

Dual formulation of the maximum entropy method applied to analytic continuation of quantum Monte Carlo data

Thomas Chuna^{1,2,*} , Nicholas Barnfield³ ,
Tobias Dornheim^{1,2} , Michael P Friedlander⁴ 
and Tim Hoheisel⁵ 

¹ Center for Advanced Systems Understanding (CASUS), D-02826 Goerlitz, Germany

² Helmholtz-Zentrum Dresden-Rossendorf (HZDR), D-01328 Dresden, Germany

³ Department of Statistics, Harvard University, Cambridge, MA 02138, United States of America

⁴ Department of Computer Science and Mathematics, University of British Columbia, Vancouver, BC V6T 1Z4, Canada

⁵ Department of Mathematics and Statistics, McGill University, Montreal, Quebec H3A 0G4, Canada

E-mail: t.chuna@hzdr.de

Received 10 March 2025; revised 23 July 2025

Accepted for publication 7 August 2025

Published 19 August 2025



CrossMark

Abstract

Many fields of physics use quantum Monte Carlo techniques, but struggle to estimate dynamic spectra via the analytic continuation of imaginary-time quantum Monte Carlo data. One of the most ubiquitous approaches to analytic continuation is the maximum entropy method (MEM). We supply a dual Newton optimization algorithm to be used within the MEM and provide analytic bounds for the algorithm's error. The optimization algorithm is freely available on github ([repository](#)). The MEM is typically used with Bryan's controversial algorithm (Rothkopf 2020 *Data* **5** 55). We present new theoretical issues that are not yet in the literature. Our algorithm has all the theoretical benefits of Bryan's algorithm without these theoretical issues the implementation of the dual Newton optimizer within the MEM is freely available on

* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

github ([repository](#)). We compare the MEM with Bryan’s optimization to the MEM with our dual Newton optimization on test problems from lattice quantum chromodynamics and plasma physics. These comparisons show that in the presence of noise the dual Newton algorithm produces better estimates and error bars; this indicates the limits of Bryan’s algorithm’s applicability. We use the MEM to investigate authentic quantum Monte Carlo data for the uniform electron gas at warm dense matter conditions and further substantiate the roton-type feature in the dispersion relation.

Keywords: inverse problems, maximum entropy method, analytic continuation, spectral function, dynamic structure factor

1. Introduction

Path integral Monte Carlo (PIMC) simulations are used to estimate imaginary-time correlation functions (ITCFs) $F(\tau)$, which are one- and two-particle Green functions. The ITCFs are associated with dynamical spectra $S(\omega)$, which are useful because they predict measurable real-time quantities. However, to extract spectra from ITCFs an analytic continuation to real time is required. This involves inverting the integral transform

$$F(\tau) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} A(\tau, \omega) S(\omega) d\omega, \quad (1)$$

where $\tau \in [0, \beta)$. Across different applications, computational budgets prohibit us from collecting sufficient data that would—even in the ideal situation—be sufficient to uniquely invert the transformation. In addition to this difficulty, the kernel takes one of the forms $A(\tau, \omega) = \exp(-\omega\tau)$ or $A(\tau, \omega) = \frac{\exp(-\omega\tau)}{1 \pm e^{-\beta\omega}}$. Thus, the analytic continuation problem is, or simplifies to, inverting the Laplace transform, which is an exponentially ill-posed problem, i.e. the singular values of the transformation matrix decay to zero exponentially fast [1]. Therefore, the inversion problem needs to be regularized in order to make the inversion well-posed. In this paper, we propose a dual Newton algorithm for solving an entropy-regularized linear least-squares problem that alleviates the difficulties associated with the analytic continuation problem. The dual Newton algorithm we describe efficiently scales to large problems and is accompanied by an error analysis that characterizes its accuracy from observable quantities.

The maximum entropy method (MEM), as proposed by Bryan [2] and subsequently applied to the analytic continuation problem by Gubernatis *et al* [3], typically serves as the benchmark for new methods [4–8]. These comparisons are complicated by the fact that Bryan’s maximum-entropy approach uses an optimization algorithm that does not solve the original entropy-regularized least-squares formulation, but only an approximation [2, 9]. Bryan’s algorithm relies on a singular value decomposition of the discrete kernel $A \in \mathbb{R}^{N_\tau \times N_\omega}$ and represents the entire kernel $A = U\Sigma V^T$ by N_s of right singular vectors in V , denoted

$$A^T \rightarrow \tilde{V}, \quad (2)$$

where $N_s < N_\tau$. This substitution eliminates the left singular vectors U and singular values Σ . These comparisons are also complicated by the fact that the relative entropy, which Bryan’s algorithm optimizes, includes a term driving the elements of the spectral function to sum to the same value as the elements of the default model (i.e. Bayesian prior). In this work, we will show that although this assumption is valid for zero-temperature systems, the MEM is often applied outside this limit. Our dual Newton algorithm addresses both of these issues.

Our comparisons between the MEM with Bryan’s algorithm and the MEM with a dual Newton algorithm reflect the theoretical advantages of our algorithm. Our dual Newton optimizer produces better solutions and uncertainty estimates. With respect to its solutions, the MEM with the dual Newton algorithm uses the transformation kernel’s entire basis so the dual Newton’s solution differs from Bryan’s solution. In particular, the additional basis vectors allow the dual Newton algorithm to estimate functional forms that are problematic for Bryan’s algorithm. With respect to the uncertainty, noise causes the posterior weighting function for MEM with Bryan’s algorithm to favor more regularized solutions than the MEM with the dual Newton algorithm. The implication is that the variance of proposed solutions diminishes and causes the MEM with Bryan’s algorithm to artificially suppress uncertainty estimates. Error estimates are of utmost importance for scientific applications where the true solution is not known. To further address this uncertainty issue, we provide analytic expressions of the error bounds for a solution obtained via the dual Newton.

The paper is partitioned as follows. Section 2 describes prior work. Section 3 formulates the optimization problem, describes the dual approach, and presents the analytic error bounds. Section 4 presents numerical comparisons between the dual Newton and Bryan’s algorithm. We conduct three investigations. First, we consider Asakawa *et al*’s zero-temperature ρ -meson benchmark, which is relevant to the field of lattice quantum chromodynamics [10]. Next, we introduce a new finite temperature benchmark, using the completed Mermin (CM) collisional uniform electron gas (UEG) dynamic structure factor (DSF) [11]. This example is relevant to the field of high energy density physics. Finally, we apply the method to authentic PIMC data of the UEG at warm dense matter (WDM) conditions. Section 5 discusses the conclusions we draw from our numerical results. Appendix gives an overview of our dual optimization algorithm.

2. Prior work

Many methods have been applied to the analytic continuation problem across different fields of physics. The Backus–Gilbert (BG) method [12, 13] and its smooth variant [14] has been used to estimate spectral densities of quarkonia. The Bayesian reconstruction method (BRM) [15, 16] and its smoothed variant [7, 17] have similarly been applied. Genetic algorithms have been used to construct spectral functions of ultracold ^4He [18–20] and ^3He [21]. Artificial neural networks have been used to estimate band structures of the Hubbard model [22, 23]. Stochastic analytic inference (SAI) methods have been used to estimate the spectral function of a Bardeen–Cooper–Schrieffer (BCS) super conductor [24], band structures of the Hubbard model [8], and spectral functions in lattice quantum chromodynamics (LQCD) [25]. The MEM has been widely applied to estimate quantities of interest in LQCD [10], Hubbard model [8], dynamical mean field Theory [26], and many more. Additionally, specialized methods have been developed for specific applications, including the Nevanlinna representation of spectral functions with noise-free data [27], the Hamburger problem formulation based on the frequency moments of a given spectral density [28], and methods using additional constraints derived from known physical properties of dynamic spectra [29, 30]. Comparative studies suggest that there is no established ‘best’ approach to the analytic continuation problem [4–8].

We discuss the more general algorithms, which take any given analytic continuation kernel and any given ITCF data. Axiomatic arguments conclude that minimizing a least-squares fidelity objective function with relative-entropy regularization does not introduce erroneous correlations into the solution [31–33], as illustrated by examples outlined by Jarrel and Gubernatis [34]. Therefore, entropic regularization, as seen in MEM, BRM, and SAI methods,

is central to our approach. Across comparisons of these three methods, the MEM estimate tends to be the smoothest of these entropic estimates. Comparing MEM and BRM, Kim *et al* [17] show that MEM estimates can be produced by the BRM if a smoothness regularization is added to the BRM objective function. A primary motivation for adding the smoothness regularization was that the BRM estimates often contained spurious wiggles not present in the MEM estimate [7]. Comparing MEM and SAI, Fuchs *et al* [8] conclude that the SAI method produces spectral functions that are less regularized and consequently show more pronounced features than the MEM, but their SAI method is orders of magnitude more computationally expensive than the MEM.

Both Kim *et al* [17] and Fuchs *et al* [8] use a MEM based on Bryan’s optimization algorithm. Bryan argues that his controversial substitution (2) is justified when the gradient of the fidelity term lies in the singular space, aiming to simplify the problem from an ill-conditioned N_ω -dimensional optimization to a well-conditioned N_s -dimensional optimization. This reduction in dimensionality and improved conditioning both decrease the computational cost. Essentially, Bryan’s approach removes singular vectors associated with small singular values, thereby eliminating flat search directions from the optimization process. However, regularizing through singular values enhances a solution’s smoothness, limiting MEM’s ability to resolve pronounced features—an observation consistent with the findings of Kim *et al* and Fuchs *et al*. Furthermore, a recent study has identified specific example functions that cannot be adequately represented by this truncated basis set [9]. This highlights a need for a maximum entropy algorithm that retains the advantages of Bryan’s method (i.e. reduced search space and improved conditioning) without excluding singular vectors.

In this work, we introduce an optimization algorithm designed to replace Bryan’s optimization algorithm. While Bryan’s algorithm uses a Levenberg–Marquardt optimizer to solve a simplified version of the primal problem [2], our approach uses a Newton–Krylov optimizer to solve the *dual problem*. Although the dual formulation of entropic regularization is well established [35–38], it has not previously been applied to the analytic continuation problem. Both the primal and dual formulations lead to the same unique solution; however, optimizing the dual formulation offers certain advantages, including guaranteed differentiability and reduced dimensionality. In particular, the dual problem is formulated within the column space of the kernel rather than its row space. For a discretized transformation kernel $A \in \mathbb{R}^{N_\tau \times N_\omega}$, this reduces the optimization dimension from N_ω to N_τ . For instance, in the work of Asakawa *et al* [10], the dual formulation would reduce the optimization dimension from $n = 600$ to $m \leq 30$ without any loss of information. Additionally, differentiability allows for the use of second-order optimizers, which outperform the Levenberg–Marquardt optimizer. In summary, the dual approach enables the application of a higher-order optimizer to a better-conditioned problem within a reduced search space, while preserving all basis vectors.

3. MEM

In this section we formulate analytic continuation, formulate dual optimization, and formulate error bounds for the dual solution. We do not discuss the MEM procedure in full detail. There are many well-organized works where the MEM is outlined [2, 10, 34]. The method assumes there exists a collection of solutions which optimize the least-squares fidelity term plus entropic regularizer at different regularizer weights α . The core of the MEM is Gull’s derivation of the Bayesian posterior, which indicates how to properly combine these solutions [39]. Importantly, the MEM procedure does not specify how the solutions are to be obtained. Typically Bryan’s algorithm, which uses a modified Levenberg–Marquardt algorithm, is used

to obtain these solutions [2]. Yet, for any given α the MEM's cost function has a unique minimum, which can be discussed independent of the optimization algorithm (e.g. Levenberg–Marquardt, ADAM, Newton, conjugate gradient) or the approach (i.e. primal or dual) used to obtain it. In this work, we leave Gull's Bayesian analysis intact and replace Bryan's primal modified Levenberg–Marquardt algorithm with a better performing dual Newton optimization algorithm. In practice, we leave our MEM code intact and a call different optimizer subroutine. An implementation of the dual Newton optimizer within the MEM is freely available on github ([repository](#)) [40].

3.1. Noisy linear model of PIMC data

In the most general form, PIMC methods generate samples of ITCFs $F(\mathbf{x}, \tau | \beta)$, where \mathbf{x} is the spatial coordinate, τ is imaginary-time, and β is the inverse temperature. By definition, the desired spectra are the Fourier transform of the correlation functions. Fourier transforming in the spatial dimension is trivial and we arrive at $F(k, \tau)$ where k is the wavenumber. However, Fourier transforming the temporal component τ is non-trivial. Most PIMC simulations are done in imaginary time $\tau = -it$. Thus the Fourier kernel $\exp(i\omega t)$ becomes a Laplace kernel, and the Laplace transform of the data is ⁶

$$F(k, \tau) \equiv \frac{1}{2\pi} \int_{-\infty}^{+\infty} \exp(-\omega\tau) S(k, \omega) d\omega + \eta(k, \tau). \quad (3)$$

The factor of 2π converts ω from rad/s to 1/s and the sampling noise $\eta(k, \tau)$ arises from the PIMC algorithm (e.g. Metropolis–Hastings). This work will focus on inverting (3), though other works continue on to produce the fermion [boson] kernels $A(\tau, \omega) = \frac{\exp(-\omega\tau)}{1 \pm e^{-\beta\omega}}$.

For finite temperature simulations, periodic boundary conditions introduce the temperature, i.e. $F(k, \tau) = F(k, \beta - \tau)$ and $T = 1/\beta$. This leads most research efforts to consider the inversion of periodic Laplace transforms (e.g. Hansen *et al* [14] or Ferre *et al* [19]), expressed as,

$$F(k, \tau) \equiv \frac{1}{2\pi} \int_0^\infty (\exp(-\omega\tau) + \exp(-\omega(\beta - \tau))) S(k, \omega) d\omega + \eta. \quad (4)$$

We have suppressed the arguments of the noise η for compactness. This equation is derived from (3) using the detailed balance condition $S(k, -\omega) = e^{-\beta\omega} S(k, \omega)$ [42].

The Laplace (3) or periodic Laplace (4) both have strengths and weaknesses. The strength of the non-periodic approach (3) is that the normalization of $S(k, \omega)$, given by

$$F(k, \tau = 0) = \frac{1}{2\pi} \int_{-\infty}^{\infty} S(k, \omega) d\omega, \quad (5)$$

is known and can be used to scale the correlation function's $\tau = 0$ entry to 1, which normalizes $S(k, \omega)$. The weakness of the Laplace transform is that the exponential suppression arising from the detail balance relation means that hundreds of numerical zeros are being fit in the $\omega < 0$ domain. By comparison, the strength of the periodic approach is that the $\omega < 0$ degrees of freedom are eliminated. The weakness of the periodic kernel is that the positive normalization of $S(k, \omega)$

⁶ Since $S(k, \omega)$ contains poles, solving (3) is an analytic continuation problem [41].

$$Z = \frac{1}{2\pi} \int_0^\infty S(k, \omega) d\omega, \quad (6)$$

is not known. To see that the positive normalization is not known, consider (4) at $\tau = 0$:

$$F(k, \tau = 0) = \frac{1}{2\pi} \int_0^\infty (1 + \exp(-\omega\beta)) S(k, \omega) d\omega. \quad (7)$$

When the temperature is zero, the second term can be neglected, but when the temperature is finite the frequency dependent term persists and the value of $F(k, \tau = 0)$ does not match the positive normalization 6. A positive normalization Z must either be supplied or searched for. The implications of not knowing the value of Z is described in the problem formulation section 3.2. In this work, we opt for the periodic Laplace transform; this choice matches with the majority of the literature we referenced in section 2.

Another practical consideration is the discretization of the transform. PIMC simulations have discrete time variables. We denote the discrete τ with index i as $\tau_i = i\Delta\tau$ where $i \in \{0, \dots, N_\tau - 1\}$; both $\Delta\tau$ and N_τ are set by the PIMC simulation. Computational capacity limits PIMC simulations to a discretization of size $N_\tau \approx 10^2 - 10^3$. Additionally, we impose the discretization of the frequency ω as $\omega_j = j\Delta\omega$ where $\Delta\omega$ is the resolution and $j \in \{1, \dots, N_\omega\}$. The resolution needs to be sufficient to resolve the relevant features in $S(k, \omega)$. Typically $N_\omega \approx 10^3 - 10^4$ is sufficient. Additionally, PIMC simulation have periodic space variable, which leads to a discrete wavenumber k . For a given k value, the discrete formulation of either (3) or (4) is given by

$$b_i = \Delta\omega \sum_j A_{i,j} x_j + \eta_i. \quad (8)$$

We dropped the k index for compactness and we have expressed the discrete kernel $A(\tau_i, \omega_j)$ as matrix $A_{i,j}$, the discrete spectra $S(k, \omega_j)$ as the vector x , the discrete noise $\eta(k, \tau_i)$ as the vector η_i , and the discrete correlation function $F(k, \tau_i)$ as the vector b_i .

3.2. Problem formulation

Equation (8) presents a numerically tractable formulation of the inversion problem. However, when $N_\tau < N_\omega$ the inversion of A becomes underdetermined, meaning the solution is not unique. Additionally, the discretized Laplace transform produces a matrix $A_{i,j}$ whose singular values decay exponentially fast, often approaching numerical zero. This rapid decay leads to severe ill-conditioning, making a direct solution via least squares both unstable and unreliable.

In such cases, regularization techniques are essential to constrain the problem and obtain a unique, stable solution. Regularization helps to mitigate the ill-posedness by incorporating prior information, ensuring the solution reflects both the data and the typical smoothness of a DSF.

3.2.1. KL-regularized least-squares. The MEM minimizes a relative entropy term alongside a generalized least-squares fidelity term. The optimization problem is formulated as

$$\min_{x \in \mathbb{R}_+^{N_\omega}} \alpha S_{\text{SJ}}(x | \mu) + \frac{1}{2} \|Ax - b\|_{C^{-1}}^2, \quad (9)$$

where $\|z\|_{C^{-1}}^2 = z^T C^{-1} z$ denotes the scaled two-norm, C represents the covariance matrix of the error in the observation vector b , and regularization parameter $\alpha \geq 0$ is a positive regularization parameter that balances the importance of fitting the observed measurements versus matching the Bayesian prior μ . The term

$$S_{\text{SJ}}(x | \mu) = \sum_{j=1}^n \mu_j - x_j + x_j \log(x_j / \mu_j) \tag{10}$$

is the Shannon–Jaynes (SJ) relative entropy⁷ between discrete distributions x and μ . Going forward we refer to (10) as the relative entropy. We adopt the standard convention that $0 \log 0 = 0$. Note that where $\mu_i = 0$ then x_i must also be zero to render $S_{\text{SJ}}(x | \mu)$ finite. The relative entropy encodes the conventional relative entropy term in addition to a linear term

$$\sum_{j=1}^n \mu_j - x_j, \tag{11}$$

which enforces a soft-constraint on the positive normalization of x with that of μ . As an option, our inversion code allows the user to omit the soft constraint in (11), optimizing instead the conventional relative entropy. In this work, we include the soft constraint to make direct comparisons.

3.2.2. Normalization. As discussed in connection with the periodic formulation (4), the normalization Z defined in (6) is not known. An additional source of uncertainty in the normalization accrues from the quantization errors in the discretization (8). The prior μ will have its own normalization $Z_\mu = \sum_i \mu_i$, but as shown by the functional (20a), the discrete normalized distribution with maximum-entropy takes the form

$$\frac{x}{Z} = \frac{\mu_j \exp\left(\sum_{i=1}^{N_\omega} A_{i,j} y_i\right)}{\sum_{\hat{j}=1}^{N_\tau} \mu_{\hat{j}} \exp\left(\sum_{i=1}^{N_\omega} A_{i,\hat{j}} y_i\right)}, \tag{12}$$

where the scalars y_i are Lagrange multipliers corresponding to the dual maximum-entropy problem; see [43, section 3] and (20a). The form of (12) is to be expected and matches the form of the grand canonical Boltzmann factor. Note that because the numerator and denominator of this expression are both linear in the prior μ , we can see that scaling prior μ_j by Z_μ has no bearing on the maximum-entropy solution x .

For any fixed normalization Z , the maximum-entropy solution has negative entropy given by the value function

$$V(Z) = \min_{x \in Z\Delta_{N_\omega}} p(x, Z | \mu) \tag{13}$$

where

$$p(x, Z | \mu) := Z S_{\text{SJ}}\left(\frac{x}{Z} \middle| \frac{\mu}{Z}\right) + \sum_{j=1}^n \mu_j - x_j + \frac{1}{2\alpha} \|Ax - b\|_{C^{-1}}^2 \tag{14}$$

⁷ Relative entropy (Kullback–Leibler Divergence) belongs to the family of f -divergences, which are often interpreted as pseudo-metrics for probability measures.

and

$$S_{\text{SJ}'}(x | \mu) = \begin{cases} \sum_{j=1}^n x_j \log(x_j / \mu_j), & \text{if } x \in \Delta_{N_\omega} \\ \infty, & \text{otherwise,} \end{cases} \quad (15)$$

is the modified relative entropy term from (10). The identity $Z S_{\text{SJ}'}\left(\frac{x}{Z} \middle| \frac{\mu}{Z}\right) = S_{\text{SJ}'}(x | \mu)$ is applied in (13) because solutions not normalized to 1 incur an infinite penalty.

We thus aim to find the normalization Z that solves the problem

$$\min_{Z > 0} V(Z). \quad (16)$$

This is equivalent to the problem (9) in terms of achieving the same optimal value.

3.2.3. Dual of KL-regularized least squares. The objective function given in (13), referred to as the *primal* problem, is nonsmooth, necessitating the use of constrained or nonsmooth optimization algorithms. Such algorithms, however, are typically less efficient and reliable than those designed for smooth unconstrained optimization [44, section 17.2]. Given the strict convexity of the KL divergence in (15), we can leverage the duality between strict convexity and differentiability [45, theorem 11.13]. The dual of the KL-regularized least-squares problem (13) is the minimization, over $y \in \mathbb{R}^{N_\tau}$, of the strongly convex and smooth objective function

$$\min_{y \in \mathbb{R}^{N_\tau}} d(y, Z | \mu), \quad (17)$$

where

$$d(y, Z | \mu) := \frac{\alpha}{2} y^T C y - \langle b, y \rangle + Z \log \exp(A^T y - \mathbf{1} | \mu) - Z \log(Z), \quad (18)$$

is the dual objective function. Here $\langle \cdot, \cdot \rangle$ indicates the inner/dot product. The vector of 1's $\mathbf{1}$ arises from (11) and the log-sum-exp function

$$\log \exp(z | \mu) \equiv \log \sum_{j=1}^n \mu_j \exp(z_j) \quad (19)$$

arises from the Legendre-Fenchel transform of the KL divergence, as detailed in Rockafellar and Wets [45, example 11.2].

A straightforward approach to optimizing (13) involves two subproblems: first, for a given Z , optimize $p(x, Z | \mu)$, then search over Z for the unique optimum. In this way, (13) corresponds to (9), where the feasible set of solutions is restricted to the Z -simplex.

A one-to-one correspondence between the primal and dual solutions is established by Fenchel-Rockafeller duality [45, theorem 11.39 and example 11.41]. This relationship is defined by the smooth mappings

$$x(y) = Z \nabla \log \exp(A^* y | \mu), \quad (20a)$$

$$y(x) = \alpha^{-1} C^{-1} (b - Ax), \quad (20b)$$

where

$$\nabla \log \exp(z | \mu) \equiv \frac{\mu \odot \exp(z)}{\mu^T \exp(z)}. \quad (20c)$$

Here, $\exp(z)$ denotes the Hadamard (elementwise) product, so $\mu \odot \exp(z)$ represents a vector with components $\mu_j \exp(z_j)$, and $\exp(z)$ denotes a vector whose components are $\exp(z_j)$.

In summary, while the primal problem (9) is nonsmooth and thus requires gradient descent optimization, our associated dual problem (17) is smooth and can be optimized with using second-order solvers (e.g. Newton's Method). Furthermore, the primal problem is formulated in a N_ω -dimensional column space and the dual problem is formulated in the lower N_τ -dimensional row space. These benefits come with no drawbacks as the primal solution can be easily recovered from the dual solution via (20a),

In [appendix](#) we describe a Newton–Krylov method that leverages the particular structure of the objective function and searches over positive normalizations Z . Our implementation of the dual Newton algorithm is freely available on github ([repository](#)) [46]. However, a variety of other algorithms are also suitable since the dual problem (17) is smooth and amenable to standard smooth optimization techniques.

3.3. Perturbation analysis

In this section, we examine how the computed primal solution is affected by various sources of error. Specifically, we analyze three key sources: first, we assess how perturbations in the dual solution influence the corresponding primal solution. Second, we analyze how perturbations in the vector of observations b influence the primal solution. Finally, we investigate how perturbations in the unknown regularization parameter α influence the primal solution. Understanding these relationships is crucial for assessing the robustness and reliability of the solutions obtained through our approach. In all cases, sensitivity analysis is performed under a fixed Z setting.

3.3.1. Relationship between dual error and primal error. For brevity in notation, we set $A_C := C^{-1/2}A$ and $b_C := C^{-1/2}b$ in this section. Let y_p be the exact optimal solution of the dual problem (17) for a fixed set of parameters $p = (b, \alpha) \in \mathbb{R}^{N_\tau} \times \mathbb{R}_{++}$, and let $x_p := x(y_p)$ be the corresponding primal solution obtained via the map (20a). Since $x(y)$ is a composition of the adjoint of the linear map A_C with $Z \log \exp$, which has a Z -Lipschitz gradient, it follows that $x(y)$ is Lipschitz continuous with Lipschitz constant $Z\|A_C\|$ [47, example 5.15]. For any parameterizations $p = (b, \alpha)$ and $p' = (b', \alpha')$, we obtain the upper bound

$$\|x_p - x_{p'}\| \leq Z\|A_C\| \cdot \|y_p - y_{p'}\|, \quad (21)$$

where $\|\cdot\|$ denotes the two-norm, and hence $\|A_C\|$ is the largest singular value of A_C . Equation (21) shows that any error in the inferred primal solution is bounded by a constant multiple of the error in the dual solution. An immediate application of this bound concerns the sequence of dual iterates y^k generated by the Newton–Krylov method described in [appendix](#). Specifically, the error in corresponding primal iterate $x(y^k)$ is bounded by the error in y^k scaled by the norm of the operator A_C :

$$\|x(y^k) - x^*\| \leq Z\|A_C\| \cdot \|y^k - y^*\|, \quad (22)$$

where (x^*, y^*) represents the exact primal–dual solution. This bound establishes the validity of the dual approach: as the dual iterates converge to the dual optimum, the corresponding primal iterates also converge to the optimum at the same rate.

3.3.2. Relationship between data perturbations and primal solution. We examine how errors or changes in the observation data b affect the solution of our primal problem. To do so,

we consider the *solution map* of the dual problem, which links the input data b and regularization parameter α to the unique optimal solution.

The dual problem can be reformulated by completing the square in the dual objective function d , which yields

$$\begin{aligned} P(b_C/\alpha) &:= \arg \min_{y \in \mathbb{R}^m} \left\{ \frac{\alpha}{2} \|y - (1/\alpha) b_C\|^2 + Z \log \exp(A_C^* y) \right\} \\ &= \arg \min_{y \in \mathbb{R}^m} d(y). \end{aligned}$$

Observe that the solution map P is the *proximal map* corresponding to the logexp function (19) [47, chapter 6]. Under mild conditions—which are met by f —the proximal map is nonexpansive [48, proposition 23.8]. This means that for any two vectors b and b' ,

$$\|y_b - y_{b'}\| \equiv \|P(b_C/\alpha) - P(b'_C/\alpha)\| \leq \frac{1}{\alpha} \|b_C - b'_C\|. \quad (23)$$

Combining this bound with (21), we deduce

$$\|x_b - x_{b'}\| \leq (Z/\alpha) \cdot \|C^{-1}\| \cdot \|A\| \cdot \|b - b'\| \quad (24)$$

for all $b, b' \in \mathbb{R}^m$. In other words, the optimal-solution map (20a) for the primal problem is $(Z\alpha^{-1}\|C^{-1}\| \cdot \|A\|)$ -Lipschitz with respect to the observation vector b .

This bound implies that if b' represents the noiseless data and b represents noisy data, then for a given α , the solution x obtained from noisy data will differ from the noiseless solution x' by no more than the right hand side of (24). As $\alpha \rightarrow \infty$, the solution becomes independent of the data b , illustrating the potential drawback of choosing highly regularized solutions. Similarly, this bound shows the solution is also independent of the data when the data is infinitely noisy, $\|C^{-1}\| \rightarrow 0$.

3.3.3. Relationship between regularization parameter and primal solution. We analyze the sensitivity of the solution map to changes in the regularization parameter α . The solution map associate the input parameters, such as α , with the unique optimal solution to the problem. Here, we examine how variations in α affect the solution.

Consider the dual solution y_α for a fixed regularization parameter α , assuming A and b are fixed. The dual problem can be reformulated by defining the function $F: \mathbb{R}^1 \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ as

$$F(\alpha, y) = \alpha y - b_C + Z A_C D_z \log \exp(z | \mu) |_{z=A^T y}, \quad (25)$$

where the partial gradient with respect to y is

$$\nabla_y F(\alpha, y) = \alpha I + Z A_C D(y) A_C^T, \quad (26)$$

with

$$D(y) = \nabla_z^2 \log \exp(z | \mu) |_{z=A_C^* y} = Z \left(\text{Diag}(x(y)) - x(y)x(y)^T \right), \quad (27)$$

and $x(y)$ defined by (20a). Since $F(\alpha, y_\alpha) = 0$ and $\nabla_y F(\alpha, y)$ is nonsingular for all $\alpha > 0$, we can apply the implicit function theorem [49, theorem 1B.1] to obtain

$$\begin{aligned} y'_\alpha &= -\nabla_y F(\alpha, y_\alpha)^{-1} \nabla_\alpha F(\alpha, y_\alpha) \\ &= -(\alpha I + Z A_C D(y_\alpha) A_C^T)^{-1} y_\alpha, \end{aligned} \quad (28)$$

where y'_α is the derivative of the solution map $\alpha \mapsto y_\alpha \in \arg \min_y d(y | \mu)$. Because the corresponding primal solution $x(y_\alpha)$ lies in the probability simplex Δ_{N_ω} , which is compact, we have $\|y_\alpha\| \leq \alpha^{-1}(\|b_C\| + \|A_C\|)$ by (20b). With $\|(\alpha I + Z A_C S(y_\alpha) A_C^T)^{-1}\| \leq \alpha^{-1}$, we derive

$$\|y'_\alpha\| \leq \frac{\|A_C\| + \|b_C\|}{\alpha^2}.$$

Written as a local perturbation bound, we obtain

$$\|y_\alpha - y_{\alpha'}\| \leq \left(\frac{\|A_C\| + \|b_C\|}{\min(\alpha'^2, \alpha^2)} \right) \cdot |\alpha - \alpha'|.$$

Combining this expression with (21), yields

$$\|x_\alpha - x_{\alpha'}\| \leq Z \cdot \|C^{-1}\|^2 \cdot \|A\| \cdot \left(\frac{\|A\| + \|b\|}{\min(\alpha'^2, \alpha^2)} \right) \cdot |\alpha - \alpha'| \quad (29)$$

for α' near α . This bound shows that as $\alpha, \alpha' \rightarrow 0$, the bound weakens significantly, indicating greater sensitivity to small values of α .

We note that uniform bounds (i.e. globally Lipschitz) are known; however, obtaining them requires more advanced concepts from convex analysis, as discussed by Bonnans and Shapiro [50, §4]. In essence, these bounds differ only by a factor of 2, meaning that, up to this factor, (29) quantifies the impact of different regularization choices α and α' on the solution.

4. Results

In this section, we consider analytic continuation problems relevant to the quark-gluon plasma (QGP) and WDM. We demonstrate that the MEM can be used with the dual Newton optimization algorithm and we compare against the ubiquitous implementation of the MEM with Bryan’s algorithm. The key goals are assessing whether sharp peaks, arising from a spectral function’s poles, can be captured and whether spurious peaks arise in relatively smooth regions. In short, we focus on uncertainty quantification.

4.1. LQCD test problems

We consider the extraction of quarkonia spectral functions from LQCD PIMC simulations. The peaks in a spectral function represent bound states for a quark/anti-quark pair and as temperature increases these peaks broaden. At some critical temperature the peaks smooth away, indicating that the quark and anti-quark cannot form a bound state. In application, knowing the melting temperature of various quarkonia allows experimental facilities like the relativistic heavy ion collider or the large hadron collider to estimate the temperature of the quark gluon plasmas they create. For more details see one of the excellent reviews of recent progress using PIMC on QGP [51–54].

4.1.1. Synthetic correlation samples from zero-temperature ρ meson. For this LQCD benchmark, we consider Asakawa, Nakahara, and Hatsuda’s seminal work which introduced the MEM with Bryan’s algorithm to quarkonia spectral function [10]. In particular, we study the synthetic correlation data generated from the ‘realistic spectral function’ of a charged ρ meson. The charged ρ meson spectral function is parameterized by the relativistic Breit–Wigner function and the relativistic Breit–Wigner’s parameters were fit to reproduce the cross section for e^+e^- annihilating into hadrons [55]. Spectral functions are by definition $S(k=0, \omega)$ and often

denoted as $\rho(\omega)$ in the LQCD community. We consider the spectral function $\rho(\omega) = \omega^2 x(\omega)$, where dimensionless x is given

$$x(\omega) = \frac{2}{\pi} \left(F_\rho^2 \frac{\Gamma_\rho(\omega) m_\rho}{(\omega^2 - m_\rho^2)^2 + \gamma_\rho^2 m_\rho^2} + \frac{1}{8\pi} \left(1 + \frac{\alpha_s}{\pi} \right) \frac{1}{1 + e^{(\omega_0 - \omega)/\delta}} \right). \quad (30)$$

The pole residue informs $F_\rho = m_\rho/g_{\rho\pi\pi}$ and

$$\Gamma_\rho(\omega) = \frac{g_{\rho\pi\pi}^2}{48\pi} m_\rho \left(1 - \frac{4m_\pi^2}{\omega^2} \right)^{3/2} \theta(\omega - 2m_\pi). \quad (31)$$

The empirical values of the constants are given by $m_\rho = 0.77$ GeV, $m_\pi = 0.14$ GeV, $g_{\rho\pi\pi} = 5.45$, $\omega_0 = 1.3$ GeV, $\delta = .02$ GeV. We multiply $x(\omega)$ with kernel $A_{i,j} = e^{\tau_i \omega_j \omega_j^2}$ to produce our noiseless synthetic correlation data b . We discretize using $\omega_j = j \Delta\omega$ where $j \in \{1, \dots, N_\omega\}$, $N_\omega = 600$, and $\Delta\omega = 0.01$ GeV as well as $\tau_i = i \Delta\tau$ where $i \in \{0, \dots, N_\tau - 1\}$, $N_\tau = 30$, and $\Delta\tau = 0.431$ GeV⁻¹. We introduce noise $\eta(\tau_i) \sim \mathcal{N}(0, \sigma_i^2)$ to the clean signal to construct synthetic samples. Following Asakawa *et al*, the variance is given by

$$\sigma_i^2 = \sigma^2 \frac{\tau_i + \epsilon}{\Delta\tau} b_i \quad (32)$$

where σ^2 is a noise level we will vary and the perturbation $\epsilon = 10^{-6}$ is introduced to ensure $\sigma_{i=0} > 0$. Since the noise is proportional to τ the signal is more precise at small τ , which helps the method to resolve features at large ω . Using this noise we generate $N_s = 1000$ samples and estimate the data b and the covariance matrix $C_{i,j} = \delta_{i,j} \sigma_i^2 / \sqrt{N_s}$ from those samples. This differs slightly from Asakawa *et al* who generated a single sample and used the true covariance matrix; our method is more reflective of what is done in practice. We normalize the data b_i and standard deviation σ_i by the first entry of the data b_0 . Anecdotally, this makes $\sum_i x_i < 1$. We use a flat prior $\mu = 0.257$ as our default model and normalize it by b_0 . We then compute a collection of solutions $x_\alpha(\omega)$ over a domain of α 's. From this larger α domain, we select a subdomain $[\alpha_{\min}, \alpha_{\max}]$ where the posterior distribution—derived by Gull [39]— $P(\alpha | b, \mu)$ and the maximum of that distribution $P(\alpha^* | b, \mu)$ are related by the bound $P(\alpha | b, \mu) \leq 0.1 \times P(\alpha^* | b, \mu)$. We compute the estimate \hat{x} from the collection of solutions as

$$\hat{x}_j = \frac{\sum_\alpha x_j(\alpha) P(\alpha | b, \mu)}{\sum_\alpha P(\alpha | b, \mu)}, \quad (33a)$$

where the index α is the α values within $[\alpha_{\min}, \alpha_{\max}]$ and x_α is the corresponding solution. The posterior weighting function is

$$P(\alpha | b, \mu) \propto \exp \left(\alpha S_{\text{SJ}}(x | \mu) - L(x | b) + \frac{1}{2} \sum_k \log \frac{\alpha}{\alpha + \lambda_k} \right). \quad (33b)$$

Here λ_k is the k th eigenvalue of the real symmetric matrix $\Lambda_{j,j'}(\alpha) = \sqrt{x_j(\alpha)} \frac{\partial^2 L}{\partial x_j(\alpha) \partial x_{j'}(\alpha)} \sqrt{x_{j'}(\alpha)}$. In the typical fashion, we estimate the error as $\pm 2 \sqrt{(x^2) - (\hat{x})^2}$.

We reconstruct the solution from data with noise parameter $\sigma^2 = 10^{-3}, 10^{-2}, 10^{-1}$ with $N_\tau = 30$ and find that as $\sigma^2 \rightarrow 0$ the MEM implemented with our dual Newton algorithm and the MEM implemented with Bryan's algorithm recover the same solution. However, our dual Newton MEM and its error estimates are more robust to noise. Bryan's MEM behavior is particularly problematic because its error bands shrink dramatically as the noise grows. This




Figure 1. Plots of $x(\omega)$, i.e. the dimensionless parameterization of the ρ -meson spectral function (30). As well as, the flat Bayesian prior, and the MEM estimates obtained using either Bryan's optimization algorithm (blue) or our dual Newton optimization algorithm (red). Plots keep number of correlator points constant (i.e. $N_\tau = 30$) and vary the noise level σ . From left to right the noise increases as $\sigma^2 = 10^{-4}, 10^{-3}, 10^{-2}$. In the leftmost plot ($\sigma^2 = 10^{-4}$), the estimates and error bands are identical except near $\omega = 0$. In the noiseless limit $\sigma^2 = 10^{-4}$ the different optimization algorithms produce similar curves.

is the exact opposite of the desired behavior. Plots are presented in figure 1, that make it clear when noise is large, MEM with Bryan's algorithm over-regularizes the solution.

To help interpret the reconstructed signals, we investigate the individual α solutions proposed by Bryan's algorithm and by the dual Newton algorithm. We find that, given the same α value, the solution obtained by the dual Newton algorithm deviates further from the default model than the that obtained by Bryan's algorithm. The proposed solutions and the posterior weighting functions are plotted in figure 2 top and middle.

Additionally, we compare the posterior weighting functions and find that the MEM with dual Newton optimization selects solutions which weight the data more (i.e. favoring solutions from lower values of α). Asakawa *et al* also noted this phenomena, but no further investigation was done to overcome it; here our approach overcomes this phenomena. In all cases the posterior weighting function associated with the dual Newton method has a sharper peak at smaller α value than Bryan's primal approach. Comparing posterior weighting function across columns, as noise increases the peak of the posterior weighting function for Bryan's primal approach moves rightward. In comparison, the posterior weighting function for the dual Newton changes little. The posterior weighting functions are plotted in figure 2 Bottom.

For the discretization discussed above, $\|A\| \approx 4 \times 10^2$, $\|C^{-1}\| \approx 1 \times 10^6$, and for all $\sigma^2 = 10^{-4}, 10^{-3}, 10^{-2}$ the Euclidean norm of b is $\|b/b_0\| \approx 1$. We numerically investigate the error bound given in (21) and find that both Bryan's algorithm and the dual newton algorithm scale better than the theoretical bounds and that the scaling is accompanied by a small coefficient $\approx 10^{-4}$. While the error scaling may be similar, the constructed solutions differ, we find that solutions produced by Bryan's algorithm's show greater resemblance to the default model μ for smaller values of α than solution's produced by the dual Newton algorithm. This manifests as increasing in $\|x_{\text{Bryan}} - x_{\text{dual N.}}\|$ with increasing α values. These results are plotted in figure 3.

4.2. WDM test problems

We consider the extraction of DSF from WDM UEG PIMC simulations. The peak in a DSF indicates a quasi-particle excited state known as a plasmon. The peak broadens with increasing temperature and tightens with increasing density. In application, having a model of the DSF allows x-ray Thomson scattering diagnostics at free electron laser facilities to estimate both

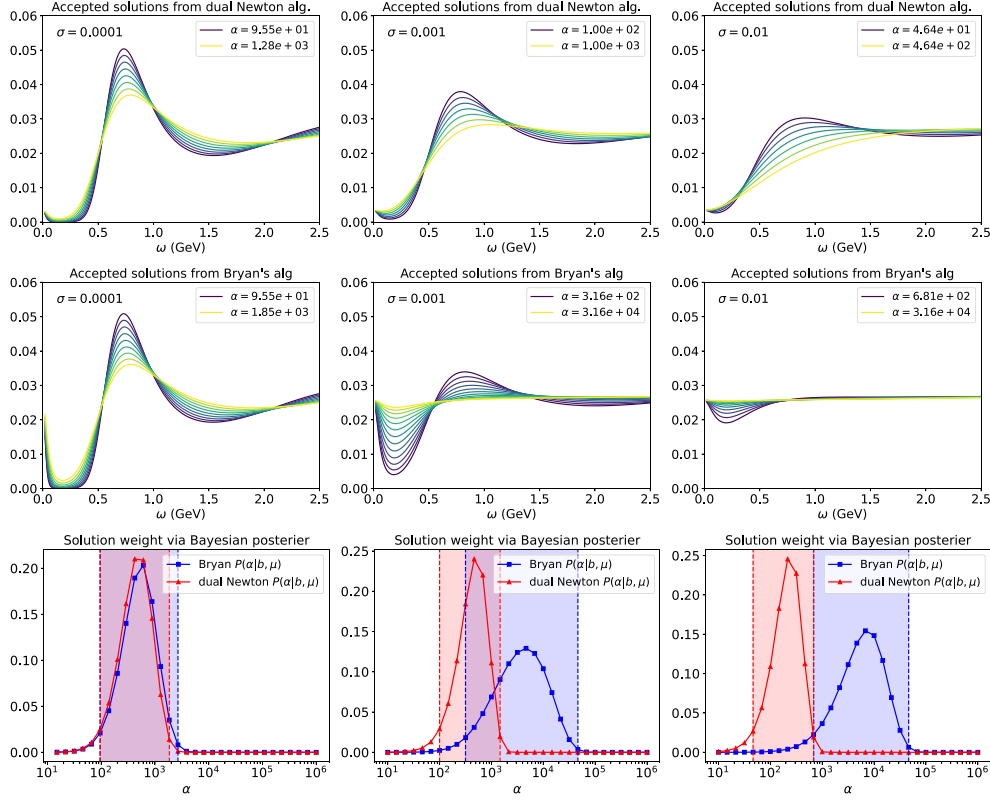


Figure 2. In rows 1 and 2 respectively, we present plots of the dimensionless solutions for $x(\omega)$ proposed by the dual Newton algorithm and Bryan’s algorithm at different regularization weights α . The legend contains only the smallest and largest α values accepted by Gull’s Bayesian posterior procedure and the color gradient indicates the changing α value; see (33) for details on the acceptance condition. All inversions use synthetic ρ -meson data generated from $N_\tau = 30$ and the noise parameter increases from left to right, i.e. $\sigma^2 = 10^{-4}$ (left column), $\sigma^2 = 10^{-3}$ (middle column), $\sigma^2 = 10^{-2}$ (right column). In the bottom row we present the posterior weighting function used to combine the various solutions. The red and blue color bands indicate the α domain $[\alpha_{\min}, \alpha_{\max}]$ accepted by the MEM with dual Newton optimizer and MEM with Bryan’s optimizer respectively.

the density and temperature of their plasmas. For more details see the review of recent progress given by Bonitz *et al* [56].

4.2.1. Synthetic correlation data from UEG DSF. The UEG [42, 57, 58] is one of the most well-studied systems in statistical physics and quantum chemistry, making it very suitable for assessing an algorithm’s performance and uncertainty quantification. A DSF can be produced from a susceptibility χ via the fluctuation dissipation theorem [42], in atomic units, as

$$S(k, \omega) = -\frac{2 \operatorname{Im}\chi(k, \omega)}{n(1 - e^{-\beta\hbar\omega})}, \quad (34)$$

where the susceptibility χ quantifies the gas’s density response to an external potential. The simplest susceptibility model satisfying the frequency sum rule and screening sum rule




Figure 3. We study the theoretical error estimates presented in (21) for the ρ -meson problem. We present the upper bound on the error in solid black; b, b' correspond to noisy ($\sigma^2 = 10^{-3}$) and noiseless data. Since $\|C^{-1}\|$, $\|A\|$, and $\|b - b'\|$ are constants, the solid black line shows a $1/\alpha$ scaling. The marked red and blue curves indicate the deviation between the algorithm's solution constructed from noisy data x and the solution constructed from noiseless data x' ; these two curves are almost identical. Lastly, we plot the deviation between Bryan's solution and the dual Newton's solution from noisy data, indicated by $\|x_{\text{Bryan}} - x_{\text{dualN}}\|$, indicating the magnitude of difference between the two solutions.

is Lindhard's model of the non-interacting Fermi gas with finite temperature corrections $\chi_{\text{Ideal}}^0(k, \omega)$.

$$\chi_{\text{Ideal}}^0(k, \omega) = 2 \int \frac{d^3k}{(2\pi)^3} \frac{f_0(k+q) - f_0(k)}{(\epsilon_{k+q} - \epsilon_k) - \hbar\omega}, \quad (35)$$

where $\epsilon_k = \hbar^2 k^2 / 2m_e$. We used Giuliani and Vignale's introductory textbook on the electron liquid [42] and Tolia *et al*'s paper on non-interacting finite temperature electron gas [59] to implement this function. The more sophisticated CM model also satisfies the frequency sum rule and screening sum rule and includes collisions that respect number and momentum [11]. For the single species electron gas the CM model is given

$$\chi_{\text{CM}}^0(k, \omega + i\nu) = \frac{\chi_{\text{Ideal}}^0(k, \omega + i\nu)}{1 - \frac{i}{\omega\tau} \left(1 - \frac{\chi_{\text{Ideal}}^0(k, \omega + i\nu)}{\chi_{\text{Ideal}}^0(k, 0)} \right) - \frac{i m \omega}{k^2 n_0 \tau} \chi_{\text{Ideal}}^0(k, \omega + i\nu)}, \quad (36)$$

where ν is the collision frequency. For both χ_{Ideal}^0 and χ_{CM}^0 the mean field interactions are included as [60–62]

$$\chi(k, \omega + i\nu) = \frac{\chi^0(k, \omega + i\nu)}{1 - v(k) \chi^0(k, \omega + i\nu)}, \quad (37)$$

where $\nu = 0$ for the ideal susceptibility. We take the particle-particle interaction $v(k)$ as a screened coulomb $v(k) = 4\pi e^2 / (k^2 + \kappa^2)$. For this UEG, we use the un-screened limit (i.e. $\kappa = 0$).

In this test problem, we generate synthetic samples by multiplying the Laplace kernel with the mean field corrected CM DSF model. We consider the UEG at density $n_e =$

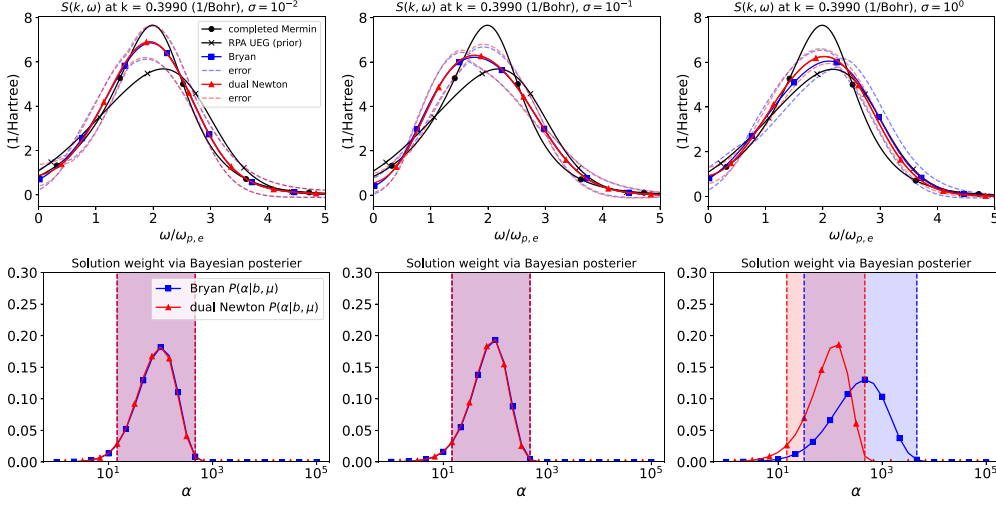


Figure 4. Plots of the estimated dynamic structure factor for the synthetic UEG example with associated posterior weighting function. From left to right, the noise level σ increases as $\sigma = 10^{-2}, 10^{-1}, 10^0$. The top row contains the plots of the DSF. The completed Mermin (CM) model is the input spectrum and desired recovery, defined in (36) with mean field correction (37) included. The RPA UEG is the Bayesian prior, defined in (35) with mean field correction (37) included. The MEM estimates are in either red or blue. On the bottom row the Bayesian posterior is plotted as a function of α . We logarithmically sample alpha in $\alpha \in [10^0 - 10^5]$. The red and blue vertical color bands indicates the α domain is selected for either MEM averaging procedure; see (33) for details on how the α domain was selected.

$1.611 \times 10^{21} \text{ cc}^{-1}$ which corresponds to a Wigner-seitz radius $r_s = 10$ Bohr. The temperature is $k_B T = 0.5 \text{ eV}$, equivalently $\beta = 1/k_B T = 54.301$ Hartree, which corresponds to a quantum degeneracy parameter of $\Theta = k_B T/E_F = 1$; these conditions match the conditions of the PIMC UEG simulation, whose correlation estimates will be considered next. For this system the electron plasma frequency is $\omega_{p,e} = \sqrt{3/r_s^3} = 0.05476 \text{ Hartree}$ and Fermi wavenumber is $k_{F,e} = (9\pi/4)^{1/3}/r_s = 0.1919 \text{ Bohr}^{-1}$. The collision frequency ν is set equal to the electron plasmon frequency $\omega_{p,e}$. We discretize our frequency as $\omega_j = j\Delta\omega$ where $j \in \{1, \dots, N_\omega\}$, $N_\omega = 1250$, $\Delta\omega = 0.0011$ Hartree. Additionally, we discretize our imaginary time as $\tau_i = i\Delta\tau$ where $i \in \{0, \dots, N_\tau - 1\}$, $N_\tau = 201$, $\Delta\tau = 0.271505$ Hartree. We introduce noise different from (32); we chose a standard deviation of

$$\sigma_i = \sigma b_i. \quad (38)$$

This noise does not grow proportionally to τ and is more reflective of our authentic PIMC data. We generate $N_s = 1000$ samples of $F(\tau)$ and from them estimate the data b and the error matrix $C_{i,j} = \delta_{i,j}\sigma_i^2/\sqrt{N_s}$. We conduct the inversions using the ideal gas with mean field correction (i.e. RPA) as the default model μ . We follow the same Bayesian posterior weighting procedure described for (33) to average our collection of solutions into an estimate.

We first consider the DSF estimate for a small wavenumber $k = 0.3990 \text{ Bohr}^{-1}$ and noise levels $\sigma = 10^{-2}, 10^{-1}, 10^0$ (see figure 4). In the limit of small noise, both the MEM with Bryan's algorithm and the MEM with dual Newton converge to the desired CM $S(k, \omega)$. However, as noise increases the peak of the posterior weighting function for Bryan's primal

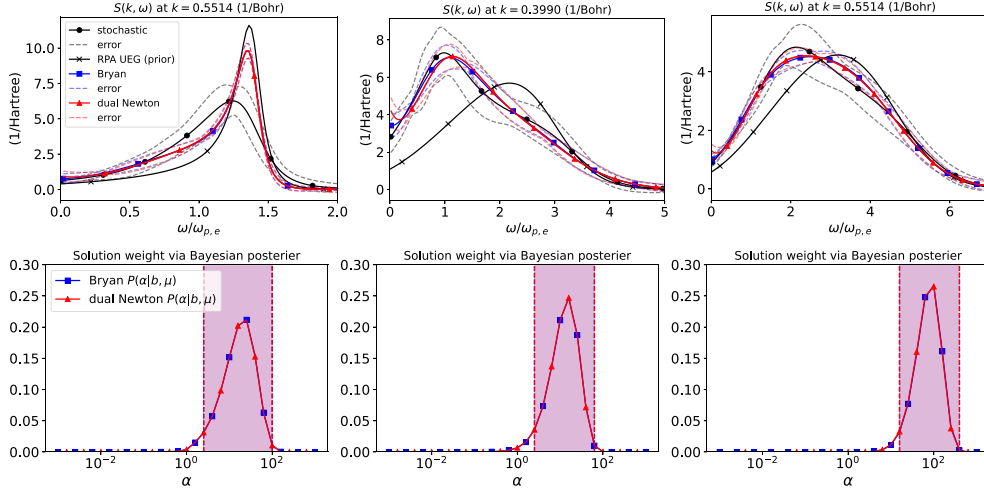


Figure 5. Top: from left to right we vary the wavenumber k considered in our DSF, for all k values, the number of correlator points is fixed at 201 and the relative noise is $\approx 10^{-3}$. The curve labeled ‘stochastic’ is Tobias *et al*’s estimate [29], not the exact solution. The RPA UEG is the Bayesian prior, defined in (35), with the mean-field correction (37) included. Bottom: plot of the posterior weighting function associated with the above estimates. We collect solutions for $\alpha \in [10^{-3} - 10^3]$. The red and blue color bands indicate the α domain kept by the MEM with dual Newton optimizer and MEM with Bryan’s optimizer respectively; see (33) for details on how the α domain was selected.

approach moves rightward. In comparison, the posterior weighting function for the dual Newton changes little. This indicates that Bryan’s algorithm breaks down, over-regularizing the solution by selecting weighting the larger α values. This is the same behavior seen in the ρ -meson investigation.

4.2.2. Authentic correlation data from UEG PIMC simulation. The authentic WDM UEG PIMC correlation data we study arises from Dornheim and Vorberger’s work [63]. They employed PIMC methods to produce Euclidean time correlation functions of the UEG at WDM conditions ($T = 0.5\text{eV}$ and $r_s = 10$). This PIMC UEG data was analyzed by Dornheim *et al* [29], who used a stochastic sampling algorithm to infer the DSF. Their stochastic sampling algorithm is a forward problem approach that relies on sampling values for a parameterized form of the susceptibility; the parameterized form is determined via sum rules and additional constraints. By comparison, our inverse method has no similar guarantee of satisfying sum rules, but will work in systems where a parameterized form of the susceptibility is not known and hence has a broader range of applicability. In this section, we include the stochastic sampling results in the slot where the CM was in synthetic results because the general shape of the DSF is physical even if the stochastic estimate is not the true solution.

We present plots of the reconstructed $S(k, \omega)$ at a small, middle and large wavenumber $k = 0.1702, 0.3990, 0.5513 \text{ Bohr}^{-1}$. These plots show excellent agreement with the stochastic method at large k , but poorer agreement at small k . Anecdotally, we found that if the Bayesian prior is sharply peaked, then the solutions deviate less from the prior. Furthermore, both MEM approaches produce solutions that have less ringing than the stochastic model; ringing at large ω is a common concern for the BRM entropic approach [7]. Plots are given in figure 5 top row.




Figure 6. Plots of the periodic Laplace transformed DSF estimates presented in the top row of figure 5 (transform defined in (4)); all curves are normalized by the ITCF data at $\tau = 0$ and symmetric about $\tau = \beta/2$. We include the imaginary time correlation function (ITCF) data obtained via PIMC UEG data for comparison as well as the associated reduced chi-square χ^2/N value. The periodic Laplace transform of the MEM DSF estimates are visually indistinguishable from the observed ITCF data.

We also present the normalized Bayesian posterior used to combine the α -solutions into an estimate. Since the error in the correlation function is small, approximately 10^{-3} , we do not expect that Bryan’s Bayesian posterior will differ from the dual Newton’s Bayesian posterior. This is reflected in figure 5 bottom row and explains why the MEM with Bryan and the MEM with dual Newton estimates are indistinguishable from the PIMC data.

We plot the correlation functions that are produced from the reconstructed DSFs and compare to the PIMC UEG data in figure 6. This helps us compare to the stochastic sampling method which accepted and averaged over DSF’s based on the proposed DSF’s ability to ‘reproduce the [correlation function] within the Monte Carlo error bars’. We see that both MEM with the dual Newton and Bryan’s optimization, as well as the stochastic estimates, produce correlation functions that are indistinguishable from the data. Comparatively, the RPA clearly disagrees with the ITCF. This indicates that the solution is not so heavily regularized that the entropic prior dominates the cost function. We also supply the associated reduced chi-squared values $\chi^2/N_\tau = (Ax - b)C^{-1}(Ax - b)/N_\tau$ in figure 6. The MEM, $\chi^2/N \approx \mathcal{O}[10^{-3}]$, vastly outperforms the stochastic approach, $\chi^2/N \approx \mathcal{O}[1]$.

We also investigate the full DSF $S(k, \omega)$ estimate and find that both Bryan’s algorithm and the dual Newton algorithm produces DSF’s that contain a dip in the dispersion relation occurring for $k \in [0.25, 0.45] \text{ Bohr}^{-1}$, which is known as a roton feature [64]. Comparing to the stochastic approach the MEM approaches do not produce as smooth of a dispersion relation. There are two abrupt jumps in the dispersion relations, i.e. two peaks that sit at larger ω than their neighboring peaks. This feature arises from two sources. Firstly, uncertainty in the data; a resampling algorithm would smoothen final estimate and propagate uncertainty into the final solution. This will be the subject of future work. Secondly, Gull’s Bayesian posterior is picking some compromise between the data and the RPA DSF and the RPA does not have a roton-type feature, so there is a tendency to be pulled upward. This serves as a reminder that all forms of entropic regularization produce a unique solution that is biased towards the Bayesian prior. Two possible ways to address bias are improving the default model, e.g. using the *static approximation* in the case of the warm dense UEG [65] or conducting the entire $S(k, \omega)$ inversion at once rather than divided into individual k inversions. Plots of the full DSF are given in figure 7.




Figure 7. Heatmaps of the dynamic structure factor $S(k, \omega)$. The $S(k, \omega)$ estimate from [63] is labeled stochastic sampling. The Bayesian prior, indicated RPA is defined in (35) with mean field correction (37) included. The estimates obtained by MEM with Bryan’s algorithm and the estimates obtained by MEM with the dual Newton optimization are also included. Dashed white lines have been placed at $k = 0.170153, 0.399044,$ and $0.551358 \text{ Bohr}^{-1}$. These lines indicate the k values plotted in figures 5 and 6. For each k value the location of the DSF peak is indicated with a black dot; together they indicate the dispersion relation.

5. Conclusions

In this work we introduce the dual formulation of the entropy regularized least squares cost function and analytic error bounds on the dual estimate. The advantage of the dual formulation is that optimizing the dual produces the same solution as optimizing primal problem, but the dual cost function has guaranteed strong convexity, differentiability, and lower dimension. This makes the dual problem robust to noise, tractable to second order optimization methods, and computationally cheaper. Of particular importance is that the dual optimization problem is formulated in the kernel’s row space, and thus the resolution of the spectral function does not impact the dimension of the dual optimization. In short, dual formulation has clear advantages for the analytic continuation problem, where the kernel is ill-conditioned and the spectral function is continuous. The error bounds establish that the dual optimization produces corresponding primal solutions that converge to the desired minimum. Furthermore, these error estimates provide insight into how a solution depends on the noise and on the regularization parameter α . Numeric tests showed that the theoretical upper bound on the error was many orders of magnitude larger than our actual error.

Many advantages of our dual approach are particular to the problem formulation that has been given in section 3.2. These advantages do not carry over to all of the regularization terms which are currently being used in the literature (e.g. Lasso L1 regularization [66]). Further, in the rare case where a closed form of the solution is known, iterative optimizations in either a dual formulation or a primal formulation will not outperform the closed solution (e.g. ridge L2 regularization, the BG method [12, 13], and its smooth variant [14]). The analytic continuation community needs to consider whether their problem formulation may benefit from a dual formulation. Further, the analytic continuation community needs to consider whether their problem (e.g. estimating a multi-peak function) will benefit. In particular, Goulko *et al* [67] established that the width of a second high frequency peak cannot be resolved without tending toward noiseless data, regardless of approach. They argue this is due to the nature of the Laplace kernel and similar results are demonstrated in [68].

We have compared our dual Newton algorithm with Bryan’s heavily cited algorithm. In section 3, we discuss two theoretical advantages of our approach. First, unlike Bryan’s

algorithms we search for the best positive normalization of the estimate. We argue that this capability is important for the periodic Laplace transform (i.e. finite temperature systems). Second, our approach optimizes in a reduced search space without sacrificing singular basis vectors. This addresses Rothkopf's results which showed that removing singular vectors prevents the MEM from representing sharply peaked functions at large ω [9]. In section 4, we demonstrate that when the noise is small both algorithms yield the same mean and variance estimates, but in presence of large noise our algorithm differs from Bryan's algorithm. Knowing that Bryan's algorithm makes an assumption that the gradient of the fidelity term lies in the singular space while our algorithm does not, these differences suggest that noise also violates Bryan's assumption. In summary, Bryan's algorithm relies on assumptions that the noise is small and the desired function is sharply peaked at small ω . As a result we echo Rothkopf's statement, 'either systematically extend the search space within the MEM or abandon the MEM in favor of one of the many modern Bayesian approaches developed over the past two decades.' We have expanded the search space within MEM and in doing so we have not seen the introduction of prominent ringing artifacts.

Future applications of this method are broad; the dual MEM routine used in this work can handle a generic kernel, so it can be readily applied to the applications mentioned in the introduction section. For the plasma physics community, future work physics-oriented work could apply this method to other number densities and temperatures to observe how the roton feature varies with these parameters. Future methods-oriented work could extend the cost function and solve for all of $S(k, \omega)$ simultaneously, rather than solving each k individually. The same sentiment is shared by the Lattice QCD community, which wants to estimate spectral functions $\rho(\omega, T)$ for all temperatures simultaneously rather than a given T . In both cases, solving for all values simultaneously would enhance smoothness across the additional variable. Additionally, for the plasma physics community this work has shown that entropic regularization can invert the Laplace transform, which matches sentiments of other fields and opens the possibility of conducting instrument deconvolution using the Laplace transform instead of the Fourier transform. Lastly, future work may wish to improve over the Newton-conjugate gradient method discussed in [appendix](#) or explore alternatives to using the discovered positive normalization of x for the default model μ .

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://github.com/chunatho/InverseProblemSolvers>.

Acknowledgments

Thomas Chuna thanks Matt King-Roskamp for his expertise and insightful discussions regarding what constitutes a proper problem formulation in mathematics. This work was partially supported by the Center for Advanced Systems Understanding (CASUS) which is financed by Germany's Federal Ministry of Education and Research (BMBF) and by the Saxon state government out of the State budget approved by the Saxon State Parliament. This work has received funding from the European Union's Just Transition Fund (JTF) within the project *Röntgenlaser-Optimierung der Laserfusion* (ROLF), Contract Number 5086999001, co-financed by the Saxon state government out of the State budget approved by the Saxon State Parliament. This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2022 research and innovation programme (Grant

Agreement No. 101076233, ‘PREXTREME’). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. The PIMC calculations were partly carried out at the Norddeutscher Verbund für Hoch- und Höchstleistungsrechnen (HLRN) under grant mvp00024, and on a Bull Cluster at the Center for Information Services and High Performance Computing (ZIH) at Technische Universität Dresden. M.P. Friedlander is partially supported by a Discovery grant from the Natural Sciences and Engineering Council of Canada.

Author contributions

Thomas Chuna  0000-0001-8400-4495

Conceptualization (equal), Investigation (lead), Software (equal), Visualization (equal), Writing – original draft (equal), Writing – review & editing (equal)

Nicholas Barnfield  0009-0003-9895-8347

Formal analysis (equal), Writing – original draft (equal)

Tobias Dornheim  0000-0001-7293-6615

Data curation (equal), Funding acquisition (equal), Writing – review & editing (equal)

Michael P Friedlander  0000-0003-0222-5222

Methodology (equal), Software (equal), Writing – original draft (equal), Writing – review & editing (equal)

Tim Hoheisel  0000-0002-0782-6302

Conceptualization (equal), Funding acquisition (equal), Methodology (equal), Writing – original draft (equal)

Appendix. Description of the dual algorithm

This appendix outlines the Newton–Krylov method [69, 70] used to solve the MEM problem presented in (9). The purpose of this method is to obtain the scaling \bar{Z} that minimizes the value function $V(Z)$ defined by (13) and to obtain a corresponding primal solution \bar{x} of (9). A high level description is written below and algorithm A1 summarizes the full method. The dual Newton algorithm described in this appendix can be found on github ([repository](#)) [46].

Duality theory establishes a direct relationship between the primal and dual solutions through smooth mappings, as given by (20). Solving the dual problem yields the same optimal value as the primal problem; however, because the dual cost function is strongly convex and differentiable, it is amenable to a simpler optimization process and fast convergence [71].

The underlying duality theory that gives rise to the smooth primal-dual maps implies that

$$V(Z) = \min \{p(x, Z | \mu/Z) | x \in Z\Delta_{N_\omega}\} = - \min \{d(y, Z | \mu/Z) | y \in \mathbb{R}^{N_\tau}\}.$$

Since the dual objective $d(y, Z | \mu/Z)$ is strongly convex in y , the dual minimizer \bar{y}_Z , which depends on the scaling Z , is unique and the Hessian $\nabla_y^2 d(\bar{y}_Z, Z | \mu/Z)$ is positive definite and hence nonsingular. Thus,

$$V(Z) = -d(\bar{y}_Z, Z | \mu/Z)$$

Algorithm 1. Dual maximum entropy method (MEM) algorithm.

Input: Prior $\mu \in \mathbb{R}_+^{N_\omega}$, observations $b \in \mathbb{R}^{N_\tau}$, operator A , covariance C^{-1} , tolerances ϵ_a, ϵ_r
Output: Optimal scaling \bar{Z} , dual solution \bar{y} , primal solution $x(\bar{y})$

```

1  $Z_0 \leftarrow \sum_{i=1}^{N_\omega} \mu_i$  //Initial guess
2  $y_0 \leftarrow 0$  //Initial dual point
//Root finding for  $V'(Z) = 0$  using hybrid method
3 while not converged do
| //Inner optimization: minimize  $d(y, Z | \mu/Z)$  with  $Z$  fixed
4  $y \leftarrow y_0$  //Warm start from previous iteration
| while  $\|\nabla_y d(y, Z | \mu/Z)\| > \epsilon_a + \epsilon_r \|b\|$  do
6 | //Trust-region Newton-Krylov (Trunk) method
| | Solve Newton system:  $\nabla_y^2 d(y, Z | \mu/Z) \Delta y = -\nabla_y d(y, Z | \mu/Z)$ 
7 | | Update  $y$  with trust-region step
8 | end
9 |  $\bar{y}_Z \leftarrow y$  //Store minimizer for current  $Z$ 
10 | Compute  $V'(Z) = -\log \langle \exp(A^* \bar{y}_Z) | \mu \rangle + \log Z + 1$ 
11 | if  $|V'(Z)| \leq \epsilon_a + \epsilon_r \|b\|$  then
12 | | break //Root found
13 | end
14 | Update  $Z$  using hybrid root-finding step
| //Combines secant/quadratic steps with bracketing if encountered
15 |  $y_0 \leftarrow \bar{y}_Z$  //Save for warm start
16 end
17  $\bar{Z} \leftarrow Z$ 
18  $\bar{y} \leftarrow \bar{y}_Z$ 
19  $x(\bar{y}) \leftarrow$  primal solution via smooth map (20a)

```

is differentiable. Applying the chain rule, one deduces that its derivative with respect to Z is

$$V'(Z) = -\log \exp(A^* \bar{y}_Z | \mu) + \log Z + 1. \quad (\text{A.1})$$

Observe that the second derivative $V''(Z) = 1/Z$ is positive for all positive scalings Z , and thus V is strictly convex and achieves its minimum value when

$$V'(Z) = 0. \quad (\text{A.2})$$

Minimizing $V(Z)$ over Z , one obtains an optimal pair (\bar{y}, \bar{Z}) and consequently, the corresponding primal solution $x(\bar{y})$.

The optimization over Z is conducted using a hybrid root-finding method [72] to obtain a root of the equation (A.2), which corresponds to the optimal Z . Each iteration of the root-finding algorithm requires solving the dual problem (17) using a trust-region Newton–Krylov method implemented by the Julia package `JSOSolvers.jl` [73]. In practice, the dual problem becomes easier to solve at each iteration because the current dual solution can be used to warm start the subsequent dual solve. The linear Krylov solver requires only matrix-vector products with the Hessian of (17) (i.e. $\nabla_y^2 d(\bar{y}_Z, Z | \mu/Z)$). The overall computational cost of the resulting algorithm is dominated by products with the operator A and its adjoint A^* .

For this class of problems, the Newton–Krylov method is guaranteed to produce a sequence of dual iterates that converge to the unique solution. Locally, Newton’s method is expected to converge quadratically, although inexact solves of the Newton equation may cause the

method to converge only superlinearly. Additionally, this computation generates a corresponding primal sequence $\{x(y^k)\}$, defined by the smooth map (20a), which converges to the unique global solution $x(y^*) \in \Delta_n$, as required. We terminate the algorithm when the gradient

$$F(y, Z) = \begin{bmatrix} \nabla_y d(y, Z | \mu) \\ V'(Z) \end{bmatrix} \quad (\text{A.3})$$

drops below absolute and relative tolerances, ϵ_a and ϵ_r , as $\|F(y, Z)\| \leq \epsilon_a + \epsilon_r \|b\|$. Our experiments all use $\epsilon_a = \epsilon_r = 10^{-6}$.

References

- [1] Epstein C L and Schotland J 2008 The bad truth about Laplace's transform *SIAM Rev.* **50** 504–20
- [2] Bryan R K 1990 Maximum entropy analysis of oversampled data problems *Eur. Biophys. J.* **18** 165–74
- [3] Gubernatis J E, Jarrell M, Silver R N and Sivia D S 1991 Quantum Monte Carlo simulations and maximum entropy: dynamics from imaginary-time data *Phys. Rev. B* **44** 6011
- [4] Kim S, Petreczky P and Rothkopf A 2018 Quarkonium in-medium properties from realistic lattice NRQCD *J. High Energy Phys.* **JHEP11(2018)088**
- [5] Tripolt R-A, Gubler P, Ulybyshev M and Smekal L V 2019 Numerical analytic continuation of Euclidean data *Comput. Phys. Commun.* **237** 129–42
- [6] Yang S, Du L and Huang Li 2023 Combining Bayesian reconstruction entropy with maximum entropy method for analytic continuations of matrix-valued Green's functions (arXiv:2401.00018)
- [7] Fischer C S, Pawłowski J M, Rothkopf A and Welzbacher C A 2018 Bayesian analysis of quark spectral properties from the Dyson-Schwinger equation *Phys. Rev. D* **98** 014009
- [8] Fuchs S, Pruschke T and Jarrell M 2010 Analytic continuation of quantum Monte Carlo data by stochastic analytical inference *Phys. Rev. E* **81** 056701
- [9] Rothkopf A 2020 Bryan's maximum entropy method—diagnosis of a flawed argument and its remedy *Data* **5** 85
- [10] Asakawa M, Nakahara Y and Hatsuda T 2001 Maximum entropy analysis of the spectral functions in lattice QCD *Prog. Part. Nucl. Phys.* **46** 459–508
- [11] Chuna T and Murillo M S 2025 Conservative dielectric functions and electrical conductivities from the multicomponent Bhatnagar–Gross–Krook equation *Phys. Rev. E* **111** 035206
- [12] Backus G and Gilbert F 1968 The resolving power of gross earth data *Geophys. J. Int.* **16** 169–205
- [13] Backus G and Gilbert F 1970 Uniqueness in the inversion of inaccurate gross earth data *Phil. Trans. R. Soc. A* **266** 123–92
- [14] Hansen M, Lupo A and Tantalo N 2019 Extraction of spectral densities from lattice correlators *Phys. Rev. D* **99** 094508
- [15] Rothkopf A 2013 Improved maximum entropy analysis with an extended search space *J. Comput. Phys.* **238** 106–14
- [16] Burnier Y and Rothkopf A 2013 Bayesian approach to spectral function reconstruction for Euclidean quantum field theories *Phys. Rev. Lett.* **111** 182003
- [17] Kim S, Petreczky P and Rothkopf A 2015 Lattice NRQCD study of s- and p-wave bottomonium states in a thermal medium with $n_f = 2 + 1$ light flavors *Phys. Rev. D* **91** 054511
- [18] Boninsegni M and Ceperley D M 1996 Density fluctuations in liquid ^4He . path integrals and maximum entropy *J. Low Temp. Phys.* **104** 339–57
- [19] Ferré G and Boronat J 2016 Dynamic structure factor of liquid He^4 across the normal-superfluid transition *Phys. Rev. B* **93** 104510
- [20] Kora Y and Boninsegni M 2018 Dynamic structure factor of superfluid ^4He from quantum Monte Carlo: maximum entropy revisited *Phys. Rev. B* **98** 134509
- [21] Dornheim T, Moldabekov Z A, Vorberger J and Militzer B 2022 Path integral Monte Carlo approach to the structural properties and collective excitations of liquid He-3 without fixed nodes *Sci. Rep.* **12** 708
- [22] Yoon H, Sim J-H and Han M J 2018 Analytic continuation via domain knowledge free machine learning *Phys. Rev. B* **98** 245101

- [23] Fournier R, Wang L, Yazyev O V and Wu Q 2020 Artificial neural network approach to the analytic continuation problem *Phys. Rev. Lett.* **124** 056401
- [24] Beach K S D 2004 Identifying the maximum entropy method as a special limit of stochastic analytic continuation (arXiv:0403055 [cond-mat])
- [25] Shu H-T, Ding H-T, Kaczmarek O, Mukherjee S and Ohno H 2015 A stochastic approach to the reconstruction of spectral functions in lattice QCD (arXiv:1510.02901)
- [26] Kotliar G, Savrasov S Y, Haule K, Oudovenko V S, Parcollet O and Marianetti C A 2006 Electronic structure calculations with dynamical mean-field theory *Rev. Mod. Phys.* **78** 865–951
- [27] Fei J, Yeh C-N and Gull E 2021 Nevanlinna analytical continuation *Phys. Rev. Lett.* **126** 056402
- [28] Tkachenko I M, Arkhipov Y V and Askaruly A 2012 *The Method of Moments and its Applications in Plasma Physics* (Akademikerverlag)
- [29] Dornheim T, Groth S, Vorberger J and Bonitz M 2018 *Ab initio* path integral Monte Carlo results for the dynamic structure factor of correlated electrons: from the electron liquid to warm dense matter *Phys. Rev. Lett.* **121** 255001
- [30] Groth S, Dornheim T and Vorberger J 2019 *Ab initio* path integral Monte Carlo approach to the static and dynamic density response of the uniform electron gas *Phys. Rev. B* **99** 235122
- [31] Shore J and Johnson R 1980 Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy *IEEE Trans. Inf. Theory* **26** 26–37
- [32] Johnson R and Shore J 1983 Comments on and correction to ‘axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy’ (jan 80 26-37)[corresp.] *IEEE Trans. Inf. Theory* **29** 942–3
- [33] Skilling J 1988 The axioms of maximum entropy *Maximum-Entropy and Bayesian Methods in Science and Engineering: Foundations* (Springer) pp 173–87
- [34] Jarrell M and Gubernatis J E 1996 Bayesian inference and the analytic continuation of imaginary-time quantum Monte Carlo data *Phys. Rep.* **269** 133–95
- [35] Rioux G, Choksi R, Hoheisel T, Marechal P and Scarvelis C 2020 The maximum entropy on the mean method for image deblurring
- [36] Maréchal P and Lannes A 1997 Unification of some deterministic and probabilistic methods for the solution of linear inverse problems via the principle of maximum entropy on the mean *Inverse Problems* **13** 135
- [37] Marechal P 1998 On the principle of maximum entropy on the mean as a methodology for the regularization of inverse problems *Submitted to the Proc. 7th Vilnius Conf. on Probability Theory and 22nd European Meeting of Statisticians* vol 18
- [38] Besnerais G L, Bercher J-F and Demoment G 1999 A new look at entropy for solving linear inverse problems *IEEE Trans. Inf. Theory* **45** 1565–78
- [39] Gull S F 1989 Developments in maximum entropy data analysis *Maximum Entropy and Bayesian Methods* (Springer) pp 53–71
- [40] Chuna T 2025 Inverse problem solvers (available at: <https://github.com/chunatho/InverseProblemSolvers>)
- [41] Ichimaru S 1991 *Statistical Plasma Physics (Basic Principles)* vol 1 (CRC Press)
- [42] Giuliani G and Vignale G 2008 *Quantum Theory of the Electron Liquid* (Cambridge University Press)
- [43] Kapur J N and Kesavan H K 1992 Entropy Optimization Principles and Their Applications *Entropy and Energy Dissipation in Water Resources (Water Science and Technology Library)* vol 9, ed V P Singh, V P Singh and M Fiorentino (Springer) pp 3–20
- [44] Nocedal J and Wright S J 2006 *Numerical Optimization* 2 edn (Springer)
- [45] Rockafellar R T and Wets R J-B 2009 *Variational Analysis* vol 317 (Springer)
- [46] Friedlander M P 2025 Dualperspective (available at: <https://github.com/MPF-Optimization-Laboratory/DualPerspective.jl>)
- [47] Beck A 2017 First-Order Methods in Optimization *Siam-Society for Industrial and Applied Mathematics*
- [48] Bauschke H H and Combettes P L 2017 Convex Analysis and Monotone Operator Theory in Hilbert Spaces *CMS Books in Mathematics* (Springer)
- [49] Rockafellar R T and Dontchev L A 2014 *Implicit Functions and Solution Mappings: A view from variational analysis (Springer Monographs in Mathematics)* (Springer)
- [50] Frédéric Bonnans J and Shapiro A 2000 *Perturbation Analysis of Optimization Problems (Springer Series in Operations Research)* (Springer)
- [51] Petreczky P 2012 Lattice QCD at non-zero temperature *J. Phys. G: Nucl. Part. Phys.* **39** 093002

- [52] Bazavov A 2015 Lattice QCD at non-zero temperature (arXiv:[1505.05543](#))
- [53] Datta S 2015 Quarkonia at finite temperature in relativistic heavy-ion collisions *Pramana* **84** 881–99
- [54] Rothkopf A 2020 Heavy quarkonium in extreme conditions *Phys. Rep.* **858** 1–117
- [55] Shuryak E V 1993 Correlation functions in the QCD vacuum *Rev. Mod. Phys.* **65** 1–46
- [56] Bonitz M, Dornheim T, Moldabekov Z A, Zhang S, Hamann P, Kählert H, Filinov A, Ramakrishna K and Vorberger J 2020 *Ab initio* simulation of warm dense matter *Phys. Plasmas* **27** 042710
- [57] Dornheim T, Groth S and Bonitz M 2018 The uniform electron gas at warm dense matter conditions *Phys. Rep.* **744** 1–86
- [58] Loos P-F and Gill P M W 2016 The uniform electron gas *Comput. Mol. Sci* **6** 410–29
- [59] Tolias P, Dornheim T and Vorberger J 2024 On the density-density correlations of the non-interacting finite temperature electron gas (arXiv:[2410.22942](#))
- [60] Bohm D and Pines D 1951 A collective description of electron interactions. I. Magnetic interactions *Phys. Rev.* **82** 625–34
- [61] Pines D and Bohm D 1952 A collective description of electron interactions: II. Collective vs individual particle aspects of the interactions *Phys. Rev.* **85** 338–53
- [62] Bohm D and Pines D 1953 A collective description of electron interactions: III. Coulomb interactions in a degenerate electron gas *Phys. Rev.* **92** 609–25
- [63] Dornheim T and Vorberger J 2020 Finite-size effects in the reconstruction of dynamic properties from *ab initio* path integral Monte Carlo simulations *Phys. Rev. E* **102** 063301
- [64] Dornheim T, Moldabekov Z, Vorberger J, Kählert H and Bonitz M 2022 Electronic pair alignment and roton feature in the warm dense electron gas *Commun. Phys.* **5** 304
- [65] Dornheim T, Cangi A, Ramakrishna K, Böhme M, Tanaka S and Vorberger J 2020 Effective static approximation: a fast and reliable tool for warm-dense matter theory *Phys. Rev. Lett.* **125** 235001
- [66] Otsuki J, Ohzeki M, Shinaoka H and Yoshimi K 2017 Sparse modeling approach to analytical continuation of imaginary-time quantum Monte Carlo data *Phys. Rev. E* **95** 061302
- [67] Goulko O, Mishchenko A S, Pollet L, Prokof'ev N and Svistunov B 2017 Numerical analytic continuation: Answers to well-posed questions *Phys. Rev. B* **95** 014102
- [68] Shao H and Sandvik A W 2023 Progress on stochastic analytic continuation of quantum Monte Carlo data *Phys. Rep.* **1003** 1–88
- [69] Kelley C T 2003 Solving Nonlinear Equations with Newton's Method *Society for Industrial and Applied Mathematics*
- [70] Knoll D A and Keyes D E 2004 Jacobian-free Newton-Krylov methods: a survey of approaches and applications *J. Comput. Phys.* **193** 357–97
- [71] Nesterov Y 2018 Lectures on Convex Optimization *Springer Optimization and Its Applications* vol 137 (Springer)
- [72] JuliaMath Contributors 2024 Roots.jl: root finding functions for Julia
- [73] Migot T, Orban D and Soares Siqueira A contributors 2023 JSOSolvers.jl: JuliaSmoothOptimizers optimization solvers