

Unsupervised representation learning of Kohn–Sham states and consequences for downstream predictions of many-body effects

Received: 8 April 2024

Bowen Hou, Jinyuan Wu  & Diana Y. Qiu  

Accepted: 18 October 2024

Published online: 02 November 2024

 Check for updates

Representation learning for the electronic structure problem is a major challenge of machine learning in computational condensed matter and materials physics. Within quantum mechanical first principles approaches, density functional theory (DFT) is the preeminent tool for understanding electronic structure, and the high-dimensional DFT wavefunctions serve as building blocks for downstream calculations of correlated many-body excitations and related physical observables. Here, we use variational autoencoders (VAE) for the unsupervised learning of DFT wavefunctions and show that these wavefunctions lie in a low-dimensional manifold within latent space. Our model autonomously determines the optimal representation of the electronic structure, avoiding limitations due to manual feature engineering. To demonstrate the utility of the latent space representation of the DFT wavefunction, we use it for the supervised training of neural networks (NN) for downstream prediction of quasiparticle bandstructures within the GW formalism. The GW prediction achieves a low error of 0.11 eV for a combined test set of two-dimensional metals and semiconductors, suggesting that the latent space representation captures key physical information from the original data. Finally, we explore the generative ability and interpretability of the VAE representation.

Recently, machine learning (ML) has emerged as a powerful tool in condensed matter and materials physics, achieving substantial progress across areas including the identification of phase transitions^{1–7}, quantum state reconstruction^{8–10}, prediction of topological order^{11,12}, symmetry^{13–19} and the study of electronic structure^{20–23}. Among these ML applications, exploring the electronic structure of real materials is of particular interest, since it allows for the extension of computationally intensive predictive quantum theories to understand the physics of larger and more complex systems, such as moiré systems^{13,24–32} and defect states^{33–39}. Within atomistic first principles theories, density functional theory (DFT)⁴⁰ is the most commonly used approach for studying the electronic structure of materials. In principle, DFT gives accurate descriptions of the ground state charge

density, but quantitative prediction of excited-state properties (including bandstructure) requires the introduction of excitations from the many-body ground state^{41,42}. However, DFT can be used as a starting point for many-body calculations, where the wavefunctions within Kohn–Sham DFT are used to construct correlation functions for the excited states^{43,44}. Therefore, harnessing the information embedded in the Kohn–Sham (KS) wavefunction⁴⁰ becomes crucial for downstream applications. Here, a key challenge lies in distilling a succinct representation of the electronic structure while preserving the essential information⁴⁵.

In contrast to the notable achievements of ML descriptors for crystal geometry and chemical composition^{46–50}, the electronic structure of materials remains extremely challenging to learn. The high

dimensionality of the KS wavefunction of real materials creates a complex data structure, making direct pattern detection challenging. Secondly, electronic structure is highly nonlinear and sensitive to both crystal configuration and intricate non-local correlations, making it difficult to develop a general ML model applicable across a broad spectrum of materials. ML models for electronic structure have been mostly confined to the study of specific subsets of materials, such as molecules⁵¹, perovskites^{52–54}, or layered transition metal dichalcogenides (TMDs)⁵⁵. In most of these approaches, the ML has focused on the prediction of single-valued properties, such as band gaps, largely due to the challenges of identifying well-defined, interpretable, and efficient representations of the electronic structure^{56–59}. Recently, the development of operator representations, which capture more nuanced data about the underlying quantum states, has provided new opportunities for the prediction of full band structures. These techniques involve the physically informed selection of specific operators as fingerprints. Examples include the early successful use of energy decomposed operator matrix elements (ENDOME), combined with radially decomposed projected density of states (RAE-PDOS), to predict quasiparticle (QP) band structures⁶⁰, and the use of spectral operator representations to predict material transparency⁶¹. However, the choice of descriptors in such approaches is informed by human physical intuition on the domain science side. This raises the question of whether more fundamental or generalizable descriptors can be learned, independent of human selection.

Variational autoencoders (VAEs)^{62,63}, a class of probabilistic models that combine variational inference with autoencoders used to compress and decompress high dimensional data, stand out as a promising tool for learning electronic structure in a way that allows for unsupervised training and thus avoids feature engineering. The effectiveness of VAE compression has been demonstrated for various applications across condensed matter physics, including quantum state compression⁶⁴, detection of critical features in phase transitions⁶⁵ and decoupled subspaces^{66,67}.

Here, we show that a well-crafted VAE is capable of representing KS-DFT wavefunctions on a manifold within a significantly compressed latent space, which is $10^3 - 10^4$ times smaller than the original input. Importantly, these succinct representations retain the physical information inherent in the initial data. To validate the efficacy of the VAE latent space in practical applications, we then build a supervised deep learning model for downstream prediction of k-resolved quasiparticle energies within the many-body GW approximation^{68,69}, whose input includes VAE representations of KS states. This model yields a mean absolute error (MAE) of 0.11 eV, when applied to a test set of two-dimensional (2D) metals and insulators, which is comparable to the intrinsic numerical error of the GW approach^{32,44,70}. Moreover, the VAE inherently predicts a smooth, physically realistic band structure, in contrast with previous models, which required an additional manually-imposed smoothing function to remove unphysical variations⁶⁰. Lastly, we address the interpretability of the VAE representation and show that the success of electronic structure representation learning and downstream GW prediction is derived from the smoothness of the VAE latent space, which corresponds to the smoothness of the wavefunction in k-space. Because of this property, the VAE can serve as a promising wavefunction generator capable of predicting realistic wavefunctions at arbitrary points in latent space, which has the potential to be used for both wavefunction interpolation in k-space and materials design and discovery across the compositional phase space.

Results and discussion

VAE for electronic structure representation

As Fig. 1a shows, the architecture of our VAE consists of two complementary sets of NN: the encoder (e_θ) and the decoder (d_θ). The encoder maps the input of high-dimensional KS wavefunction moduli

$|\phi_{nk}(\mathbf{r})| \in \mathbb{R}^{R_x \times R_y \times R_z}$ to a low-dimensional vector ($\boldsymbol{\mu} \in \mathbb{R}^p$, $\boldsymbol{\sigma} \in \mathbb{R}^p$) of variational mean and deviation in the latent space. Here, $p \ll R_x \times R_y \times R_z$, where p represents the dimension of the latent space, and $R_{x,y,z}$ is the dimension of the real-space wavefunction $\phi_{nk}(\mathbf{r})$ in the unit cell, where the wavefunction is indexed by the band n and the crystal momentum (k -point) \mathbf{k} . Conversely, the decoder, mirroring the structure of the encoder, has an opposite function, which is to reconstruct the input wavefunction from its low-dimensional latent vector ($\boldsymbol{\mu}$, $\boldsymbol{\sigma}$). In essence, the VAE can be conceptualized as an information bottleneck for electronic structure, where the encoder acts as a data refinement process, discarding redundant information and noise from the input wavefunction. Meanwhile, the decoder ensures that vital physical information is preserved for wavefunction recovery. Since the KS state is essentially a spatial distribution with local patterns, convolutional neural networks (CNN)^{71–74} provide a practical framework for the first two layers of the encoder and decoder.

By imposing a distribution of latent variables close to a standard normal distribution $N(0, I)$ during training, the VAE is capable of a smooth mapping from the connected KS states $|n, \mathbf{k}\rangle$ and $|n, \mathbf{k} + \Delta\mathbf{k}\rangle$ to nearby points in the latent space, enabling generative power after training. We define the total VAE loss function as:

$$\mathcal{L} = \frac{1}{T} \sum_{nk} |||\phi_{nk}| - d_\theta(e_\theta(|\phi_{nk}|))||^2 + \beta \cdot \frac{1}{T} \sum_{nk} \sum_{j=1}^p \frac{1}{2} (-1 + \sigma_{nk}^{(j)2} + \mu_{nk}^{(j)2} - \log \sigma_{nk}^{(j)2}) \quad (1)$$

where the first term is the mean squared error (MSE) of the reconstructed wavefunctions, and the second term is the Kullback-Leibler (KL) divergence, which forces the latent space to approach a normal $N(0, I)$ distribution. The parameter β is the weight of KL divergence, tuning the degree of regularization of the latent space. T is the total number of KS states in the training set and, j is an index in the latent space.

Additionally, we design our VAE to respect rotational and translational symmetries of the crystal so that downstream prediction of physical observables will not vary with the choice of unit cell. As shown in Fig. 1a, b, we insert a global average pooling (GAP)⁷⁵ layer between the CNN and the dense layer in the encoder. The GAP layer allows CNNs to accept inputs of varying sizes and produce a fixed-length output, which is crucial for the subsequent fully connected layers in the encoder. By aggregating the spatial information of the feature maps, the GAP layer also enforces translational invariance with respect to the unphysical degree of freedom that arises from the choice of the origin of the unit cell. Then, we include a circular padding technique in the CNNs⁷⁶ to account for the periodic boundary conditions of the crystal system. The first layer of the encoder features a discrete rotational CNN⁷⁷ along the lattice constant directions, which select out the combined feature maps with maximum summation value to preserve invariance to the choice of lattice vectors and direction. With such a CNN, given the same material, data from any continuously complete periodicity of a specific KS state will generate the same latent space. This allows us to capture properties associated with physical symmetries of the crystal without explicitly including symmetry labels. Finally, to ensure that the VAE-generated wavefunction's size aligns with the input, allowing us to define a proper loss function, we include an adaptive layer as the final layer in our decoder.

By feeding DFT wavefunctions to our VAE model, we can interpret the vector ($\boldsymbol{\mu}$, $\boldsymbol{\sigma}$) in latent space as a low-dimensional effective representation associated with the individual KS state. This representation has several advantages: (i) it is smooth due to the smooth nature of the neural network; (ii) it can handle translational invariance and periodic boundary conditions, (iii) it can handle any symmetry and cell size; (iv) even when trained with unit cell data, it can be extended to supercells beyond a single unit cell. More details and benchmarks regarding the VAE model are given in Supplementary Notes 1–3.

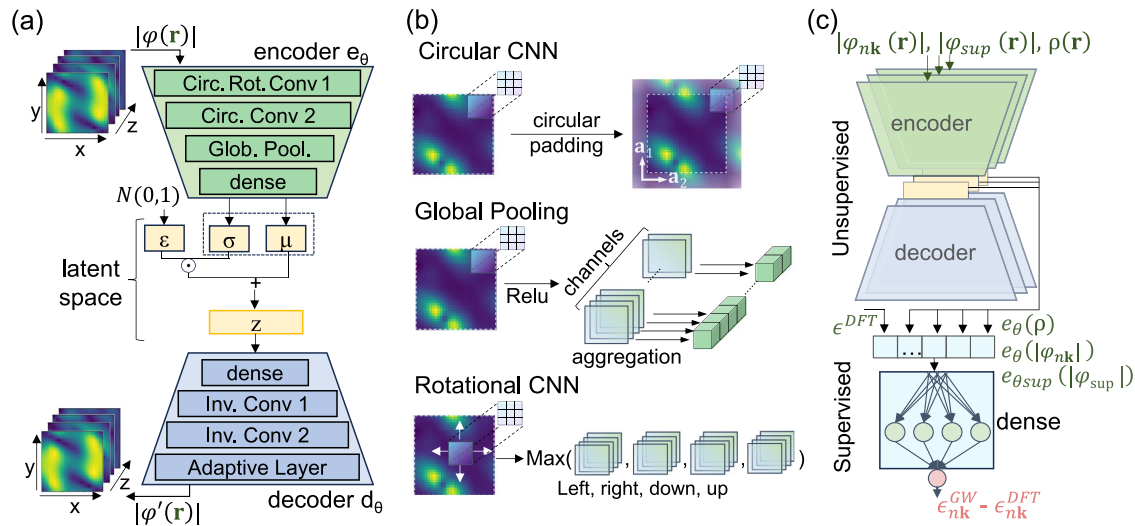


Fig. 1 | Neural network architecture of variational autoencoders (VAE) for crystal systems. **a** Schematic of the VAE. The encoder (green trapezoid) consists of a rotational convolutional neural network (CNN), a circular CNN, a global pooling layer and one flattened dense layer, which are explicitly presented in **(b)**. The encoder maps a real space wavefunction to a latent space vector of variational mean μ and variance σ . Z is the sampled latent vector, drawn from a variational Gaussian distribution using a reparameterization trick⁶³. The decoder (blue trapezoid) has symmetric neural networks (NN) structures. The latent space serves as an information bottleneck for the VAE as its dimensions are $10^3 - 10^4$ smaller than the input and output (represented by the colormaps of the wavefunction in real space). **b** The circular CNN layer includes circular padding techniques based on the periodic boundary conditions. The width of the circular padding equals the size of the CNN kernel so that the output of the CNN has the same periodicity as the input. Global pooling layers output the average value of each channel from the CNN

feature map. Rotational CNN layers scan the input along four directions and outputs the max of (left, right, up, down) feature maps to the next layer. The detail of model parameters and additional benchmarks are listed in Supplementary Note 1. **c** Schematic of the overall semi-supervised learning model, including both the unsupervised VAE and supervised dense NN. The VAE inputs are the Kohn–Sham (KS) wavefunction modulus $|\varphi_{nk}(\mathbf{r})|$, all super states $|\varphi_{sup}(\mathbf{r})|$ and charge density $\rho(\mathbf{r})$ in real space. The input layer of the supervised dense NN is comprised of density functional theory (DFT) energies, denoted as ϵ^{DFT} , along with low dimensional effective representations of $\varphi_{nk}(\mathbf{r})$, $\varphi_{sup}(\mathbf{r})$ and $\rho(\mathbf{r})$ denoted as $e_{\theta}(\varphi_{nk}(\mathbf{r}))$, $e_{\theta_{sup}}(\varphi_{sup}(\mathbf{r}))$ and $e_{\theta_{\rho}}(\rho(\mathbf{r}))$ respectively. These representations are encoded within the VAE latent space (yellow square) through an encoder with parameters θ , θ_{sup} and θ_{ρ} , which are unsupervisedly trained for all KS wavefunctions, super states and charge density.

Neural network for downstream prediction of GW band structures

Due to its ground-state nature, DFT calculations can yield inaccurate single-particle band structures and tend to underestimate the band gap of semiconductors compared to experiment^{78–80}. To incorporate the many-electron correlation effects that are missing in DFT and obtain accurate quasiparticle (QP) bandstructures, one can replace the effective single-particle exchange correlation potential in DFT with non-local and energy-dependent self energy calculated within the GW approximation in many-body perturbation theory^{43,68,81,82}. This approach proves highly effective for computing QP bandstructures in a wide range of materials⁴³. However, in practice, constructing a GW self-energy, even for small systems^{70,81,83–86}, is much more computationally expensive than DFT, and this remains a bottleneck to the broader adoption of the GW approach for high throughput studies. Therefore, in the context of understanding materials' excited-state properties, a natural approach is to use low-fidelity techniques like density functional theory (DFT) to try and predict the results of high-fidelity, computationally intensive many-body calculations. Currently, most ML work in this field focuses on the use of indirect predictors for model training, using, for instance, crystal geometry, chemical composition, and the DFT bandgap as input^{46–49,56,57}. These methods, however, are generally limited to single value prediction of the band gap, and only effective within specific subsets of materials, such as inorganic solids⁵⁹. Operator fingerprint methods like ENDOME and RAD-PDOS extend predictive capability to k -resolved band structures, but the selection of features in these models heavily relies on human intuition and can introduce unphysical fluctuations in band structures, which are then smoothed in post-processing⁶⁰. There have also been ML models dedicated to predicting the screened dielectric function^{87–89}, which can speed up GW calculations but don't provide information that is generalizable to different systems.

Due to the challenging nature of capturing non-local frequency-dependent correlations in the electronic structure, we select the downstream prediction of GW bandstructures as proof of principle of the effectiveness of our VAE latent vector representation of the KS wavefunctions. We develop a supervised deep NN on top of the VAE latent space to predict many-body GW corrections. The goal of this NN is to successfully learn the diagonal part of GW self-energy Σ_{nk}^{GW} :

$$\Sigma_{nk}^{GW} = \frac{i}{2\pi V} \sum_{\mathbf{G}\mathbf{G}'} \sum_{\bar{\mathbf{q}}} \sum_m^{all} \int_{-\infty}^{\infty} d\omega' W_{\mathbf{G}\mathbf{G}'}(\bar{\mathbf{q}}, \omega') \times \frac{\rho_{nk, m\mathbf{k}-\mathbf{q}}(\mathbf{G}) \rho_{nk, m\mathbf{k}-\mathbf{q}}^*(\mathbf{G}')}{\epsilon_{n, \mathbf{k}} + \omega' - \epsilon_{m, \mathbf{k}-\mathbf{q}} + i\eta \cdot \text{sgn}(\epsilon_{m, \mathbf{k}-\mathbf{q}} - \mu)} \quad (2)$$

where \mathbf{G} is a reciprocal lattice vector; \mathbf{q} is the difference between any two \mathbf{k} -vectors $\mathbf{k} - \mathbf{k}'$ and is integrated over the Brillouin zone (BZ); ϵ is KS-DFT energy; and $\rho_{nk, m\mathbf{k}-\mathbf{q}}(\mathbf{G}) = \langle n\mathbf{k} | e^{i(\mathbf{q} + \mathbf{G})\mathbf{r}} | m\mathbf{k} - \mathbf{q} \rangle$. $W_{\mathbf{G}\mathbf{G}'}$ is the screened Coulomb interaction calculated within the random phase approximation, and ω is the frequency dependence of the self energy. Notably, the calculation of the GW self energy includes a sum over infinite bands, m . In practice, the sum over states is treated as a convergence parameter, and the number of bands included in the summation is of the same order as the number of reciprocal lattice vectors $\bar{\mathbf{G}}$ included in W .

The diagonal self-energy matrix element for a specific state $|n, \mathbf{k}\rangle$ shown in Eq.(2) can be expressed as: $\Sigma_{nk}^{GW} = f(\varphi_{nk}, \epsilon_{nk}, \Phi, \epsilon, \rho)$, where $\varphi_{nk} \in \mathbb{C}^{R_x \times R_y \times R_z}$ and $\Phi \in \mathbb{C}^{N \times R_x \times R_y \times R_z}$ are the KS state vectors; $\rho \in \mathbb{R}^{R_x \times R_y \times R_z}$ is the charge density; N is the total number of occupied and unoccupied states $|n', \mathbf{k}'\rangle$ used in GW calculations. $\epsilon_{nk} \in \mathbb{R}^1$ and $\epsilon \in \mathbb{R}^N$ are the corresponding DFT eigenvalues. Here, we note that Φ includes φ_{nk} and ϵ includes ϵ_{nk} , but we include both terms explicitly in the

function f to make later steps more transparent. The output $\Sigma_{nk}^{GW} \in \mathbb{R}^1$ is the GW self energy. Due to the closed-form expression of the self-energy, mathematically, we expect that a simple dense NN can acquire an understanding of the non-linear mapping from the KS wavefunction and energies to Σ_{nk}^{GW} , which consists of the mapping of $(\mathbb{C}^{(N+1) \times (R_x \times R_y \times R_z)}, \mathbb{R}^{R_x \times R_y \times R_z}, \mathbb{R}^{(N+1)} \rightarrow \mathbb{R}^1)^{90,91}$.

However, two significant challenges prevent the application of a simple NN model such as $\hat{\Sigma}_{nk}^{GW} = \hat{f}_{NN}(\phi_{nk}, \varepsilon_{nk}, \Phi, \varepsilon, \rho, \eta)$, where $\hat{\Sigma}_{nk}^{GW} \in \mathbb{R}^1$ is the NN predicted self-energy, and $\hat{\eta}$ are the parameters of the model. (i) Due to the high computational cost of the GW algorithm, only a limited subset of GW energies near the Fermi level can be exactly calculated and used for the supervised learning training set, denoted as $\Sigma_{nk}^{train} \in \mathbb{R}^{N_{train}}$. As a result, $N_{train} \ll (N+1) \times (R_x \times R_y \times R_z)$, and overfitting is inevitable (see Supplementary Fig. 8). (ii) The curse of dimensionality makes it formidable to learn the nonlinear mapping from a high-dimensional sparse Φ wavefunction space to the self-energy Σ_{nk}^{train} space. To address these challenges, we adopt the manifold assumption^{92–95} that the DFT electronic wavefunction in real space can be modeled as lying on a low-dimensional manifold. Additionally, we assume that two DFT wavefunctions mapped to nearby points on the manifold should have comparable contributions to the final GW energy corrections. If these two assumptions hold true, then the VAE is ideal for downstream prediction of GW self energies.

Here, to capture the non-local correlation, the pseudobands approximation^{96–98} is further employed for all states. That is, DFT wavefunctions with close energies are summed into effective super states (see Supplementary Note 6). We assume that the manifold assumption also applies to the super states and charge density, and the VAE is used to further remove redundant information in them. Eventually, our semi-supervised, physics-informed model reads:

$$\hat{\Sigma}_{nk}^{GW} = \hat{f}_{NN}(e_{\theta}(\phi_{nk}(\mathbf{r})), \varepsilon_{nk}, e_{\theta_{sup}}(\Phi_{sup}(\mathbf{r})), \varepsilon_{sup}, e_{\theta_{\rho}}(\rho), \eta) \quad (3)$$

where e_{θ} , $e_{\theta_{sup}}$ and $e_{\theta_{\rho}}$ are encoders exclusively trained for KS states, super states and charge density respectively. The schematic workflow of our model is shown in Fig. 1b. The output of the NN is the predicted GW diagonal self-energies $\hat{\Sigma}_{nk}^{GW} = \hat{\varepsilon}_{nk}^{GW} - \varepsilon_{nk}^{DFT}$, and the inputs are the wavefunction of $|n, \mathbf{k}\rangle$, super bands corresponding to occupied and unoccupied states on a uniform k-grid, and the ground state energies. These inputs are identical to the input of an explicit GW calculation to eliminate bias due to feature selection.

Application to dataset of 2D materials

To benchmark the predictive power of our model, we select 302 materials with unit cells of 3 to 4 atoms, including both metals and insulators, from the Computational 2D Materials Database (C2DB)^{99–101}, for training and validation. For the unsupervised VAE training, our dataset is comprised of 68,384 DFT electronic states sampled on a $6 \times 6 \times 1$ uniform k-grid, which are randomly split into 90% training set and 10% test set. For supervised training of GW corrections, the entire GW energy dataset is randomly partitioned into two subsets: 10% (2201 electronic states) is allocated as the test set, while the remaining 90% (19801 electronic states) is designated as the training set. We further completely exclude a small subset of 30 materials from the training set for the VAE and three monolayer TMD materials—MoS₂, WS₂, and CrS₂—from both the VAE and GW training datasets. Comprehensive details regarding dataset, training dynamics and additional benchmarks are listed in Supplementary Note 5.

Figure 2a shows three DFT electronic wavefunctions, A, B, and C, for monolayer MoS₂, which is completely excluded from the training data, within the x-y plane of the unit cell, which correspond to the dark blue circles shown in Fig. 2e. Figure 2b shows the VAE-reconstructed wavefunction of states A, B and C, which are nearly identical to the original wavefunction after recovery from the low-dimensional latent

space with high coefficients of determination, R^2 , of 0.96, 0.94 and 0.91 respectively. Figure 2c shows the VAE variational mean vector $\hat{\mu}$ for states A, B and C respectively. Overall, the R^2 of the VAE-reconstructed wavefunction is 0.92 across the test set when compared with the amplitude of the original DFT wavefunction, and the normalization of the wavefunction is preserved within 1% (see Supplementary Fig. 10). Remarkably, the VAE-reconstructed wavefunctions reproduce the ground state charge density, the physically relevant quantity in DFT, with a high $R^2 = 0.99$ even though the total charge density is never used in the VAE training. The results demonstrate that the original electronic wavefunction in the high-dimensional real space can be effectively compressed by $10^3 - 10^4$ times into a representative vector in $\mathbb{R}^{p=30}$ in a way that still preserves the vital information needed for reconstruction. Our autonomously determined representation is over 100 times more compact than previous electronic fingerprint approaches⁶⁰.

Figure 2d shows the comparison between the GW corrections calculated explicitly and predicted by the downstream ML model, where orange (blue) dots represent the test (training) sets. The model yields an MAE of 0.06 eV ($R^2 = 0.96$) and 0.11 eV ($R^2 = 0.94$) for the training and test sets respectively, confirming that the VAE-learned representations of the DFT wavefunctions contain sufficient information to describe the non-local GW self-energy. We emphasize that the GW correction is more than a simple scissor shift. The average standard deviation of GW corrections across all bands and k-points in each material is 0.54 eV, which is much larger than the prediction accuracy of 0.11 eV (see Supplementary Fig. 4). Figure 2e shows the ML predicted GW band structure of monolayer MoS₂, which agrees remarkably with results obtained by explicit GW calculation (red circles in the bandstructure). Due to the generative power of the VAE latent space⁶³, even in the absence of electronic states along the high-symmetry path and the material MoS₂ in the training set, our model can accurately predict a smooth GW bandstructure along $\Gamma - M - K - \Gamma$ for MoS₂. Figure 2f shows how each individual input affects the accuracy of the GW NN model. We find that excluding the latent space representation of the wavefunction $|n, \mathbf{k}\rangle$ in training significantly reduces the R^2 value by 0.5, directly showing the importance of VAE representation of individual KS states in GW prediction. GW calculations can also be accomplished through the self-consistent Sternheimer equation, utilizing solely the occupied electronic states^{102–104}, so in principle, inclusion of empty state information through the use of super states is not strictly necessary. Here, excluding the super states encoding the empty states reduces the R^2 value by 0.05. More details regarding training dynamics are presented in Supplementary Figs. 5–7.

Interpretability and generative power of the latent space

Next, to understand the generative power of our model, we open up the black box of the VAE and explore the meaning of the latent space obtained from unsupervised learning. We utilize a 3D t-distributed stochastic neighbor embedding (t-SNE) to visualize the VAE latent space for the electronic structure. As depicted in Fig. 3a, the circles linked by blue dashed lines represent the latent points of the first valence band states along the Γ -K-M- Γ high symmetry path of monolayer MoS₂. These points form a continuous and enclosed trajectory in the latent space, corresponding to the smooth closed path in the k-space shown in the lower bandstructure of Fig. 3b. For comparison, the latent trajectories from two other TMD monolayers WS₂ (blue) and CrS₂ (red), are also shown in Fig. 3a. The similarity in the electronic structures of these three TMDs (see Fig. 3b) is encoded in the similarities of the paths in latent space. To further demonstrate the generative power of the VAE, Fig. 3b shows the continuous evolution of 15 VAE generated real-space KS wavefunction moduli from M-K in the first valence band of monolayer MoS₂. The generated wavefunctions are constructed using the VAE decoder, which processes sampled points along a smooth curve from K to M in the latent space, denoted by the blue arrows in Fig. 3a. Notably, even though the VAE training set is only

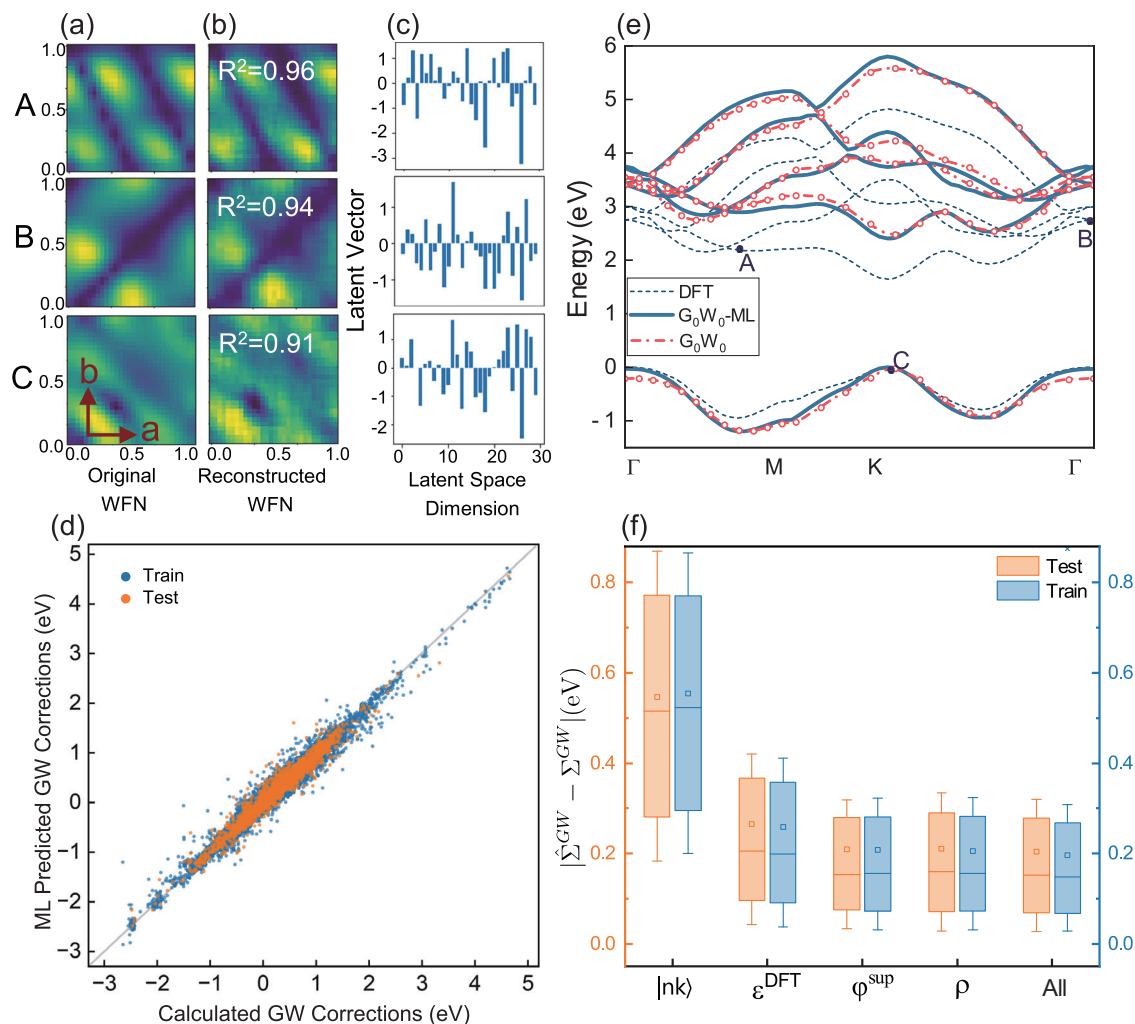


Fig. 2 | Application of crystal variational autoencoders (VAE) and downstream GW predictions. **a** Density functional theory (DFT) calculated wavefunction of states A, B and C in MoS₂ (see (e)) used as input to VAE encoder, in crystal coordinates in units of the lattice vector. WFN denotes wavefunction. **b** VAE reconstructed wavefunctions of states A, B and C through latent space decoding. **c** Low-dimensional variational mean latent space for states A, B and C. **(d)** Parity plot comparing the exact calculated values (x-axis) to the Machine learning (ML) predicted values (y-axis) of the GW correction for individual states. Blue (orange) dots represent training (test) sets. The mean absolute error (MAE) for the training set and test set are 0.06 and 0.11 eV respectively. The total number of data points in training (test) set are 19801(2201). **e** ML predicted GW band structures (blue solid curve) and calculated PBE band structures (blue dashed line) for monolayer MoS₂. The red circles are the exactly calculated quasiparticle (QP) energies from GW self-

energy. The red dashed lines are the interpolated GW band structures. **f** Whisker plot of ML predicted GW error without utilizing representations of KS states $|nk\rangle$, DFT energies ϵ^{DFT} , superstates ϕ^{super} or charge density ρ . "All" denotes using all information. The orange (blue) boxes represent the test (training) set. The total number of data points under each box are 19801(2201) for training(test) sets. Each box plot displays the absolute error distribution of the self energies $|\Sigma_{nk}^{GW} - \Sigma_{nk}^{GW}|$, highlighting the median (box center line, i.e. Q2), 25th – 75th percentiles of dataset (lower and upper boundary of box, i.e. Q1 and Q3), mean (small square within box) and outlier cutoff (lower and upper whisker mark are defined as $Q1 - (Q3 - Q1) \times 0.2$ and $Q1 + (Q3 - Q1) \times 0.2$, so all datapoints beyond the range are considered outliers). The training process spans 1,000 epochs. Source data are provided as a Source Data file.

comprised of states sampled on a uniform k -point grid, and these three TMDs are entirely excluded from the training set, the generated wavefunction in Fig. 3b can still have a high R^2 around 0.9, showing the VAE accurately generates wavefunctions that the training has never seen. The underlying reason is that the three trajectories in Fig. 3a lie on the smooth latent surface learned by the well-trained VAE, where similar states are mapped to neighboring points in the latent space. Therefore, the smoothness and regularity of the VAE latent space can be used to quantify the extrapolative and generative power of our model and suggest that any sampled points from the learned latent surface are physically meaningful and can be reconstructed back to a wavefunction.

Additionally, to explain the success of our model in predicting GW corrections at arbitrary k -points, we investigate how the supervised training at certain k -points affects the GW prediction at untrained k -

points in the BZ. Figure 3c illustrates the average GW MAE for different k -points, as predicted by a dense NN that is only trained with GW energies at the Γ point (left panel) or a single finite momentum k -point (right panel). Notably, the prediction error tends to be lower for k -points in proximity to the trained k -point, contrasting with higher errors for those farther away from the trained k -point. This trend originates from the smoothness of the VAE latent space, indicating that two states that are close in the latent space should contribute comparably to the GW corrections, as they are also expected to be close in k -space as shown in Fig. 3b. This explains the capability of our model to accurately predict k -resolved GW band structures when trained with limited and uniform k -grid data.

Finally, to further verify the manifold assumption for electronic structure in real materials and explore how GW energies correlate to the VAE-coded representation, we apply the 3D t-SNE on top of the

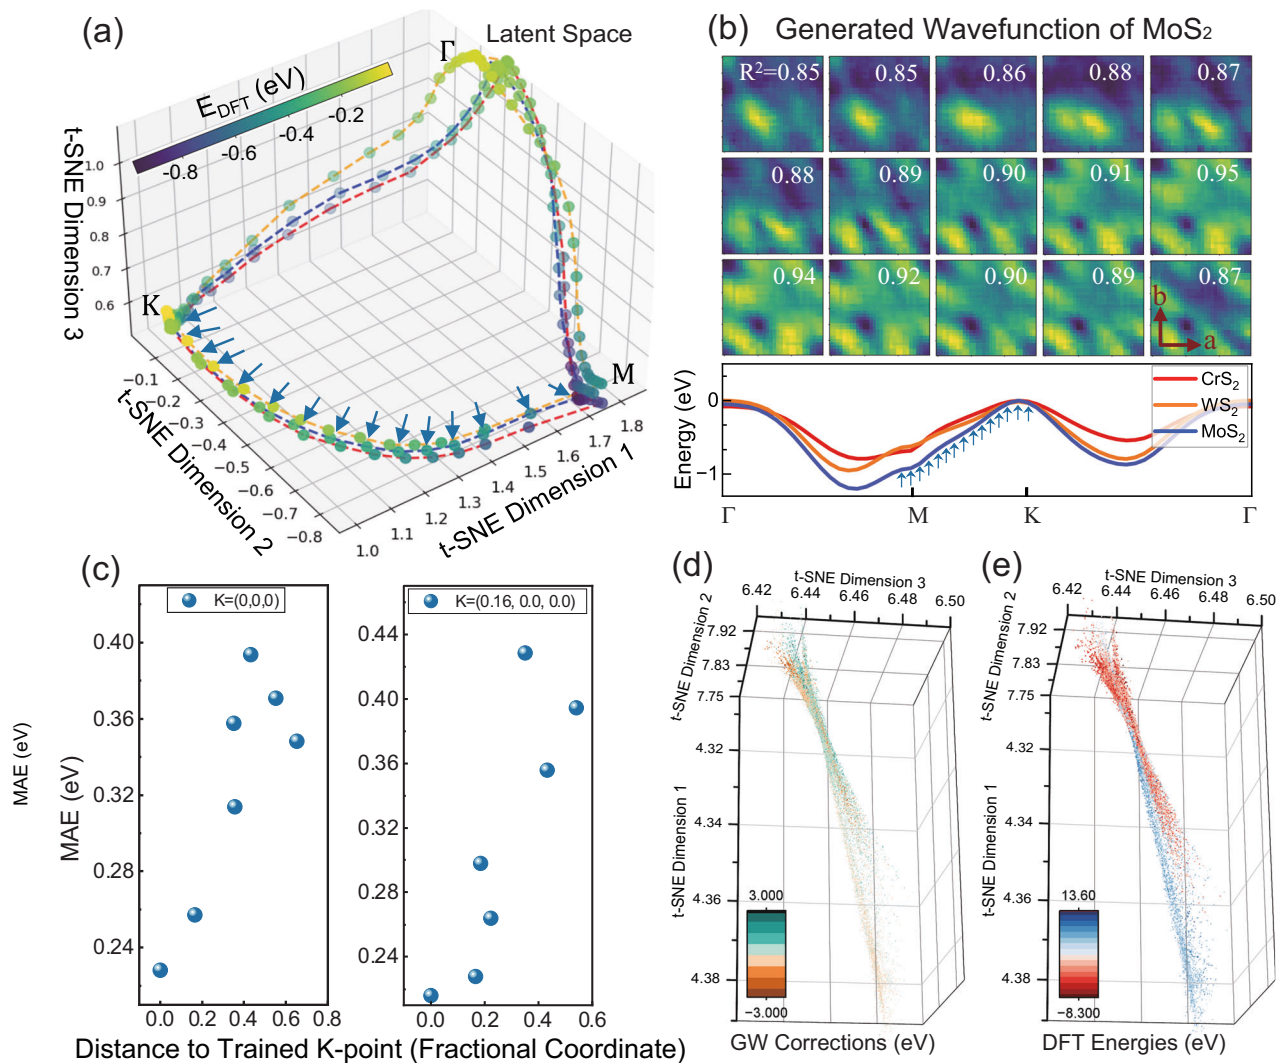


Fig. 3 | Interpretation and generative ability of low dimensional representation of Kohn-Sham (KS) states learnt by variational autoencoder (VAE). **a** 3D t-distributed stochastic neighbor embedding (t-SNE) visualization of the μ latent space for KS wavefunctions in the first valence band along $\Gamma - M - K - \Gamma$ for the transition metal dichalcogenide (TMD) monolayers MoS₂, WS₂ and CrS₂, as shown in the lower band structures in **(b)**. The dots (latent points) are mapped from the high-dimensional electronic wavefunction with color coded by Density Functional Theory (DFT) energies. The blue, orange and red dashed lines connect latent points corresponding to the KS states from the first valence band of MoS₂, WS₂ and CrS₂ respectively. Blue arrows denote latent points of 15 states along $M - K$ from MoS₂, as shown in **(b)**. **b** Generated real-space wavefunction moduli decoded from latent points, which are sampled from the smooth latent curve of monolayer MoS₂, for each KS state from $M - K$ in the first valence band. The inset white number

represents the R^2 value, showing the correlation with the explicitly calculated wavefunction. The order of states are in row-major order and from left to right. The lower figure represents the DFT bandstructures of monolayer MoS₂ (Blue), WS₂ (Orange) and CrS₂ (Red). The blue arrows indicate the KS states shown in the upper part. **c** Mean absolute error (MAE) of GW prediction for different k-points and generative power for k-point interpolation. Here, the model is trained exclusively using GW energies at the Γ (left panel) point or $\mathbf{k}=(0.1667, 0, 0)$ (right panel). The x-axis is the distance from the trained k-point to the untrained k-points in reciprocal space. **d-e** 3D t-SNE components of VAE latent μ vector of 22002 KS wavefunction with GW labels, whose color are mapped to the GW correction $\epsilon^{GW} - \epsilon^{DFT}$ corrections **(d)** and DFT energies ϵ^{DFT} **(e)** respectively. Source data are provided as a Source Data file.

latent mean μ space for all 22002 states with GW labels, whose colors are coded to their calculated GW-correction (Fig. 3d) and DFT energies (Fig. 3e). The t-SNE analysis provides direct evidence for the manifold assumption that the VAE effectively maps the latent μ space of the electronic wavefunction to a smooth manifold even for different materials. More intriguingly, even though the unsupervised VAE learning procedure is entirely independent of GW labels, the magnitude of both the GW corrections and DFT energies exhibit a distinct pattern with different branches in the manifold corresponding to the magnitude of different GW corrections or DFT energies, as opposed to a random distribution. As a result, the t-SNE analysis serves as a crucial validation, demonstrating that unsupervised representation learning can effectively capture the inherent statistical correlation between GW

energies and the electronic wavefunction. Thus, the unsupervised VAE acts as a form of pre-learning, which can significantly lower the barrier for the subsequent supervised learning of GW self energies.

In summary, we demonstrate that a properly designed VAE model can unsupervisedly learn KS DFT wavefunctions, compressing them as a low-dimensional latent space, in a way that preserves physical observables like total charge density and fundamental information needed for downstream prediction of excited-state properties. Since our model autonomously determines the crucial information for preservation through the unsupervised reconstruction of wavefunction data, it can establish a low-dimensional representation that avoids limitations due to feature engineering and selection. The representations preserve the translational invariance, discrete rotational

invariance, and periodic boundary conditions of crystals with different cell sizes and symmetry. Our VAE model achieves a R^2 of 0.92 when reconstructing the wavefunctions in the test set. To further test the effectiveness of the VAE representation of KS states, we train a supervised dense NN for downstream prediction of GW self-energies on top of the latent space of KS states. The resulting model demonstrates a remarkably low MAE of 0.11 eV on a test set and can be used to predict both arbitrary k-points and materials held back from the training set. While other ML models have been used to predict GW bandstructure, the main advantage of our approach lies in our ability to interpret the smoothness of the latent space in relation to the completeness of our training set across both k-space and the space of material chemical and structural composition. This eliminates unphysical oscillations seen in previous models and allows us to confidently evaluate the generalizability of our model to states outside the training set improving its generative power. The smooth evolution of KS states in k-space can be mapped to a smooth trajectory in the latent space, and the sampled points from this continuous trajectory are physically meaningful and can also be reconstructed, enabling the generation of wavefunctions (and related physical observables) at uncalculated points at no additional cost. Here, in this first demonstration, we focus on predicting GW bandstructures of 2D materials, but we expect this framework to be generalizable to other crystalline systems and downstream applications, given expanded training sets. Notably, VAE wavefunction generation suggests a general route to building ML models for many post-Hartree-Fock (HF) corrections, including quantum Monte Carlo (QMC)¹⁰⁵, Density Matrix Downfolding¹⁰⁶, coupled cluster singles and doubles (CCSD)^{107,108}, and complete active space configuration interaction (CASCI)^{109,110}, all of which rely on the construction of higher-order wavefunctions, Green's functions, and density matrices from DFT wavefunctions.

Methods

This section describes the details of construction of the translationally invariant variational autoencoder (VAE) model, downstream neural networks, and first principles density functional theory (DFT) and GW dataset. The mathematical proof of invariance, training dynamics, ground state applications of the latent space, additional benchmarks, and analysis of the pseudobands approximation are given in the Supplementary Notes. The atomic structure of materials used in this study are from the open database: Computational 2D Materials Database (C2DB)⁹⁹⁻¹⁰¹.

Translationally invariant VAE

Our model is developed utilizing the Pytorch package with parameters shown in Table 1¹¹¹. The first column categorizes each layer by its functionality, alongside the respective layer indices. Specifically, 'Conv2d' denotes a two-dimensional CNN layer, 'RotationalConv2d' denotes the two-dimensional rotational CNN layer. A kernel size of 3×3 and 320 channels per CNN layer are sufficient to recover KS states with an overall R^2 score exceeding 0.87. 'ReLU' stands for the rectified linear unit, an activation function renowned for its widespread application in neural networks¹¹². 'MaxPool2d' indicates the max pooling layer, while the 'Flatten' layer plays a pivotal role by transforming the multi-dimensional outputs from preceding layers into a format suitable for the subsequent fully connected layers. 'AdaptiveAvgPool2d' indicates the GAP layer by specifying output as 1×1 . The 'Linear' layer refers to the fully connected layer. 'Sigmoid' is mentioned as a non-linear activation function that generates an output ranging from 0 to 1¹¹³, which is ideal for modeling the reconstruction of wavefunction moduli. The architecture's layout is further detailed by listing the number of parameters contained within each model layer in the second column.

The input grid size (resolution) of wavefunctions in real space within a primitive unit cell depends on both the crystal lattice and the

Table 1 | VAE Parameters

Layer (type)	Parameter number
RotationalConv2d-1	115,520
MaxPool2d-3	0
Conv2d-4	921,920
ReLU-5	0
AdaptiveAvgPool2d-6	0
Flatten-7	0
Linear-8	19,260
Linear-9	48,608
ReLU-10	0
Unflatten-11	0
ConvTranspose2d-12	17,340
ReLU-13	0
ConvTranspose2d-14	9,640
Sigmoid-15	0
Total Trainable Parameters	1,132,288

The table shows the detail of each layer of variational autoencoder (VAE).

kinetic energy cutoff of the plane wave basis in the density functional theory (DFT) calculations. However, even if the number of trainable parameters in the convolutional neural network (CNN) layers remains constant regardless of the size of the input grid, the fully connected dense layer always requires the same input size from the output of the CNN layer. To address this, a commonly used preprocessing technique is cropping/resizing before passing input to CNN. However, this straightforward method does not adequately account for the freedom to choose the origin of the unit cell and the invariance of physical properties under discrete translational and rotational symmetries.

The main bottlenecks faced by using a simply rigid and uniform grid are (i) lack of preservation of translational invariance, (ii) lack of preservation of discrete rotational symmetry, (iii) improper handling of periodic boundary conditions (PBC), and (iv) lost information about unit cell size and symmetry. To address these issues and improve the generalizability of the VAE, we integrate a circular CNN layer, global average pooling layer, and rotational CNN technique in the encoder as shown in Fig. 1a, b. As a result, the latent vector $\mathbf{z}_i \in \mathbb{R}^P$ of KS state $|i\rangle$ are invariant to the unphysical degree of freedom of arbitrary translation and lattice constant selection. In addition, given the same crystal structure, \mathbf{z}_i are a representation of $|i\rangle$ in the whole real space, so it is identical for different supercells constructed from the same unit cell. i.e.,

$$\mathbf{z}_i = f_{NN}(\mathbf{z}_{1,i}) = f_{NN}(\mathbf{z}_{m,n,i})$$

where $\mathbf{z}_{m,n,i} \in \mathbb{R}^C$ are the output of the last CNN layer, C is the number of channels, and m, n are arbitrary positive integers defining the number of primitive cells in a supercell. f_{NN} is the fully connected neural network following the CNN. The mathematical demonstration, model parameters, and corresponding benchmark are detailed in Supplementary Note 1.

For the downstream prediction of GW bandstructures, we employ a vanilla neural network including three fully connected layers with 13,801 trainable parameters. To reduce the risk of training set dependence, we also adapt k-fold cross validation for our evaluating model, where the data set is evenly split into 10 subsets (see Supplementary Fig. 6c). To visualize the latent space learnt by the VAE, we employed 3D t-distributed stochastic neighbor embedding (t-SNE). The perplexity in t-SNE calculations is set as 10,000 (22,002 datapoints in dataset) to balance the preservation of the global and local structure of the high dimensional data.

DFT and GW dataset

To benchmark the predictive power of our model, we select 302 materials from the Computational 2D Materials Database (C2DB)^{99–101}, which was also used in the work of Knøsgaard et al.⁶⁰, for training and validation. The dataset includes both metals and semiconductors across all crystal systems (we note that previous ML for GW prediction in this database was restricted to the subset of semiconducting materials) with the number of atoms ranging from 3 to 4. For the unsupervised VAE training, our dataset is comprised of 68,384 DFT electronic states sampled on a $6 \times 6 \times 1$ uniform k-grid for 302 2D materials, which are randomly split into 90% training set and 10% test set. We include the same number of conduction states as valence states for each material in the VAE training, ensuring a comprehensive understanding of the electronic structure of both occupied and unoccupied states. For supervised training of GW corrections, the entire GW energy dataset is randomly partitioned into two subsets: 10% (2201 electronic states) is allocated as the test set, while the remaining 90% (19801 electronic states) is designated as the training set. We further completely exclude a small subset of 30 materials from the training set for the VAE and three monolayer TMD materials—MoS₂, WS₂, and CrS₂—from both the VAE and GW training datasets. The DFT and GW datasets are generated using GPAW^{114,115}. Comprehensive details regarding dataset, training dynamics and additional benchmarks are listed in Supplementary Note 5.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The data (KS wavefunction and GW energies of 302 materials) generated in this study have been deposited in the Materials Data Facility repository under <https://doi.org/10.18126/rpa1-yp91>. Source data are provided with this paper.

Code availability

Code associated with this work is available from the public github repository <https://github.com/bwhou1997/VAE-DFT> with <https://doi.org/10.5281/zenodo.13617033>¹¹⁶ and from Code Ocean¹¹⁷.

References

- Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nat. Phys.* **13**, 431–434 (2017).
- Broecker, P., Carrasquilla, J., Melko, R. G. & Trebst, S. Machine learning quantum phases of matter beyond the fermion sign problem. *Sci. Rep.* **7**, 8823 (2017).
- Ch'ng, K., Carrasquilla, J., Melko, R. G. & Khatami, E. Machine learning phases of strongly correlated fermions. *Phys. Rev. X* **7**, 031038 (2017).
- Zhang, Y. & Kim, E.-A. Quantum loop topography for machine learning. *Phys. Rev. Lett.* **118**, 216401 (2017).
- van Nieuwenburg, Evert P. L., Liu, Y.-H. & Huber, Sebastian D. Learning phase transitions by confusion. *Nat. Phys.* **13**, 435–439 (2017).
- Costa, N. C., Hu, W., Bai, Z. J., Scalettar, R. T. & Singh, R. R. P. Principal component analysis for fermionic critical points. *Phys. Rev. B* **96**, 195138 (2017).
- Ponte, P. & Melko, R. G. Kernel methods for interpretable machine learning of order parameters. *Phys. Rev. B* **96**, 205146 (2017).
- Carrasquilla, J., Torlai, G., Melko, R. G. & Aolita, L. Reconstructing quantum states with generative models. *Nat. Mach. Intell.* **1**, 155–161 (2019).
- Lohani, S., Searles, T. A., Kirby, B. T. & Glasser, R. T. On the experimental feasibility of quantum state reconstruction via machine learning. *IEEE Trans. Quant. Eng.* **2**, 1–10 (2021).
- Lohani, S., Kirby, B. T., Brodsky, M., Danaci, O. & Glasser, R. T. Machine learning assisted quantum state estimation. *Mach. Learn. Sci. Technol.* **1**, 035007 (2020).
- Scheurer, M. S. & Slager, R.-J. Unsupervised machine learning and band topology. *Phys. Rev. Lett.* **124**, 226401 (2020).
- Zhang, P., Shen, H. & Zhai, H. Machine learning topological invariants with neural networks. *Phys. Rev. Lett.* **120**, 066401 (2018).
- Gong, X. et al. General framework for E(3)-equivariant neural network representation of density functional theory Hamiltonian. *Nat. Commun.* **14**, 2848 (2023).
- Gasteiger, J., Becker, F. & Günnemann, S. Gemnet: Universal directional graph neural networks for molecules. *Adv. Neural Inform. Process. Syst.* **34**, 6790–6802 (2021).
- Batzner, S. et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **13**, 2453 (2022).
- Qiao, Z. et al. Informing geometric deep learning with electronic interactions to accelerate quantum chemistry. *Proc. Natl Acad. Sci. USA* **119**, e2205221119 (2022).
- Xie, T. & Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* **120**, 145301 (2018).
- Unke, O. T. & Meuwly, M. PhysNet: A neural network for predicting energies, forces, dipole moments, and partial charges. *J. Chem. Theory Comput.* **15**, 3678–3693 (2019).
- Li, H. et al. Deep-learning density functional theory Hamiltonian for efficient ab initio electronic-structure calculation. *Nat. Comput. Sci.* **2**, 367–377 (2022).
- Xia, R. & Kais, S. Quantum machine learning for electronic structure calculations. *Nat. Commun.* **9**, 4195 (2018).
- Fiedler, L. et al. Predicting electronic structures at any length scale with machine learning. *npj Comput. Mater.* **9**, 115 (2023).
- Jackson, N. E. et al. Electronic structure at coarse-grained resolutions from supervised machine learning. *Sci. Adv.* **5**, eaav1190 (2019).
- Schütt, K. T. et al. How to represent crystal structures for machine learning: towards fast prediction of electronic properties. *Phys. Rev. B* **89**, 205118 (2014).
- Guo, H., Zhang, X. & Lu, G. Shedding light on moiré excitons: a first-principles perspective. *Sci. Adv.* **6**, eaab5638 (2020).
- Cao, Y. et al. Unconventional superconductivity in magic-angle graphene superlattices. *Nature* **556**, 43–50 (2018).
- Andrei, E. Y. et al. The marvels of moiré materials. *Nat. Rev. Mater.* **6**, 201–206 (2021).
- Liu, D., Luskin, M. & Carr, S. Seeing moiré: convolutional network learning applied to twistronics. *Phys. Rev. Res.* **4**, 043224 (2022).
- Willhelm, D. et al. Predicting Van der Waals heterostructures by a combined machine learning and density functional theory approach. *ACS Appl. Mater. Interfaces* **14**, 25907–25919 (2022).
- Sobral, J. A., Obernauer, S., Turkel, S., Pasupathy, A. N. & Scheurer, M. S. Machine learning the microscopic form of nematic order in twisted double-bilayer graphene. *Nat. Commun.* **14**, 5012 (2023).
- Yang, H. et al. Identification and structural characterization of twisted atomically thin bilayer materials by deep learning. *Nano Lett.* **24**, 2789–2797 (2024).
- Tritsaris, G. A., Carr, S. & Schleder, G. R. Computational design of moiré assemblies aided by artificial intelligence. *Appl. Phys. Rev.* **8**, 031401 (2021).
- Louie, S. G., Chan, Y.-H., da Jornada, F. H., Li, Z. & Qiu, D. Y. Discovering and understanding materials through computation. *Nat. Mater.* **20**, 728–735 (2021).
- Stolz, S. et al. Spin-stabilization by coulomb blockade in a Vanadium Dimer in WSe₂. *ACS Nano* **17**, 23422–23429 (2023).
- Frey, N. C., Akinwande, D., Jariwala, D. & Shenoy, V. B. Machine learning-enabled design of point defects in 2D materials for

- quantum and neuromorphic information processing. *ACS Nano* **14**, 13406–13417 (2020).
35. Wang, C., Tan, X. P., Tor, S. B. & Lim, C. S. Machine learning in additive manufacturing: state-of-the-art and perspectives. *Addit. Manuf.* **36**, 101538 (2020).
36. Arrigoni, M. & Madsen, G. K. H. Evolutionary computing and machine learning for discovering of low-energy defect configurations. *npj Comput. Mater.* **7**, 71 (2021).
37. Cubuk, E. D. et al. Identifying structural flow defects in disordered solids using machine-learning methods. *Phys. Rev. Lett.* **114**, 108001 (2015).
38. Schattauer, C., Todorović, M., Ghosh, K., Rinke, P. & Libisch, F. Machine learning sparse tight-binding parameters for defects. *npj Comput. Mater.* **8**, 116 (2022).
39. Wu, X., Chen, H., Wang, J. & Niu, X. Machine learning accelerated study of defect energy levels in Perovskites. *J. Phys. Chem. C* **127**, 11387–11395 (2023).
40. Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138 (1965).
41. Mattuck, R. D. *A Guide to Feynman Diagrams in the Many-Body Problem* 2nd edn, Vol. 464 (Dover Publication, 1992).
42. Cohen, M. L. & Louie, S. G. *Fundamentals of condensed Matter Physics* Vol. 460 (Cambridge University Press, 2016).
43. Aryasetiawan, F. & Gunnarsson, O. The GW method. *Rep. Prog. Phys.* **61**, 237 (1998).
44. Reining, L. The GW approximation: content, successes and limitations. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **8**, e1344 (2018).
45. Carleo, G. et al. Machine learning and the physical sciences. *Rev. Modern Phys.* **91**, 045002 (2019).
46. Schleder, G. R., Padilha, A. C., Acosta, C. M., Costa, M. & Fazzio, A. From DFT to machine learning: recent approaches to materials science—a review. *J. Phys. Mater.* **2**, 032001 (2019).
47. Li, Y., Dong, R., Yang, W. & Hu, J. Composition based crystal materials symmetry prediction using machine learning with enhanced descriptors. *Comput. Mater. Sci.* **198**, 110686 (2021).
48. Sonpal, A., Afzal, M. A. F., An, Y., Chandrasekaran, A. & Halls, M. D. Benchmarking machine learning descriptors for crystals. In *Machine Learning in Materials Informatics: Methods and Applications* (ed. An, Y.) 1416 (American Chemical Society, 2022).
49. Patala, S. Understanding grain boundaries – the role of crystallography, structural descriptors and machine learning. *Comput. Mater. Sci.* **162**, 281–294 (2019).
50. Himanen, L. et al. Dscribe: Library of descriptors for machine learning in materials science. *Comput. Phys. Commun.* **247**, 106949 (2020).
51. Venturella, C., Hillenbrand, C., Li, J. & Zhu, T. Machine learning many-body green's functions for molecular excitation spectra. *J. Chem. Theory Comput.* **20**, 143–154 (2024).
52. Tao, Q., Xu, P., Li, M. & Lu, W. Machine learning for perovskite materials design and discovery. *Npj Comput. Mater.* **7**, 23 (2021).
53. Yilmaz, B. & Yildirim, R. Critical review of machine learning applications in perovskite solar research. *Nano Energy* **80**, 105546 (2021).
54. Matur, M. N., Nagappan, N., Rath, S. & Thomas, T. Prediction of nature of band gap of perovskite oxides (ABO₃) using a machine learning approach. *J. Mater.* **8**, 937–948 (2022).
55. Ryu, B., Wang, L., Pu, H., Chan, M. K. & Chen, J. Understanding, discovery, and synthesis of 2D materials enabled by machine learning. *Chem. Soc. Rev.* **51**, 1899–1925 (2022).
56. Lee, J., Seko, A., Shitara, K., Nakayama, K. & Tanaka, I. Prediction model of band gap for inorganic compounds by combination of density functional theory calculations and machine learning techniques. *Phys. Rev. B* **93**, 115104 (2016).
57. Rajan, A. C. et al. Machine-learning-assisted accurate band gap predictions of functionalized MXene. *Chem. Mater.* **30**, 4031–4038 (2018).
58. Na, G. S., Jang, S., Lee, Y.-L. & Chang, H. Tuplewise material representation based machine learning for accurate band gap prediction. *J. Phys. Chem. A* **124**, 10616–10623 (2020).
59. Zhuo, Y., Mansouri Tehrani, A. & Brgoch, J. Predicting the band gaps of inorganic solids by machine learning. *J. Phys. Chem. Lett.* **9**, 1668–1673 (2018).
60. Knøsgaard, N. R. & Thygesen, K. S. Representing individual electronic states for machine learning GW band structures of 2D materials. *Nat. Commun.* **13**, 468 (2022).
61. Zadoks, A., Marrazzo, A. & Marzari, N. Spectral operator representations. *arXiv* <https://doi.org/10.48550/arXiv.2403.01514> (2024).
62. Jimenez, D. R., Shakir M., Daan W. Stochastic backpropagation and approximate inference in deep generative models. *arXiv* <https://doi.org/10.48550/arXiv.1401.4082> (2014).
63. Kingma D. P., Welling M. Auto-encoding variational bayes. *arXiv* <https://doi.org/10.48550/arXiv13126114> (2013).
64. Rocchetto, A., Grant, E., Strelchuk, S., Carleo, G. & Severini, S. Learning hard quantum distributions with variational auto-encoders. *npj Quant. Inform.* **4**, 28 (2018).
65. Yin, J., Pei, Z. & Gao, M. C. Neural network-based order parameter for phase transitions and its applications in high-entropy alloys. *Nat. Comput. Sci.* **1**, 686–693 (2021).
66. Szoltdra, T., Sierant, P., Lewenstein, M. & Zakrzewski, J. Unsupervised detection of decoupled subspaces: many-body scars and beyond. *Phys. Rev. B* **105**, 224205 (2022).
67. Wetzel, S. J. Unsupervised learning of phase transitions: from principal component analysis to variational autoencoders. *Phys. Rev. E* **96**, 022140 (2017).
68. Hybertsen, M. S. & Louie, S. G. Electron correlation in semiconductors and insulators: band gaps and quasiparticle energies. *Phys. Rev. B* **34**, 5390–5413 (1986).
69. Hedin, L. New method for calculating the one-particle green's function with application to the electron-gas problem. *Phys. Rev.* **139**, A796–A823 (1965).
70. Qiu, D. Y., da Jornada, F. H. & Louie, S. G. Optical spectrum of MoS₂: many-body effects and diversity of exciton states. *Phys. Rev. Lett.* **111**, 216805 (2013).
71. Krizhevsky A., Sutskever I., Hinton G. E. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inform. Process. Syst.* **60**, 84–90 (2012).
72. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv* <https://doi.org/10.48550/arXiv14091556> (2014).
73. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
74. Szegedy, C. et al. Going deeper with convolutions. *arXiv* <https://doi.org/10.48550/arXiv.1409.4842> (2015).
75. Lin, M., Chen, Q. & Yan, S. Network in network. *arXiv* <https://doi.org/10.48550/arXiv13124400> (2013).
76. Schubert, S., Neubert, P., Pöschmann, J. & Protzel, P. Circular convolutional neural networks for panoramic images and laser data. In *2019 IEEE Intelligent Vehicles Symposium (IV)* (IEEE, 2019).
77. Jaderberg M., Simonyan K., Zisserman A. Spatial transformer networks. *Adv. Neural Inform. Process. Syst.* <https://doi.org/10.48550/arXiv.1506.02025> (2015).
78. Cohen, A. J., Mori-Sánchez, P. & Yang, W. Challenges for density functional theory. *Chem. Rev.* **112**, 289–320 (2012).
79. Burke, K. Perspective on density functional theory. *J. Chem. Phys.* **136**, 150901 (2012).

80. Becke, A. D. Perspective: Fifty years of density-functional theory in chemical physics. *J. Chem. Phys.* **140**, 18A301 (2014).
81. Deslippe, J. et al. BerkeleyGW: A massively parallel computer package for the calculation of the quasiparticle and optical properties of materials and nanostructures. *Comput. Phys. Commun.* **183**, 1269–1289 (2012).
82. Rojas, H. N., Godby, R. W. & Needs, R. J. Space-time method for Ab initio calculations of self-energies and dielectric response functions of solids. *Phys. Rev. Lett.* **74**, 1827–1830 (1995).
83. Bruneval, F. & Gonze, X. Accurate G. W. self-energies in a plane-wave basis using only a few empty states: towards large systems. *Phys. Rev. B* **78**, 085125 (2008).
84. Berger, J., Reining, L. & Sottile, F. Ab initio calculations of electronic excitations: collapsing spectral sums. *Phys. Rev. B* **82**, 041103 (2010).
85. Shih, B.-C., Xue, Y., Zhang, P., Cohen, M. L. & Louie, S. G. Quasiparticle band gap of ZnO: high accuracy from the conventional G^0W^0 approach. *Phys. Rev. Lett.* **105**, 146401 (2010).
86. Samsonidze, G., Jain, M., Deslippe, J., Cohen, M. L. & Louie, S. G. Simple approximate physical orbitals for GW quasiparticle calculations. *Phys. Rev. Lett.* **107**, 186404 (2011).
87. Dong, S. S., Govoni, M. & Galli, G. Machine learning dielectric screening for the simulation of excited state properties of molecules and materials. *Chem. Sci.* **12**, 4970–4980 (2021).
88. Morita, K., Davies, D. W., Butler, K. T. & Walsh, A. Modeling the dielectric constants of crystals using machine learning. *J. Chem. Phys.* **153**, 024503 (2020).
89. Zauchner, M. G., Horsfield, A. & Lischner, J. Accelerating GW calculations through machine-learned dielectric matrices. *npj Comput. Mater.* **9**, 184 (2023).
90. Hornik, K. Approximation capabilities of multilayer feedforward networks. *Neural Netw.* **4**, 251–257 (1991).
91. Cybenko, G. Approximation by superpositions of a sigmoidal function. *Math. Control Sign. Syst.* **2**, 303–314 (1989).
92. Lin, B., He, X. & Ye, J. A geometric viewpoint of manifold learning. *Appl. Inform.* **2**, 3 (2015).
93. Bengio, Y., Courville, A. & Vincent, P. Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1798–1828 (2013).
94. Psenka, M., Pai, D., Raman, V., Sastry, S. & Ma, Y. Representation learning via manifold flattening and reconstruction. *J. Mach. Learn. Res.* **25**, 132 (2024).
95. Pandey, A., Fanuel, M., Schreurs, J. & Suykens, J. A. Disentangled representation learning and generation with manifold optimization. *Neural Comput.* **34**, 2009–2036 (2022).
96. Altman, A. R., Kundu, S. & da Jornada, F. H. Mixed stochastic-deterministic approach for many-body perturbation theory calculations. *Phys. Rev. Lett.* **132**, 086401 (2024).
97. Del Ben, M. et al. Large-scale GW calculations on pre-exascale HPC systems. *Comput. Phys. Commun.* **235**, 187–195 (2019).
98. Gao, W., Xia, W., Gao, X. & Zhang, P. Speeding up GW calculations to meet the challenge of large scale quasiparticle predictions. *Sci. Rep.* **6**, 36849 (2016).
99. Gjerding, M. N. et al. Recent progress of the computational 2D materials database (C2DB). *2D Mater.* **8**, 044002 (2021).
100. Haastrup, S. et al. The computational 2D materials database: high-throughput modeling and discovery of atomically thin crystals. *2D Mater.* **5**, 042002 (2018).
101. Rasmussen, A., Deilmann, T. & Thygesen, K. S. Towards fully automated GW band structure calculations: what we can learn from 60.000 self-energy evaluations. *npj Comput. Mater.* **7**, 22 (2021).
102. Giustino, F., Cohen, M. L. & Louie, S. G. GW method with the self-consistent Sternheimer equation. *Phys. Rev. B* **81**, 115105 (2010).
103. Umari, P., Stenuit, G. & Baroni, S. GW quasiparticle spectra from occupied states only. *Phys. Rev. B* **81**, 115104 (2010).
104. Govoni, M. & Galli, G. Large scale GW calculations. *J. Chem. Theory Comput.* **11**, 2680–2696 (2015).
105. Foulkes, W. M., Mitas, L., Needs, R. & Rajagopal, G. Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.* **73**, 33 (2001).
106. Zheng, H., Changlani, H. J., Williams, K. T., Busemeyer, B. & Wagner, L. K. From real materials to model hamiltonians with density matrix downfolding. *Front. Phys.* **6**, 43 (2018).
107. Bartlett, R. J. & Musiał, M. Coupled-cluster theory in quantum chemistry. *Rev. Mod. Phys.* **79**, 291–352 (2007).
108. Purvis, G. D. & Bartlett, R. J. A full coupled-cluster singles and doubles model: the inclusion of disconnected triples. *J. Chem. Phys.* **76**, 1910–1918 (1982).
109. Roos, B. O., Taylor, P. R. & Sigbahn, P. E. A complete active space SCF method (CASSCF) using a density matrix formulated super-CI approach. *Chem. Phys.* **48**, 157–173 (1980).
110. Knowles, P. J. & Handy, N. C. A new determinant-based full configuration interaction method. *Chem. Phys. Lett.* **111**, 315–321 (1984).
111. Paszke, A. et al. Pytorch: An imperative style, high-performance deep learning library. *Adva. Neural Inform. Process. Syst.* **32**, 8026–8037 (2019).
112. Hinton, G. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv* <https://doi.org/10.48550/arXiv.1207.0580> (2012).
113. Bishop C. M., Nasrabadi N. M. *Pattern Recognition and Machine Learning* 1st edn, Vol. 778 (Springer, 2006).
114. Mortensen, J. J. et al. GPAW: An open python package for electronic structure calculations. *J. Chem. Phys.* **160**, 092503 (2024).
115. Larsen, A. H. et al. The atomic simulation environment—a python library for working with atoms. *J. Phys. Condensed Matter* **29**, 273002 (2017).
116. Hou, B., Wu, J. & Qiu D. Y. Code fore unsupervised representation learning of Kohn-Sham states and consequences for downstream predictions of many-body effects'. *Zenodo* <https://doi.org/10.5281/zenodo.13617033> (2024).
117. Hou, B., Wu, J. & Qiu D. Y. Code ocean for 'unsupervised learning of individual Kohn-Sham states: interpretable representations and consequences for downstream predictions of many-body effects'. *arXiv* <https://doi.org/10.48550/arXiv.2404.14601> (2024).

Acknowledgements

We gratefully acknowledge helpful discussions with Mr. Xingzhi Sun at Yale University. This work was primarily supported by the U.S. DOE, Office of Science, Basic Energy Sciences under Early Career Award No. DE-SC0021965 (DYQ). Excited-state code development was supported by Center for Computational Study of Excited-State Phenomena in Energy Materials (C2SEPEM) at the Lawrence Berkeley National Laboratory, funded by the U.S. DOE, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division, under Contract No. DE-CO2-05CH11231 (DYQ). The calculations used resources of the National Energy Research Scientific Computing (NERSC), a DOE Office of Science User Facility operated under contract no. DE-AC02-05CH11231; the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS), which is supported by National Science Foundation grant number ACI-1548562; and the Texas Advanced Computing Center (TACC) at The University of Texas at Austin.

Author contributions

B.H. and D.Y.Q. conceived and designed the workflow of the project. B.H. and J.W. developed the Python code for VAE model and downstream prediction NNs. All authors analyzed the data and wrote the manuscript. D.Y.Q. supervised the work and helped interpret the results.

Competing interests

The authors declare no competing interests

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-024-53748-7>.

Correspondence and requests for materials should be addressed to Diana Y. Qiu.

Peer review information *Nature Communications* thanks Shivesh Pathak and the other anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

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

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

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024