

Non-Hertz-Millis scaling of the anti-ferromagnetic quantum critical metal via scalable Hybrid Monte Carlo

Received: 8 June 2022

Accepted: 23 March 2023

Published online: 03 May 2023

 Check for updatesPeter Lunts ^{1,2} ✉, Michael S. Albergo ^{3,5} & Michael Lindsey^{4,5}

A key component of the phase diagram of many iron-based superconductors and electron-doped cuprates is believed to be a quantum critical point (QCP), delineating the onset of antiferromagnetic spin-density wave order in a quasi-two-dimensional metal. The universality class of this QCP is believed to play a fundamental role in the description of the proximate non-Fermi liquid behavior and superconducting phase. A minimal model for this transition is the O(3) spin-fermion model. Despite many efforts, a definitive characterization of its universal properties is still lacking. Here, we numerically study the O(3) spin-fermion model and extract the scaling exponents and functional form of the static and zero-momentum dynamical spin susceptibility. We do this using a Hybrid Monte Carlo (HMC) algorithm with a novel auto-tuning procedure, which allows us to study unprecedentedly large systems of 80×80 sites. We find a strong violation of the Hertz-Millis form, contrary to all previous numerical results. Furthermore, the form that we do observe provides good evidence that the universal scaling is actually governed by the analytically tractable fixed point discovered near perfect “hot-spot” nesting, even for a larger nesting window. Our predictions can be directly tested with neutron scattering. Additionally, the HMC method we introduce is generic and can be used to study other fermionic models of quantum criticality, where there is a strong need to simulate large systems.

Quantum critical phenomena play an important role in condensed matter physics¹. Of particular interest are quantum phase transitions in metals, since they are ubiquitous in strongly correlated materials displaying exotic quantum phenomena, most notably high-temperature superconductivity.

These phase transitions are notoriously difficult to study theoretically². This is due to the presence of an extensive number of gapless fermionic modes on the Fermi surface and the strong coupling between these modes and the transition order parameter. These difficulties render nearly all analytical perturbative approaches

uncontrolled. On the numerical side, these difficulties are manifested in a large amount of entanglement and nearly ubiquitous sign problems, making it hard for controlled numerical techniques such as tensor networks and quantum Monte Carlo (QMC) to make progress.

The most common of such phase transitions in Nature is the onset of antiferromagnetic (AF) spin-density wave (SDW) order in a metal. This transition exists in many material classes of interest, such as electron-doped cuprates³, iron-based materials⁴, and heavy fermion compounds⁵, in which it is believed to generically be a continuous transition. It is often accompanied by a superconducting ‘dome,’

¹Joint Quantum Institute and Department of Physics, University of Maryland, College Park, MD 20742, USA. ²Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA. ³Center for Cosmology and Particle Physics, New York University, New York, NY 10003, USA. ⁴Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. ⁵These authors contributed equally: Michael S. Albergo, Michael Lindsey. ✉ e-mail: plunts@umd.edu

where the maximal T_c occurs near the putative zero-temperature critical point. This makes the (near-critical) SDW fluctuations a strong candidate for the ‘glue’ of Cooper pairs in those materials.

Due to its importance, the theory of this phase transition has received a considerable amount of attention in the last three decades^{6–13}. Very early on it was believed to be well described by Hertz-Millis theory^{14,15}, although it was soon realized that, due to the dimensionality, the arguments of Hertz-Millis theory are invalid^{6–8}. Crucially, the large- N expansion was shown to fail⁹, which left the theory without a controlled approach until the introduction of the fully local ϵ -expansion^{10,11}. Using these controlled results as guidance, in ref. 12 it was shown that there exists a parameter regime where the theory naturally develops a small control parameter, w , which is a ratio of velocities, without the need for any dimensional modification. It is then possible to compute observables perturbatively in w , giving the only fully controlled analytical calculation of the universal low-energy data for the unmodified problem. However, the parameter regime, or ‘basin of attraction,’ of this solution could not be determined from the arguments of ref. 12. It is therefore not clear whether this solution exists for physically relevant parameter values, or only in a minuscule slice of parameter space. One of the central goals of this work is to answer this question.

In the last decade, the SDW transition in metals has also been studied extensively using numerical techniques. The seminal work of ref. 16 introduced a microscopic two-band model of the effective field theory for this transition, which crucially lacks a sign problem. This has led to many studies of this model with Determinantal Quantum Monte Carlo (DQMC)^{17–25}, as well as works that have studied other sign-problem-free models of quantum criticality^{26–32}.

A recent such DQMC work²¹ focused on the critical scaling of the spin susceptibility, in particular seeking a comparison to the predictions of ref. 12 by tuning the UV value of the nesting parameter, ν (explained in detail in Section “Theoretical analysis near perfect nesting”) close to the value at the fixed point of ref. 12, $\nu = 0$. However, at criticality, the spin susceptibility was actually observed to have a (nearly-perfect) Hertz-Millis form, contradicting the theoretical finding^{6–8} that the Hertz-Millis arguments are not valid in two dimensions. Importantly, the maximal system size studied in ref. 21 was $V = L \times L = 14 \times 14$. With such a small system size, it is often difficult to convincingly extract long-wavelength behavior of a critical system. Therefore, even though ref. 21 is the current state-of-the-art for this problem, it is very desirable to revisit the problem at much larger L .

Although DQMC provides a numerically exact and unbiased way to study the properties of these phase transitions, it is severely hindered in its ability to simulate systems with large spatial volume V by its computational scaling of at least $-\beta V^\beta$, where β is the inverse temperature (in the case of small fermion density, this can be reduced by exploiting the low-rank structure of the fermionic determinant^{33,34}).

The true scaling may in fact be worse due to the need to take smaller steps in high dimensions in order to maintain a nonvanishing Metropolis acceptance probability, as well as the presence of ‘critical slowing down,’ discussed further below. Improving the computational scaling with respect to V is of great interest in the study of quantum criticality. Indeed, in order to extract scaling properties near a quantum critical point (QCP), it is crucial to be ‘close enough’ to the thermodynamic limit $V \rightarrow \infty$. This ‘close enough’ is never possible to determine a priori, and, in principle, due to potential semi-stable fixed points, there is never any reason to expect that it has been reached, unless the observed scaling matches a predicted result.

In this paper, we use a different QMC method to study the microscopic model of ref. 16, namely Hybrid Monte Carlo (HMC), which is sometimes referred to as Hamiltonian Monte Carlo. This method is the main numerical tool in the study of lattice quantum chromodynamics (LQCD). Its primary advantage over DQMC is the potential for improved computational scaling with respect to V , which we explain in more detail below. HMC has seen a recent revival in condensed matter physics^{35–37}. Several works have used it to study the half-filled Hubbard model on various lattices (square, honeycomb, hexagonal)^{37–43}, electron-phonon models^{37,44}, as well as extended and long-range Hubbard models of graphene^{45–49}. Nearly all of these results point to extremely favorable scaling with V . However, to the best of our knowledge, HMC has not yet been applied to a model of quantum criticality in the presence of a Fermi surface.

Using our large-scale simulations, we find that the critical theory does in fact strongly deviate from the Hertz-Millis prediction. By tuning the nesting parameter, ν , closer to the fixed-point value of $\nu = 0$, we observe a systematic reduction of the dynamical critical exponent z below the value of $z = 2$ predicted by Hertz-Millis. The prediction of ref. 12 is that $z \rightarrow 1^*$ as $\nu \rightarrow 0$. Additionally, we find that the momentum dependence of the critical spin susceptibility is $O(2)$ -symmetric at intermediate momenta, as predicted by Hertz-Millis, but at lower momenta the symmetry gets reduced to C_4 , as predicted in ref. 12. These two findings provide strong numerical evidence that the critical point theory is governed by the fixed point of ref. 12, even at values of ν that are appreciable. This is summarized in Table 1. Our predictions for the dynamical spin susceptibility can be directly tested with neutron scattering.

In our largest computations we reach system sizes $N_\tau \times L \times L$ given by $N_\tau = 200$, $L = 80$ and $N_\tau = 800$, $L = 20$, where N_τ is the number of slices in the imaginary time direction. Going to such large scales turns out to be necessary for extracting accurate critical scaling behaviors at the critical point.

Now we give an overview of the computational scaling properties of our approach that allow us to achieve these results. Within the HMC algorithm, the number of ‘integration steps’ per effective sample can enjoy scaling as low as $O(d^{1/4})$, where $d = O(\beta V)$ is the dimension of our

Table 1 | Summary of results

	$z(\theta)$	Symmetry	Functional form
This work	$\in (1.665(29), 1.953(35))$ for $\theta \in (0.5^\circ, 8^\circ)$	C_4	$\chi^{-1}(\omega) \sim \omega ^\Delta$ $\chi^{-1}(q_x) \sim q_x ^{z\Delta}$ $\chi^{-1}(\mathbf{q}) \sim q_x ^{z\Delta} + q_y ^{z\Delta}$
			$z\Delta \in (1.466(14), 1.521(15))$ $\Delta \in (0.880(12), 0.779(12))$ for $\theta \in (0.5^\circ, 8^\circ)$
Ref. 12	$z \rightarrow 1^*$ as $\theta \rightarrow 0$	C_4	$\chi^{-1}(\omega) \sim \omega ^\Delta$ $\chi^{-1}(q_x) \sim q_x ^{z\Delta}$ $\chi^{-1}(\mathbf{q}) \sim$ not determined
			$z\Delta, \Delta \rightarrow 1^*$ as $\theta \rightarrow 0$
Hertz-Millis	$z = 2 \forall \theta$	$O(2)$	$\chi^{-1}(\omega) \sim \omega $ $\chi^{-1}(q_x) \sim q_x^2$ $\chi^{-1}(\mathbf{q}) \sim q_x^2 + q_y^2$

The main results of this paper, which pertain to the critical spin susceptibility, $\chi(\omega, \mathbf{q})$. We list its dynamical critical exponent z , spatial symmetry, and functional form, as compared to that predicted in ref. 12 and by Hertz-Millis theory, all for the varying nesting angle θ we study (here we use θ instead of the nesting parameter $\nu = \tan(\theta)$ that we use in the rest of the paper). The forms quoted for ref. 12 are neglecting the logarithmic flow of ν .

discretized field to be sampled⁵⁰. This scaling may be worse in the presence of ‘critical slowing down’ for critical models, yielding a number of integration steps per effective sample of potentially $O(\beta^{1/4+z_1} V^{1/4+z_2})$, where $z_1, z_2 > 0$. Even in the presence of criticality, z_1 and z_2 may be reduced via the choice of the ‘metric’ within HMC, as discussed below. For the critical model, our algorithm in fact achieves $z_2 \approx 0$ and $z_1 \lesssim 0.5$. Each of the integration steps within HMC requires the solution of linear system of size $O(\beta V)$, which constitutes the bottleneck for the algorithm. Our approach for solving the linear system in fact scales linearly with respect to V for fixed β and permits fast GPU implementation. In conjunction with the scaling $z_2 \approx 0$ observed above, this performance yields overall wall clock scaling with exponent approximately 5/4 with respect to V . See Section “Numerical Performance” for further details.

In our implementation we in fact develop several augmentations of the basic HMC algorithm. First, we introduce an auto-tuning procedure, which tunes our hyperparameters in an initial warmup phase. This procedure is common practice in statistics and industry applications of HMC^{51,52} but to the best of our knowledge has not yet been fully applied in condensed matter physics. Moreover, relative to such works, our translation-invariant physical setting allows us to tune the aforementioned HMC metric with operations that scale linearly in βV , up to log factors.

In the Results Section, we introduce the low-energy effective theory of the SDW transition, both in the continuum and on the lattice, we review the theoretical results near perfect nesting of ref. 12 and derive its predictions for the observables that are computed numerically in this work, and we show our main results. In the Methods Section we give a brief review of HMC, followed by a detailed presentation of our implementation and its numerical performance.

Results

Continuum effective action

Our starting point is the theory that describes the low-energy degrees of freedom near a metallic AF SDW quantum critical point in two spatial dimensions. These are the order parameter for the transition, which is the collective spin excitation, and the electrons near points on the Fermi surface called ‘hotspots’ that are connected by the AF ordering wavevector. For concreteness, we consider a single band with C_4 symmetry and an ordering wavevector equal to (π, π) , as shown in Fig. 1. Generically for such a Fermi surface there are four pairs of coupled hotspots. The Euclidean-time action for this low-energy theory is then given by

$$\begin{aligned}
 S = & \sum_{n=1}^4 \sum_{m=\pm} \sum_{\sigma=\uparrow,\downarrow} \sum_{j=1}^{N_f} T \sum_{\omega_k} \int \frac{d\mathbf{k}}{(2\pi)^2} \\
 & \psi_{n,\sigma j}^{(m)*}(\mathbf{k}) [i\omega_k + e_n^m(\mathbf{k}; \nu)] \psi_{n,\sigma j}^{(m)}(\mathbf{k}) \\
 & + T \sum_{\omega_q} \int \frac{d\mathbf{q}}{(2\pi)^2} \frac{1}{2} \left[\frac{\omega_q^2}{c^2} + (q_x^2 + q_y^2) + r \right] \boldsymbol{\phi}(\mathbf{q}) \cdot \boldsymbol{\phi}(-\mathbf{q}) \\
 & + \frac{g}{\sqrt{N_f}} \sum_{n=1}^4 \sum_{\sigma,\sigma'=\uparrow,\downarrow} \sum_{j=1}^{N_f} \int d\tau \sum_{\mathbf{r}} \\
 & \left[\boldsymbol{\phi} \cdot \psi_{n,\sigma j}^{(+)*} \boldsymbol{\tau}_{\sigma,\sigma'} \psi_{n,\sigma' j}^{(-)} + \text{h.c.} \right] + \frac{u}{4} \int d\tau \sum_{\mathbf{r}} (\boldsymbol{\phi} \cdot \boldsymbol{\phi})^2.
 \end{aligned} \tag{1}$$

In Eq. (1), $k = (\omega_k, \mathbf{k})$ consists of the fermionic Matsubara frequency and the two-dimensional momentum $\mathbf{k} = (k_x, k_y)$. T is the temperature. The $\psi_{n,\sigma}^{(m)}$ correspond to electrons at the hot spots labeled by $n \in \{1, 2, 3, 4\}$ and $m \in \{+, -\}$ (cf. Fig. 1), and spin $\sigma \in \{\uparrow, \downarrow\}$. The axes are chosen as follows: \hat{k}_x is in the direction from hot spot (2, +) to (2, -) and \hat{k}_y is in the direction from hot spot (1, -) to (1, +). Given this choice of axes, the ordering wave vector connecting the paired hot spots is $\mathbf{Q}_{\text{AF}} = (\pm\sqrt{2}\pi\hat{k}_x, \pm\sqrt{2}\pi\hat{k}_y)$ up to the reciprocal lattice vectors

$\sqrt{2}\pi(\hat{k}_x \pm \hat{k}_y)$. See Fig. 1 for details. The linearized electron dispersions are given by $e_1^\pm(\mathbf{k}; \nu) = -e_3^\pm(\mathbf{k}; \nu) = \nu k_x \pm k_y$, $e_2^\pm(\mathbf{k}; \nu) = -e_4^\pm(\mathbf{k}; \nu) = \mp k_x + \nu k_y$, where the momentum \mathbf{k} is measured relative to each hot spot. For any $\nu \neq 0$, the curvature of the Fermi surface ($O(k^2)$ terms) can be ignored, since the linearized dispersions at the coupled hotspots are not parallel to each other, and therefore the problem is still fully two-dimensional. The component of the Fermi velocity along \mathbf{Q}_{AF} has been set to one by rescaling \mathbf{k} . ν is the component of Fermi velocity that is perpendicular to \mathbf{Q}_{AF} . It controls the degree of nesting between coupled hot spots and can be written as $\nu = \tan \theta$, where θ is the nesting angle (c.f. Fig. 1). $\boldsymbol{\phi}(\mathbf{q})$ is the three-component boson field that describes the AFM collective mode in the fundamental representation of $O(3)$, with frequency ω_q and momentum \mathbf{q} . Note that while the collective spin is centered at \mathbf{Q}_{AF} , $\boldsymbol{\phi}$ is centered at zero, since \mathbf{Q}_{AF} is already incorporated into the hotspot label. c is the boson velocity, and u is its quartic interaction. r is the squared mass of the boson, as is tuned drive the boson to criticality. g is the Yukawa coupling between the boson and the electrons, which scatters the electrons between hot spot pairs via the spin-spin interaction. These couplings represent all the relevant and marginal terms obeying the symmetries of the problem. $\boldsymbol{\tau}$ consists of the three generators of the $SU(2)$ group. We have also generalized the action from one to N_f fermion flavors. This generalization was used in the large N_f expansion in previous renormalization group studies^{7,9}, and we use this general form in subsequent sections.

Sign-problem-free UV completion

The action in Eq. (1) captures the universal low-energy properties of the system associated with the divergent correlation length near the phase transition. Since the electrons far away from the hotspots do not enter into the theory, the precise shape of the Fermi surface beyond the momentum cutoffs does not play a role in the critical phenomena of this theory (the complete validity of this has recently been called into question in Ref. 53). Therefore, we can change the band structure while ensuring that the action of Eq. (1) is not modified. One such UV completion was given in ref. 16 and has the real-space action of

$$\begin{aligned}
 S = & \int d\tau \sum_{\sigma=\uparrow,\downarrow} \sum_{\alpha=x,y} \sum_{j=1}^{N_f} \sum_{\mathbf{r},\mathbf{r}'} \\
 & \psi_{\alpha,\sigma j,\mathbf{r}}^* [(\partial_\tau - \mu)\delta_{\mathbf{r},\mathbf{r}'} - t_{\alpha,\mathbf{r},\mathbf{r}'}] \psi_{\alpha,\sigma j,\mathbf{r}'} \\
 & \int d\tau \sum_{\mathbf{r}} \left[\frac{1}{c^2} (\partial_\tau \boldsymbol{\phi}_{\mathbf{r}})^2 + \frac{1}{2} (\nabla \boldsymbol{\phi}_{\mathbf{r}})^2 + \frac{r}{2} (\boldsymbol{\phi}_{\mathbf{r}})^2 + \frac{u}{4} (\boldsymbol{\phi}_{\mathbf{r}})^4 \right] \\
 & + \frac{g}{\sqrt{N_f}} \sum_{\sigma,\sigma'=\uparrow,\downarrow} \sum_{j=1}^{N_f} \int d\tau \sum_{\mathbf{r}} e^{i\mathbf{Q}_{\text{AF}} \cdot \mathbf{r}} \\
 & \boldsymbol{\phi}_{\mathbf{r}} \cdot [\psi_{x,\sigma,\mathbf{r}}^* \boldsymbol{\tau}_{\sigma,\sigma'} \psi_{y,\sigma',\mathbf{r}} + \text{h.c.}].
 \end{aligned} \tag{2}$$

Here, the number of bands has been doubled ($\alpha = x, y$ represents a band index). The boson now scatters electrons between the two bands. The band-dependent hopping amplitude $t_{\alpha,\mathbf{r},\mathbf{r}'}$ still respects the C_4 symmetry of the lattice, provided that the bands are also interchanged. The location of the hotspots and the dispersion linearized about them is unchanged from the one-band model. Note that the axes of this model are rotated 45° relative to the model of Eq. (1). The advantage of this model is that it now contains an inter-band anti-unitary symmetry that guarantees the positivity of the fermionic determinant¹⁶ and makes it amenable to sign-problem-free Monte Carlo simulations^{17–21,24,25}.

We study the lattice model defined in Eq. (2) using HMC. Our method is described in detail in Section “Methods”. As explained there, in order for our algorithm to function, we must work with an even N_f , and we choose $N_f = 2$. Unless stated otherwise, we fix the parameter values to be $u = 0, g = 0.7\sqrt{2} \approx 1, c = 3$. We study five different Fermi




Fig. 1 | The Fermi surface. **a** The first Brillouin zone of the model in Eq. (1). Note the axes rotated by 45° . The shaded region corresponds to the occupied states. The SDW ordering wavevector \mathbf{Q}_{AF} is denoted by red arrows. The hot spots are the red points connected by \mathbf{Q}_{AF} . At each hot spot, the linearized Fermi surface is shown with a blue line. **b** A close-up of linearized dispersions of the paired hot spots ($1, \pm$). One has been shifted by \mathbf{Q}_{AF} , to show their crossing. The angle between them θ is called the ‘nesting angle’, and its tangent $\nu = \tan(\theta)$ is called the ‘nesting parameter’.

surfaces, with the following values of the nesting parameter: $\nu_1 \approx 0.149$, $\nu_2 \approx 0.072$, $\nu_3 \approx 0.036$, $\nu_4 \approx 0.018$, $\nu_5 \approx 0.0092$, corresponding to nesting angles of $\theta_1 \approx 8.5^\circ$, $\theta_2 \approx 4.13^\circ$, $\theta_3 \approx 2.05^\circ$, $\theta_4 \approx 1.03^\circ$, $\theta_5 \approx 0.53^\circ$, respectively. We illustrate the Fermi surface corresponding to $\nu = \nu_1$ and $\nu = \nu_5$ in Fig. 2. The non-zero hopping amplitudes and chemical potentials corresponding to each are $t_{h,x} = t_{v,y} = 1$ and

$$\tilde{t}_{v,x}^{(1)} = -\tilde{t}_{h,y}^{(1)} = 0.45, \quad \mu_x^{(1)} = -\mu_y^{(1)} = -0.47 \quad (3)$$

$$\tilde{t}_{v,x}^{(2)} = -\tilde{t}_{h,y}^{(2)} = 0.48, \quad \mu_x^{(2)} = -\mu_y^{(2)} = -0.46 \quad (4)$$

$$\tilde{t}_{v,x}^{(3)} = -\tilde{t}_{h,y}^{(3)} = 0.498, \quad \mu_x^{(3)} = -\mu_y^{(3)} = -0.44 \quad (5)$$

$$\tilde{t}_{v,x}^{(4)} = -\tilde{t}_{h,y}^{(4)} = 0.505, \quad \mu_x^{(4)} = -\mu_y^{(4)} = -0.44 \quad (6)$$

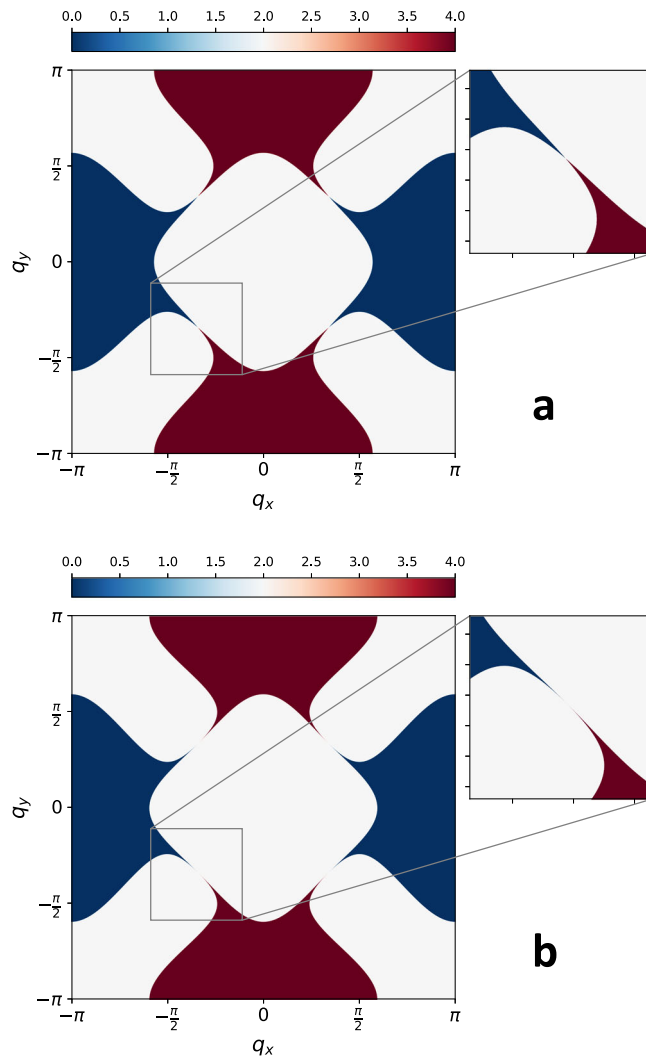


Fig. 2 | Occupation number. The total occupation number $n_q = \sum_{\alpha,\sigma} \langle \psi_{\alpha,\sigma}^\dagger \psi_{\alpha,\sigma} \rangle$ of one flavor of fermions for the free theory, summed over spin and band degrees of freedom. Here, we have incorporated the phase shift $e^{i\mathbf{Q}_{AF} \cdot \mathbf{r}}$ into the $\alpha = y$ band. The nesting values shown are (a) $\nu = \nu_1$ and (b) $\nu = \nu_5$, which are the two extremes we study. Note that the Fermi surface at the hot-spots is actually connected, but due to the very small nesting parameter, that is hard to show graphically.

$$\tilde{t}_{v,x}^{(5)} = -\tilde{t}_{h,y}^{(5)} = 0.5085, \quad \mu_x^{(5)} = -\mu_y^{(5)} = -0.44. \quad (7)$$

Here, $t_{h,\alpha}$, $t_{v,\alpha}$ denote the nearest-neighbor hopping amplitudes in the x and y directions, respectively, which are the same for all nesting parameter values; $\tilde{t}_{h,\alpha}^{(i)}$, $\tilde{t}_{v,\alpha}^{(i)}$ denote the next-nearest-neighbor hopping amplitudes in the x and y directions, respectively, for nesting parameter value i ; $\mu_\alpha^{(i)}$ are the chemical potentials for the nesting parameter value i .

The reason for choosing these specific parameter values is to make contact with ref. 21, where the authors in turn tried to make contact with ref. 12. Of course, it is important to scan the values of u, g, c to check the stability of our results and look for new behavior. However, such a detailed study is beyond the scope of the present work, and is left for future studies.

Theoretical analysis near perfect nesting

Although in general the theory in Eq. (1) cannot be understood analytically using a controlled approach, there exists a parameter regime where a controlled solution in the IR can be obtained¹². This parameter




Fig. 3 | Renormalization group flow. RG flow of the theory in Eq. (1), computed using the epsilon expansion, reproduced from ref. 11. The axes are $w, \bar{x} = x/10, \bar{\lambda} = 10\lambda$. Initially, there is a fast flow (blue lines) to a fixed one-dimensional manifold (dashed orange line). The flow along this manifold is slow, and towards $w = 0$. The flow of κ is excluded, since its flow towards zero is the fastest.

regime is primarily characterized by a small nesting parameter $\nu \ll 1$ and an effective coupling of intermediate strength, leading to strong correlations. We start by reviewing this IR fixed point of refs. 11,12,54.

At criticality, the UV theory is best described in terms of the ratios $\lambda \equiv \frac{g^2 c^2}{\nu}, x \equiv \frac{g^2}{c}, \kappa \equiv u c^2$, and $w = \frac{\nu}{c}$. (We note that compared to refs. 11,12,54, the action of Eq. (1) has $\phi \rightarrow \phi/c$). The renormalization group flow of the UV theory to the IR fixed point is two-fold and is shown in Fig. 3, which is reproduced from ref. 11. Initially, there is a fast (algebraically in the running energy scale μ) flow of λ and x to w -dependent $\mathcal{O}(1)$ values and κ to zero. Once the first step of the flow is complete, the couplings will keep flowing along the one-dimensional manifold defined by the fixed point values λ, x, κ . This one-dimensional manifold defines the only remaining free coupling of the theory. We note that since $\lambda \sim \mathcal{O}(1)$ is the effective coupling of the UV theory, this manifold represents the theory in the strongly correlated regime. Crucially, due to all the coupling inter-dependencies, along the manifold w becomes a function of ν only. We can choose to parametrize the manifold by $w(\nu)$. This choice is useful, since it turns out that theory on the manifold has a perturbation theory that is organized in powers of $w(\nu)$ (along with powers of $\log(w)$). However, since $w(\nu)$ is only a function of ν , such that $w(\nu) \rightarrow 0$ with $\nu \rightarrow 0$, perturbation theory can also be done in ν itself.

Finally, once the theory sits on the one-dimensional manifold, the remaining flow is towards $w(\nu) \rightarrow 0$. However, unlike the first part of the flow, the flow of $w(\nu)$ in this second part happens at a rate that is sub-logarithmic in μ . Therefore, for all practical purposes we can take $w(\nu)$ (and therefore ν) to be fixed (scale independent) on this manifold. An important point is that due to the initial fast flow, the value of ν on the manifold is different than the bare (UV) value ν_B . However, given that the bare parameters of the theory are tuned to minimize the length of this initial flow, we assume that $\nu \approx \nu_B$ in the analysis of Section “Results”.

If the value of $w(\nu)$ is small enough, the theory can be studied using perturbation theory in $w(\nu)$ to a finite order. Focusing on the critical spin susceptibility $\chi(\omega, \mathbf{q})$, its leading order in $w(\nu)$ behavior is given by (here we use the axes of Eq. (1))

$$\chi^{-1}(\omega, \mathbf{q}) = |\omega| + c(w(\nu)) (|q_x| + |q_y|). \tag{8}$$

Here, $c(w(\nu))$ is a new emergent boson velocity, not related to the bare value c , and dependent on the only coupling in the theory. It has its own expansion in powers of $w(\nu)$ ¹². This form breaks the O(2) spatial

symmetry (a $|\mathbf{q}|$ dependence) of the spin susceptibility down to the C_4 symmetry of the Fermi surface. This is a key prediction of the small ν fixed point theory¹². Equation (8) can be thought of as the tree-level susceptibility at the fixed point, and perturbations to it are computed in powers of $w(\nu)$. The deviations of the scaling are encoded in the dynamical critical exponent z and the anomalous boson dimension η_ϕ , which were computed to leading order in ref. 12,

$$z = 1 + \frac{3}{4\pi N_f} w(\nu) \tag{9}$$

$$\eta_\phi = \frac{1}{2\pi N_f} w(\nu) \log\left(\frac{1}{w(\nu)}\right). \tag{10}$$

Reference¹² also computed the dependence of $w(\nu)$ on ν to lowest order in ν ,

$$w = 4\sqrt{N_f} \sqrt{\frac{\nu}{\log(1/\nu)}}. \tag{11}$$

Together these relations give the leading contributions to z, η_ϕ for small ν .

In order to make a connection between the leading-order perturbation theory in $w(\nu)$ results of ref. 12 and the theory of Eq. (2) at finite ν , we need to use the non-perturbative renormalization group (RG) equation. This equation for $\chi(\omega, \mathbf{q})$ at criticality ($r = r_c$) is given by

$$\left[\Delta(\eta_\phi, z) + \omega \frac{\partial}{\partial \omega} + \frac{1}{z} \mathbf{q} \cdot \frac{\partial}{\partial \mathbf{q}} + T \frac{\partial}{\partial T} + \frac{1}{z} \left(\frac{\partial}{\partial L} \right) \right] \chi(\omega, \mathbf{q}; T, L) = 0, \tag{12}$$

where $\Delta(\eta_\phi, z) = \frac{1-2\eta_\phi-(z-1)}{z}$. Here, we have used the fact that $\nu \approx \nu_B$ is approximated as being constant as a function of scale (otherwise its beta function would enter in Eq. (12)). Since at criticality the correlation length is absent from Eq. (12), we have included T, L^{-1} as relevant energy scales. We note that for a scaling theory different from that of ref. 12, the RG equation would be the same, with the only modification being the form of Δ . The only requirement is that the flow of the couplings can be ignored, which is a good approximation when the couplings flow at most logarithmically in the running scale.

We study systems of fixed size and temperature. We focus on two limits of Eq. (12). In the first limit, $\mathbf{q} = 0$, the solution is given by

$$\chi(\omega) = |\omega|^{-\Delta(\eta_\phi, z)} f(\omega/T, \omega L^z), \tag{13}$$

where f is some non-universal crossover function. Here, T and $1/L$ act as IR cutoffs to the critical scaling. In other words, there is some effective IR frequency cutoff given roughly by $\omega_{\text{IR}} \sim \max(\frac{2\pi}{\beta}, (\frac{2\pi}{T})^z)$, and the function $f(\omega/T, \omega L^z)$ is a finite constant for $\omega \gg \omega_{\text{IR}}$. On the other hand, the scaling equation itself only holds below some non-universal UV cutoff Λ_ω . Therefore, the pure algebraic scaling of $\chi(\omega, \mathbf{q} = 0) \sim |\omega|^{-\Delta(\eta_\phi, z)}$ can be observed in an intermediate frequency window, $\omega_{\text{IR}} \ll \omega \ll \Lambda_\omega$.

Now we turn to the case of $\omega = 0$. For simplicity, we first also set $q_y = 0$. Analogously to $\chi(\omega)$, $\chi(q_x)$ is given by

$$\chi(q_x) = |q_x|^{-z\Delta(\eta_\phi, z)} g(q_x/T^{1/z}, q_x L). \tag{14}$$

Here, g is another non-universal crossover function. The critical scaling $\chi(q_x) \sim |q_x|^{-z\Delta(\eta_\phi, z)}$ holds in the intermediate region of $q_{\text{IR}} \ll q_x \ll \Lambda_q$, where $q_{\text{IR}} = \max((\frac{2\pi}{\beta})^{1/z}, \frac{2\pi}{L})$ and Λ_q are IR and UV momentum cutoffs, respectively. Once the intermediate-scale algebraic scaling of both

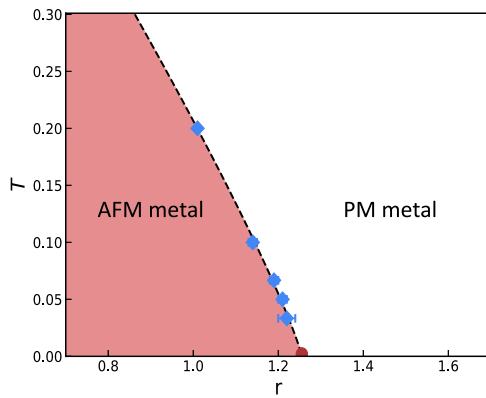


Fig. 4 | Phase diagram in T and r . The red (white) region is a metal with AFM (no magnetic) order. The blue points denote the intersection of curves of the Binder cumulant for different system sizes at fixed temperatures. The error bars denote the finite resolution in r that we have in those intersection points. The phase diagram for all nesting parameters studied is the same up to our resolution. Superconductivity is absent for all the parameters we study.

$\chi(\omega)$ and $\chi(q_x)$ are known, z is computed directly from their ratio: $z\Delta(\eta_\phi, z)/\Delta(\eta_\phi, z)$.

When we turn on q_y , the RG equation does not say anything about the functional form of the dependence of $\chi(\mathbf{q})$ on q_x and q_y , which must be deduced by other means. All that is known from the tree-level form of Eq. (8) is that the spatial dependence must be C_4 -symmetric. Jumping ahead to Section “Results”, the Monte Carlo data suggests the form

$$\chi^{-1}(\mathbf{q}) = (|q_x|^{z\Delta(\eta_\phi, z)} + |q_y|^{z\Delta(\eta_\phi, z)}) g((|q_x| + |q_y|)/T^{1/z}, (|q_x| + |q_y|)L), \quad (15)$$

which we take as our conjecture.

As an aside, we note that this way of extracting critical exponents from a system at fixed L, β is not the usual one of finite size scaling. Although the latter is more systematic, it would require the simultaneous scaling of L and $\beta \propto L^z$. Since we do not know what z is a priori, this requires an additional scan over its potential values. Also, the z value we find later on is always $z > 1.6$, which implies the need for a large β in the finite size scaling. Both of these are very computationally expensive, and we therefore opt for the present method.

Numerical results

The phase diagram of the theory in Eq. (2) as a function of r and T is shown in Fig. 4.

At large values of r the system is a paramagnetic (PM) metal. As a function of decreasing r , for any finite T (below some large temperature) there is a sharp crossover to a SDW metal. This crossover will not become a true second-order phase transition in the thermodynamic limit, since the dimensionality and symmetry of the order parameter prevent the system from ordering due to the Mermin-Wagner theorem. However, if we infinitesimally couple stacks of this two-dimensional system (as is the situation in many relevant experiments), order will be stabilized and the transition will become a true one. We can identify the sharp location of this potential second-order transition, $r_c(T)$, by studying finite-size scaling diagnostics, such as the Binder cumulant (c.f. Supplementary Note 1 for details). The line of these transition points terminates at $T=0$ and $r=r_c$, which is the quantum critical point. This is the familiar picture of quantum criticality.

For all nesting parameters we see no superconductivity down to the lowest measured temperatures. This is consistent with the arguments outlined in ref. 21 that the superconducting transition

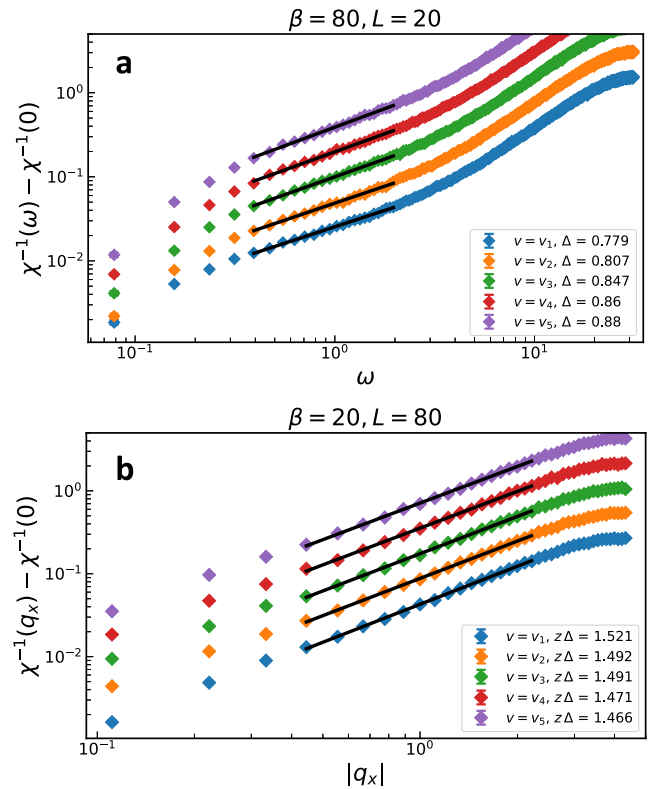


Fig. 5 | Scaling of critical spin susceptibility. The (a) dynamic and (b) static spin susceptibility, shown on log-log plots for all nesting parameters v_i studied. The error bars are given by the one sigma statistical uncertainties from the stochastic calculation. The different curves are shifted relative to each other for visual clarity. We show only the largest (a) β and (b) L values that we simulated in order to extract the intermediate ω and q_x regimes of power-law scaling. For each curve, we choose the regions that look the most straight by eye (if there are two such regions, as in $\chi^{-1}(\omega)$, we choose the one with smaller ω or q_x). Choosing slightly different boundaries leads to slightly different exponents, but the trend with v is always the same.

temperature T_c of the spin-fermion model is suppressed with decreasing v as $T_c \sim g^2 \sin(\arctan(v))$.

To study the critical scaling we tune $r=r_c$ and vary $\beta \equiv 1/T$. Following Section “Theoretical analysis near perfect nesting”, we look at the dependence of the spin susceptibility on frequency and momentum, $\chi(\omega)$, $\chi(q_x)$, $\chi(\mathbf{q})$. Here, we use the axes of Eq. (1), in order to compare to Section “Theoretical analysis near perfect nesting”. From the scaling of $\chi(\omega)$ and $\chi(q_x)$ in the intermediate scaling regions we obtain the exponents $\Delta(\eta_\phi, z)$ and $z\Delta(\eta_\phi, z)$, as illustrated in Fig. 5. From $\Delta(\eta_\phi, z)$ and $z\Delta(\eta_\phi, z)$ we determine z and η_ϕ , which are plotted in Fig. 6 as functions of v . An important point is that the intermediate regime of power-law scaling in Fig. 5 is chosen by eye, and different choices yield slightly different forms of $z(v)$ and $\eta_\phi(v)$. The specific choice we use in Fig. 5 is described in Supplementary Note 1, along with how we compute the corresponding error bars. However, crucially, for all (reasonable) choices of the scaling region boundaries, $z(v)$ monotonically decreases with decreasing v , starting from a value slightly less than $z=2$ (which is the Hertz-Millis prediction) for v_1 . Similarly, $\eta_\phi(v)$ monotonically increases from $\eta_\phi \approx -0.75$ with decreasing v .

In addition to the dependence of $z(v)$ and $\eta_\phi(v)$ on v , we determine the spatial symmetry of $\chi(\mathbf{q})$ and its functional form. In Fig. 7 we show density plots of $\chi^{-1}(\mathbf{q})$ for $v=v_5$. We plot contours at various values. We can see that at large momenta, the symmetry of the contours is that of the lattice, C_4 . At intermediate momenta, the contours are circles, indicating an O_2 symmetry, indicating that those momenta are small enough that $\cos(k_x) + \cos(k_y) \approx 2 - k^2/2$ is a good approximation. However, at smaller momenta the contours transform again to a

C_4 -symmetric form, with the maxima/minima now rotated by 45° . In Supplementary Note 1, we show more density plots at smaller L , in order to illustrate that this effect gets more pronounced as L increases and is therefore a true long-wavelength effect present in the thermodynamic limit.

It is not possible to provably determine the exact functional form, as there are many forms that obey the C_4 symmetry. However, $|q_x|^{z\Delta} + |q_y|^{z\Delta}$ seems to fit very well. We illustrate this in Fig. 8, along with a comparison to $|q|^{z\Delta}$. We therefore conjecture that $\chi^{-1}(\mathbf{q}) \sim |q_x|^{z\Delta} + |q_y|^{z\Delta}$ is the correct long-wavelength form.

These two features of the boson susceptibility—i.e., the monotonic forms of $z(\nu)$, $\eta_\phi(\nu)$ and the symmetry reduction of $O(2) \rightarrow C_4$ in the long-wavelength limit—provide strong numerical evidence in favor of the critical scaling at the SDW QCP being governed by the theory of ref. 12 for all the nesting values we study.

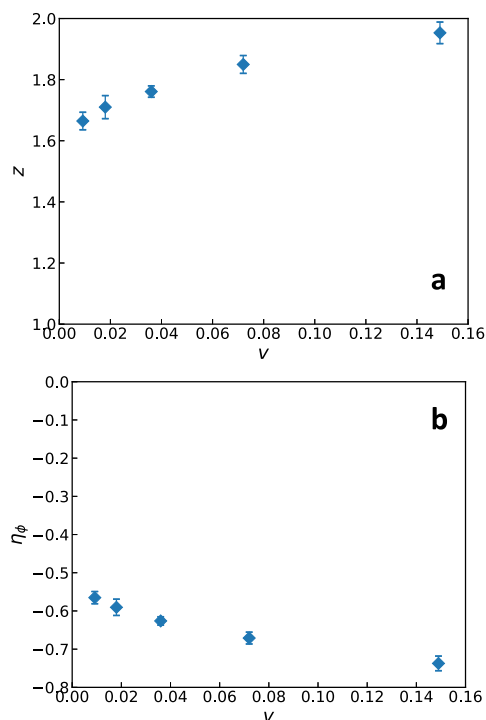


Fig. 6 | Critical exponents. The (a) dynamical critical exponent z and (b) the anomalous boson dimension η_ϕ as a function of ν . The error bars are computed using a bootstrap method on the fitted data, allowing for duplicates (c.f. Supplementary Note 1).




Fig. 7 | Density plots of the static critical spin susceptibility. All are for $\nu = \nu_s$, $\beta = 20$ and $L = 80$. We are interested in showing the equal-density contours (we do not show a color bar as it is irrelevant for this purpose). The three plots are different levels of zooming in to $\mathbf{q} = 0$. In **a** we show the entire Brillouin zone, along with three dashed-line equal-density contours (passed through a Gaussian filter for smoothness), at large, intermediate, and small momenta. We can see that the largest contour has a C_4 -symmetric form, due to the lattice. In **b** we zoom in to focus on the

inner two contours, which we overlay with two red circles at radii 0.63, 1.33, in order to illustrate their symmetry. The larger circle overlays with the contour extremely well, indicating an $O(2)$ symmetry at those momenta. The smaller circle overlays with its contour poorly, since that contour is much more square-like, indicating that at smaller momenta the symmetry is again C_4 . To show this more convincingly, in **c** we zoom in further to enlarge this smallest contour.

Despite the similarities, there are two parts of the data that are seemingly in contrast to the predictions of ref. 12. The first is the fact that η_ϕ is negative, whereas Eq. (10) predicts that as $\eta_\phi \rightarrow 0^+$ as $\nu \rightarrow 0$. However, Eq. (10) is only the leading order correction, and it particular, in obtaining it, terms of $\mathcal{O}(w(\nu))$ were ignored in ref. 12. Therefore, it is perfectly possible that the actual leading-order coefficient behaves as $\eta_\phi \sim w(\nu)(\log(1/w(\nu)) - a)$, where $a > 0$ is a constant, which would lead to a sign-change in the leading-order behavior when $1 \gg w(\nu) > e^{-a}$. On the other hand, the leading term of z from Eq. (9) does not suffer from a similar problem, and its sign can be trusted at small $w(\nu)$.

The second part of the data that is seemingly in contrast to the predictions of ref. 12 is that $z(\nu)$ and $\eta_\phi(\nu)$ do not convincingly exhibit the limiting behaviors $z \rightarrow 1$ and $\eta_\phi \rightarrow 0$ as $\nu \rightarrow 0$. We conjecture the following explanation. As noted in Section “Theoretical analysis near perfect nesting”, we assume that the renormalized value of ν is close to the bare one. However, as we decrease ν we are not tuning g , so the bare value of $\lambda = g^2 c^2 / \nu$ is increasing, not staying fixed. This increases the RG-time needed to flow to the one-dimensional manifold discussed in Section “Theoretical analysis near perfect nesting”. The large bare value of the effective coupling λ might be renormalizing ν to larger values, before the one-dimensional manifold it reached. Remediating this requires decreasing g with decreasing ν such that λ remains fixed. This task is beyond the scope of this paper, and we leave it for future work.

Finally, we show the fermion occupation function at criticality in Fig. 9 (as well as Supplementary Note 1), where we can see the renormalization of the hot-spots by the Yukawa interaction as well as the gap appearing at the hot-spots in the ordered phase. Unfortunately, within our spatial resolution it is not possible to measure the renormalized nesting parameter ν , and therefore confirm the conjecture from the above paragraph.

Discussion

In this work, we study the critical theory of the $O(3)$ spin-fermion model as a model of the antiferromagnetic transition in two-dimensional metals. We use a novel HMC method that we extensively develop. Below we discuss some consequences of our work and outlook.

Implications of theoretical results

Our study of the critical spin susceptibility $\chi(\omega, \mathbf{q})$ provides one of the few controlled numerical studies of a model where there is an observed violation of Hertz-Millis scaling. The critical exponents and C_4 -symmetric form of $\chi(\omega, \mathbf{q})$ can be directly tested experimentally with neutron scattering in materials where the bare hot-spot nesting value ν_B is small.




Fig. 8 | Scaling form of static critical spin susceptibility. $\chi^{-1}(q) - \chi^{-1}(0)$ plotted against (a) $|q_x|^{2\Delta} + |q_y|^{2\Delta}$ and (b) $|q|^{2\Delta}$. Both plots are for $\beta = 20$ and $L = 80$. The error bars are again the one sigma uncertainties. The curves for different ν_i are shifted to separate them for visual clarity. The initial region, up to $q^{2\Delta} - 1.7$ is fit with both

forms. The C_4 is much better in this region; see the solid lines in a and b, and compare the coefficients of determination R^2 from a with R_1^2 from b. However, over larger momenta (up to $q^{2\Delta} - 3.6$) the $O(2)$ fit it much better; see the dashed line as well as R_2^2 in b.

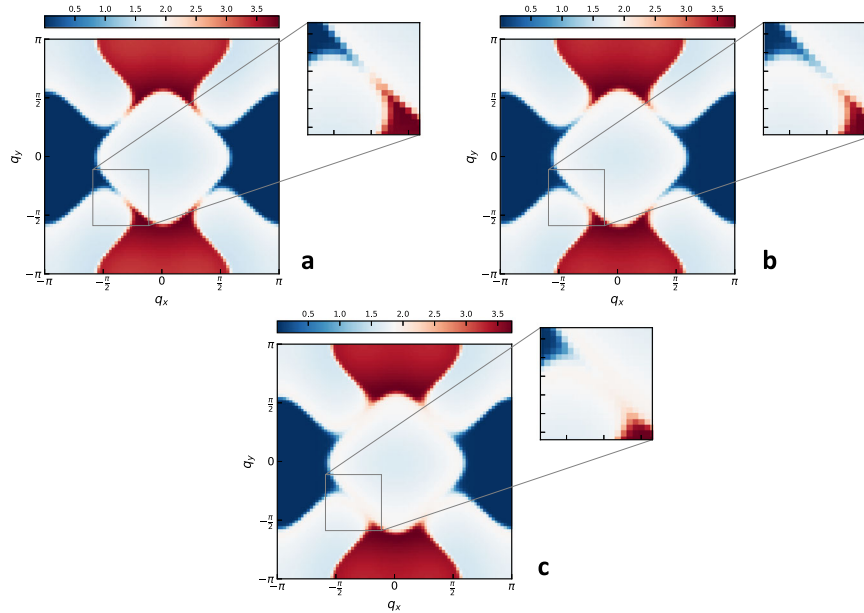


Fig. 9 | Measured occupation number. The total occupation number $n_{\mathbf{q}} = \sum_{\alpha, \sigma} \langle \psi_{\alpha, \sigma, \mathbf{q}}^\dagger \psi_{\alpha, \sigma, \mathbf{q}} \rangle$ for the (a) free theory, and the interacting theory (b) at criticality ($r = r_c \approx 1.26$) and (c) slightly in the ordered phase ($r = 1.15$). This explicitly

shows the renormalization of the Fermi surface near the hot spots due to fluctuations of the order parameter. All plots are for $\nu = \nu_1$ and $\beta = 20, L = 80$.

Additionally, we provide strong evidence that the physics of the quantum critical metal is governed by the fixed point of ref. 12, even for appreciable nesting parameter values ν . If one believes this evidence, the implication is that the theory of ref. 12, valid at small ν , augmented with numerical calculations at larger ν , can be used to generate a fully controlled quantitative solution to the problem. For example, non-equilibrium properties such as conductivity, which are generally very difficult to resolve with imaginary time methods, but are the most common observable used to experimentally identify the quantum critical fan, can be computed using small ν perturbation theory about the fixed point of ref. 12 and supplemented with non-perturbative numerical calculation of the critical exponents. If valid, this approach would provide an unprecedented advance in the understanding of experimentally observed strange metals proximate to itinerant antiferromagnetism, characterized by, e.g., linear-in-temperature DC resistivity and present in many materials displaying unconventional superconductivity.

Future directions with HMC

As noted in the main text, Fig. 6 does not convincingly show that $z \rightarrow 1$ as $\nu \rightarrow 0$, and the most likely reason for this is the increasing value of $\lambda = g^2 c^2 / \nu$ as we decrease ν . It is therefore desirable to revisit this problem, while keeping λ fixed by tuning g . We intend to address this in a future work.

In this work, we focus on only two critical exponents of the spin-fermion model, z and η_{ϕ} , for the most straightforward comparison to the nearly-nested fixed point theory of ref. 12. There are several other

pieces of critical data that remain to be examined for a more comprehensive comparison. The two most important ones are the critical exponent ν that relates the correlation length to the deviation from the critical point, and the scaling of the fermion Green’s function. Measuring the critical scaling of thermodynamic quantities like specific heat is also an important task. One can also tune g to a large enough value to see the onset of superconductivity in order to track the dependence of T_c on ν and compare to the prediction of $T_c(\nu)$ made in ref. 12. Finally, here we only study the parameter values $u = 0, c = 3$, and five small values of ν . An obvious question is how the conclusions of this work change with varying parameters u, c , as well as larger values of ν .

The HMC algorithm we use in this work is quite general and can be applied to many other theories of interest. The main restrictions are that the configuration (boson) field be continuous and that the number of fermions can be doubled without qualitatively changing the physics of the model (this is not necessary if the fermion matrix is purely real). Some immediate candidates are other sign-problem-free ‘designer’ models of various flavors of quantum criticality, such as easy-plane XY^7 and Z_2^{31} antiferromagnetic transitions in metals, the Ising nematic transition in metals²⁶, fermions coupled to emergent gauge fields^{55–61}, Kondo lattice physics^{62–64}, disorder-averaged criticality⁶⁵, Gross–Neveu–Yukawa criticality⁶⁶, and models of flat band physics^{67–70}. Other classes of models are Hubbard-type models and electron-phonon models, which have been recently studied^{35–47,49,71,72} using various HMC and DQMC algorithms. We note that for theories with dynamical gauge fields of continuous gauge

groups or non-linear-sigma models, HMC is less efficient due to the trouble it has in sampling the various topological sectors via molecular dynamics. In these cases, perhaps an HMC augmented with generative model sampling might make for a very efficient sampling⁷³. This issue is not present in the present model, since the magnitude of the $O(3)$ bosonic field is not restricted.

Potential algorithmic improvements

The implementation could benefit from more sophisticated approaches to solving the linear system arising from each integration step within the HMC algorithm. Currently, the computational cost of our preconditioned CG approach enjoys ideal scaling in the limit of large L . However, the scaling with respect to N_τ for $\Delta\tau$ fixed, i.e., the zero-temperature limit, is superlinear with respect to $\Delta\tau$. Several approaches to similar linear solves have been introduced in the LQCD and condensed matter literatures that could pay dividends here. These include Hasenbuch preconditioning^{40,74}, multi-grid preconditioners and solvers⁷⁵, and non-iterative ideas based on Schur complements⁴³. These approaches should also help for larger values of the Yukawa coupling, when the preconditioner used in this work would be less effective. Moreover, a mixed-precision implementation of CG on the GPU could provide significant speedup⁷⁶.

There is also room for further investigation into the HMC hyperparameter auto-tuning procedure, as there exist alternative criteria for tuning ε and n_{leap} . For example, one could instead tune ε based on the average acceptance rate or based on the stability of the numerical integrator. Meanwhile, the tuning of n_{leap} is by far the most expensive part of the warmup phase. The current state-of-the-art approach in the statistics literature for tuning n_{leap} is called the No U-Turn Sampler (NUTS)⁷⁷. Implementing NUTS in our setting might offer a similar advantage.

Methods

QMC path integral with fermions

In order to simulate the action of Eq. (2) we put the system on a finite square lattice of size L with periodic boundary conditions. The lattice spacing is set to one. The imaginary time direction is also discretized into N_τ slices, with the imaginary time spacing, or Trotter step, labeled as $\Delta\tau$. We set $\Delta\tau=0.1$. Fermions and boson obey anti-periodicity and periodicity, respectively, in the time direction. As in any path integral quantum Monte Carlo method involving fermions, we start by writing out the partition function and performing the Grassmann integral over the fermionic fields:

$$\begin{aligned} Z &= \int d\boldsymbol{\phi} e^{-S_B(\boldsymbol{\phi})} \int d\psi^* d\psi e^{-S_F(\psi, \psi^*) - S_{FB}(\boldsymbol{\phi}, \psi, \psi^*)} \\ &= \int d\boldsymbol{\phi} e^{-S_B(\boldsymbol{\phi})} \det(D(\boldsymbol{\phi}))^{N_f}. \end{aligned} \tag{16}$$

Here, $D(\boldsymbol{\phi})$ is the fermion matrix of size $2 \cdot 2 \cdot N_\tau \cdot L \cdot L = 4N_\tau L^2$, where the two factors of 2 correspond to the band and spin indices. Since the fermionic action is diagonal in the flavor index j , N_f enters trivially as a power of the determinant. Component-wise $D(\boldsymbol{\phi})$ is given by

$$\begin{aligned} D(\boldsymbol{\phi})_{(\alpha, s, \tau, x, y), (\alpha', s', \tau', x', y')} &= \delta_{\alpha, \alpha'} \delta_{s, s'} \left\{ \delta_{x, x'} \delta_{y, y'} (\delta_{\tau, \tau'} [-1 - \Delta\tau \mu_\alpha] \right. \\ &\quad \left. + \delta_{\tau+1, \tau'} [1 - 2\delta_{\tau, N_\tau-1}]) - \delta_{\tau, \tau'} \Delta\tau t_{\alpha, (x, y), (x', y')} \right\} \\ &\quad + \delta_{x, x'} \delta_{y, y'} \delta_{\tau, \tau'} \Delta\tau g e^{iQ_{\text{AF}} \cdot (x, y)} \boldsymbol{\phi}(\tau, x, y) \cdot \boldsymbol{\tau}_{\alpha, \alpha'} \sigma_{\alpha, \alpha'}^{(x)}, \end{aligned} \tag{17}$$

which is read directly from the action in Eq. (2). Here, τ is in the range $\{0, \dots, N_\tau - 1\}$, and x, y are in the range $\{0, \dots, L - 1\}$. $\delta_{a,b}$ is the Kronecker delta function. $\sigma_{\alpha, \alpha'}^{(x)}$ is the usual Pauli matrix, but acting on the band indices instead of the spin indices.

Stochastic formulation of the determinant

In the rest of this section, we will use the notation $\boldsymbol{\phi}$ (not bold) to denote the bosonic field, to emphasize its interpretation alternatively as either a vector of length $3N_\tau L^2$ or an array of size $3 \times N_\tau \times L \times L$, depending on context, as opposed to its equivalent representation as 3-component field $\boldsymbol{\phi} = \boldsymbol{\phi}(\tau, x, y)$.

In the usual BSS algorithm⁷⁸, the determinant in Eq. (16) is rewritten as

$$\det(D(\boldsymbol{\phi})) = \det \left(\mathbf{1} + \prod_{l=0}^{N_\tau-1} B_l(\boldsymbol{\phi}) \right), \tag{18}$$

where $\mathbf{1}$ and the $B_l(\boldsymbol{\phi})$ are square matrices of size $2 \cdot 2 \cdot L \cdot L = 4L^2$. This reduces the cost of exact computation of the determinant from the naive scaling of $\mathcal{O}((4N_\tau L^2)^3)$ to the improved scaling of $\mathcal{O}(N_\tau (4L^2)^3)$ for every configuration $\boldsymbol{\phi}$. Although this is better than the direct computation of the larger determinant, the scaling with respect to L is nonetheless quite severe.

As an alternative to the exact computation of the determinant, we use the pseudofermion method^{79,80} to evaluate it stochastically. This method is based on the expression

$$\det(A) \propto \int D\varphi D\varphi^* e^{-\varphi^* A^{-1} \varphi}, \tag{19}$$

where φ is an auxiliary complex bosonic field (called a ‘pseudofermion’) and A is any Hermitian positive definite matrix.

Although the anti-unitary symmetry of ref. 16 guarantees that $\det(D(\boldsymbol{\phi}))$ is always non-negative, $D(\boldsymbol{\phi})$ itself is not necessarily positive definite. However, we can write

$$\begin{aligned} \det(D)^{N_f} &= (\det(D) \det(D)^*)^{N_f/2} \\ &= (\det(D) \det(D^\dagger))^{N_f/2} = \det(DD^\dagger)^{N_f/2}. \end{aligned} \tag{20}$$

The matrix $D(\boldsymbol{\phi})D(\boldsymbol{\phi})^\dagger$ is guaranteed to be positive definite. Note that when N_f is odd, the integral in Eq. (19) comes with a fractional power, which limits us to the case of even N_f . From here on, we set $N_f=2$, and moreover this value is used in our simulations. Rewriting the determinant in this way, the partition function becomes

$$Z = \int d\boldsymbol{\phi} d\varphi^* d\varphi e^{-(S_B(\boldsymbol{\phi}) + S_{PF}(\boldsymbol{\phi}, \varphi))}, \tag{21}$$

where $S_{PF}(\boldsymbol{\phi}, \varphi) = \varphi^* (D(\boldsymbol{\phi})D(\boldsymbol{\phi})^\dagger)^{-1} \varphi$. Note that the dimension of φ is the same as $D(\boldsymbol{\phi})$, i.e., $4N_\tau L^2$. This new partition function defines a joint distribution $p(\boldsymbol{\phi}, \varphi)$, which can be sampled with Markov chain Monte Carlo (MCMC).

Solving the linear system

While Eq. (21) introduces a matrix inverse to the action, the expression can be evaluated efficiently by treating the application of $(DD^\dagger)^{-1}$ to φ as the solution to a linear system:

$$(DD^\dagger)\boldsymbol{\eta} = \varphi, \tag{22}$$

where $\boldsymbol{\eta}$ is the unknown. A wide class of iterative solvers are available to tackle this problem efficiently. Additionally, other non-iterative solvers have recently been employed in the application of HMC to the Hubbard model⁴³. Here, we use the conjugate gradient (CG) method⁸¹, a commonly used technique with practical advantages for Hermitian positive definite systems.

Choice of preconditioner. Iterative solvers can often be preconditioned with a transformation that improves the conditioning of the

linear system. For any preconditioner P , the equation $P^{-1}(DD^\dagger)\eta = P^{-1}\phi$ has the same solution as the original system but may be better conditioned. Within the CG algorithm, we choose a preconditioner as follows.

Note that $D(\phi)$ as defined in Eq. (17) can be split into two terms,

$$D(\phi) = K + V(\phi), \tag{23}$$

where the first term, K , i.e. the fermionic kinetic energy, consists of all terms in Eq. (17) that are independent of ϕ . Since K is a translation-invariant operator (modulo fermionic imaginary-time antiperiodicity), it can be diagonalized and inverted efficiently using fast Fourier operations and then used to construct an efficient preconditioner.

To wit, let \mathcal{F} denote the discrete Fourier transform with respect to the spacetime lattice indices (τ, x, y) , and let \mathcal{T} be the unitary diagonal ‘twist’ operator defined by

$$\mathcal{T}_{(\alpha,s,\tau,x,y),(\alpha',s',\tau',x',y')} = \delta_{(\alpha,s,\tau,x,y),(\alpha',s',\tau',x',y')} e^{i\pi\tau}. \tag{24}$$

Note that \mathcal{T} transforms imaginary-time-periodic fields to imaginary-time-antiperiodic fields, and as such we can write

$$K = \tilde{\mathcal{F}} \hat{K} \tilde{\mathcal{F}}^*, \tag{25}$$

where $\tilde{\mathcal{F}} := \mathcal{T}\mathcal{F}$ is the ‘twisted’ Fourier transform and \hat{K} is diagonal.

Our preconditioner is then defined by

$$P := (KK^\dagger)^{-1} = \tilde{\mathcal{F}} \hat{P} \tilde{\mathcal{F}}^*, \tag{26}$$

where $\hat{P} = \hat{K}\hat{K}^\dagger$ is diagonal. As such, P can be applied with computational cost that is linear (up to log factors) in the spacetime volume $N_t L^2$. In the limit where the coupling constant $g \rightarrow 0$, this is the perfect preconditioner, $P = (DD^\dagger)^{-1}$, achieving convergence in a single CG iteration. More generally, the preconditioner is effective in addressing poor conditioning due to the unboundedness of the differential operator term of D as the imaginary time step $\Delta\tau$ is refined, in analogy with, e.g., Laplacian preconditioners⁸² in the numerical PDE literature.

See also ref. 35 for a similar approach to preconditioning based on a noninteracting model. Note that such a preconditioner may not be satisfactory in general in the strong coupling limit, and we discuss possible future improvements in Section “Potential algorithmic improvements”. For the couplings studied in this paper, we do nonetheless observe improvement due to the preconditioner, as detailed in Supplementary Note 5.

HMC

Hybrid/Hamiltonian Monte Carlo (HMC) is a Markov chain Monte Carlo (MCMC) method in which numerical integration of Hamilton’s equations in a fictitious phase space produces sequential updates in a Markov chain^{50,83}. Historically used as tool for the study of lattice quantum chromodynamics (LQCD), HMC is currently the state-of-the-art method in the field, where it has yielded sub-1% errors on over $\sim 10^4$ degrees of freedom^{75,84}.

We now give a brief overview of HMC as applied to theories with fermions. As explained above, Eq. (21) defines a joint probability distribution,

$$p(\phi, \varphi) = \frac{1}{Z} e^{-S_B(\phi) - S_{FF}(\phi, \varphi)}, \tag{27}$$

that needs to be sampled to estimate observables. We adopt a Gibbs sampler approach to sampling the two fields ϕ, φ , wherein we update each one conditioned on the other in alternating fashion.

The φ sample can be drawn exactly from the conditional distribution for fixed ϕ , which is the complex normal distribution

$\mathcal{N}(0, D(\phi)D(\phi)^\dagger)$. This is achieved by sampling a standard complex Gaussian vector χ and setting $\varphi = D(\phi)\chi$.

Because the action is not quadratic in ϕ , sampling of this field remains non-trivial but can be handled with MCMC. We employ the HMC method for this purpose. The first step is to introduce a field π that acts as an artificial conjugate momentum to the field ϕ . Together, ϕ and π define phase space coordinates for a fictitious Hamiltonian system

$$\tilde{H}(\phi, \varphi, \pi) = K(\pi; M) + S(\phi, \varphi) = \frac{1}{2} \pi^\top M^{-1} \pi + S(\phi, \varphi), \tag{28}$$

consisting of ‘kinetic energy’ and ‘potential energy’ terms. The kinetic energy specifically is defined with respect to a choice of metric M , also called a mass matrix, which is given to be positive definite. The choice of M will be discussed later. This Hamiltonian allows us to define a new joint distribution $p(\phi, \varphi, \pi) \sim e^{-\tilde{H}(\phi, \varphi, \pi)}$, compatible with our target distribution after marginalization of π .

The new joint distribution can now be efficiently sampled by drawing π directly as a normal random vector with covariance M and then integrating the equations of motion of \tilde{H} ,

$$\frac{d\phi}{dt} = M^{-1}\pi \quad \text{and} \quad \frac{d\pi}{dt} = -\frac{\partial S(\phi, \varphi)}{\partial \phi}. \tag{29}$$

Here, t is a fictitious ‘time’ variable. The trajectory defined by Eqs. (29) is also known as the molecular dynamics trajectory. Once a new sample of (ϕ, π) is obtained, its ϕ component is taken as the new sample of $p(\phi, \varphi)$.

With these tools, we can construct the Markov chain update for ϕ as follows:

1. Generate a momentum sample from $\pi \sim \mathcal{N}(0, M)$.
2. Approximately integrate the molecular dynamics (29) with initial condition (ϕ, π) for some time to produce new configuration (ϕ', π') .
3. Accept the new configuration (ϕ', π') with probability $\alpha = \min\left(1, e^{-\{\tilde{H}(\phi', \pi') - \tilde{H}(\phi, \pi)\}}\right)$.

Exact integration of Eqs. (29) is not feasible and must therefore be performed numerically in discrete time. A suitable choice of numerical integration scheme is leapfrog integration⁸⁵, which preserves phase space volume and ensures that the final Metropolis accept/reject step preserves detailed balance.

Scale-invariant sampling

In high dimensions, local MCMC methods encounter difficulties in the setting of ‘poorly scaled,’ i.e., severely anisotropic, distributions. In order to guarantee a nonvanishing acceptance rate without correcting for this difficulty, the local moves in the MCMC sampler are constrained by the smallest scale present in the problem, and as such the autocorrelation time with respect to the large scales must grow accordingly. For a critical lattice model, the smallest scale present in the distribution is not bounded away from zero in the thermodynamic limit, and this difficulty must be addressed to maintain constant autocorrelation time in this limit. Existing approaches based on affine invariance are unsuitable in this context due either to a curse of dimensionality⁸⁶ or to unacceptable (i.e., at least quadratic) scaling⁸⁷ with respect to the spacetime volume $N_t L^2$. Relative to these approaches, we are able to exploit the known a priori structure of our model, i.e., translation invariance, to achieve a fast scale-invariant sampler. Specifically, within the context of HMC, we learn an optimal metric M ‘online.’ In contrast with⁴⁴, beyond translation invariance we do not assume any a priori functional form for the metric, as any a priori choice may struggle near criticality. Since M is updated adaptively, it is then important to tune the HMC hyperparameters ε and n_{leap}

adaptively in a scale-invariant fashion. For this task we adapt best practices from the statistics literature (as implemented, e.g., in ref. 52) to our setting, as we shall discuss below. These quantities (i.e., M , ε , and n_{leap}) are all determined during a warm-up phase, after which they are fixed for the duration of the run.

Online metric estimation. HMC as described above allows for an arbitrary choice of metric or mass matrix M . When the underlying probability distribution to be sampled is poorly scaled, the metric must be adapted to the scaling of the distribution to maintain a constant autocorrelation time. An idealized choice for M is given by the inverse of the exact covariance matrix $\Sigma_{\text{true}} = \langle \phi \phi^T \rangle$ of the underlying distribution p , i.e., by $M = \Sigma_{\text{true}}^{-1}$. In the case where p is the Gaussian $\mathcal{N}(0, \Sigma_{\text{true}})$, such a choice is equivalent to rescaling p to the standard (isotropic) Gaussian distribution $\mathcal{N}(0, \mathbf{1})$. More generally, such a choice can be viewed as rescaling p to a distribution with unit covariance matrix.

There are two obstacles to achieving this idealized choice. First, the exact covariance Σ_{true} is not known a priori. Second, even if it were known or estimated as $\Sigma \approx \Sigma_{\text{true}}$, the cost of even storing this matrix scales as $\sim (VN_\tau)^2$, which already imposes a computational bottleneck. Worse still, within the HMC algorithm, we must generate samples from $\mathcal{N}(0, \Sigma^{-1})$, which in general scales as $\sim (VN_\tau)^2$, i.e., as the cost of a Cholesky factorization of Σ .

The first of these obstacles can be tackled via online estimation of the covariance matrix Σ_{true} from samples. To wit, we can choose some initialization for Σ (typically $\Sigma = \mathbf{1}$) and run the HMC algorithm to produce a batch of S samples $\phi^{(1)}, \dots, \phi^{(S)}$ of the bosonic field, then reset $\Sigma \leftarrow \Sigma_{\text{est}} := \frac{1}{S} \sum_{s=1}^S \phi^{(s)} \phi^{(s)T}$ and repeat the procedure until Σ converges.

There is an appearance of circular reasoning to such an online estimation scheme in that it assumes the ability to draw samples $\phi - p$ via our MCMC sampler (which may be bottlenecked by a long autocorrelation time) in order to improve the autocorrelation of the MCMC method itself! However, it is important to realize that in the initial iteration, the autocorrelation time for the smallest eigenmodes of Σ_{true} is not long. Indeed, to get a nonvanishing acceptance probability, the step size in HMC must be tuned to be on the order of the smallest eigenvalue of the covariance. Due to anisotropy, this may result in very slow MCMC mixing in the highest eigendirections, but this does not interfere with our ability to estimate the smallest eigenmodes. After the first iteration, having corrected for the smallest eigenmodes, the procedure proceeds to correct the next smallest, etc.

Despite its theoretical appeal, there are two problems with this approach, from the point of view of computational scaling. First, to get an estimate $\Sigma_{\text{est}} \approx \Sigma_{\text{true}}$ of fixed accuracy requires a size-extensive number S of samples, due to the size-extensive numerical rank of Σ_{true} . Second, as mentioned earlier, the cost of storing and factorizing Σ_{est} is prohibitive. In fact, both problems can be addressed by exploiting translation invariance.

Indeed, since Σ_{true} is translation-invariant, it is equivalently diagonal in Fourier space—or more precisely, (3×3) -block-diagonal with respect to the vector index $l = 1, 2, 3$ of ϕ . Therefore $\Sigma_{\text{true}} = \mathcal{F} \hat{\Sigma}_{\text{true}} \mathcal{F}^T$, where \mathcal{F} represents the discrete Fourier transform with respect to the spacetime lattice indices and $\hat{\Sigma}_{\text{true}}$ is block-diagonal. As such the block-diagonal can be obtained as the expectation $\text{diag}([\hat{\Sigma}_{\text{true}}]_{ll}) = \langle \hat{\phi}_l \circ \hat{\phi}_l \rangle$, where samples $\hat{\phi}$ are obtained as $\hat{\phi} = \mathcal{F}\phi$ from samples ϕ and ‘ \circ ’ indicates the entrywise product. In practice, for simplicity we consider only the true diagonal of $\hat{\Sigma}$, which is sufficient to address the increasing anisotropy due to large volume. Concretely, we therefore form

$$\hat{\sigma} \leftarrow \hat{\sigma}_{\text{est}} := \frac{1}{S} \sum_{s=1}^S \hat{\phi} \circ \hat{\phi} \quad (30)$$

and let $\Sigma = \mathcal{F} \hat{\Sigma} \mathcal{F}^T$, where $\hat{\Sigma}$ is the diagonal matrix with diagonal $\hat{\sigma}$.

In order to draw sample momenta $\pi \sim \mathcal{N}(0, M)$ within the HMC algorithm, observe that $M = \Sigma^{-1} = (\mathcal{F} \hat{\Sigma}^{-1/2}) (\mathcal{F} \hat{\Sigma}^{-1/2})^*$. Therefore we can simply draw $z \sim \mathcal{N}(0, \mathbf{1})$ and let $\pi = \mathcal{F}(\hat{\sigma}^{-1/2} \circ z)$, where $\hat{\sigma}^{-1/2}$ indicates the entrywise inverse square root of $\hat{\sigma}$. The total cost of drawing such a sample is linear in the spacetime volume $L^2 N_\tau$, up to a log factor, as is the cost of evaluating the HMC energy of Eq. (28).

Since the off-diagonal elements of $\hat{\Sigma}$ are known a priori to be zero (and implicitly set to zero automatically), we only need $O(1)$ effective samples to estimate $\hat{\sigma}$ to fixed accuracy. Hence the total cost of each iterative update of $\hat{\Sigma}$ scales linearly in the spacetime volume $L^2 N_\tau$, up to log factors.

Online HMC hyperparameter tuning. Within the HMC algorithm, once the metric M is fixed, it remains to choose (subordinate to this choice) appropriate values for the integration step size $\varepsilon > 0$ and the number n_{leap} of leapfrog integration steps per proposal.

First we consider the choice of ε . Roughly speaking, one wants to choose ε as large as possible without sabotaging the acceptance rate. Quantitatively, let $\alpha(\varepsilon)$ denote the expected acceptance probability for a single leapfrog integration step of size ε . More precisely, recalling that $\alpha = \alpha(\phi, \pi, \phi', \pi')$ denotes the acceptance probability of a proposed move $(\phi, \pi) \rightarrow (\phi', \pi')$, we may define $\alpha(\varepsilon) = \langle \alpha(\phi, \phi') \rangle_{\varepsilon, n_{\text{leap}}=1}$, where the statistical expectation is computed by sampling $(\phi, \pi) - p$ and then sampling (ϕ', π') as a proposal from ϕ according to the HMC algorithm with step size ε and $n_{\text{leap}} = 1$. Then one wants to maximize ε subject to the inequality

$$(1 - \alpha(\varepsilon/2))^2 \leq 2(1 - \alpha(\varepsilon)). \quad (31)$$

The left-hand side estimates the probability of accepting two steps of size $\varepsilon/2$, whereas the right-hand side estimates the probability of accepting one step of size ε . Since the former strategy costs twice as many linear solves $(DD^*)^{-1}\phi$ as the latter, the latter is computationally preferable as long as the inequality holds.

In practice, we maintain an estimate for $\alpha(\varepsilon)$ and $\alpha(\varepsilon/2)$ based on an empirical average over a recent history of samples, and we increase or decrease ε when Eq. (31) is satisfied or violated, respectively.

Now we turn to the choice of n_{leap} . Within the actual HMC algorithm, we choose n_{leap} uniformly at random from the set $\{1, \dots, n_{\text{max}}\}$, independently for each iteration, where n_{max} is a hyperparameter to be determined adaptively. This ‘jittering’ procedure is standard in the statistics community⁵². The hyperparameter n_{max} is determined as the maximizer of the expected squared jump distance (ESJD)⁵¹

$$\text{ESJD}(n) := \left\langle (\phi - \phi')^T M (\phi - \phi') \alpha(\phi, \phi') \right\rangle_{\varepsilon, n_{\text{leap}}=n}. \quad (32)$$

Importantly, the metric M is needed here to correctly define the squared distance $|\phi - \phi'|_M^2 = (\phi - \phi')^T M (\phi - \phi')$.

In practice, similarly to $\alpha(\varepsilon)$, we maintain an estimate for $\text{ESJD}(n)$ based on an empirical average over recent history, and we choose n_{max} to maximize it.

During the warmup phase, the three components of ε , n_{max} and M are each tuned several times, until they stop changing appreciably.

Note that since the acceptance probability in HMC is directly determined by the energy conservation error in the integration of Hamilton’s equations (29), alternative higher-order integration schemes besides leapfrog can be considered to improve acceptance rates at the price of more expensive integration steps, cf. for example^{88,89}.

Estimation of observables

Here we describe the methods we use to compute physical observables directly in Fourier space, while keeping the computational cost near-linear in βV .

Bosonic observables. Bosonic observables can be computed directly as a sample average over the bosonic field variable ϕ . Specifically of interest is the SDW susceptibility

$$\chi(\omega, \mathbf{q}) = \int_0^\beta \sum_{\mathbf{r}} \langle \phi(\tau, \mathbf{r}) \cdot \phi(0, \mathbf{0}) \rangle e^{i\omega\tau - i\mathbf{q}\cdot\mathbf{r}} d\tau. \quad (33)$$

The SDW susceptibility can be estimated at linear cost (up to log factors) as

$$\chi(\omega, \mathbf{q}) = \langle |\hat{\phi}(\omega, \mathbf{q})|^2 \rangle \Delta\tau, \quad (34)$$

where we recall that $\hat{\phi} = \mathcal{F}\phi$ is the spacetime discrete Fourier transform of ϕ .

Fermionic observables. Fermionic observables require more care to estimate while maintaining almost-linear scaling. Of particular interest are the two-point correlator (in Fourier space)

$$G_{\alpha,s}(\omega, \mathbf{q}) = \int_0^\beta \sum_{\mathbf{r}} \langle \psi_{\alpha,s}(0, \mathbf{0}) \psi_{\alpha,s}(\tau, \mathbf{r})^* \rangle e^{i\omega\tau - i\mathbf{q}\cdot\mathbf{r}} d\tau. \quad (35)$$

and the superfluid (SF) density, which is defined in Supplementary Note 3.

The key observation is that all fermionic observables of interest can be phrased in terms of the expectation value of the diagonal of a matrix for which it is possible to perform efficient matrix-vector multiplications, i.e., as

$$\langle \text{diag}(O_\phi) \rangle, \quad (36)$$

where the statistical average is taken with respect to the bosonic density for ϕ and where O_ϕ is an operator that depends on the bosonic field. Rather than compute the diagonal entries individually, it is more efficient to recover them simultaneously via the identity

$$\text{diag}(O_\phi) = \mathbb{E}[\nu \odot (O_\phi \nu)], \quad (37)$$

where ν is a random vector with independent entries that take values ± 1 each with probability 1/2 and ‘ \mathbb{E} ’ indicates the expectation with respect to this distribution over ν . As before, ‘ \odot ’ indicates the entrywise product of vectors. This identity defines a randomized matrix-free algorithm for recovering a matrix diagonal from only $O(1)$ matrix-vector multiplications, which has appeared before in various works^{90,91}. Moreover, after taking traces of both sides, one recovers the famous Hutchinson trace estimator⁹². During the preparation of this work, this randomized diagonal estimator has also appeared in ref. 44 for the same purpose of computing fermionic observables, though their approach to computing quartic fermionic observables such as the SF density differ somewhat from ours.

Since we must average $\text{diag}(O_\phi)$ over the bosonic distribution for ϕ , we can in fact obtain a consistent estimator by independently drawing a single vector $\nu^{(s)}$ for each bosonic sample $\phi^{(s)}$, $s=1, \dots, S$, where S is the sample size of our empirical average and estimating

$$\text{diag}(O_\phi) \approx \frac{1}{S} \sum_{s=1}^S \nu^{(s)} \odot O_\phi \nu^{(s)}. \quad (38)$$

Further details of the implementation of the approach to computing fermionic observables (i.e., the specification of O_ϕ for observables of interest) can be found in Supplementary Note 3.

Numerical performance

Here we summarize the numerical performance of our algorithm, with more details provided in Supplementary Note 6. As noted in the Introduction, in the presence of a critical slowing down, the HMC algorithm can require a number of integration steps per effective sample of $O(\beta^{1/4+z_1} V^{1/4+z_2})$, where z_1, z_2 are to be determined empirically. To extract these exponents, we benchmark our algorithm across spacetime lattice sizes at the critical parameters of the theory, scaling with respect to $V=L^2$ and N_τ separately. We track the growth of the integrated autocorrelation time τ_{int} of the total SDW susceptibility $\chi \equiv \chi(0, 0)$ at criticality, and use it to quantify the wallclock time and number of HMC integration steps per effective sample: $\tau_{\text{int}} \times n_{\text{leap}}$.

We find that our algorithm exhibits constant scaling of τ_{int} with respect to both lattice volume and inverse temperature. For the exponents, we find $z_1 \approx 0.5, z_2 \approx 0$. This implies an absence of critical slowing down with respect to the lattice volume V for the auto-tuned HMC algorithm presented in this work.

The actual wall clock time in turn depends on the cost of the linear solves of the form Eq. (22). With the preconditioned CG approach described in Section ‘‘Choice of preconditioner’’, we observe near linear scaling in wall clock time per effective sample with respect to V , while the scaling is superlinear with respect to N_τ . The linear solve can be approached with more advanced techniques, as well as more refined GPU parallelism. These directions are discussed in Section ‘‘Potential algorithmic improvements’’.

Data availability

The data analyzed in the current manuscript is available from the corresponding author upon reasonable request.

Code availability

All the custom codes used in this study can be requested from the corresponding author.

References

1. Sachdev, S., *Quantum phase transitions* (Cambridge University Press, 2011).
2. Lee, S.-S. Recent developments in non-Fermi liquid theory. *Ann. Rev. Condens. Matter Phys.* **9**, 227–244, <https://doi.org/10.1146/annurev-conmatphys-031016-025531> (2018).
3. Greene, R. L., Mandal, P. R., Poniatowski, N. R., & Sarkar, T. The strange metal state of the electron-doped cuprates. *Ann. Rev. Condens. Matter Phys.* **11**, 213–229 <https://doi.org/10.1146/annurev-conmatphys-031119-050558> (2020).
4. Paglione, J. & Greene, R. L. High-temperature superconductivity in iron-based materials. *Nat. Phys.* **6**, 645–658 <https://doi.org/10.1038/nphys1759> (2010).
5. Stockert, O. & Steglich, F. Unconventional quantum criticality in heavy-fermion compounds. *Ann. Rev. Condens. Matter Phys.* **2**, 79–99 <https://doi.org/10.1146/annurev-conmatphys-062910-140546> (2011).
6. Abanov, A., Chubukov, A. V., & Schmalian, J. Quantum-critical theory of the spin-fermion model and its application to cuprates: normal state analysis. *Adv. Phys.* **52**, 119 <https://doi.org/10.1080/0001873021000057123> (2003).
7. Abanov, A. & Chubukov, A. V. Spin-Fermion model near the quantum critical point: one-loop renormalization group results. *Phys. Rev. Lett.* **84**, 5608 <https://doi.org/10.1103/PhysRevLett.84.5608> (2000).
8. Abanov, A. & Chubukov, A. Anomalous scaling at the quantum critical point in itinerant antiferromagnets. *Phys. Rev. Lett.* **93**, 255702 <https://doi.org/10.1103/PhysRevLett.93.255702> (2004).

9. Metlitski, M. A. & Sachdev, S. Quantum phase transitions of metals in two spatial dimensions. II. Spin density wave order. *Phys. Rev. B* **82**, 075128 <https://doi.org/10.1103/PhysRevB.82.075128> (2010).
10. Sur, S. & Lee, S.-S. Quasilocal strange metal. *Phys. Rev. B* **91**, 125136 <https://doi.org/10.1103/PhysRevB.91.125136> (2015).
11. Lunts, P., Schliefl, A., & Lee, S.-S. Emergence of a control parameter for the antiferromagnetic quantum critical metal. *Phys. Rev. B* **95**, 245109 <https://doi.org/10.1103/PhysRevB.95.245109> (2017).
12. Schliefl, A., Lunts, P., & Lee, S.-S. Exact critical exponents for the antiferromagnetic quantum critical metal in two dimensions. *Phys. Rev. X* **7**, 021010 <https://doi.org/10.1103/PhysRevX.7.021010> (2017).
13. Lunts, P. & Patel, A. A. Many-body chaos in the antiferromagnetic quantum critical metal. *Phys. Rev. B* **100**, 235104 <https://doi.org/10.1103/PhysRevB.100.235104> (2019).
14. Hertz, J. A. Quantum critical phenomena. *Phys. Rev. B* **14**, 1165 <https://doi.org/10.1103/PhysRevB.14.1165> (1976).
15. Millis, A. J. Effect of a nonzero temperature on quantum critical points in itinerant fermion systems. *Phys. Rev. B* **48**, 7183 <https://doi.org/10.1103/PhysRevB.48.7183> (1993).
16. Berg, E., Metlitski, M. A., & Sachdev, S. Sign-problem-free quantum Monte Carlo of the onset of antiferromagnetism in metals. *Science* **338**, 1606–1609 <https://doi.org/10.1126/science.1227769> (2012).
17. Gerlach, M. H., Schattner, Y., Berg, E., & Trebst, S. Quantum critical properties of a metallic spin-density-wave transition. *Phys. Rev. B* **95**, 035124 <https://doi.org/10.1103/PhysRevB.95.035124> (2017).
18. Wang, X., Schattner, Y., Berg, E., & Fernandes, R. M. Superconductivity mediated by quantum critical antiferromagnetic fluctuations: the rise and fall of hot spots. *Phys. Rev. B* **95**, 174520 <https://doi.org/10.1103/PhysRevB.95.174520> (2017).
19. Wang, X., Wang, Y., Schattner, Y., Berg, E., & Fernandes, R. M. Fragility of charge order near an antiferromagnetic quantum critical point. *Phys. Rev. Lett.* **120**, 247002 <https://doi.org/10.1103/PhysRevLett.120.247002> (2018).
20. Schattner, Y., Gerlach, M. H., Trebst, S., & Berg, E. Competing orders in a nearly antiferromagnetic metal. *Phys. Rev. Lett.* **117**, 097002 <https://doi.org/10.1103/PhysRevLett.117.097002> (2016).
21. Bauer, C., Schattner, Y., Trebst, S., & Berg, E. Hierarchy of energy scales in an O(3) symmetric antiferromagnetic quantum critical metal: a Monte Carlo study. *Phys. Rev. Res.* **2**, 023008 <https://doi.org/10.1103/PhysRevResearch.2.023008> (2020).
22. Li, Z.-X., Wang, F., Yao, H., & Lee, D.-H. What makes the T_c of monolayer FeSe on SrTiO₃ so high: a sign-problem-free quantum Monte Carlo study. *Sci. Bull.* **61**, 925–930 <https://doi.org/10.1007/s11434-016-1087-x> (2016).
23. Li, Z.-X., Wang, F., Yao, H., & Lee, D.-H. Nature of the effective interaction in electron-doped cuprate superconductors: a sign-problem-free quantum Monte Carlo study. *Phys. Rev. B* **95**, 214505 <https://doi.org/10.1103/PhysRevB.95.214505> (2017).
24. Berg, E., Lederer, S., Schattner, Y., & Trebst, S. Monte Carlo studies of quantum critical metals. *Ann. Rev. Condensed Matter Phys.* **10**, 63–84 <https://doi.org/10.1146/annurev-conmatphys-031218-013339> (2019).
25. Li, Z.-X. and Yao, H. Sign-problem-free Fermionic quantum Monte Carlo: developments and applications. *Ann. Rev. Condensed Matter Phys.* **10**, 337–356 <https://doi.org/10.1146/annurev-conmatphys-033117-054307> (2019).
26. Schattner, Y., Lederer, S., Kivelson, S. A., & Berg, E. Ising nematic quantum critical point in a metal: a Monte Carlo Study. *Phys. Rev. X* **6**, 031028 <https://doi.org/10.1103/PhysRevX.6.031028> (2016).
27. Lederer, S., Schattner, Y., Berg, E., & Kivelson, S. A. Superconductivity and non-Fermi liquid behavior near a nematic quantum critical point. *Proc. Natl Acad. Sci.* **114**, 4905–4910 <https://doi.org/10.1073/pnas.1620651114> (2017).
28. Klein, A., Chubukov, A. V., Schattner, Y., & Berg, E. Normal state properties of quantum critical metals at finite temperature. *Phys. Rev. X* **10**, 031053 <https://doi.org/10.1103/PhysRevX.10.031053> (2020).
29. Xu, X. Y., Sun, K., Schattner, Y., Berg, E., & Meng, Z. Y. Non-Fermi liquid at (2+1)D ferromagnetic quantum critical point. *Phys. Rev. X* **7**, 031058 <https://doi.org/10.1103/PhysRevX.7.031058> (2017).
30. Xu, X. Y., Klein, A., Sun, K., Chubukov, A. V., & Meng, Z. Y. Identification of non-Fermi liquid fermionic self-energy from quantum Monte Carlo data. *npj Quantum Mater.* **5**, 65 <https://doi.org/10.1038/s41535-020-00266-6> (2020).
31. Liu, Z. H., Pan, G., Xu, X. Y., Sun, K., & Meng, Z. Y. Itinerant quantum critical point with fermion pockets and hotspots. *Proc. Natl Acad. Sci.* **116**, 16760 <https://doi.org/10.1073/pnas.1901751116> (2019).
32. Liu, Y. et al. Dynamical exponent of a quantum critical itinerant ferromagnet: a Monte Carlo study. *Phys. Rev. B* **105**, L041111 <https://doi.org/10.1103/PhysRevB.105.L041111> (2022).
33. Motta, M., Shee, J., Zhang, S., & Chan, G. K.-L. Efficient ab initio auxiliary-field quantum Monte Carlo calculations in Gaussian bases via low-rank tensor decomposition. *J. Chem. Theory Comput.* **15**, 3510–3521 <https://doi.org/10.1021/acs.jctc.8b00996> (2019).
34. He, Y.-Y., Shi, H., and Zhang, S. Reaching the continuum limit in finite-temperature ab initio field-theory computations in many-fermion systems. *Phys. Rev. Lett.* **123**, 136402 <https://doi.org/10.1103/PhysRevLett.123.136402> (2019).
35. Scalettar, R. T., Scalapino, D. J., Sugar, R. L., & Toussaint, D. Hybrid molecular-dynamics algorithm for the numerical simulation of many-electron systems. *Phys. Rev. B* **36**, 8632 <https://doi.org/10.1103/PhysRevB.36.8632> (1987).
36. Hirsch, J. E. Stable monte carlo algorithm for fermion lattice systems at low temperatures. *Phys. Rev. B* **38**, 12023 <https://doi.org/10.1103/PhysRevB.38.12023> (1988).
37. Beyl, S., Goth, F., & Assaad, F. F. Revisiting the hybrid quantum Monte Carlo method for Hubbard and electron-phonon models. *Phys. Rev. B* **97**, 085144 <https://doi.org/10.1103/PhysRevB.97.085144> (2018).
38. Buividovich, P., Smith, D., Ulybyshev, M., & von Smekal, L. Hybrid Monte Carlo study of competing order in the extended fermionic Hubbard model on the hexagonal lattice. *Phys. Rev. B* **98**, 235129 <https://doi.org/10.1103/PhysRevB.98.235129> (2018).
39. Wynen, J.-L., Berkowitz, E., Körber, C., Lähde, T. A., & Luu, T. Avoiding ergodicity problems in lattice discretizations of the Hubbard model. *Phys. Rev. B* **100**, 075141 <https://doi.org/10.1103/PhysRevB.100.075141> (2019).
40. Krieg, S., Luu, T., Ostmeyer, J., Papaphilippou, P., & Urbach, C. Accelerating hybrid Monte Carlo simulations of the Hubbard model on the hexagonal lattice. *Comput. Phys. Commun.* **236**, 15–25 <https://doi.org/10.1016/j.cpc.2018.10.008> (2019).
41. Ostmeyer, J. et al. Semimetal–Mott insulator quantum phase transition of the Hubbard model on the honeycomb lattice. *Phys. Rev. B* **102**, 245105 <https://doi.org/10.1103/PhysRevB.102.245105> (2020).
42. Ostmeyer, J. et al. Antiferromagnetic character of the quantum phase transition in the Hubbard model on the honeycomb lattice. *Phys. Rev. B* **104**, 155142 <https://doi.org/10.1103/PhysRevB.104.155142> (2021).
43. Ulybyshev, M., Kintscher, N., Kahl, K., & Buividovich, P. Schur complement solver for Quantum Monte-Carlo simulations of strongly interacting fermions. *Comput. Phys. Commun.* **236**, 118–127 <https://doi.org/10.1016/j.cpc.2018.10.023> (2019).
44. Cohen-Stead, B. et al. Fast and scalable quantum monte carlo simulations of electron-phonon models. *Phys. Rev. E* <https://doi.org/10.48550/ARXIV.2203.01291> (2022).
45. Drut, J. E. & Lähde, T. A. Is graphene in vacuum an insulator? *Phys. Rev. Lett.* **102**, 026802 <https://doi.org/10.1103/PhysRevLett.102.026802> (2009).

46. Hands, S. & Strouthos, C. Quantum critical behavior in a graphene-like model. *Phys. Rev. B* **78**, 165423 <https://doi.org/10.1103/PhysRevB.78.165423> (2008).
47. Ulybyshev, M. V., Buividovich, P. V., Katsnelson, M. I., & Polikarpov, M. I. Monte Carlo study of the semimetal-insulator phase transition in monolayer graphene with a realistic interelectron interaction potential. *Phys. Rev. Lett.* **111**, 056801 <https://doi.org/10.1103/PhysRevLett.111.056801> (2013).
48. Brower, R., Rebbi, C., & Schaich, D. Hybrid Monte Carlo simulation on the graphene hexagonal lattice. *High Energy Phys. Lattice* <https://doi.org/10.48550/ARXIV.1204.5424> (2012).
49. Smith, D. & von Smekal, L. Monte Carlo simulation of the tight-binding model of graphene with partially screened Coulomb interactions. *Phys. Rev. B* **89**, 195429 <https://doi.org/10.1103/PhysRevB.89.195429> (2014).
50. Neal, R. MCMC using Hamiltonian dynamics. In: *Handbook of Markov Chain Monte Carlo*, (eds. Brooks, S., Gelman, A., Jones, G. & Meng, X.-L.) (Chapman & Hall/CRC Press, 2011) Chap. 5.
51. Pasarica, C. & Gelman, A. Adaptively scaling the metropolis algorithm using expected squared jumped distance. *Stat. Sin.* **20**, 343 <http://www.jstor.org/stable/24308995> (2010).
52. Carpenter, B. et al. Stan: a probabilistic programming language. *J. Stat. Softw.* **76**, <https://doi.org/10.18637/jss.v076.i01> (2017).
53. Borges, F., Borisso, A., Singh, A., Schlieff, A., & Lee, S. S. Field-theoretic functional renormalization group formalism for non-Fermi liquids and its application to the antiferromagnetic quantum critical metal in two dimensions. *Annals of Physics* **450**, 169221 (2023).
54. Schlieff, A., Lunts, P., & Lee, S.-S. Noncommutativity between the low-energy limit and integer dimension limits in the ϵ expansion: a case study of the antiferromagnetic quantum critical metal. *Phys. Rev. B* **98**, 075140 <https://doi.org/10.1103/PhysRevB.98.075140> (2018).
55. Gazit, S., Randeria, M., & Vishwanath, A. Emergent Dirac fermions and broken symmetries in confined and deconfined phases of Z_2 gauge theories. *Nat. Phys.* **13**, 484–490 <https://doi.org/10.1038/nphys4028> (2017).
56. Assaad, F. F. & Grover, T. Simple Fermionic model of deconfined phases and phase transitions. *Phys. Rev. X* **6**, 041049 <https://doi.org/10.1103/PhysRevX.6.041049> (2016).
57. Gazit, S., Assaad, F. F., Sachdev, S., Vishwanath, A., & Wang, C. Confinement transition of Z_2 gauge theories coupled to massless fermions: Emergent quantum chromodynamics and $SO(5)$ symmetry. *Proc. Natl Acad. Sci.* **115**, E6987–E6995 <https://doi.org/10.1073/pnas.1806338115> (2018).
58. Hohenadler, M. & Assaad, F. F. Fractionalized metal in a Falicov-Kimball model. *Phys. Rev. Lett.* **121**, 086601 <https://doi.org/10.1103/PhysRevLett.121.086601> (2018).
59. Xu, X. Y. et al. Monte Carlo study of lattice compact quantum electrodynamics with fermionic matter: the parent state of quantum phases. *Phys. Rev. X* **9**, 021022 <https://doi.org/10.1103/PhysRevX.9.021022> (2019).
60. Chen, C., Xu, X. Y., Qi, Y., & Meng, Z. Y. Metal to orthogonal metal transition. *Chinese Phys. Lett.* **37** <https://doi.org/10.1088/0256-307x/37/4/047103>.
61. Gazit, S., Assaad, F. F., & Sachdev, S. Fermi surface reconstruction without symmetry breaking. *Phys. Rev. X* **10**, 041057 <https://doi.org/10.1103/PhysRevX.10.041057> (2020).
62. Hofmann, J. S., Assaad, F. F., & Grover, T. Fractionalized Fermi liquid in a frustrated Kondo lattice model. *Phys. Rev. B* **100**, 035118 <https://doi.org/10.1103/PhysRevB.100.035118> (2019).
63. Sato, T., Assaad, F. F., & Grover, T. Quantum Monte Carlo simulation of frustrated Kondo lattice models. *Phys. Rev. Lett.* **120**, 107201 <https://doi.org/10.1103/PhysRevLett.120.107201> (2018).
64. Sato, T., Hohenadler, M., Grover, T., McGreevy, J., & Assaad, F. F. Topological terms on topological defects: a quantum Monte Carlo study. *Phys. Rev. B* **104**, L161105 <https://doi.org/10.1103/PhysRevB.104.L161105> (2021).
65. Patel, A. A., Guo, H., Esterlis, I., & Sachdev, S. Universal t-linear resistivity in two-dimensional quantum-critical metals from spatially random interactions. Preprint at <https://arxiv.org/abs/2203.04990> (2022).
66. Lang, T. C. & Läuchli, A. M. Quantum Monte Carlo simulation of the chiral Heisenberg gross-Neveu-Yukawa phase transition with a single Dirac cone. *Phys. Rev. Lett.* **123**, 137602 <https://doi.org/10.1103/PhysRevLett.123.137602> (2019).
67. Hofmann, J. S., Berg, E., & Chowdhury, D. Superconductivity, charge density wave, and supersolidity in flat bands with tunable quantum metric. Preprint at <https://arxiv.org/abs/2204.02994> (2022).
68. Zhang, X., Pan, G., Zhang, Y., Kang, J., & Meng, Z. Y. Momentum space quantum monte carlo on twisted bilayer graphene. *Chin. Phys. Lett.* **38**, 077305 <https://doi.org/10.1088/0256-307x/38/7/077305> (2021).
69. Fang, S.-C., Liu, G.-K., Lin, H.-Q., & Huang, Z.-B. Quantum Monte Carlo study of magnetic ordering and superconducting pairing symmetry in twisted bilayer graphene. *Phys. Rev. B* **100**, 115135 <https://doi.org/10.1103/PhysRevB.100.115135> (2019).
70. Hofmann, J. S., Khalaf, E., Vishwanath, A., Berg, E., & Lee, J. Y. Fermionic Monte Carlo study of a realistic model of twisted bilayer graphene. *Phys. Rev. X* **12**, 011061 <https://doi.org/10.1103/PhysRevX.12.011061> (2022).
71. Brower, R. C., Schaich, D. & Rebbi, C. Hybrid Monte Carlo simulation on the graphene hexagonal lattice. *PoS Lattice* **2011**, 056 (2012).
72. Zhang, C. et al. Bipolaronic high-temperature superconductivity. *Phys. Rev. X* **13**, 011010 (2023).
73. Kanwar, G. et al. Equivariant flow-based sampling for lattice gauge theory. *Phys. Rev. Lett.* **125**, 121601 <https://doi.org/10.1103/PhysRevLett.125.121601> (2020).
74. Hasenbusch, M. & Jansen, K. Speeding up the Hybrid-Monte-Carlo algorithm for dynamical fermions. *Nucl. Phys. B Proc. Suppl.* **106**, 1076–1078 [https://doi.org/10.1016/S0920-5632\(01\)01933-8](https://doi.org/10.1016/S0920-5632(01)01933-8) (2002).
75. Clark, M. A. et al. Accelerating Lattice QCD Multigrid on GPUs Using Fine-Grained Parallelization. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '16 (IEEE Press, 2016).
76. Clark, M., Babich, R., Barros, K., Brower, R., & Rebbi, C. Solving lattice QCD systems of equations using mixed precision solvers on GPUs. *Comput. Phys. Commun.* **181**, 1517–1528 <https://doi.org/10.1016/j.cpc.2010.05.002> (2010).
77. Hoffman, M. D. & Gelman, A. The No-U-Turn Sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**, 1593 (2014) <http://jmlr.org/papers/v15/hoffman14a.html>.
78. Blankenbecler, R., Scalapino, D. J., & Sugar, R. L. Monte Carlo calculations of coupled boson-fermion systems. I. *Phys. Rev. D* **24**, 2278 <https://doi.org/10.1103/PhysRevD.24.2278> (1981).
79. Gattringer, C. & Lang, C. B. Quantum chromodynamics on the lattice. *Lect. Notes Phys.* **788**, 1 <https://doi.org/10.1007/978-3-642-01850-3> (2010).
80. Fucito, F., Marinari, E., Parisi, G., & Rebbi, C. A proposal for Monte Carlo simulations of fermionic systems. *Nucl. Phys. B* **180**, 369–377 [https://doi.org/10.1016/0550-3213\(81\)90055-9](https://doi.org/10.1016/0550-3213(81)90055-9) (1981).
81. Fletcher, R. & Reeves, C. M. Function minimization by conjugate gradients. *Comput. J.* **7**, 149–154 <https://doi.org/10.1093/comjnl/7.2.149> (1964).
82. Gergelits, T., Mardal, K.-A., Nielsen, B. F., & Strakoš, Z. Laplacian preconditioning of elliptic PDEs: Localization of the eigenvalues of the discretized operator. *SIAM J. Numer. Anal.* **57**, 1369 (2019).

83. Duane, S., Kennedy, A., Pendleton, B. J., & Roweth, D. Hybrid Monte Carlo. *Phys. Lett. B* **195**, 216–222 [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X) (1987).
84. Lüscher, M. Stochastic locality and master-field simulations of very large lattices. *EPJ Web Conf.* **175**, 01002 <https://doi.org/10.1051/epjconf/201817501002> (2018).
85. Hockney, R. & Eastwood, J. *Computer simulation using particles* (CRC Press, 2021) <https://doi.org/10.1201/9780367806934>.
86. Goodman, J. & Weare, J. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science* **5**, 65 (2010).
87. Garbuno-Inigo, A., Nüsken, N., & Reich, S. Affine Invariant Interacting Langevin Dynamics for Bayesian Inference. *SIAM J. Appl. Dyn. Syst.* **19**, 1633 (2020).
88. Kennedy, A. & Pendleton, B. Cost of the generalised hybrid Monte Carlo algorithm for free field theory. *Nucl. Phys. B* **607**, 456–510 [https://doi.org/10.1016/S0550-3213\(01\)00129-8](https://doi.org/10.1016/S0550-3213(01)00129-8) (2001).
89. Calvo, M., Sanz-Alonso, D., & Sanz-Serna, J. HMC: Reducing the number of rejections by not using leapfrog and some results on the acceptance rate. *J. Comput. Phys.* **437**, 110333 <https://doi.org/10.1016/j.jcp.2021.110333> (2021).
90. Bekas, C., Kokiopoulou, E., & Saad, Y. An estimator for the diagonal of a matrix. *Appl. Numer. Math.* **57**, 1214–1229 (2007).
91. Tang, J. M. & Saad, Y. A probing method for computing the diagonal of a matrix inverse. *Numer. Linear Algebra Appl.* **19**, 485 (2012).
92. Hutchinson, M. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Commun. Stat.* **19**, 433 (1990).

Acknowledgements

We are extremely indebted to Bob Carpenter for teaching us the auto-tuning procedure used in industry applications of HMC and employed in this paper. We thank Phiala Shanahan for insightful discussions. P.L. thanks Subir Sachdev, the attendees of the Aspen Center for Physics conference ‘New Directions in Strong Correlation Physics: From Strange Metals to Topological Superconductivity,’ and in particular Sung-Sik Lee for useful discussions. We thank Sung-Sik Lee for comments on the manuscript. We thank Nick Carriero and the other members of the Scientific Computing Core at Flatiron Institute for their assistance in running the codes. M.A. and M.L. are grateful for the hospitality of the Center for Computational Quantum Physics at the Flatiron Institute. The Flatiron Institute is a division of the Simons Foundation. This work was partially completed at the Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1607611. P.L. is supported by the Simons Foundation. M.A. is supported by the National Science Foundation under the award PHY-2141336. M.L. is supported by the National Science Foundation under Award No. 1903031.

Author contributions

P.L. conceived and directed the project, wrote the majority of the code, and collected and analyzed the data. M.A. provided the initial HMC code design, conceived and led the GPU implementation, wrote the code for several observables, and helped with the data analysis. M.L. conceived the ideas of using translation invariance in the metric, the stochastic estimators for the fermionic observables and the choice of preconditioner, worked out the expression for the superfluid density, and wrote the code for several observables. All authors extensively discussed the project and co-wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-023-37686-4>.

Correspondence and requests for materials should be addressed to Peter Lunts.

Peer review information *Nature Communications* thanks Florian Goth, SangEun Han, and the other, anonymous, reviewer for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023