

Machine-learning-assisted determination of electronic correlations from magnetic resonance

Anantha Rao ¹, Stephen Carr ^{2,3}, Charles Snider ², D. E. Feldman ^{2,3}, Chandrasekhar Ramanathan ⁴ and V. F. Mitrović ²¹Department of Physics, University of Maryland, College Park, Maryland 20742, USA²Department of Physics, Brown University, Providence, Rhode Island 02912-1843, USA³Brown Theoretical Physics Center, Brown University, Providence, Rhode Island 02912-1843, USA⁴Department of Physics and Astronomy, Dartmouth College, Hanover, New Hampshire 03755, USA

(Received 9 May 2023; accepted 16 August 2023; published 30 October 2023)

In the presence of strong electronic spin correlations, the hyperfine interaction imparts long-range coupling between nuclear spins. Efficient protocols for the extraction of such complex information about electron correlations via magnetic response are not well known. Here we study how machine learning can extract material parameters and help interpret magnetic response experiments. A low-dimensional representation that classifies the strength and range of the interaction is discovered by unsupervised learning. Supervised learning generates models that predict the spatial extent of electronic correlations and the total interaction strength. Our work demonstrates the utility of artificial intelligence in the development of new probes of quantum systems, with applications to experimental studies of strongly correlated materials.

DOI: [10.1103/PhysRevResearch.5.043098](https://doi.org/10.1103/PhysRevResearch.5.043098)

I. INTRODUCTION

Quantum phases of electronic matter, such as superconductors and spin liquids, have promising applications, but they are challenging to study due to their fragile nature [1,2]. Luckily, the hyperfine interaction offers a nonintrusive coupling of the electronic spin structure to nuclear spins, providing an alternate probing mechanism of the electronic phase. As nuclear magnetic resonance (NMR) probes the nuclear spins with low-frequency pulses relative to typical electronic energies, measurement of the nuclear magnetization should not disturb the electronic ground state *a priori*. Most NMR techniques in correlated systems measure how short-range electronic spin susceptibility changes the spin dynamics through dephasing, energy dissipation, or spectral (Knight) shifts (e.g., T_2 , T_1 , and K) [3–11].

Measurement of decay processes has its limitations, however. Magnetic relaxation becomes difficult to measure in a number of different situations: in frustrated materials with a highly degenerate energy spectra, if repeated measurements heat the sample, or if the magnetic resonance echo fails to satisfy time-reversal symmetry [12–15]. Some of these confounding features have been recently reproduced by large-scale simulations of thousands of nuclear spins with long-range interactions [16]. As the details of the nuclear interaction are inherited from the electronic spin susceptibility, aspects of the electronic spin-spin correlations that were previously inaccessible to NMR can be estimated after appropriate analysis. The most successful approach required a

careful sweep of pulse angles to capture the anisotropy of the electronic spin-spin correlations. But no diagnostic tool was developed to distinguish between exponential (short-range) or power-law (long-range) decay of the interaction strength in space, although it was noted that the two types of decay produced different magnetization responses.

Here we revisit large lattice simulations of interacting nuclear spins and ask a direct question: How much information about the electronic correlations can be extracted from a single NMR experiment on an unknown material? Although we will consider only traditional spin-echo protocols at ideal pulsing, this work develops a framework for testing the efficacy of magnetic probes of electronic features. After applying standard machine-learning techniques to a large dataset of time-series simulations [17–19], if the initial parameters are predicted at a rate better than random guessing, then one can surmise that some amount of information is obtained by the proposed experiment. Moreover, a guide for interpreting real measurements is developed from feature ranking of the simulated data [20,21]. In this case, we compare the “automatically” generated approach to the previous analytical treatment [16] and highlight improvements that the data-driven approach provides.

In Fig. 1, we provide a conceptual outline of this methodology, illustrating both an interacting spin system and the inverse problem of material discovery. Our work is organized as follows: In Sec. II, we discuss the methods employed for simulations and learning the electronic susceptibility. In Sec. III, we use machine learning applications to see what information about the material is accessible from the echo measurements and generate predictive models. Finally, in Sec. IV, we discuss the results in the greater context of magnetic measurements of strongly correlated systems.

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




FIG. 1. A schematic for material discovery via NMR. Top: A magnetic resonance experiment in a fixed field B_0 allows for the controlled evolution of many nuclei, with interactions between spins mediated by the electron susceptibility χ . This susceptibility is sometimes strongly polarized along either the out-of-plane (α_z) or in-plane (α_{xy}) spin axes. Bottom: Magnetization time-series data is generated by spin-echo simulations, represented by the map $M_{sim} : \chi \rightarrow M_{xy}$, and a machine learning (ML) model is trained to predict the parameters of χ from them.

II. METHODS

A. NMR simulations

The combined nuclei-electron system can be understood by considering a Hamiltonian of the form

$$H = H_e + H_n + H_{hf}, \quad (1)$$

where H_e and H_n are the Hamiltonians of the electrons and nuclei, respectively, and H_{hf} is the interaction between the two species, mediated by the hyperfine interaction. As we are interested in just the nuclear time dynamics (NMR probes do not operate at frequencies relevant to electrons), we integrate out the degrees of freedom associated with the electrons to obtain a Hamiltonian that consists of just the nuclei and an electron-mediated nuclear-nuclear interaction. We then apply a mean-field approximation to the nuclear spins, replacing direct spin-spin interactions with a mean magnetization on each nuclei M_i , to arrive at [16]

$$H_{mf}(i) = -v_i I_i^z - \sum_{d=x,y,z} \alpha_d I_i^d M_i^d, \quad (2)$$

where the nuclei are labeled by index i , I_i^d is the spin operator along the d axis for nuclei i , and v_i are the (Zeeman) precession frequencies of the nuclei. The second term represents the mean-field interaction, with α_d the effective electron-mediated coupling strength along nuclear spin axis d . Our goal is to extract the parameters of this Hamiltonian solely from the shape of the total in-plane magnetization curve during a spin-echo pulse sequence. In our simulations, the echos occur at a time $t = 2\tau = 200 \mu\text{s}$, as shown in Fig. 2(a). The spin-echo protocol consists of polarizing the nuclear spins along the z direction, applying an in-plane 90° magnetic pulse at $t = 0$ and then applying a refocusing 180° magnetic pulse at $t = \tau$.

Note that the α_d depend on both the hyperfine interaction tensor and the electronic susceptibility tensor. M_i^d is the effective local magnetic field felt by nuclei i along axis d due




FIG. 2. Overview of the spin-echo dataset. (a) The echo responses over all 15 000 simulations. The refocused echo due to spin realignment is observed at $t = \tau = 200 \mu\text{s}$. The region around the echo, $t = [167, 233] \mu\text{s}$, is highlighted as an “Area of Interest” that will serve as a truncated dataset alongside the full time series. (b) Absolute value of the three magnetization kernels K : Gaussian (pink) with $\xi = 8$, power (brown) with $p = 1.47$, and RKKY (green) with $\gamma = 13$. (c) Time-domain echo responses for the Gaussian, power-law, and RKKY radial kernel forms. (d) Time-domain responses within the area of interest. (e) Spin-echo responses in frequency domain, relative to the resonance frequency.

to the mean-field interaction. The overall strength of the electronic susceptibility is encoded by the α variables, while the spatial structure of the susceptibility enters into the formula for M_i^d . The effective local magnetization for a nucleus at site r_i is defined as the sum

$$\mathbf{M}_i = \frac{1}{\kappa} \sum_j K(r_{ij}) \langle \mathbf{I}_j \rangle, \quad (3)$$

with K the “kernel” for the interaction that encodes the spatial structure of the electronic spin correlations, $r_{ij} \equiv r_j - r_i$ the distance to nuclei j , and $\kappa = \sum_j K(r_{0j})$ is a normalization constant. We also introduce an effective length scale for the kernel, L , given by the weighted average of r :

$$L = \frac{1}{\kappa} \sum_j r_{0j} K(r_{0j}). \quad (4)$$

We study three forms of K . First, a short-range Gaussian that depends on a correlation length ξ , $K(r) = e^{-(r/\xi)^2}$, motivated by the susceptibility expected from a gapped spin excitation. Second, a long-range form given by a power p , $K(r) = r^{-p}$, motivated by a gapless spin excitation. Finally, the RKKY form expected from electron-mediated spin-spin interactions in a simple metal which is also dependent on a length γ , $K(x) = x^{-4}(x \cos x - \sin x)$ for $x = 2(r/\gamma)$. In Fig. 2(b), the three functional forms for K are plotted using parameters that yield similar length scales for comparison. We note that the Gaussian form decays rapidly, while the alternating sign in the tail of the RKKY form leads to an effective cancellation of any net magnetization far away from the nuclei. In contrast, the power law has a characteristic long-range tail. One of the most important parameters to extract in this model is the total effective interaction of the net magnetization on a single nuclear spin. This is given by an integral of the interaction α over the entire lattice, which we call the “weight” W . As we have already normalized K , this is simply given by

$$W = \alpha_x + \alpha_y + \alpha_z. \quad (5)$$

The time-dependent magnetization $M(t)$ is then obtained by averaging over all nuclei in a large interacting two-dimensional (2D) lattice. The time dynamics of the spins are calculated from Eq. (2) using a massively parallel GPU code [16]. The code implements discretized time propagation of a density matrix following the Lindbladian formulation of quantum dynamics. We performed calculations on a 100×100 square lattice of spin- $\frac{1}{2}$ nuclei, with a time step $dt = 160$ ns, delay time $\tau = 100$ μ s, and with precession frequencies ν_i sampled randomly from a Lorentzian distribution with a full width at half maximum linewidth of 8 kHz. We sampled α_d and the kernel-specific variables $\{\xi, p, \gamma\}$ such that the kernel integral W lies in the range [0.1,0.3] MHz. This range of W is chosen to avoid echos that have no interaction-induced features ($W < 0.1$) and cases where the echos are completely dephased (e.g., decay to zero) at our chosen τ value ($W > 0.3$). Specifically, α_z and $\alpha_x = \alpha_y$ were sampled randomly in a range of [0.01,0.1] MHz, $\xi \in [8, 32]$, $p \in [1, 3]$, and $\gamma \in [13, 52]$. A random sample was simulated if the associated W was in the range [0.1,0.3]; otherwise, that random parameter selection was discarded and a new selection was generated.

TABLE I. F_1 scores of different standard models applied to the problem of classifying the nuclear interaction as a long-range (power-law) or short-range (non-power-law) functional. Sr. No is the internal index for that model.

Sr.No	Algorithm	F_1 score
1	Logistic regression	0.795
2	K -nearest neighbors	0.793
3	Support vector support (SVM)	0.796
4	Kernel SVM	0.69
5	SGD classifier	0.79
6	Decision tree	0.756
7	Random forests	0.82

This process was looped until three datasets consisting of 5000 samples each were obtained for the three kernel types (15 000 simulations in total). Although the database shows uniform distributions of α_d and W , the L distributions were nonuniform and kernel-type dependent due to the different ways in which each uniformly sampled length-scale parameter enters into their respective kernel form [see Fig. 5(b)]. For additional details of the simulation and spin-echo protocol, see Appendix A.

B. Machine learning

We have used a combination of supervised and unsupervised learning techniques to look for patterns within the simulated time-dependent magnetization that reflect the material properties. We employed two unsupervised learning techniques that perform dimensionality reduction, principal component analysis (PCA) and variational autoencoders (VAEs). Without any reference to the initial material parameters, PCA identifies which subspace of the time-series basis provides the best metric for distinguishing the response curves [22]. VAEs offer a more sophisticated, nonlinear approach to identifying important subspaces but are more computational expensive and prone to overfitting [23]. Essentially, the VAE is trained to reproduce echo responses, and the center of the network yields a two-dimensional representation for each echo, called a latent space. One can also pick a point in this two-dimensional latent space and then use just the decoder side of the VAE to simulate how the magnetization response depends on the discovered latent space.

Within a supervised setting, our primary dataset was used to train models to predict one of five characteristic properties of each curve: the strength of the interaction along the plane ($\alpha_x = \alpha_y$), the strength perpendicular to the plane (α_z), the form of the radial kernel (Gaussian, power law, or RKKY), the integral of the interaction over the entire lattice (W), and the average length scale of the interaction over the lattice (L), a dimensionless quantity. Models built for each predictor take the average magnetization time series from a spin-echo protocol as input and predict the property of interest.

For both supervised and unsupervised learning, each echo in the dataset is normalized such that each time series only contains values between (0,1). Since we are most interested in characterizing the behavior near the peak of the echo, we also consider a truncated dataset which contains only the magnetic




FIG. 3. Results of unsupervised learning with variational autoencoders. (a) The full echo response dataset, plotted in the coordinates of latent axis 1 (LA1) and latent axis 2 (LA2) obtained by the VAE. Each of the three columns show the same two-dimensional latent space, but with different material properties chosen for the color scale. These three properties are the kernel type (1, Gaussian; 2, power; 3, RKKY), the total integral of the kernel W , and the characteristic length scale of the system L . (b) Absolute magnetization responses $|M(t)|$ generated by sampling the latent space along latent axis 2 (LA2). The sampled points are $(x, 0)$, with x equally spaced between -2.5 and 2.5 . [(c) and (d)] Same as in (a) and (b) but for a VAE trained on the truncated area of interest instead of the full time domain. The latent space sampling in (d) is now performed along a vertical slice, $(0, x)$ in the latent axis coordinates. The colors of the curves in (b) and (d) correspond to the average value of W at each selected x value, following the colorbar used in the W column of (a) and (c).

response near the spin-echo [$t \in [\frac{4}{3}\tau, \frac{7}{3}\tau]$], labeled “Area of Interest” in Fig. 2(a)]. The rapid decay of magnetization at the very start of the simulation is known as the free induction decay.

To interpret the results of our supervised classification and regression models, we use two different featurization techniques. The first is the straightforward choice: Simply use the evenly spaced magnetization values $M(t)$ within the full or truncated time domain. Pointwise feature ranking will reveal which time contributes the most information to the model prediction. However, it would remain unclear if the magnetization’s magnitude, slope, or curvature drives predictions. This motivates the other choice: multiscale polynomial featurization [24]. We partition the region of interest into n equally sized sections with $n = 3, 5, \text{ or } 10$ and fit the magnetization to a cubic-polynomial in each section:

$$M(x) \sim c_0 + c_1x + c_2x^2 + c_3x^3. \quad (6)$$

The variable $x \equiv t - t_c$ is introduced to center the time axis at the middle of the section (t_c), and the coefficients of the polynomial c_i in each section act as the new feature set. The coefficients track the local constant, linear, quadratic, and cubic behaviors of the average magnetization in each section and feature ranking highlights which of these is most critical for characterization.

III. RESULTS

A. Unsupervised learning

We begin with unsupervised dimensionality reduction. Such techniques are unlikely to provide any regression capability but can generate useful classifiers for material parameters. More importantly, if they find structure in the output

which corresponds to a material parameter, then it can help guide traditional analysis of the physical system.

The VAE successfully finds a useful (classifying) low-dimensional space while the PCA fails. Although three clusters were returned by a K -nearest neighbors algorithm after PCA, these clusters showed no clear correspondence to any of the material parameters (Appendix B 1). When applied to the full time-domain dataset, the two-dimensional latent-space generated by the VAE contains features that correspond to the total interaction strength W , the interaction range L , and the kernel type [Fig. 3(a)]. The unsupervised technique has organized simulations by increasing W along its first latent space axis (LA1). This is in agreement with the results of Ref. [16], where W was noted as the most important material parameter in determining the echo shape. Both the kernel type and the length scale L were also roughly organized within the two-dimensional space. Smaller L values were grouped in a U shape near the point $(0, 2)$, along with many of the power-law kernel types. This is partially because the sampled parameters for the power-law kernel type lead to distribution of L which was lower than the other two types, shown later in Fig. 5(b). Similar results are obtained after applying the VAE to only the windowed area of interest near the echo, but now the variable W is ordered along LA2 instead [Fig. 3(c)].

We ran specific points in the latent space through the decoder network, generating characteristic $M(t)$ curves in each case. As the first axis of the latent space (LA1) corresponded more to W than the second (LA2), we sample the points $(x, 0)$ with $x \in [-2.5, 2.5]$ in the latent space. The echo responses at low LA1 (corresponding to small W) exhibit a clean echo peak while those at large LA1 (large W) show a suppressed peak at the echo pulse and significant ringing after time τ (180° pulse) and 2τ (echo). This again follows the intuition developed in the previous analytical work and proves that unsupervised learning can generate scientifically significant

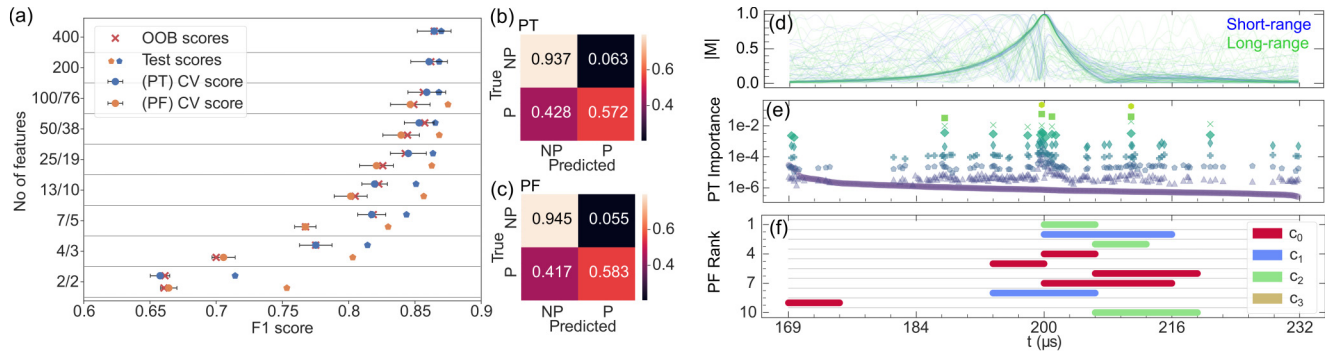


FIG. 4. Classification of the interaction's kernel form. Power (P) vs nonpower (NP), which correspond respectively to the long-range and short-range interaction kernels. (a) The total number of features vs F_1 scores of the classification models for both pointwise (PT) and polynomial (PF) featurization. Error bars show the standard deviation in F_1 scores from 10 random forests trained on the same data with different initial seeds. CV, fivefold cross-validation F_1 score; OOB, out-of-bag sample scores. (b) Confusion matrix for a random forest model trained on all 400 PT features. (c) Confusion matrix for a random forest model trained on all 76 PF features. (d) Echo responses obtained from the short-range (NP, Gaussian, and RKKY, in blue) and long-range (P, power law, in green) kernels. (e) Feature importance obtained from the Gini impurity for each PT model. The marker type and color indicates the number of features used in the training (blue denotes more features and green denotes fewer features). (f) Top 10 ranked features from the full (76 element) polynomial featurization. The color of the bar indicates the coefficient order. The bar location and width corresponds to the time window that was used for the polynomial fit.

interpretations of magnetic resonance measurements. These insights are gained effectively “for free,” requiring only a suitably sampled dataset and access to basic machine learning software. In situations with more material parameters, or for protocols with multiple pulse durations and axes, it can be challenging to make analytic progress due to the complicated nature of the resulting Hamiltonians. Unsupervised learning can immediately check which key material parameters control physical observables and yield the representative features of the response curves as that parameter is varied.

B. Classification of the spatial extent of interactions

The main shortcoming of the previous analytical work was an inability to predict if a material had short-range (exponential) or long-range (power-law) decay in the nuclear coupling [16]. By comparing simulations across the three different kernel forms with similar interaction weights W , subtle differences in the resulting $M(t)$ curves were noticed. But for a randomly sampled material, distinguishing between kernel forms seemed impossible. Note that this classification of short vs long range should not be confused with the value of L . Rather, short vs long describes if the interaction is exponentially localized or not.

Here we utilize random forest models to generate a classification model for the spatial extent of the interactions: short-range vs long-range. Since the dataset is imbalanced, we partition it into a training and test set and oversample the training set's minority class (short-range, e.g., power law) [25]. By oversampling only the training data, none of the information in the test data is double counted, avoiding the introduction of synthetic observations. While our models were unable to distinguish between the RKKY and Gaussian kernel, classification of power-law vs non-power-law interactions showed significantly better performance than random guessing [Figs. 4(b) and 4(c)]. Such a classification is important, as

a power-law decay indicates that the susceptibility arises from a gapless spin system, while an exponential decay suggest the presence of a gap in the excitation spectra. Although the RKKY type interaction arises from spin interactions in a metal (gapless), its oscillating power-law tail on average contributes zero total magnetization. In the mean-field simulations, this makes it act more similarly to the exponentially localized (Gaussian) interaction, but using a full many-body method may change this.

For classifying non-power-law data, the two featurization methods (pointwise vs polynomial) yield nearly identical confusion matrices, but the correlation between the two models' decisions were not checked. The predictions of non-power-law simulations are nearly perfect, with very few false positives (top row of the confusion matrices). However, for power-law simulations the decisions are only slightly better than random guessing, with a false-positive rate just above 0.4 (lower row of the confusion matrices).

We contrast the performance between the pointwise and polynomial featurization schemes as a function of the number of retained features [Fig. 4(a)]. To identify the essential features of the classification model, we train successive classification models that only have access to features above the median feature importance of the previous model. This allows us to create a sequence of successively simpler models. We continue this process for both featurization models until only two features remains, but because the pointwise data start with more features than the polynomial model the pairings are not of identical feature count. The models trained on pointwise features perform slightly better within each pair but also retain roughly 30% more features than the polynomial model. Although the test scores for the two featurizations are quite similar despite this discrepancy, the fivefold cross-validation scores are significantly worse for the polynomial featurization (PF), hinting at severe overfitting. When we restrict the models to using only the two most important features

obtained by this scheme, the polynomial model performs slightly better.

The two most critical time-value features are the magnitude of the magnetization at the echo pulse ($2\tau = 200 \mu\text{s}$) and the magnitude in the postecho shoulder at approximately $210 \mu\text{s}$. For classifying the range of the interaction, the relative size of the echo's decay (the value at 2τ) to the intensity of the postecho refocusing (shoulder) encodes key information about the range of the interaction. The polynomial featurization ranking in Fig. 4(f) shows that the slope and curvature at the echo pulse being the most important, with the mean values near the echo ranked as the fourth and fifth most important instead. Overall, this is fairly consistent with the theoretical understanding of gained from study of the mean-field model [16]. Short-range interactions led to more strongly suppressed echoes and reduced postecho ringing, which are encoded in the slope and curvature near the echo and the postecho shoulder. Nevertheless, we are now armed with a model that can make predictions instead of relying on pure intuition.

As mentioned previously, the most common error in our classifier were false positives for the power-law radial kernel [Figs. 4(b) and 4(c)]. Motivated by the latent structure identified by the VAE, we also train separate models for different ranges of the total interaction strength W . That is, we can refine the classification models by first using a robust prediction of W , as W and the radial-decay type both play a role in determining the shape of $M(t)$. We train three families of random forest models, each constrained only to consider simulations that have W in the range $[0.1, 0.17]$, $[0.17, 0.23]$, or $[0.23, 0.3]$, all in units of MHz. With this approach, the false-positive error rate is reduced to about 0.2 for both the pointwise and polynomial featurization (Appendix C5), which is half the error rate when not using a W filter. As the total interaction strength W explains a large amount of the $M(t)$ variations, using W screening leads to better classification models of the other physical parameters.

C. Prediction of the effective interaction strength

Finally, we now discuss models that take the $M(t)$ curves as input and return predictions of the value of the material's interaction parameters, with results shown in Fig. 5(e). The prediction of α_x and W are very accurate, with R^2 values of approximately 0.9 and 0.8, respectively. This further justifies W 's use as a screening step before using a kernel-type classifier. The models were unable to make good predictions of the out-of-plane coupling prefactor (α_z). This is not surprising, as our dataset consists of only simulations where the interaction-free magnetization lies entirely in-plane (due to the 90° – 180° pulsing sequence). When there is no out-of-plane magnetization, the α_z term does not affect the spin dynamics [Eq. (2)]. What the poor training of α_z confirms is that, even with significant in-plane interactions ($\alpha_x = \alpha_y$), the out-of-plane net magnetization remains very small, and therefore the value of α_z has little effect on resulting spin echo. As one third of the value of W is obtained from α_z , this explains the slightly worse performance of learning W compared to learning α_x directly.

To check for any cross-correlation between the kernel type and the four studied parameters of Fig. 5, we also ran the regression models after partitioning the dataset based on the

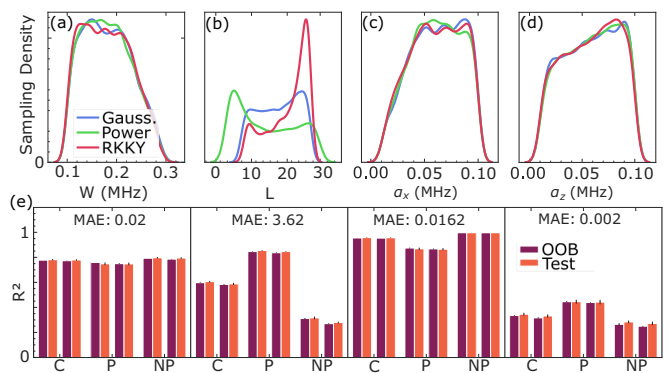


FIG. 5. Results of regression from random forest models. [(a)–(d)] Distribution of the simulation parameters present in our randomly sampled database, grouped by kernel type. (e) The out-of-bag (OOB) score and test set scores for regression models of the parameters W , L , $\alpha_x = \alpha_y$, and α_z trained on different datasets. The first set of four columns in each panel are R^2 scores for models trained on the complete dataset (C), the second set are trained on a dataset with only power-law simulations (P), and the third are trained on a dataset with only non-power-law simulations (NP). For each dataset, the first two columns are for a model trained on pointwise features while the last two columns are for a model trained on polynomial features. MAE denotes mean absolute error.

range of interaction (power “P” vs nonpower “NP”). As before, we oversample the minority class (P) in the training set due to the mismatch in dataset sizes (5000 vs 10000). The overall performance does not change significantly [Fig. 5(e)], as predictions slightly improve for the model trained on only short-range interactions (nonpower, NP) but worsen for those trained on long-range (power, P) model. Of the four trained parameters, only L shows significant deviation between the P and NP datasets. The model successfully predicts L for the power-law simulations but performs poorly on the non-power-law simulations. However, the analytic approach provided *no* method for the accurate prediction of L [16]. To see an R^2 value above 0.8, even when restricted to just power-law simulations, is a large improvement gained by the use of machine learning.

To extract physical intuition for the proposed regression schemes, we also analyze the feature ranking for each predictor variable in Appendix C6. We find that the most important features for predicting W are the average value (c_0) in the pre-echo region ($t < 2\tau$) and the slope and curvature of the magnetization right at the echo pulse (c_1, c_2). We also find that the relative feature importance changes when considering just the power-law or non-power-law decays. It indicates that short-range and long-range interactions affect the magnetization response in different ways, even at similar W . This encourages further development of better pulse sequences for studying materials with large electronic spin susceptibility, as even an unrefined spin-echo protocol is already capturing some scale-dependent behaviors.

IV. CONCLUSION

We have used data-driven approaches to understand the effects of strong electronic spin correlations on nuclear spin

dynamics. By applying machine learning methods to these spin-echo simulations, an effective probe of some spin susceptibility properties were developed. Unsupervised learning and feature-ranking of classification models provided insight into which physical parameters best describe the interacting system. The difficult task of classifying the time series based on the radial kernel was performed adequately by our random-forest models, with better performance achieved after sectioning the data into three bins of interaction magnitude W , yielding an F_1 score of 0.9.

From a broader perspective, we have demonstrated that machine learning can be used to develop inference techniques when applied to magnetic resonance experiments. Developing combined theoretical and experimental tools in this manner can improve our understanding of real materials not only by providing better traditional probes but also probes designed on real-time feedback between measurement and a machine-learning model [26–28]. A simple experiment could be run to estimate the total interaction strength W , and then a more complicated sequence that is tailored for the resulting W could be run to extract accurate information about the effective interaction length scale L and its normalized strength α_i .

Revisiting this problem with a deep learning algorithm may provide better predictions of L and α_i but requires a much larger dataset. On the other hand, designing specific pulse sequences which can access different aspects of the spin dynamics may lead to a larger improvement in our predictions. The development of highly specialized pulse sequences in order to capture the symmetry of certain physical parameters is a well developed field for on-site or nearest-neighbor interactions [29–31]. Similar protocols for systems with long-range nuclear coupling are unexplored. Reinforcement learning could possibly be leveraged to develop such sequences.

The adoption of machine learning as an interpretable method to analyze and infer the properties of complex materials is an ongoing task. As evidenced by our work, it can decode some of the wealth of information obtainable from magnetic resonance experiments.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grant No. OIA-1921199. A.S.R. acknowledges support from the Google Summer of Code program and the KVPY fellowship by Department of Science and Technology (Government of India). The calculations were conducted using computational resources and services at the Center for Computation and Visualization, Brown University.

APPENDIX A: SIMULATION OF NUCLEAR SPINS

To achieve a spin echo in our simulations, all spins begin in alignment along the z axis, and at time $t = 0$ an I^x pulse rotates the spins about the x axis by an angle θ . At time $t = \tau$, an I^y pulse rotates them by 2θ about the y axis. When $\theta = 90^\circ$, the second pulse inverts the spins and further time propagation begins to cancel any accumulated phases from the variations in each spin's resonant frequency ν . This forms a spin echo at $t = 2\tau$. We simulate this process on a 100×100 lattice of

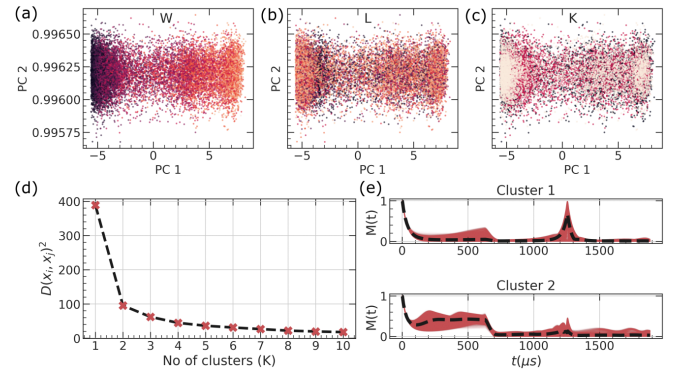


FIG. 6. PCA and K -means results. [(a)–(c)] The 2D PCA of the centered echo-response dataset. Each plot is colored based on the (a) total kernel integral W , (b) average length scale L , or (c) the type of radial kernel K . There is no discernible difference or pattern amongst the principal components. (d) K -means clustering reveals that the optimum number of clusters in the echo-response must be three from the sum of squared error (SSE) scree plot (e) Magnetization curves for all simulations in the two clusters, which correspond to low (high) kernel integrals W .

spin-1/2 interacting nuclei [16] for different values of $(\alpha_x = \alpha_y, \alpha_z, K, l_K)$. α_i is the strength of the interaction along axis $i = \{x, y, z\}$, K is the type of radial kernel and l_K is a parameter that defines the form of the specified kernel (with l_K either ξ , p , or γ).

To preprocess the simulated data before usage by a machine learning algorithm, each echo-response time series is normalized such that all values of net magnetization lie between (0,1). Since we are interested in the behavior of such systems at the echo pulse, we also generate a dataset which consists of only a narrow time window (of width $67 \mu s$) around the echo pulse (centered at $t = 200 \mu s$).

APPENDIX B: UNSUPERVISED LEARNING

1. Principal component analysis

PCA is a dimensionality reduction technique that generates a low-dimensional representation of a large dataset by finding the uncorrelated variables that maximize variance. The task of identifying these new variables can be mapped to an eigenvalue problem where the principal components (eigenvectors) and their contribution to the dataset can be obtained from singular value decomposition of the (centered) data matrix. We used the scikit-learn package to perform PCA on our datasets. Figure 6 depicts a 2D PCA of the echo-response data.

2. K -means clustering

K -means clustering is an unsupervised learning technique often used to identify distinct, nonoverlapping clusters in a dataset. The primary objective is to minimize the variation within each “cluster,” which is the mean distance between every (vectorized) datapoint assigned to that cluster. This is achieved by minimizing the objective function:

$$D(K, X) = \underset{C_1, \dots, C_k}{\text{minimize}} \left[\sum_{k=1}^K \frac{1}{|C_k|} \sum_{i=1}^{P_k} (x_i - \mu_k)^2 \right], \quad (B1)$$

where K is the number of clusters, $x \in C_K \subset X$ are one of the P_k elements of the k th cluster, and μ_k is the mean value of the k th cluster. Here, we first specify the desired number of clusters K , and then the K -means algorithm assigns each observation to exactly one of the K clusters.

The result for a K -means cluster with varying K for our dataset is given in Fig. 6(d). We find that $K = 3$ gives the last useful clustering, as increasing beyond that gives little improvement in the objective. In Fig. 6(e) we show the typical magnetization curves for the two clusters for the optimized $K = 2$ case, showing clear differences between cluster 1 (large W) and cluster 2 (small W).

3. Variational autoencoders

VAEs [23] are generative models based on layered neural networks. Assume our dataset can be described as a set of independent and identically distributed data points, $X = \{x^{(i)}\}$ with $x^{(i)} \in R^n$. Further, assume these data are sampled from a distribution with Gaussian distributed latent variables z and model parameters θ , $p_\theta[x^{(i)}|z]$. Then, finding the exact posterior density $p_\theta[z|x^{(i)}]$, e.g., solving the inverse problem of obtaining latent variables from elements of the dataset, is often an intractable problem. VAEs approximate the true posterior distribution with a tractable approximate model $q_\phi[z|x^{(i)}]$, with parameters ϕ , and provide an efficient procedure to sample efficiently from $p_\theta[x^{(i)}|z]$. In practice, a VAE is a network composed of three main components. An encoder (1) projects the input into a latent space (2), and then a decoder (3) attempts to reconstruct the input from the latent representation. After the network is trained, one can sample according to the original distribution by dropping the encoder and sampling the latent space directly. The model is trained by minimizing (over θ and ϕ) the cost function:

$$J_{\theta, \phi, x^{(i)}} = -\mathbb{E}_{z \sim q_\phi(z|x^{(i)})} \{\log p_\theta[x^{(i)}|z]\} + \beta D_{\text{KL}}\{q_\phi[z|x^{(i)}] || p_\theta(z)\}. \quad (\text{B2})$$

The first term, the reconstruction loss consists of the expected negative log-likelihood of the i th data point and favors choices of (θ, ϕ) that lead to accurate reconstructions of the input. The second term, the regularization loss D , is the Kullback-Leibler divergence between the encoder's distribution $q_\phi[z|x^{(i)}]$ and the Gaussian prior on z . A full treatment and derivations of the variational objective are given in Ref. [23].

In our setup, we use a β -VAE, where $\beta \neq 1$ is an additional weight (prefactor) assigned to the KL-divergence term in the loss function. This modifies the relative importance of the reconstruction and regularization losses. Such β tuning has been useful in identifying latent structures in highly nonlinear datasets [32]. In our simulations, we use $\beta = 2$. In addition, to prevent overfitting, we partition the dataset into training and a validation sets and train until the reconstruction loss on the validation set stops improving for seven epochs. Two different VAE models are trained for the full echo-response and the time-window echo-response data.

For the full echo-response data, the encoder has the following number of nodes in successive layers: (1882, 512, 128, 16, (2, 2)) where (2,2) is the latent space containing the results of the encoder. The regularization term

constrains this space to represent the mean and log-variance of a multivariate standardized Gaussian. The decoder has the reversed architecture, (2, 16, 128, 512, 1882), with the input to the decoder sampled from the encoder's resulting latent space via the reparametrization trick [23]. All layers in both the encoder and decoder uses a LReLU (nonlinear leaky rectified linear units) activation function with the leak set to 0.01, except for the final layer which uses a sigmoid activation function.

The architecture for the echo-window data is similar, with an encoder of shape (400, 128, 32, 16, (2, 2)), a decoder of shape (2, 16, 32, 128, 400), and using the same activation functions.

APPENDIX C: SUPERVISED LEARNING

1. Random forest

A random forest model is an ensemble technique that utilizes multiple decision trees to learn a dependent variable. For classification, the dependent variable is determined via a majority vote from multiple trees each trained on randomly chosen ("bootstrapped") samples and features from the training dataset. This method of averaging over many decision trees reduces overfitting errors from a single model applied to the entire training dataset.

Within a random forest classification model, each tree is generated by iterative branching of yes-no decisions, known as a binary tree. Eventually, the branching process terminates, and each end node is assigned a specific category (in our case, kernel type K). Ideally, each branch would put all members of the same category into the same child node. This goal is formulated as an optimization problem by minimizing the Gini impurity of each node via modification of the binary decision parameters. The Gini impurity for a dataset with C categories is defined as

$$1 - \sum_{j=1}^C p_j^2, \quad (\text{C1})$$

where p_j represents the fraction of data which is of category j . If the resulting dataset is entirely of one type (often called a "pure" node), for example, $p_1 = 0$ and $p_2 = 1$, then the Gini impurity will be zero. But if the dataset is exactly mixed, $p_1 = p_2 = 0.5$, then the Gini impurity will be 0.5.

For regression, our training criteria is instead the mean-squared error,

$$\text{MSE} = \frac{1}{N} \sum_{j=1}^N [y_i - f(x_i)]^2, \quad (\text{C2})$$

where y_i represents the label associated with the i th input x_i and f represents the regression model. Then each parameter of a tree is optimized via a gradient descent algorithm to minimize the error.

2. Performance metrics

To measure the quality of our classification models, we adopt the F_1 score, with results for all trained models given in Table I. The F_1 score can be understood as the harmonic mean

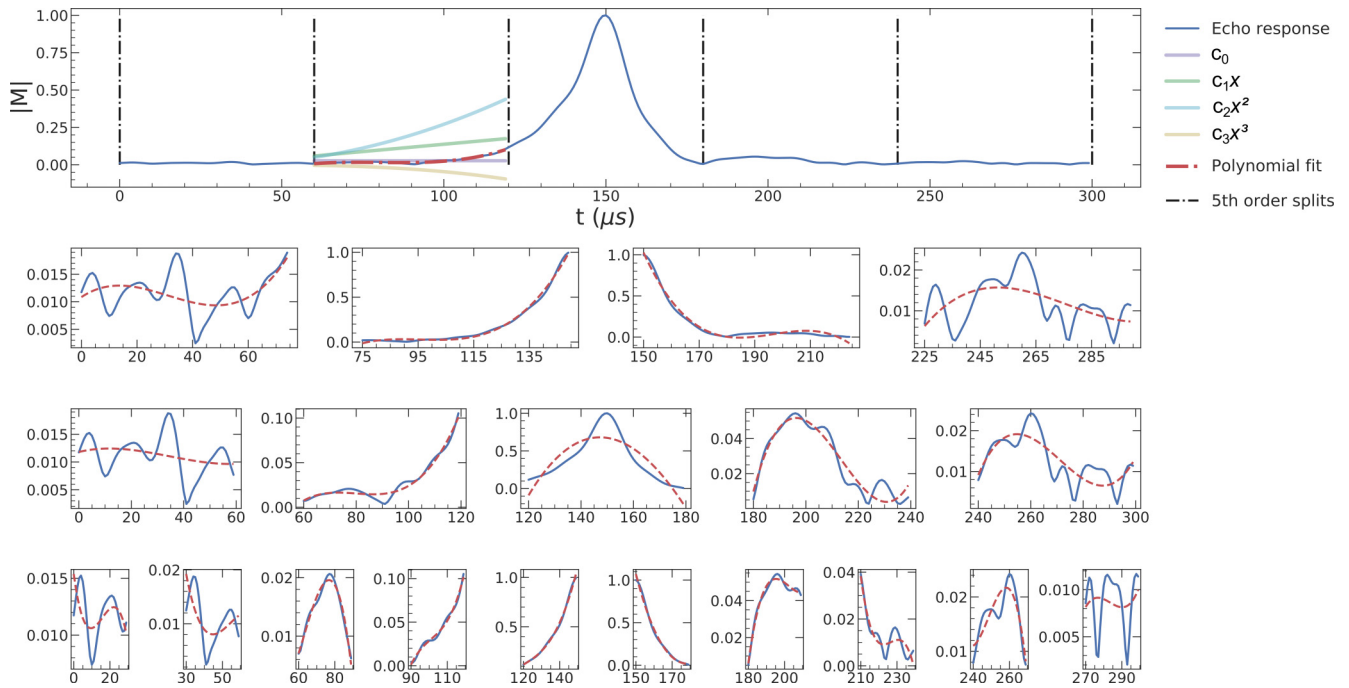


FIG. 7. Multiscale polynomial featurization. Here we perform polynomial featurization on 4, 5, and 10 equally spaced sections of a sample echo response. The top row is the full echo, while each subsequent row shows evenly spaced sections of the full echo. The polynomial fit is given by the dashed red line. The coefficients of the cubic polynomial fit to each section serve as new features for our classification and regression models. The initial three hundred pointwise features are transformed into $4 \times (4 + 5 + 10) = 76$ features (i.e., the number of polynomial parameters times the number of unique sections).

between two common classification metrics, the precision (Pr) and recall (Rc) scores. All three are defined below based on the relative rates of true positives (TP), false positives (FP), and false negatives (FN) from the model:

$$\begin{aligned} \text{Pr} &= \frac{\text{TP}}{\text{TP} + \text{FP}} \\ \text{Rc} &= \frac{\text{TP}}{\text{TP} + \text{FN}} \end{aligned} \tag{C3}$$

$$F_1 = 2 \times \frac{\text{Pr} \times \text{Rc}}{\text{Pr} + \text{Rc}}.$$

To assess the quality of our regression models, we use the R^2 score:

$$R^2 = 1 - \frac{\sum_i [y_i - f(x_i)]^2}{\sum_i (y_i - \bar{y})^2}, \tag{C4}$$

where y_i are the true values, $f(x_i)$ are the predicted values, and \bar{y} is the mean of all values in the training set.

3. Out-of-bag estimation

When optimizing a random forest model, the decision trees are repeatedly fit to bootstrapped subsets of the observations such that each “bagged” tree uses around two-thirds of the total observations [33]. The remaining one-third of the observations not used to fit a given bagged tree constitute the out-of-bag (OOB) observations. We then predict the response for a given observation using each of the trees in which that observation was OOB. This effectively allows one to treat each element of the dataset as both a training and testing element, as the OOB labeling system ensures predictions of that observation do not use any trees that were trained on it. In order to obtain a single prediction over many trees, we average over all predicted responses (regression) or take a majority vote (classification). After an OOB prediction is obtained for each element of the dataset, the overall OOB MSE or classification error is computed. The resulting OOB

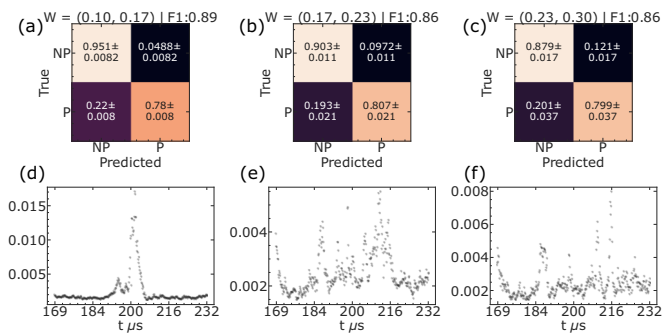


FIG. 8. Classification of kernel type based on the total value of the kernel integral from the time-series data. [(a)–(c)] Confusion matrices for the results of the models trained on each unique W range. [(d)–(f)] Feature importance for each model based on Gini impurity.

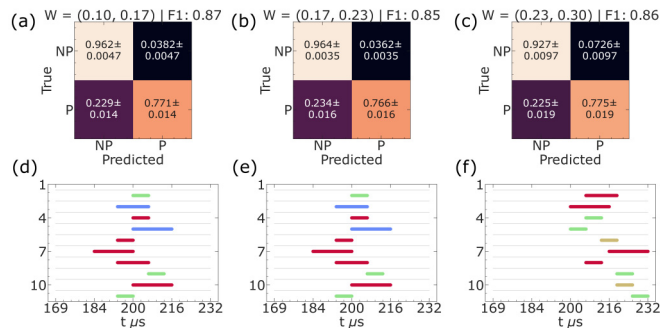


FIG. 9. Classification of kernel type based on the total value of the kernel integral from the multiscale polynomial featurization of echo responses. [(a)–(c)] Confusion matrices for the results of the models trained on each unique W range. [(d)–(f)] The top 10 features in each model, averaged over 10 iterations with different random seeds. The most important features are plotted at the top, with the width of the bar giving the time window of that polynomial fit and the color giving the power of the coefficient.

error is a valid estimate of the test error, since the response for each observation is predicted using only the trees that were not fit using that observation. In our work, we utilize the `oob_score` functionality in Sklearn [34].

4. Polynomial featurization

To turn pointwise data into polynomial features, we simply subdivide the time series into N equal sections and then fit a cubic polynomial $\sum_i c_i x^i$ in each section. This process is shown for an example spin-echo simulation in Fig. 7.

Polynomial featurization can provide clearer physical insight into the behavior of complex magnetic responses. Specifically, the curvature and slope of each feature is captured in the c_2 and c_1 terms, while the mean and any nonlinear or ringing behavior is captured by the c_0 and c_3 coefficients. Feature ranking these coefficients reveals the region and shape of the response that is most important when predicting a physical property.




FIG. 10. Feature ranking for W regression models. [(a)–(c)] The median pointwise feature importances when predicting W using (a) the complete dataset, (b) a dataset with only long-range interactions (power), and (c) one with only short-range interactions (nonpower). [(d)–(f)] The top 10 polynomial features for each model, averaged over 10 iterations with different random seeds.




FIG. 11. Feature ranking for L regression models. [(a)–(c)] The median pointwise feature importances when predicting L using (a) the complete dataset, (b) a dataset with only long-range interactions (power), and (c) one with only short-range interactions (nonpower). [(d)–(f)] The top 10 polynomial features for each model, averaged over 10 iterations with different random seeds.

5. Classification with W screening

In Fig. 8 we present the confusion matrices and relative feature importance for the power-law classifier models trained within reduced ranges of the total interaction strength (W) value, using the time-series features. In Fig. 9 we show the same results but now for the polynomial features. All models significant improvement by reducing the false-error rate of power-law simulations (the results for the non-power-law simulations are similar to the nonscreened models).

6. Feature importance for regression models

In Fig. 10, we present the relative feature importances of both the time-series and the polynomial featurizations for a W regression model. Included are results for the model trained on the complete dataset (“All Kernels”) as well as for models trained on only the power-law kernels (“Long Range”) and the nonpower kernels (“Short Range”). Figures 11 and 12 show the relative feature importances for the L and α_x regression models, respectively. In all models, the data near the echo time $\tau = 200 \mu s$ are the most important, with the pre-echo shoulder near $\tau = 170 \mu s$ also of relevance in some.




FIG. 12. Feature ranking for α_x regression models. [(a)–(c)] The median pointwise feature importances when predicting $\alpha_x = \alpha_y$ using (a) the complete dataset, (b) a dataset with only long-range interactions (power), and (c) one with only short-range interactions (nonpower). [(d)–(f)] The top 10 polynomial features for each model, averaged over 10 iterations with different random seeds.

- [1] M. R. Norman, Fragile superconductivity at high magnetic fields, *Proc. Natl. Acad. Sci. USA* **118**, e2100372118 (2021).
- [2] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, *Science* **374**, 1242 (2021).
- [3] L. B. Ioffe, S. Kivelson, and A. I. Larkin, Spin correlations and NMR relaxation rates in strongly correlated electron systems, *Phys. Rev. B* **44**, 12537 (1991).
- [4] A. Rigamonti, F. Borsa, and P. Carretta, Basic aspects and main results of NMR-NQR spectroscopies in high-temperature superconductors, *Rep. Prog. Phys.* **61**, 1367 (1998).
- [5] Y. Fagot-Revurat, M. Horvatić, C. Berthier, P. Ségransan, G. Dhalenne, and A. Revcolevschi, NMR evidence for a magnetic soliton lattice in the high-field phase of CuGeO_3 , *Phys. Rev. Lett.* **77**, 1861 (1996).
- [6] C. Berthier, M. Horvatić, M.-H. Julien, H. Mayaffre, and S. Krämer, Nuclear magnetic resonance in high magnetic field: Application to condensed matter physics, *C. R. Phys.* **18**, 331 (2017).
- [7] V. F. Mitrović, H. N. Bachman, W. P. Halperin, A. P. Reyes, P. Kuhns, and W. G. Moulton, Pseudogap in $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ from NMR in high magnetic fields, *Phys. Rev. B* **66**, 014511 (2002).
- [8] G. Koutroulakis, M. D. Stewart, V. F. Mitrović, M. Horvatić, C. Berthier, G. Lapertot, and J. Flouquet, Field evolution of coexisting superconducting and magnetic orders in CeCoIn_5 , *Phys. Rev. Lett.* **104**, 087001 (2010).
- [9] H. Alloul, NMR in strongly correlated materials, *Scholarpedia* **10**, 30632 (2015).
- [10] I. Vinograd, R. Zhou, M. Hirata, T. Wu, H. Mayaffre, S. Krämer, R. Liang, W. N. Hardy, D. A. Bonn, and M.-H. Julien, Locally commensurate charge-density wave with three-unit-cell periodicity in $\text{YBa}_2\text{Cu}_3\text{O}_y$, *Nat. Commun.* **12**, 3274 (2021).
- [11] N. Janša, A. Zorko, M. Gomilšek, M. Pregelj, K. W. Krämer, D. Biner, A. Biffin, C. Rüegg, and M. Klanjšek, Observation of two types of fractional excitation in the Kitaev honeycomb magnet, *Nat. Phys.* **14**, 786 (2018).
- [12] K. Ishida, M. Manago, K. Kinjo, and Y. Maeno, Reduction of the ^{17}O knight shift in the superconducting state and the heat-up effect by NMR pulses on Sr_2RuO_4 , *J. Phys. Soc. Jpn.* **89**, 034712 (2020).
- [13] A. Pustogow, Y. Luo, A. Chronister, Y.-S. Su, D. A. Sokolov, F. Jerzembeck, A. P. Mackenzie, C. W. Hicks, N. Kikugawa, S. Raghu, E. D. Bauer, and S. E. Brown, Constraints on the superconducting order parameter in Sr_2RuO_4 from Oxygen-17 nuclear magnetic resonance, *Nature (Lond.)* **574**, 72 (2019).
- [14] I. Vinograd, S. P. Edwards, Z. Wang, T. Kissikov, J. K. Byland, J. R. Badger, V. Taufour, and N. J. Curro, Inhomogeneous Knight shift in vortex cores of superconducting FeSe , *Phys. Rev. B* **104**, 014502 (2021).
- [15] V. F. Mitrović, G. Koutroulakis, M. Klanjšek, M. Horvatić, C. Berthier, G. Knebel, G. Lapertot, and J. Flouquet, Comment on “Texture in the superconducting order parameter of CeCoIn_5 revealed by nuclear magnetic resonance,” *Phys. Rev. Lett.* **101**, 039701 (2008).
- [16] S. Carr, C. Snider, D. E. Feldman, C. Ramanathan, J. B. Marston, and V. F. Mitrović, Signatures of electronic correlations and spin-susceptibility anisotropy in nuclear magnetic resonance, *Phys. Rev. B* **106**, L041119 (2022).
- [17] D.-W. Li, A. L. Hansen, C. Yuan, L. Bruschweiler-Li, and R. Brüschweiler, DEEP picker is a deep neural network for accurate deconvolution of complex two-dimensional NMR spectra, *Nat. Commun.* **12**, 5229 (2021).
- [18] M. Cordova, M. Balodis, B. S. de Almeida, M. Ceriotti, and L. Emsley, Bayesian probabilistic assignment of chemical shifts in organic solids, *Sci. Adv.* **7**, eabk2341 (2021).
- [19] C. Cobas, NMR signal processing, prediction, and structure verification with machine learning techniques, *Magn. Reson. Chem.* **58**, 512 (2020).
- [20] I. Bratko, Machine learning: Between accuracy and interpretability, in *Learning, Networks and Statistics*, edited by G. Della Riccia, H.-J. Lenz, and R. Kruse (Springer Vienna, Vienna, 1997), pp. 163–177.
- [21] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [22] I. T. Jolliffe and J. Cadima, Principal component analysis: A review and recent developments, *Philos. Trans. R. Soc. A* **374**, 20150202 (2016).
- [23] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, *arXiv:1312.6114* [stat.ML] (2022).
- [24] S. B. Torrisi, M. R. Carbone, B. A. Rohr, J. H. Montoya, Y. Ha, J. Yano, S. K. Suram, and L. Hung, Random forest machine learning models for interpretable X-ray absorption near-edge structure spectrum-property relationships, *npj Comput. Mater.* **6**, 109 (2020).
- [25] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, SMOTE: synthetic minority over-sampling technique, *J. Artif. Intell. Res.* **16**, 321 (2002).
- [26] S. Mavadia, V. Frey, J. Sastrawan, S. Dona, and M. J. Biercuk, Prediction and real-time compensation of qubit decoherence via machine learning, *Nat. Commun.* **8**, 14106 (2017).
- [27] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, Machine learning at the energy and intensity frontiers of particle physics, *Nature (Lond.)* **560**, 41 (2018).
- [28] K. Kutsukake, Y. Nagai, T. Horikawa, and H. Banba, Real-time prediction of interstitial oxygen concentration in czochralski silicon using machine learning, *Appl. Phys. Expr.* **13**, 125502 (2020).
- [29] M. H. Levitt, Symmetry-based pulse sequences in magic-angle spinning solid-state NMR, in *Encyclopedia of Nuclear Magnetic Resonance*, edited by D. M. Grant and R. K. Harris, Advances in NMR, Vol. 9 (John Wiley & Sons, Ltd, Chichester, 2002), pp. 165–196.
- [30] M. H. Levitt, Symmetry in the design of NMR multiple-pulse sequences, *J. Chem. Phys.* **128**, 052205 (2008).
- [31] I. Schwartz, J. Scheuer, B. Tratzmiller, S. Müller, Q. Chen, I. Dhand, Z.-Y. Wang, C. Müller, B. Naydenov, F. Jelezko, and M. B. Plenio, Robust optical polarization of nuclear spin baths using Hamiltonian engineering of nitrogen-vacancy center quantum dynamics, *Sci. Adv.* **4**, eaat8978 (2018).
- [32] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, beta-VAE: Learning basic visual concepts with a constrained variational

- framework, in *International Conference on Learning Representations* (OpenReview.net, Toulon, France, 2017).
- [33] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: with Applications in R* (Springer, New York, 2013), pp. 303–335.
- [34] 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 Édouard Duchesnay, Scikit-learn: Machine learning in python, *J. Mach. Learn. Res.* **12**, 2825 (2011).