

Defect diffusion graph neural networks for materials discovery in high-temperature, clean energy applications

Lauren Way^{1,†}, Catalin D. Spataru^{1,†}, Reese Jones¹, Dallas R. Trinkle², Andrew J. E. Rowberg³, Joel B. Varley³, Robert B. Wexler⁴, Christopher M. Smyth⁵, Tyra C. Douglas⁵, Sean R. Bishop⁵, Elliot Fuller¹, Anthony H. McDaniel¹, Stephan Lany⁶, and Matthew D. Witman^{1,†,*}

¹Sandia National Laboratories, Livermore, CA 94551, USA

² Department of Materials Science and Engineering, University of Illinois, Urbana-Champaign, IL 61801, USA

³Laboratory for Energy Applications for the Future (LEAF) and Quantum Simulations Group, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA

⁴Department of Chemistry and Institute of Materials Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130, USA.

⁵Sandia National Laboratories, Albuquerque, NM 87123, USA

⁶National Renewable Energy Laboratory, Golden, CO 80401, USA

**mwitman@sandia.gov*

† These authors contributed equally.

Abstract

The migration of crystallographic defects dictates material properties and performance for a plethora of technological applications. Density functional theory (DFT)-based nudged elastic band (NEB) calculations are a powerful computational technique for predicting defect migration activation energy barriers, yet they become prohibitively expensive for high-throughput screening of defect diffusivities. Without introducing hand-crafted (i.e., chemistry- or structure-specific) descriptors, we propose a generalized deep learning approach to train surrogate models for NEB energies of vacancy migration by hybridizing graph neural networks with transformer encoders and simply using pristine host structures as input. With sufficient training data, computationally efficient and simultaneous inference of vacancy defect thermodynamics *and* migration activation energies can be obtained to compute temperature-dependent vacancy diffusivities and to down-select candidates for more thorough DFT analysis or experiments. Thus, as we specifically demonstrate for potential water-splitting materials, candidates with desired defect thermodynamics, kinetics, and host stability properties can be more rapidly targeted from open-source databases of experimentally validated or hypothetical materials.

Introduction

Vacancies are the primary mediator for the diffusion of atoms in ions in a large number of solid materials,^{1,2} including crystalline oxides. The study of vacancy mobility and diffusion in bulk materials has myriad technological applications, including in devices that depend on low or high mobility. Both metal^{3,4} and non-metal^{5,6} oxides are of interest. Passivation layers, thermal and electrical insulators, as well high conductivity thermoelectric,

electronic, and electrochemical devices depend on stable atomic structures, which can undergo degradation due to vacancy migration when exposed to high temperatures and temperature gradients. A particularly consequential application is the geological confinement of hazardous waste^{7,8} where diffusion can impact the long-term stability of the confinement environment. Vacancy migration can also play a role in other degradation processes, such as corrosion.^{9,10} Diffusion is strongly affected by temperature, vacancy concentration, and other factors,^{1,2} and

significant diffusional anisotropy⁸ can serve either as a challenge or a benefit in technological design.

To guide materials design and optimization for targeted vacancy defect diffusion properties, first-principles calculations of defect migration activation energies^{11–14} based on nudged elastic band (NEB) calculations^{15,16} are often utilized to provide critical atomistic insights. However, they become extremely expensive due to their reliance on density functional theory (DFT), particularly due to the need for large host supercell structural representations to avoid spurious interactions between the periodic defect images. Even more difficult is the prediction of, for example, temperature-dependent vacancy diffusivity tensors,¹⁷ because one must compute the vacancy migration activation energy of all plausible symmetrically inequivalent paths, of which there can be many for compositionally or structurally complex materials. To address the large computational cost of these calculations, surrogate modeling approaches like cluster expansion (CE) have been employed to predict defect migration activation energies in a variety of contexts^{18–20} for specific types of structures and chemistries. Meanwhile, other approaches to mitigate this difficulty have focused on *a priori* estimation of low energy paths to lower the number of first-principles calculations needed subsequently.²¹

Substantial work has recently been devoted to the use of various machine learning surrogate models for vacancy or interstitial defect thermodynamic properties,^{22–27} and defect migration activation energies have been well correlated with physical descriptors.^{28–31} Here, we extend these efforts and derive *generalized* structure-property surrogate models of NEB-calculated defect migration (i.e., yielding both formation energies *and* activation energies) that can be applied across any structure type or chemistry space with an accuracy that is principally limited by quantity and diversity of NEB training data. Critically, the only input to the model is the relaxed *host structure* and a tuple of indices that represent the start and end atomic sites for the vacancy migration event; i.e., relaxed structures for each NEB image are not used. The basis for this model is a graph neural network architecture, whose convolution functional form could in principle be taken from one of the many variants recently proposed,^{32–36} which is then hybridized with a transformer encoder³⁷ to enforce various physical, NEB-required symmetry constraints in the model’s output.

To begin, we create a training database of neutral oxygen vacancy migration activation energies across a diverse structural and chemical space of metal oxides. Our trained model can then screen Materials Project (MP) structures³⁸ to predict the activation energies of all possible vacancy diffusion paths below a distance cutoff on the path length, and subsequently compute temperature-dependent diffusivity tensors from this data.¹⁷ In addition to generally identifying materials with interesting diffusivity characteristics, utility in a specific materials discovery application is demonstrated via a multi-

dimensional down-select of materials for thermochemical water-splitting (TCH)^{39–42} based on predicted oxygen vacancy defect thermodynamics, diffusivity, and host stability. While this dataset considers only neutral oxygen vacancies and does not delve into more complex defect-mediated diffusion pathways^{31,43} (e.g. interstitial-facilitated or knock-out diffusion), we nonetheless expect to identify promising materials that can be investigated on a case-by-case basis with more accurate methods. We therefore also conclude with some perspectives on how the proposed method can be improved with better accuracy and greater applicability for future materials discovery problems.

Results and Discussion

High-throughput NEB for oxygen vacancy diffusion. We first queried MP (using `mp_api` v0.33.3) for all oxides that only contain cations in the space of { Al, Ba, Bi, Ca, Cs, Ga, Ge, Hf, In, K, Li, Mg, Na, Nb, Rb, Si, Sn, Sr, Ta, Ti, Y, Zr }, yielding a total of ~3700 unique structures. Selection of this chemical space will be motivated later, where we describe our screening for TCH candidate materials. To compute oxygen vacancy diffusivity via a simple lattice site diffusion mechanism in a single oxide, one must first define an O-O neighbor cutoff radius (5.0 Å in this work) to identify all possible oxygen vacancy migration paths, then compute the migration activation energy (via NEB) for all symmetrically inequivalent paths. Depending on its structural/chemical complexity, a single oxide may easily contain $O(10^2)$ or more inequivalent O-O paths; thus, considering the expense of NEB calculations, exhaustively obtaining the migration activation energies needed for diffusivity predictions in a single material can be arduous. Such an approach would be infeasible for the $O(10^4)$ or more materials typically investigated in high-throughput screening studies.

An alternative approach is to screen materials using a surrogate model, which nonetheless requires a training dataset of oxygen vacancy migration activation energies computed with NEB. For an N atom host supercell structure, C_h (where h denotes “host”), with total energy E_h , we define an oxygen vacancy migration path by a tuple of site indices corresponding to the start and end sites, $p = (s, e)$. A vacancy is created at the start or end site and the atomic positions are relaxed at constant volume, yielding structures C_s and C_e with total energies E_s and E_e . These two structures and energies, which are themselves independent of the path by which they are connected, provide the fixed endpoints for an n -image NEB calculation, which yields a sequence of energies,

$$\{E_{p_i}\} = \{E_{p_0=s}, E_{p_1}, \dots, E_{p_n}, E_{p_{n+1}=e}\}. \quad (1)$$

Here $p_{i \in \{1, \dots, n\}}$ denotes the index of the image along an n -image NEB calculation, and $p_0 = s$ and $p_{n+1} = e$.

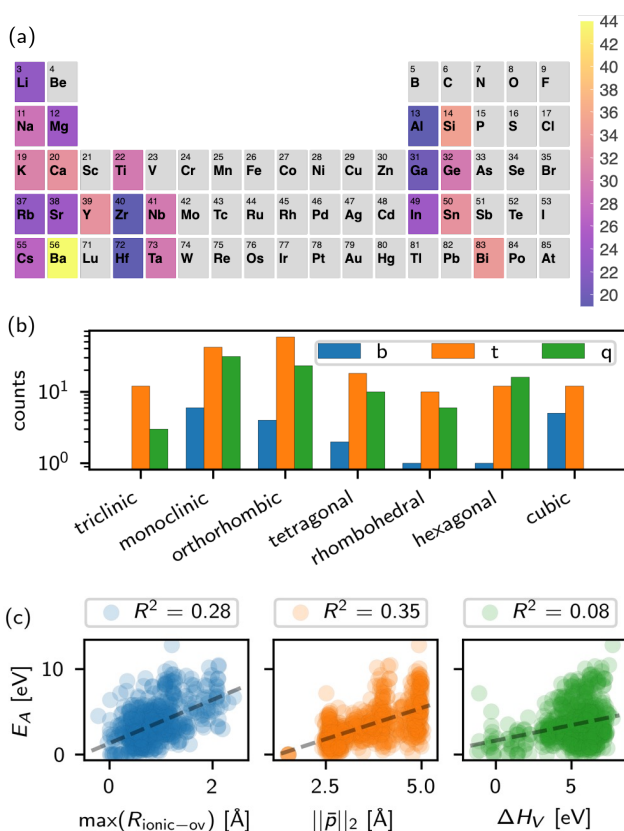


Figure 1: (a) Number of times an element appears in the 272 training data compounds with at least one converged NEB calculation. (b) Distribution of training data compounds among crystal systems and binary, ternary, or quaternary oxide chemical systems. (c) R^2 performance of some simple single variable linear regression models for E_A . Dashed black lines represent the line of best fit.

Relative formation energies are computed via

$$\Delta H_{p_i} = E_{p_i} - E_h + \mu_{\text{O}}^{\text{ref}} \quad (2)$$

where $\mu_{\text{O}}^{\text{ref}}$ is the reference oxygen chemical potential, such that the first and last values in the sequence correspond to relaxed vacancy formation energies. Therefore, $\Delta H_{p_0} = \Delta H_s$ and $\Delta H_{p_{n+1}} = \Delta H_e$ and are path independent. The forward and reverse migration activation energies can then be computed,

$$\begin{aligned} E_A^f &= \max(\{\Delta H_{p_1}, \dots, \Delta H_{p_n}\}) - \Delta H_{p_0} \\ E_A^r &= \max(\{\Delta H_{p_1}, \dots, \Delta H_{p_n}\}) - \Delta H_{p_{n+1}} \end{aligned} \quad (3)$$

Note that we compute only neutral oxygen vacancy formation and migration activation energies; the merits and limitations of this choice for our desired materials screening applications will be discussed later.

From the ~ 3700 queried oxides, we *a priori* prioritized structures and paths for NEB calculations to ensure diversity of chemistry, structure, and path lengths, as shown in Figure 1 (see Supplementary Section 1 for more details on

DFT-based NEB calculations and training data structure and path selection). We obtained 622 converged NEB calculations spanning 272 structures, each of which represents a unique chemical system. Figure 1a summarizes the chemical diversity of the training data, with color-coding used to represent the number of compounds in which each element appears. Figure 1b counts the number of compounds for each crystal system, grouping them into bins corresponding to binary (b), ternary (t), and quaternary (q) chemical systems.

Before delving into development of complex deep learning-based regression models, it is worthwhile to investigate whether simple linear regression models, crafted with physically-motivated features, may provide sufficient accuracy and to provide a benchmark against more complex models. Figure 1c indicates that single variable linear regression models based on simple descriptors can be derived with R^2 as high as ~ 0.3 . Here, $\max(R_{\text{ionic-ov}})$ denotes the maximum in ionic radii overlap (using default pymatgen⁴⁴ values) between any atom in the crystal structure and the diffusing O atom along the vector of the linearly interpolated path, $\bar{p} = \bar{r}_s + \eta(\bar{r}_e - \bar{r}_s)$, $\eta \in [0, 1]$, between end and start site coordinates, \bar{r}_e and \bar{r}_s , respectively, with $\|\bar{p}\|_2$ denoting the path length. Some interesting qualitative insights are observed, such as an intuitive lower bound to E_A for increasing $\max(R_{\text{ionic-ov}})$ and $\|\bar{p}\|_2$. However, these simple models are generally not accurate enough (all have $R^2 \leq 0.35$) for downstream diffusivity calculations in materials screening applications.

d²GNN surrogate model. We now seek a better surrogate model for Equation (2), i.e., a single-shot predictor of all NEB energies for a vacancy migration event from one crystallographic site to another. The model should facilitate rapid inference of NEB energies for all symmetrically inequivalent paths in not just one structure but the potentially tens of thousands oftentimes examined in high-throughput screening studies. We therefore seek to construct a model, parameterized by weights θ , that relies only on the relaxed host crystal structure, \mathcal{C}_h , and the path tuple p as input,

$$\{\Delta H_{p_i}\} = f_{\text{d}^2\text{GNN}}(\mathcal{C}_h, p; \theta). \quad (4)$$

The model must make predictions of the final energies without knowledge of the final structure;^{24,45} i.e., any structural information derived from the relaxed vacancy crystal structures (\mathcal{C}_s or \mathcal{C}_e) or image crystal structures (\mathcal{C}_{p_i}) cannot be utilized as input to the model. Figure 2a conceptualizes this procedure, whereby symmetrically distinct paths in a unit cell are projected within a supercell, decorated with path image nodes, and subsequently used to predict all sequence energies directly in one shot.

This one-to-many prediction must also address several subtle symmetry constraints, also illustrated schematically in Figure 2a. For symmetrically inequivalent paths, $p \neq p' \neq p''$, within the same crystal structure that share

symmetrically equivalent end points, the predicted vacancy formation energies in different sequence predictions must be equivalent,

$$\Delta H_{p_0} = \Delta H_{p'_0} = \Delta H_{p''_0}, \quad (5)$$

and therefore independent of the NEB images themselves. The converse is not true, however. Image energies may differ widely between symmetrically inequivalent paths that share equivalent endpoints,

$$\{\Delta H_{p_1, \dots, p_n}\} \neq \{\Delta H_{p'_1, \dots, p'_n}\} \neq \{\Delta H_{p''_1, \dots, p''_n}\}. \quad (6)$$

Finally, mirror symmetry should be preserved such that reversing p yields a reversed sequence of energies,

$$\text{rev}(\{\Delta H_{p_i}\}) = f_{\text{d}^2\text{GNN}}(\mathcal{C}_h, \text{rev}(p); \theta) \quad (7)$$

These constraints will also be addressed in the following model construction, which summarizes the architectures visualized in Figure 2(b-c).

(Step 1): Graph encoding. A graph is constructed for a given symmetrically distinct migration path that consists of $N + n$ nodes for the relaxed host crystal structure's N crystallographic sites, $\mathcal{V} = \{v_1, \dots, v_N\}$, and the migration path's n images, $\mathcal{P} = \{v_{p_1}, \dots, v_{p_n}\}$. In other words, \mathcal{V} must be identical across migration path graphs sharing the same base structure, but \mathcal{P} will differ between those paths that are symmetrically distinct. Image nodes' coordinates are linearly interpolated along the displacement vector between v_s and v_e .

(Step 2): Embedding. A crystal node's initial feature vector, $\mathbf{v}_i^{(0)} = f_{\text{embed}}(\mathbf{x}_i)$ is generated by an embedding function on a feature vector that depends only on the elemental identity of site i . In this case, \mathbf{x}_i is the ground state electron configuration utilized by SpookyNet,³⁶ and f_{embed} is a feed forward neural network layer. An image node's initial feature vector is equivalent to that of the path's endpoints such that $\mathbf{v}_{p_i}^{(0)} = \mathbf{v}_s^{(0)} = \mathbf{v}_e^{(0)}$.

(Step 3): Convolutions with directed message passing. GNNs generally use a series of convolution, or message-passing, functions to iteratively update nodes' feature vectors, from which a final property prediction can be made. The challenge here lies in deriving a convolution function that ensures that Equation (5) and Equation (6) are respected. In other words, to preserve Equation (5), feature vectors corresponding to crystal sites (which will ultimately be used to predict the vacancy formation energies, ΔH_{p_0} and $\Delta H_{p_{n+1}}$) cannot be influenced by feature vectors corresponding to image nodes (which will be used to predict the image energies, $\Delta H_{p_{\{0, \dots, n\}}}$). To ensure Equation (6), the converse must be true and the image energies clearly must depend on the path to which they belong, and therefore the endpoints and their vacancy formation energies. These considerations are also shown schematically in Figure 2.

We therefore propose that all node feature vectors are updated according to an identical convolution (or message passing) function, but that neighbor lists for v_i and

v_{p_i} differ in construction and are *directed*. Specifically, a crystal node's neighbor list only consists of other crystal nodes, $\{v_j \in \mathcal{V}\}$, while a path image node's neighbor list may consist of any nodes (either image or crystal), $\{v_j \in \mathcal{V} \cup \mathcal{P}\}$. A series of t convolution functions, $f_{\text{conv}}^{(t)}$, are identically applied to update both crystal and path image nodes,

$$\begin{aligned} \mathbf{v}_i^{(t+1)} &= f_{\text{conv}}^{(t)}\left(\mathbf{v}_i^{(t)}, \{v_j^{(t)} \in \mathcal{V}\}\right) \\ \mathbf{v}_{p_i}^{(t+1)} &= f_{\text{conv}}^{(t)}\left(\mathbf{v}_{p_i}^{(t)}, \{v_j^{(t)} \in \mathcal{V} \cup \mathcal{P}\}\right). \end{aligned} \quad (8)$$

A multitude of convolution architectures with varying complexity have been proposed in recent years.³²⁻³⁶ While f_{conv} in Equation (8) could in principle be based on a variety of these methodologies, given our small training dataset, we utilize a low parameter complexity CGCNN³² convolution function for practical model training purposes. More details of the CGCNN convolutional function and the small modifications used herein are provided in Supplementary Section 4.

(Step 4): Sequence extraction. Graph neural networks for global property predictions often apply a pooling function to aggregate crystal node feature vectors and produce a crystal size-invariant feature vector for downstream manipulation. Or, in the case of local property predictions like defect formation energies,²⁴ extraction of a single node isolates the feature vector of interest for downstream manipulation and property prediction. Here, after t' total convolutions, we extract an ordered sequence of nodes, $\mathbf{X}^{(t')} = \{\mathbf{v}_0^{(t')}, \mathbf{v}_{p_1}^{(t')}, \dots, \mathbf{v}_{p_n}^{(t')}, \mathbf{v}_{n+1}^{(t')}\}$ that features the NEB trajectory for downstream manipulation. Any crystallographic sites other than $p_0 = s$ and $p_{n+1} = e$ are discarded.

(Step 5): Sequence-to-sequence translation. Next we apply a sequence-to-sequence transformation block to preserve Equation (7). This could in principle adopt a variety of forms, although here we choose to exploit the self-attention mechanism of a (low parameter complexity) Transformer Encoder³⁷ to update the sequence values, $\mathbf{X}^{(t'')} = f_{\text{TE}}(\mathbf{X}^{(t')})$.

(Step 6): NEB energy decoder. The final NEB energy sequence, $\{\Delta H_{p_i}\}$, is predicted by one or more decoders, here chosen as a single-input/single-output multi-layer perceptron (MLP) as shown in Figure 2. The challenge lies in assigning a frame of reference to the image nodes and their relation to the energy sequence, i.e., visualized in "Path decoration" in Figure 2a. Figure 2b shows the first architecture variation we tested, where the NEB energies are predicted by the same decoder block. However, this formalism does not guarantee the satisfaction of Equation (5), since the self-attention mechanism of the transformer encoder enforces that $\mathbf{v}_{p_0}^{(t'')}$ and $\mathbf{v}_{p_{n+1}}^{(t'')}$ depend on $\mathbf{v}_{p_i=\{1, \dots, n\}}^{(t'')}$. A fundamentally better solution is proposed in Figure 2c, where more than one MLP decoder block is used. Conceptually, a first MLP is used to decode only the image feature vectors, while a second MLP

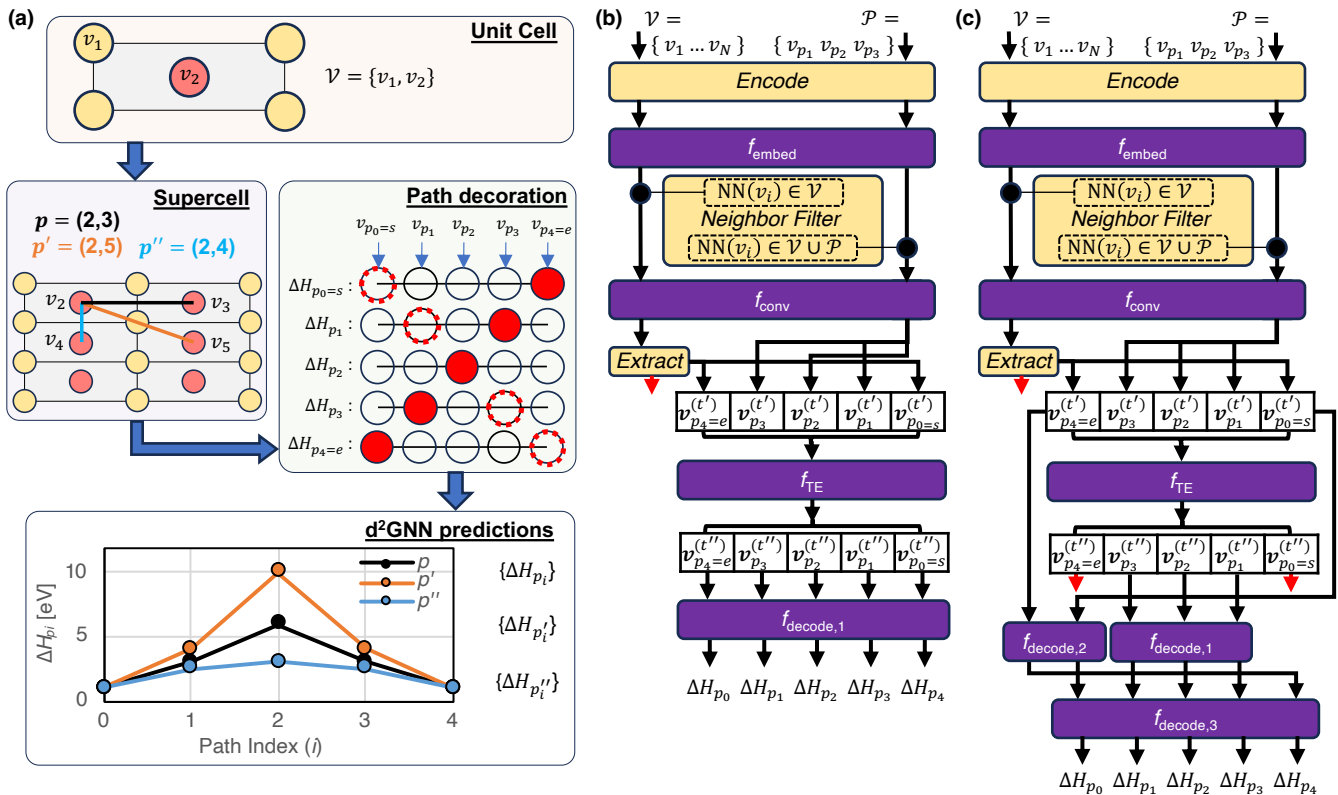


Figure 2: (a) d²GNN utilizes only the relaxed host structure as input. In the depicted unit cell with only one inequivalent red crystal site node (v_2), the supercell contains multiple inequivalent paths ($p \neq p' \neq p''$) between this symmetrically equivalent endpoint. A given path is decorated with nodes $v_{p_1 \dots p_3}$ that break the supercell symmetry. The ordering of the energy sequence ($0 \dots n + 1$) is with respect to the vacancy motion (dashed red circle), which is a reverse of the index that represents the occupation of the diffusing species (filled red circle) at a given image coordinate ($n + 1 \dots 0$). (b-c) The d²GNN architecture consists of six general steps: graph encoding of the crystal and image nodes, an embedding function for the node features, a nearest neighbor filter that enforces directed message passing in the convolution function, a sequence extraction, a sequence-to-sequence translation via Transformer Encoder (TE), and a decoder step. Purple indicates steps involving learnable parameters, while yellow denotes parameterless operations. Black arrows represent the flow of information, while red arrows represent a discard of features. The improvement of (c) over (b) involves physically constraining the model to preserve Equation (5) via independent decoders for the end states and the images.

is used to decode the endpoint feature vectors, but using only the pre-transformed sequence of feature vectors. Meanwhile, $v_{p_0}^{(t'')}$ and $v_{p_{n+1}}^{(t'')}$ from the second MLP are discarded. A third and final MLP decoder predicts the final energies. This second architecture, which preserves Equations (5) to (7), yields a large improvement in the model performance (see Supplementary Section 4).

Model performance and validation. The d²GNN model performance is evaluated by nested ($K = 10$, $L = 10$)-fold cross validation. In this scheme, unique structures are randomly split into $K = 10$ outer splits (90% train, 10% test), and all paths in a given structure are assigned accordingly.⁴⁶ Thus no paths from the same structure may be split between the train and test sets. For each k^{th} outer train split, the data is randomly divided into $L = 10$ inner splits (90% train, 10% test). A model is

trained for each inner train split (using 10% of the data as validation sets for early stopping), then all inner models are used to predict the outer test set example, $\{\Delta H_{p_i}\}_L$. We take the bootstrapped ensemble average as the final predicted value, $\Delta H_{p_i} = \langle \{\Delta H_{p_i}\}_L \rangle$, and the standard deviation, $\sigma_{p_i} = \sigma(\{\Delta H_{p_i}\}_L)$, as a heuristic metric for the prediction uncertainty.^{46,47}

Figure 3a-b show d²GNN vs. DFT test set parity for vacancy formation energies and path image energies, respectively. Figure 3c shows d²GNN vs. DFT predicted E_A , where some non-systematic error cancellation between the vacancy formation energies and the path image energies reduces the R^2 relative to the individual ΔH_{p_i} predictions. Here we show cross-validated test predictions only for stable materials with hull energies $E_H = 0$ eV/atom, since model performance begins to degrade for test set predictions of unstable materials with $E_H \gtrsim 0.025$

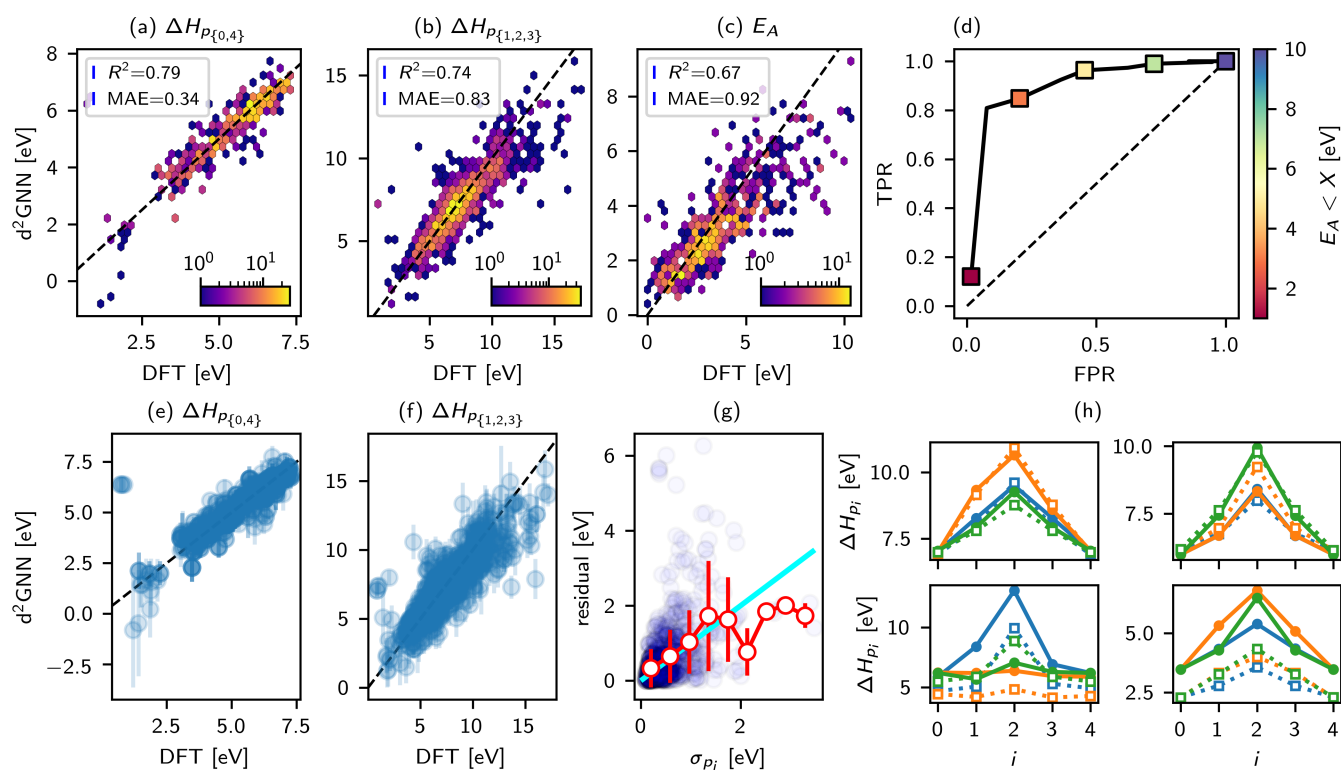


Figure 3: (a) Test set predictions for the vacancy formation energies. (b) Test set predictions for the image energies. (c) Test set predictions for the activation energies. (d) True positive rate vs. false positive rate when using the regression model as a classifier for a given threshold of $E_A < X$. (e-f) Test set predictions for the vacancy formation energies and image energies, respectively, with error bars corresponding to σ_{p_i} . (g) Residual error for all predictions vs. σ_{p_i} . After binning σ_{p_i} , red circles indicate the average of residuals within a given σ_{p_i} bin, and red error bars the standard deviation of residuals within a given σ_{p_i} bin. (h) Examples of two structures with well predicted NEB energies (top) and two poorly predicted NEB energies (bottom).

eV/atom, as discussed in the Supplementary Section 4. more).

Figure 3d shows the true positive rate (TPR) vs. false positive rate (FPR) when using the regression model as a binary classifier to identify whether a given path's E_A is less than a chosen threshold. For low thresholds of $E_A \sim 1$ eV, the model has about a 10 times higher TPR than FPR. We therefore expect the model to differentiate accurately between fast- and slow-diffusing materials, even if it is not able to provide quantitative comparisons among the most exceptional candidates (i.e., those with connected diffusion pathways with $E_A \ll 1$ eV). Figure 3e,f show the test set predictions with error bars corresponding to σ_{p_i} , which is a useful uncertainty metric *on average* that is correlated with the model's residual for a given prediction (Figure 3g).⁴⁷ Finally, Figure 3h shows two test set materials with low mean absolute error (MAE) predictions (top row) and two materials with high MAE predictions (bottom row). Detailed investigations on potential sources of error and uncertainty in the d²GNN predictions and the underlying NEB calculations themselves are discussed in Supplementary Section 3 and Supplementary Section 4 (e.g., paths with high atomic overlap, end points corresponding to split vacancies, and

High-throughput screening. As shown in Figure 4a for one example material, we first enumerate all possible inequivalent O-O paths in the crystal structure (up to 5 Å), then use d²GNN to predict $\{\Delta H_{p_i}\}$ for all p , and finally assign the predicted activation energies to each possible symmetrically equivalent path in the migration path network. A critical aspect of our screening is the bootstrapped uncertainty metric, σ_{p_i} , which we use to eliminate materials with uncertainty above a chosen threshold. From hereon we only show materials predictions where the average uncertainty across all paths in a given material, $\{p\}$, is below a threshold, $\langle \sigma_{p_i} \rangle_{\{p\}} < 0.75$ eV. Figure 4b shows the distribution of all computed activation energies and vacancy formation enthalpies ($\Delta H_V = \Delta H_{p_{0,4}}$) across these screened compounds, both of which are highly skewed. The raw screening data, including the high uncertainty predictions, also reveals anomaly/outlier predictions that can be readily identified as out-of-distribution materials with respect to our training data; these are therefore discarded on the basis of uncertainty (Supplementary Section 4).

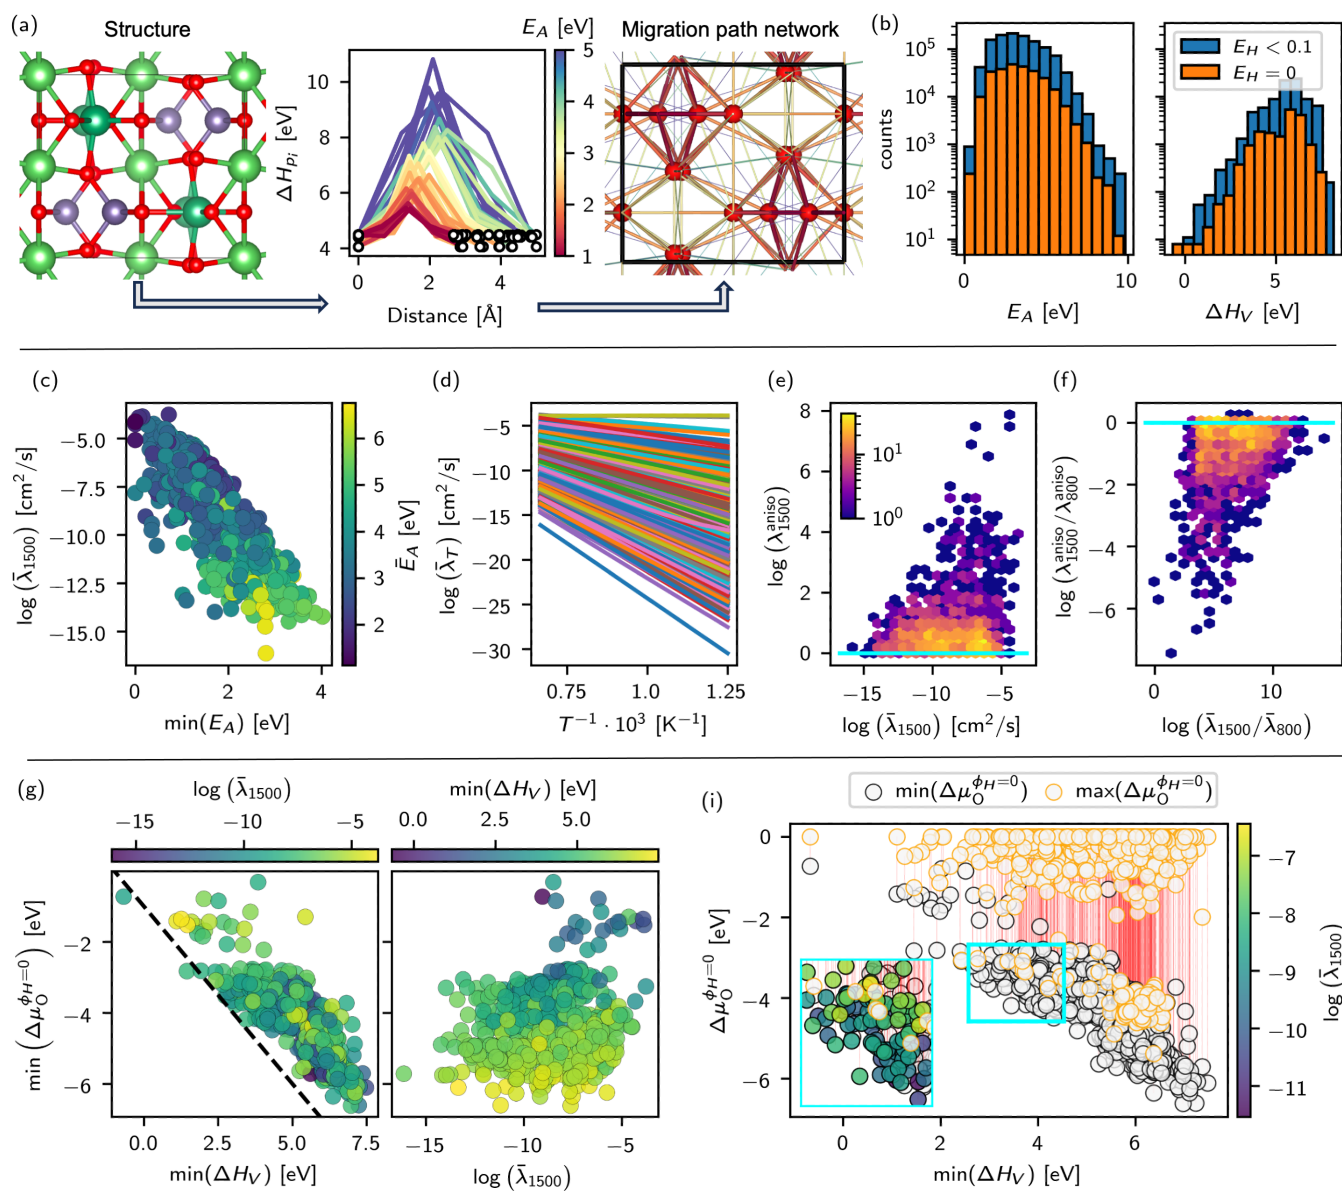


Figure 4: (a) Visualization of the mp-6456 structure (LiNbGeO_5), the predicted $\{\Delta H_{p_i}\}$ for all paths less than 5 Å (color-coded by E_A) with white circles for $\Delta H_{P_{\{0,4\}}}$, and migration path network (red spheres are oxygen atoms, edges are paths color-coded by E_A). (b) Distribution of E_A and ΔH_V across all screened structures, color-coded by E_H threshold. For only near hull materials $E_H < 0.025$ eV/atom: (c) Mean diffusivity vs. $\min(E_A)$ for all structures, color-coded by \bar{E}_A . (d) Arrhenius plot of the mean diffusivity vs. inverse temperature. (e) High-temperature diffusion anisotropy vs. mean diffusivity and (f) the anisotropy ratio between high and low temperatures vs. the change in mean diffusivity, where the colorbar represents the material count. For materials with non-null $\Delta\mu_{\text{O}}^{\phi_{\text{H}}=0}$: (g) minimum host stability vs minimum vacancy formation energy, color-coded by mean diffusivity. The dashed black line depicts $y = -x - 1$, an approximate lower bound to the host vs. defect stability correlation. (i) min host stability (gray circles) and max host stability (orange circles) are connected by a red line. Within the cyan inset, materials satisfying the strict set of TCH down-selection criteria are color-coded by mean diffusivity.

Our goal is then to predict the temperature-dependent oxygen vacancy diffusivity tensor, D , in a high-throughput manner to identify materials with useful (i.e., fast, slow, anisotropic) diffusivities and to elucidate relevant correlations between vacancy defect thermodynamics, kinetics, host stability, and chemistry. Given the va-

cancy sites, their energies, the connectivity of sites, and transition state energies, the master equation for the diffusivity is solved directly at each temperature T . The direct solution,¹⁷ implemented in the open source Onsager software (<https://github.com/DallasTrinkle/Onsager>), is more computationally efficient than a Monte

Carlo approach, and accuracy is limited only by round-off error. Using the d²GNN predictions, we can therefore obtain the temperature dependent oxygen vacancy diffusivity tensors, D_T , for each screened material. Units for all reported diffusivity values are in cm²/s.

Because our model is less accurate for meta-stable and unstable materials that are less likely to be able to be synthesized, we show D_T ($T \in [800, 1500]$ K) predictions in Figure 4c-f only for materials with $E_H < 0.025$ eV/atom. We denote the eigenvalues of D_T as λ_T , their average as the mean diffusivity, $\bar{\lambda}_T$, and the ratio between the maximum and minimum eigenvalues as the diffusion anisotropy, $\lambda_T^{\text{aniso}} = \lambda_T^{\text{max}}/\lambda_T^{\text{min}}$. For a given material, we plot $\bar{\lambda}_{1500}$ vs. the minimum activation energy $\min(E_A)$ and color-coded by the average of all paths, \bar{E}_A . As expected, a small $\min(E_A)$ is a necessary but insufficient criteria for a large diffusivity, as up to 8 orders of magnitude change in diffusivity can be observed for materials with $\min(E_A) < 0.5$ eV. This indicates the importance of a screening model that can estimate activation energies of all possible (including high E_A) paths in a structure as well. Figure 4d shows an Arrhenius plot of diffusivity and the orders of magnitude change among different materials and across temperatures, some of which deviate from Arrhenius behavior in the modeled temperature range. Figure 4e reveals that most materials have near-isotropic diffusivities, but some can have λ_T^{max} as much as 8 orders of magnitude greater than λ_T^{min} ; furthermore, among those with the highest mean diffusivities (e.g., $> 10^{-6}$ cm²/s), a higher percentage of materials tend also to have large λ_T^{aniso} . This anisotropy is also highly temperature dependent, i.e., some materials that are isotropic at high temperature become highly anisotropic at lower temperatures (Figure 4f).

Among our screened materials, we can isolate those with interesting predicted diffusivity properties (see Supplementary Section 6), a sampling of which is shown in Figure 5 and summarized in Table 1. These materials were selected for (a) large mean diffusivity, (b) small mean diffusivity, and (c) large mean *and* highly anisotropic diffusivities, respectively. Materials (a) and (c) respectively contain a 3d-connected and 2d-connected network of low E_A paths, while material (b) only possesses a periodically connected migration network with relatively high E_A .

Defect thermodynamics, kinetics, and host stability correlations. Some materials discovery exemplars necessitate identifying candidates with simultaneous target metrics for vacancy defect thermodynamics, diffusivity, *and* host stability, all of which may be interdependent. Defining the oxygen chemical potential difference with respect to the reference state, $\Delta\mu_{\text{O}} = \mu_{\text{O}} - \mu_{\text{O}}^{\text{ref}}$, and using MP-computed formation energies for all oxides in the screening space, we compute a stability range, denoted $\Delta\mu_{\text{O}}^{\phi_{\text{H}}=0} = [\Delta\mu_{\text{O}} \mid \phi_{\text{H}}(\Delta\mu_{\text{O}}) = 0]$. This yields the $\Delta\mu_{\text{O}}$ range over which a given material is stable, i.e., its grand energy above the hull, ϕ_{H} , equals zero. Figure 4g shows a

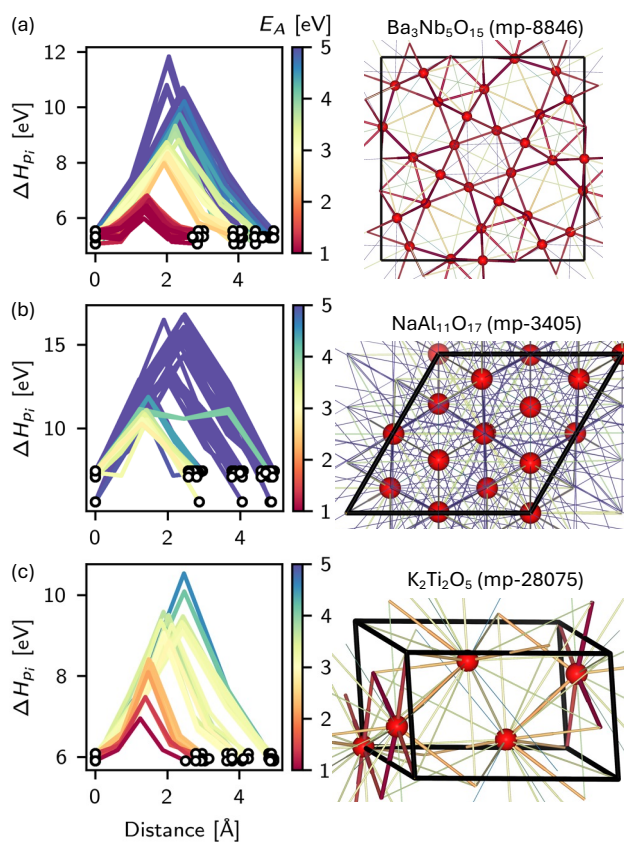


Figure 5: For three example structures with (a) large mean diffusivity, (b) small mean diffusivity, and (c) large mean *and* highly anisotropic diffusivities, we plot ΔH_{p_i} for all paths less than 5 Å (color-coded by E_A) and visualize the migration path network, with red spheres representing oxygen sites and edges representing possible migration paths (also colored by their E_A).

strong correlation between the minimum value of the stability range, $\min(\Delta\mu_{\text{O}}^{\phi_{\text{H}}=0})$, vs. the minimum vacancy formation energy in a given material, with an approximate lower bound of $\min(\Delta\mu_{\text{O}}^{\phi_{\text{H}}=0}) \approx -1$ eV $- \min(\Delta H_V)$. This trade-off constrains the space in which materials with desired host and defect stability can be designed. Meanwhile, neither criteria strongly correlates with the diffusivity $\bar{\lambda}_{1500}$, which can be tuned more independently.

Thermochemical hydrogen production materials screening. One example where all three down-selection criteria are important is the identification of metal oxides for thermochemical hydrogen (TCH) production, i.e., water-splitting via a direct two-step thermal redox cycle.^{39,41} Criterion #1 is that $\Delta H_V \in [2.5, 4.0]$ eV,⁴⁸ which has been addressed by a variety of surrogate model-based screenings. Criterion #2 necessitates that $\Delta\mu_{\text{O}}^{\phi_{\text{H}}=0}$ span, or at least intersect, the typical TCH redox oxygen chemical potential range between thermal reduction and water-splitting, $\Delta\mu_{\text{O}}^{\text{TCH}} = [-3.0, -2.5]$ eV.^{49,50} Finally, Criterion #3 for screening TCH materials, hitherto unexplored

Formula	ID	$\min(E_A)$	\bar{E}_A	$\log(\bar{\lambda}_{1500})$	$\log(\lambda_{1500}^{\text{aniso}})$
mp-8846	Ba ₃ Nb ₅ O ₁₅	0.6	3.3	-5.1	0.2
mp-3405	NaAl ₁₁ O ₁₇	2.8	6.6	-16.1	0.2
mp-28075	K ₂ Ti ₂ O ₅	1.0	3.0	-5.7	3.5

Table 1: A summary of the diffusivity related predictions (formula, MP ID, $\min(E_A)$, \bar{E}_A , $\max(E_A)$, $\bar{\lambda}_{1500}$, and $\lambda_{1500}^{\text{aniso}}$ for the three example structures shown in Figure 5.

via high-throughput screening, is that the oxygen vacancy diffusivity, shown in Figure 4i, be as fast as possible. The $\bar{\lambda}_{1500}$ color-coded materials in the cyan inset box are the small fraction of screened compounds that satisfy both Criteria #1 and #2. As shown by the red line connecting $\min(\Delta\mu_{\text{O}}^{\phi_H=0})$ and $\max(\Delta\mu_{\text{O}}^{\phi_H=0})$, many materials possess an excellent stability range for TCH, but this is inextricably correlated with vacancy formation energies that are too high, which eliminates a large swath of candidates. Only $O(10^2)$ candidates survive these downselection criteria, as examined in more detail in Supplementary Section 7.

The selection of the cationic chemical space in the present work excludes the $3d$ transition metals, with the exception of Ti, which occurs mostly as Ti^{4+} in a d^0 configuration. On one hand, this selection allowed us to limit the scope of this initial work for d²GNN to non-magnetic DFT calculations. On the other hand, this selection also addresses the hypothesis that high-performance materials for TCH might be found outside the realm of $3d$ transition metal oxides, where the vast majority of recent research efforts has been focused.⁵¹ While numerous discoveries of $3d$ oxides have been made over the past decade, especially including manganates,^{40–42} the general trend remains that these oxides split water only under dilute $\text{H}_2:\text{H}_2\text{O}$ conditions.^{41,52} This drawback can be attributed to an insufficient reduction entropy, resulting from the presence of repulsive defect interactions that reduce the (atomic) configurational entropy contribution⁵³ and from the lack of electronic entropy contributions.^{49,54}

The strong electron correlations in the $3d$ shell cause localized electronic states, which nevertheless experience significant interactions with the ligands,⁵⁵ as well as the existence of multiple oxidation states. As a consequence, the excess electrons introduced by the O vacancy defects tend to be strongly bound to the defect site, resulting in a large ionization energy, which precludes the electronic entropy contribution.⁵⁶ While there could be exceptions to this trend for suitable combinations of the transition metal electron configuration and the ligand coordination environment,⁵⁷ a potentially rewarding alternative materials search strategy could focus on cations with more delocalized non-magnetic electronic states, as included in the present work, where a smaller ionization energy could enable water splitting under much more concentrated $\text{H}_2:\text{H}_2\text{O}$ mixtures,⁵⁶ which is currently achieved only in CeO_2 . While the present work does not directly address the ionization energy and resulting electronic en-

ergy, it does provide the basis for selecting candidate materials for additional electronic structure studies with beyond-DFT methods that are required to address this question.

Conclusions

We have provided a generalized graph neural network modeling framework that can perform surrogate model predictions for NEB energies in arbitrary crystal structures and vacancy migration paths. Relying only on the host structure input and the specification of the end points corresponding to the vacancy migration hop, the model facilitates rapid inference on vacancy thermodynamics (formation energies) and kinetics (migration activation energies), while respecting relevant symmetry constraints thanks to an intermediary Transformer encoder and MLP decoder architecture. For a given material, all predicted activation energies of all paths up to a certain cutoff can be integrated with a first-principles, temperature-dependent calculation of the diffusivity tensor. Many materials can then be screened in a high-throughput manner to identify candidates with anomalous/interesting vacancy diffusivity properties. As a proof-of-concept, we conducted high-throughput screening of a large chemical space of metal oxides, combining neutral vacancy defect thermodynamics, kinetics, and host stability, to identify promising candidates for thermochemical hydrogen production via two-step water-splitting redox cycles.

Several outstanding challenges remain and point to exciting areas for improvement in future work. The chemical space across which the model has been trained is limited to 14 cations, so expanding the high-throughput NEB calculations to more diverse (but more challenging for DFT) materials like oxides with $3d$ transition metals are rare earths will improve its applicability for materials discovery across different domains. Similarly, our DFT training data only consists of NEB trajectories for oxygen vacancy migration, and this could be extended to different element types to again improve applicability in other materials discovery/modeling domains. The current form of the model also only accounts for vacancy migration events occurring between lattice oxygen sites, while hops between interstitial sites are often favorable,³¹ and could be accounted for in future data collection and an adjusted modeling framework. As discussed in more detail in Supplementary Section 2, this initial demonstra-

tion only focused on neutral oxygen vacancies, which often diffuse more slowly and have higher formation energies than charged vacancies (e.g., V_{O}^{2+}) in materials with band gaps. In such cases, we anticipate that our predictions represent at worst a lower bound on the vacancy diffusivity. Therefore, it would be beneficial in future work to repeat certain calculations with charged vacancies and to estimate the effect of charge on diffusivity predictions in a large-scale screening.⁵⁸ Finally, with the advent of machine learning methods that can model potential energy surfaces of inorganic materials across large portions of the periodic table,^{33,59–61} it will be interesting to compare the relative accuracy of this method to that of NEB calculations performed using machine learned force fields, which may not have been explicitly trained on high energy migration barriers.

Acknowledgements

The authors gratefully acknowledge research support from the Laboratory Directed Research and Development (LDRD) program at Sandia National Laboratories and the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Fuel Cell Technologies Office through the HydroGEN Consortium. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration (DOE/NNSA) under contract DE-NA0003525. This written work is authored by an employee of NTESS. The employee, not NTESS, owns the right, title and interest in and to the written work and is responsible for its contents. The Alliance for Sustainable Energy, LLC, operates the National Renewable Energy Laboratory (NREL) for the DOE under Contract No. DE-AC36-08GO28308. Any subjective views or opinions that might be expressed in the written work do not necessarily represent the views of the U.S. Government. The publisher acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this written work or allow others to do so, for U.S. Government purposes. The DOE will provide public access to results of federally sponsored research in accordance with the DOE Public Access Plan. The work at the Lawrence Livermore National Laboratory was performed under the auspices of the U.S. Department of Energy (DOE) under Contract No. DE-AC52-07NA27344.

Data Availability

Data will be made available for peer review and publicly released upon publication.

Code Availability

Data will be made available for peer review and publicly released upon publication.

References

- (1) H. Mehrer, *Diffusion in solids: fundamentals, methods, materials, diffusion-controlled processes*, Springer Science & Business Media, 2007, vol. 155.
- (2) P. Shewmon, *Diffusion in solids*, Springer, 2016.
- (3) R. Freer, *Journal of materials science*, 1980, **15**, 803–824.
- (4) J. Carrasco, N. Lopez and F. Illas, *Physical review letters*, 2004, **93**, 225502.
- (5) M. Martin, in *Diffusion in Condensed Matter: Methods, Materials, Models*, Springer, 2005, pp. 209–247.
- (6) J. Pelleg, *Diffusion in ceramics*, Springer, 2016.
- (7) T. Ishigaki, S. Yamauchi, K. Kishio, J. Mizusaki and K. Fueki, *Journal of Solid State Chemistry*, 1988, **73**, 179–187.
- (8) J. A. Van Orman and K. L. Crispin, *Reviews in Mineralogy and Geochemistry*, 2010, **72**, 757–825.
- (9) P. Kofstad, *Elsevier Applied Science Publishers, Crown House, Linton Road, Barking, Essex IG 11 8 JU, UK, 1988.*, 1988.
- (10) P. Kofstad, *Oxidation of metals*, 1995, **44**, 3–27.
- (11) M. Y. Yang, K. Kamiya, B. Magyari-Köpe, M. Niwa, Y. Nishi and K. Shiraiishi, *Appl. Phys. Lett.*, 2013, **103**, DOI: 10.1063/1.4819772.
- (12) J.-H. Hur, *Sci. Rep.*, 2020, **10**, 5405.
- (13) A. Roy, P. Singh, G. Balasubramanian and D. D. Johnson, *Acta Mater.*, 2022, **226**, 117611.
- (14) A. J. E. Rowberg, H. S. Slomski, N. Kim, N. A. Strange, B. P. Gorman, S. Shulda, D. S. Ginley, K. E. Kweon and B. C. Wood, *Chem. Mater.*, 2024, **36**, 6464–6474.
- (15) G. Henkelman, B. P. Uberuaga and H. Jónsson, *J. Chem. Phys.*, 2000, **113**, 9901–9904.
- (16) D. Sheppard, P. Xiao, W. Chemelewski, D. D. Johnson and G. Henkelman, *J. Chem. Phys.*, 2012, **136**, DOI: 10.1063/1.3684549.
- (17) D. R. Trinkle, *Philos. Mag.*, 2016, **96**, 2714–2735.
- (18) A. Van der Ven, G. Ceder, M. Asta and P. D. Tapesch, *Phys. Rev. B*, 2001, **64**, 184307.
- (19) T. Rehman, M. Jaipal and A. Chatterjee, *J. Comput. Phys.*, 2013, **243**, 244–259.
- (20) C. Li, T. Nilson, L. Cao and T. Mueller, *Phys. Rev. Mater.*, 2021, **5**, 013803.

- (21) J.-X. Shen, H. H. Li, A. Rutt, M. K. Horton and K. A. Persson, *npj Comput. Mater.*, 2023, **9**, 99.
- (22) R. B. Wexler, G. S. Gautam, E. B. Stechel and E. A. Carter, *J. Am. Chem. Soc.*, 2021, **143**, 13212–13227.
- (23) K. Choudhary and B. G. Sumpter, *AIP Adv.*, 2023, **13**, arXiv:2205.08366.
- (24) M. D. Witman, A. Goyal, T. Ogitsu, A. H. McDaniel and S. Lany, *Nat. Comput. Sci.*, 2023, **3**, 675–686.
- (25) M. H. Rahman, P. Gollapalli, P. Manganaris, S. K. Yadav, G. Pilania, B. DeCost, K. Choudhary and A. Mannodi-Kanakithodi, *arXiv*, 2023, arXiv:2309.06423.
- (26) I. Mosquera-Lois, S. R. Kavanagh, A. M. Ganose and A. Walsh, *npj Comput. Mater.*, 2024, **10**, 121.
- (27) S. R. Kavanagh, A. G. Squires, A. Nicolson, I. Mosquera-Lois, A. M. Ganose, B. Zhu, K. Brlec, A. Walsh and D. O. Scanlon, *J. Open Source Softw.*, 2024, **9**, 6433.
- (28) Y.-L. Lee, J. Kleis, J. Rossmeisl, Y. Shao-Horn and D. Morgan, *Energy Environ. Sci.*, 2011, **4**, 3966.
- (29) G. M. Lu, M. Witman, S. Agarwal, V. Stavila and D. R. Trinkle, *Phys. Rev. Mater.*, 2023, **7**, 105402.
- (30) C. López, A. Emperador, E. Saucedo, R. Rurali and C. Cazorla, *Mater. Horizons*, 2023, **10**, 1757–1768.
- (31) J. Meng, M. S. Sheikh, R. Jacobs, J. Liu, W. O. Nachlas, X. Li and D. Morgan, *Nat. Mater.*, 2024, DOI: 10.1038/s41563-024-01919-8.
- (32) T. Xie and J. C. Grossman, *Phys. Rev. Lett.*, 2018, **120**, 145301.
- (33) C. Chen, W. Ye, Y. Zuo, C. Zheng and S. P. Ong, *Chem. Mater.*, 2019, **31**, 3564–3572.
- (34) C. W. Park and C. Wolverton, *Phys. Rev. Mater.*, 2020, **4**, 063801.
- (35) K. Choudhary and B. DeCost, *npj Comput. Mater.*, 2021, **7**, 185.
- (36) O. T. Unke, S. Chmiela, M. Gastegger, K. T. Schütt, H. E. Saucedo and K.-R. Müller, *Nat. Commun.*, 2021, **12**, 7273.
- (37) A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser and I. Polosukhin, *Advances in Neural Information Processing Systems*, ed. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, *Advances in Neural Information Processing Systems*, 2017, vol. 30.
- (38) A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder and K. A. Persson, *APL Mater.*, 2013, **1**, 011002.
- (39) W. C. Chueh, C. Falter, M. Abbott, D. Scipio, P. Furler, S. M. Haile and A. Steinfeld, *Science (80-.)*, 2010, **330**, 1797–1801.
- (40) A. H. McDaniel, E. C. Miller, D. Arifin, A. Ambrosini, E. N. Coker, R. O’Hayre, W. C. Chueh and J. Tong, *Energy & Environmental Science*, 2013, **6**, 2424–2428.
- (41) D. R. Barcellos, M. D. Sanders, J. Tong, A. H. McDaniel and R. P. O’Hayre, *Energy Environ. Sci.*, 2018, **11**, 3256–3265.
- (42) R. B. Wexler, G. Sai Gautam, R. T. Bell, S. Shulda, N. A. Strange, J. A. Trindell, J. D. Sugar, E. Nygren, S. Sainio, A. H. McDaniel, D. Ginley, E. A. Carter and E. B. Stechel, *Energy Environ. Sci.*, 2023, **16**, 2550–2560.
- (43) Y. Knausgård Hommedal, M. Etzelmüller Bathen, V. Mari Reinertsen, K. Magnus Johansen, L. Vines and Y. Kalmann Frodason, *J. Appl. Phys.*, 2024, **135**, DOI: 10.1063/5.0205866.
- (44) 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, *Comput. Mater. Sci.*, 2013, **68**, 314–319.
- (45) J. Gibson, A. Hire and R. G. Hennig, *npj Comput. Mater.*, 2022, **8**, 211.
- (46) M. Witman and P. Schindler, *ChemRxiv*, 2024, 10.26434/chemrxiv-2024-bmw1n.
- (47) G. Palmer, S. Du, A. Politowicz, J. P. Emory, X. Yang, A. Gautam, G. Gupta, Z. Li, R. Jacobs and D. Morgan, *npj Comput. Mater.*, 2022, **8**, 115.
- (48) A. A. Emery, J. E. Saal, S. Kirklin, V. I. Hegde and C. Wolverton, *Chem. Mater.*, 2016, **28**, 5621–5634.
- (49) S. Lany, *The Journal of Chemical Physics*, 2018, **148**, 071101.
- (50) S. L. Millican, J. M. Clary, C. B. Musgrave and S. Lany, *Chem. Mater.*, 2022, **34**, 519–528.
- (51) J. T. Tran, K. J. Warren, S. A. Wilson, C. L. Mulich, C. B. Musgrave and A. W. Weimer, *WIREs Energy and Environment*, 2024, **13**, e528.
- (52) S. Zhai, J. Rojas, N. Ahlborg, K. Lim, M. F. Toney, H. Jin, W. C. Chueh and A. Majumdar, *Energy & Environmental Science*, 2018, **11**, 2172–2178.
- (53) A. Goyal, M. D. Sanders, R. P. O’Hayre and S. Lany, *PRX Energy*, 2024, **3**, 013008.
- (54) S. S. Naghavi, A. A. Emery, H. A. Hansen, F. Zhou, V. Ozolins and C. Wolverton, *Nature communications*, 2017, **8**, 285.
- (55) C. N. R. Rao, *Annual Review of Physical Chemistry*, 1989, **40**, 291–326.
- (56) S. Lany, *J. Am. Chem. Soc.*, 2024, **146**, 14114–14127.

- (57) S. Lany, *Journal of Physics: Condensed Matter*, 2015, **27**, 283203.
- (58) J.-X. Shen, L. F. Voss and J. B. Varley, *J. Appl. Phys.*, 2024, **135**, DOI: 10.1063/5.0203124.
- (59) K. Choudhary, B. DeCost, L. Major, K. Butler, J. Thiyagalingam and F. Tavazza, *Digit. Discov.*, 2023, **2**, 346–355.
- (60) B. Deng, P. Zhong, K. Jun, J. Riebesell, K. Han, C. J. Bartel and G. Ceder, *Nat. Mach. Intell.*, 2023, **5**, 1031–1041.
- (61) I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon, W. J. Baldwin, F. Berger, N. Bernstein, A. Bhowmik, S. M. Blau, V. Cărare, J. P. Darby, S. De, F. Della Pia, V. L. Deringer, R. Elijošius, Z. El-Machachi, F. Falcioni, E. Fako, A. C. Ferrari, A. Genreith-Schriever, J. George, R. E. A. Goodall, C. P. Grey, P. Grigorev, S. Han, W. Handley, H. H. Heenen, K. Hermansson, C. Holm, J. Jaafar, S. Hofmann, K. S. Jakob, H. Jung, V. Kapil, A. D. Kaplan, N. Karimitari, J. R. Kermode, N. Kroupa, J. Kullgren, M. C. Kuner, D. Kuryla, G. Liepuoniute, J. T. Margraf, I.-B. Magdău, A. Michaelides, J. H. Moore, A. A. Naik, S. P. Niblett, S. W. Norwood, N. O’Neill, C. Ortner, K. A. Persson, K. Reuter, A. S. Rosen, L. L. Schaaf, C. Schran, B. X. Shi, E. Sivonxay, T. K. Stenczel, V. Svahn, C. Sutton, T. D. Swinburne, J. Tilly, C. van der Oord, E. Varga-Umbrich, T. Vegge, M. Vondrák, Y. Wang, W. C. Witt, F. Zills and G. Csányi, *arXiv*, 2023, arXiv:2401.00096.