How to Address Nonlinearity and Estimate Uncertainty with Polynomial Fitting in Hubbard U Calculations for DFT+U

Alan Antony,[1, a)] David D. O'Regan,[1, b)] and Lórien MacEnulty[1, 2, c)]

[1]*School of Physics, CRANN Institute and AMBER Research Centre, Trinity College Dublin, The University of Dublin, Ireland*
[2]*Modeling Division, Institute of Condensed Matter and Nanosciences, Université Catholique de Louvain, Chemin des Étoiles 8, Louvain-la-Neuve 1348, Belgium*

[a)]Corresponding author: antonyal@tcd.ie
[b)]david.o.regan@tcd.ie
[c)]lorien.macenulty@uclouvain.be

Abstract. Finite-difference linear response calculations for the Hubbard U and Hund's J parameters in DFT+U—a corrective method aiming to improve the accuracy of DFT calculations of strongly correlated electron systems like Mott insulators—typically involve linear regression of small response datasets. Many subspaces of material systems, however, are found to exhibit nonlinear response (e.g., curvature) that may disrupt the extraction of the required linear component. For these systems, low-order polynomial regressions can capture curvature in linear response datasets, permitting the quantification of uncertainty on the Hubbard parameters through the unbiased root-mean-square error. What remains unclear in this context is the optimal strategy practitioners can employ to avoid the overfitting that occurs at higher polynomial degrees. In this work, we determine uncertainties on the Hubbard parameters and investigate polynomial degree selection strategies based on leave-one-out cross validation (LOOCV) and repeated K-fold cross validation (RKFCV). We show that LOOCV is a practical technique, not least due to its simple closed-form formula, given the small perturbation-occupancy datasets typically used in DFT+U. For larger datasets, LOOCV exhibits high variance in its estimated errors, whereas RKFCV provides more stable estimates. In practical calculations, we have found that fitting to a quadratic or cubic polynomial before taking the slope at zero perturbation of the resulting fit satisfactorily addresses the issue of nonlinear response.

INTRODUCTION

Density functional theory (DFT) is a computational framework for determining the properties of materials. In DFT the total energy is reformulated as a functional of the electron density, which includes an approximated exchange-correlation (XC) term. Common XC functionals are the local density approximation (LDA) and generalized gradient approximation (GGA), and although more accurate functionals exist, they have higher computational cost. DFT using such (semi-)local functionals tends to incorrectly predict the electronic structure of strongly correlated materials, where strong electron interactions restrict electronic motion [1]. The DFT+U(+J) method introduces subspace-tailored corrective terms, the strengths of which are set by the Hubbard U and Hund's J parameters. DFT+U was originally designed to better capture correlation effects such as on-site Coulomb interactions and intra-atomic exchange coupling. U and J are often determined semiempirically, although this precludes the application of DFT+U+J to materials that have not been discovered or synthesized. This semiempiricism is limiting within high-throughput discovery-oriented workflows, where structures are randomly generated and screened for certain properties.

Alternatively, U and J can be determined ab initio through various techniques, among them minimum-tracking linear response [2] (see Appendix A), cRPA [3], and the focus of this work: the self-consistent field (SCF) linear response method [4], wherein U is determined for a chosen subspace by evaluating its response to a range of small subspace potential perturbations (α). In this formalism, suppressing matrix inversion for simplicity, U is defined as

$$\mathrm{U} = \chi_0^{-1} - \chi^{-1} = \underbrace{\left(\left.\frac{d(n_0^{\uparrow}+n_0^{\downarrow})}{d\alpha}\right|_{\alpha=0}\right)^{-1}}_{\chi_0} - \underbrace{\left(\left.\frac{d(n^{\uparrow}+n^{\downarrow})}{d\alpha}\right|_{\alpha=0}\right)^{-1}}_{\chi}, \tag{1}$$

where χ_0 (χ) is the response of the unscreened (screened) total occupancy n_0^{σ} (n^{σ}) on spin channel σ due to a minor subspace potential perturbation α in the absence (presence) of electronic screening. J is determined analogously [5, 6, 7], although we relegate all details of this analogy to Appendix A to focus the ensuing discussion on the Hubbard U. In practice, χ_0 and χ are extracted as the slopes of ordinary least-squares (OLS) linear regression, typically conducted on 5-9 perturbation-total-occupancy data points. However, many system-functional-subspace combinations exhibit non-negligible nonlinear response [7, 8, 9]. Nonlinearity in the response often occurs near to empty or full subspace filling, and hence it appears to be a drawback of overly optimized subspace choices. It also arises generally for larger perturbation strengths, which are sometimes necessary to improve the response's signal-to-noise ratio.

How to address nonlinearity in this context is of contemporary practical interest. Here, we develop a straightforward nonlinear response postprocessing method, designed expressly for uncomplicated integration into linear response (LR) practitioners' existing Hubbard U workflows. Leveraging the simplicity of low-order polynomial regressions, we apply well-attested statistical techniques to real perturbation-occupancy datasets, providing suggestions on how to regress response and avoid under- or overfitting. Our inquiries fit into a broader discussion on obtaining uncertainty estimates for the Hubbard parameters, which quantify the noisiness of the underlying response even when it is linear.

METHODOLOGY

Consider a screened or unscreened LR dataset comprising $5 \leq k \leq 9$ (α_i, $N_i = n_i^{\uparrow} + n_i^{\downarrow}$) pairs, indexed from $i = 1$ to k. Assuming no uncertainty in α_i, one may use OLS to fit a polynomial function $f^{(p)}(\alpha) = \sum_{q=0}^{p} c_q \alpha^q$ of degree p and coefficients $\{c_q\}$ to the data, about which the data points are normally distributed with zero mean and constant variance. The behavior of this regression may be characterized by the unbiased mean squared error (MSE) of the curve—the sum of squared residuals, differences between the data point and the curve, divided by degrees of freedom of the curve ($k-(p+1)$), the square root of which yields the unbiased root-mean-square (RMS) error. The uncertainty on the Hubbard U and its constituent response functions are then obtained as [10]

$$\sigma_{\mathrm{U}} = \sqrt{\left(\frac{\sigma_\chi}{\chi^2}\right)^2 + \left(\frac{\sigma_{\chi_0}}{\chi_0^2}\right)^2}, \quad \sigma_\chi = \sqrt{\underbrace{\frac{\sum_{i=1}^{k}[N_i - f^{(p)}(\alpha_i)]^2}{k-p-1}}_{\text{unbiased MSE}} (\boldsymbol{X}^\top \boldsymbol{X})^{-1}_{22}}, \quad \boldsymbol{X} = \begin{pmatrix} 1 & \alpha_1 & \alpha_1^2 & \ldots & \alpha_1^p \\ 1 & \alpha_2 & \alpha_2^2 & \ldots & \alpha_2^p \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ 1 & \alpha_k & \alpha_k^2 & \ldots & \alpha_k^p \end{pmatrix}, \quad (2)$$

where $\boldsymbol{X}$ is the design matrix, and the $(\boldsymbol{X}^\top \boldsymbol{X})^{-1}_{22}$ term encodes the effect of the distribution of α_i on c_1 (i.e., the first-order coefficient of the polynomial function), and so on the value of χ or χ_0. The $(\boldsymbol{X}^\top \boldsymbol{X})^{-1}_{22}$ term, which stems from the formulation of the variance-covariance matrix for OLS, in the simplest case provides the physical units and spacing of the perturbations α.

While Eq. (2) effectively quantifies the uncertainty on the Hubbard U, it does not permit the identification of the most appropriate polynomial degree to fit the LR data. OLS minimizes the sum of the squared residuals, but as p increases, the function gains additional degrees of freedom, allowing it to bend closer to data points and necessarily shrinking the (biased) RMS error. For noisy datasets, then, increasing p causes the curve to better fit the noise, minimizing the RMS error but overfitting the curve [11]. The unbiased RMS error, which features in the definition of σ_χ, introduces a penalty for higher p to address this, although it is unclear whether minimizing the unbiased RMS or σ_χ corresponds to the optimal p. Given the typically small datasets aggregated for the first-principles determination of U, we test the efficacy of leave-one-out cross validation (LOOCV), a diagnostic for overfitting that quantifies a regression's ability to predict unseen data. This property may be minimized to identify the optimal p regressing a particular dataset.

In the above dataset of k (α_i, N_i) points, LOOCV removes a point and uses the remaining $k-1$ points to compute the parameters of a polynomial regression of degree p. The difference between the omitted point and the point as predicted by the curve is squared, then recorded as the validation error of the curve. This process is repeated iteratively, changing the omitted point until all k points have been used for validation. The mean of the validation errors is then calculated, generating a validation error quantifier for the curve. By repeating this process for different p, one may map the relation between validation error and polynomial degree; the smallest validation error is identified as the optimal curve. Fortunately, an analytical formula encapsulates this process entirely and is given by

$$\sigma_{\mathrm{LOOCV}} = \frac{1}{k} \sum_{i=1}^{\mathrm{k}} \frac{(N_i - f^{(p)}(\alpha_{\mathrm{i}}))^2}{(1 - \boldsymbol{H}_{ii})^2}, \quad (3)$$

where $\boldsymbol{H}_{ii}$ is the ith diagonal element of the hat matrix $\boldsymbol{H} = \boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1}\boldsymbol{X}^\top$ [11]. $\boldsymbol{H}_{ii}$ weights the residual of the ith data point by its contribution to the regression parameters. Overfitting due to high polynomial degree models can lead to instability in the diagonals of the hat matrix, resulting in deviations in σ_{LOOCV}.

If more LR datapoints are available, practitioners can consider using the typically more stable generalization of LOOCV called repeated K-fold cross validation (RKFCV). RKFCV randomly partitions the dataset into K equally sized folds. The p-degree polynomial regressions are trained on $K-1$ folds and validation tested on the final fold, which produces a validation error defined as the sum of the square of residuals across all data points in the fold. This

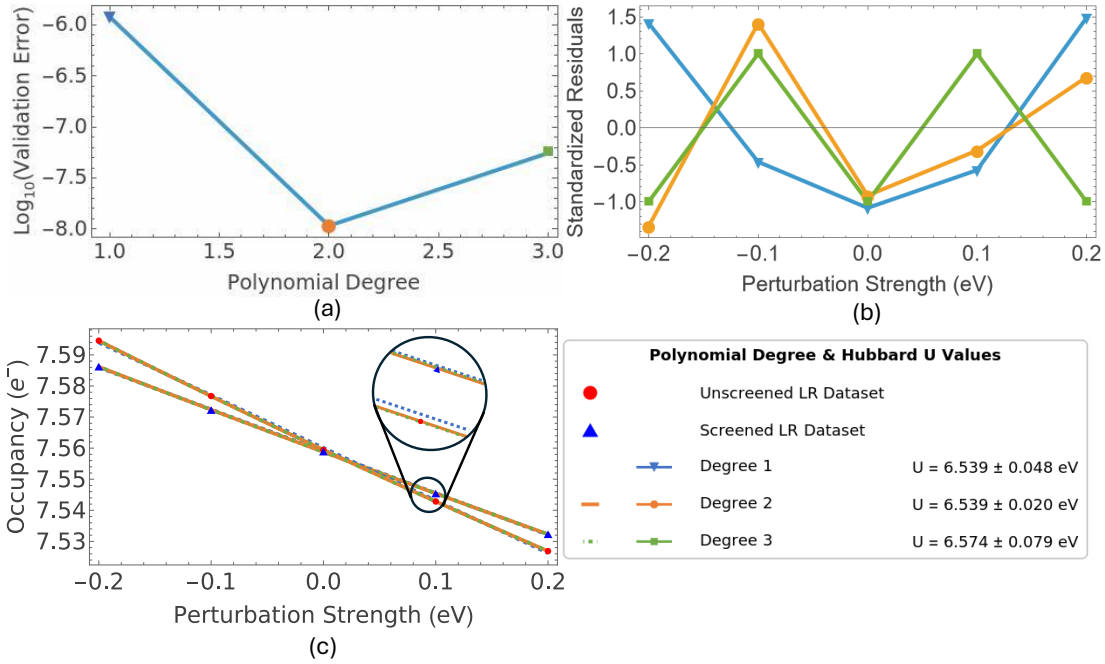

FIGURE 1: For the five-point LR dataset, perturbations conducted on Ni $3d$ sites in AF_I NiO: (a) LOOCV validation error (quantity of Eq. (3) shown on a log scale) of the unscreened data and (see bottom right panel) associated Hubbard U parameters as a function of polynomial degree p using the LOOCV procedure; (b) standardized residuals as a function of perturbation strength α for $p = 1,\ 2$, and 3 [13]; and (c) unscreened and screened occupancies versus perturbation overlaid with regressed curves for $p = 1,\ 2$, and 3.

process is repeated iteratively, changing out the validation fold each time. The mean of the validation errors across all folds is determined. The dataset is then randomly partitioned again into K folds, and the process repeats. The mean of the validation error across all repeats is calculated and mapped to and optimized with respect to degree p.

To test the effectiveness of the LOOCV method in an LR context, we analyze regressions up to degree $k-2$—the maximum polynomial degree possible, as each training fold contains at most $k-1$ points—for two ABINIT-generated [12] perturbation-total-occupancy LR datasets. The first is a typically sized LR set of five data points, and the second is much larger, with 39 data points. We repurpose DFT calculations performed on Ni $3d$ subspaces in antiferromagnetic (AF) NiO, two of which were initially published in [7]. In addition to the calculation details in Appendix B, we refer the reader to this article for more precise computational details. The two cross-validation methods are then assessed using a residual analysis based on the principle that residuals should be randomly distributed about the regression curve, not imbued with polynomial symmetry coming from the regression. To this end, we first investigate the randomness of the residuals produced by the optimal degree as predicted by LOOCV [13]. For the large perturbation-occupancy response dataset, the LOOCV-predicted optimal degree is then compared to the optimal degrees predicted by RKFCV with different validation set (fold) sizes to assess the variability of the LOOCV technique. In the RKFCV applications, if the dataset was not evenly divisible, some points were left out. Details of the linear response datasets used in this investigation can be seen in Appendix B.

RESULTS

As shown in Fig. 1(a), LOOCV demonstrated minimized validation error at $p = 2$ for both the unscreened and screened values of the small dataset. Indeed, the standardized residuals in Fig. 1(b) demonstrate noticeable randomness, in particular, asymmetry across the $\alpha = 0$ axis for $p = 2$ only, corroborating the LOOCV prediction. Standardized residuals are plotted to facilitate comparison between the residuals of the different curves, as they were often of different magnitude, and to identify potential outliers in the data. Please consult Sec. 2.2 of [13] for more details and a description of the standardization process. The associated Hubbard U—where $p = 2$ was used to regress both the

unscreened and screened response data—was identical to that at $p = 1$; however, the $p = 2$ uncertainty is much lower, even minimized compared to the other degrees tested. It is unclear whether there is a correlation between minimized σ_U and minimized validation error, particularly, as based on our method, minimized σ_U does not necessarily mean that σ_χ and σ_{χ_0} were simultaneously minimized via $p = 2$ regression. For example, for a nearly identical dataset to the above, differing in its magnetic ordering, lattice parameter, and k-point sampling (AF_{II} NiO instead of AF_I), LOOCV identified $p = 1$ for the screened dataset but $p_{max} = 3$ for the unscreened one. Systematic testing across a broad range of LR datasets and reevaluation of assumptions underlying the behavior of χ_0 and χ are necessary to solidify the efficacy of taking σ_U as an indicator of optimal p. Although the value of U is shown in Fig. 1 to vary only subtly with p, an example of large nonlinear response and hence p dependence is shown in Fig. (2) of [7].

Concurrence of the LOOCV validation error minimum at the maximum testable polynomial degree is an indication that the dataset is perhaps too small, as higher-order polynomial degrees would need to be tested to verify it as a minimum. This, in addition to the standard statistical benefits of additional points, leads us to recommend performing 6+ perturbations for rigorous LR demonstrating nonlinearity. Not too many data points should be used, however. For the 39-point dataset, LOOCV predicted a minimum validation error at degree 14 [Fig. 2(a)]. This behavior is attributable to the high variance of LOOCV. Each model, for a given degree, is trained on $k-1$ data points, making the models highly correlated, as $k-2$ datapoints are shared between any model. This results in high correlation between the predicted validation errors and hence a high variance in the final averaged validation error quantifier. While LOOCV is a nearly unbiased estimator of the generalization error (i.e., with an infinite dataset, the LOOCV validation error would closely align with the generalization error of the curve), this property has a tradeoff with high variance in its estimates. It can yield unstable estimates of the optimal degree, particularly when testing high-order polynomial degrees. RKFCV can mitigate this issue, trading variance for bias by increasing the validation set size [11]. As a result, in Fig. 2(a) we observe that the larger validation set sizes afforded by RKFCV provided more stable validation error curves, which predicted an optimal degree of 5, even as we changed the number of folds K. As shown in Fig. 2(b), convergence to a reliable optimal degree required validation set sizes with around 15%-20% of the total dataset. While the residuals [Fig. 2(c)] clearly demonstrate polynomial symmetry for $p = 2$, $p = 5$ exhibits more noticeable randomness, albeit to a degree that is not a priori distinguishable from the randomness of $p = 14$ residuals.

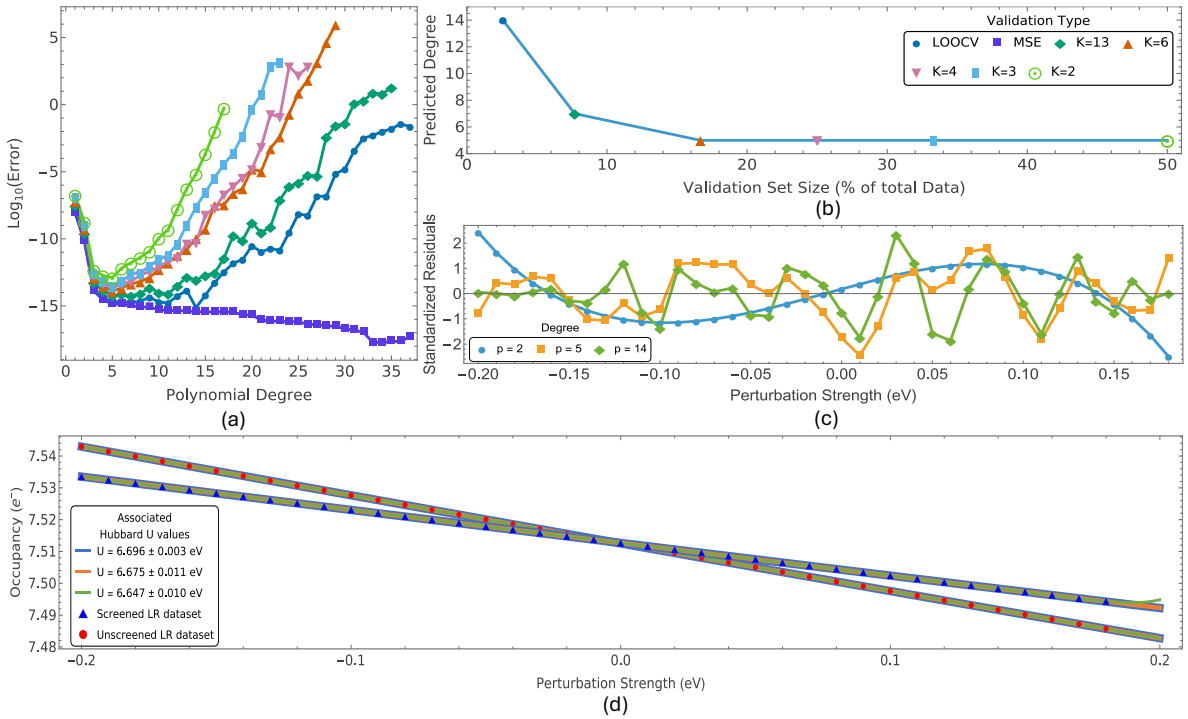

FIGURE 2: For the 39-point, unscreened LR dataset: (a) LOOCV error using the LOOCV procedure, RKFCV error using the RKFCV procedure using K validation sets, and MSE [shown in Eq. (2)] as a function of polynomial degree; (b) RKFCV-predicted optimal degree versus validation set size; (c) standardized residuals of fitted polynomials [13], which are shown in (d) for unscreened and screened response data with $p = 2$, $p = 5$, and $p = 14$.

CONCLUSION

This work investigated uncertainty quantification for the Hubbard U and Hund's J parameters, paying particular attention to polynomial regression and cross-validation-based methodologies to address nonlinearity in the response of finite-difference DFT calculations. We have found that fitting to a quadratic or cubic polynomial—as the LOOCV procedure identified and residual analysis confirmed for some archetypal LR datasets conducted on Ni 3*d* of NiO—before taking the slope of the resulting fit at zero perturbation satisfactorily addresses the issue of nonlinear response. With its simple closed-form formula, LOOCV was shown to be a useful and lightweight metric for selection of a well-fitted polynomial degree for the small datasets typical to LR, while RKFCV was found to be a more reliable method for polynomial optimization for larger datasets owing to validation set sizes approaching 15%–20% of the full dataset. In practical calculations, even for well-behaved systems with sensible DFT+U subspaces, nonlinear but low-order polynomial fitting before taking the slope at zero perturbation of the resulting fit is often necessary, and this satisfactorily addresses the issue of nonlinear response. The optimal polynomial degree, which is easily assessed via Eq. (3), is expected to depend somewhat for a given system on details of pseudopotential, basis, the DFT+U projector set and how their nonorthogonality is treated, as well as the supercell used to isolate the perturbed atom in the case of solids. Future research pathways include (i) identification and application of a suitable quantifier of the heteroskedasticity and symmetry about the zero-perturbation axis characteristic to model residuals, ultimately to automate and sophisticate their inspection; (ii) a systematic extension of this type of assessment to minimum-tracking linear response [2]; and (iii) comparison of the cross-validation methods with the uncertainty on the Hubbard parameter itself as defined in Eq. (2) in terms of efficiency in finding the optimal polynomial degree.

ACKNOWLEDGMENTS

This research was supported by Taighde Éireann – Research Ireland under Award No. 12/RC/2278_P2 and the European Regional Development Fund. Calculations were performed on the Boyle cluster, maintained by the Trinity Research IT and funded through grants from the European Research Council and Taighde Éireann.

REFERENCES

1. S. A. Tolba, K. M. Gameel, B. A. Ali, H. A. Almossalami, and N. K. Allam, in Density Functional Calculations - Recent Progresses of Theory and Application, edited by G. Yang (IntechOpen, London, 2018), Chap. 1.
2. E. B. Linscott, D. J. Cole, M. C. Payne, and D. D. O'Regan, Phys. Rev. B **98**, 235157 (2018).
3. M. Springer and F. Aryasetiawan, Phys. Rev. B **57**, 4364 (1998).
4. M. Cococcioni and S. de Gironcoli, Phys. Rev. B **71**, 035105 (2005).
5. O. K. Orhan and D. D. O'Regan, Phys. Rev. B **101**, 245137 (2020).
6. D. S. Lambert and D. D. O'Regan, Phys. Rev. Res. **5**, 013160 (2023).
7. L. MacEnulty and D. D. O'Regan, Phys. Rev. B **108**, 245137 (2023).
8. H. J. Kulik and N. Marzari, J. Chem. Phys. **133**, 114103 (2010).
9. M. Sun, T. Wu, B. Chen, Z. Li, Q. Lu, L. Lu, H. H. Wong, C. H. Chan, and B. Huang, Adv. Intell. Discovery , e202500160 (2025).
10. L. MacEnulty, Interrogation and refinement of DFT+U+J energetics in magnetic systems, Ph.D. thesis, Trinity College Dublin, (2025).
11. G. James, D. Witten, T. Hastie, R. Tibshirani, and J. Taylor, An Introduction to Statistical Learning (Springer Texts in Statistics, 2023), pp. 204–209.
12. M. J. Verstraete, J. Abreu, G. E. Allemand, B. Amadon, G. Antonius, M. Azizi, L. Baguet, C. Barat, L. Bastogne, R. Béjaud, et al., J. Chem. Phys. **163** (2025).
13. V. Verma, J. Inf. Syst. Eng. Manag. **10**, 1 (2025).
14. G. C. Moore, M. K. Horton, E. Linscott, A. M. Ganose, M. Siron, D. D. O'Regan, and K. A. Persson, Phys. Rev. Mater. **8**, 014409 (2024).
15. B. Himmetoglu, R. M. Wentzcovitch, and M. Cococcioni, Phys. Rev. B **84**, 115108 (2011).
16. L. MacEnulty, J. P. A. de Mendonça, R. Poloni, and D. D. O'Regan, J. Chem. Theory Comput. **22**, 240 (2026).
17. L. MacEnulty, M. Giantomassi, B. Amadon, G.-M. Rignanese, and D. D. O'Regan, Electron. Struct. **6**, 037003 (2024).
18. J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).
19. F. Jollet, M. Torrent, and N. Holzwarth, Comput. Phys. Commun. **185**, 1246 (2014).

Appendix A: Determination and uncertainty quantification for the Hund's J and minimum-tracking LR parameters

In the SCF LR formalism, the Hund's J is defined [6, 14] as

$$\mathrm{J} = -\left(\chi_{0_M}^{-1} - \chi_M^{-1}\right) = -\underbrace{\left(\left.\frac{d(n_0^{\uparrow} - n_0^{\downarrow})}{d\beta}\right|_{\beta=0}\right)^{-1}}_{\chi_{0_M}} + \underbrace{\left(\left.\frac{d(n^{\uparrow} - n^{\downarrow})}{d\beta}\right|_{\beta=0}\right)^{-1}}_{\chi_M}, \tag{4}$$

where β is a potential perturbation applied positively to the spin-up channel and negatively to the spin-down (in equal magnitude), and χ_{0_M} and χ_M are, respectively, the unscreened and screened magnetic responses to that perturbation. This formula is a modification of that proposed in [15]. In the present work, the response matrices are treated as scalars. The uncertainty on Hund's J and its constituent response functions, arising from polynomial regression on a dataset $\{(\beta_i,\ M_i = n_i^{\uparrow} - n_i^{\downarrow})\}$, are then obtained from Eq. (2), replacing all occurrences of χ_0 with χ_{0_M}, χ with χ_M, and α with β.

The equivalent Hubbard parameters in the minimum-tracking (MT) LR formalism [2] are instead

$$\mathrm{U} = \frac{1}{2}\frac{d(V_{\mathrm{Hxc}}^{\uparrow} + V_{\mathrm{Hxc}}^{\downarrow})}{dN} \qquad \mathrm{J} = -\frac{1}{2}\frac{d(V_{\mathrm{Hxc}}^{\uparrow} - V_{\mathrm{Hxc}}^{\downarrow})}{dM}\ , \tag{5}$$

where $V_{\mathrm{Hxc}}^{\sigma}$ is the subspace-averaged Hartree and exchange-correlation potential on spin-channel σ. Finite-differences response data within MT LR can also, of course, benefit from LOOCV analysis to help pick the optimal fitting polynomial degree. In this case, one may address non-linearity in the response by fitting degree p polynomial regressions $f^{(p)}\left(N(\alpha)\right) = \sum_{q=0}^{p} c_q N(\alpha)^q$, where c_q are the polynomial coefficients determined via OLS fitting to the response data sets. For the Hubbard U, the response data comprises k occupancy-Hxc potential pairs $\left(N(\alpha_i), V_{\mathrm{Hxc}_i}^{+}\right)$, where $V_{\mathrm{Hxc}_i}^{+} = V_{\mathrm{Hxc}_i}^{\uparrow} + V_{\mathrm{Hxc}_i}^{\downarrow}$. The Hubbard U is then calculated following Eq. (5) through the evaluation of the derivative of the regression at $\alpha = 0$,

$$\mathrm{U} = \frac{1}{2}\sum_{q=0}^{p} q\ c_q\ N(0)^{q-1}\ . \tag{6}$$

The Hubbard U is thus a multivariate function with respect to the fitted coefficients. It is important (even numerically imperative) to assert that the fitted polynomial coefficients are covariate, meaning their uncertainties do not vary independently. Therefore, the error on the MT LR Hubbard U is found to be [10]

$$\sigma_{\mathrm{U}}^{(p)} = \frac{1}{2}\sqrt{\sum_{q=0}^{p}\sum_{r=0}^{p} q\ r\ C_{q+1,r+1}\ N(0)^{q+r-2}}\ , \tag{7}$$

where we use the unbiased standard deviation and the $k \times (p+1)$ design matrix $\mathbf{X}$ with elements $X_{i,q+1} = N(\alpha_i)^q$ to compute the covariance matrix $\mathbf{C}$, featuring matrix elements

$$C_{q+1,r+1} = \frac{\sum_{i=1}^{m}\left[V_{\mathrm{Hxc}_i}^{+} - f^{(p)}\left(N(\alpha_i)\right)\right]^2}{m-p-1}\left(\mathbf{X}^{\top}\mathbf{X}\right)^{-1}_{q+1,r+1}\ . \tag{8}$$

Uncertainty on the Hund's J may be calculated by replacing all instances of α with β, total occupancy $N(\alpha)$ with subspace magnetization $M(\beta)$, and $V_{\mathrm{Hxc}_i}^{+}$ with $V_{\mathrm{Hxc}_i}^{-} = V_{\mathrm{Hxc}_i}^{\uparrow} - V_{\mathrm{Hxc}_i}^{\downarrow}$.

In practice, the uncertainty incorporating the covariance between polynomial coefficients, described by Eqs. (7) and (8), we found to be reasonably identical to the regression error obtained by shifting the $N(\alpha_i)$ values by $-N(0)$ and overlooking the matter of covariance [16]. In this case, one may evaluate the derivative about a zero-perturbation axis, rendering the Hubbard parameter a singly-variate function of c_1 and the uncertainty of which may be evaluated with an analogous formula to that in Eq. (2).

Appendix B: Computational Details

We analyze two LR datasets featuring DFT simulations of antiferromagnetically ordered NiO, aggregated using what is now the `lrUJ` post-processing utility [17]. Both datasets comprise α perturbations (for the Hubbard U), the strengths of which were evenly spaced between -2.0 and 2.0 eV, conducted on the Ni $3d$ subspaces using the Perdew-Burke-Ernzerhof (PBE) GGA [18] XC functional and the Jollet-Torrent-Holzwarth PBE projector-augmented wave (PAW) pseudopotentials (Version 1.1) [19]. Calculations were conducted in ABINIT, where `dmatpuopt`$= 2$ and the total energy relaxed to 1×10^{-7} Ha.

The 5-point dataset features calculations performed on a 64-atom supercell of Type I antiferromagnetic (AF_I) NiO with an experimental lattice parameter of 4.168 Å. In addition to the aforementioned settings, `diemix` was set to 0.15 to facilitate convergence, the plane-wave cutoff energy was set to 33.1 Ha (44.1 Ha for the PAW augmentation sphere), and Monkhorst-Pack grid k-point sampling set at $3 \times 3 \times 3$. The nearly identical dataset performed on AF_{II} NiO differed in magnetic ordering, experimental lattice parameter (4.170 Å), and MP k-point sampling ($5 \times 5 \times 5$). Please consult [7] and references for additional computational details.

The 39-point dataset features calculations performed on a 4-atom unit cell of AF_{II} NiO with a lattice parameter of 4.191 Å. Here, `diemix` was set to its default of 0.45, the plane-wave cutoff energy was set to 22.0 Ha (42.3 Ha for the PAW augmentation sphere), and Monkhorst-Pack grid k-point sampling set at $10 \times 10 \times 10$.