Supplementary information for: Non-Arrhenius Li-ion transport and grain-size effects in argyrodite solid electrolytes

Yongliang Ou,[1,2,*] Lena Scholz,[3,†] Sanath Keshav,[3] Yuji Ikeda,[1] Marvin Kraft,[4,5] Sergiy Divinski,[6] Rafael Gómez-Bombarelli,[2] Wolfgang G. Zeier,[4,5] Felix Fritzen,[3,‡] and Blazej Grabowski[1,§]

[1] *Institute for Materials Science, University of Stuttgart, Stuttgart, Germany*
[2] *Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA*
[3] *Institute of Applied Mechanics, University of Stuttgart, Stuttgart, Germany*
[4] *Institute of Inorganic and Analytical Chemistry, University of Münster, Münster, Germany*
[5] *Institute of Energy Materials and Devices (IMD), IMD-4: Helmholtz-Institut Münster, Forschungszentrum Jülich, Münster, Germany*
[6] *Institute of Materials Physics, University of Münster, Münster, Germany*

(Dated: October 21, 2025)

Supplementary Note 1: Details of parameterizing moment tensor potentials

Supplementary Fig. 1 shows an overview of the parameterization scheme for moment tensor potentials, with extended information shown in Supplementary Fig. 2. The concept of quality-level-based active learning was proposed and tested in a previous study [1] for the same material system, i.e., argyrodites. At the beginning, ab initio molecular dynamics simulations are performed on the ideal bulk structure. Considering the high computational cost of ab initio molecular dynamics simulations, the simulation temperature is set relatively high (up to the melting point), and the simulation time is set relatively short (a few picoseconds). The high-temperature molecular dynamics simulations enable a quick exploration of the phase space for a fixed chemical configuration by thermal vibrations, and the resulting trajectories provide the basis for the moment tensor potential training. The initial optimization of the moment tensor potential parameters is essential in calculating the extrapolation grade based on the D-optimality criterion [2]. Since the ab initio molecular dynamics trajectories contain structurally similar configurations, only a small number of configurations from the ab initio molecular dynamics are chosen for the initial moment tensor potential training, and a pretaining step is introduced to include more configurations evaluated as extrapolation, i.e., with an extrapolation grade larger than one. Unlike standard active learning schemes, where a high level of moment tensor potential is directly used for active learning of complex structures, a low level of moment tensor potential is used for standard active learning, and the targeted materials' properties (for example, formation energy or diffusion coefficient) are evaluated at the end of active learning with the obtained low-level moment tensor potentials. If the accuracy of the targeted materials' properties are not satisfied, a higher level of untrained moment tensor potential is used, and the pretaining, the standard active learning, and the evaluation processes are repeated. During the step-by-step training of moment tensor potentials with an increasing level, the configurations selected by active learning and labeled by density-functional theory are accumulated in the training set. The final moment tensor potential is obtained when the accuracy of the targeted material's property is satisfied.

In a previous paper [1] focusing on grain boundaries in Li_6PS_5Cl, only three small grain boundaries are considered, and the full targeted grain boundary structure is used for active learning. This process can be classified as *global-AL*. When the target structure is large, e.g., a large grain boundary structure with more than 30 000 atoms, the density-functional theory calculation of the target structure selected by active learning becomes problematic. As a result, we have incorporated the quality-level-based concept with active learning of the local atomic environments (*local-AL*). The *local-AL* process proposed here has four main steps (Supplementary Fig. 3).

Sampling. The moment tensor potential obtained after pretaining is used to run molecular dynamics simulations of the full target structure. For each snapshot resulting from molecular dynamics simulations, the extrapolation grade calculated based on the local atomic environment is evaluated for each atom, and the maximum local grade of all atoms is focused on. Hyperparameters related to the threshold local grade, i.e., to select the snapshot or stop the molecular dynamics simulations, need to be set. When the maximum local grade of a snapshot is larger than the set threshold, the snapshot is added to and accumulated in a sampling set.

Extraction. In this step, the number of atoms in the simulation cell, which is used for labeling (energy and forces calculation with density-functional theory) later on, is significantly reduced. First, the atom that has the maximum local grade is centered in the simulation cell. The local atomic environment of the centered atom is obtained by cutting out the atoms within a cubic box around the centered atom. The cubic box forms a new simulation cell with periodic boundary conditions. Next, atoms closely positioned due to the periodic boundary conditions are removed until the stoichiometry of atoms in the simulation cell is ensured.

Relaxation. This step aims to avoid unfavorable interactions of atoms at the boundary of the simulation box due to the periodic boundary conditions. To protect the local atomic environment that has been sampled, a

* yongliang.ou@imw.uni-stuttgart.de
† scholz@mib.uni-stuttgart.de
‡ fritzen@mib.uni-stuttgart.de
§ blazej.grabowski@imw.uni-stuttgart.de

protected region in the form of a cubic box is defined around the central atom, and the atoms within the protected region are fixed during the relaxation. The rest of the atoms in the simulation cell are relaxed to minimize the atomic forces. To reduce computational cost, density-functional theory calculations were performed with looser numerical convergence parameters. In this step, only the final relaxed structure enters the next step. Energies and forces obtained during structural relaxation do not enter the training set, maintaining the consistency of the data accuracy of the training set.

Labeling. The energies and forces of the relaxed structure are calculated by density-functional theory with tight convergence parameters. Because only single-point density-functional theory calculations are performed, the computational cost is kept low. The resulting energies and forces of all atoms in the simulation cell with the corresponding relaxed structures are added to and accumulated in the training set.

Updating. The moment tensor potential, either initialized or obtained from the previous round of active learning, was retrained using the updated training set. The resulting moment tensor potential is then used for the next round of *local-AL*, starting from *Sampling*.

With this method, active learning of large and complex structures becomes possible. Since ab initio molecular dynamics simulations are required only for the bulk unit cell, the proposed scheme is highly efficient.

Supplementary Note 2: Consistency of continuum simulation models

A: Methodology

We employed two representations for grain boundaries (GBs) in polycrystalline solid electrolytes: (i) A volume-resolved GB model, in which the GB domain is explicitly meshed, and (ii) A thickness-collapsed model, in which two-dimensional submanifolds represent GBs. The volume-resolved approach is feasible only when the GB width exceeds the numerical mesh size and ensures a sufficient resolution to capture the behavior within the GB domain. Hence, this model is limited to nanograin setups. In contrast, the collapsed GB model assumes a "thin" GB—typical for micrograins—so there is a clear separation of length scales for grain size and GB width. Due to the conceptually different assumptions concerning the GB representation, a straightforward quantitative comparison of their results is non-trivial. Here, we systematically study the respective applicability regimes required to guarantee physically sound results. In particular, a transition grain size will be determined at which one should switch from the volume-resolved to the thickness-collapsed model.

To isolate the intrinsic behavior of the two continuum models and eliminate any geometrical artifacts, we employed a highly idealized microstructure: a two-grain Voronoi tessellation (see Extended Data Fig. 9b). This configuration is perfectly isotropic, features symmetric GB junctions, and keeps the computational cost low. The collapsed approach allows for a parametric variation of the GB width on a single mesh. In contrast, changing the GB width in a resolved model requires rebuilding the three-dimensional voxel image each time to properly identify GB voxels. Therefore, we generated resolved polycrystalline models for a fixed spatial resolution where the GB width is a multiple of the voxel size. Because the considered Voronoi tessellation was fixed, increasing the GB width reduces the effective grain size. Designing a corresponding collapsed model that matches the GB width and GB volume fraction of the resolved model is then straightforward. This was done by rescaling the GB width parameter together with the simulation cell length. Simulations were then carried out for both GB representations independently for a range of bulk and GB diffusivities, D^{bulk} and $D^{\mathrm{GB}} := D^{\mathrm{GB}}_{\parallel} = D^{\mathrm{GB}}_{\perp}$. By repeating the procedure at several mesh resolutions, we can (i) Compare the results of both models across different grain size regimes depending on the inputs $D^{\mathrm{GB}}/D^{\mathrm{bulk}}$, and (ii) Assess how the spatial resolution influences the accuracy of the resolved approach. Results are shown in Extended Data Fig. 9c. From these systematic studies, we aim to (i) Identify the transition grain size at which the modeling strategy should be switched (Supplementary Note 2 B), and (ii) Formulate a practical guideline for selecting a mesh resolution that ensures convergence and accurate results for the resolved approach (Supplementary Note 2 C).

B: Transition grain size and agreement of the models

The true diffusion behavior of a polycrystalline material cannot be measured in detail, so the following considerations are based on two assumptions: (i) The resolved approach yields accurate results in the nanograin regime, and (ii) The collapsed model reliably reproduces diffusion in the micrograin regime. When both models are applied across the entire range of grain sizes, we therefore expect a transition point in between the nanograin and micrograin limits where the agreement is significantly improved compared to the two extremes. This is confirmed by the relative deviation of the macroscopic diffusivities shown in Extended Data Fig. 9c, which shows small deviations for medium-sized grains. Consequently, we define the transition grain size as the grain size at which the two models are most consistent, regardless of the GB-to-bulk diffusion contrast $D^{\mathrm{GB}}/D^{\mathrm{bulk}}$, i.e., we minimize the sum of squared relative deviations. The transition grain size (indicated by the dashed lines in Extended Data Fig. 9c) depends strongly on the spatial resolution of the resolved model, which will be discussed in Supplementary Note 2 C. Upon mesh refinement, the transition size converges to about 100 nm. This value was therefore adopted as a reference for all continuum-model results presented in this study. For the computations that cover the entire grain size range (see Extended Data Fig. 5 and Extended Data Fig. 6d), both the resolved and the collapsed representations were evaluated slightly beyond the identified transition grain size to demonstrate the consistency of the two approaches.

The obtained transition grain size of 100 nm is specific to the considered GB width of 2.5 nm that was assumed

for the given geometry. Nevertheless, the same procedure can be applied to any GB width or microstructure. Converting the transition grain size to a GB volume fraction yields approximately 6.3%, which serves as a GB width-independent indicator of the regime change.

Even in the vicinity of the transition grain size, the two approaches do not achieve perfect agreement (Extended Data Fig. 9c). This residual mismatch stems from inherent differences between the two GB representations: (i) The explicit representation of GB junction domains in the volume-resolved approach versus their implicit treatment in the collapsed approach, (ii) The volume corrections required for the collapsed approach, and (iii) The underlying microstructures considered for the two approaches are not identical but only match key characteristics. Nevertheless, the deviation is modest—typically below 5%—at the transition grain size.

Large deviations occur primarily for blocking GBs, i.e., when the GB diffusivity is lower than the bulk diffusivity ($D^{\mathrm{GB}} < D^{\mathrm{bulk}}$). In this regime, the less-conductive GBs induce pronounced concentration jumps across the GB. These jumps and their effects are especially difficult to capture if not resolved well, as in the following two cases: (i) In volume-resolved simulations of micrograins, a GB may be represented by only a few voxel layers, leading to an under-resolved GB domain and noticeable staircasing effects. (ii) In the collapsed approach for nanograins, the zero-width assumption neglects a substantial GB volume. The impact of resolution in case (i) on concentration and flux fields is shown in Extended Data Fig. 9d for the smallest GB width relative to the grain size. Notably, the volume fraction varies substantially with resolution, even when the GB width remains fixed. This leads to differences in the flux magnitude in the presented results obtained with the volume-resolved approach.

When GB diffusion dominates in the polycrystalline model ($D^{\mathrm{GB}} > D^{\mathrm{bulk}}$), results from the volume-resolved and thickness-collapsed approaches converge for micrograins. In this diffusion regime, GBs do not introduce concentration jumps. Rather, it provides additional diffusion pathways along the GBs. Because the grains are large, the contribution of these pathways to the overall transport is modest, so the agreement between the two models persists even when the GB width is resolved by only a few voxel layers. Therefore, the volume-resolved approach can be employed reliably for grain sizes beyond the estimated transition grain size (100 nm).

C: Resolution requirements for volume-resolved GBs

When a resolved microstructure is generated for a prescribed GB width for different resolutions, the resulting GB volume fraction varies slightly because the voxel-based discretization only approximates the GB domain. This prevents a direct comparison of the macroscopic diffusivities across different resolutions, since the GB volume fraction itself influences the resulting macroscopic diffusivities. Therefore, the convergence behavior of the volume-resolved approach cannot be evaluated unambiguously. Extended Data Fig. 9c shows that the mesh resolution has a noticeable impact on the accuracy of the results, considering the corresponding results from the collapsed approach as a resolution-independent reference.

In the polycrystalline model, the relative GB width decreases with increasing grain size. Consequently, for a given resolution, the number of voxel layers representing a GB varies with grain size. Extended Data Fig. 9c shows that only a few layers were used to resolve a GB for micrograins, and more layers were used for nanograins. The resolution effect is most pronounced for blocking GBs ($D^{\mathrm{GB}} < D^{\mathrm{bulk}}$) combined with large grains, i.e., when the resolved model under-resolves the GB. The convergence of the relative error curves with increasing resolution implies a practical rule: The GB width should be at least ten times the voxel length. For a given geometry, the GB width in the unit cube is first adjusted to satisfy the prescribed ratio between GB width and grain size. The resolution is then selected such that the voxel size is less than one-tenth of the GB width. In contrast, as discussed in Supplementary Note 2 B, results with an enhancing GB setup ($D^{\mathrm{GB}} > D^{\mathrm{bulk}}$) exhibit almost no sensitivity to mesh resolution. In those cases, a coarser discretization is sufficient: A GB width of more than five times the voxel length yields mesh-independent results. In the present study, all volume-resolved simulations were conducted on a $512 \times 512 \times 512$ voxel grid, which exceeds the resolution required to resolve GBs in polycrystalline models with nanograins (grain size less than 100 nm) while avoiding prohibitive computational cost.

[1] Y. Ou, Y. Ikeda, L. Scholz, S. Divinski, F. Fritzen, and B. Grabowski, Atomistic modeling of bulk and grain boundary diffusion in solid electrolyte Li_6PS_5Cl using machine-learning interatomic potentials, Phys. Rev. Mater. **8**, 115407 (2024).

[2] E. Podryabinkin, K. Garifullin, A. Shapeev, and I. Novikov, MLIP-3: Active learning on atomic environments with moment tensor potentials, J. Chem. Phys. **159**, 084112 (2023).

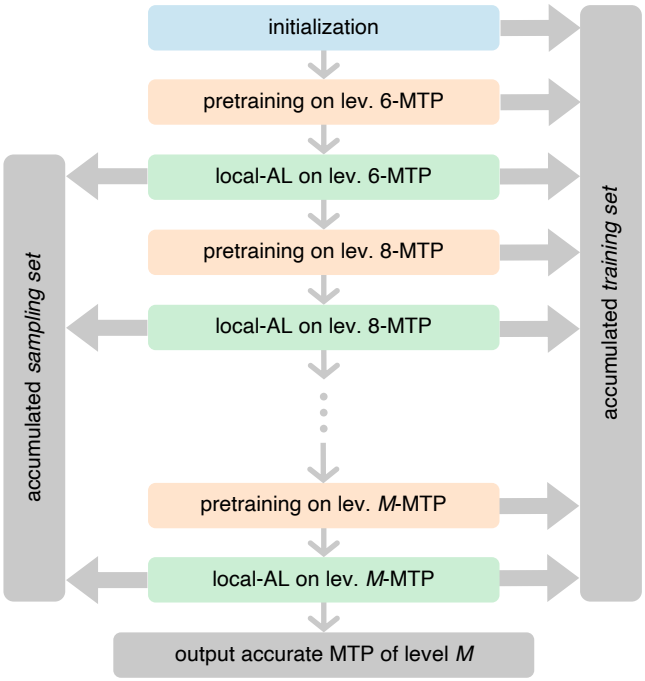

Supplementary Fig. 1. **Overview of the parameterization scheme for moment tensor potentials (MTP).** The "quality-level-based" concept is emphasized with the changing MTP level (lev.).

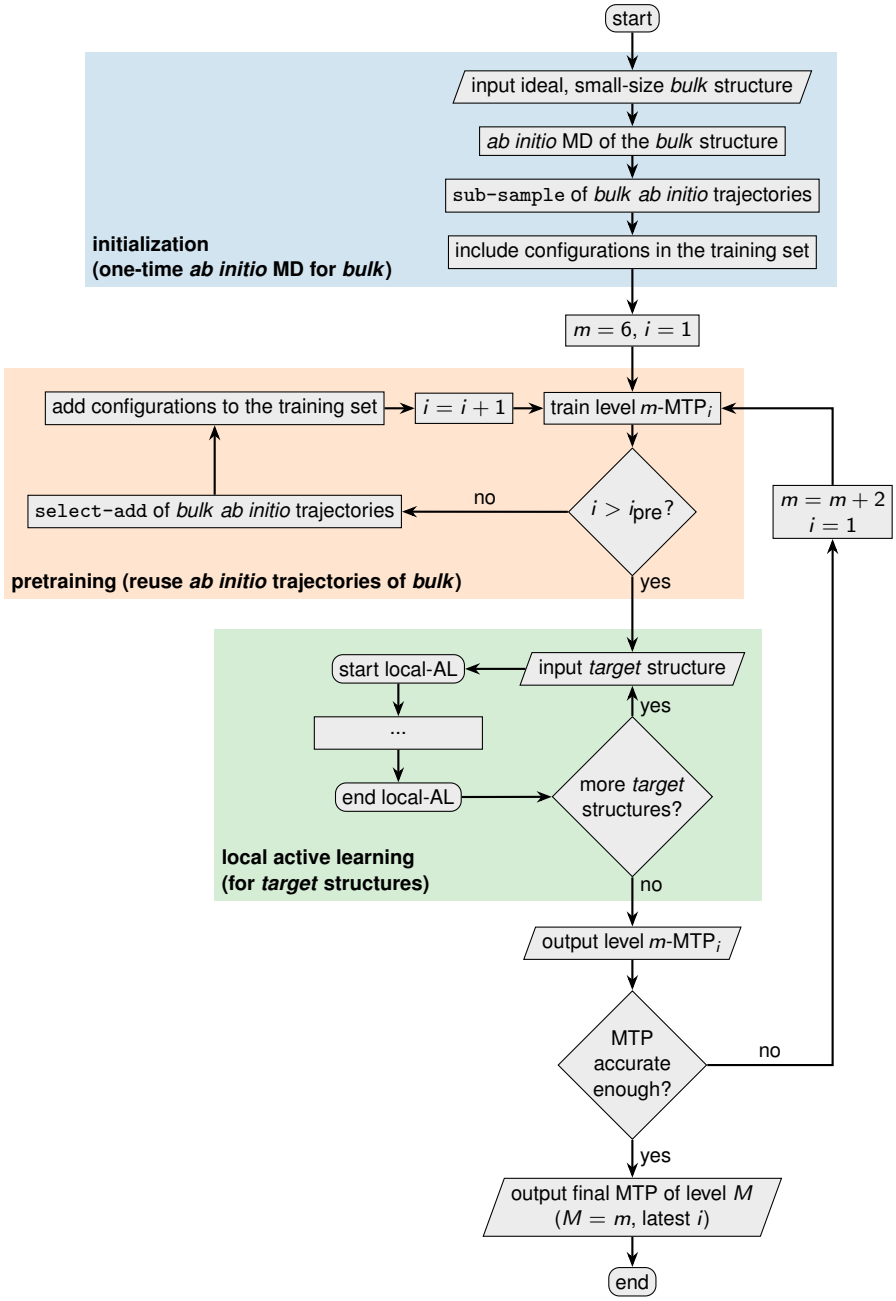

Supplementary Fig. 2. **Complete workflow of the parameterization scheme for moment tensor potentials (MTP).** MD stands for molecular dynamics simulations, while AL stands for active learning.

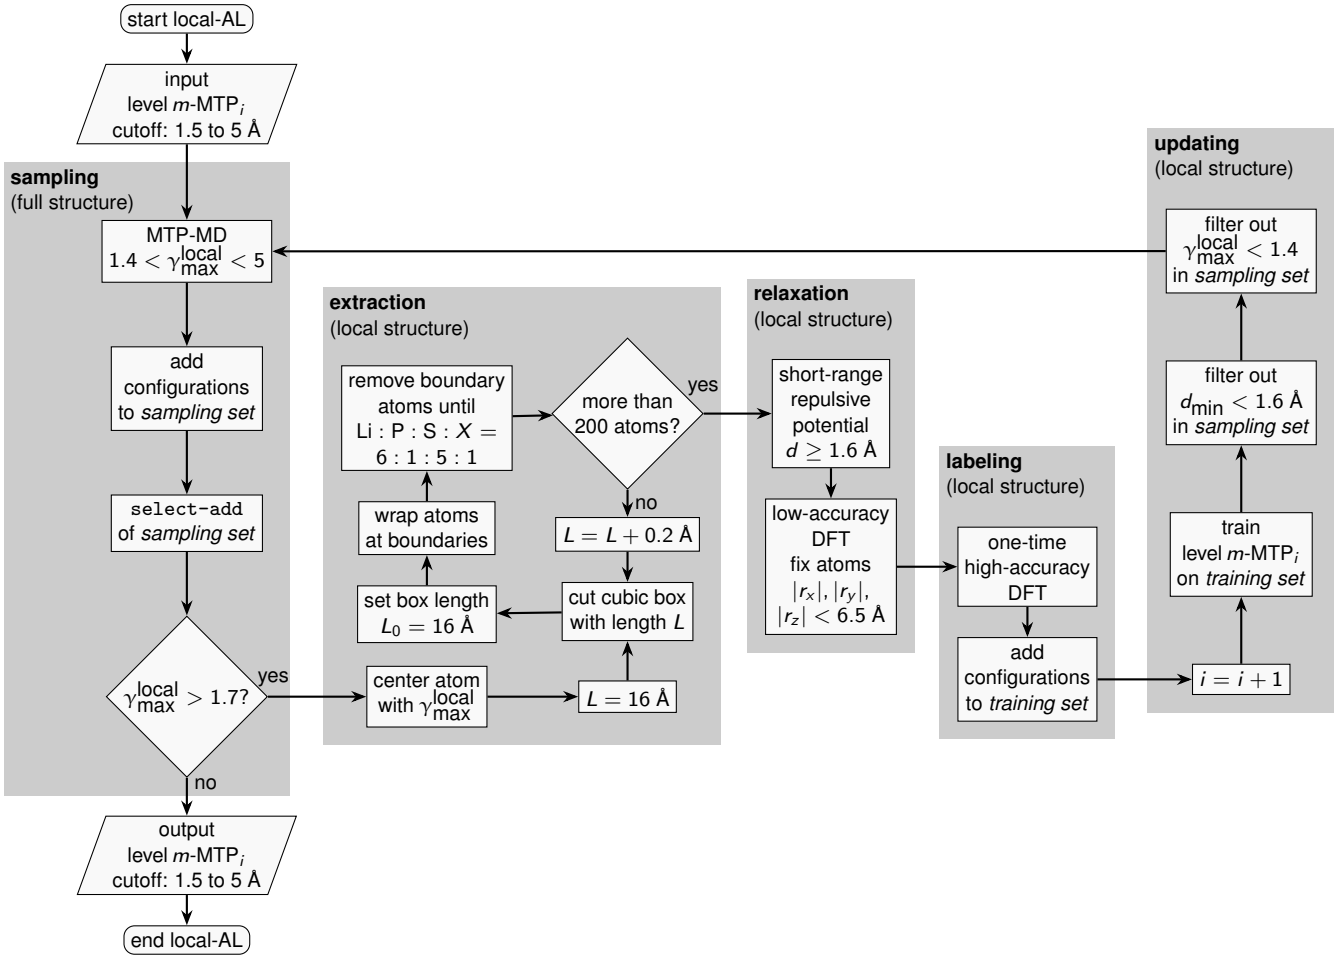

Supplementary Fig. 3. **Details of active learning (AL) for local structures.** Parameters used for Li_6PS_5X with $X \in \{Cl, Br, I\}$ are shown. MTP-MD and DFT refer to molecular dynamics simulations with moment tensor potentials and density-functional theory, respectively. The extrapolation grade and atomic distance are denoted by γ and d, respectively.

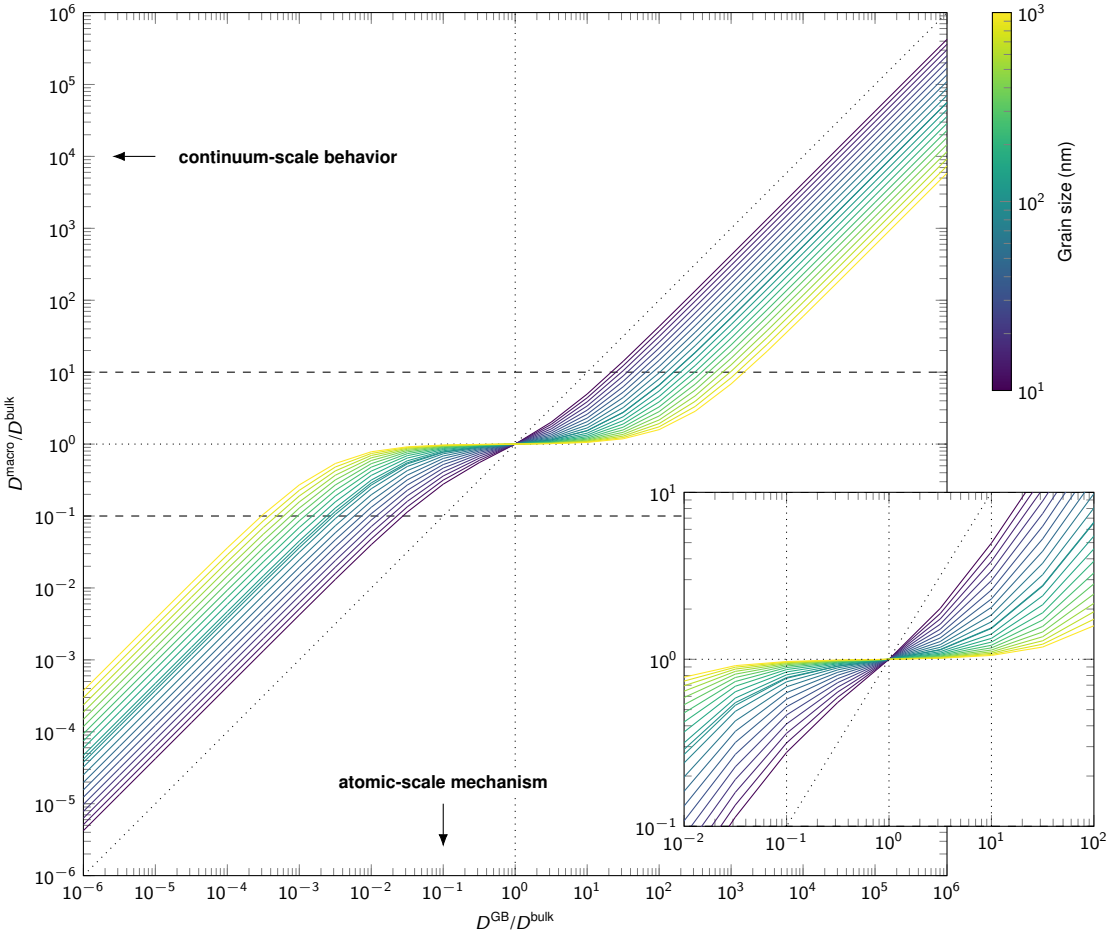

Supplementary Fig. 4. **Precomputed map linking atomic-scale mechanisms to continuum-scale behavior in polycrystalline solid electrolytes.** In a polycrystalline model with isotropic GBs, the x-axis indicates the assumed GB diffusivity relative to the bulk, while the y-axis shows the resulting macroscopic diffusivity of the polycrystal relative to the bulk. Line color represents the average grain size of the polycrystalline model, with the grain boundary width of 2.5 nm. Dash lines indicate a one-order-of-magnitude change in macroscopic diffusivity. The inset shows an enlarged section of the precomputed map.