still need to specify the density function $\rho(x)$ and the number of grid points N as before, we expect the effect on the results to be weaker than in the fixed-grid scenario because the data are now allowed to influence the positions of the grid points during the sampling.

It is worth noting that sampling grid points has been done before in the context of the average spectrum method in Refs. [8, 11]. In these papers, the spectrum was represented as a superposition of delta functions whose both weights and positions were sampled. However, it was implicitly assumed that sampling the positions was a technical detail and that the result would be the same as the typical average spectrum method with fixed positions [6, 7, 10]. This was probably due to the mistaken belief in the existence and uniqueness of the functional integral of Eq. (5). As shown in ASM1, this functional integral does not exist and the result depends on the discretization grid. In Sec. IV we will show that the results, in fact, depend on whether the grid points are sampled and rather improves significantly. But before that, we will describe an efficient algorithm for performing the sampling and how the spectra of different grids are averaged.

A. Sampling algorithm

The multidimensional integral in Eq. (15) is evaluated using a Monte Carlo sampling algorithm. We start from some initial spectrum on an initial grid. In practice, we use the same grid as in the fixed-grid case, i.e., we choose the grid points at regular intervals based on the density $\rho(x)$ as in ASM1. We also use the nonnegative least squares (NNLS) solution on that grid as the initial spectrum [22].

The spectral integrals are then updated on the current grid using the blocked-mode sampling introduced in ASM1. Given the spectral integrals, the grid points are updated one at a time using the Metropolis-Hastings algorithm explained below. All samples of $\bar{\mathbf{f}}$ and \mathbf{x} are stored during the sampling, and the average spectrum $f_{\text{ASM}}(\rho, N; x)$ can be evaluated later at any point x by evaluating each sampled model $f(\bar{\mathbf{f}}, \mathbf{x}; x)$ at that point and averaging the results.

The Metropolis-Hastings algorithm for sampling grid points has the following acceptance ratio

$$r = \frac{e^{-\chi^2(x'_i)/2} \rho(x'_i) q(x'_i \rightarrow x_i)}{e^{-\chi^2(x_i)/2} \rho(x_i) q(x_i \rightarrow x'_i)}, \quad (16)$$

where $q(x_i \rightarrow x'_i)$ is the proposal distribution of moving grid point i from an old position x_i to a new position x'_i . As a proposal distribution, we use a Gaussian approximation of the data factor $e^{-\chi^2/2}$. It is derived by writing the data fit as a function of x'_i and expanding the kernel around x_i . Keeping only terms to second order in x'_i , we obtain

$$\chi^2(x'_i) \approx \mathbf{r}^T \mathbf{r} - 2[\mathbf{r}^T \partial \mathbf{K}_i \bar{\mathbf{f}}_i][x'_i - x_i] + [\bar{\mathbf{f}}_i^2 \partial \mathbf{K}_i^T \partial \mathbf{K}_i - \bar{\mathbf{f}}_i \mathbf{r}^T \partial^2 \mathbf{K}_i][x'_i - x_i]^2, \quad (17)$$

where $\mathbf{r} := \mathbf{g} - \sum_j \mathbf{K}(x_j) \bar{f}_j$ is the old residual vector and $\partial \mathbf{K}_i := \partial_x \mathbf{K}(x_i)$, $\partial^2 \mathbf{K}_i := \partial_x^2 \mathbf{K}(x_i)$ are kernel derivatives. By completing the squares, the data fit can be written in the following suggestive form:

$$\chi^2(x'_i) \approx (x'_i - \mu_\chi)^2 / \sigma_\chi^2 + \text{const.}, \quad (18)$$

where $\mu_\chi := x_i + \sigma_\chi^2 \bar{f}_i^T \partial \mathbf{K}_i$ is the mean of the Gaussian approximation and $\sigma_\chi^{-2} := \bar{f}_i^T \partial \mathbf{K}_i^T \partial \mathbf{K}_i - \bar{f}_i^T \partial^2 \mathbf{K}_i$ is its width. Using this as a proposal probability gives, in general, a high acceptance rate because the data are typically the dominating factor in the acceptance ratio (16). However, when a grid point is far away from zero or its weight is very small, the Gaussian (18) is quite wide and the prior density $\rho(x)$ becomes important. To account for such cases, we combine the data Gaussian with another one, centered at the old position x_i , whose width we choose equal to the width of the prior density. The last width is used by itself in the proposal probability should the data fit have negative curvature so that the approximation (18) fails.

Samples are produced by iteratively sampling all the grid points \mathbf{x} followed by sampling all the spectral integrals $\bar{\mathbf{f}}$. The grid points are sampled one after the other in a random order using the above algorithm. The spectral integrals are sampled using blocked-mode sampling with a random block size. The movement of grid points implies that the kernel matrix is changing, so the singular value decomposition (SVD) of its blocks should be recalculated in each iteration after all the grid points have been updated. This decomposition costs $\mathcal{O}(N_b^2 M)$ where N_b is the block size and M is the data vector size. Since this is computationally expensive, we restrict the block size to a maximum value. For a fixed grid, this leads to more correlation between samples because the global updates using larger blocks are skipped. For a released grid, however, the effect is not as severe because the movement of the grid points compensates for the lack of these global updates. We found that a maximum block size of 32 provides a good balance between the cost of each sample and the correlation between samples. Due to grid points sampling and SVD, the computational cost of a single sample of the average spectrum method using released grid points is larger than that of the fixed grid. However, we found that the samples of the released-grid calculations are much less correlated than those on a fixed-grid so that the total cost is effectively similar. The execution time for any of the results presented later in this paper does not exceed 2 minutes on a typical modern laptop.

B. Binning and averaging

Averaging samples requires evaluating the sampled spectra on some fixed grid. We call this grid the binning grid to distinguish it from the sampled one \mathbf{x} . Let us denote its intervals, the bins, as \mathcal{B}_i . The binning would be different depending on the mapping $f(\bar{\mathbf{f}}, \mathbf{x}; x)$. Assuming the delta functions representation of Eq. (11), the i th bin average is computed as

$$f_i \approx \frac{1}{L} \sum_{k=1}^L \frac{1}{\text{len}(\mathcal{B}_i)} \sum_{x_j^k \in \mathcal{B}_i} \bar{f}_j^k, \quad (19)$$

where the superscript k is indexing the samples and L is the total number of samples.

Alternatively, we can assume a constant value inside each interval \mathcal{I}_j^k of the k th grid sample, as done in Eq. (12). This implies that the corresponding spectral integral \bar{f}_j^k should be split proportionally among the bins that intersects this interval

$$f_i \approx \frac{1}{L} \sum_{k=1}^L \frac{1}{\text{len}(\mathcal{B}_i)} \sum_{j=1}^N \frac{\text{len}(\mathcal{B}_i \cap \mathcal{I}_j^k)}{\text{len}(\mathcal{I}_j^k)} \bar{f}_j^k. \quad (20)$$

This type of binning can be thought of as a linear interpolation of (19) and thus leads to an average with less noise from binning. Nevertheless, whatever binning we use, the averages are similar when using reasonably large grid sizes N . For simplicity, we typically use the delta binning.

Note that the bin size also affects the statistical error of its average. Larger bins contain more sampled grid points and thus have lower fluctuations and less noisy averages. Therefore, in practice we use the following binning, which gives roughly equal error bars across the bins: aggregate all the grid samples and choose the bins such that each bin contains roughly the same number of grid points. When we want to compare to the fixed-grid ASM, we use, however, the fixed grid for binning, making comparisons easier. Also note that the error-bars of a binning grid will not change when doubling the number of grid points and halving the number of samples.

IV. DENSITY DEPENDENCE

To study the effect of releasing the grid points, we choose a test case first introduced in Ref. [23], which we studied further using fixed-grids in ASM1. We want to reconstruct an optical conductivity given by

$$\sigma(\omega) = \frac{1}{1 + (\omega/\Gamma_e)^6} \sum_{p=0, \pm 1} \frac{W_{|p|}}{1 + [(\omega + \text{sgn}(p)\epsilon_{|p|})/\Gamma_{|p|}]^2}, \quad (21)$$

where the sum has three terms: A peak of weight $W_0 = 0.3$ and width $\Gamma_0 = 0.6$ centered at zero, and two peaks of weight $W_1 = 0.2$ and width $\Gamma_1 = 1.2$ centered at $\omega = \pm \epsilon_1 = \pm 3$. All terms are multiplied by a damping factor with $\Gamma_e = 4$ for a faster decay at large frequencies. To get the necessary data for analytic continuation, we compute the imaginary-frequency correlation function $\Pi(i\omega_m)$ analytically using Eq. (2) on the first 60 Matsubara frequencies $\omega_m = 2m\pi/\beta$ with inverse temperature $\beta = 15$. We then add relative Gaussian noise with a standard deviation of 10^{-3} to simulate the noise in real QMC data.

In Fig. 1, we compare ASM results using fixed and released grid points. The same grid densities and numbers of grid points were used in both cases: Uniform grid densities with increasing cutoffs: 8, 16, 32, and 64 and correspondingly increasing number of points $N = 32, 64, 128$, and 256. It is clear that using a fixed uniform grid biases the results significantly, leading to more pronounced spurious features as the cutoff increases. This completely disappears when the grid points are released so that the cutoff has negligible effect on the result.

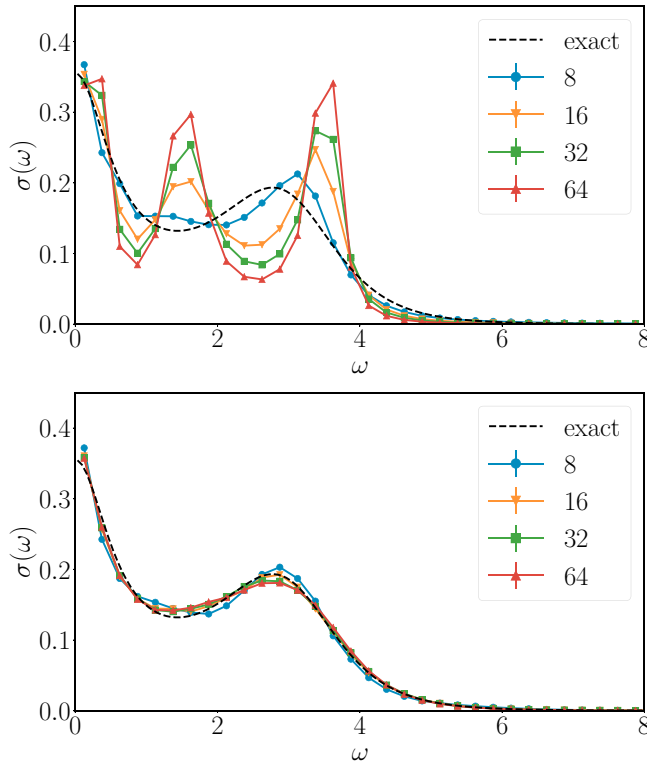


FIG. 1. Optical conductivity $\sigma(\omega)$ obtained using fixed-grid ASM (top) and released-grid ASM (bottom). Uniform grid densities with increasing cutoff ω_{\max} (label) are used. The number of grid points is proportional to the cutoff: $N = 4 \omega_{\max}$. For ease of comparison, the samples of the released calculations are binned and averaged on the grids of the corresponding fixed calculations.




FIG. 2. Optical conductivity $\sigma(\omega)$ obtained using fixed-grid ASM (top) and released-grid ASM (bottom). The following grid densities (labels) are used: Uniform ($\omega_{\max} = 8$), Gaussian ($\alpha = 4$), exponential ($\beta = 3$), and a Lorentzian ($\gamma = 2.5$). The number of grid points is $N = 64$. For ease of comparison, the samples of the released calculations are binned and averaged on the grids of the corresponding fixed calculations.

In Fig. 2, we also show results for different grid densities with comparable widths. Also here, by releasing the grid points, the influence of the shape of the grid density is reduced significantly. This is understandable as now the grid points can move to the region where the spectrum is concentrated, allowing for the data to override the prior information encoded in the grid density.

From this we might conclude that sampling the grid points solves the bias problem. We know, however, that in the absence of data except for a sum rule on the spectrum, averaging spectra on any grid just gives a result proportional to the grid density function $\rho(x)$. Thus, also released-grid ASM must have a default model bias.

To see this bias clearly, we consider a somewhat contrived test case: we seek to recover a Gaussian spectral function of width 0.5 from the Green's function data computed using (1). The data are evaluated on 60 τ -points equally spaced in the interval $[0, \beta]$ with $\beta = 50$. As before, we add relative Gaussian noise with a standard deviation of 10^{-3} . One might think that reconstructing such a featureless Gaussian should be a trivial and boring task. However, it turns out that avoiding spurious features in this setting is more challenging than anticipated. In Fig. 3, we show ASM results using fixed and released grids. We use a Gaussian grid density with different widths: 0.5, 1.0, 2.0. When the default model (grid density) width equals the width of the exact spectrum 0.5, both methods give

perfect results as expected. As the width increases, spurious features start to develop quickly. Although these features are milder for released-grid than for fixed-grid ASM, they are still clearly visible, in contrast with the earlier optical conductivity case. This indicates that the data here are weaker in forcing the grid points to stay in the frequency region of interest. A practical solution for choosing the best width is to use the same criterion as the one employed in choosing the fixed-grid in ASM1: We choose the grid density with the best fit to the data, which in this case would single out the width 0.5. Still, this approach requires considering a number of grid densities $\rho(x)$ of different widths and choosing the best one. In the next section, we describe a method that avoids choosing a specific width altogether. It is worth mentioning that the bias caused by the default model is not unique to ASM. Also MaxEnt, which is known for its smooth results, produces in this case spurious features when the width of the default model is larger than it is supposed to be. They are, however, somewhat milder than those of the two flavors of ASM (see Fig. 4).

V. AVERAGING WIDTH OF GRID DENSITY

Instead of finding the width giving the best fit to the data, we propose here an alternative that is more in the spirit of the average spectrum method: Instead of choosing some




FIG. 3. Spectral function $A(\omega)$ obtained using fixed-grid ASM (top) and released-grid ASM (bottom). Gaussian grid densities with increasing width (label) are used. The number of grid points is $N = 512$. For ease of comparison, the samples of the released calculations are binned and averaged on the grids of the corresponding fixed calculations.

width *a priori*, we average over the width parameter of the grid density. Integrating over this parameter in (15) requires evaluating the integral

$$\int dw \int d\mathbf{x} \prod_{i=1}^N \rho(w; x_i) \int_0^\infty d\bar{\mathbf{f}} e^{-\chi^2[\bar{\mathbf{f}}, \mathbf{x}]/2} f(\bar{\mathbf{f}}, \mathbf{x}; x). \quad (22)$$

One way could be sampling the width directly using Metropolis-Hasting, but that would be inefficient: updating the width would change the prior probabilities for all the grid points. So for a large number of grid points, one would be forced to take very small updates of the width to achieve a reasonable acceptance rate. The more grid points, the less efficient the sampling is. There is, however, a much more elegant and efficient solution.

We notice that unlike the grid points \mathbf{x} and spectral integrals $\bar{\mathbf{f}}$, the width parameter w is not directly related to the data. Therefore, the above expression can be rearranged such that the integral over the width is a function of the grid points only

$$\underbrace{\int d\mathbf{x} \int dw \prod_{i=1}^N \rho(w; x_i)}_{=: P(\mathbf{x})} \int_0^\infty d\bar{\mathbf{f}} e^{-\chi^2[\bar{\mathbf{f}}, \mathbf{x}]/2} f(\bar{\mathbf{f}}, \mathbf{x}; x). \quad (23)$$




FIG. 4. Spectral function $A(\omega)$ obtained using MaxEnt with Gaussian default models used in Fig. 3. The results are obtained using Bryan's method implemented in Ref. [24]. Other methods for choosing the regularization parameter, i.e., classic and historic MaxEnt give indistinguishable results.

We can then perform the width integral $P(\mathbf{x})$ analytically, e.g., for the family of density functions

$$\rho_q(w; x) \propto \frac{1}{w} \exp \left[-\frac{1}{q} \left(\frac{|x|}{w} \right)^q \right] \quad \text{where } q > 0, \quad (24)$$

which are known as the *exponential power distributions* and include the Gaussian distribution ($q = 2$), the exponential distribution ($q = 1$) and the uniform distribution ($q \rightarrow \infty$). The integral over the width then becomes

$$\begin{aligned} P(\mathbf{x}) &\propto \int_0^\infty dw \frac{1}{w^N} \exp \left[-\sum_i \frac{1}{q} \left(\frac{|x_i|}{w} \right)^q \right] \\ &= \int_0^\infty dw \frac{1}{w^N} \exp \left[-\frac{1}{q} \frac{\|\mathbf{x}\|_q^q}{w^q} \right], \end{aligned} \quad (25)$$

where the L_q -norm [25] is defined by

$$\|\mathbf{x}\|_q := \left(\sum_i |x_i|^q \right)^{1/q}. \quad (26)$$

We use this norm to make the following change of variable:

$$z := \frac{\|\mathbf{x}\|_q}{w} \Rightarrow \frac{dw}{dz} = -\frac{\|\mathbf{x}\|_q}{z^2}, \quad (27)$$

so that the integral over the new variable z becomes independent of the grid points, leaving us with

$$P(\mathbf{x}) \propto \frac{1}{\|\mathbf{x}\|_q^{N-1}}. \quad (28)$$

From this, we see that weighting the grid points vector by a power of its L_q -norm is equivalent to using a q th power exponential function as a grid density and integrating flat over the width parameter. More specifically, using the L_2 -norm is equivalent to using a Gaussian grid density and integrating over its standard deviation. Similarly, using the L_1 -norm is equivalent to using an exponential density and integrating over the scale parameter while using the L_∞ -norm corresponds to using a uniform density and integrating over the cutoff. Note that using a reciprocal prior in the width integral, as




FIG. 5. Spectral function $A(\omega)$ obtained using width-sampling ASM for the same problem as in Figs. 3 and 4. An exponential ($q = 1$) and a Gaussian grid density ($q = 2$) with 512 grid points is used. The samples are binned on a uniform grid.

appropriate for scale parameters [26, p. 109], instead of a flat one leads to a very similar result where the power of the norm is $N-2$ instead of $N-1$. The difference between the two choices is insignificant in practice.

Substituting (28) back into the original expression (23) and absorbing the integral over z in the overall normalization constant we get the formal expression for this extended version of the average spectrum method

$$f_{\text{ASM}}(\rho_q, N; x) \propto \int d\mathbf{x} \frac{1}{\|\mathbf{x}\|_q^{N-1}} \int_0^\infty d\bar{f} e^{-\frac{1}{2}\chi^2[\bar{f}, \mathbf{x}]} f(\bar{f}, \mathbf{x}; x). \quad (29)$$

We can easily adapt the sampling algorithm for released-grid ASM discussed in Sec. III A to evaluate this expression: We simply replace the prior density function in the acceptance ratio by the power ratio of the norms of grid samples. We only need to choose a reasonable value for the width parameter w of the proposal distribution which can be easily estimated, e.g., from the nonnegative least squares (NNLS) solution [22]. This value need not be very close to the width of the exact model. We ran calculations where it was an order of magnitude off. Obviously, its choice only affects the acceptance ratio of grid points sampling.

In Fig. 5, we show the results of width-sampling ASM for the same test case as in Fig. 3. We did the calculations using both exponential (L_1 -norm) and Gaussian grid densities (L_2 -norm). The results of both calculations are in excellent agreement with the exact spectrum without the need for fine-tuning the exact value of the width parameter. Width-sampling ASM produces even better agreement than MaxEnt (Fig. 4).

In Fig. 6, we show the histogram of the scaled L_2 -norm of grid samples from Gaussian densities ($q = 2$). We calculated the scaled norm of a grid sample as $\hat{w} := \sqrt{\sum_i x_i^2}/N$. Notice how the method automatically finds the optimal value of 0.5 and averages around it to the extent allowed by the noise in the data.

For completeness, we also report the effect of width averaging on the optical conductivity test case. The earlier results with released grid points were already very good and showed no noticeable dependence on the width, so it does not come as




FIG. 6. Histogram of the scaled L_2 -norm of grids sampled for the case $q = 2$ of Fig. 5. The scaled L_2 -norm of a grid \mathbf{x} is calculated as the standard deviation of its points $\hat{w} := \sqrt{\sum_i x_i^2}/N = \sqrt{\sum_i x_i^2}/N$. The widths are centered around the width of the exact model (0.5).

a surprise that the results with width averaging are as good, as can be seen from Fig. 7.

Note that, unlike the fixed- and released-grid average spectrum method, width-sampling ASM does not have a default model: In the absence of data, except for a sum rule, the method does not give a result. This is by design: For convergence, width-averaging requires the data to provide information about the width. Having the least informative prior, width-sampling ASM thus is the least biased of the methods discussed here as may be seen by comparing Figs. 3 and 4 to Fig. 5.

VI. GRID SIZE DEPENDENCE

In ASM1, we discussed the dependence of the fixed-grid average spectrum method on the number of grid points and found that it plays the role of a regularization parameter: As the number of grid points increases, the results change and get more biased towards the grid density. The same behavior still applies to the released-grid ASM but is much weaker. This is in line with the overall reduction in the dependence on the grid density we saw earlier. For example, we




FIG. 7. Optical conductivity $\sigma(\omega)$ obtained using width-sampling ASM for the same problem as in Figs. 1 and 2. An exponential ($q = 1$) and a Gaussian grid density ($q = 2$) with 512 grid points is used. The samples are binned on a uniform grid.




FIG. 8. Dependence of the optical conductivity $\sigma(\omega)$ obtained using fixed-grid ASM (top) and released-grid ASM (bottom) on the grid size N (label). A Lorentzian grid density with parameter $\gamma = 2.5$ is used. For ease of comparison, the samples of the released calculations are binned and averaged on the grids of the corresponding fixed calculations.

compare in Fig. 8 the N -dependence of the optical conductivity test case using fixed and released grid points. While the fixed-grid calculations show a significant variation, released-grid ASM results hardly change with the number of grid points except for small differences near $\omega = 0$. We know that using a much larger number of grid points, even the released-grid method will eventually show a more pronounced dependence on the grid size. We did not, however, observe this dependence in this test case for any reasonable value of N . To see it in released-grid ASM or its width-sampling extension, we need to look at yet another case where the exact spectrum is significantly different from the singly peaked default models we typically use.

To this end, let us take a spectral function composed of four Gaussian peaks symmetric around zero. Two of the peaks are narrow with width 0.1 and weight 0.15 and located at frequencies ± 0.5 . The other two are wider with width 0.5 and weight 0.35 and located further out at ± 2 . As with the previous spectral function, we generate the Green's function data using (1) on 60 τ -points for $\beta = 50$. We add relative Gaussian noise with a standard deviation of 10^{-2} .

In the top panel of Fig. 9, we plot the results of width-sampling ASM using a Gaussian grid density ($q = 2$) and different grid sizes. As the grid size increases, the results get smoother and the peaks get wider demonstrating the regularizing effect of the grid size even in the released-grid




FIG. 9. Grid size dependence of the spectral function $A(\omega)$ of four peaks using width-sampling ASM. A Gaussian grid density ($q = 2$) is used. The samples are binned on a uniform grid. The bottom panel shows histograms of the data fits of sampled spectra.

case. The behavior is similar using a fixed width. A better understanding is gained by checking the fit histograms of the sampled spectra in the bottom panel. As the grid size increases beyond $N = 128$, the fit histograms, introduced in ASM1, shift to the right and get wider showing a systematic bias towards spectra of worse fits. On the other end, we observed that when the grid size is lower than $N = 16$, the data fit gets extremely poor due to the large discretization error. We did not include results using such low grid sizes to avoid cluttering the plots.

The qualitative behavior of the fit in the average spectrum methods is depicted in Fig. 10. For very low grid sizes, the




FIG. 10. Schematic diagram of the average fit behavior in ASM as the number of grid points N changes.

discretization error dominates, leading to a very bad fit. Once the grid size is large enough such that the discretization error becomes negligible relative to the noise on the data, increasing the grid size leads to more bias and worse fit. This dependence on the grid size starts out slowly and then accelerates till the average spectrum approaches the default model (grid density) for very large N .

Therefore, a weak dependence on the grid size indicates a better grid density, and when the density matches the exact model, the sweet spot extends to infinity. This behavior applies equally to both the fixed and the released-grid ASM. However, there are two main differences. The *discretization* error for released-grid ASM is normally less than that of the fixed grid. Also, the region of low N -dependence extends further because the bias towards the grid density is much reduced compared to fixed-grid ASM. In general, the results of released-grid ASM show weak dependence for the typical grid sizes we use. A good recipe for choosing N is to use the largest value for which the fit does not get substantially worse. For the test case of Fig. 9 this would be $N = 128$.

It may be worth mentioning that we tried sampling the grid size using a flat prior. We found that the method chooses a high number of points when the default model is highly compatible with the data as in Fig. 8. But when the default model does not fit the data very well, the method chooses a very low number of points. This happened in cases like the one shown in Fig. 9. In the end, we decided to keep the grid size as an independent parameter of the method that we use for checking the reliability of the results: A strong dependence on the grid size indicates a strong bias towards the default model and away from data.

Finally, we comment on the smoothing effect observed for the satellite peaks of Fig. 9. The reason for this smoothing should become clear by examining the integral equation (1), which relates this spectrum to its data. Due to the exponential factor $e^{-\omega\tau}$, features of the spectrum that lie to the far right, where $\omega \gg 1/\tau$, have almost no contribution to the data points evaluated at τ or larger. Similarly, features that lie to the far left, such that $\omega \ll -1/\tau$, have almost no contribution to the data points evaluated at $\beta - \tau$ or smaller. As a result, fewer data points will contain information about spectral features as these features move away from zero. In general, the data have a weaker dependence on the tail of the spectrum than the features near zero. For ASM methods, this means that there are many spectra with different tails that would give rise to similar data. One obvious remedy would be to obtain higher-quality data with more points, especially at the edges of the interval $[0, \beta]$. When this is not feasible, no unbiased method will be able to resolve such features fully *without prior knowledge*.

VII. SUMMARY AND DISCUSSION

The results of the average spectrum method can strongly depend on the discretization grid. One approach for handling this dependence is choosing a grid density that fits the data

recently. We showed that sampling the grid points helps to reduce the bias dramatically and thus obviates, in many cases, the need to search for the best grid. But in some cases, the width of the grid density still influences the results noticeably. We, therefore, went one step further and averaged over the results for grids of different widths. Remarkably, for a large family of grid densities, we could perform this additional sampling analytically, incurring only a negligible computational overhead.

The approach used here for handling the grid dependence led us to the following hierarchy of average spectrum methods where each method extends the previous one by averaging over the relevant parameters

$$\underbrace{\int dw \int d\mathbf{x} \prod_{i=1}^N \rho_q(w; x_i)}_{f_{\text{ASM}}(\rho_q, w, N; \mathbf{x})} \underbrace{\int_0^\infty d\tilde{\mathbf{f}} e^{-\frac{1}{2}\chi^2[\tilde{\mathbf{f}}, \mathbf{x}]} f(\tilde{\mathbf{f}}, \mathbf{x}; x)}_{f_{\text{ASM}}(\mathbf{x}; x)}. \quad (30)$$

It is obvious that the approach can be taken further by varying over other parameters of the grid density. Had we, for example, observed a strong dependence on the functional form of the grid density, e.g., the parameter q in (29), we would have to average the results of different values using suitable weights.

From this perspective, the functional-integral-based average spectrum method can be seen as a general framework for obtaining data-compatible spectra in the context of analytic continuation and similar spectral reconstruction problems. Given a certain parametrization of the spectrum, the most straightforward solution is estimating these parameters by fitting the data. In many situations, this may be enough to single out a small region of the parameter space with tolerable variations in the spectrum. When observing a noticeable sensitivity of the results to a parameter, one should consider averaging over the results of different values of this parameter to smooth out details not supported by the data. In light of this, the fixed-grid method of ASM1 can be itself seen as an extension of the nonnegative least squares method (NNLS), where instead of finding the spectral integrals that fits the data best, it averages over, giving equal weights to all spectral integrals fitting the data equally.

In this paper, we could discuss only a few test cases, each of which was introduced for illustrating specific aspects of the average spectrum methods. Further test cases and applications of these extensions to real-world problems are discussed in Ref. [27]. To encourage further testing of the ASM methods and their use for practical problems, we present a highly optimized implementation of all ASM methods and make it freely accessible through an online interface found in Ref. [28], where the user can upload data in a convenient format, submit calculations with few clicks, and get the results interactively. The interface also facilitates comparing different results to apply our recipe for choosing the optimal parameters, if needed.

- [3] O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, *Phys. Rev. B* **81**, 155107 (2010).
- [4] M. Jarrell, in *Correlated Electrons: From Models to Materials*, edited by E. Pavarini, E. Koch, F. Anders, and M. Jarrell (Forschungszentrum Jülich, Jülich, Germany, 2012).
- [5] D. Bergeron and A.-M. S. Tremblay, *Phys. Rev. E* **94**, 023303 (2016).
- [6] S. R. White, in *Computer Simulation Studies in Condensed Matter Physics III*, edited by D. P. Landau, K. K. Mon, and B.-B. Schüttler (Springer, Heidelberg, Germany, 1991), pp. 145–153.
- [7] A. W. Sandvik, *Phys. Rev. B* **57**, 10287 (1998).
- [8] K. S. D. Beach, [arXiv:cond-mat/0403055](https://arxiv.org/abs/cond-mat/0403055).
- [9] K. Vafayi and O. Gunnarsson, *Phys. Rev. B* **76**, 035115 (2007).
- [10] O. F. Syljuåsen, *Phys. Rev. B* **78**, 174429 (2008).
- [11] S. Fuchs, T. Pruschke, and M. Jarrell, *Phys. Rev. E* **81**, 056701 (2010).
- [12] A. W. Sandvik, *Phys. Rev. E* **94**, 063308 (2016).
- [13] K. Ghanem and E. Koch, *Phys. Rev. B* **101**, 085111 (2020).
- [14] H. J. Vidberg and J. W. Serene, *J. Low Temp. Phys.* **29**, 179 (1977).
- [15] K. S. D. Beach, R. J. Gooding, and F. Marsiglio, *Phys. Rev. B* **61**, 5147 (2000).
- [16] A. Östlin, L. Chioncel, and L. Vitos, *Phys. Rev. B* **86**, 235107 (2012).
- [17] J. Schött, I. L. M. Locht, E. Lundin, O. Grånäs, O. Eriksson, and I. Di Marco, *Phys. Rev. B* **93**, 075104 (2016).
- [18] L.-F. Arsenault, A. Lopez-Bezanilla, O. A. von Lilienfeld, and A. J. Millis, *Phys. Rev. B* **90**, 155136 (2014).
- [19] L.-F. Arsenault, R. Neuberger, L. A. Hannah, and A. J. Millis, *Inverse Probl.* **33**, 115007 (2017).
- [20] H. Yoon, J.-H. Sim, and M. J. Han, *Phys. Rev. B* **98**, 245101 (2018).
- [21] R. Fournier, L. Wang, O. V. Yazyev, and Q. S. Wu, *Phys. Rev. Lett.* **124**, 056401 (2020).
- [22] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (SIAM, Philadelphia, 1995).
- [23] O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, *Phys. Rev. B* **82**, 165125 (2010).
- [24] R. Levy, J. LeBlanc, and E. Gull, *Comput. Phys. Commun.* **215**, 149 (2017).
- [25] Strictly speaking, this expression does not define a norm when $q < 1$ because it violates the triangle inequality. Nevertheless, our results still hold even in that case.
- [26] C. S. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial*, 2nd ed. (Oxford University Press, Oxford, 2006).
- [27] K. Ghanem, Stochastic Analytic Continuation: A Bayesian Approach, Ph.D. thesis, RWTH Aachen University, 2017.
- [28] www.spektra.app.