

CHCRUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Predicting impurity spectral functions using machine learning

Erica J. Sturm, Matthew R. Carbone, Deyu Lu, Andreas Weichselbaum, and Robert M. Konik

Phys. Rev. B 103, 245118 — Published 14 June 2021

DOI: 10.1103/PhysRevB.103.245118

Predicting impurity spectral functions using machine learning

Erica J. Sturm,^{1,*} Matthew R. Carbone,^{2,†} Deyu Lu,³ Andreas Weichselbaum,¹ and Robert M. Konik¹

¹Condensed Matter Physics and Materials Science Division,

Brookhaven National Laboratory, Upton, New York 11973, USA

²Department of Chemistry, Columbia University, New York, New York 10027, USA

³Center for Functional Nanomaterials, Brookhaven National Laboratory, Upton, New York 11973, USA

(Dated: May 12, 2021)

The Anderson Impurity Model (AIM) is a canonical model of quantum many-body physics. Here we investigate whether machine learning models, both neural networks (NN) and kernel ridge regression (KRR), can accurately predict the AIM spectral function in all of its regimes, from empty orbital, to mixed valence, to Kondo. To tackle this question, we construct two large spectral databases containing approximately 410k and 600k spectral functions of the single-channel impurity problem. We show that the NN models can accurately predict the AIM spectral function in all of its regimes, with point-wise mean absolute errors down to 0.003 in normalized units. We find that the trained NN models outperform models based on KRR and enjoy a speedup on the order of 10^5 over traditional AIM solvers. The required size of the training set of our model can be significantly reduced using furthest point sampling in the AIM parameter space, which is important for generalizing our method to more complicated multi-channel impurity problems of relevance to predicting the properties of real materials.

I. INTRODUCTION

Describing the physics of strongly correlated quantum many-body systems in real material systems is a signature challenge. In weakly correlated systems like simple metals, semiconductors, and band insulators, the physics is single-particle in nature and tools like Landau-Fermi liquid theory work well. However in materials where correlations are not weak, the single-particle picture is typically insufficient to describe the physics at low energy scales where emergent, completely novel phenomena can arise. Such physics is of greatest interest because they gift correlated materials with exceptional properties ranging over high temperature superconductivity [1, 2], colossal magnetoresistance [3], heavy fermion behavior [4], immense thermopower [5, 6], and huge volume collapses [7] to name but a few.

Measuring the response functions of applied weak external stimuli represents a key means to probe the properties of a strongly correlated material. However, of all the properties of a strongly correlated system, the response functions are the most difficult to ascertain theoretically. The response functions require not only knowledge of the ground state properties of a correlated material, but detailed knowledge of its excited state structure together with matrix elements of the observables of interest (for example, electric and heat currents). Many different theoretical approaches exist for resolving this difficult problem. Here, our motivational focus is one technique that has shown great promise for being able to categorically describe wide classes of correlated materials: dynamical mean field theory (DMFT) [8, 9].



FIG. 1: Cartoon of a typical Anderson impurity model with a set of physical input parameters (left). The physical properties are fed as training data into a neural network (center), which then learns to predict the spectral function (right).

DMFT is a Green's function method [10] that sums over an infinite set of Feynman diagrams consistent with the self-energy of the single-particle Green's function being local. In practice, performing this infinite summation amounts to solving a self-consistent quantum impurity problem. Besides lying at the heart of DMFT, quantum impurity problems are interesting many-body systems in and of themselves. They also describe magnetic impurities in metallic systems [11, 12], engineered quantum dots in bosonic or fermionic environments [13–15], and boundary edge modes in topologically non-trivial systems [16]. They also generically experience low-energy dynamically generated phenomena that are beyond perturbation theory such as the Kondo effect [17, 18] and the attendant Abrikosov-Suhl resonance that appears at low temperatures and frequencies in the quantum impurity's spectral response function.

Different techniques are available to find the spectral function of a quantum impurity problem. Included among those that are numerically exact are continuoustime quantum Monte Carlo simulations [19, 20], the nu-

 ^{*} esturm@bnl.gov; These authors contributed equally to this work
 † mrc2215@columbia.edu; These authors contributed equally to this work

merical renormalization group (NRG) [21], and the density matrix renormalization group [22]. The number of channels in the impurity problem determines the complexity of computing the spectral function (in the context of DMFT, the number of channels in the effective impurity model matches the number of bands involved in the underlying material). Single channel impurity models are relatively cheap to solve numerically while multiple channel impurity problems are exponentially more challenging. For example, five band f-electron materials have associated impurity models requiring petascale computational resources to accurately solve [23]. Furthermore, DMFT embeds quantum impurity model solutions in a self-consistent loop, requiring *multiple* solutions for final convergence.

In this light, we ask if machine learning (ML) approaches offer an alternative to expensive many-body simulations of impurity response functions. This question was first posed in Ref. 24, where kernel ridge regression (KRR) models were trained on a small database of about 5000 spectral functions computed at imaginary frequencies. A focus of this study was to understand the optimal parameterization of the spectral function for training purposes, finding that a representation in terms of Legendre polynomials worked best. More recently Ref. 25 used a set of neural networks to train a spectral solver for a quantum impurity connected to a bath of six sites, and each network was trained to predict the spectral function at a single frequency.

In the work presented herein, we investigate whether an individual model, whether it be KRR or a neural network, can predict the impurity response function in regimes where the relevant energy scales are separated by orders of magnitude and where temperature and magnetic field are also parameters. To this end, we have constructed large (~ 10^5) databases of high fidelity spectral functions in the thermodynamic limit using NRG. While large, our trial set size is commensurate with datasets used for other quantum physical/chemical applications [26–28]. We examine the dependence of these results on the training set size, as generating large training sets for the multi-channel impurity problems, the problem of ultimate interest, is much more computationally intensive. We use NRG to create our databases as it is an approach that (i) can reliably reach arbitrary, exponentially small dynamically generated energy scales; (ii) computes spectral properties directly on the real frequency axis, and (iii) can work with arbitrary temperatures in an efficient, systematic manner [29–33]. To the best of our knowledge, such high quality NRG databases do not exist for the AIM model.

II. QUANTUM IMPURITY MODEL

In this work we consider the most elementary of quantum impurity models, the single impurity Anderson model (SIAM) [34] with a fixed hybridization function.

While we aspire to use ML algorithms [35, 36] to study multi-channel impurity problems, we begin here with the simpler SIAM test environment. The SIAM Hamiltonian is given by:

$$\hat{H} = \hat{H}_{imp} + \sum_{\sigma} \int d\varepsilon \sqrt{\frac{\Gamma(\varepsilon)}{\pi}} (\hat{d}^{\dagger}_{\sigma} \hat{c}_{\varepsilon\sigma} + \text{H.c.}) + \int d\varepsilon \varepsilon \, \hat{c}^{\dagger}_{\varepsilon\sigma} \hat{c}_{\varepsilon\sigma}.$$
 (1)

Here, $\hat{H}_{imp} = \varepsilon_{d\sigma} \hat{n}_{d\sigma} + U \hat{n}_{d\uparrow} \hat{n}_{d\downarrow}$, where $\hat{d}^{\dagger}_{\sigma}$ creates a particle with spin $\sigma = \pm 1$ at the impurity *d*-level at energy $\varepsilon_{d\sigma} = \varepsilon_d - \frac{\sigma}{2} B$ with *B* an external magnetic field representing the Zeeman splitting. Double occupation of the impurity levels pays a Coulombic energy penalty, *U*, as measured by $U \hat{n}_{d\uparrow} \hat{n}_{d\downarrow}$. The coupling of the bath of electrons, $\hat{c}^{(\dagger)}_{\varepsilon\sigma}$, to the impurity is described by the hybridization function, $\Gamma(\varepsilon) = \pi \rho_{\varepsilon} V_{\varepsilon}^2$, with ρ_{ε} the density of states, and V_{ε} the corresponding hopping matrix element at ε . For simplicity in this work we use a featureless hybridization function, $\Gamma(\varepsilon) = \vartheta(D - |\varepsilon|) \Gamma$, with constant strength $\Gamma = 1$ for $|\varepsilon| < D$ and zero elsewhere (we take all energies in units of the half-bandwidth D = 1, unless specified otherwise, as well as $\hbar = k_B = 1$).

III. DATABASE CONSTRUCTION

Inspired by earlier work [24], we started with impurity parameters, U, Γ , and ε_d , which we expanded to include external parameters B and temperature T. For a single data point, these inputs will henceforth be written as the ordered set $x_p \equiv (U, \Gamma, \varepsilon_d, B, T)_p$. We randomly select values for the parameters within predefined physically motivated domains. x_p is mapped onto the scaled target $y_p(\omega) \equiv \pi \Gamma_p A_p(\omega) \in [0,1]$, where $A(\omega)$ is the corresponding spectral function which we compute via the NRG. Two disjoint sets were generated: the "Anderson" (\mathcal{D}^{A}) and "Kondo" (\mathcal{D}^{K}) sets of sizes $|\mathcal{D}^{A}| \approx 600$ k and $|\mathcal{D}^{K}| \approx 410k$. Both sets produce related physics and span similar regions of the 5D input hyperspace (see Fig. 2). As the ML results are similar, we present the Anderson Set results unless otherwise specified. Details on the input parameter generation including histograms and results on the Kondo Set can be found in the Appendices [37].

In practice, spectra were examined between $\omega = \pm 0.8$ to circumvent band edge artifacts, and within that window the spectra are sampled on a refined mixed linearlogarithmic frequency grid ω_i for $i = 1, \ldots, M$ with M = 333. Finally, all regions of the input hyperspace were sufficiently represented according to the trials' smallest physical energy (SPE) scale, defined as $E_0 \equiv \max(T, T_{\rm K}, |B|)$, where $T_{\rm K}$ is the Kondo temperature [37].

Each dataset contains data points (x_p, y_p) , and is partitioned into disjoint training (\mathcal{R}^{Ω}) , cross-validation (\mathcal{V}^{Ω}) ,



FIG. 2: Principal component analysis performed on the Kondo (top) and Anderson (bottom) spectral datasets. In each of the four columns the same PCA data are colored in terms of different parameters, left to right: temperature, Kondo temperature, magnetic field strength, and E_0 .

and testing (\mathcal{T}^{Ω}) sets, with $\Omega \in \{A, K\}$ for the Anderson and Kondo datasets, respectively. We use approximately 2% of the data for cross-validation, another 2% for testing, and the rest for training [38]. In an effort to evaluate how the data selection method of the minimal training sets affects the results, we also define two subsets of \mathcal{R} , both of size 50k. The first is a randomly down-sampled subset $\mathcal{R}_r \subset \mathcal{R}$. The second is a furthest-points downsampled (FPS) subset $\mathcal{R}_f \subset \mathcal{R}$ constructed by first selecting a random point, and then iteratively sampling the next furthest point in the remainder of the scaled, 5D parameter-space. Models are hyper parameter-tuned using the validation set, and all results presented in this work correspond to the testing sets, both of which are consistent regardless of the model or training set.

In order to visually evaluate the integrity of the generated datasets, we performed a principal component analysis (PCA) in the spectral space, reducing the dimension of each spectrum from M to 2. The results are plotted in Fig. 2, and color-coded with respect to parameters of the SIAM that are most directly relevant for the physical low-energy regime, $T, T_{\rm K}, |B|$, and the derived E_0 . Overall, the smooth color gradients observed in the PCA plots suggest that the physically-relevant input parameters can be mapped continuously to prominent spectral features. Additionally, one can identify interesting trends within the physical parameter regimes. In the plots involving E_0 , the lowest values cluster at the top and towards the left, while the largest E_0 scales concentrate at the bottom towards the right/left for the Kondo/Anderson dataset. This suggests that the dynamically generated Kondo peak at low energies and the higher energy side peaks correspond to different spectral features, as already understood from domain knowledge. Despite the overall similarity between the Kondo and Anderson PCA plots, there are subtle differences. For example, in the Kondo set, the input parameters were generated on a grid [37] leading to streaks in the PCA plots while such streaks are

absent in the uniformly sampled Anderson set. These PCA plots qualitatively confirm the physical intuition that a well-defined mapping exists between the input parameters, which determine the physics of the system, and the spectral functions. This suggests that machine learning algorithms are well suited to modeling the featuretarget mapping.

IV. MODEL RESULTS & DISCUSSION

We use the mean absolute error (MAE) to characterize model performance. The MAE between a ground truth NRG spectrum $y_p(\omega)$ and ML-predicted $\hat{y}_p(\omega)$ is defined as an average over the testing set,

$$\delta y_p = \frac{1}{M} \sum_{i=1}^{M} |y_p(\omega_i) - \hat{y}_p(\omega_i)|, \quad \overline{\delta y} = \frac{1}{|\mathcal{T}|} \sum_{y_p \in \mathcal{T}} \delta y_p, \quad (2)$$

as displayed in Table I.

We introduce a measure of the spectral data variation, referred to as the baseline error, as the MAE of the test set against its average spectrum, $\overline{y}(\omega) \equiv$





set, \mathcal{T}^{A} , from the model trained on \mathcal{R}^{A} . Data correspond to the best example of each of the five pentiles of the data top to bottom, respectively. The system parameters $x_{p} \equiv (U, \Gamma, \varepsilon_{d}, B, T)_{p}$ are specified within each panel. As the algorithms have no notion of the ω -grid on which the spectra are defined, we present

the "ML-grid" (left) which represents the spectral functions, $y_p = \pi \Gamma_p A_p$, on a uniformly-spaced grid (emphasizing how the algorithms "see" the targets).

TABLE I: Summary of average and standard deviations of mean absolute error [cf. Eq. (2)] of all trained models as computed on test sets, \mathcal{T}^{Ω} . The training sets used are shown in parenthesis, e.g., MLP($\mathcal{R}_{\rm f}$) indicates the MLP trained on the down-sampled 50k trial $\mathcal{R}_{\rm f}$ sets but still evaluated on the appropriate \mathcal{T} . All KRR($\mathcal{R}_{\rm r}$) training attempts yielded poor predictions with $r^2 < 0$. DC-KRR($\mathcal{R}_{\rm F}$) trained on the full training set \mathcal{R} but in sequentially-sampled subsets ordered by the FPS algorithm as explained in Eq. (3).

Ω	Baseline	$\mathrm{MLP}(\mathcal{R})$	$\mathrm{MLP}(\mathcal{R}_r)$	$\mathrm{MLP}(\mathcal{R}_{\mathrm{f}})$	$\mathrm{KRR}(\mathcal{R}_{\mathrm{r}})$	$\mathrm{KRR}(\mathcal{R}_{\mathrm{f}})$	$\mathrm{DC}\text{-}\mathrm{KRR}(\mathcal{R}_\mathrm{F})$
А	0.126 ± 0.054	0.003 ± 0.003	0.091 ± 0.070	0.014 ± 0.013	0.084 ± 0.068	0.021 ± 0.023	0.017 ± 0.026
Κ	0.234 ± 0.076	0.003 ± 0.002	0.152 ± 0.141	0.010 ± 0.009	0.152 ± 0.145	0.019 ± 0.021	0.035 ± 0.033

 $\frac{1}{|\mathcal{T}|} \sum_{y_p \in \mathcal{T}} y_p(\omega).$ In a dataset with a high degree of variance, such as \mathcal{R}^A or \mathcal{R}^K , we expect the baseline error to be quite large; our goal is to train a machine learning model to learn the mapping between the input and output effectively, and thus significantly outperform this baseline. We begin our investigation with a multi-layer perceptron (MLP), a deep learning model capable of capturing highly non-linear relations in high-dimensional data. We observe a superb performance from the MLP trained on \mathcal{R} , where the results outperform the baseline by factors of roughly 40 and 80 on the Anderson $(\overline{\delta y} = 0.003)$ and Kondo $(\overline{\delta y} = 0.003)$ testing sets, respectively.

From each pentile of the Anderson test set, the best representative examples of the model predictions are presented in Fig. 3. We first note that all important spectral features are well-reproduced in these examples including the peak heights, widths, and locations of the sharp central peak and the broader side peaks. As expected, models trained using all of \mathcal{R} present the best results. However, models trained using the FPS subset, \mathcal{R}_{f} , which only constitutes about 10% of each of the full training sets, also perform surprisingly well, indicating that even a moderately sized training set can result in accurate predictions, if the sampling of the input-parameter space is well-spanned. This is critically important for more complex physical problems where generating training data becomes much more expensive. By comparison, models trained on a randomly sampled subset, \mathcal{R}_r , of the same size performs roughly an order of magnitude worse than \mathcal{R}_{f} , and even barely outperforms the baseline.

We also examine analytical results from the KRR model [39, 40]. However, this method scales cubically with training set size $|\mathcal{R}_x|$ with $x \in \{r, f\}$, as it requires a full matrix inversion, making it intractable to use the full \mathcal{R} at once. We mitigate this problem in two ways: using down-sampled training sets of \mathcal{R}_r or \mathcal{R}_f , and a divideand-conquer KRR (DC-KRR) algorithm [40, 41]. Details regarding KRR and DC-KRR algorithms and hyperparameters can be found in [37]. Models trained using \mathcal{R}_f perform an order of magnitude better, and have higher coefficients of determination (r^2) than the \mathcal{R}_r counterparts, in agreement with our earlier findings for the MLP.

The final KRR models trained on \mathcal{R}_{f} are superior to the baseline error by factors of roughly 6 and 10 for the

Anderson and Kondo sets, respectively. Interestingly, both the MLP and KRR models trained with the down-sampled \mathcal{R}_r achieved nearly identical results, and both are essentially indistinguishable from the baselines, indicating there is insufficient information contained in the randomly-downsampled training sets to train successful models.

In contrast, DC-KRR uses the full data set by partitioning \mathcal{R} into S consecutive (ordered), disjoint subsets $\mathcal{R}_{f}^{(s)}$, such that

$$\mathcal{R}_{\rm F} = \left(\mathcal{R}_{\rm f}^{(1)}, \mathcal{R}_{\rm f}^{(2)}, ..., \mathcal{R}_{\rm f}^{(S)}\right), \quad |\mathcal{R}_{\rm f}^{(s)}| \approx 50 \rm k, \quad (3)$$

where the union $\bigcup_{s} \mathcal{R}_{f}^{(s)}$ is equivalent to the full training set \mathcal{R} with respect to the earlier analysis, and it holds that $\mathcal{R}_{f} \equiv \mathcal{R}_{f}^{(1)}$. We then train an independent KRR model for each $\mathcal{R}_{f}^{(s)}$ and average the resulting learned parameters. Because the KRR models performed poorly with \mathcal{R}_{r} we only present the DC-KRR whose subsets were indexed according to the FPS algorithm in Table I. Here, both models exceed the baseline average by a factor of roughly 7. Despite training on the full data set, the DC-KRR performs comparably to the KRR(\mathcal{R}_{f}) model.

V. CONCLUSION

In summary we have shown that ML algorithms can predict state-of-the-art Anderson impurity model spectra to overall quantitative accuracy at a speedup of 10^5 over NRG. We have found that the use of furthest points sampling can significantly reduce the required amount of training data to achieve satisfactory accuracy. While KRR and DC-KRR are effective for small datasets or large datasets properly divided into chunks of small datasets, our results imply that deep learning algorithms trained on sufficient amount of data are superior for predicting both the single and many-body features of a SIAM spectral function. Future work will expand the physical model to include additional impurity parameters, most importantly a structured hybridization function and more channels, and thus examine the viability of a ML algorithm in the context of a DMFT self-consistent loop.

ACKNOWLEDGMENTS

EJS and RK were supported by the U.S Department of Energy, Office of Science, Basic Energy Sciences as a part of the Computational Materials Science Program. MRC acknowledges support from the US Department of Energy through the Computational Science Graduate Fellowship (DOE CSGF) under Grant No. DE-FG02-97ER25308. AW was supported by the U.S. Department of Energy, Office of Basic Energy Sciences. This research used resources of the Center for Functional Nanomaterials, which is a U.S. DOE Office of Science Facility, and the Scientific Data and Computing Center, a component of the Computational Science Initiative, at Brookhaven National Laboratory under Contract No. DE-SC0012704. We are grateful to Cole Miles, Kipton Barros, and Laura Classen for thoughtful discussions.

Appendix A: Dataset construction

We use the Numerical Renormalization Group (NRG, [21, 29]) to compute the impurity spectral functions. The NRG obtains the spectral data directly on the real-frequency axis for arbitrary temperatures where we use the fdm-NRG approach [31]. We use a typical discretization parameter of $\Lambda = 2$, with the discrete data subsequently smoothened using standard log-Gaussian broadening schemes. We also compute the local self-energy $\Sigma^{imp}(\omega)$ to improve spectral resolution of the NRG data [42].

Each NRG trial takes a set of randomly sampled system parameters $x_p \equiv (U, \Gamma, \varepsilon_d, B, T)_p$, and generates a spectral function $\pi \Gamma A(\omega)$. Here the factor of $\pi \Gamma$ maintains normalization, such that, e.g., $\pi \Gamma A(0) \leq 1$ can be interpreted as transmission probability in transport measurements. [12, 43, 44] The hybridization strength Γ is derived from the hybridization function chosen to be featureless, i.e. $\Gamma(\omega) = \Gamma \vartheta(D - |\omega|)$, [21, 29] where the half-bandwidth D := 1 sets the unit of energy throughout, unless specified otherwise.

Due to the artificial sharp cutoff at the band edge caused by our choice of constant hybridization strength. we only consider the energy window $\omega = [-0.8, 0.8] D$ to avoid artefacts at the band edge. We do not examine regions, and hence also ignore spectral weight beyond the band edge. Within the specified window, we coarsegrain the essentially continuous NRG data to the same fixed frequency grid to be used in the machine learning (ML) algorithms. This grid is chosen such that it is linearly spaced for larger frequencies (67 points for $\omega \in [0.1, 0.8]$), and logarithmically spaced for smaller frequencies (99 points for $\omega \in [10^{-5}, 10^{-1}]$). The same grid is mirrored for negative frequencies. With the addition of a single point at $\omega = 0$ to bridge the logarithmic grid from positive to negative we have our final 333 coarsegrain frequency points. The grid ω_i is the same for every trial, which thus maps the five physical parameters to the

normalized spectral function,

$$x_p \mapsto y_{p,i} \equiv \pi \Gamma_p A_p(\omega_i).$$
 (A1)

The five physical parameters x_p are referred to as the *input features*, and the 333 y_p values as the *output targets* in our ML algorithms.

The smallest non-zero ω value in our logarithmic frequency grid is 10^{-5} , which is well above the NRG's minimum set at $\sim 10^{-6}$ as determined by our Wilson chain length of L = 50 and discretization parameter $\Lambda = 2$. Therefore, to ensure that features are captured within our frequency grid for ML, any NRG trials with smallest physical energy (SPE) E_0 scales smaller than 10^{-4} are removed. The SPE scale of each trial is defined by:

$$E_0 = \max\left(|B|, T, T_K\right) , \qquad (A2)$$

with the Kondo temperature T_K determined according to Haldane's formula [45],

$$T_K = \min\left(0.575, \sqrt{\frac{U\Gamma}{2}}\right) \exp\left(\frac{\pi\varepsilon_d(\varepsilon_d+U)}{2U\Gamma}\right)$$
 (A3)

By definition of being the smallest physical energy scale, all physical features in the spectral data are at least as broad as the SPE. For example, low-energy Kondo features (described by T_K which is defined at B = T = 0) are physically smeared out at the energy scale of |B| or T if these are larger than T_K .

1. Physical Parameter Selection

The five physical parameters for the single impurity Anderson model (SIAM) in this work are collected as an ordered set $x_p \equiv (U, \Gamma, \varepsilon_d, B, T)_p$. Inspired by the earlier work of Arsenault *et. al.* [24], we began with all five parameters sampled on pre-selected grids (referred to as 'on-grid data' below). This initial set contained just over 1.45 million trials (see blue bars in Fig. 4). However, we eventually found it favorable to supplement these with 329k additional trials with random parameter values from preset ranges ('off-grid') provided that their E_0 values fell within the targeted range of $[10^{-4}, 10^{-1}]$ (see green bars in Fig. 4). Ultimately we removed any trials where the E_0 did not fall in the desired energy range (see right panel in Fig. 4).

We then enforced the additional requirement on the hybridization strength Γ ,

$$\Gamma > \min\left(\frac{|B|}{5}, \frac{U}{40}\right) \tag{A4}$$

in order to avoid extremely narrow features in the spectral data comparable to or below the frequency grid spacing chosen for ML. From a physical perspective, this would correspond to an essentially decoupled and hence trivial impurity. Finally, a randomly selected subset of trials that satisfied the above requirements were selected for a data set of ~ 411 k trials. We refer to this dataset



FIG. 4: SPE histograms for the Kondo data set. The original *on-grid* data (blue) of 1.45 million trials was supplemented by an additional 329k *off-grid* trials (green) to ensure a more even E_0 distribution on a log-scale (left panel). Trials with SPE values outside our range of interest $E_0 \in [10^{-4}, 10^{-1}]$ were discarded, which gave rise to a final SPE distribution (right panel). Bin widths are the same in both panels.

as the "Kondo Set" because many parameter sets have large Coulomb energies U comparable to or larger than the half bandwidth (i.e. $U \gtrsim 1$). The SPE distribution for the Kondo set is presented in Fig. 5, while the parameter value distribution is shown in black in Fig. 8.

In contrast, the "Anderson Set" was designed with an even SPE distribution in mind, while also keeping U < 0.5 right from the start as shown in black in the leftmost panels of Fig. 7. To be specific, the parameters for this data set were sampled in the following manner: (i) randomly choose an SPE value with a flat distribution on the log-scale in the range $E_0 \in [10^{-4}, 10^{-1}]$ (ii) randomly select which of the three parameters |B|, T, or T_K takes that E_0 value, then (iii) ensure that the other two parameters are smaller than E_0 , using the range $[0, E_0/10]$. Both positive and negative values for B are computed in all cases. With T_K fixed, (iv) the remaining values for U, ε_d , and Γ are sampled based on Eq. (A3). Specifically, the values for Γ and U are assigned at random on a linear scale below bandwidth, while still ensuring Eq. (A4). This fixes the prefactor in Eq. (A3). By taking the logarithm, we then solve for ε_d . If no valid solutions exist for ε_d given the choices in U and Γ , we go back and repeat step (iv). With this basic algorithm we generated a total of about 600k trials with the desired flat SPE distribution as demonstrated in Fig. 6. The parameter distribution for the Anderson Set is shown in black in Fig. 7.

2. Split selection

The full datasets for the Anderson and Kondo datasets are labeled \mathcal{D}^{Ω} with $\Omega \in \{A, K\}$. Each contain about 500k of data points (x_p, y_p) that are split into disjoint training (\mathcal{R}^{Ω}) , validation (\mathcal{V}^{Ω}) , and testing (\mathcal{T}^{Ω}) sets. Since the Anderson and Kondo sets themselves are disjoint and also mostly dealt with on an equal footing, we generally suppress the subscript Ω for readability unless stated otherwise. The two pairs of validation and testing sets are fixed for the entirety of this work. The splits were generated as follows:

- 1. The testing set, \mathcal{T}^{Ω} , contains roughly 2% of the total data and is selected by randomly down-sampling \mathcal{D}^{Ω} . It is evaluated at the end of the pipeline, and represents the most unbiased representation of the model performance. All results are evaluated on the testing set, unless explicitly stated otherwise.
- 2. The validation set, \mathcal{V}^{Ω} , also contains roughly 2% of the total data, and is also selected by randomly down-sampling \mathcal{D}^{Ω} . It is used only to tune model hyperparameters.
- 3. The full training set, \mathcal{R} (with Ω indexing suppressed), contains 96% of the total data and is used to train the model. The different versions of the training set are explained below.
 - The randomly sampled training set \mathcal{R}_r contains 50k data points. It is selected by randomly down-sampling \mathcal{R} .
 - The furthest points-sampled (FPS) training set $\mathcal{R}_{\rm f}$ contains 50k data points. It is selected via the algorithm presented in Section A1.4.
 - For divide-and-conquer DC-KRR only, we generate a collection of *ordered* disjoint subsets $\mathcal{R}_{\rm F} \equiv \left(\mathcal{R}_{\rm f}^{(1)}, \mathcal{R}_{\rm f}^{(2)}, \dots, \mathcal{R}_{\rm f}^{(S)}\right)$, where the



FIG. 5: SPE histogram of 411,267 data point Kondo set on linear and log-scale (left and right panel respectively), with approximately equal distribution across trials at lower energies only where E_0 is dominated by |B|, T, or T_K . Regularly spaced spikes in right panel are due to the *on-grid* portion of the Kondo set.



FIG. 6: SPE histogram of 599,578 data point Anderson set on linear and log-scale (left and right panel respectively), with approximately equal distribution across trials where E_0 is dominated by |B|, T, or T_K .

set itself $\mathcal{R} = \bigcup_{s=1}^{S} \mathcal{R}_{f}^{(s)}$ (i.e., the unordered union of the subsets are equivalent to the full training set). The sequence is ordered in the sense that the $\mathcal{R}_{f}^{(s)}$ are generated consecutively using the algorithm presented in Section A1.4. As an example, the sampling of $\mathcal{R}_{f}^{(s)}$ excludes samples already included in $\{\mathcal{R}_{f}^{(1)}, \ldots, \mathcal{R}_{f}^{(s-1)}\}$, and the first point in $\mathcal{R}_{f}^{(s)}$ is furthest-sampled from the last point in $\mathcal{R}_{f}^{(s-1)}$. Note also that $\mathcal{R}_{f}^{(1)} = \mathcal{R}_{f}$.

3. Preprocessing input features (symlog scaling)

Anderson and Kondo-type models frequently exhibit dynamically generated, exponentially small energy scales. Hence the parameters that enter the SPE in Eq. A2, namely $T_{\rm K}$, B, and T, were sampled on a logarithmic scale that stretches over several orders of magnitude. Therefore, prior to model training, one is naturally led to also apply a (symmetric) logarithmic ("symlog") rescaling to the input features x_p . Feature scaling was necessary for the MLP and the analytical methods when $\mathcal{R}_{\rm R}$ was used, as described elsewhere in this document.

Given that the value for $T_{\rm K}$ itself is not a bare Hamiltonian parameter, we apply the following logarithmic scal-

ing to Γ , B, and T:

$$\begin{split} & \Gamma \to \log_{10} \ \Gamma \\ & B \to \operatorname{symlog}_{10}(B) \equiv \operatorname{sgn}(B) \log_{10} |B| \qquad \text{(A5)} \\ & T \to \log_{10} \ T \ , \end{split}$$

where U and ε_d are left on the linear scale. Here $\Gamma, T > 0$ are always chosen non-zero, yet possibly exponentially small. Also, we always have |B| < 1, and in the case that B = 0, the trial's B value is reset to an order of magnitude smaller than the smallest non-negative B value in the set.

After performing the scaling step in Eq. A6, the mean (μ) and standard deviation (σ) are computed for each of the 5 components $x_p^{(i)} \in x_p$ in the current training set. Then we use $x_p^{(i)} \rightarrow \frac{x_p^{(i)} - \mu^{(i)}}{\sigma^{(i)}}$ for every trial from the training, validation, and testing sets.

4. Furthest points sampling algorithm

Here we employ a relatively simple algorithm for sampling data points in an arbitrary dimensional space. Note that sampling is performed on scaled features as to treat their notion of distance on equal footing. The algorithm is defined below:

- 1. Select a random point $x_p \in \mathcal{R}$ (here, x_p represents the 5-dimensional input parameter vector, i.e., the feature space). It is labeled the *current point*, and added as first point to the set of sampled points $\mathcal{R}_{\rm F}$.
- 2. Find the point that is furthest away from the current point that has not yet been selected, and add it to the set of sampled points. We define *distance* by $|| \cdot ||_{L_d}$ with d = 2.
- 3. Repeat step 2 until the desired size of $\mathcal{R}_{\rm F}$ is reached.

5. Parameter Distributions

In this section we present the input feature distributions. In Fig. 7 we show the full Anderson set with its two down-sampled subsets which were plotted with a degree of translucency to exhibit the concentration of the down-sampled trials. One can immediately observe that the furthest point sampled data is more diffuse whereas the random points tend to concentrate in several regions in each distribution. In Fig. 8 we present the same data for the full Kondo set. Here, we can clearly see the striations due to the 'on-grid' trials as opposed to the clusters of "off-grid" trials explained in Section A1.

Appendix B: Machine learning models

In this section, we provide a brief overview of some machine learning theory, along with details of the specific model architectures and hyperparameters chosen in this work. For a detailed overview of practical machine learning technique in the domains of physics and chemistry, we refer the interested reader to the overview by Wang *et al.* [46] and for details on feed-forward neural networks, to the overview by Cheng and Titterington [47]. For details regarding the analytical models, see texts by Mohri *et al.* [48] and Hastie *et al.* [49]. For the particular applications of these methods explored in this work, we refer our readers to Zhang *et al.* [40] and You *et al.* [41].

1. Neural networks

In this work we use a fully-connected feed-forward neural network known as a multi-layer perceptron (MLP) as the sole deep learning model. An MLP consists of sequential layers of nodes, where each node takes as input the output from all previous layers, and outputs a single value, known as an activation. The activations of the *l*th layer are computed by the general matrix equation $\mathbf{a}^{l} = f(W^{l}\mathbf{a}^{l-1})$, where \mathbf{a}^{l} is the vector of activations from the *l*th layer. W^l is the *l*th weight matrix mapping the output from the (l-1)th layer to the input of the *l*th layer. The function f is called the activation function. It applies a differentiable, element-wise non-linearity to the output, allowing the weights to learn highly-nonlinear representations of the input, parameterized by the neural network weights. These weights are learned during training, in which a numerical optimization procedure tries to find the minimum distance between ground truth training data points and the predictions. The specific model hyperparameters, or non-learned parameters of the model, are discussed in Section A2.2.

2. Neural network training

Here we present the training protocols we used to train and hyperparameter-tune the neural networks used in this work. We find the optimal set of hyperparameters by using a combination of grid search and hand-tuning. The best model is determined by training models with different hyperparameter combinations on the same training set, evaluating on the validation set \mathcal{V} , and selecting the model with the lowest MAE.[50] That model, finally, is evaluated on the testing set \mathcal{T} which yields the presented data, unless indicated otherwise.

3. Common neural network training parameters

All models share the following hyperparameters:



FIG. 7: The input feature distribution for the full Anderson set is presented in black. Elements from the down-sampled subsets of the furthest point and random point algorithms each representing approximately 10% of the full set are superimposed in color at 10% opacity to exhibit trial density. The top row shows the furthest point distribution in green while the bottom row shows the random points in red. See Sections A1.3 and A1.4 for additional information regarding the symmetric logarithm (symlog₁₀) procedure and furthest-point sampling algorithm respectively.



FIG. 8: The input feature distribution for the full Kondo set (black). The down-sampled subsets each representing about 10% of the full data set are superimposed in color at 10% opacity to demonstrate density. The top row shows the furthest point algorithm's subset in green, while the bottom row has the random point subset in red. See Sections A1.3 and A1.4 for additional information regarding the symmetric logarithm (symlog₁₀) procedure and furthest-point sampling algorithm respectively. The horizontal line at symlog₁₀(B) = -8 in the right-most panels is caused by setting our $B \equiv 0 \rightarrow 10^{-8}$ as explained in Section A1.3.

- $|| \cdot ||_{L_1}$ loss, the mean absolute error (MAE).
- Adam optimizer [51] with a starting learning rate of 10^{-3} .
- ReLU activation function.
- A scheduler which decreases the learning rate when the validation loss plateaus. This scheduler has a patience of 10 epochs, decrease factor of 0.5, and minimum learning rate of 10⁻⁷.
- Training batch size of $16\,384$ (2^{14}) .
- For simplicity, we constrain all hidden layers to be of the same size.
- Total of 5000 epochs.

4. Best neural network hyperparameters

As evaluated on \mathcal{V} , we present the best hyperparameter combinations in Table II. Note that hyperparameter tuning is a highly non-convex problem, and it is possible that better combinations exist.

5. Kernel ridge regression

Regression methods are trained to find the line of best fit. However, not all problems lend themselves well to a linear-fit decision boundary. In those cases it may be easier to compute the parameters' higher dimensional dual space by way of a kernel where a linear fit may now be possible [48]. This approach is known as kernel ridge regression (KRR) [39, 41].

For a given trial (x_p, y_p) of N total trials, the generic form of a KRR minimization algorithm reads:

$$\min_{\alpha} \frac{1}{N} \left(\sum_{p=1}^{N} \left(y_p - f_p \right)^2 + \lambda \| f \|_{\mathcal{H}}^2 \right)$$
(B1a)

$$f_p = \sum_{p'=1}^{N} \alpha_{p'} k(x_{p'}, x_p)$$
(B1b)

where the first term in Eq. B1a is the usual linear regression mean squared error (MSE) cost function between the model's kernel-based prediction f_p and ground truth value y_p [40]. While several kernels exist, in this work we exclusively use the Laplacian kernel [53] given by $K_{pp'} \equiv k(x_p, x_{p'}) = \exp\left(-\frac{1}{\sigma}||x_p - x_{p'}||_1\right)$ with input feature vectors x_p and $x_{p'}$ (5-dimensional in this work). The exponential argument is the L1 norm (Manhattan Distance) divided by the kernel radius $\sigma > 0$ which determines how similar x_p is to $x_{p'}$. The target vectors y_p are of size M = 333 in the present case.

The second term in Eq. (B1a) is the regularization term which helps prevent over-fitting. It includes two factors: the strength $\lambda \geq 0$, and the Hilbert space norm $||f||_{\mathcal{H}}$ defined as $||f||_{\mathcal{H}} := \langle f, f \rangle_{\mathcal{H}}^{1/2}$. [40] Both the kernel radius σ and regularization strength λ are tunable hyperparameters. Table III shows the selected hand-tuned hyperparameters. In this work we find that scaling the data by a symmetric logarithm procedure for B, T, and Γ only improved the analytical algorithm results when the model is fit to \mathcal{R}_r . Hence, no such scaling is applied when training with \mathcal{R}_f . Throughout, we apply the standard normalization of the data, such that mean and standard deviations of all training data are 0 and 1, respectively [cf. Section A1.3].

Solving Eq. (B1) for the weight matrix α requires an expensive inversion of the kernel matrix, $k(x_p, x_{p'})$ which scales as $O(N^3)$ in time and $O(N^2)$ in memory for N data points (x_p, y_p) . This expense can be mitigated in several ways, including the divide-and-conquer approach, detailed below. In this work all KRR trials were modelled on 50k training trials (see Section A3).

The r^2 score (R-squared) value determines the quality of a trained regression model. A perfect model would achieve $r^2 = 1$, so the closer a model is to 1, the better the fit. A score of $r^2 = 0$ indicates that a constant model predicts the same result regardless of input. A negative score indicates that the model is arbitrarily worse than a constant representing the mean value. In this work we found that it was possible to generate acceptable models for the furthest-point sampled training data $\mathcal{R}_{\rm f}$, but the models fit with random point sampled $\mathcal{R}_{\rm r}$ training set were so poor that they had negative r^2 values. Therefore, if training set size is a limiting factor for a brute-force affordable KRR model (here inverting a full matrix of dimension 50k), then one must chose the training set with great care.

6. Divide-and-conquer kernel ridge regression

Divide-and-conquer kernel ridge regression (DC-KRR) is one of several approaches one may take to mitigate the poor scaling of a traditional KRR algorithm explained above [40]. In DC-KRR the total number $N \approx 500k$ of training trials is subdivided into $S \sim 10$ subsets with an equal number $n \approx 50k$ training trials in each (with minor variations for the Kondo as compared to the Anderson set). Then, a separate KRR model is fitted for each, and the trained weights (α in Eq. (B1)) are saved. After training, a prediction may be acquired by computing (for each subset) the product of the weights and the kernel of the queried value and subset's training trials. The final results are averaged for the final prediction as shown in Eq. B2.

Given input parameters x_P corresponding to ground truth target y_P , one may obtain a DC-KRR prediction \hat{y}_P via:

$$\hat{y}_P = \frac{1}{S} \sum_{s=1}^{S} \ker(\mathcal{R}_f^s, x_P) \alpha_s \tag{B2}$$

TABLE II: The best hyperparameters for the neural networks for Anderson (upper rows) and Kondo (lower rows) as found by a combination of hand tuning and grid search. Training and validation losses are displayed in Fig. 10a and Fig. 11a for the Anderson and Kondo sets respectively. Testing set error distributions are shown in Fig. 10b and Fig. 11b for the Anderson and Kondo sets respectively. In the rightmost column (Examples), we provide the figure number where randomly drawn samples from each of the respective testing sets may be viewed. Both NRG ground truth (black) and predicted (red) spectra are displayed. Summaries of where the example results are shown are

Training set	hidden layer size	Number of hidden layers	Dropout [52]	Examples
\mathcal{R}^{A}	256	8	0	12a
$\mathcal{R}_{\mathrm{r}}^{\mathrm{A}}$	256	4	0.05	12b
$\mathcal{R}_{\mathrm{f}}^{\mathrm{A}}$	256	8	0	12c
\mathcal{R}^{K}	256	8	0	13a
$\mathcal{R}_{\mathrm{r}}^{\mathrm{K}}$	256	4	0.05	13b
$\mathcal{R}_{\mathrm{f}}^{\mathrm{K}}$	256	8	0	13c

given in Tables IV and V.

where α_s is a matrix of the learned algorithm weights described in Eqn. B1 and ker(\mathcal{R}_f^s, x_P) indicates the kernel between all $x_p \in \mathcal{R}_f^s$ and the queried x_P input. DC-KRR has hyperparameters λ and σ which function in the same manner as described in Section 2.5 and we use the Laplacian kernel here as well. Final hyperparameter values are given in Table III, which are applied universally for all subsets. Fig. 9 shows the individual r^2 values for each subset in the DC-KRR algorithm, whereas the final averaged r^2 value is reported in Table III.

This approach reduces the cost of training with all data points in \mathcal{R} and making the approach easily parallelizable. However, the overall scaling has not changed in this implementation; DC-KRR simply permits the use of the full training set. Additionally, it should be noted that the composition of the individual subsets matters. As DC-KRR works by training a multitude of standard KRR models, the same pitfalls that can negatively affect a single KRR model can also detract from the individual DC-KRR subset models which are then averaged, thus potentially compounding the error. In this work we chose to apply the furthest- point ordering on the training data prior to splitting \mathcal{R} set into equal sized subsets as the KRR trials performed better with this pre-processing step.

7. Best KRR and DC-KRR hyperparameters

In Fig. 9 we present the r^2 values for each subset s for the validation \mathcal{V} and testing \mathcal{T} sets in DC-KRR. In all cases there is excellent agreement between a given subset's r^2 score for the validation and test sets. By the very construction of the underlying FPS, the last group (s = S) collects the rest of the data set, hence is no longer strictly FPS. This introduces a certain bias towards late patches, as seen in Fig. 9. The values for the Anderson set are tightly distributed. In contrast, the Kondo sets exhibit significant deterioration towards the last subsets s = 7 and 8, negatively impacting the overall average. Despite this deterioration, however, the final r^2 values



FIG. 9: DC-KRR r^2 values for each subset within the Anderson (squares) and Kondo (circles) $\mathcal{R}_{\rm F}$ sets. Closed markers indicate the validation r^2 scores, and open markers indicate that of the test sets. Final DC-KRR values are the averages across all relevant subsets.

remain quite high for both datasets, indicating reliable models.

Appendix C: Results

In this section, we present figures demonstrating various representations of each model's performance. These figures are connected to the different models in Table IV for the Anderson results, and Table V for the Kondo results. For clarity, we highlight the differences between each dataset; note that all results that follow were hyperparameter-tuned on the same \mathcal{V} and all results that follow in this section correspond to the testing set (\mathcal{T}) results as evaluated on the best model, which is determined by the MAE on \mathcal{V} . TABLE III: The hand-tuned hyperparameters σ and λ [cf. Eqs. (B1)], together with the r^2 scores of the validation (\mathcal{V}) and testing sets (\mathcal{T}) used in this work for kernel ridge regression (KRR) and divide-and-conquer KRR (DC-KRR). While several values were examined for each of these hyperparameters for each set of training data, only the final choices are presented. We also differentiate furthest-point sampled $\mathcal{R}_{\rm f}$ from random point sampled $\mathcal{R}_{\rm r}$ training data. For the training of the DC-KRR models, the full training set \mathcal{R} was used, yet partitioned into S disjoint subsets $\{\mathcal{R}_f^s\}$ as described in Eqn. 3 in the main text. Values for the DC-KRR r^2 entries are the result of averaging the r^2 values from the subsets shown in Fig. 9. In the rightmost column (Examples), we provide the figure number where randomly drawn samples from each of the respective testing sets may be viewed. Both NRG ground truth (black) and predicted (red) spectra are displayed. Summaries of where the example results are shown are given in Tables IV and V. Only the best models have associated result examples in Appendix A3. Figs. 14a and 14b demonstrate regression approach error distributions for the Anderson and Kondo sets respectively.

		1	Anderson						Kondo				
KRR	Scaling	σ	λ	$r_{\mathcal{V}}^2$	r_T^2	$MAE(\mathcal{T})$	Examples	σ	λ	$r_{\mathcal{V}}^2$	r_T^2	$MAE(\mathcal{T})$	Examples
$\mathcal{R}_{ m r}$	None	100	0	-0.267	-0.255	0.102 ± 0.096	N/A	100	0	-0.192	-0.204	0.182 ± 0.187	N/A
$\mathcal{R}_{ m r}$	Symlog	$6.\overline{6}$	0.1	0.277	0.282	0.085 ± 0.068	15a	10	0	0.179	0.188	0.152 ± 0.145	16a
\mathcal{R}_{f}	None	1	0.01	0.929	0.932	0.021 ± 0.023	15b	1	0.01	0.979	0.979	0.019 ± 0.021	16b
\mathcal{R}_{f}	Symlog	1	0	0.909	0.915	0.024 ± 0.026	N/A	1	0.01	0.972	0.972	0.023 ± 0.024	N/A
DC-KRR													
\mathcal{R}_{F}	None	1	0	0.950	0.950	0.0175 ± 0.026	15c	1	0	0.902	0.899	0.035 ± 0.033	16c

- The full training set refers to training on all available training data, \mathcal{R} .
- The FPS training set refers to training on only the 50k furthest-points sampled data, $\mathcal{R}_{f} = \mathcal{R}_{f}^{(1)}$.
- The dataset \mathcal{R}_r refers to training on only the 50k randomly down-sampled data,

For all datasets, we present the following results:

- Random 25: a randomly-selected 25 samples plotted on the ML grid.
- Training info: loss and learning rate plots when applicable.
- Distribution: error distributions.

	$\mathrm{MLP}(\mathcal{R})$	$\mathrm{MLP}(\mathcal{R}_{\mathrm{r}})$	$\mathrm{MLP}(\mathcal{R}_{\mathrm{f}})$	$\mathrm{KRR}(\mathcal{R}_{\mathrm{r}})$	$\mathrm{KRR}(\mathcal{R}_\mathrm{f})$	$DC-KRR(\mathcal{R}_F)$
Training info	10a	10a	10a	_	_	_
Distribution	$10\mathrm{b}$	$10\mathrm{b}$	$10\mathrm{b}$	14a	14a	14a
Random 25	12a	12b	12c	15a	$15\mathrm{b}$	15c

TABLE IV: A quick-reference for the Anderson results presented in this section. The table entries are figure labels.See text for description for the three rows.

TABLE V: A quick-reference for the Kondo results presented in this section. The table entries are figure labels. Seetext for description for the three rows.

	$\mathrm{MLP}(\mathcal{R})$	$\mathrm{MLP}(\mathcal{R}_{\mathrm{r}})$	$\mathrm{MLP}(\mathcal{R}_\mathrm{f})$	$\mathrm{KRR}(\mathcal{R}_{\mathrm{r}})$	$\mathrm{KRR}(\mathcal{R}_\mathrm{f})$	$DC\text{-}KRR(\mathcal{R}_F)$
Training info	11a	11a	11a	-	_	-
Distribution	11b	11b	11b	14b	14b	14b
Random 25	13a	13b	13c	16a	$16\mathrm{b}$	16c



FIG. 10: (a) MLP training and validation losses and learning rates plotted as a function of epochs for the Anderson set. (b) Histogram for MLP errors as measured by the average MAE values, $\overline{\delta y}$. Black represents the errors when trained with the full set \mathcal{R}^{A} ; red represents the errors when trained with the furthest point sampled subset \mathcal{R}^{A}_{f} ; blue represents the errors when trained with the random-sampled subset \mathcal{R}^{A}_{r} .



FIG. 11: (a) MLP training and validation losses and learning rates plotted as a function of epochs for the Kondo set. (b) Histogram for MLP errors as measured by the average MAE values, $\overline{\delta y}$. Black represents the errors when trained on the full Kondo set \mathcal{R}^{K} ; red represents the errors with the model trained with the furthest-point sampled subset \mathcal{R}^{K}_{f} ; blue represents the model errors when trained with the random-sampled subset \mathcal{R}^{K}_{r} .



FIG. 12: Random samples from the Anderson testing set from the MLP models trained on different training sets: (a) the full Anderson set; (b) the randomized subset; (c) the furthest-point sampling subset.



FIG. 13: Random samples from the Kondo testing set from the MLP models trained on different training sets: (a) the full Kondo set; (b) the randomized subset; (c) the furthest-point sampling subset.



FIG. 14: (a) Anderson and (b) Kondo testing set histograms for the the KRR models trained with the random \mathcal{R}_r and furthest-point ordered \mathcal{R}_f down-sampled subsets in red and blue respectively. The DC-KRR model trained with the chunked furthest-point ordered set \mathcal{R}_F is shown in black.



FIG. 15: Random samples from the Anderson testing set from the KRR models for (a) the random subset $\mathcal{R}_{\rm r}$, (b) the furthest point subset $\mathcal{R}_{\rm f}$, and (c) the DC-KRR model with all trials ordered according to the furthest point algorithm $\mathcal{R}_{\rm F}$.



FIG. 16: Random samples from the Kondo testing set from the KRR models for (a) the random subset \mathcal{R}_r , (b) the furthest point subset \mathcal{R}_f , and (c) the DC-KRR model with all trials ordered according to the furthest point algorithm \mathcal{R}_F .

- P. A. Lee, N. Nagaosa, and X.-G. Wen, Rev. Mod. Phys. 78, 17 (2006).
- [2] E. Fradkin, S. A. Kivelson, and J. M. Tranquada, Rev. Mod. Phys. 87, 457 (2015).
- [3] A. P. Ramirez, J. Phys. Condens. Matter 9, 8171 (1997).
- [4] Q. Si and F. Steglich, Science **329**, 1161 (2010).
- [5] C. C. Homes, Q. Du, C. Petrovic, W. H. Brito, S. Choi, and G. Kotliar, Sci. Rep. 8 (2018).
- [6] A. Chikina, J.-Z. Ma, W. H. Brito, S. Choi, P. Sémon, A. Kutepov, Q. Du, J. Jandke, H. Liu, N. C. Plumb, M. Shi, C. Petrovic, M. Radovic, and G. Kotliar, Phys. Rev. Research 2, 023190 (2020).
- [7] A. McMahan, C. Huscroft, R. Scalettar, and E. Pollock, J. Comput. Aided Mol. Des. 5, 131 (1998).
- [8] W. Metzner and D. Vollhardt, Phys. Rev. Lett. 62, 324 (1989).
- [9] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. 68, 13 (1996).
- [10] P. R. C. Kent and G. Kotliar, Science **361**, 348 (2018).
- [11] W. J. de Haas, J. de Boer, and G. J. van d
 en Berg, Physica 1, 1115 (1934).
- [12] T. A. Costi, L. Bergqvist, A. Weichselbaum, J. von Delft, T. Micklitz, A. Rosch, P. Mavropoulos, P. H. Dederichs, F. Mallet, L. Saminadayar, and C. Bäuerle, Phys. Rev. Lett. **102**, 056802 (2009).
- [13] A. J. Leggett, S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg, and W. Zwerger, Rev. Mod. Phys. 59, 1 (1987).
- [14] S. M. Cronenwett, T. H. Oosterkamp, and L. P. Kouwenhoven, Science 281, 540 (1998).
- [15] C. Latta, F. Haupt, M. Hanl, A. Weichselbaum, M. Claassen, W. Wuester, P. Fallahi, S. Faelt, L. Glazman, J. v. Delft, H. E. Türeci, and A. Imamoglu, Nature 474, 627 (2011).
- [16] I. K. Drozdov, A. Alexandradinata, S. Jeon, S. Nadj-Perge, H. Ji, R. J. Cava, B. A. Bernevig, and A. Yazdani, Nat. Phys. 10, 664 (2014).
- [17] P. W. Anderson, Journal of Physics C: Solid State Physics 3, 2436 (1970).
- [18] J. Kondo, Progress of Theoretical Physics 32, 37 (1964).
- [19] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. 83, 349 (2011).
- [20] H. Aoki, N. Tsuji, M. Eckstein, M. Kollar, T. Oka, and P. Werner, Rev. Mod. Phys. 86, 779 (2014).
- [21] K. G. Wilson, Rev. Mod. Phys. 47, 773 (1975).
- [22] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).
- [23] C. Melnick, P. Smon, K. Yu, N. D'Imperio, A.-M. Tremblay, and G. Kotliar, "Accelerated impurity solver for dmft and its diagrammatic extensions," (2020), ArXiv:2010.08482 [cond-mat.str-el].
- [24] L.-F. Arsenault, A. Lopez-Bezanilla, O. A. von Lilienfeld, and A. J. Millis, Phys. Rev. B 90, 155136 (2014).
- [25] N. Walker, S. Kellar, Y. Zhang, and K.-M. Tam, "Neural network solver for small quantum clusters," (2020), arXiv:2008.12331 [cond-mat.str-el].
- [26] T. Xie and J. C. Grossman, Physical Review Letters 120, 145301 (2018).
- [27] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Nature Physics 13, 435 (2017).
- [28] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and

G. E. Dahl, in *Proceedings of the 34th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 70, edited by D. Precup and Y. W. Teh (PMLR, International Convention Centre, Sydney, Australia, 2017) pp. 1263–1272.

- [29] R. Bulla, T. A. Costi, and T. Pruschke, Rev. Mod. Phys. 80, 395 (2008).
- [30] F. B. Anders and A. Schiller, Phys. Rev. Lett. 95, 196801 (2005).
- [31] A. Weichselbaum and J. von Delft, Phys. Rev. Lett. 99, 076402 (2007).
- [32] K. M. Stadler, Z. P. Yin, J. von Delft, G. Kotliar, and A. Weichselbaum, Phys. Rev. Lett. **115**, 136401 (2015).
- [33] S.-S. B. Lee, J. von Delft, and A. Weichselbaum, Phys. Rev. Lett. **119**, 236402 (2017).
- [34] P. W. Anderson, Phys. Rev. **124**, 41 (1961).
- [35] J. B. Rigo and A. K. Mitchell, Phys. Rev. B 101, 241105 (2020).
- [36] D. Hendry and A. E. Feiguin, Phys. Rev. B 100, 245123 (2019).
- [37] "Supplemental Material (url will be inserted by publisher).".
- [38] The Anderson and Kondo data sets are completely disjoint, i.e., they have separate training, validation and testing sets, and are independently trained and evaluated. As such we often suppress the index Ω for brevity.
- [39] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, J. Mach. **12**, 2825 (2011).
- [40] Y. Zhang, J. C. Duchi, and M. J. Wainwright, J. Mach. 16, 3299 (2013).
- [41] Y. You, J. Demmel, C.-J. Hsieh, and R. Vuduc, in *Proceedings of the 2018 International Conference on Super-computing ICS '18* (ACM Press, New York, New York, USA, 2018) pp. 307–317.
- [42] R. Bulla, A. C. Hewson, and T. Pruschke, J. Phys. Condens. Matter 10, 8365 (1998).
- [43] Y. Meir, N. S. Wingreen, and P. A. Lee, Phys. Rev. Lett. 66, 3048 (1991).
- [44] D. Goldhaber-Gordon, H. Shtrikman, D. Mahalu, D. Abusch-Magder, U. Meirav, and M. A. Kastner, Nature **391**, 156 (1998).
- [45] F. D. M. Haldane, Phys. Rev. Lett. 40, 416 (1978).
- [46] A. Y.-T. Wang, R. J. Murdock, S. K. Kauwe, A. O. Oliynyk, A. Gurlo, J. Brgoch, K. A. Persson, and T. D. Sparks, Chem. Mater. (2020).
- [47] B. Cheng and D. M. Titterington, Stat. Sci 9, 2 (1994).
- [48] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foun*dations of Machine Learning, 2nd ed. (The MIT Press, Cambridge, MA, 2018).
- [49] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. (Springer-Verlag, New York, New York, USA, 2009).
- [50] We remind the reader that the validation and testing sets are fixed for each \mathcal{D}^{Ω} for the entirety of this work.
- [51] D. P. Kingma and J. Ba, arXiv preprint arXiv:1412.6980 (2014).
- [52] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever,

and R. Salakhutdinov, The journal of machine learning

research **15**, 1929 (2014). [53] M. Rupp, Int. J. Quantum Chem. **115**, 1058 (2015).