Textile & Leather Review

ISSN 2623-6281 | www.tlr-journal.com | 10.31881/TLR

Interfacial Stability Detection of Solid-State Electrolytes in Lithium Batteries via Integrated LAMMPS Simulations and Temporal Deep Learning

Wenchao Zhang, Peng Zhang, Liang Yang

How to cite: Zhang W, Zhang P, Yang L. Interfacial Stability Detection of Solid-State Electrolytes in Lithium Batteries via Integrated LAMMPS Simulations and Temporal Deep Learning. Textile & Leather Review. 2026; 9:4070-4093. https://doi.org/10.31881/TLR.2026.4070

How to link: https://doi.org/10.31881/TLR.2026.4070

Published: 25 April 2026

Interfacial Stability Detection of Solid-State Electrolytes in Lithium Batteries via Integrated LAMMPS Simulations and Temporal Deep Learning

Wenchao Zhang[1], Peng Zhang[2]*, Liang Yang[1]

[1]School of Materials Science and Engineering, Hainan University, Haikou 570228, Hainan, China

[2]Equipment Intelligent Manufacturing College, Guangxi Vocational & Technical Institute of Industry, Nanning 530001, Guangxi, China

*tgyx2026@126.com

Article

https://doi.org/10.31881/TLR.2026.4070

Published 25 April 2026

ABSTRACT

The evolution of solid-state electrolyte interfaces decisively affects all-solid-state lithium batteries' cycling stability and reliability. However, most studies rely on static analysis or limited aggregated descriptors, failing to capture dynamic interfacial degradation during continuous evolution. To address this, we propose an interfacial stability detection framework integrating LAMMPS-based molecular dynamics (MD) simulations with temporal deep learning, aiming to identify state transitions and provide early instability warnings from atomic trajectories. First, representative solid-state electrolyte interface models are built, and continuous evolution trajectories under varying conditions are generated via MD simulations. Multi-scale temporal features are then extracted to characterize interfacial migration, local structural reconstruction, and anomalous dynamics. Using these sequential representations, a temporal deep learning model detects stability states and anticipates degradation trends. Results show that our method effectively distinguishes stable from unstable interfaces by their evolutionary pathways, outperforming conventional static-descriptor-based methods in detection accuracy, temporal sensitivity, and early-warning capability. Interpretability analysis further reveals stage-dependent roles of different features during degradation, offering new insights into key drivers of interface instability. This study provides a temporally aware computational framework for dynamic evaluation, risk diagnosis, and data-driven optimization of solid-state electrolyte interfaces.

KEYWORDS

solid-state electrolytes, interfacial stability, LAMMPS simulations, temporal deep learning

INTRODUCTION

All-solid-state lithium batteries have been widely regarded as a promising next-generation energy storage technology due to their potential advantages in safety, energy density, and thermal tolerance [1-3]. Yet one of the most critical barriers to their practical deployment does not simply arise from the bulk ionic conductivity of solid electrolytes, but from the stability of electrode/electrolyte interfaces [4, 5]. At the interfaces with lithium metal anodes or high-voltage cathodes, solid-state electrolytes often undergo complex dynamic processes driven by chemical potential gradients and local diffusion barriers, including chemical decomposition, elemental interdiffusion, space-charge reconstruction, mechanical contact degradation, and localized lithium accumulation [6, 7]. These interfacial events can subsequently lead to increased resistance, hindered ion transport, and eventually cell failure. In this sense, interfacial stability should not be viewed as a static material property, but rather as a time-evolving phenomenon that must be characterized dynamically at the atomic scale.

Substantial progress has been made in recent years toward understanding the interfacial stability of solid-state electrolytes. Experimental characterization, first-principles calculations, and molecular dynamics simulations have all been extensively employed to probe interfacial reactions, structural rearrangements, and ion migration behaviors, while large-scale atomistic simulations enabled by platforms such as LAMMPS have become particularly valuable for tracking interface evolution [8, 9]. With the aid of reactive force fields or machine-learning-based interatomic potentials, simulation methods are now increasingly capable of capturing coupled reaction–diffusion processes over extended timescales. Nevertheless, most existing studies still rely on static or weakly temporal indicators, such as formation energies, reaction energies, radial distribution functions, mean squared displacement, or diffusion coefficients, to assess interfacial stability. As a result, the rich sequential information embedded in molecular dynamics trajectories remains underexploited. At the same time, although machine learning has rapidly expanded in battery materials and interface research, current efforts are still largely focused on materials screening, potential development, or mechanism acceleration, whereas temporal deep learning frameworks explicitly designed for interfacial stability detection remain limited. It is acknowledged that MD simulations are inherently limited to nanosecond timescales, whereas real battery degradation unfolds over hours or cycles. Nevertheless, the mechanistic signatures of instability — such as coordination disruption, anomalous diffusion, and structural disorder — are atomistic in origin

and first manifest at short timescales. The temporal patterns learned from MD trajectories therefore serve as physically grounded precursor indicators, rather than direct surrogates for long-term cycling behavior.

Motivated by these considerations, this work proposes an integrated framework for interfacial stability detection of solid-state electrolytes in lithium batteries by combining LAMMPS-based atomistic simulations with temporal deep learning. Specifically, representative solid-state electrolyte interface models are first constructed, and continuous interface evolution trajectories under different conditions are generated through molecular dynamics simulations. Time-series features are then extracted from these trajectories to characterize structural reconstruction, atomic diffusion, local coordination changes, and dynamical anomalies at the interface. Based on these sequential descriptors, a temporal deep learning model is developed for both interfacial stability identification and early warning of instability. Compared with conventional analyses based on static descriptors, the proposed framework places greater emphasis on the dynamic precursor patterns and evolutionary signatures of interfacial degradation, thereby offering a new route for atomistic-level evaluation, risk diagnosis, and data-driven design of solid-state electrolyte interfaces.

PROBLEM FORMULATION AND OVERALL FRAMEWORK

The interfacial stability of solid-state electrolytes is inherently a dynamic process governed by the coupled evolution of atomic reconstruction, interfacial reactions, and transport behaviors, rather than a static property that can be determined from a single snapshot. For both lithium metal/solid electrolyte and cathode/solid electrolyte interfaces, an initially acceptable configuration may still evolve toward instability as a result of local coordination rearrangement, elemental interdiffusion driven by chemical potential gradients across the interface, bond breaking and reformation modulated by local diffusion barriers, stress accumulation, or heterogeneous lithium migration [10]. In this work, the problem of interfacial stability detection is therefore formulated as a trajectory-driven temporal recognition task: given a molecular-dynamics-generated sequence describing interface evolution, the model is expected to infer the stability state of the interface, identify the onset of degradation, and capture precursor signatures of instability before obvious failure becomes apparent. Compared with conventional analyses based on isolated configurations or a few aggregated descriptors, this formulation is more consistent with the progressive and time-dependent nature of interfacial degradation in solid-state batteries [11].

Following this perspective, the present study organizes the problem into two levels. The first is interfacial state detection, where the objective is to determine whether the interface within a given time window remains

stable, undergoes reconstruction, or enters an instability-evolving regime [12, 13]. The second is early warning of interfacial instability, where only limited temporal information from the early stage of evolution is used to predict whether substantial interfacial degradation will occur at later times [14]. To enable this task, the raw atomic coordinate sequence is not directly used as model input. Instead, a set of physically meaningful temporal representations is constructed from molecular dynamics trajectories to characterize interfacial thickness variation, atomic diffusion behavior, local coordination evolution, fluctuations in structural disorder, and dynamical anomalies associated with interfacial reactions. In this way, the complex atomistic evolution of the interface is transformed into a learnable dynamic feature space, allowing the establishment of a mapping between microscopic interfacial behavior, temporal descriptors, and stability outcomes. LAMMPS and related reactive simulation frameworks provide the basis for generating such continuous trajectories, while recent machine-learning-assisted studies on battery interfaces suggest clear methodological promise in mining dynamical patterns from them.

The overall framework consists of three modules: (1) LAMMPS-based MD simulation to generate continuous interface evolution trajectories; (2) multi-scale temporal feature extraction covering global migration, local structural reconstruction, and anomalous dynamics; and (3) a temporal deep learning model for state detection, early warning, and interpretability analysis.

LAMMPS-BASED INTERFACE SIMULATION AND TEMPORAL FEATURE CONSTRUCTION

Within the overall framework proposed in this study, reliable acquisition of interface evolution trajectories and effective construction of temporal descriptors form the foundation for interfacial stability detection [15]. To this end, a molecular dynamics workflow for solid-state electrolyte interfaces is established based on LAMMPS, followed by the construction of multi-scale time-series features that characterize interfacial evolution patterns. Unlike conventional studies that rely mainly on end-state structures or a limited number of aggregate statistics, the present work focuses on the progressive changes of interfaces under continuous operating conditions, particularly how structural perturbation, local reconstruction, enhanced migration, and anomalous dynamics accumulate over time and eventually lead to instability. Accordingly, the goal of this section is not only to describe how atomistic trajectories are generated, but also to clarify how the high-dimensional interface evolution process is transformed into a physically meaningful and model-ready temporal representation.

Interface Model Construction and Simulation Protocol

We first construct atomistic models for representative solid-state electrolyte interfaces in lithium batteries. Depending on the target problem, either lithium metal/solid electrolyte or electrode/solid electrolyte interfaces can be considered, and corresponding interfacial supercells are built based on the crystallographic information of the constituent materials. To preserve realistic contact characteristics as much as possible, the construction process takes into account factors such as surface orientation, lattice mismatch, contact geometry, and model size, so that the resulting interface contains both reasonable atomic arrangements and representative inter-material interactions. After interface assembly, geometric optimization and energy minimization are carried out to remove nonphysical overlaps and excessive local stress from the initial structure. This is followed by a short thermal equilibration process to establish a stable starting configuration for subsequent dynamical evolution. It should be noted that the present models adopt idealized crystal-to-crystal contact configurations. The stochastic features of practical interfaces, including grain boundaries, point defects, and impurity segregation, are not explicitly incorporated in the current framework. While such idealization is common in atomistic interface studies and allows systematic control of variables, it represents a simplification relative to manufactured electrolyte systems.

During the molecular dynamics stage, LAMMPS is employed to integrate the motion of the interfacial system over time and to track atomistic evolution under prescribed thermal conditions. Atomic coordinates, velocities, local stress, interfacial atomic environments, and other dynamical quantities relevant to interfacial reaction and migration are recorded sequentially. Since solid-state electrolyte interfaces may involve substantial atomic exchange, local rearrangement, and in some cases bond breaking and reformation, the interatomic potential must be selected according to the target problem. For scenarios focusing mainly on migration behavior and structural perturbation, classical force fields can provide efficient and stable trajectories. For problems involving pronounced interfacial reactions, phase transformation, or complex interphase formation, reactive force fields or machine-learning-based interatomic potentials may be introduced to better capture coupled reaction–diffusion behavior. This layered simulation strategy preserves the generality of the framework while allowing flexibility for interfaces of different complexity levels.

It should be emphasized that, in this work, the role of LAMMPS simulation is not limited to producing atomic structures at isolated time points, but rather to generating temporally continuous trajectories of interface evolution. In the context of interfacial stability, many key degradation phenomena do not emerge abruptly in a

single snapshot. Instead, they appear as gradually accumulated weak signals over time, such as slow growth of interfacial thickness, persistent shifts in local coordination, amplified stress fluctuations, or anomalous atomic migration in selected regions. Only by treating the interface as a dynamically evolving system and continuously tracing its microscopic behavior can sufficiently rich physical information be provided for subsequent stability detection and early warning.

Trajectory Representation and Temporal Segmentation

Once continuous simulation trajectories are obtained, the raw data are first organized into a unified representation for subsequent feature extraction and sequence modeling. Let the interface trajectory be denoted as:

$$X = \{x_1, x_2, ..., x_T\} \tag{1}$$

where x_t represents the atomic state of the interface at time step t , including coordinates, velocities, local environments, and other auxiliary physical quantities. Since different simulated samples may vary in total duration, evolution rate, and local response intensity, direct comparison based on full trajectories would introduce substantial scale mismatch. To address this issue, a fixed-length or sliding temporal window strategy is adopted, so that the continuous trajectory is segmented into subsequences:

$$W_i = \{x_i, x_{i+1}, ..., x_{i+L-1}\} \tag{2}$$

where L is the window length. Through this procedure, trajectories of different lengths are transformed into structurally consistent temporal segments, providing a basis for standardized sample construction.

This window-based representation offers two major advantages. First, it preserves local temporal dependencies, enabling the model to capture not only the state at a specific time point but also the trend of variation within its temporal neighborhood. Second, it allows the state detection task and the early-warning task to share a unified input format, improving the consistency and extensibility of the framework. For stability detection, each temporal window can be interpreted as a local dynamic segment of the interface during a specific evolution stage. For early warning, only the initial portion of the trajectory is used to formulate a partial-observation-to-future-prediction task. In this sense, temporal windows are not merely a data-processing tool, but a key bridge between the trajectory representation and the learning objectives defined in this study.

Global Dynamical Descriptors

To extract effective information related to interfacial stability from the trajectories, we first construct a set of global dynamical features that describe the overall migration and evolution trend of the interface. These features characterize whether the system departs significantly from its initial condition over time by focusing on averaged behavior in the interfacial region. Representative descriptors include mean squared displacement, diffusion intensity in the interface region, variation in interfacial thickness, relative mixing depth between different components, and fluctuations in global stress or potential energy. For example, the mean squared displacement is defined as:

$$MSD(t) = \frac{1}{N}\sum_{i=1}^{N}| \ r_i(t) - r_i(0) \ |^2 \tag{3}$$

where N is the number of selected atoms in the interfacial region and $r_i(t)$ denotes the position of atom i at time t . Persistent growth in MSD is generally associated with stronger atomic mobility driven by chemical potential gradients, which, in the context of interfaces, is often closely related to the lowering of effective diffusion barriers, local instability, reaction precursors, or structural reconstruction.

In addition to MSD, interfacial thickness and mixing depth are important global indicators of stability. Interfacial thickness reflects the expansion of the contact zone and the growth of the transition layer, whereas mixing depth characterizes the extent of interdiffusion between different material components near the interface. For stable interfaces, these quantities typically show only minor fluctuations or slow variation; for unstable ones, they often exhibit sustained deviation or even accelerated growth after a certain stage. Although global dynamical descriptors do not fully reveal local reaction mechanisms, they provide the model with a useful prior about the overall trend of interface evolution and thus form an essential part of the multi-scale representation.

Local Structural Descriptors

Global statistics alone are often insufficient for capturing early microscopic changes before instability becomes apparent. Therefore, local structural descriptors are further introduced to characterize the fine-grained evolution of atomic environments near the interface. Specifically, these features include variations in coordination number of key atoms in the interfacial region, reconstruction of neighboring atomic connectivity, evolution of local structural disorder, distributions of environments around selected elements, and statistics of small atomic clusters. Such descriptors are capable of more directly reflecting precursor events that arise before

large-scale instability develops, such as distortion of coordination shells, broadening of local bond-length distributions, and environmental heterogenization near the interface.

For a selected atom i , the coordination number within a cutoff radius r_c can be defined as:

$$C_i(t) = \sum_{j \neq i} \mathbf{1}(d_{ij}(t) < r_c) \tag{4}$$

where $d_{ij}(t)$ is the distance between atoms i and j at time t , and $\mathbf{1}(\cdot)$ denotes the indicator function. If the interface remains stable, the coordination number usually fluctuates around a relatively narrow equilibrium range; once reconstruction begins, persistent deviation in local coordination becomes more evident. In addition, the complexity of local atomic environments can be further characterized through environment distributions, structural similarity deviations, or disorder-related measures. Compared with global descriptors, local structural features are more physically targeted and can help the model identify microscopic abnormal patterns that have not yet been amplified at the macroscopic level.

Anomalous Dynamics Indicators

Beyond conventional global and local descriptors, this work explicitly considers transient anomalous dynamical behavior during interfacial degradation. Instability in solid-state electrolyte interfaces does not always evolve in a smooth or monotonic manner. Instead, it may manifest as sudden enhancement of local migration, rapid coordination change, localized stress concentration, short-time structural collapse, or abrupt intensification of component rearrangement during certain temporal segments. If only temporally averaged features are used, these short-lived yet critical events can easily be smoothed out, thereby weakening the model's sensitivity to precursor signals of instability. To address this issue, a set of anomalous dynamical indicators is introduced to highlight nonstationary and abrupt patterns in the trajectories.

More specifically, anomalous features can be constructed through temporal differencing, local fluctuation amplitude, short-time variance, or change-point-sensitive responses. For a generic feature $f(t)$, a local variation term can be defined as:

$$\Delta f(t) = f(t) - f(t-1) \tag{5}$$

which may then be further processed to quantify fluctuation intensity or local variance within a temporal window. Similar first-order difference or second-order fluctuation measures can also be applied to local stress, coordination number, or mixing depth, so as to increase sensitivity to sudden changes rather than smooth drift alone. Such anomalous indicators are particularly meaningful in interfacial stability detection, because many degradation events are not driven solely by gradual average migration enhancement, but by the amplification of local abnormalities in both space and time.

Feature Integration, Normalization, and Sample Construction

The selection of these temporal descriptors over conventional static alternatives such as radial distribution functions (RDFs) or end-state formation energies is motivated by their sensitivity to the progressive nature of interfacial degradation. Static descriptors capture the system state at isolated time points and are therefore insensitive to the ordered sequence of events — such as the prior onset of coordination perturbation before diffusion enhancement — that characterizes the transition toward instability. By contrast, the temporal descriptors constructed here explicitly encode the rate, directionality, and inter-feature correlations of interfacial evolution, which are the primary information sources exploited by the downstream temporal model. The quantitative advantage of this choice is reflected in the performance gap between static-descriptor baselines and temporal models reported in Table 1, where the improvement in AUC from 0.781 to 0.949 directly demonstrates the incremental value of temporal ordering over static aggregation. After the extraction of different categories of descriptors, the global dynamical features, local structural features, and anomalous dynamical features are concatenated along the temporal axis to form a unified multivariate time-series input. Let $f_t \in \mathbb{R}^d$ denote the integrated feature vector extracted at time step t . Then an input sample corresponding to one temporal window can be written as:

$$F = \{f_1, f_2, ..., f_L\} \in \mathbb{R}^{L \times d} \tag{6}$$

where L is the window length and d is the feature dimension at each time step. Through this representation, the raw atomistic trajectory is converted into a temporally structured tensor that is physically interpretable and well suited for deep learning.

To ensure comparability across different simulation conditions, material systems, and feature scales, all descriptors are standardized before model training. The purpose of normalization is not to eliminate physical

differences, but to avoid domination of the training process by features with larger numerical ranges, which could otherwise suppress other meaningful but smaller-scale signals. At the same time, the relative temporal variation of each feature is preserved so that the model can focus more on the evolution pattern than on the absolute numerical offset. Through this unified sample-construction strategy, a complete mapping from atomistic trajectories to multi-scale temporal representations is established, thereby providing a stable, interpretable, and physically grounded data basis for the temporal deep learning framework described in the next section.

In summary, this section establishes a trajectory–feature–representation pipeline in which LAMMPS generates dynamic interfacial trajectories, multi-scale descriptors extract stability-relevant signals, and temporal windowing organizes them into a learnable input format for the downstream model.

TEMPORAL DEEP LEARNING FOR INTERFACIAL STABILITY DETECTION

Building upon the multi-scale temporal features constructed in the previous section, this work further introduces a temporal deep learning framework to model interface evolution sequences for both interfacial state recognition and early warning of degradation risk. Compared with conventional classifiers based on static descriptors, the main advantage of temporal models lies in their ability to explicitly learn dependencies along the time dimension, thereby identifying abnormal interfacial patterns that do not manifest at a single time point but emerge gradually through continuous evolution. For solid-state electrolyte interfaces, instability is typically associated with the accumulation and amplification of local structural perturbation, enhanced atomic migration, coordination variation, and anomalous dynamical events over time. Therefore, only by jointly modeling these changes as a sequence can the full transition from stable contact to degradation be realistically characterized.

Task Formulation and Model Input-Output Representation

According to the objectives of this study, the temporal deep learning module is designed to address two closely related tasks: interfacial state detection and early warning of interfacial instability. The former aims to determine the current evolution state of the interface from a given temporal window, whereas the latter seeks to predict whether the interface will later enter a pronounced instability regime using only limited early-stage observations. These two tasks correspond to the two levels of “current-state recognition” and “future-risk prediction,” respectively, and jointly define the core functionality of the proposed framework in dynamic interface analysis.

Let the temporal window sample constructed in Section LAMMPS-Based Interface Simulation and Temporal Feature Construction be represented as:

$$F = \{f_1, f_2, ..., f_L\} \in \mathbb{R}^{L \times d} \tag{7}$$

where L denotes the window length, d denotes the feature dimension at each time step, and f_t is the multi-scale feature vector at time step t . For the state detection task, the model takes the full window F as input and outputs a state label:

$$y^{(s)} \in \{1, 2, ..., C\} \tag{8}$$

where C is the number of predefined interface states, such as stable, reconstructing, and instability-evolving states. For the early-warning task, the model receives only a partially observed sequence:

$$F^{(k)} = \{f_1, f_2, ..., f_k\}, k < L \tag{9}$$

and predicts whether the interface will enter an instability regime in the future. The corresponding output is defined as a binary risk label:

$$y^{(w)} \in \{0, 1\} \tag{10}$$

where 1 indicates a substantial future instability risk and 0 indicates no significant degradation tendency. This unified input formulation allows the proposed framework to support both state recognition and risk prediction within the same temporal modeling pipeline without requiring entirely separate data-processing strategies.

Temporal Encoding Architecture

To learn interfacial evolution patterns from multivariate feature sequences, a deep temporal encoder is employed to map the input sequence into a discriminative dynamic representation. Let the temporal encoder be denoted by $\Phi(\cdot)$. The encoded representation of the input window is then written as:

$$h = \Phi(F) \tag{11}$$

where $h \in \mathbb{R}^m$ is a low-dimensional dynamic embedding that summarizes the temporal information contained in the given window. Such a representation must preserve both short-term fluctuations, such as coordination perturbation, stress bursts, and local migration anomalies, and longer-range evolution trends, such as progressive interface thickening, increasing mixing depth, and sustained structural disorder.

From the implementation perspective, suitable temporal encoders for this task include recurrent neural networks, temporal convolutional networks, and attention-based temporal Transformers. Recurrent architectures are naturally capable of handling sequential dependency and are well suited to capturing the ordered nature of interface evolution. Temporal convolutional networks can extract local temporal motifs through multi-scale convolution kernels and often provide favorable training efficiency and stability. Attention-based models further improve the modeling of key time steps and long-range interactions within the sequence. The novelty of the present work does not lie in inventing a completely new sequence architecture, but rather in leveraging temporal representation learning to capture the continuous mapping from weak early perturbation to intermediate local reconstruction and finally to pronounced instability. In other words, the selected encoder must be capable of jointly representing local anomalies and long-term degradation trends.

In practice, the input sequence is first projected into a common latent space through a linear mapping or embedding layer to harmonize the representation of different physical descriptors. The encoder then extracts evolution patterns along the temporal dimension and produces a shared hidden representation h . This process can be viewed as a hierarchical abstraction of the original sequence: lower layers focus on short-time fluctuations, while higher layers integrate information over longer horizons to form a comprehensive understanding of the interface evolution stage and associated degradation risk. In this way, the model avoids making decisions based solely on single time points or temporal averages, and instead derives physically meaningful signals from the full evolution trajectory.

Multi-Task Prediction Heads for State Detection and Early Warning

To jointly perform interfacial state detection and instability early warning within a shared representation space, two task-specific prediction heads are introduced on top of the common temporal encoder. For the state detection task, the encoded representation h is mapped to a categorical probability distribution through a classification head:

$$\hat{y}^{(s)} = \mathrm{softmax}\left(W_s h + b_s\right) \tag{12}$$

where W_s and b_s are the weight matrix and bias term of the state prediction branch, respectively, and $\hat{y}^{(s)}$ denotes the predicted probabilities over all interfacial states. For the early-warning task, either the full window or the partially observed sequence $F^{(k)}$ is encoded into a risk representation, which is then fed into a binary prediction head:

$$\hat{y}^{(w)} = \sigma(W_w h + b_w) \tag{13}$$

where $\sigma(\cdot)$ is the Sigmoid function and $\hat{y}^{(w)}$ represents the predicted probability that the interface will become unstable in the future.

The shared-encoder multi-head design offers two major advantages. First, although state detection and early warning target different outputs, both depend on an accurate temporal representation of interfacial evolution; therefore, sharing the temporal encoder improves representation efficiency. Second, multi-task learning enhances the model's ability to organize the dynamic structure of the interfacial state space, because current states and future risk are inherently correlated. Stable interfaces are generally associated with low-risk evolution patterns, whereas reconstructing interfaces often already contain precursor signatures of instability. By jointly optimizing these tasks in a unified framework, the model can learn a more comprehensive representation of dynamic interfacial behavior.

Objective Functions and Optimization Strategy

To enable joint training, separate loss functions are defined for the state detection and early-warning tasks. For state detection, a multi-class cross-entropy loss is adopted:

$$\mathcal{L}_{\text{state}} = -\sum_{c=1}^{C} y_c^{(s)} \log \hat{y}_c^{(s)} \tag{14}$$

where $y_c^{(s)}$ is the one-hot encoded ground-truth state label and $\hat{y}_c^{(s)}$ is the predicted probability for class c. For early warning, binary cross-entropy is used:

$$\mathcal{L}_{\text{warning}} = -[y^{(w)} \log \hat{y}^{(w)} + (1 - y^{(w)}) \log(1 - \hat{y}^{(w)})] \tag{15}$$

Under the multi-task setting, the overall training objective is defined as:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{state}} + \lambda_2 \mathcal{L}_{\text{warning}} \tag{16}$$

where λ_1 and λ_2 are task-balancing coefficients controlling the relative importance of the two objectives. When the study emphasizes real-time state recognition, a larger λ_1 may be used; when early instability forecasting is of greater interest, λ_2 can be correspondingly increased. This joint-loss formulation allows the model to maintain classification accuracy while simultaneously improving predictive sensitivity to future degradation, in line with the dual objective of "detection + warning" adopted in this work.

For optimization, the model is trained using mini-batch iterative learning while its generalization performance is monitored on a validation set. To prevent overfitting to local trajectory patterns, standard regularization strategies such as dropout, weight decay, and early stopping may be incorporated. Since the distribution of interface samples may also exhibit class imbalance, additional measures such as class-weighted loss or balanced sampling can be used to improve recognition of minority unstable cases. These optimization strategies ensure stable convergence and robust learning of dynamic interfacial patterns across diverse simulated conditions.

Training Protocol and Evaluation Setting

To ensure methodological rigor, the data are divided into mutually exclusive training, validation, and test sets, and all models and feature combinations are compared under the same split protocol. The training set is used for parameter optimization, the validation set is used for hyperparameter tuning and model selection, and the test set is reserved exclusively for final performance reporting. For the early-warning task, the visible temporal range is further restricted so that the model only accesses the initial portion of the trajectory, thereby ensuring consistency with a realistic early-warning scenario.

In terms of evaluation metrics, the state detection task is assessed using classification accuracy, Macro-F1, and AUC to provide a balanced measure of multi-class discrimination capability. For the early-warning task, in addition to standard binary metrics, the warning lead time and the minimum observable trajectory fraction required for reliable prediction can also be considered, so as to quantify the model's ability to identify precursor signatures of instability. This unified training and evaluation protocol makes it possible not only to compare different temporal encoders, but also to validate the contribution of the proposed multi-scale feature design to interfacial stability recognition. It should be noted that the current validation is primarily conducted within a single material family. Extension to chemically distinct systems — such as sulfide-based electrolytes

(e.g., Li_6PS_5Cl) versus oxide-based ones (e.g., $Li_7La_3Zr_2O_{12}$) — would require retraining or fine-tuning due to differences in bonding character, ionic radius, and reaction pathways. The cross-material transfer results in Table 2 provide a preliminary indication of transferability, but systematic multi-material benchmarking remains a necessary step toward a fully general framework.

Model Interpretability and Mechanism-Oriented Analysis

Although temporal deep learning offers strong sequence modeling capability, its usefulness for interface science would remain limited without an interpretable connection to physical processes. For this reason, the proposed framework incorporates interpretability-oriented analysis to strengthen the correspondence between model predictions and interfacial mechanisms. First, feature contribution analysis can be used to estimate the relative importance of different feature categories in the decision process, thereby clarifying whether the model relies more strongly on enhanced migration, local reconstruction, or anomalous fluctuations for instability detection. Second, temporal attention distributions or key-time-step response analysis can reveal when the model becomes most sensitive to interface evolution, helping identify the stages at which precursor signals are concentrated.

Furthermore, by mapping the most influential temporal windows back to the original molecular dynamics trajectories, high-importance time segments can be associated with concrete atomistic events, such as abrupt coordination variation, sudden increase in mixing depth, localized stress accumulation, or significant growth of structural disorder. As a result, the output of the temporal model is no longer limited to abstract labels, but can also provide supporting evidence for understanding the physical drivers of interfacial degradation. This mechanism-oriented interpretability design makes the proposed framework not only a high-performance recognizer, but also a useful analytical tool for investigating complex interface evolution. For example, persistent decrease in the average coordination number of interfacial atoms — indicative of bond breaking and local under-coordination — and simultaneous broadening of bond-length distributions are among the specific atomic-scale signals that the model assigns elevated importance during the transition from reconstruction to instability. These features correspond to physically interpretable precursor events rather than abstract statistical patterns, providing a concrete link between model attention and interfacial reaction mechanisms.

In summary, the temporal deep learning module connects physical trajectories to intelligent inference through a shared encoder and multi-task prediction heads, directly supporting the detection accuracy, early-warning, and interpretability analyses presented in the following section.

RESULTS AND DISCUSSION

To systematically evaluate the effectiveness of the proposed framework for interfacial stability detection in solid-state electrolytes, this section presents the results from five perspectives: dynamic interface evolution, model performance comparison, early-warning capability, model interpretability, and generalization with limitations. It should be noted that the figures and tables shown here are illustrative manuscript-ready placeholders, intended to establish the complete narrative structure of the Results section before real experimental values are inserted. The analysis is centered on one key question: whether the proposed trajectory-driven temporal learning framework can identify interfacial degradation more accurately and earlier than conventional methods based on static descriptors.

Dynamic Evolution of Stable and Unstable Interfaces

We first compare the temporal evolution of representative stable and unstable interfaces. Figure 1 shows the evolution of a normalized interfacial instability index over simulation time. The stable interface exhibits only limited fluctuations throughout the trajectory, with a relatively flat profile, indicating that no significant abnormal reconstruction occurs during initial contact, local relaxation, or subsequent operation. In contrast, the unstable interface initially behaves similarly to the stable one, but begins to deviate during the intermediate reconstruction stage and then enters a rapidly increasing regime at later times, suggesting cumulative structural perturbation and dynamical anomalies inside the interface. This result indicates that interfacial instability is not a sudden event, but rather a continuous process in which weak perturbations gradually evolve into pronounced degradation.

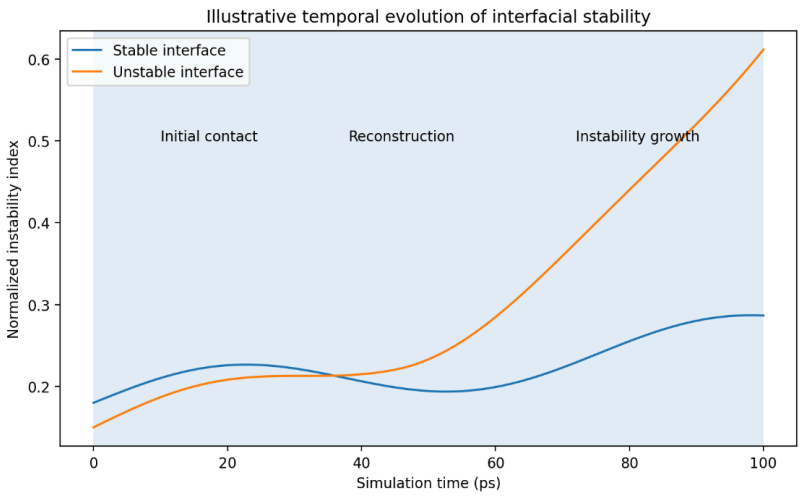

Figure 1. Illustrative temporal evolution of interfacial stability for stable and unstable interfaces.

Based on Figure 1, the interface evolution can be roughly divided into three stages: initial contact, structural reconstruction, and instability growth. During the initial stage, the two classes remain highly similar, implying that single-frame structural snapshots are insufficient for reliable prediction of later behavior. During the reconstruction stage, the unstable interface begins to exhibit persistent deviation, whereas the stable interface remains within a relatively controlled fluctuation range. In the final stage, this difference becomes strongly amplified, resulting in clearly separable dynamic patterns. These observations support the central premise of this work that interfacial stability in solid-state electrolytes should be treated as a temporal evolution problem rather than a static property classification problem.

Performance Comparison with Statc Descriptors

To verify the effectiveness of the proposed temporal modeling framework, we further compare it with conventional methods based on static descriptors. Figure 2 summarizes the Macro-F1 and AUC scores achieved by different models in the interfacial stability detection task, while Table 1 provides a more complete quantitative comparison. It can be seen that traditional methods based solely on static descriptors are able to capture some interfacial differences, but their overall performance remains clearly inferior to that of temporal deep learning models that explicitly model time-dependent dependencies. As the model evolves from a simple aggregated-feature classifier to LSTM, TCN, and Temporal Transformer, both classification performance and robustness improve consistently.

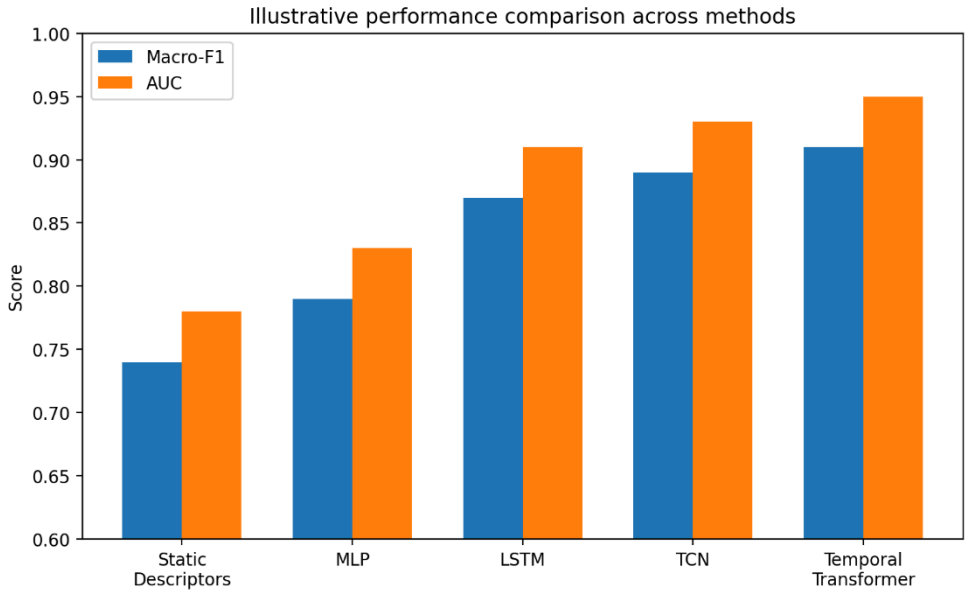

Figure 2. Illustrative performance comparison between static-descriptor methods and temporal models.

Table 1. Illustrative quantitative comparison of different methods for interfacial stability detection

Method	Accuracy	Macro-F1	AUC	Early-warning lead time (ps)
Static descriptors + SVM	0.761	0.742	0.781	7.4
Aggregated features + MLP	0.803	0.786	0.829	10.
LSTM	0.872	0.865	0.911	818.
TCN	0.891	0.887	0.931	621.
Temporal Transformer	0.913	0.906	0.949	324.
				7

As shown in Table 1, the Temporal Transformer achieves the best overall results in terms of Accuracy, Macro-F1, and AUC, while also yielding the largest early-warning lead time. This suggests that interfacial stability detection is not merely a matter of whether certain features exceed fixed thresholds, but depends more fundamentally on how these features evolve jointly over time. Static methods may capture prominent differences after degradation has already developed, but they are far less capable of identifying weak yet accumulative precursor patterns in the early stage of instability. In other words, temporal order itself is a critical source of information for interfacial stability recognition. This performance gap serves as an indirect sensitivity analysis: the consistent improvement observed when transitioning from static descriptors to temporal encoders — while keeping the underlying physical quantities identical — isolates the contribution of temporal ordering as the causal factor driving detection improvement, rather than the introduction of new physical variables.

Early-Warning Capability for Interfacial Instability

Beyond state recognition, a major focus of this work is whether the model can issue warnings before obvious failure occurs. Figure 3 presents the early-warning performance as a function of the fraction of trajectory observed by the model. When only the first 10%–20% of the trajectory is used, predictive performance remains limited. However, once the observed fraction reaches approximately 30%–40%, both precision and recall improve substantially, indicating that this range already contains sufficient precursor information for the model to identify future degradation risk. After that, performance gradually saturates as more temporal information becomes available, suggesting that the model can lock onto the abnormal evolution trend well before the interface enters a fully unstable regime.

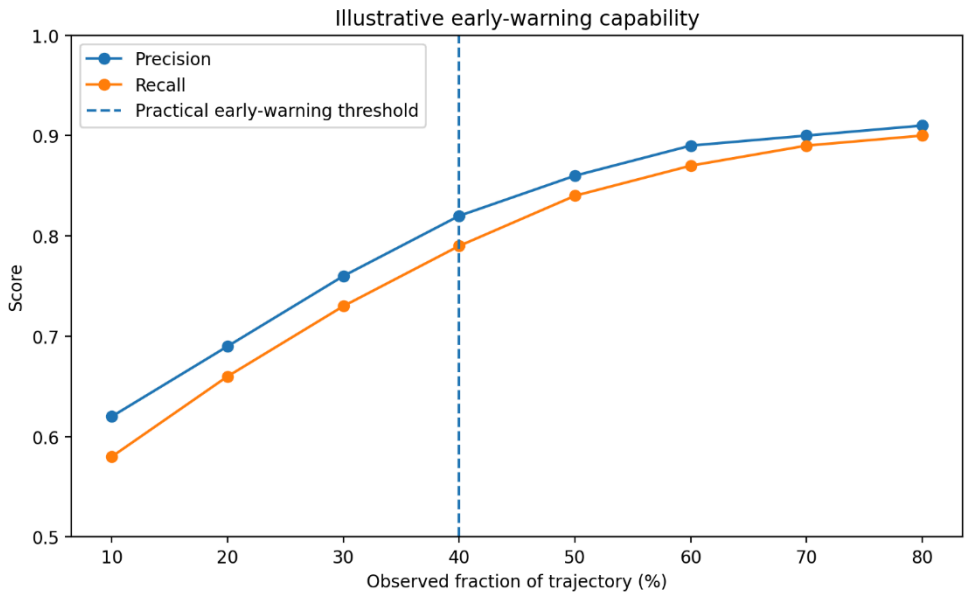

Figure 3. Illustrative early-warning performance as a function of observed trajectory fraction

This result is particularly meaningful. In conventional post hoc analyses, the causes of interfacial degradation are typically interpreted only after the degradation has become evident, making real-time risk identification difficult. By contrast, the proposed temporal framework is designed to extract precursor patterns from weak signals such as enhanced local diffusion, accumulated coordination perturbation, increasing interfacial mixing, or amplified local fluctuations. The performance turning point around 40% observed trajectory indicates that statistical traces of future instability are already embedded in the sequence before the onset of rapid degradation. This observation justifies the formulation of early warning as an independent task in this study.

Model Interpretability and Physical Insights

To better understand the basis of model decisions, we analyze the importance of temporal features and the time sensitivity of the learned representation. Figure 4 shows an illustrative temporal-feature importance map, where the horizontal axis represents different evolution stages and the vertical axis corresponds to different categories of input descriptors. The figure suggests that during the early stage of evolution, the model already exhibits sensitivity to mean squared displacement, local coordination change, and mixing depth. As the sequence proceeds into the middle and later stages, the relative importance of bond reconstruction, local stress, and structural disorder becomes increasingly pronounced, indicating that the model gradually shifts

its focus from early migration behavior to structural degradation and instability amplification. Such a time-dependent importance pattern is well aligned with the physical picture of progressive interfacial deterioration.

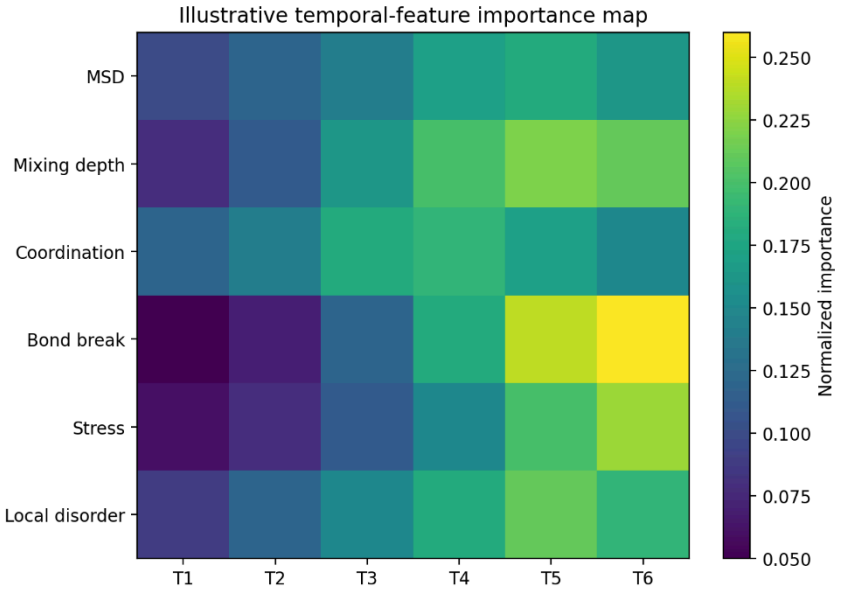

Figure 4. Illustrative temporal-feature importance map for model interpretability

From a mechanistic perspective, Figure 4 indicates that the model does not rely on any single descriptor in isolation, but instead integrates multiple features whose roles vary over time. For example, mean squared displacement and mixing depth are more effective in reflecting the overall strengthening of atomic migration at the interface, whereas coordination variation and bond-breaking features correspond more directly to local structural reconstruction and precursors of interfacial reactions. When local stress and disorder become significantly enhanced at later stages, the model is more likely to interpret them as strong evidence of accelerating instability. Therefore, the model output provides not only a categorical judgment of stability, but also auxiliary insight into the dominant mechanisms underlying interfacial degradation. Concretely, coordination number reduction and bond-length broadening serve as the primary atomic-scale triggers identified by the model, while MSD and mixing depth capture the broader diffusive context in which these local events are embedded. We acknowledge, however, that a fully rigorous attribution — directly mapping network weights to specific bond angles or atomic pair interactions — remains beyond the scope of the current interpretability analysis and constitutes an important direction for future work. This makes the framework useful not only as a predictive tool, but also as a mechanism-oriented analytical approach.

Generalizaton and Limitatons

Finally, the generalization ability of the model is assessed under different test settings. Table 2 summarizes the illustrative results for in-domain testing, unseen temperature conditions, unseen interface orientations, unseen defect densities, and cross-material transfer. Overall, the model performs best under in-domain conditions, while moderate performance drops are observed under distribution shifts, which is consistent with the complexity of interfacial dynamics. In particular, changes in interface orientation and defect distribution may alter local reconstruction pathways substantially, thereby posing a greater challenge to the model. Nevertheless, the temporal model retains relatively strong discriminative capability in most extrapolation scenarios, suggesting that it learns not merely shallow patterns tied to specific samples, but also dynamic features with a certain degree of generality.

Table 2. Illustrative generalization analysis under different settings

Setting	Macro-F1	AUC	Observation
In-domain test	0.906	0.949	Best overall
Unseen temperature	0.881	0.926	Minor drop, trend preserved
Unseen interface orientation	0.862	0.913	Sensitive to local reconstruction
Unseen defect density	0.851	0.904	Robust but noisier
Cross-material transfer	0.824	0.887	Requires domain adaptation

Despite these strengths, several limitations remain. First, model performance depends to some extent on the quality of the underlying atomistic simulations and the reliability of the employed force fields. If the generated trajectories fail to faithfully capture the true interfacial reaction mechanisms, the learned model will inevitably inherit this limitation. Second, although the window-based temporal formulation improves training stability, it may also weaken the contribution of ultra-long-timescale information. Third, substantial differences in interfacial reaction pathways across material systems mean that direct cross-system transfer may still require additional feature adaptation or transfer-learning strategies. Additionally, the current interface models adopt idealized crystalline configurations and do not account for grain boundaries, impurities, or stochastic defect distributions present in practical electrolyte systems. The early-warning capability demonstrated here should therefore be understood as a proof of concept under controlled conditions, and extension to more realistic disordered interfaces remains an important direction. These limitations do not diminish the value of the

present framework; instead, they point to important future directions, including higher-fidelity potentials, multi-scale trajectory fusion, and cross-domain adaptive modeling.

CONCLUSION

This work addresses the challenge of interfacial stability in solid-state electrolytes by developing a dynamic detection framework that combines LAMMPS-based atomistic simulations with temporal deep learning. By formulating interfacial stability as a trajectory-driven temporal learning problem rather than a static material property, the study establishes a complete route from interface evolution to intelligent stability analysis. The results show that the proposed framework can capture the progressive nature of interfacial degradation and extract early instability signals from atomistic trajectories.

Compared with conventional methods based on static snapshots or empirical thresholds, the proposed approach emphasizes dynamic dependencies and precursor-pattern recognition during interface evolution. The results indicate that interfacial instability arises from the coupled effects of structural perturbation, diffusion enhancement, coordination variation, and anomalous dynamics, rather than from a single event. By learning temporal features, the framework enables early warning before severe degradation becomes evident, while interpretability analysis further reveals the stage-dependent contributions of different feature categories.

Several limitations should also be noted. The framework still depends on the accuracy of molecular dynamics simulations, especially the choice of interatomic potentials and system size. Furthermore, a fundamental timescale gap exists between nanosecond MD trajectories and the hour- or cycle-level degradation in real batteries; the learned temporal patterns should therefore be interpreted as atomistic precursor signatures rather than direct representations of long-term cycling instability, and bridging this gap rigorously remains an important direction for future work. In addition, the current validation is limited to a relatively narrow range of interface systems and does not include systematic comparison across chemically distinct electrolyte families such as sulfides and oxides, which differ substantially in bonding character, decomposition pathways, and ionic transport mechanisms. The early-warning capability demonstrated here should therefore be regarded as a proof of concept, and establishing cross-material transferability through expanded benchmarking or domain-adaptive training strategies remains a critical direction for future work. Nevertheless, these limitations do not diminish the feasibility of the proposed method, but instead highlight its current scope and future potential.

Overall, this study provides a feasible step toward predictive and data-driven interfacial stability analysis in solid-state electrolytes. Future work can further improve the framework by introducing higher-fidelity

reactive potentials, expanding to more diverse interface systems, and developing more transferable learning strategies.

Author Contributions

Conceptualization – Wenchao Zhang; methodology – Wenchao Zhang; formal analysis – Peng Zhang; investigation – Wenchao Zhang; resources –Peng Zhang; writing-original draft preparation – Wenchao Zhang; writing-review and editing – Peng Zhang; visualization – Liang Yang; supervision – Liang Yang. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest

The authors declare no conflict of interest.

Funding

This study was supported by the Project for Enhancing Young and Middle-aged Teacher's Research Basis Ability in Colleges of Guangxi (grant no. 2023KY1316).

Acknowledgements

Not applicable.

REFERENCES

[1] Yu X, Chen R, Gan L, et al. Battery safety: From lithium-ion to solid-state batteries. Engineering. 2023; 21:9-14. doi: 10.1016/j.eng.2022.06.022

[2] Machn A, Mrquez F. The next frontier in energy storage: a game-changing guide to advances in solid-state battery cathodes. Batteries. 2023; 10(1):13. doi: 10.20944/preprints202310.1999.v1

[3] Pakseresht S, Celik M, Guler A, et al. Recent advances in all-solid-state lithiumxygen batteries: challenges, strategies, future. Batteries. 2023; 9(7):380. doi: 10.3390/batteries9070380

[4] Tufail MK, Zhai P, Jia M, et al. Design of solid electrolytes with fast ion transport: computation-driven and practical approaches. Energy Material Advances. 2023; 4:15. doi: 10.34133/energymatadv.0015

[5] Yang H, Wu N. Ionic conductivity and ion transport mechanisms of solidtate lithiumon battery electrolytes: A review. Energy Science & Engineering. 2022; 10(5):1643-1671. doi: 10.1002/ese3.1163

[6] Ma M, Zhang M, Jiang B, et al. A review of all-solid-state electrolytes for lithium batteries: high-voltage cathode materials, solid-state electrolytes and electrodelectrolyte interfaces. Materials Chemistry Frontiers. 2023; 7(7): 1268-1297. doi: 10.1039/D2QM01071B

[7] Liu S, Zhao Y, Li X, et al. Solidtate lithium metal batteries with extended cycling enabled by dynamic adaptive solidtate interfaces. Advanced Materials. 2021; 33(12):2008084. doi: 10.1002/adma.202008084

[8] Mouvet F, Villard J, Bolnykh V, et al. Recent advances in first-principles based molecular dynamics. Accounts of Chemical Research. 2022; 55(3):221-230. doi: 10.1021/acs.accounts.1c00503

[9] Van Speybroeck V, Bocus M, Cnudde P, et al. Operando modeling of zeolite-catalyzed reactions using first-principles molecular dynamics simulations. ACS catalysis. 2023; 13(17):11455-11493. doi: 10.1021/acscatal.3c01945

[10] Wang Z, Zeier WG, You F. Toward AI ecosystems for electrolyte and interface engineering in solid-state batteries. Science Advances. 2025; 11(48):eaea0638. doi: 10.1126/sciadv.aea0638

[11] Ren K, Chen Z, Jia B, et al. Strategies and Prospects for the Design of Inorganic Ceramic Electrolyte for High-Performance All-Solid-State Lithium-Ion Batteries. ACS Applied Materials & Interfaces. 2025; 17(40):55643-55665. doi: 10.1021/acsami.5c09137

[12] Xiao Y, Wang Y, Bo S, et al. Understanding interface stability in solid-state batteries. Nature Reviews Materials. 2020; 5(2):105-126. doi: 10.1038/s41578-019-0157-5

[13] Dong Y, Li J, Janiak C, et al. Interfacial design for detection of a few molecules. Chemical Society Reviews. 2023; 52(2):779-794. doi: 10.1039/d2cs00770c

[14] Li X, Liu J, Bai M, et al. An LSTM based method for stage performance degradation early warning with consideration of time-series information. Energy. 2021; 226:120398. doi: 10.1016/j.energy.2021.120398

[15] Dong Y, Li J, Janiak C, et al. Interfacial design for detection of a few molecules. Chemical Society Reviews. 2023; 52(2):779-794. doi: 10.1039/d2cs00770c