

Latent Diffusion Models for Virtual Battery Material Screening and Characterization

Deepalaxmi Rajagopal, Arnd Koeppe,* Adrian Cierpka, and Britta Nestler

Efficient characterization of battery materials is fundamental to understanding the underlying electrochemical mechanisms and ensuring the safe operation of batteries. In this work, an innovative data-driven multimodal generative method is proposed to accelerate the characterization and screening of battery materials. This approach leverages a variant of the latent diffusion model, which combines a variational autoencoder (VAE) and a denoising U-Net. The VAE maps microscale information from characterization techniques, such as atomic force microscopy

(AFM), into a common latent space, and the denoising U-net, conditioned on battery properties, guides the screening of battery materials. Together, the data-driven properties of material space, enriched with battery functional properties and formulated in a common latent space, achieve the accurate translation of information from AFM to meaningful material descriptors and accelerate the screening of battery materials to meet the functional needs of the battery system under consideration.

1. Introduction

Understanding how the atomic structure, arrangement of crystallographic domains, and composition at nano-, micro-, and mesoscales affect the capacity and performance of electrode materials is crucial for optimizing the efficiency of battery systems. Several characterization techniques are employed, each providing distinct characteristics to create a comprehensive picture of a single material configuration.^[1] Atomic force microscopy (AFM)^[2,3] can provide the complete morphological profile of electrodes during the charging and discharging stages at high resolutions. AFM enables the probing of electrochemical reactions, current distribution, surface potentials, ion transfers, and other properties of battery surfaces.^[4,5] This method allows for studying nanoscale processes developed at the interface of the electrodes and electrolytes. Thus, AFM is indispensable in nanoscience, especially when used in situ for resolution where the microscope tip is functionalized to reveal surfaces' atomic and electronic molecular structure.^[6–9] The current AFM machines utilize an automated tip preparation method with machine learning (ML) techniques to provide the operator with information about the type of tip functionalization required and tip quality,

thereby reducing the time spent on tip preparation. As a result, these machines can process thousands of experimental samples daily.^[10] However, high-resolution AFM experiments produce unlabeled data, making interpretation tedious and challenging for human experts.^[11]

Understanding how the atomic structure, arrangement of crystallographic domains, and composition at nano-, micro-, and mesoscales affect the capacity and performance of electrode materials is crucial for optimizing the efficiency of battery systems. Several characterization techniques are employed, each providing distinct characteristics to create a comprehensive picture of a single material configuration.^[1] Atomic force microscopy (AFM)^[2,3] can provide the complete morphological profile of electrodes during the charging and discharging stages at high resolutions. AFM enables the probing of electrochemical reactions, current distribution, surface potentials, ion transfers, and other properties of battery surfaces.^[4,5] This method allows for studying nanoscale processes developed at the interface of the electrodes and electrolytes. Thus, AFM is indispensable in nanoscience, especially when used in situ for resolution where the microscope tip is functionalized to reveal surfaces' atomic and electronic molecular structure.^[6–9] The current AFM machines utilize an automated tip preparation method with machine learning (ML) techniques to provide the operator with information about the type of tip functionalization required and tip quality, thereby reducing the time spent on tip preparation. As a result, these machines can process thousands of experimental samples daily.^[10] However, high-resolution AFM experiments produce unlabeled data, making interpretation tedious and challenging for human experts.^[11]

Comprehending the encoded information about sharp image contrasts in high-resolution AFM techniques necessitates an interpretation of the various factors contributing to these contrasts.^[12] These sharp contrasts in AFM images are crucial for understanding the distinct features and nature of image contrasts

D. Rajagopal, A. Koeppe, A. Cierpka, B. Nestler
Institute of Applied Materials (IAM-MMS)
Karlsruhe Institute of Technology (KIT)
Straße am Forum 7, 76131 Karlsruhe, Germany
E-mail: arnd.koeppe@kit.edu

D. Rajagopal, A. Koeppe, A. Cierpka, B. Nestler
Institute of Nanotechnology (INT)
Karlsruhe Institute of Technology (KIT)
Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen,
Germany

© 2025 The Author(s). Batteries & Supercaps published by Wiley-VCH GmbH. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

to comprehend the intermolecular regions in AFM images fully. Interpreting experimental AFM images, particularly in the case of molecular systems, is challenging due to the complexity of the structure, chemical composition, internal corrugation, and the corresponding interplay with operational parameters of AFM imaging techniques.^[13] The current process of interpreting AFM images is laborious and challenging, highlighting the need for a more efficient and comprehensive solution. One approach to comprehensively understanding these image features and contrasts is the use of a simple mechanical model known as probe particle AFM (PPAFM), which encompasses all the essential elements of AFM imaging. The PPAFM model,^[14–16] developed by Hapala et al., is an open-source software that simulates high-resolution AFM and various scanning probe microscopy techniques with submolecular resolution. It is designed to generate AFM images captured using a diverse range of modified tips across varying tip-sample distances. The software models the deflection of the probe particle connected to the tip. This probe particle features a flexible tip apex, commonly functionalized with a carbon monoxide (CO) molecule, but other possibilities, such as Xe, Cl[−], H₂O, and others, are also used to examine interactions at the molecular level.

High-resolution AFM differentiates between atoms within a molecule by leveraging the variations in the interaction between the AFM tip and the sample and the specific chemical properties of the materials under examination.^[17,18] While characteristics in 2D features, like sharper vertices, are readily identifiable through visual observation,^[19] discerning 3D information to distinguish between atoms of similar shape but with distinct tip-sample interactions at molecular sites poses significant challenges. Data-driven approaches based on deep learning can be implemented to address these challenges and understand the intricate salient features underlying large image-based material datasets. Deep generative models automate the extraction of insights from characterization data, enabling phase segmentation, synthetic data generation, and material optimization.^[20–24] The ML methods have facilitated predicting molecular structures and chemical compositions by analyzing 3D stacks of constant-height AFM images. This progress opens up new possibilities for battery material research, offering a powerful tool to comprehend and foresee the behavior of battery materials at a molecular level. The previous data-driven approaches in AFM analysis have focused on deterministic convolutional neural networks (CNNs) to map force maps to individual material descriptors. These approaches have demonstrated effectiveness in reconstructing particular features and facilitating nanoscale characterization. However, they are limited in their capacity to perform generative modeling tasks, such as sampling, interpolation, and synthesizing novel material configurations.^[25–27]

On the other hand, the functional properties of materials play a significant role in the screening and design of battery materials tailored for specific applications. The average voltage of a material plays a crucial role in defining its operating voltage, which can enhance energy production and inspire new system designs. Gravimetric and volumetric capacities evaluate the total charge a material can store and the energy held per unit weight and unit

volume. Furthermore, volume change refers to the expansion and contraction of electrode materials during charge and discharge cycles, which is crucial for achieving long-term stability and high performance. Finding the right balance in these properties enhances battery performance, increases energy storage, and reduces size for various applications.^[28] To develop battery materials that improve performance and safety, researchers often conduct numerous trial-and-error experiments based on characterization and screening. ML-assisted methods can accelerate the effective translation of materials informatics from characterization techniques and the inverse design of battery materials based on the functional properties required for any specific application.

Our research will leverage this potential by proposing a “BattGen” framework, a generative multimodal methodology based on latent diffusion model (LDM) to virtually screen and characterize battery materials based on their functional properties, such as average voltage, volume change, gravimetric and volumetric capacity, and gravimetric and volumetric energy. The proposed framework is demonstrated on data for battery materials sourced from the Materials Project database,^[29] which contains battery electrode materials and their corresponding battery functional properties based on different working ions.^[30–32] The chemical formula of the electrode at the end of the discharge cycle of around 3,154 electrodes is utilized to create input for virtual AFM simulations using the PPAFM model, thereby constructing a material database. In this proof-of-concept work, we utilize virtual 3D stack AFM images as input, as image-like data is well-suited for a convolutional neural network-based LDM. Through this approach, we translate the 3D stack AFM data into 2D image-based descriptors containing relevant information about the battery electrode material under study. Additionally, the learned reduced space, which includes latent features of the 3D AFM stack conditioned with battery functional properties, will facilitate the screening of target battery materials. The framework developed integrates atomic-scale structural data with compositional features to facilitate the early virtual screening of battery materials. Although it does not cover all experimental and device-level factors, it emphasizes key battery material properties to streamline the selection of promising candidates before conducting laboratory tests.

2. Virtual Battery Material Screening and Characterization

2.1. Overview on LDM

Diffusion models are developed based on the principles of nonequilibrium thermodynamics.^[33] These diffusion probabilistic models demonstrated cutting-edge performance in density estimation and sampling quality.^[34,35] The diffusion models have gained attention for addressing challenges faced by popular deep generative models, such as vanilla generative adversarial neural network (GAN), which exhibit poor training stability and struggle to capture the entire data distribution.^[36–38] The diffusion

models can be seen as a composition of two processes, the forward diffusion and the reverse diffusion. For better visualization, **Figure 1** illustrates the forward and reverse diffusion processes in image space using the type map descriptor of an example battery material obtained from the Materials Project database. The forward diffusion process involves a series of diffusion steps, utilizing a Markov chain of T steps to introduce noise into the data gradually. The Markov chain assumption states that the system evolves in such a way that the probability distribution of the current state \mathbf{x}_t depends only on the immediately preceding state \mathbf{x}_{t-1} , and not on the full history. Formally, this is written as $q(\mathbf{x}_t|\mathbf{x}_{t-1}, \dots, \mathbf{x}_0) = q(\mathbf{x}_t|\mathbf{x}_{t-1})$, which defines a discrete-time Markov chain. **Figure 1** shows that for a given sample \mathbf{x}_0 sampled from the real data distribution $q(\mathbf{x})$, the forward process adds Gaussian noise for T steps and obtains $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T$. For each step of the forward process, the Gaussian noise with variance β_t is added to \mathbf{x}_{t-1} to compute \mathbf{x}_t with distribution $q(\mathbf{x}_t|\mathbf{x}_{t-1})$. This process can be expressed as

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \quad (1)$$

with $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$,

where $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ represents the conditional probability distribution of the data at time t given the data at the previous timestep. The notation $\mathbf{x}_{1:T} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$ represents the sequence of \mathbf{x} values from time step 1 to T . The forward noising process is repeated until the distribution of \mathbf{x}_T becomes nearly indistinguishable from pure Gaussian noise, denoted as $\mathcal{N}(0, \mathbf{I})$. Each step of noise addition during this forward diffusion process is Gaussian, which ensures that the overall distribution remains Gaussian even after multiple iterations. According to the central limit theorem, when independent random variables are combined, the resulting distribution approaches a Gaussian distribution as the number of noise addition steps increases. Following this, the LDM learns a reverse diffusion process that restores the data's original structure, resulting in a flexible and manageable generative model (**Figure 1**). In contrast to the forward diffusion process, the reverse diffusion process aims to train the denoising network using a U-Net architecture. U-Nets are CNNs that feature an encoder-decoder structure resembling the letter U.^[39,40] The reverse diffusion process aims to recursively remove noise from \mathbf{x}_t . It moves backward along the time chain as t decreases from T to 0. By conditioning the U-Net on timestep t , the neural network learns to predict the Gaussian parameters, such as the mean

$\mu_\theta(\mathbf{x}_t, t)$ and the covariance matrix $\Sigma_\theta(\mathbf{x}_t, t)$, for each time step.^[35,41] From a mathematical point of view, the reverse diffusion process can be expressed as

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t), \quad (2)$$

with $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t))$,

where $p_\theta(\mathbf{x}_{0:T})$ denotes the reverse diffusion trajectory. The reverse diffusion kernel $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ is defined by the mean $\mu_\theta(\mathbf{x}_t, t)$ and the covariance matrix $\Sigma_\theta(\mathbf{x}_t, t)$. Both are parametrized with θ of the reverse distribution. The reverse diffusion process learns to optimize the denoising parameters θ through a multi-step chain in the reverse process. The denoising parameters θ are optimized through training a neural network-based diffusion model. This model aims to discover the reverse transitions that maximize the likelihood of achieving the highest probability $p_\theta(\mathbf{x}_t)$ for each sample \mathbf{x}_t . The intractability of $p_\theta(\mathbf{x}_t)$ for all possible reverse trajectories is solved by minimizing the variational upper bound. After learning the reverse diffusion process, the model parameters θ have been fine-tuned to generate new data by sampling from random Gaussian noise.^[33,42] The LDM is a modified version of diffusion models involving diffusion in the latent space, representing the lower-dimensional form of input data. LDM utilizes a trained encoder to encode full-size images into a lower-dimensional representation known as the latent space. It then carries out the forward and reverse diffusion processes within the latent space, and afterward, a trained decoder upsamples the latent representation back into full-size images.^[43] LDM reduces the computational burden of reverse diffusion by leveraging the latent space, resulting in faster denoising of images, improved speed, and enhanced overall model stability and robustness. In this work, the reduced-dimensional material space from AFM force maps allows the translation of information from characterization techniques to understandable material image descriptions, and encoding battery properties in the learned latent space expedites the guided screening of battery materials.

2.2. BattGen Framework: Screen and Characterize

Figure 2 illustrates the entire BattGen framework, based on the LDM methodology, for screening and characterizing battery materials according to target battery functional properties. The implemented LDM architecture comprises a variational




Figure 1. Visualization of forward and reverse diffusion process on the image-based material descriptor type map encoded with different colored spheres based on van der Waals radius and type of atom.

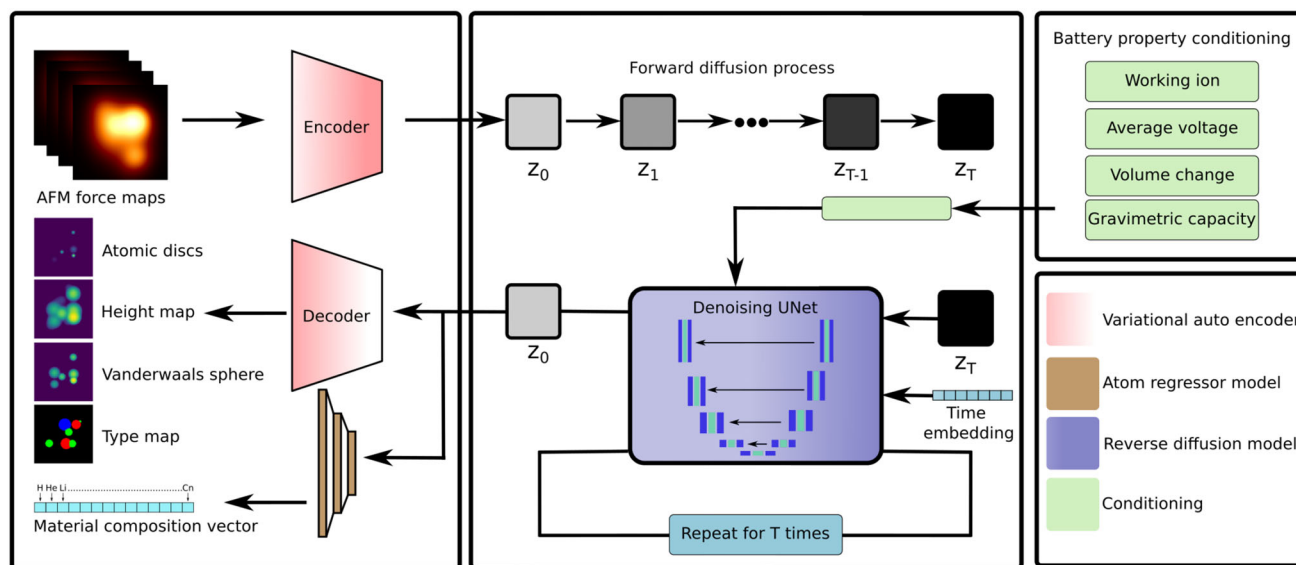


Figure 2. Overview of the BattGen framework. The first frame shows the encoder and decoder of the translational model, which maps AFM force maps to corresponding battery material descriptors. The middle frame highlights the diffusion model in the learned latent space, guided by battery functional properties for the virtual screening of target battery materials.

autoencoder (VAE) with a patch discriminator, which compresses the input image in pixel space to a lower-dimensional latent space that is perceptually equivalent to the actual data space. The VAE architecture includes convolutional layers that compress images and extract meaningful latent features during the encoding process. During the decoding process, the encoded latent features are upsampled. Using residual blocks enables the model to learn deep hierarchical features while maintaining information between layers, which helps preserve meaningful latent representations. Attention layers allow the model to focus on the global features present in the input. The patch discriminator acts as a critic, distinguishing between the generated feature maps and their corresponding ground truth images. This study utilizes the entire VAE with a patch discriminator as a translational model to translate AFM force maps into their respective material image descriptors. For diffusion in the latent space learned by the VAE, the U-Net architecture with skip connections is implemented with convolution, attention, and residual blocks. The implemented U-Net is a symmetric architecture with input and output of the same spatial size, utilizing skip connections between encoding and decoding blocks of corresponding feature dimensions. These skip connections in the U-Net architecture help preserve local details from the encoder to the decoder, which in turn enables the U-Net to produce high-quality image generation. In this work, as the first stage of training the LDM, a VAE is trained adversarially with a patch discriminator to efficiently translate 3D AFM images into 2D descriptors of materials scanned in AFM, ensuring that the original and reconstructed samples are accurately distinguished. Based on the original implementation of the LDM, the training of the VAE is enhanced by utilizing both perceptual losses^[44] and an image-conditioned patch-based adversarial objective from Pix2Pix GAN to facilitate better image-to-image translation.^[45] Finally, the denoising U-net conditioned with

the corresponding battery functional properties of the material is trained using the learned latent representation from the VAE. For the conditional generation of battery materials based on target battery functional properties, the symmetric concatenation of these properties at the start of each down, middle, and up block in the U-net helps move the learned latent distribution toward the conditioned properties. The U-net training is optimized using the mean squared error loss function as an objective. In the forward diffusion process, the number of time steps T is set to 1000, as used in the original implementation Ho et al.^[41] which is found empirically sufficient to disrupt the original data structure and allow for gradual noise injection necessary for stable training of the reverse process. The sequence β_1, \dots, β_T determines the variance of the Gaussian noise added at each step of the forward process, increasing linearly from 10^{-4} to 0.02, which are small relative to data normalized to the range $[-1, 1]$. The quality of images generated by the translation model is evaluated using the Fréchet inception distance (FID), a quantitative metric for measuring image quality. The FID is a metric that evaluates the quality of generated images in comparison to ground-truth images. It utilizes a pretrained CNN, specifically the Inception Net, to model the generated and ground truth images as mean and covariance matrices of multivariate Gaussian distributions. A lower FID indicates a smaller distance between the generated and ground truth data distributions, suggesting better image quality.^[46,47]

2.3. Data Generation Pipeline

The primary data for this study is sourced from the Materials Project.^[29] For this analysis, the source data includes geometrical information on the battery material composition under study and its corresponding functional properties, such as average




Figure 3. Training data generation pipeline. The source data for this study is obtained from the battery explorer feature of the Materials Project database. Atomic coordinates of the battery formulas and their functional properties are extracted using the Materials Project Application Programming Interface. The extracted data is organized and managed through the Kadi4Mat programming interface. KadiAI and CIDS enable the ML process in this study.

voltage, gravimetric capacity, and volume change. The molecular Cartesian coordinates of atoms in the battery discharge formula are used as input to generate a 3D AFM simulation stack at different tip-sample distances. The corresponding 2D image descriptors, such as height maps, van der Waals spheres, atomic discs, and multimap sphere elements (type map), are computed, containing chemical and physical information about the battery discharge formula. The computed data are preprocessed to a normalized range for efficient training. **Figure 3** visualizes the implemented data generation pipeline. We utilized the Kadi4Mat data platform to manage our data and ML workflows. Kadi4Mat is a shared repository and electronic lab notebook developed by our group^[48,49] and is deployed in several instances across Europe. This platform collects and organizes data from simulations and experiments along with relevant information about that data. For our study, we used the computational intelligence and data science (CIDS) tools and KadiAI,^[50] which are part of the Kadi4Mat ecosystem, to set up and run our ML processes (Figure 3).

3. Results and Discussion

3.1. Data

The battery materials extracted from the Materials Project database are defined based on the following battery systems: lithium (Li), magnesium (Mg), zinc (Zn), calcium (Ca), potassium (K), sodium (Na), aluminum (Al), yttrium (Y), rubidium (Rb), and cesium (Cs). The extracted materials are filtered to remove duplicates in battery compositions, yielding 3,167 unique battery material configurations. In order to increase the diversity of data for training, the atomic coordinates of the obtained material configuration are rotated using the Fibonacci spiral method in sphere tangent space, creating a quasirandom distribution without clustering. From this process, we selected the 30 rotations with the highest entropy for the AFM simulations. This entropy-based selection process quantifies the depth variation of projected atomic positions to identify structurally rich orientations for AFM simulation. To ensure that the principal features are fully captured within the

scan area, AFM image validity is evaluated by identifying the pixels of maximum intensity and calculating its minimum distance to the image boundaries using Minimum distance to edge = $\min(x, \text{width} - 1 - x, y, \text{height} - 1 - y)$. Here (x, y) denotes the coordinates of the pixel exhibiting maximum intensity, while the image dimensions are defined as a height and width of 128 pixels each. AFM images with a minimum distance to the edge equal to zero are considered invalid, indicating that the main feature lies outside the defined scan dimensions used for the simulation of the entire dataset. During the AFM simulations, 13 material configurations, along with 390 corresponding rotations that consistently produced invalid AFM images, were removed to prevent the inclusion of data lacking learnable information for training purposes. As a result, the final dataset comprises 3,154 samples, each multiplied by 30 applied rotations, resulting in a total of 94,640 samples that simulate the AFM measurements of the initial materials. For the AFM simulation, the maximum scan dimension is set to 128×128 , and each augmented molecule is scanned for ten different sample distances. The auxiliary map function of the PPAFM model allows the generation of corresponding material descriptors, such as height maps, van der Waals sphere, atomic discs, and multimap sphere elements (type map).^[25–27] For stable training, input data is normalized using min-max scaling to map values to a target range such as $[0, 1]$ or $[-1, 1]$. The 3D force maps and the material descriptors are preprocessed to a range of $[-1, 1]$ to facilitate efficient training of the LDM. The conditioning battery functional properties are preprocessed to a range of $[0, 1]$ to enable guided screening of battery materials. Each material configuration's atom count is encoded into a material composition vector. This vector maintains a constant length of 112, with each index corresponding to the sequence of elements in the periodic table. The vector is populated with the atom count at the corresponding index, while the remaining indices are filled with zeroes. This vector is then utilized in a regression model to learn the chemical space from the compressed latent space. The stratified split strategy is used to divide the dataset into training, validation, and test sets with a ratio of 0.7 for training, 0.15 for validation, and 0.15 for testing, based on the working ion type. The implemented split strategy prevents the overlap of identical materials across the training,

validation, and test datasets, ensuring balanced representation of each working ion-based material. To improve model performance and generalization, the training set includes sample weighting based on normalized inverse frequency for the model to focus on underrepresented working ion-based materials during training.

3.2. Translation of 3D AFM Image Stacks to Material Descriptor

The trained VAE component of the implemented LDM is used to translate the 3D stack of simulated AFM images with force scans of the battery material obtained at different tip-sample distances to corresponding material descriptors such as height map, van der Waals sphere, atomic discs, and type map. The above VAE model is trained adversarially with a patch discriminator to discriminate between the original and generated images. The implemented adversarial patch discriminator in this work consists of two convolutional branches to handle grayscale descriptors like height map, van der Waals sphere, and atomic discs, and RGB-based descriptors like type map, which encodes red, green, and blue colors to represent different atoms based on van der Waals radius. In the type map descriptor, the red, green, and blue channels are used to classify the different types of atomic species relevant to the study. The blue channel signifies the working ion species, the green channel represents gas phase elements, and the red channel corresponds to all other elements up to atomic number 112 in the periodic table. The image representation of these descriptors can be found in the Figure 2. This change in the general architecture of the patch discriminator to compute the patches of the input image based on grayscale and RGB images helped to improve the reconstructions of VAE without any wrong prediction of colors, especially in the case of a type map and also distinguished the local feature of force maps with the help of receptive fields defined at each convolution step of the discriminator. **Figure 4** shows the accurate translation of the image descriptors of a sample from test datasets with the different working ions of the battery. The translation model exhibits a low FID score of 2.301 when converting 3D AFM scans into grayscale descriptors, indicating the generation of outputs with high fidelity. For type map translations, the model achieves an FID score of 26.845, which suggests a moderate level of fidelity and acceptable image quality. In both cases, the relatively low FID scores imply that the model generates outputs that are statistically close to the underlying distributions of the real data.

3.3. Screening of Battery Materials

The virtual screening of battery materials is facilitated by conditioning the average voltage, gravimetric and volumetric capacity, gravimetric and volumetric energy density, volume change, and working ion to the learned latent space from a translational model as input to the U-Net part of LDM, along with the time embedding. The latter directs the prediction of the noise level added at each time step during the forward diffusion process. The above-mentioned battery characteristic properties as

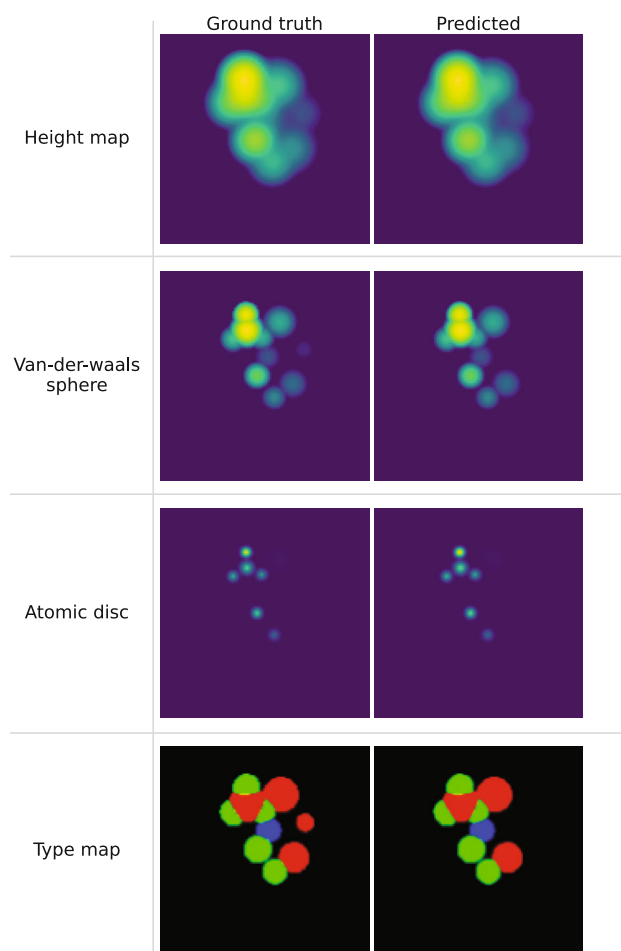


Figure 4. Comparison plot between image-based material descriptor ground truth and corresponding predictions using the implemented translation model.

one of the inputs to the U-Net help to direct the screening of target battery electrode materials at discharged state based on the possible working ions such as lithium (Li), magnesium (Mg), zinc (Zn), calcium (Ca), potassium (K), sodium (Na), aluminum (Al), yttrium (Y), rubidium (Rb), and cesium (Cs). Demonstrating the implemented virtual screening tool involves providing pure random Gaussian noise, as well as working ion and battery properties as inputs to the trained U-Net model. The model then screens various material configurations based on the specified battery properties. In this case, the initial condition of the Gaussian noise is crucial for determining the types of material descriptors that are generated. **Figure 5** shows the predictions of the LDM under certain seeding conditions and rotations on test samples that are not used in training with different working ion and battery functional properties. To evaluate the influence of each conditioning variable on the generated outputs, a sensitivity analysis is conducted utilizing the L2 distance metric. This methodology quantifies output variations following perturbations of each feature by ± 0.05 while maintaining the remaining variables constant. Since all continuous features are min–max normalized to the $[0, 1]$ range, the value ± 0.05 corresponds to

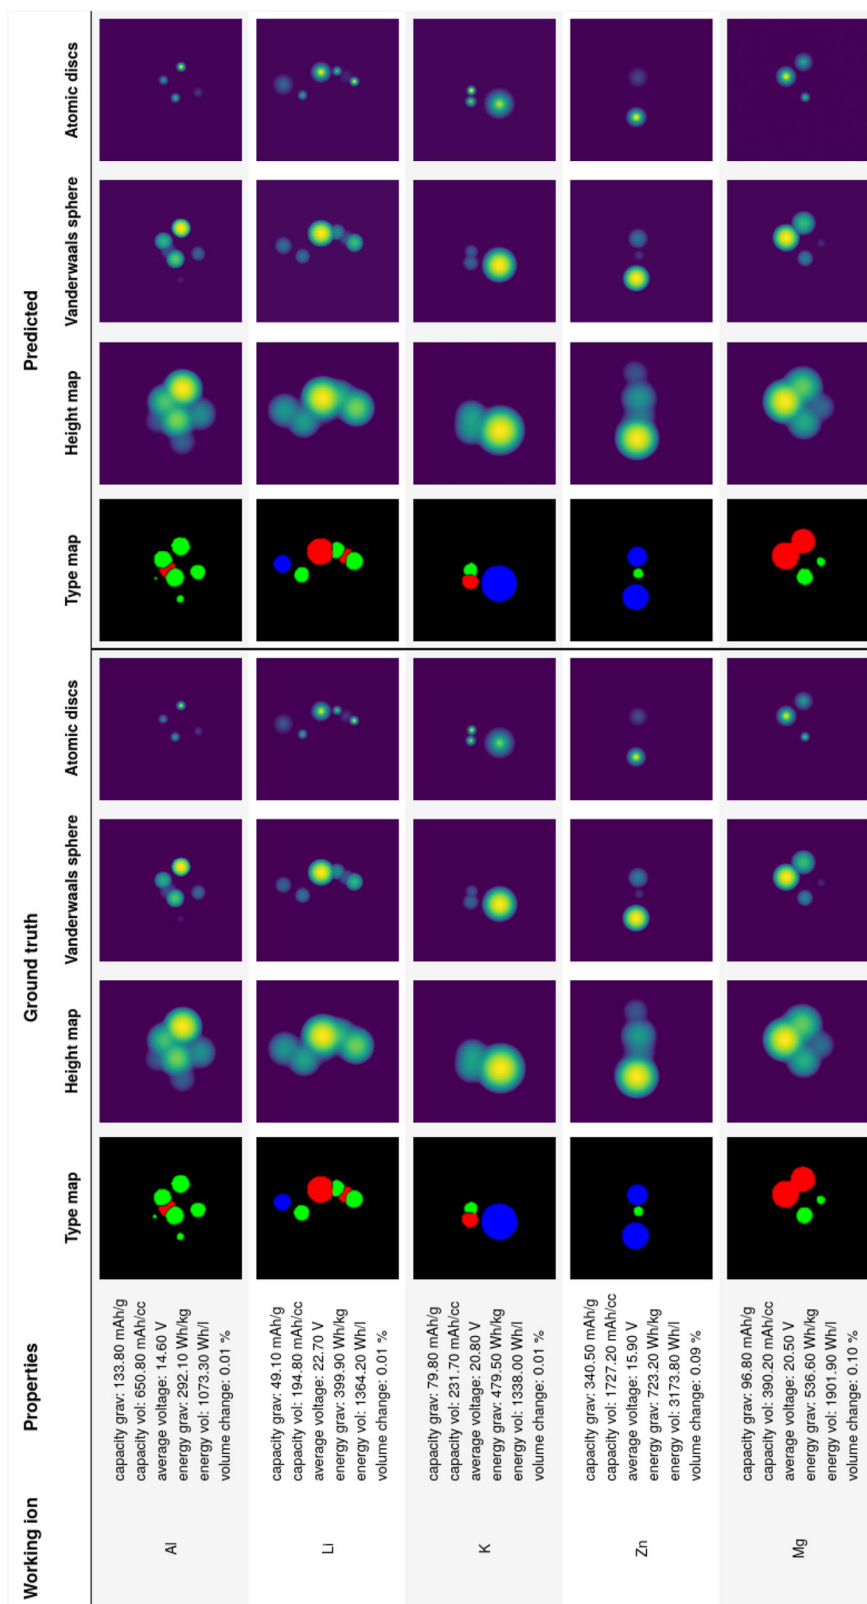


Figure 5. Demonstration of the implemented virtual screening tool to predict the image descriptors of the materials from the test dataset. Here, the latent space, configured with battery contextual properties such as working ion, gravimetric capacity, volumetric capacity, average voltage, gravimetric energy, volumetric energy, and volume change, is used to input the U-net part of LDM to predict the material image descriptors. The predicted image descriptors match the corresponding ground truths exactly.

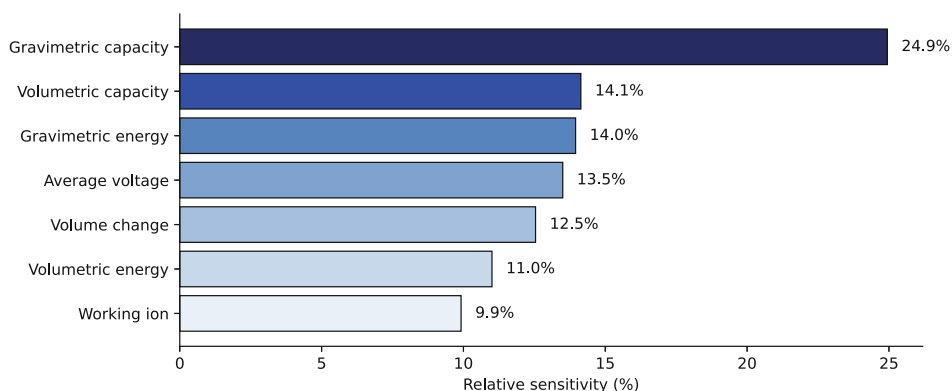


Figure 6. Relative sensitivity of conditioning variables in the conditional LDM. The influence of each input condition on the generated outputs is quantified using a normalized L2 distance metric. The sensitivity analysis indicates that gravimetric capacity and volumetric capacity have a significant impact on the model's performance in screening battery materials.

a 5% relative change and allows consistent comparison across features. Using both positive and negative perturbations enables the capture of any asymmetry in the model's response. For categorical features such as the working ion, a discrete substitution approach is applied by replacing the original ion with all valid alternatives and averaging the resulting L2 distances. The relative sensitivities are derived by normalizing the resulting distances. As shown in **Figure 6**, gravimetric and volumetric capacities show the highest influence, followed by gravimetric energy density and average voltage. The volumetric energy density, volume change, and working ion contribute less, indicating reduced influence on generation. These results highlight the dominant role of gravimetric and volumetric properties in the screening of battery materials.

3.4. Atom Regression Model to Predict the Chemical Information in AFM Images

The implemented atom regression model is designed to decode chemical information from scanned AFM images. We map the

learned latent space of the translational model to represent the chemical space specific to each AFM image. This model is trained using the learned latent space from the translational model. A balanced combination of convolutional and dense layers enables the model to effectively filter and extract relevant features from the latent space representation, allowing for the prediction of the material composition vector. This vector corresponds to atomic numbers up to 112, effectively capturing information about the elements within that atomic number. The last layer of the atom regression model uses softplus activation to avoid negative values in the predicted material composition vector. Additionally, the trained atom regression model can be used to predict the material composition of battery materials that are virtually screened using the trained U-net part of LDM. The model is trained for 400 training epochs and optimized using the mean squared error loss function. **Figure 7** visualizes the performance of the atom regression model to predict certain atoms based on working ions. The predicted atom count in the material composition vector matches exactly with the atom count in most




Figure 7. Graphical visualization of the performance of the atom regression model in the prediction of atoms based on the working ion of the battery in the material composition vector. The distribution of true and predicted atom counts aligns well, with notable deviations observed for aluminum (Al), rubidium (Rb), yttrium (Y), and cesium (Cs). The width of each violin represents the kernel density estimation of the distribution. The thin dashed lines in each violin represent the 25th and 75th percentiles, while the thick dashed line represents the median.

Atom	Metrics		
	MAE	MSE	R2 score
Li	0.021	0.0236	0.9989
Na	0.0181	0.1593	0.9707
Mg	0.0138	0.03	0.9996
Al	0.0475	0.2687	0.8097
K	0.0211	0.1054	0.849
Ca	0.0089	0.011	0.9931
Zn	0.0072	0.0073	0.9989
Rb	0.0169	0.0738	0.7284
Y	0.0158	0.0411	0.9015
Cs	0.0115	0.0705	0.5087

cases, except in aluminum (Al), rubidium (Rb), yttrium (Y), and cesium (Cs). **Table 1** shows the evaluation metrics of the atom regressor model for the atoms based on working ions, and **Figure 8** shows the prediction error distribution plot of the atom regression model over different working ions. The model demonstrates high predictive accuracy, achieving a coefficient of determination (R2) score of 0.9996 for magnesium (Mg), 0.9989 for lithium (Li) and zinc (Zn), and 0.9931 for calcium (Ca). However, the model performance for cesium (Cs) and rubidium (Rb), with R2 scores of 0.5087 and 0.7284, respectively, indicates challenges in capturing their underlying patterns due to their lower frequency in the training samples. For aluminum (Al) and potassium (K), the model exhibits intermediate performance, with R2 scores of 0.8097 and 0.849, respectively. The model generalizes well for most elements, but certain atoms will need additional samples for improved accuracy.

4. Concluding Remarks and Outlook

In this work, we proposed a virtual screening tool to filter target battery materials with enhanced safety, cyclability, and long-life performance capabilities. The intermediate VAE with a patch discriminator translates the given AFM scans of battery materials into related image descriptors to characterize the chemical space. The results show that training a VAE with discriminators and perceptual loss yields better generation quality. The learned latent space from the translation model is utilized to train the LDM for virtual screening of battery materials. The reverse diffusion process of the LDM configured the latent space effectively with battery functional properties. Through the described latent space configuration, we demonstrate virtually screened battery materials. The trained LDM generates material descriptors from pure Gaussian noise, conditioned on battery properties that correspond precisely to the actual material descriptors for the given battery properties under specific seeding conditions. Furthermore, the atom regressor model trained with the learned latent space from the LDM directly predicts the chemical space of virtually screened battery materials. The BattGen framework accelerates the translation of material informatics from AFM-based characterization techniques, promoting the inverse design of battery materials based on functional properties tailored to meet specific application requirements. While AFM does not provide direct elemental composition, its high-resolution structural and mechanical contrast offers valuable nanoscale insights into battery electrode surfaces. Although focused on a simplified cathode composition, the framework is adaptable to experimental AFM data and more complex material systems through transfer learning and domain adaptation. The screening component of the BattGen framework enables the identification of battery

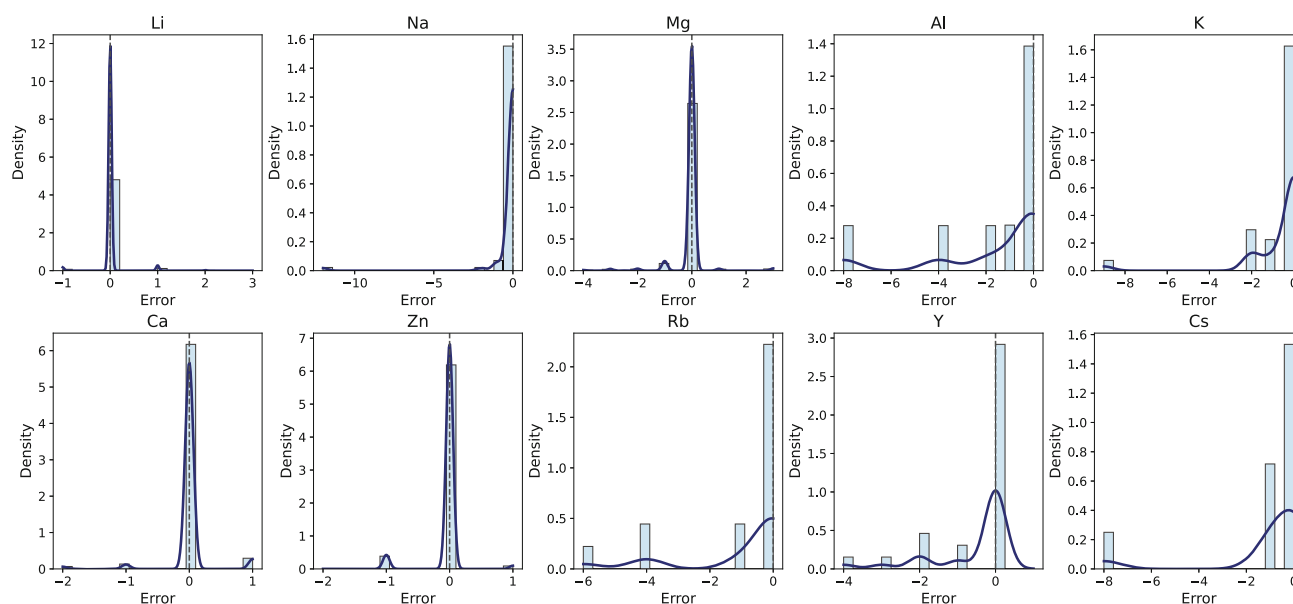


Figure 8. Prediction error distributions for individual atomic species in the atom regression model. Each subplot presents a histogram and kernel density estimate of the errors, with a dashed line indicating zero error. While most atoms are predicted accurately, larger deviations are observed for aluminum (Al), rubidium (Rb), yttrium (Y), and cesium (Cs).

material based on functional metrics. The implemented screening approach prioritizes materials that are promising for a specific use case, thereby narrowing the candidate list for further computational or experimental validation. The implemented application-driven screening approach enables application-targeted exploration of the materials space, especially where extensive density functional theory (DFT) datasets are not available. Future work will focus on incorporating DFT-based predictions of battery properties to govern the second-stage screening of optimized battery materials for a specific use case.^[50–54] The translation component of the current framework exhibits strong performance in analyzing a simulated AFM scan focused on a simple composition of a battery material. However, the direct application to experimental AFM images necessitates further development of the implemented framework. Since the experimental AFM data frequently present complex surface features, regions with multiple compositions, noise from instrumental artifacts, environmental factors, and varying imaging conditions that are not sufficiently captured by the simulated training data. Future work will focus on bridging the gap between the real and simulation domains through the following steps: a preprocessing and augmentation pipeline, along with fine-tuning through transfer learning using a limited set of labeled experimental images. This transfer learning approach will enhance the model's ability to generalize to experimental AFM data.

Acknowledgements

D.R. and A.K. contributed equally to this work. The authors appreciate the financial support for this research from multiple organizations. The authors would like to specifically acknowledge the Deutsche Forschungsgemeinschaft for the “Cluster of Excellence” POLiS (project number 390874152). Additionally, the authors thank the German Federal Ministry of Education and Research (BMBF) for the AQuaBP project (grant number 03XP0315B) and the Helmholtz Association for the MTET program (no: 38.02.01).

Open Access funding enabled and organized by Projekt DEAL.

Code Availability

The code for data preprocessing, model definition, and training scripts is made publicly available through Kadi4Mat along with the data on Zenodo as a FAIR data object in the RO-Crate format: <https://doi.org/10.5281/zenodo.14775779>.

The code for our machine-learning module (CIDS, v3.2a) is available open-source on GitLab: <https://gitlab.com/intelligent-analysis/cids>.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

The data that support the findings of this study are openly available in [Zenodo] at [<https://doi.org/10.5281/zenodo.14775779>], reference number [14775779].

Keywords: artificial intelligence · characterizations · generative models · neural networks · virtual screenings

- [1] D. P. Finegan, I. Squires, A. Dahari, S. Kench, K. L. Jungjohann, S. J. Cooper, *ACS Energy Lett.* **2022**, *7*, 4368.
- [2] G. Binnig, C. F. Quate, C. Gerber, *Phys. Rev. Lett.* **1986**, *56*, 930.
- [3] J. Loos, *Adv. Mater.* **2005**, *17*, 1821.
- [4] E. M. Tennyson, C. Gong, M. S. Leite, *ACS Energy Lett.* **2017**, *2*, 2761.
- [5] M. Krieg, G. Fläschner, D. Alsteens, B. M. Gaub, W. H. Roos, G. J. Wuite, H. E. Gaub, C. Gerber, Y. F. Dufrène, D. J. Müller, *Nat. Rev. Phys.* **2019**, *1*, 41.
- [6] J. Zhang, R. Wang, X. Yang, W. Lu, X. Wu, X. Wang, H. Li, L. Chen, *Nano Lett.* **2012**, *12*, 2153.
- [7] Z. Zhang, S. Said, K. Smith, R. Jervis, C. A. Howard, P. R. Shearing, D. J. Brett, T. S. Miller, *Adv. Energy Mater.* **2021**, *11*, 2101518.
- [8] T. M. Arruda, M. Heon, V. Presser, P. C. Hillesheim, S. Dai, Y. Gogotsi, S. V. Kalinin, N. Balke, *Energy Environ. Sci.* **2013**, *6*, 225.
- [9] Z. Deng, X. Lin, Z. Huang, J. Meng, Y. Zhong, G. Ma, Y. Zhou, Y. Shen, H. Ding, Y. Huang, *Adv. Energy Mater.* **2021**, *11*, 2000806.
- [10] B. Alldritt, F. Urtev, N. Oinonen, M. Aapro, J. Kannala, P. Liljeroth, A. S. Foster, *Comput. Phys. Commun.* **2022**, *273*, 108258.
- [11] B. Huang, Z. Li, J. Li, *Nanoscale* **2018**, *10*, 21320.
- [12] L. Gross, F. Mohn, N. Moll, P. Liljeroth, G. Meyer, *Science* **2009**, *325*, 1110.
- [13] J. Carracedo-Cosme, P. Hapala, R. Pérez, *Mach. Learn.: Sci. Technol.* **2024**, *5*, 025025.
- [14] N. Oinonen, A. V. Yakutovich, A. Gallardo, M. Ondráček, P. Hapala, O. Krejčí, *Comput. Phys. Commun.* **2024**, *305*, 109341.
- [15] P. Hapala, R. Temirov, F. S. Tautz, P. Jelínek, *Phys. Rev. Lett.* **2014**, *113*, 226101.
- [16] P. Hapala, G. Kichin, C. Wagner, F. S. Tautz, R. Temirov, P. Jelínek, *Phys. Rev. B* **2014**, *90*, 085421.
- [17] C.-S. Guo, M. A. Van Hove, R.-Q. Zhang, C. Minot, *Langmuir* **2010**, *26*, 16271.
- [18] K. Kaiser, F. Schulz, J. F. Maillard, F. Hermann, I. Pozo, D. Peña, H. J. Cleaves, A. S. Burton, G. Danger, C. Afonso, S. Sandford, L. Gross, *Meteorit. Planet. Sci.* **2022**, *57*, 644.
- [19] M. Ellner, P. Pou, R. Pérez, *ACS Nano* **2019**, *13*, 786.
- [20] Y. Zhao, P. Altschuh, J. Santoki, L. Griem, G. Tosato, M. Selzer, A. Koeppe, B. Nestler, *Acta Mater.* **2023**, *253*, 118922.
- [21] L. C. Griem, A. Koeppe, A. Greß, T. Feser, B. Nestler, *Acta Mater.* **2025**, *296*, 121220.
- [22] R. Kulagin, P. Reiser, K. Truskovskiy, A. Koeppe, Y. Beygelzimer, Y. Estrin, P. Friederich, P. Gumbsch, *Adv. Eng. Mater.* **2023**, *25*, 2300048.
- [23] Y. Ji, A. Koeppe, P. Altschuh, D. Rajagopal, Y. Zhao, W. Chen, Y. Zhang, Y. Zheng, B. Nestler, *Comput. Mater. Sci.* **2024**, *232*, 112628.
- [24] Y. Ji, A. Koeppe, P. Altschuh, L. Griem, D. Rajagopal, B. Nestler, *Phys. Scr.* **2024**, *99*, 076010.
- [25] B. Alldritt, P. Hapala, N. Oinonen, F. Urtev, O. Krejci, F. F. Canova, J. Kannala, F. Schulz, P. Liljeroth, A. S. Foster, *Sci. Adv.* **2020**, *6*, eaay6913.
- [26] F. Priante, N. Oinonen, Y. Tian, D. Guan, C. Xu, S. Cai, P. Liljeroth, Y. Jiang, A. S. Foster, *ACS Nano* **2024**, *18*, 5546.
- [27] N. Oinonen, C. Xu, B. Alldritt, F. F. Canova, F. Urtev, S. Cai, O. Krejčí, J. Kannala, P. Liljeroth, A. S. Foster, *ACS Nano* **2021**, *16*, 89.
- [28] T. Yu, C. Wang, H. Yang, F. Li, *J. Energy Chem.* **2024**, *90*, 191.
- [29] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, K. A. Persson, *APL Materials* **2013**, *011002*.
- [30] F. Zhou, M. Cococcioni, C. A. Marianetti, D. Morgan, G. Ceder, *Phys. Rev. B - Condens. Matter Mater. Phys.* **2004**, *70*, 235121.
- [31] L. Wang, T. Maxisch, G. Ceder, *Chem. Mater.* **2007**, *19*, 543.
- [32] S. Adams, R. P. Rao, *Phys. Status Solidi A* **2011**, *208*, 1746.
- [33] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, S. Ganguli, *International conference on machine learning* **2015**, *37*, 2256.
- [34] D. Kingma, T. Salimans, B. Poole, J. Ho, *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 21696.

- [35] P. Dhariwal, A. Nichol, *Adv. Neural Inf. Process. Syst.* **2021**, 34, 8780.
- [36] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, *Commun. ACM* **2020**, 63, 139.
- [37] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, A. C. Courville, *Adv. Neural Inf. Process. Syst.* **2017**, 30, 5767.
- [38] L. Metz, B. Poole, D. Pfau, J. Sohl-Dickstein, arXiv preprint arXiv:1611.02163 **2016**.
- [39] A. Koeppe, F. Bamer, B. Markert, *PAMM* **2018**, 18, e201800283.
- [40] A. Koeppe, F. Bamer, B. Markert, *Comput. Methods Appl. Mech. Eng.* **2020**, 366, 113088.
- [41] J. Ho, A. Jain, P. Abbeel, *Adv. Neural Inf. Process. Syst.* **2020**, 33, 6840.
- [42] F.-A. Croitoru, V. Hondru, R. T. Ionescu, M. Shah, *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, 45, 10850.
- [43] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)* **2022**, 1, 10674.
- [44] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, O. Wang, *Proceedings of the IEEE Conf. on computer vision and pattern recognition* **2018**, 1, 586.
- [45] P. Isola, J.-Y. Zhu, T. Zhou, A. A. Efros, *Proceedings of the IEEE conference on computer vision and pattern recognition* **2017**, 1, 5967.
- [46] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, *Adv. Neural Inf. Process.* **2017**, 30.
- [47] H. Alqahtani, M. Kavakli-Thorne, G. Kumar, F. SBSSTC, *International Conf. on Information Technology and Applications (ICITA)* **2019**, 7, 2.
- [48] N. Brandt, L. Griem, C. Herrmann, E. Schoof, G. Tosato, Y. Zhao, P. Zschumme, M. Selzer, *Data Sci. J.* **2021**, 20, 8.
- [49] L. Griem, P. Zschumme, M. Laqua, N. Brandt, E. Schoof, P. Altschuh, M. Selzer, *Data Sci. J.* **2022**, 21, 16.
- [50] A. Koeppe, the CIDS team, CIDS and KadiAI, v3.2a **2025**, <https://gitlab.com/intelligent-analysis/cids/-/tree/v3.2a>.
- [51] F. T. Bülle, N. R. Mathiesen, A. J. Nielsen, T. Vegge, J. M. Garcia-Lastra, I. E. Castelli, *Batteries Supercaps* **2020**, 3, 488.
- [52] N. Mishra, R. Boral, T. Paul, *Mater. Today Phys.* **2025**, 51, 101634.
- [53] C. I. Azambou, F. H. K. Djioko, O. O. Obiukwu, P. K. Tsobnang, E. E. Kalu, I. T. Kenfack, E. E. Oguzie, *Mater. Today Commun.* **2023**, 35, 105738.
- [54] W. Hu, H. Wang, W. Luo, B. Xu, C. Ouyang, *Solid State Ion.* **2020**, 347, 115257.

Manuscript received: January 31, 2025

Revised manuscript received: July 25, 2025

Version of record online: