



CHORUS

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

Decoding structure-spectrum relationships with physically organized latent spaces

Zhu Liang, Matthew R. Carbone, Wei Chen, Fanchen Meng, Eli Stavitski, Deyu Lu, Mark S. Hybertsen, and Xiaohui Qu

Phys. Rev. Materials **7**, 053802 — Published 16 May 2023

DOI: [10.1103/PhysRevMaterials.7.053802](https://doi.org/10.1103/PhysRevMaterials.7.053802)

Decoding Structure-Spectrum Relationships with Physically Organized Latent Spaces

Zhu Liang,¹ Matthew R. Carbone,² Wei Chen,¹ Fanchen Meng,¹ Eli Stavitski,³ Deyu Lu,^{1,*} Mark S. Hybertsen,^{1,*} and Xiaohui Qu^{1,*}

¹Center for Functional Nanomaterials, Brookhaven National Laboratory, Upton, New York 11973, USA

²Computational Science Initiative, Brookhaven National Laboratory, Upton, New York 11973, USA

³National Synchrotron Light Source II, Brookhaven National Laboratory, Upton, New York 11973, USA

* dlu@bnl.gov, mhyberts@bnl.gov, xiaqu@bnl.gov

KEYWORDS

structure-spectrum relationship, interpretable machine learning, interpretable latent space, X-ray absorption spectroscopy, X-ray absorption near-edge structure, autoencoder, deep learning

ABSTRACT

A new semi-supervised machine learning method for the discovery of structure-spectrum relationships is developed and then demonstrated using the specific example of interpreting X-ray absorption near-edge structure (XANES) spectra. This method constructs a one-to-one mapping between individual structure descriptors and spectral trends. Specifically, an adversarial autoencoder is augmented with a novel “rank constraint” (RankAAE). The RankAAE methodology produces a continuous and interpretable latent space, where each dimension can track an individual structure descriptor. As a part of this process, the model provides a robust and quantitative measure of the structure-spectrum relationship by decoupling intertwined spectral contributions from multiple structural characteristics. This makes it ideal for spectral interpretation and the discovery of new descriptors. The capability of this procedure is showcased by considering five local structure descriptors and a database of over fifty thousand simulated XANES spectra across eight first-row transition metal oxide families. The resulting structure-spectrum relationships not only reproduce known trends in the literature, but also reveal unintuitive ones that are visually indiscernible in large data sets. The results suggest that the RankAAE methodology has great potential to assist researchers to interpret complex scientific data, test physical hypotheses, and reveal new patterns that extend scientific insight.

I. INTRODUCTION

Structure-property relationships in materials encode fundamental physical knowledge and enable a constructive design process to meet functional goals and guide materials discovery [1–3]. Considering materials beyond the simplest of molecules or crystals, e.g. amorphous materials, thin films, and nano catalysts, the full description of the atomic-scale structure generally involves many separate quantities. Similarly, a collection of properties can also be a complex data set. Consequently, the discovery and representation of structure-property relationships pose significant challenges in the raw form of a “many-to-many” map. Traditionally, expert intuition has been used

to identify a few simple structure descriptors that can be related to specific trends in properties through physical arguments and experimental probes [4]. The emergence of high-throughput experiments and expanding databases of computed materials properties invites the application of new, data-driven approaches to discover deeper, more comprehensive structure-property relationships [1–3]. Significant progress has been made in applying machine learning methods to physical datasets in recent years [1–3,5–7]. However, the development of interpretable models [8–12], highly desirable for use by physical science domain experts, remains an ongoing challenge in the field.

Mapping the trends in a complex dataset, such as a collection of physical spectra measured for a set of materials, to a reduced dimensional space is a common and productive first step. For example, an autoencoder [13,14] (AE) can map each input spectrum to a point in a latent space of low dimensionality while simultaneously training a decoder to perform the inverse mapping of a point in the latent space back to a spectrum. However, the latent space variables do not inherently have a physical interpretation [9,10,15–19]. Specifically for the structure-property relationship problem, the data-driven discovery of correlative structure descriptors remains largely unsolved.

X-ray absorption spectroscopy (XAS) illustrates the complexity of structure-spectrum relationships. XAS is a widely used technique for materials characterization [20,21] due to its element specificity and its sensitivity to the local chemical environment of the absorbing atom [22]. Furthermore, modern synchrotron facilities enable high-throughput experimentation and in situ materials discovery methodologies [23]. In XAS experiments, a core electron is excited by an incident photon to empty states. The X-ray absorption spectrum exhibits steps, called edges, at clearly distinguishable energies where the absorption coefficient rises sharply. These are classified into K-, L-, and M-edges corresponding to $n=1, 2,$ and $3,$ where n is the principal quantum number of the core electron. For a given absorbing element, the XAS spectrum is typically divided into two regions relative to the edge. The region extending up to roughly 50 eV above the edge is referred to as the X-ray Absorption Near Edge Structure (XANES), and beyond that is the Extended X-ray Absorption Fine Structure (EXAFS) [20].

XANES and EXAFS encode different types of information. EXAFS contains information about the radial distribution of first shell neighbors, and potentially second and third shell neighbors if the signal-to-noise ratio is high enough. The analysis of EXAFS is well developed, with robust approaches available for extracting structural information [24,25]. In the framework of multiple scattering theory [26] (see Figure 1a), the EXAFS signal is dominated by the sum of direct back-scattering terms. In contrast, XANES depends on the interference of a more diverse set of interactions, such as the three-body path illustrated in Figure 1a. Furthermore, it contains rich information about the electronic structure, specifically dipole-allowed transitions to low-lying empty states. As a result, XANES spectral features can be correlated with local physical and structure descriptors, such as oxidation state (OS), coordination number (CN), and local symmetry [20,21,26,27]. However, XANES is more difficult to analyze than EXAFS. An in-depth analysis typically requires details of the electronic states of the material as well as an accurate physical model of the many-body core-hole effects [26–29]. Because of this inherent complexity, it is challenging to unravel the structure-spectrum relationships in the analysis of XANES.

Structure-spectrum relationships enable the extraction of structural information from XANES spectra measured on new materials, a core application of X-ray measurement in science and engineering. Consequently, the development of methods to infer structure descriptors from spectra

is a central research problem in XANES analysis [21,30–35]. In contrast to the well-posed problem of computing a spectrum from a given material structure, this is an “inverse problem.” The inverse problem may not always have a unique solution. Multiple structures may be consistent with a measured spectrum. Solving the inverse problem relies on identifying a set of robust structure descriptors that correlate to distinct spectral features or fingerprints. The search for such relationships has inspired decades of research in X-ray spectroscopy. In 1920, Bergengren discovered that the edge position strongly correlates with the OS of the absorbing atom (see Figure 1b) [30]. Since then, OS has been a common focus of the K-edge XANES analysis [21,32,33,36]. Other important structure descriptors include local symmetry (e.g., tetrahedral versus octahedral) and CN of the cation, which are correlated with the position and intensity of the pre-edge peak in early 3d transition metal oxides (Figure 1c) [21,31,35,37]. However, this empirical relationship shows strong system dependence and does not apply to late 3d transition metals with vanishing pre-edge features, as illustrated in Figure 1d [37–39]. Beyond the first coordination shell, it is significantly more difficult to identify useful structure descriptors that have a significant contribution to the spectrum features.

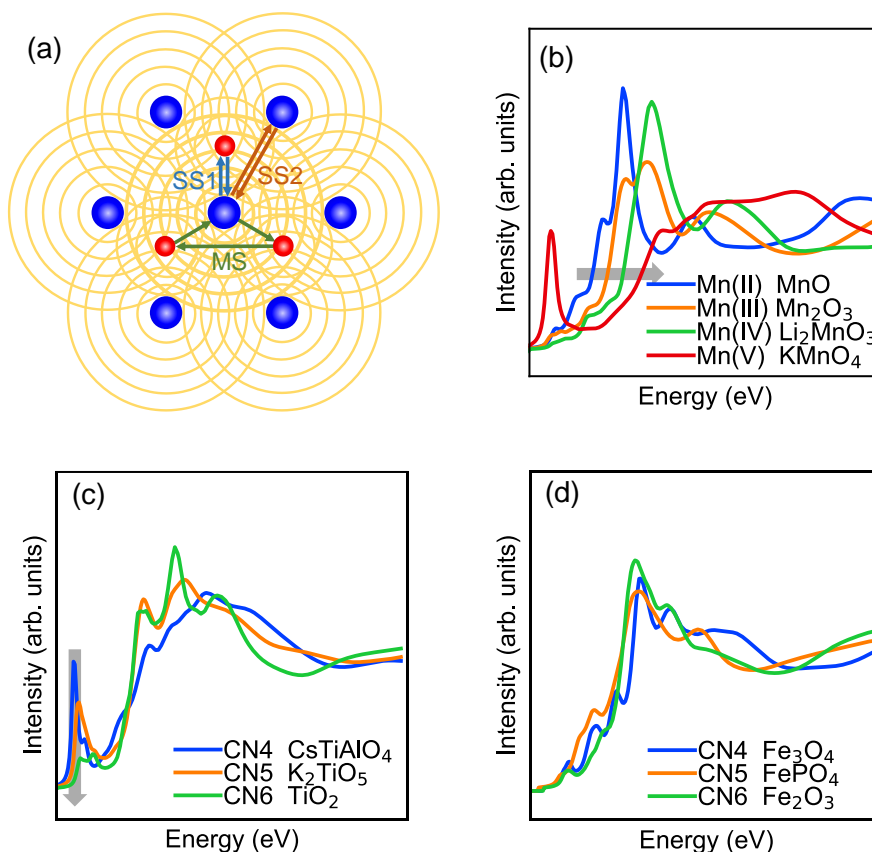


Figure 1. Real space scattering picture of X-ray absorption and exemplary spectral trends in K-edge XANES. (a) The final state electron amplitude on the central atom includes the sum of scattering contributions from all neighboring atoms. Scattering processes include: SS1 – single scattering from the first coordination shell, SS2 – single scattering from the second coordination shell, and MS – an exemplary multiple scattering path. (b) Simulated Mn K-edge XANES for a series of crystals in which the oxidation state (OS) of Mn varies showing the shift of the main edge (gray arrow). (c) Simulated Ti K-edge XANES for three oxide crystals with varying titanium coordination numbers (CN) showing the change in pre-edge feature intensity (gray arrow). CN4, CN5, and CN6 designate 4, 5, and 6-coordinated motifs. (d) Same for Fe K-edge XANES but showing no simple trend.

Despite the utility of existing structure descriptors (e.g., OS and CN), there is a strong demand to discover or engineer a sufficiently complete set of structure descriptors to support local structure prediction in complex materials spaces. The significance of a structure descriptor strongly depends on the underlying materials and the physical properties of interest. For example, the pre-peak features in transition metal K-edge XANES are sensitive to the distortion of the cation octahedral cage caused by substrate [40] or pressure [41]. Based on studies of Ti K-edge XANES and Li K-edge electron energy loss spectra, local distortions are identified as structure descriptors that play an important role in understanding the phase transformation and the fast lithium ion transport in lithium titanate [42–44]. To capture the spectral trends of nanostructures, bond-length-based structure descriptors have been investigated, which include metal-metal bond lengths in metallic (Pd K-edge) [18] and bimetallic nanoclusters (Pd K-edge and Au L₃-edge) [45], average Fe-O bond length in Fe oxide clusters (Fe K-edge) [34], and Co-C and Co-O bond lengths in a Co single atom catalyst complex (Co K-edge) [46]. In addition, the chemical composition (e.g., hydrogen content) can also serve as a good descriptor for Pd nanoparticle catalysts exposed to H₂ [18]. These examples illustrate the material specificity of the problem. Despite extensive studies [32,33,47,48], a robust and widely applicable approach for identifying structure descriptors still does not exist.

Intuitively, one can label a spectral dataset with various structure descriptors and use visual inspection to search for qualitative trends. This empirical approach can succeed when there is an obvious trend, such as the correlation of edge shift to OS in Mn K-edge XANES (Figure 1b) or the correlation of pre-edge intensity to CN in Ti K-edge XANES (Figure 1c). However, this approach fails when the spectral trend is too complex to identify by visual inspection, e.g., the trend of CN in Fe K-edge XANES (Figure 1d). Nonetheless, in the latter situation, there is great potential to use data analytics tools to extract statistically meaningful trends [31,33,38,48–55].

In this study, we focus on methods capable of learning a latent space that simultaneously represents spectra and correlates with structure descriptors. The resulting map is an interpretable, multi-dimensional structure-spectrum relationship. We identify key technical challenges that must be met to achieve this goal and develop new data analytics tools to address them. Specifically, we build on the adversarial autoencoder (AAE) [56] by adding a new rank constraint that drives each latent space variable to track a target structure descriptor. Taken together, our rank-constrained adversarial autoencoder (RankAAE) method captures spectral variations along latent space dimensions and correlates them with physically interpretable structure descriptors, as illustrated in Figure 2. We show that the method works across a multidimensional latent space incorporating a set of structure descriptors in a single training procedure.




Figure 2. Conceptual diagram illustrating the key features of the rank-constrained adversarial autoencoder (RankAAE) methodology. The encoder maps spectra to points in a latent space with axes that correlate with target physical structure descriptors due to the rank constraint, such as coordination and oxidation. The decoder maps points in that latent space back to physical spectra. Pathways in the latent space represent trends in structure descriptors that map to trends in spectral signatures, as illustrated for CN and OS. Taken together, the result is a multi-dimensional structure-spectrum model.

To demonstrate the RankAAE method, we apply it to transition metal cation K-edge XANES from eight $3d$ transition metal oxide families (Ti-Cu). A large database of simulated spectra is calculated for crystal structures drawn from the Materials Project database, which exhibits diverse local cation chemical environments. This represents a broad coverage of materials compositions, including binary, ternary and quaternary oxides. The database includes over fifty thousand individual cation spectra. With this database, we use the RankAAE methodology to build structure-spectrum relationships and characterize the effectiveness of the method. We show how it can be used by a domain expert to explore candidate structure descriptors and their corresponding spectral trends. In particular, this method can be used to compare the performance of different combinations of structure descriptors in multi-dimensional models for structure-spectrum relationships. For the transition metal oxides under study, the RankAAE method enabled us to both recover historical structure-spectrum relationships [21,30,31,35–37,57] and to reveal new ones hidden in the data.

XANES, the specific example chosen in this work as a case study, is a spectroscopy technique that is a widely used tool in materials science and its interpretation has proven to be non-trivial. The spectra result from complex physical processes. The relationship to material structures can be convoluted and difficult to establish. The present work demonstrates a new tool that lowers the barrier to the discovery of structure-spectrum relationships with data-driven techniques. The complexity of the XANES problem is representative of many physical science datasets. The application of the RankAAE methodology to other spectroscopy problems is straightforward and the generality of the methodology will inform data driven techniques for other classes of physical science datasets as well.

We henceforth summarize the structure of our manuscript. Latent space methods are briefly reviewed in Section II. The structure-spectrum database and key characteristics are described in Section III. Section IV presents the main results demonstrating the use of the RankAAE method. The results are discussed in Section V, including the physical interpretation of the emerging trends. Conclusions appear in Section VI. Technical details of the methodology are described in Appendix A.1 and A.2.

II. LATENT SPACE METHODS

Dimensionality reduction techniques pertinent to the structure-spectrum problem are briefly described here. In particular, we place our RankAAE method in the context of previous work.

One example of a simple, linear method commonly used to extract trends in data is principal component analysis (PCA). Liu *et al.* applied PCA to the Cu K-edge of Cu_xPd_y bimetallic nanoparticles and identified patterns correlated with Cu_xPd_y cluster types [58]. Carbone *et al.* performed PCA on K-edge XANES of eight $3d$ transition metal families and identified clear patterns, where the data points in the reduced dimensional space are distributed into identifiable clusters ordered according to the label of the absorbing cation CN=4, 5, and 6 [38]. Similar patterns are shown here for Fe OS and CN in Figure 3(a, b), where the simulated spectra were drawn from the database used in the present study (see Appendix A.3). Fe OS exhibits a meaningful pattern in the reduced-dimensional space, where the gradient of the OS is roughly along the horizontal direction. Similarly, the ordered PCA pattern of Fe CN is roughly along the vertical direction, in stark contrast to the vague trend in the raw spectra of Figure 1d. The PCA decomposition of Fe K-edge XANES spectra in Figure 3 provides strong evidence that OS and CN are good structure descriptors. However, the variation of OS or CN is highly non-linear with respect to the PCA axes, and the patterns of the two descriptors are thus intertwined, as illustrated schematically in Figure

3c (left). This example illustrates the limitations of standard linear dimensionality reduction methods, such as PCA, for quantitative structure-spectrum mapping.

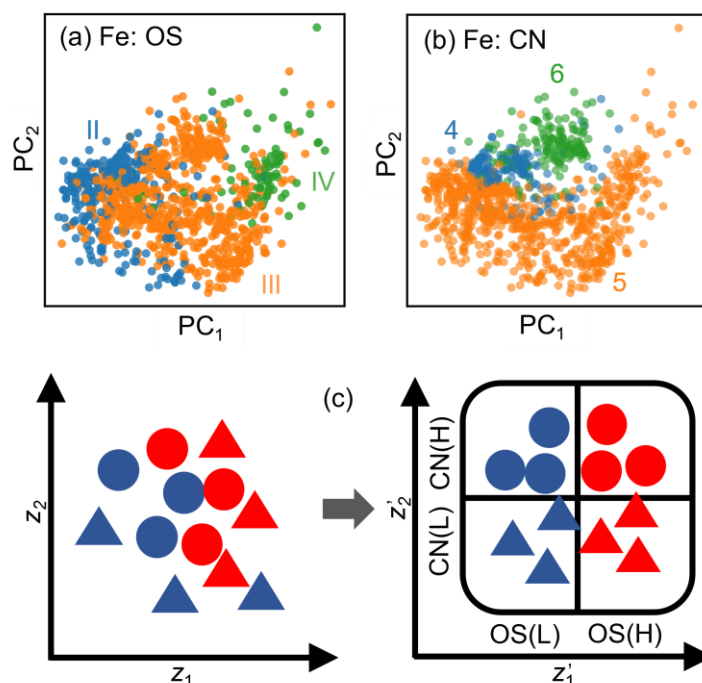


Figure 3. Dimensionality reduction with correlation to structure descriptors. (a, b) Principal component analysis of simulated Fe K-edge XANES spectra. Data points are colored according to the indicated structure descriptors: (a) OS of Fe; (b) CN of Fe. (c) Schematic illustration of feature entanglement in the spectrum latent space $\{z_1, z_2\}$ (left) and the goal of transforming the latent space into $\{z'_1, z'_2\}$ (right), where z'_1 and z'_2 align with structure descriptors OS and CN, respectively. OS is represented by color (blue: low OS; red: high OS), and CN is represented by shape (triangle: low CN; circle: high CN).

A more promising dimensionality reduction approach uses autoencoder (AE)-based methods [13,14]. An AE generally consists of two neural networks trained in tandem to approximate the identity function: an encoder that compresses vector data sets to a lower dimensional latent space and a decoder that reverses this mapping (see the blue dashed box in Figure 4). An AE can perform sophisticated non-linear data compression and capture statistically relevant information in the latent space. For example, Routh *et al.* trained an autoencoder to transform simulated Pd K-edge XANES spectra of small Pd clusters to a latent representation and revealed a strong correlation between structure descriptors (CN, interatomic distance, and hydrogen content) and latent space variables [18]. However, in studies that require continuous sampling of the latent space, the mapping created by a basic AE can be problematic. For example, when applied to image reconstruction problems, some regions of the latent space do not decode to sensible images [59]. In a chemical science example, the latent space of an AE trained on SMILES encodings of molecules [60] exhibited “dead areas” that did not decode to valid molecules [17]. In the present work, the spectrum reconstruction from the decoder can yield unphysical spectra (See Figure S1a in the Supplemental Material [61]). These failures are due to the lack of regularization of the latent space. Some points in the latent space may not correspond to regions in the neighborhood of any previously seen training data, causing the decoder to produce unpredictable

results. Variational (VAEs) [10,17,62], Wasserstein (WAE) [9] and adversarial autoencoders (AAEs, green dashed box in Figure 4) [56,63,64] tackle this problem by regularizing the latent space during training, forcing the training data coverage of the latent space to be more complete. Although different in technical detail, all result in models capable of performing robust data generation (i.e., generative models). The latent space can then be sampled continuously, with decoded signals that remain valid in the target application, e.g., yielding physical spectra. For example, a VAE has been used to analyze spectral functions [19], and VAEs, WAE and AAEs all have been used as generative models to search for new molecules/materials from the latent space [9,17,63,64].

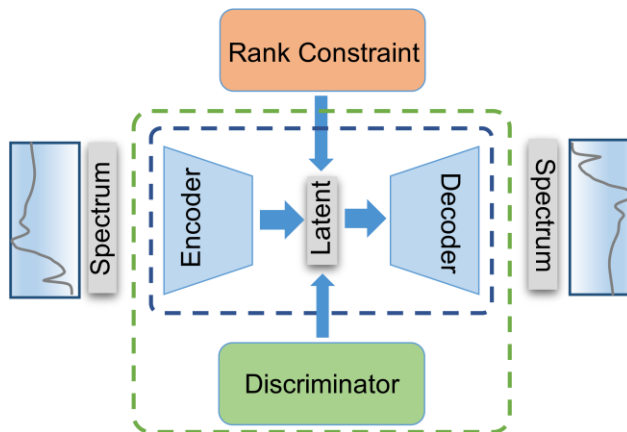


Figure 4. Network structure of the RankAAE method. A standard autoencoder (blue dashed rectangle) creates a latent space representation of the data that faithfully reproduces the spectra in the training dataset. An adversarial autoencoder (green dashed rectangle) introduces a discriminator on top of the regular autoencoder that regularizes the latent space. RankAAE further adds a rank constraint (orange) to the latent space to establish correlations between specific latent space dimensions and target physical descriptors.

Robust, data-driven methodologies to uncover structure-spectrum relationships and discover new descriptors must address three key technical challenges:

1. *Spectrum validity*: a data-driven method is needed to construct a set of latent variables that serves as a proxy for the spectrum. In practice, the latent space needs to be regularized to ensure that all data points in the latent space correspond to physical spectra.
2. *Structure mapping*: the method needs to establish quantitative mappings between the latent space and structure descriptors, such that the statistical importance of various descriptors can be assessed and compared.
3. *Feature disentanglement*: the method must disentangle spectral contributions driven by different underlying structural and chemical characteristics.

For example, in the case of XANES, the spectral variation reflects the net effects from multiple sources that are often convoluted in overlapping energy ranges. Even with prior knowledge from domain-specific expertise, it is currently impossible to disentangle the impact of multiple structure descriptors. To achieve disentanglement, the latent space needs to be further organized, such that each latent space dimension aligns with a specific structure descriptor. Currently, none of the off-the-shelf data analytics tools simultaneously satisfies these three requirements of *spectrum validity*, *structure mapping*, and *feature disentanglement*.

New methods need to be developed to simultaneously address all three challenges. We focus on VAE/AAE methods. Because they are generative models in nature, they already satisfy the *spectrum validity* requirement. Joint training of a property prediction network with a VAE can address the *structure mapping* requirement [17], however it does not solve the *feature disentanglement* problem. Specifically, even with the addition of joint training, each latent variable does not represent a target property directly since it can be entangled with multiple physical descriptors. To construct a direct latent variable-structure descriptor mapping, the latent space of a VAE/AAE needs to be reorganized to align each dimension with a structure descriptor. Figure 3c (right) depicts this idea: the goal is to have the latent variable z'_1 only depend on OS (symbol color) monotonically, while z'_2 only depends on CN (symbol shape) monotonically. To accomplish this, additional constraints need to be engineered and applied to a VAE or AAE.

The RankAAE method, developed in this work, accomplishes this goal with a novel “rank” constraint specifically designed to organize the latent space in the way depicted in Figure 3c. The implementation balances the goal of aligning each latent space variable with a target structure descriptor while minimizing the impact on the other, statistical characteristics of the latent space representation learned with the AAE. The technical details are described in Appendix A.1 and A.2. Applied to a multi-dimensional latent space, the interplay of different structure descriptors is disentangled and the RankAAE method associates each dimension of the latent space with a well-defined structure descriptor.

III. DATABASE PREPARATION AND CHARACTERIZATION

For our database of transition metal oxides, we simulate the XANES spectra for each symmetrically non-equivalent transition metal absorber using the multiple-scattering code, FEFF9 [26,65]. Details of the database development and computational methods are given in Appendix A.3.




Figure 5. Structure and XANES spectral descriptors. (a) Schematic local structure around the central cation emphasizing the cation-oxygen network. Shaded areas are first coordination shells of cation (light red triangle) and oxygen (light blue triangle). (b) Schematic spectral descriptors applicable to typical cation K-edge XANES spectra in oxides: E_{edge} , I_{pre} , I_{main} , C_{pit} , E_{main} , E_{post} and I_{post} .

Since XANES cation K-edge spectra probe local chemical environments, we investigate the structure descriptors that encode them. For the transition metal oxides, the cation-oxygen network characteristics are of particular importance (Figure 5a), although most crystal structures in the database also incorporate counterions that influence key chemical and structural attributes of the local transition metal and its environment. For simplicity, we restrict our attention to sites where the nearest neighbor shell contains only oxygen. In addition to the two well-known descriptors

(OS and CN), we consider several other descriptors directed at capturing distortions in the nearest-neighbor shell of the absorbing cation and the influence of the second nearest-neighbor shell. The structure descriptors used in machine learning potentials and structure-property relationship studies provide guidance [66,67]. The considered descriptors are detailed in Appendix A.4, but we outline them here briefly.

- 1) The average of the coordination numbers for the nearest-neighbor oxygen atoms (OCN),
- 2) The second-nearest-neighbor coordination number (CN-2),
- 3) The spread in the cation-oxygen bond lengths expressed as the nearest-neighbor radial standard deviation (NNRS),
- 4) The spread in the cation centered bond angles expressed as the nearest-neighbor angular standard deviation (NNAS),
- 5) The minimum oxygen-oxygen distance on the edges of the nearest neighbor polyhedron (MOOD), and
- 6) The point group symmetry order (PGSO).

To facilitate the discussion of structure-spectrum relationships, we adopt a set of spectral descriptors that capture the main spectral characteristics seen in a typical transition metal oxide XANES K-edge spectrum (Figure 5b). Described in more detail in Appendix A.5, these basic metrics will be referenced throughout our work. We briefly outline them below:

- 1) The position of the absorption edge (E_{edge}).
- 2) The pre-peak intensity (I_{pre}).
- 3) The main peak position (E_{main}).
- 4) The main peak intensity (I_{main}).
- 5) The post edge curvature (C_{pit}).
- 6) The second peak position (E_{post}).
- 7) The absorption intensity at E_{post} (I_{post}).

IV. RESULTS

We present our results in two parts. First, we illustrate the use of the RankAAE method and our validation of its performance for one family of materials, Vanadium (V) oxides. Then, we illustrate the performance of the method across the full set of transition metal oxides considered in this study.

A. Application and validation of RankAAE for vanadium oxides

The full scope of the V K-edge XANES data across the present vanadium oxide database is shown in Figure 6a. To illustrate the structure-spectrum relationships in this raw data, the spectra are colored according to the values of five structure descriptors in sequence from the top of the figure: OS, CN, OCN, NNRS, and MOOD, respectively. In particular, we note that OS and CN are categorical variables because their valid values are limited to a few integer numbers. Visual inspection of the examples in Figure 6a shows color concentrations indicative of correlations between structure descriptors and spectral features, as one expects on physical grounds. For example, the position of the edge and the main peak (E_{edge} and E_{main}) shift with OS, in line with the trend observed in experimental spectra [21,31,35,37]. Pre-edge peak intensity (I_{pre}) and main peak position and intensity (E_{main} and I_{main}) change with CN, but E_{main} and I_{main} are also affected by OCN and MOOD. Higher energy features above the main peak show similar characteristics.

This illustrates the entanglement of the contributions from different structure descriptors to trends in spectral features. This diverse dataset exemplifies the challenges that domain experts face when trying to expand the scope of structure descriptors in the analysis. While the overall trends in spectral features for OS and CN follow prior domain experience and physical models, the trends for additional spectral descriptors are too ambiguous or complicated to draw clear conclusions. A specific structure descriptor correlates with spectral features in several locations across the spectrum and each of those features can be affected by multiple structure descriptors. Taken together trends are obscured and physical interpretation is non-trivial.

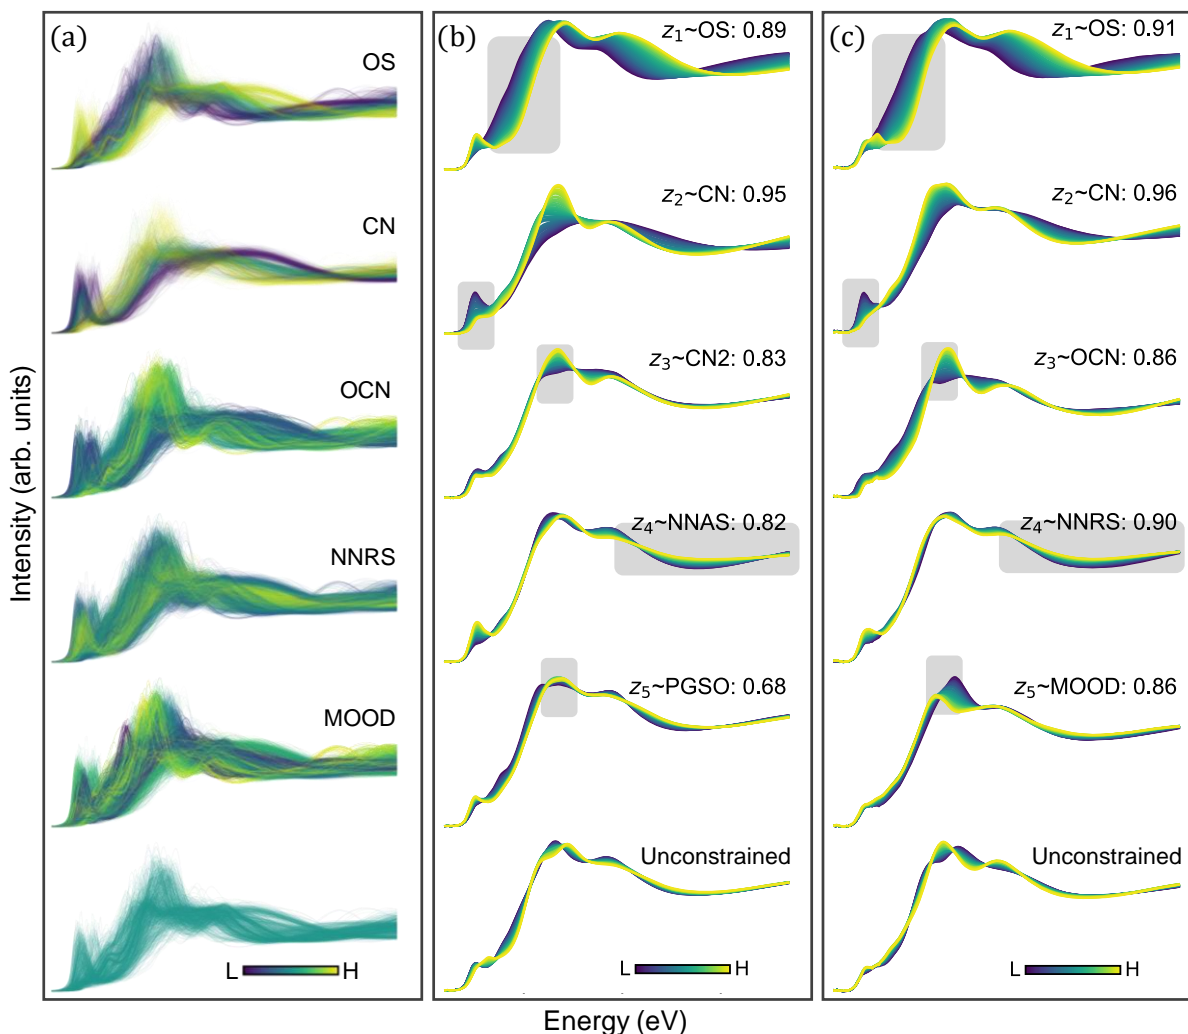


Figure 6. Correlating structure descriptors to spectral trends for V K-edge XANES from vanadium oxides. (a) Simulated spectra of vanadium oxides colored according to the value of specific structure descriptors: OS, CN, OCN, NNRS, MOOD, and no designation. (b, c) Reconstructed spectra from trained RankAAE models. As described in the text, each subplot maps the spectral trend associated with one of the six dimensions in the latent space. Spectra are colored by the value of the corresponding latent variable. Five of the latent space variables are constrained to correlate with target structure descriptors while the sixth variable is unconstrained. In (b), an initial set of structure descriptors is chosen for training: OS, CN, CN-2, NNAS, and PGSO. In (c) a final set of structure descriptors is chosen: OS, CN, OCN, NNRS, and MOOD. For each descriptor, the degree of correlation to the corresponding latent space variable is quantified by an F1 score (for OS and CN) or the Spearman rank correlation coefficient [68] (for the other descriptors), annotated next to each plot. The shaded area marks the location of the primary spectral variation characteristic of the spectral trend.

We illustrate the use of the RankAAE method (Figure 4, Appendix A.1 and A.2) through application to this vanadium oxide dataset. The RankAAE model is trained to map the spectra in the database to a latent space (center of Figure 4) of chosen dimensionality N through an autoencoding procedure. For most of the results presented in this study we take $N = 6$. As a special type of AAE, the latent space resulting from the trained RankAAE model is regularized and each reconstruction maps to a physical spectrum (Figure S1b in the Supplemental Material [61]). In contrast, a standard AE model can yield wider variance and unphysical features in the spectra, such as excessive noise and even dramatic, sharp spikes (Figure S1a in the Supplemental Material [61]). In addition, the rank constraint is enforced for selected latent space variables to align with chosen structure descriptors, as described in Appendix A.2. Unconstrained dimensions represent residual characteristics of the spectral data set beyond the constrained dimensions. Training is repeated from scratch for each set of chosen structure descriptors. The details of dataset splitting (training, validation, and test) are listed in Table S1.

The regularized and aligned latent space resulting from the RankAAE method allows us to present spectral trends in an easily interpretable fashion. Reconstructed spectra are mapped from the latent space with the trained decoder. Once the model is trained, each point in the latent space decodes to a spectrum. In addition, due to our training procedure, each point should correlate to a set of structure descriptor values according to the imposed constraints. If the method is successful, the latent space will encode structure-spectrum relationships. By tracing a path through the latent space, we sample that mapping, tracking the correlated structure descriptor values together with the associated reconstructed spectra. In the simplest version of this picture, we can track the evolution of spectra along each axis of the latent space while holding the other latent space values constant, e.g., equal to zero (Figure S2a in the Supplemental Material [61]). With the correlation to target structure descriptors, this can isolate specific spectral features that correlate to that structure descriptor.

Given the complexity of the dataset and the goal of quantifying the structure-spectrum relationship in a multidimensional descriptor space, we want to sample the latent space more systematically, not just along isolated, one-dimensional paths. To this end, for each latent variable z_i , we average over spectra corresponding to different values of $\{z_j\}_{j \neq i}$. In this way, the evolution of spectra along direction z_i is statistically representative of the full dataset. Appendix A.6 describes this averaging procedure. Figures S2a and S2b in the Supplemental Material [61] compare the sampling along each isolated axis with the averaging procedure. The two approaches result in consistent trends. We adopt the averaging procedure in this work for its statistical rigor.

For the vanadium oxide dataset, Figures 6b and 6c show a series of trends for two trained RankAAE models respectively, each with five constrained latent space dimensions and one unconstrained dimension. Each subpanel shows the reconstructed spectral evolution along one dimension of the latent space. The smooth variation in spectra highlights the success in regularizing the latent space. The distinctive variations from panel to panel provide the input for interpreting the trends in terms of the target structure descriptors and assessing their relative utility.

It is not known *a priori* which structure descriptors capture the most significant spectral variation given a diverse dataset of materials structures. Figure 6b shows our “initial guess” for these structure descriptors based on prior domain knowledge. We evaluate the results according to several criteria. First, we monitor the quantitative correlation between the latent space variable and

the target descriptor as detailed below. Second, we look for emerging spectral trends in distinct regions of the XANES spectrum. To be most useful for applications, those trends should be distinguishable. Finally, we monitor the amplitude of the spectral variation captured by the unconstrained latent space variable z_6 . In essence, the last metric captures the spectral variance outside the scope of that driven by the target structure descriptors. Reducing it should indicate improvements in the completeness of the descriptor set.

For the results in Figure 6b, we see that the target descriptors are captured with generally high correlation (inset values in each panel). Furthermore, for each descriptor, a primary trend in the spectra can be identified (shaded area). In addition, there are other spectral trends associated with each descriptor across the energy range. For example, CN is correlated with both pre-edge intensity and main peak intensity. Thus, while it is convenient to isolate localized spectral features for discussion, each structure descriptor is associated with extended fingerprints with contributions from spectral features across the full energy range. We also note an overlap in the trends. For example, OS, CN, and NNAS all correlate to pre-edge intensity. Finally, the unconstrained dimension exhibits a relatively small spectral variation.

The process of refining the choices of structure descriptors is iterative and still requires a “human in the loop.” As an example, we retain the well-known OS and CN descriptors that clearly show a strong, systematic spectral trend and consider alternatives for other structure descriptors. Figure S3 in the Supplemental Material [61] shows the change in spectral trend as each of those structure descriptors is replaced one by one. The OCN descriptor improves over CN-2 with a somewhat higher correlation to the latent space variable and qualitatively supports a stronger spectral trend in the shape of the primary peak. Introducing NNRS retains the correlation to the pit curvature while reducing the impact on pre-edge peak intensity. The MOOD descriptor exhibits a much-improved correlation to the latent space variable compared to PSGO (0.86 versus 0.68). In this example, the spectral variation attributable to the unconstrained latent space dimension remains relatively small. The spectral trend for the final structure descriptor set is presented in Figure 6c. In comparing the trends with the initial result in Figure 6b, the interaction between the descriptor choices can be seen. The pre-edge intensity trend for OS has been suppressed, and the spectral trend around the main peak shape for CN has been concentrated. Also, some of the overlaps in the spectral trends have been reduced. In particular, the pre-edge peak intensity is now more specifically correlated to the CN descriptor, and to a smaller degree to NNRS. This set is likely not unique, but it captures most of the spectral variance in the vanadium oxide database.

A different perspective on descriptor development is to consider how the model evolves as descriptors are added one by one. This is illustrated in Figure S4 in the Supplemental Material [61] for the five descriptors used in Figure 6c. The quantitative correlation between each structure descriptor and the corresponding latent space variable is stable. The qualitative spectral trends are also similar, although we observe a clear sharpening of them. For example, the initial pre-edge feature in the OS trend systematically reduces to essentially zero. The CN trend consolidates into clear pre-edge and main peak features. At the same time, the OCN trend is capturing main peak shape changes. Finally, we also see the stepwise reduction in the amplitude of the unconstrained latent variables as additional structure descriptors are added. This strongly supports the notion that each new structure descriptor is extending the model to capture additional variation in the spectra. It also illustrates how the RankAAE approach overcomes the complexity of intertwined spectra trends in the raw spectral data (Figure 6a). As the set of structure descriptors incorporated into the trained model is modified, a set of distinguishable spectra trends emerges.

From the fully trained models, we can identify regions of the spectra with “primary spectral trends,” each one linked to a specific latent variable z_i and the associated structure descriptor (Figures 6b and 6c). Other significant concentrations of spectral variations constitute “secondary spectral trends.” Together, they represent a spectral fingerprint. For the specific model for structure-spectrum relationships shown in Figure 6c, we identify the following trends:

Oxidation state (OS): The edge position shifts to higher energy with the increase of atomic charge (Figure 6c, z_1 and Figure 5b, E_{edge}). In addition, there is a prominent secondary spectral trend: the higher energy portion of the spectrum also shifts horizontally (Figure 5b, E_{post}).

Cation coordination number (CN): The pre-edge peak intensity decreases sharply with the increase of CN (Figure 6c, z_2 and Figure 5b, I_{pre}). In addition, the intensity change after the main peak constitutes an identifiable secondary spectral trend (Figure 6b, I_{post}).

Oxygen coordination number (OCN): The intensity of the main peak increases with the increase of OCN (Figure 6c, z_3 and Figure 5b, I_{main}). Overall, OCN caused smaller but more focused changes than CN and alters the main peak shape.

Standard deviation in the nearest neighbor bond length (NNRS): The oscillation in absorption intensities becomes weaker as NNRS increases, especially at higher energies (Figure 6c, z_4 and Figure 5b, C_{pit}). In addition, there is also a small variation in the pre-edge peak. In general, this is a mild contribution to the total signal.

Minimum oxygen-oxygen distance (MOOD): This contribution to the total signal is even smaller than the other descriptors. However, the variation at the main peak is strong: the position of the main peak shifts to lower energy with increasing MOOD (Figure 6c, z_5 and Figure 5b, E_{main}).

Unconstrained latent variable z_6 : By design, an AAE disentangles its latent space. In the case of the RankAAE model presented here, z_6 is disentangled from and contains information supplementary to the constrained variables $\{z_1, z_2, \dots, z_5\}$. In other words, the plot for z_6 in Figure 6c represents *all other* spectral trends not captured by the five structure descriptors.

Having illustrated the characteristics of RankAAE models, we now describe quantitative validation. First, we characterize the correlation between target descriptors and latent space dimensions. To be clear, the value of a latent variable does not equal the value of the structure descriptor directly. However, each latent variable is nearly monotonically correlated to a specific structure descriptor, due to the soft constraint in our loss function. For the same model presented in Figure 6c, latent space values are plotted against calculated structure descriptor values in Figure 7 for the test set in scatter plots for each of the five descriptors. The correlations are visually apparent. They are quantified using the F1 score [69] for categorical variables (OS and CN) and the Spearman rank correlation coefficient [68] (SRCC) for continuous variables (OCN, NNRS, and MOOD). In order to match datatypes in the computation of the F1 score, the corresponding latent variable is mapped to a categorical variable by thresholding. An F1 score ranges from 0 for poor classification performance to 1 for perfect classification performance. An SRCC value of 0 occurs when there is no correlation; a value of 1 indicates a perfect, monotonic, positive correlation. Furthermore, the degree of correlation depends on the choice of descriptor, as illustrated in Figure S3 in the Supplemental Material [61]. As shown by the annotations in Figure 6c and summarized in Figure S5a in the Supplemental Material [61], all the F1 scores are above 0.91 and all the SRCCs are above 0.86, highlighting the ability of the RankAAE methodology to drive latent variables to track structure descriptors in a nearly monotonic fashion.

Next, the emergent spectral trends encoded in the reconstructed spectra in the RankAAE model in Figure 6c are compared to features in the ground truth spectra, i.e., the spectra directly from the data set. For simplicity, we focus on the primary spectral trends that have been identified in Figure 6c. We compute these primary spectral descriptors for each of the underlying spectra in the test set and present them in Figure 7. Visual inspection verifies that the evolution of the spectral descriptors tracks both the latent variables and the structure descriptors. To quantify these relationships, the F1 scores and SRCCs between the spectral descriptor in the data set and the latent space variable are computed (Figure S5b in the Supplemental Material [61]). The coefficients are smaller than those for the structure descriptors (Figure S5a in the Supplemental Material [61]). However, they are still significant, especially considering that only one local spectra feature has been used. Also, the additional fluctuations associated with the spectral features as emergent in a multi-dimensional space play a role. Finally, the reduction of the correlation analysis to single dimensions is an oversimplification. For example, for z_5 , the correlation with spectral descriptor (E_{main}) is only clear after categorizing the data points according to the CN because the main peak variation is also affected by CN. This effect is shown in Figure S6 in the Supplemental Material [61] using the test portion of the vanadium oxide dataset as an example. The data points aggregate to a strip for each CN category. Inside each strip, the correlation between z_5 and E_{main} is much clearer, especially for CN5 and CN6, than when considering all the data in aggregate. This illustrates that a local spectral feature, such as E_{main} , can still be affected by multiple structure descriptors.

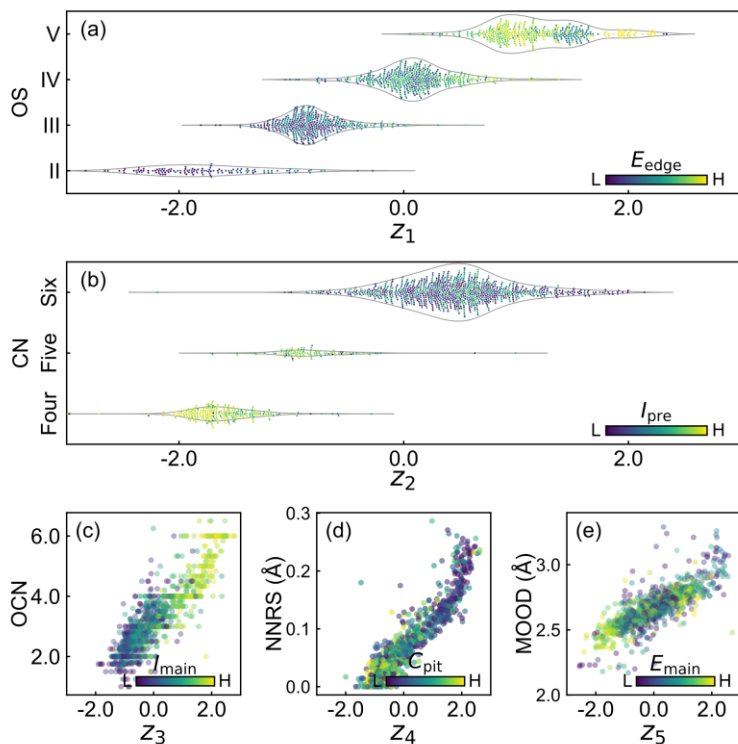


Figure 7. Correlation between the RankAAE latent space variable (x-axis), target structure descriptor (y-axis), and emergent spectrum descriptor (color) for the vanadium oxide data set. Violin plots are shown for categorical variables OS (a) and CN (b). Scatter plots are shown for continuous variables OCN (c), NNRS (d), and MOOD (e).

Taken together, these validation results demonstrate that the trained RankAAE model provides multidimensional, physical structure-spectrum relationships for the vanadium oxide dataset.

B. RankAAE performance across the full set of transition metal oxides

The characteristics, trends, validation, and the final set of five structure descriptors shown for the specific example of V oxides carry over to the full set of 3d transition metal families considered in this work. The spectral data colored by the structure descriptors reveals trends with OS and CN, but ambiguity remains when considering the larger set of descriptors (Figure S7 in the Supplemental Material [61]). RankAAE models are trained on the datasets of Ti, Cr, V, Mn, Fe, Co, Ni, and Cu oxides. An independent model is created for each material family. The same set of structure descriptors is used for the entire series to facilitate comparisons. The reconstructed spectra versus each latent space dimension exhibit smooth trends (Figure S8 in the Supplemental Material [61]). The amplitude of spectral variation for the unconstrained latent space variable z_6 varies across the series. This suggests an opportunity to further fine-tune the choice of structure descriptors. Examining the spectral trends, the same set of spectral features can be tracked, although the specific energy ranges (relative to the edge) and the relative importance of the primary versus secondary vary.

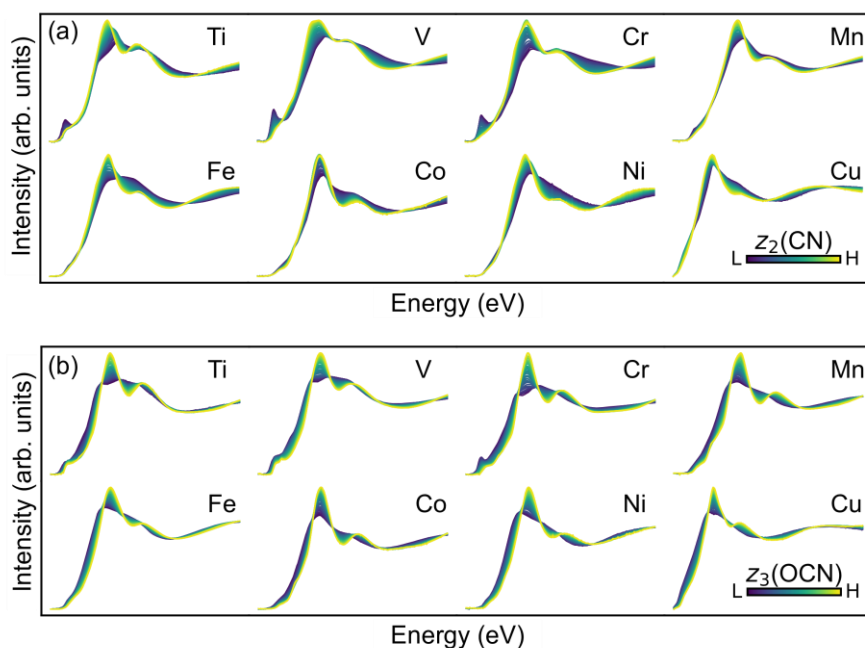


Figure 8. RankAAE derived spectrum variation trends for CN (a) and OCN (b) over Ti, V, Cr, Mn, Fe, Co, Ni, and Cu oxides data sets. The spectra are colored by underlying latent variables z_2 and z_3 targeting CN and OCN, respectively.

To illustrate these trends further, Figure 8a shows the trend of the spectral variations for CN across the series of first-row transition metals. The RankAAE models reveal two distinct primary trends. First, like for V oxides, the pre-edge peak variation is the primary spectral trend for Ti, Cr, and Mn oxides. The pre-edge peak intensity is strongest for V and Cr, intermediate for Ti, and small, but discernable for Mn. Second, the late transition metals (Fe, Co, Ni, and Cu) exhibit minimal intensity in the pre-edge peak region of the XANES K-edge spectrum [37,38]. For these metals, the secondary spectral trend overweighs the primary trend: the second peak intensity above

the main edge, which is at least 5 eV away from the main peak (Figure 5b, I_{post}), correlates with CN.

As a second example, the spectral trend for OCN is shown in Figure 8b. There is a universal pattern: spectral intensity changes at the main peak. For Cu oxides, the intensity variation is localized to that around the main peak. It is also dominant for Mn oxides, Co oxides and Ni oxides. Other cases show more extensive, secondary spectral variation. Overall, the spectral features associated with OCN are distinguishable from those driven by CN. However, Cr oxide is an exception, where the trend associated with OCN is very similar to the trend for CN.

Validation of all the RankAAE models using the test sets for each material family shows strong correlations between latent space values and the target structure descriptors as well as the identified spectral features (Figure S9 in the Supplemental Material [61]). Quantitatively the F1 scores and SRCCs for structure descriptor to latent space values are high, as shown as insets in Figure S8 in the Supplemental Material [61] for each material and descriptor, as well as in summary form in Figure S5a in the Supplemental Material [61]. Most F1 scores are above 0.90 with a minimum of 0.82. Most SRCCs are above 0.85 with a minimum of 0.80. This indicates that the latent space in the RankAAE models captures the changes in structure descriptors sufficiently across the entire set of 3d transition metal families. OS and CN exhibit high correlations in all cases. The correlation for the three new structure descriptors varies across the metal series, but is very strong in general, showing their potential for motif characterization in XANES analysis. Quantitative validation against the primary spectral descriptors (Figure S5b in the Supplemental Material [61]) supports the qualitative results presented in Figure S9 in the Supplemental Material [61]. Some specific cases do highlight the limitations of relying on a single spectral descriptor. For example, the low correlation (SRCC=0.13) for the Mn oxide CN descriptor with the pre-edge intensity spectral descriptor traces to the relatively low intensity of the pre-edge feature, as noted above. However, the SRCC increases to 0.44 if I_{post} is used. Consistently, I_{post} captures CN better as a spectral descriptor for late transition metals (Mn-Cu).

V. DISCUSSION

The RankAAE model described here creates a multidimensional, continuous spectral representation that maps a target set of structure descriptors to XANES K-edge spectral characteristics. The *spectrum validity* is a key prerequisite for its interpretability. A meaningful structure-spectrum relationship can only be established on a properly regularized latent space. Each point in the latent space must reconstruct to a physical XANES spectrum. As illustrated in the vanadium K-edge data set in Figure S1a in the Supplemental Material [61], a basic AE model does not meet the *spectrum validity* requirement, as discussed in the Introduction. In fact, the AE model does reconstruct the spectra very well for latent space points close to areas well represented in the original training data set, evidenced by the dominance of “normal” spectra in Figure S1a in the Supplemental Material [61]. However, there are regions with values that fall in the valid range overall, but for which the reconstruction can become catastrophically poor. These latent space points fall where the model is interpolating in regions sparsely represented in the training original dataset [59]. In contrast, the latent space of the RankAAE model is regularized by the adversarial constraint, which enforces *spectrum validity*, as shown in Figure S1b in the Supplemental Material [61], consistently produces physically meaningful data points.

Analysis of the reconstructed spectra along specific dimensions in the latent space of the

RankAAE associated with each structure descriptor reveals well-defined spectral trends. For use by domain experts, it is essential that these trends are robust, and that there is a reasonable workflow to allow the expert to assess which choices of a structure descriptor are capturing the essential spectral variations, and that the trends are physically interpretable.

The workflow for developing multidimensional models demonstrated here exhibits stability in the spectral trends that emerge. We revisit the test sequence where we start with a single constrained latent space variable and one free dimension. We also add a latent space variable that is constrained to a pseudo-random number sequence as a structure descriptor (NOISE) to probe the limit of the noise level in deriving spectral trends. Thus, we start with $N = 3$, and add dimensions together with constraints in succession. The root of this chain is OS, followed by CN, OCN, NNRS, and finally MOOD. For the vanadium oxide dataset, following the same analysis as presented in Figure 6c, the appropriately averaged, reconstructed spectra along each dimension are shown in Figure S4 in the Supplemental Material [61]. Several important points emerge. First, by visual inspection, the trends shown for each dimension remain qualitatively consistent as new dimensions and structure descriptors are added to the model. For example, for OS, the edge shift and horizontal translation remain stable with different combinations of structure descriptors, although the amplitude of the changes varies quantitatively. In particular, in the test using OS as the only structure descriptor, the pre-edge variation is not localized. Instead, it appears in two latent variables (z_1 and unconstrained z_3). With the addition of CN and more structure descriptors, it is localized to z_2 . This reflects the resolution of entangled trends as further structure descriptors are added. Second, the effectiveness of the constraint in forcing the latent space variables to track the target descriptors is nearly independent of the dimensionality of the model, as shown by the F1 scores and SRCCs that quantify the correlation between latent space variables and target structure descriptors (Figure S4 in the Supplemental Material [61]). Third, the noise spectral trend always appears in a very narrow range independent of dimensionality and the choices of constraint on other latent dimensions. Finally, the amplitude of the unconstrained spectral trend systematically goes down as new dimensions and structure descriptors are added to the model. This gives a qualitative guide indicating that each additional descriptor is capturing new spectral variation. With the full complement of five structure descriptors in the model, the spectral variation associated with the unconstrained variable is close to the noise level in vanadium oxides.

To go beyond this qualitative description raises the non-trivial question of how many distinct structure descriptors can be extracted from a given set of spectral data covering a chosen energy range in XAS. For the EXAFS region, the analysis has incorporated considerations of information content based on Fourier analysis or Nyquist's theorem, as discussed both early in the development of the field [70] and in recent reviews [71,72]. While this provides one estimate of the amount of near neighbor shell information that can be extracted, it is often an upper bound for many reasons, including noise in experimental spectra. For the XANES region, the situation is more complex, e.g., due to the role of multi-body scattering paths. For a given chosen spectral range, it is an open question as to whether general guidelines for information content can be developed. In the present work, our choice of a 50 eV range is empirical, reflecting a common understanding of the scope of the XANES region in the field. Our analysis demonstrates the extraction of five descriptors for the example of transition metal oxides, but further research may reveal opportunities to extract more descriptors and provide more quantitative guidelines.

Another important aspect of the workflow for a domain expert is to distinguish between different choices of target descriptors. Several criteria have already been illustrated, including

distinguishable spectral trends, quantitative correlations between target structure descriptors and constrained latent space variables, and the amplitude of the unguided spectral trend. Here we further analyze the degree to which each new descriptor results in a qualitatively distinguishable spectral trend. In order to capture the energy dependence of each trend when generating new spectra by traversing the latent space, we examine a “differential spectrum”. We specifically represent the total dynamic range of the trend by taking the difference between the last and first spectrum corresponding to the 95th percentile and 5th percentile values in the range of the latent space variable. We revisit the comparison of trends for different structure descriptors presented in Figures 6 and Supplemental Material S3 [61] for vanadium oxides through the differential spectra shown in Figure S10 in the Supplemental Material [61]. This representation of the trends also exhibits the qualitative stability of trends for OS and CN as other structure descriptors are changed. Further, as previously highlighted, the overall amplitude of the unguided trend is reduced as the choice of structure descriptors is refined. As the other structure descriptors are swapped out, the qualitative changes in the trends can be seen. For example, OCN more narrowly isolates the change in the main peak intensity while CN-2 is convolved with the shift of the main peak energy. Also, the trend for MOOD captures changes to the main peak shape, position, and intensity as compared to the relative spread-out trend for PGSO. The spectral trends are an emergent characteristic of the RankAAE models, so they are not constrained to be linearly independent. Nonetheless, in the workflow, the degree of overlap can be monitored through the intercorrelation measured by the Cosine similarity [73], with the maximum values noted in each panel of Figure S10 in the Supplemental Material [61]. Overall, as the workflow proceeds, the degree of overlap is reduced. In our final model, the CN descriptor and unconstrained dimension result in the most independent trends while OS, OCN, NNRS, and MOOD derived trends remain correlated with each other to a moderate degree by this measure. Finally, this analysis points in the direction of future investigations to refine spectral fingerprints beyond the simplification of the specific, localized spectral descriptors in Figure 5b.

Turning to physical interpretation, each of the primary and secondary spectral trends identified above with structure descriptors (Figure 6c) can be understood. Among all the structure-spectrum relationships, the correlation between OS and edge position is probably the most well-known [30]. Conceptually, it corresponds to a basic physical principle: the energy cost to excite a core electron increases as the cation becomes more positively charged. The secondary spectral trend, the correlated shift of the post-edge intensity to higher energy, corresponds to the famous “molecular ruler” [20], also referred to as Natoli’s rule: [21,74,75] the local bond length is inversely linked to the position of a peak. Here the ionic radius changes inversely with the atomic charge for the $3d$ transition metal cation, driving the metal-oxygen bond length. These two empirical rules were proposed independently. The unified spectral trend for OS that emerges from the RankAAE models captures these correlated effects of the cation oxidation state. Because of the inverse correlation between OS and ionic radius, similar trends would be expected upon replacing OS with ionic radius in other studies if that were more appropriate for the description of material structures.

The pre-edge peak is linked to the breaking of local inversion symmetry in the first coordination shell around the cation. The onset of the K-edge spectrum involves low-energy empty states on the cation that are typically of $3d$ orbital character. The $s \rightarrow d$ transition is formally forbidden for an ideal octahedral motif for 6-fold coordinated cations that exhibits inversion symmetry [38]. In lower coordination motifs, e.g., tetrahedral (4-fold coordination), as well as distorted atomic motifs indicated by non-zero NNRS, there is no local inversion symmetry. The

hybridization with cation $4p$ orbitals leads to dipole-allowed transitions [31,37]. This is the basic physical explanation for the pre-edge spectral trends associated with low CN and high NNRS, fully consistent with the original observations of Farges *et al.*, and subsequent work [31,37].

Drawing on the multiple scattering representation of the absorption process (Figure 1a), the number of scatterers is an important factor that affects the absorption intensity, while path lengths and interference effects determine the energy position where the absorption intensity changes. The changes in main peak position, intensity, and shape are a key part of the spectral trend for both CN and OCN (Figure 6c, z_2 and z_3). The overall increase of intensity with increasing CN tracks the number of scatterers. For CN, the sharpening of the main peak shape with more weight at lower energy as well as the shift of the post peak features to lower energy are further examples of Natoli's rule. As coordination increases, so do the metal-oxygen bond lengths. Correspondingly, the path lengths get longer.

The OCN trend (Figure 6c, z_3) modulates the intensity around the main peak and to a lesser extent the post-peak region. The OCN descriptor correlates with the number of atoms in the second coordination shell, and it covers a broader range of atoms that can contribute to three- and four-body scattering paths (Figure 1a). With an increase in the number of such paths induced by the increase of such atoms, interference effects are stronger. But, as compared to the CN trend, the energy location differs because of interatomic distances. OCN essentially counts the atoms in the second coordination shell while CN counts the first coordination shell (Figure 5a), and hence is related to larger interatomic distance. It results in a longer scattering path length (Figure 1a, SS2 versus SS1), longer wavelength, and consequently a lower energy photoelectron. As a result, the OCN induces spectral variation only at lower energy while CN also induces variations at higher energies (Figure 6c z_3 vs z_2). While OCN and second shell coordination number count the same set of atoms, it is worth noting that the definition of cutoff radius for the second coordination shell is ambiguous. In contrast, OCN is well-defined in established algorithms [76]. As a result, better latent space-structure descriptor correlation can be observed for OCN (Figure 6c z_3 versus Figure 6b z_3).

Interestingly, there are several isosbestic points in the spectral trends for CN and OCN shown in Figure 6 and Figure 8. In general, isosbestic points in spectral trends indicate that the change is a linear combination of spectra contributions from two distinct sources. For XANES, changes in physical structure affect both the number of scatters and the path lengths that together determine the net spectral change, so the emergence of these isosbestic points may be surprising. However, in this example, varying CN and OCN primarily changes the number of scattering paths. The length of the scattering paths remains similar, making it plausible that the net effect is to change only the intensity of the features in the spectra with isosbestic points in between.

The NNRS descriptor measures the spread in the cation-oxygen bond lengths in the first coordination shell (Figure 5a). The scattering path lengths are spread out with larger NNRS. Correspondingly, there is less reinforcement of constructive and destructive interference among the paths as a function of photoelectron energy. This, in turn, reduces the definition of the peaks and valleys in the spectrum. Specifically, as the NNRS increases, the peak-to-valley height decreases. This is expressed, for example, in the local curvature in the post-peak spectral region, quantified in the C_{pit} spectral feature (Figure 5b). The curvature is higher (deeper, better-defined valley) when NNRS is smaller (Figure 6c z_4). In the database and this analysis, a larger NNRS value represents crystal structures with low symmetry and a spread of local bond lengths. It is

worth noting that C_{pit} is a specific spectral descriptor that can be clearly discerned for most transition metals. Nonetheless, for specific metals, e.g., Mn in Figure S5b in the Supplemental Material [61], the correlation between C_{pit} and z_4 is low with an SRCC of 0.05. As suggested by the spectral trends in Figure S8 in the Supplemental Material [61], the curvature at the main peak correlates with z_4 better where the SRCC increases to 0.33.

The variation in the local angles between the cation-oxygen bonds is another measure of spread in low symmetry local motifs. It flows through to the path lengths for three-body scattering paths (Figure 1a). In our comparison of different specific structure descriptors, we found that MOOD, the distance minimum between two oxygens in the first coordination shell, was the most effective metric for the RankAAE models (Figure 5a) and can be related to the 3-body scattering path (Figure 1a, MS) in an X-ray absorption process. Based on a few examples in Figure S11 in the Supplemental Material [61], the bending of the two axial oxygens towards the equatorial plane is one of the most intuitive structure characteristics that the underlying latent variable z_5 tracks. The distortion associated with the bending is expected to change the 3-body scattering path length. The triangular path of MS is relatively long and leads to oscillation at relatively low energy. In the case of vanadium oxides, it is located around the main peak. With increasing oxygen-oxygen distance, the corresponding wavelength increases and pushes the main peak position to lower energy.

VI. CONCLUSION

The approximately fifty-seven thousand simulated XANES K-edge data considered in this study span eight $3d$ transition metal oxide families and exemplify complex data sets that domain experts are highly motivated to study. Our novel rank constraint introduces physical interpretability into the latent space representation of data produced by an AAE. Specifically, latent variables are guided to track specific structure descriptors. From there, the reconstruction from different regions in the latent space produces spectral variations reflecting the underlying changes in structure descriptors. The correlations between structure descriptors and spectral trends are confirmed to be both qualitatively and quantitatively observed in the ground truth data.

For the specific example of XANES analysis, we have shown that our method significantly extends the scope of spectral features a domain expert can analyze systematically, from local metrics such as edge position to extended fingerprints that can encompass the full energy range of XANES spectra. Leveraging the prior knowledge of domain experts, these spectral trends can be used to gain further insight into features in materials structures and their relationship to spectral trends. Also, as illustrated by the set of new structure descriptors explored in this work, it is straightforward for a domain expert to explore new trends and new structure descriptors using the RankAAE method driven by their intuition.

From a machine-learning perspective, we also highlight the challenges associated with the extreme diversity of our datasets. First, the data is completely unbiased towards the specific methods used in this study. It consists of materials data compiled independently by contributors to the Materials Project over many years, driven by goals unrelated to the scope of this work. Second, the target signal, namely a specific transition metal K-edge XANES spectrum, is sampled from a set of diverse crystal structures with widely varying, albeit physically realistic, local structure motifs and associated atomic species. Consequently, the data encompasses a broad variety of physical complexities that affect the spectrum. These include different local chemical states of the transition metal cation, scattering amplitudes that differ among atomic species, and the distinct set

of scattering paths (number of paths contributing, number of atoms in each path, and overall length of each path) associated with each cation structure motif. This means that our methods have numerically parsed a tremendous amount of diversity to identify the key trends we present. Thus, the successes that we demonstrate in this work are quite encouraging and highlight the potential of this, and related methods, for unearthing new spectral trends for impact on physical science and engineering.

In summary, we demonstrate the capability of RankAAE using the XANES structure-spectrum relationship as an example. However, the RankAAE is a *general* framework for enhancing the interpretability of an AAE. In principle, it is also applicable to other kinds of spectroscopic data, and in general, to *any* dataset in which one is attempting to discover physical correlations between target signals and associated driving characteristics.

VII. APPENDIX A: METHODS

A.1 Structure of RankAAE and training

AAE [56], VAE [62] and WAE [9] regularize the latent space to a target distribution. The VAE constrains its latent space by a KL divergence penalty to impose a prior distribution. The WAE performs similar regularization by a Wasserstein distance [77]. The AAE does not require a specified functional form of the penalty function to the prior distribution. It instead matches the aggregated posterior of the latent vector with the prior distribution by a discriminator via an adversarial training procedure. The AAE is reported to generate a better data manifold [56,59,78,79]. In this work, a Gaussian distribution is used as the prior distribution to guide the training. We have adopted the gradient reversal layer from the domain adversarial neural network to implement the adversarial training mechanism [80].

The encoder, decoder, and discriminator (Figure 4) are all composed of 5 fully connected layers, with a hidden layer size of 64. Both Dropout and Batch Normalization are used. A Parametric Rectified Linear Unit (PReLU) [81] function is added after every fully connected layer with two exceptions: (1) there is no activation function after the last layer in the encoder; (2) the last activation function in the decoder is Softplus.

We incorporate the mutual information regularization term from the Dual Adversarial AutoEncoder [82] to maximize the information shared between the latent space and the spectrum data. We have also used a smooth loss in the early stage of the training to speed up the convergence of network parameters.

With the performance evaluated on the validation set, the hyperparameters (learning rate, batch size, noise-based data augmentation, gradient reversal amplitude, hidden layer size, the number of hidden layers, and other aspects of the network structure) are optimized for vanadium oxides. The performance is balanced between the appearance of the spectral trends and the correlation of the latent variables with structure descriptors.

A.2 Rank constraint

We developed a novel regularization term based on the Kendall rank correlation coefficient (KRCC) [83] to organize the latent space such that each dimension has a clear physical meaning. It enforces a monotonic dependence while being a soft enough constraint to allow variation in the

final value. KRCC is built upon the concept of concordant and discordant pairs. In the application to RankAAE, a latent variable for spectrum i can be denoted as z_i and the corresponding structure descriptor can be denoted as p_i . Any pair of (z_i, p_i) and (z_j, p_j) are said to be concordant if both $z_i > z_j$ and $p_i > p_j$ or both $z_i < z_j$ and $p_i < p_j$; otherwise, they are said to be discordant. KRCC describes the monotonic dependence by the ratio of the number of concordant/discordant pairs. The counting is simplified using a $\text{sgn}()$ function, i.e., $\text{sgn}(z_i - z_j)$ and $\text{sgn}(p_i - p_j)$. However, the gradient of the $\text{sgn}()$ function is 0 except at 0, which is detrimental to the purpose of neural network parameter optimization. We have modified KRCC to a loss function with a proper gradient everywhere by removing $\text{sgn}()$ function around $(z_i - z_j)$. The loss function $L_c(z)$ for each constrained dimension of the latent space is defined by:

$$\begin{aligned}
 f(z_i, z_j) &= (z_i - z_j)\text{sgn}(p_i - p_j) \\
 N_{\text{con}} &= \sum_{i=1}^n \sum_{j=1}^n \Theta(f(z_i, z_j)) \\
 N_{\text{dis}} &= \sum_{i=1}^n \sum_{j=1}^n \Theta(-f(z_i, z_j)) \\
 f_{\text{con}}(z) &= \frac{1}{\max[N_{\text{con}}, N_{\text{dis}}]} \sum_{i=1}^n \sum_{j=1}^n \Theta(f(z_i, z_j)) f(z_i, z_j) \\
 f_{\text{dis}}(z) &= \frac{1}{N_{\text{dis}}} \sum_{i=1}^n \sum_{j=1}^n \Theta(-f(z_i, z_j)) f(z_i, z_j) \\
 L_c(z) &= -\frac{1}{n(n-1)} N_{\text{dis}} (f_{\text{con}}(z) + f_{\text{dis}}(z)) \tag{1}
 \end{aligned}$$

where z is a latent variable, p is a structure descriptor, n is the number of samples in a minibatch in network training, i and j are sample indices in a minibatch, $\text{sgn}()$ is a function to extract the sign of a real number, $\Theta()$ is a Heaviside step function being 1 for positive values and 0 for other values. Abbreviations “con” and “dis” designate concordant and discordant. N_{con} and N_{dis} represent the number of concordant and discordant pairs in the minibatch for the constrained dimension of the latent space. The minimum value of $N_{\text{dis},m}$ is clamped to 1 to avoid the singularity. The prefactors for the concordant and discordant pair contributions to the loss function are scaled to be comparable early in the training and then to have a diminishing impact as more data points become concordant. Overall, we find that this approach minimizes the effect of rank constraint on the distribution of the latent space variable, which is defined by the AAE algorithm [56].

A.3 Data acquisition and preprocessing

The crystal structures are pulled from the Material Project [57]. Each crystal structure is a local minimum in energy with lattice parameters and atomic coordinates [84]. All the crystal structures containing each target transition metal atom and oxygen are included. The XANES K-Edge spectra

are simulated using the FEFF9 program [65] for the resulting database of about forty thousand unique atomic sites identified in the oxides for the eight transition metals considered here with the condition that the first neighbor shell of the transition metal site only contains oxygen. The breakdown of the dataset according to transition metal is given in Table S1.

In this work, XANES is defined as the spectrum with energy ranging from the start of the edge of the spectrum to 50 eV above the edge. To create feature vectors that are directly comparable for each metal considered, the spectra are interpolated to a fixed energy grid, equally spaced with 256 points. For machine learning, each data set is partitioned with a random selection of 70%, 15%, and 15% of the spectra to form training, validation, and test sets, respectively. The size of the sets is given in Table S1 for each transition metal.

A.4 Structure descriptor calculation

Eight structure descriptors have been explored in this study:

- 1) OS: oxidation state of a cation represented by an integer. OS is determined using the maximum a posteriori (MAP) estimation method in *pymatgen* [85–87].
- 2) CN: coordination number in the first coordination shell around a specified transition metal cation. CN is determined using the CrystalNN algorithm in the *pymatgen* package [76,86].
- 3) CN-2: coordination number in the second coordination shell for each target transition metal cation. CN-2 is estimated by the number of atoms that falls in the layer between radius r_1 and r_2 , chosen for each transition metal based on the radial distribution of the atoms.
- 4) OCN: oxygen coordination number. OCN is computed as an average of the coordination numbers for the nearest neighbor oxygen atoms to the target transition metal cation. The coordination number of each oxygen atom is determined using the CrystalNN algorithm in the *pymatgen* package [76,86].
- 5) NNRS: the standard deviation of the bond lengths from the target transition metal cation to the nearest neighbor oxygen atoms.
- 6) NNAS: the standard deviation of the bond angles with the cation as the vertex and two nearest neighbors as endpoints.
- 7) MOOD: the minimum oxygen-oxygen distance on the edges of the nearest neighbor polyhedron encompassing the target transition metal cation. The polyhedron is determined by the CrystalNN algorithm in the *pymatgen* package [76,86].
- 8) PSGO: the point group symmetry order (PGSO) represented by an integer, computed as the total number of symmetry operations for the point group for the full crystal space group determined by the *pymatgen* package [76,86].

A.5 Spectrum descriptor calculation

The spectrum descriptors shown in Figure 5b are defined as:

- 8) E_{edge} : the position of the absorption edge.
- 9) I_{pre} : pre-peak intensity, the maximum intensity among the peaks occurring at an energy lower than the edge.
- 10) E_{main} : main peak position, the energy value at the first peak after the edge.
- 11) I_{main} : main peak intensity. For numerical stability purposes, we use the average over a 1 eV window centered at E_{main} .
- 12) C_{pit} : post edge curvature. The position of the pit is found by a local minimum of the

spectrum in the portion at least 20 eV above the edge. The curvature is computed as the second order derivative. To avoid numerical instabilities, the curvature is averaged over a 10 eV window centered at the pit.

- 13) E_{post} : second peak position. A second peak is identified by a local maximum of the spectrum. To avoid ambiguities in peak counting, we considered $E_{\text{edge}} + 15$ eV for transition metals Ti through Mn and $E_{\text{edge}} + 20$ eV for Fe through Cu as a proxy for E_{post} . The numerical test in Figure S5b in the Supplemental Material [61] confirms the effectiveness of this approach.
- 14) I_{post} : absorption intensity at E_{post} .

A.6 Spectral variation plot generation

Specific latent space samples are fed to the decoder to generate reconstructed spectral variation sequences. For each latent variable, 50 points on an equally spaced grid are sampled. From the range of the latent space values encoded from the test set spectra, the latent space grid is chosen to extend from the 5th to the 95th percentile values. For Figure S2a in the Supplemental Material [61], a spectrum is reconstructed for each of the 50 values of the specified latent space variable, setting the others equal to zero. For Figure S2b in the Supplemental Material [61] and all of the other reconstructed spectral trends, an averaging procedure is applied to sample the other latent space variables. Specifically, the average of reconstructed spectra for each value of the target latent space variable is carried out over 10000 samples drawn from a multivariate Gaussian distribution to fill the rest latent variables. There are 50 spectra in each spectral trend presented.

A.7 Model selection

For each dataset, we repeat the training 100 times with different random seeds for the neural network parameters, and all the discussions are based on the best one. The performance for each model is assessed by:

$$S = -\max_{i,j} |\rho_{ij}| + \sum_i |\rho'_i| \quad (2)$$

where ρ_{ij} is the Spearman rank correlation coefficient (SRCC) between a guided latent variable z_i and an unconstrained latent variable z_j , ρ'_i is the correlation score (F1 score for OS and CN, SRCC for other structure descriptors) between a constrained latent variable z_i and the target structure descriptor. To drive the performance metrics to have equal contributions, the z-scores are computed across the 100 models by subtracting the average and dividing the standard deviation for each ρ_{ij} and ρ'_i . Then the z-scores [88] are used to evaluate S in Eq. (2). A good model maximizes the latent-structure correlation and minimizes the intercorrelation between the unconstrained latent variables and the other dimensions. Therefore, a larger S indicates better performance. For this study, it is sufficient to consider the two terms with equal weight. However, for other studies, the assessment of model performance can be further refined by scaling the two terms differently. Table S2 details the statistics of different loss terms (prior to z-score normalization) that enter Eq. (2), as well as the standard reconstruction error attributed to the autoencoder. For the current study, the reconstruction error is small and it does not vary significantly. Hence, we did not include reconstruction error in the model selection. The other terms vary within reasonable bounds. For further insight, the correlation plots for the worst model out of 100 runs by this criterion for the case of V oxide are shown in Figure S12 in the Supplemental Material [61]. Encouragingly, the correlations for that model are still reasonably

good.

ACKNOWLEDGMENT

This research is based upon work supported by the U.S. Department of Energy, Office of Science, Office Basic Energy Sciences, under Award Number FWP PS-030. This research also used the Theory and Computation facility of the Center for Functional Nanomaterials (CFN), which is a U.S. Department of Energy Office of Science User Facility, at Brookhaven National Laboratory under Contract No. DE-SC0012704.

REFERENCES

- [1] J. Yano, K. J. Gaffney, J. Gregoire, L. Hung, A. Ourmazd, J. Schrier, J. A. Sethian, and F. M. Toma, *The Case for Data Science in Experimental Chemistry: Examples and Recommendations*, Nat. Rev. Chem. **6**, 357 (2022).
- [2] R. Batra, L. Song, and R. Ramprasad, *Emerging Materials Intelligence Ecosystems Propelled by Machine Learning*, Nat. Rev. Mater. **6**, 655 (2020).
- [3] Z. Wang, Z. Sun, H. Yin, X. Liu, J. Wang, H. Zhao, C. H. Pang, T. Wu, S. Li, Z. Yin, et al., *Data- Driven Materials Innovation and Applications*, Adv. Mater. **34**, 2104113 (2022).
- [4] T. Le, V. C. Epa, F. R. Burden, and D. A. Winkler, *Quantitative Structure–Property Relationship Modeling of Diverse Materials Properties*, Chem. Rev. **112**, 2889 (2012).
- [5] B. Sanchez Lengeling and A. Aspuru Guzik, *Inverse Molecular Design Using Machine Learning: Generative Models for Matter Engineering*, Science **361**, 360 (2018).
- [6] Z. Chen, N. Andrejevic, T. Smidt, Z. Ding, Q. Xu, Y. Chi, Q. T. Nguyen, A. Alatas, J. Kong, and M. Li, *Direct Prediction of Phonon Density of States with Euclidean Neural Networks*, Adv. Sci. **8**, 2004214 (2021).
- [7] J. Yang, L. Tao, J. He, J. R. McCutcheon, and Y. Li, *Machine Learning Enables Interpretable Discovery of Innovative Polymers for Gas Separation Membranes*, Sci. Adv. **8**, eabn9545 (2022).
- [8] P. Y. Lu, S. Kim, and M. Soljačić, *Extracting Interpretable Physical Parameters from Spatiotemporal Systems Using Unsupervised Learning*, Phys. Rev. X **10**, 031056 (2020).
- [9] Z. Rao, P. Y. Tung, R. Xie, Y. Wei, H. Zhang, A. Ferrari, T. P. C. Klaver, F. Körmann, P. T. Sukumar, A. K. Silva, et al., *Machine Learning–Enabled High-Entropy Alloy Discovery*, Science **378**, 78 (2022).
- [10] D. Flam Shepherd, T. Wu, X. Gu, A. Cervera Lierta, M. Krenn, and A. Aspuru Guzik, *Learning Interpretable Representations of Entanglement in Quantum Optics Experiments Using Deep Generative Models*, Nat. Mach. Intell. **4**, 544 (2022).
- [11] S. H. Wang, H. S. Pillai, S. Wang, L. E. K. Achenie, and H. Xin, *Infusing Theory into Deep Learning for Interpretable Reactivity Prediction*, Nat. Commun. **12**, 5288 (2021).
- [12] A. E. A. Allen and A. Tkatchenko, *Machine Learning of Material Properties: Predictive and Interpretable Multilinear Models*, Sci. Adv. **8**, eabm7185 (2022).
- [13] M. A. Kramer, *Nonlinear Principal Component Analysis Using Autoassociative Neural Networks*, AIChE J. **37**, 233 (1991).
- [14] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016).
- [15] Z. Pei, K. A. Rozman, Ö. N. Doğan, Y. Wen, N. Gao, E. A. Holm, J. A. Hawk, D. E. Alman, and M. C. Gao, *Machine- Learning Microstructure for Inverse Material Design*, Adv. Sci. **8**, 2101207 (2021).

- [16] V. K. Ramaswamy, S. C. Musson, C. G. Willcocks, and M. T. Degiacomi, *Deep Learning Protein Conformational Space with Convolutions and Latent Interpolations*, Phys. Rev. X **11**, 011052 (2021).
- [17] R. Gómez Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández Lobato, B. Sánchez Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru Guzik, *Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules*, ACS Cent. Sci. **4**, 268 (2018).
- [18] P. K. Routh, Y. Liu, N. Marcella, B. Kozinsky, and A. I. Frenkel, *Latent Representation Learning for Structural Characterization of Catalysts*, J. Phys. Chem. Lett. **12**, 2086 (2021).
- [19] C. Miles, M. R. Carbone, E. J. Sturm, D. Lu, A. Weichselbaum, K. Barros, and R. M. Konik, *Machine Learning of Kondo Physics Using Variational Autoencoders and Symbolic Regression*, Phys. Rev. B **104**, 235111 (2021).
- [20] S. P. Cramer, *X-Ray Spectroscopy with Synchrotron Radiation: Fundamentals and Applications* (Springer Nature Switzerland, Cham, 2020).
- [21] J. Timoshenko and B. Roldan Cuenya, *In Situ / Operando Electrocatalyst Characterization by X-Ray Absorption Spectroscopy*, Chem. Rev. **121**, 882 (2021).
- [22] G. S. Henderson, F. M. F. de Groot, and B. J. A. Moulton, *X-Ray Absorption Near-Edge Structure (XANES) Spectroscopy*, Rev. Mineral. Geochem. **78**, 75 (2014).
- [23] F. Meirer and B. M. Weckhuysen, *Spatial and Temporal Exploration of Heterogeneous Catalysts with Synchrotron Radiation*, Nat. Rev. Mater. **3**, 324 (2018).
- [24] K. V. Klementev, *Extraction of the Fine Structure from X-Ray Absorption Spectra*, J. Phys. Appl. Phys. **34**, 209 (2001).
- [25] M. Newville, P. Liviņš, Y. Yacoby, J. J. Rehr, and E. A. Stern, *Near-Edge X-Ray Absorption Fine Structure of Pb: A Comparison of Theory and Experiment*, Phys. Rev. B **47**, 14126 (1993).
- [26] J. J. Rehr and R. C. Albers, *Theoretical Approaches to X-Ray Absorption Fine Structure*, Rev. Mod. Phys. **72**, 621 (2000).
- [27] J. J. Rehr and A. L. Ankudinov, *Progress in the Theory and Interpretation of XANES*, Coord. Chem. Rev. **249**, 131 (2005).
- [28] S. I. Bokarev and O. Kühn, *Theoretical X-ray Spectroscopy of Transition Metal Compounds*, WIREs Comput. Mol. Sci. **10**, e1433 (2020).
- [29] P. Norman and A. Dreuw, *Simulating X-Ray Spectroscopies and Calculating Core-Excited States of Molecules*, Chem. Rev. **118**, 7208 (2018).
- [30] J. Bergengren, *Über die Röntgenabsorption des Phosphors*, Z. Für Phys. **3**, 247 (1920).
- [31] F. Farges, G. E. Brown, and J. J. Rehr, *Ti K-Edge XANES Studies of Ti Coordination and Disorder in Oxide Compounds: Comparison between Theory and Experiment*, Phys. Rev. B **56**, 1809 (1997).
- [32] M. Wang and Z. Feng, *Pitfalls in X-Ray Absorption Spectroscopy Analysis and Interpretation: A Practical Guide for General Users*, Curr. Opin. Electrochem. **30**, 100803 (2021).
- [33] A. A. Guda, S. A. Guda, K. A. Lomachenko, M. A. Soldatov, I. A. Pankin, A. V. Soldatov, L. Braglia, A. L. Bugaev, A. Martini, M. Signorile, et al., *Quantitative Structural Determination of Active Sites from in Situ and Operando XANES Spectra: From Standard Ab Initio Simulations to Chemometric and Machine Learning Approaches*, Catal. Today **336**, 3 (2019).
- [34] A. A. Guda, S. A. Guda, A. Martini, A. N. Kravtsova, A. Algasov, A. Bugaev, S. P. Kubrin, L. V. Guda, P. Šot, J. A. van Bokhoven, et al., *Understanding X-Ray Absorption Spectra by*

- Means of Descriptors and Machine Learning Algorithms*, Npj Comput. Mater. **7**, 203 (2021).
- [35] J. Wong, F. W. Lytle, R. P. Messmer, and D. H. Maylotte, *K -Edge Absorption Spectra of Selected Vanadium Compounds*, Phys. Rev. B **30**, 5596 (1984).
- [36] H. Singh, M. Topsakal, K. Attenkofer, T. Wolf, M. Leskes, Y. Duan, F. Wang, J. Vinson, D. Lu, and A. I. Frenkel, *Identification of Dopant Site and Its Effect on Electrochemical Activity in Mn-Doped Lithium Titanate*, Phys. Rev. Mater. **2**, 125403 (2018).
- [37] T. Yamamoto, *Assignment of Pre-Edge Peaks in K-Edge x-Ray Absorption Spectra of 3d Transition Metal Compounds: Electric Dipole or Quadrupole?*, X-Ray Spectrom. **37**, 572 (2008).
- [38] M. R. Carbone, S. Yoo, M. Topsakal, and D. Lu, *Classification of Local Chemical Environments from X-Ray Absorption Spectra Using Supervised Machine Learning*, Phys. Rev. Mater. **3**, 033604 (2019).
- [39] W. E. Jackson, F. Farges, M. Yeager, P. A. Mabrouk, S. Rossano, G. A. Waychunas, E. I. Solomon, and G. E. Brown, *Multi-Spectroscopic Study of Fe(II) in Silicate Glasses: Implications for the Coordination Environment of Fe(II) in Silicate Melts*, Geochim. Cosmochim. Acta **69**, 4315 (2005).
- [40] J. C. Woicik, E. L. Shirley, C. S. Hellberg, K. E. Andersen, S. Sambasivan, D. A. Fischer, B. D. Chapman, E. A. Stern, P. Ryan, D. L. Ederer, et al., *Ferroelectric Distortion in SrTiO₃ Thin Films on Si (001) by x-Ray Absorption Fine Structure Spectroscopy: Experiment and First-Principles Calculations*, Phys. Rev. B **75**, 140103 (2007).
- [41] J. P. Itié, B. Couzinet, A. Polian, A. M. Flank, and P. Lagarde, *Pressure-Induced Disappearance of the Local Rhombohedral Distortion in BaTiO₃*, Europhys. Lett. EPL **74**, 706 (2006).
- [42] W. Zhang, M. Topsakal, C. Cama, C. J. Pelliccione, H. Zhao, S. Ehrlich, L. Wu, Y. Zhu, A. I. Frenkel, K. J. Takeuchi, *Multi-Stage Structural Transformations in Zero-Strain Lithium Titanate Unveiled by in Situ X-Ray Absorption Fingerprints*, J. Am. Chem. Soc. **139**, 16591 (2017).
- [43] A. I. Frenkel, Y. Feldman, V. Lyahovitskaya, E. Wachtel, and I. Lubomirsky, *Microscopic Origin of Polarity in Quasiamorphous BaTiO₃*, Phys. Rev. B **71**, 024116 (2005).
- [44] W. Zhang, D. H. Seo, T. Chen, L. Wu, M. Topsakal, Y. Zhu, D. Lu, G. Ceder, and F. Wang, *Kinetic Pathways of Ionic Transport in Fast-Charging Lithium Titanate*, Science **367**, 1030 (2020).
- [45] N. Marcella, Y. Liu, J. Timoshenko, E. Guan, M. Luneau, T. Shirman, A. M. Plonka, J. E. S. van der Hoeven, J. Aizenberg, C. M. Friend, et al., *Neural Network Assisted Analysis of Bimetallic Nanocatalysts Using X-Ray Absorption near Edge Structure Spectroscopy*, Phys. Chem. Chem. Phys. **22**, 18902 (2020).
- [46] S. Xiang, P. Huang, J. Li, Y. Liu, N. Marcella, P. K. Routh, G. Li, and A. I. Frenkel, *Solving the Structure of “Single-Atom” Catalysts Using Machine Learning - Assisted XANES Analysis*, Phys. Chem. Chem. Phys. **24**, 5116 (2022).
- [47] 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).
- [48] C. Zheng, K. Mathew, C. Chen, Y. Chen, H. Tang, A. Dozier, J. J. Kas, F. D. Vila, J. J. Rher, L. F. J. Riper, et al., *Automated Generation and Ensemble-Learned Matching of X-Ray Absorption Spectra*, Npj Comput. Mater. **4**, 12 (2018).
- [49] J. Timoshenko, D. Lu, Y. Lin, and A. I. Frenkel, *Supervised Machine-Learning-Based*

- Determination of Three-Dimensional Structure of Metallic Nanoparticles*, J. Phys. Chem. Lett. **8**, 5091 (2017).
- [50] C. Chen, Y. Zuo, W. Ye, X. Li, and S. P. Ong, *Learning Properties of Ordered and Disordered Materials from Multi-Fidelity Data*, Nat. Comput. Sci. **1**, 46 (2021).
- [51] C. Zheng, C. Chen, Y. Chen, and S. P. Ong, *Random Forest Models for Accurate Identification of Coordination Environments from X-Ray Absorption Near-Edge Structure*, Patterns **1**, 100013 (2020).
- [52] J. Timoshenko and A. I. Frenkel, “*Inverting*” *X-Ray Absorption Spectra of Catalysts by Machine Learning in Search for Activity Descriptors*, ACS Catal. **9**, 10192 (2019).
- [53] A. Yamashita, T. Nagata, S. Yagyu, T. Asahi, and T. Chikyow, *Direct Feature Extraction from Two-Dimensional X-Ray Diffraction Images of Semiconductor Thin Films for Fabrication Analysis*, Sci. Technol. Adv. Mater. Methods **2**, 23 (2022).
- [54] Y. Suzuki, T. Taniai, K. Saito, Y. Ushiku, and K. Ono, *Self-Supervised Learning of Materials Concepts from Crystal Structures via Deep Neural Networks*, Mach. Learn. Sci. Technol. **3**, 045034 (2022).
- [55] S. Tetef, N. Govind, and G. T. Seidler, *Unsupervised Machine Learning for Unbiased Chemical Classification in X-Ray Absorption Spectroscopy and X-Ray Emission Spectroscopy*, Phys. Chem. Chem. Phys. **23**, 23586 (2021).
- [56] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, *Adversarial Autoencoders*, ArXiv:1511.05644 (2016).
- [57] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, et al., *Commentary: The Materials Project: A Materials Genome Approach to Accelerating Materials Innovation*, APL Mater. **1**, 011002 (2013).
- [58] Y. Liu, A. Halder, S. Seifert, N. Marcella, S. Vajda, and A. I. Frenkel, *Probing Active Sites in Cu_xPd_y Cluster Catalysts by Machine-Learning-Assisted X-Ray Absorption Spectroscopy*, ACS Appl. Mater. Interfaces **13**, 53363 (2021).
- [59] D. Berthelot, C. Raffel, A. Roy, and I. Goodfellow, *Understanding and Improving Interpolation in Autoencoders via an Adversarial Regularizer*, ArXiv:1807.07543 (2018).
- [60] D. Weininger, *SMILES, a Chemical Language and Information System. 1. Introduction to Methodology and Encoding Rules*, J. Chem. Inf. Model. **28**, 31 (1988).
- [61] See Supplemental Material at [URL Will Be Inserted by Publisher] for Additional Figures and Tables.
- [62] D. P. Kingma and M. Welling, *Auto-Encoding Variational Bayes*, ArXiv:1312.6114 (2014).
- [63] O. Prykhodko, S. V. Johansson, P. C. Kotsias, J. Arús Pous, E. J. Bjerrum, O. Engkvist, and H. Chen, *A de Novo Molecular Generation Method Using Latent Vector Based Generative Adversarial Network*, J. Cheminformatics **11**, 74 (2019).
- [64] S. H. Hong, S. Ryu, J. Lim, and W. Y. Kim, *Molecular Generative Model Based on an Adversarially Regularized Autoencoder*, J. Chem. Inf. Model. **60**, 29 (2020).
- [65] J. J. Rehr, J. J. Kas, F. D. Vila, M. P. Prange, and K. Jorissen, *Parameter-Free Calculations of X-Ray Spectra with FEFF9*, Phys Chem Chem Phys **12**, 5503 (2010).
- [66] F. Musil, A. Grisafi, A. P. Bartók, C. Ortner, G. Csányi, and M. Ceriotti, *Physics-Inspired Structural Representations for Molecules and Materials*, Chem. Rev. **121**, 9759 (2021).
- [67] T. Lam Pham, H. Kino, K. Terakura, T. Miyake, K. Tsuda, I. Takigawa, and H. Chi Dam, *Machine Learning Reveals Orbital Interaction in Materials*, Sci. Technol. Adv. Mater. **18**, 756 (2017).
- [68] S. Kokoska and D. Zwillinger, *CRC Standard Probability and Statistics Tables and Formulae*,

- Student Edition*, 1st Edition (CRC Press, New York, 2000).
- [69] A. Tharwat, *Classification Assessment Methods*, Appl. Comput. Inform. **17**, 168 (2021).
- [70] P. A. Lee, P. H. Citrin, P. Eisenberger, and B. M. Kincaid, *Extended X-Ray Absorption Fine Structure—Its Strengths and Limitations as a Structural Tool*, Rev. Mod. Phys. **53**, 769 (1981).
- [71] S. Bordiga, E. Groppo, G. Agostini, J. A. van Bokhoven, and C. Lamberti, *Reactivity of Surface Species in Heterogeneous Catalysts Probed by In Situ X-Ray Absorption Techniques*, Chem. Rev. **113**, 1736 (2013).
- [72] A. Kuzmin and J. Chaboy, *EXAFS and XANES Analysis of Oxides at the Nanoscale*, IUCrJ **1**, 571 (2014).
- [73] G. Salton, *Automatic Text Processing* (Addison-Wesley Longman Publishing Co., Inc., USA, 1988).
- [74] A. Bianconi, M. Dell’Ariccia, A. Gargano, and C. R. Natoli, *Bond Length Determination Using XANES*, in *XAFS and Near Edge Structure* (Springer-Verlag, Berlin Heidelberg, 1983), pp. 57–61.
- [75] J. J. Rehr, A. Ankudinov, and S. I. Zabinsky, *New Developments in NEXAFS/EXAFS Theory*, Catal. Today **39**, 263 (1998).
- [76] H. Pan, A. M. Ganose, M. Horton, M. Aykol, K. A. Persson, N. E. R. Zimmermann, and A. Jain, *Benchmarking Coordination Number Prediction Algorithms on Inorganic Crystal Structures*, Inorg. Chem. **60**, 1590 (2021).
- [77] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf, *Wasserstein Auto-Encoders*, ArXiv:1711.01558 (2019).
- [78] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther, *Autoencoding beyond Pixels Using a Learned Similarity Metric*, ArXiv:1512.09300 (2016).
- [79] Q. Jin, Y. Ma, F. Fan, J. Huang, X. Mei, and J. Ma, *Adversarial Autoencoder Network for Hyperspectral Unmixing*, IEEE Trans. Neural Netw. Learn. Syst. (2021).
- [80] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, *Domain-Adversarial Training of Neural Networks*, in *Domain Adaptation in Computer Vision Applications*, edited by G. Csurka (Springer International Publishing, Cham, 2017), pp. 189–209.
- [81] K. He, X. Zhang, S. Ren, and J. Sun, *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*, ArXiv:1502.01852 (2015).
- [82] P. Ge, C. X. Ren, D. Q. Dai, J. Feng, and S. Yan, *Dual Adversarial Autoencoders for Clustering*, IEEE Trans. Neural Netw. Learn. Syst. **31**, 1417 (2020).
- [83] M. G. Kendall, *A New Measure of Rank Correlation*, Biometrika **30**, 81 (1938).
- [84] P. Hohenberg and W. Kohn, *Inhomogeneous Electron Gas*, Phys. Rev. **136**, B864 (1964).
- [85] R. Bassett and J. Deride, *Maximum a Posteriori Estimators as a Limit of Bayes Estimators*, Math. Program. **174**, 129 (2019).
- [86] S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson, and G. Ceder, *Python Materials Genomics (Pymatgen): A Robust, Open-Source Python Library for Materials Analysis*, Comput. Mater. Sci. **68**, 314 (2013).
- [87] M. O’Keefe and N. E. Brese, *Atom Sizes and Bond Lengths in Molecules and Crystals*, J. Am. Chem. Soc. **113**, 3226 (1991).
- [88] W. Mendenhall and T. Sincich, *Statistics for Engineering and the Sciences* (Pearson Prentice-Hall, 2007).