

General Learning of the Electric Response of Inorganic Materials

Bradley A. A. Martin^{1,*}, Alex M. Ganose², Venkat Kapil³, Tingwei Li⁴, and Keith T. Butler^{1,†}¹Department of Chemistry, *University College London*, London WC1E 6BT, United Kingdom²Department of Chemistry, *Imperial College London*, London W12 0BZ, United Kingdom³Department of Physics and Astronomy, *University College London*, London WC1H 0AH, United Kingdom⁴Department of Materials, *Imperial College London*, London SW7 2AZ, United Kingdom

(Received 10 February 2026; accepted 22 June 2026; published 28 July 2026)

Dielectric response governs how materials interact with electric fields and light. Still, first-principles prediction of polarization $\mathbf{P}$, Born effective charges Z^* , and polarizability α remains too costly for large-scale screening and long-time dynamics. We introduce MACE-Field, a field-aware, $O(3)$ -equivariant interatomic potential that learns a single electric-enthalpy functional $\mathcal{F}(\{\mathbf{R}\}, \mathbf{E})$ and obtains $\mathbf{P}$, Z^* , and α by exact differentiation. A uniform field couples to latent equivariant features inside the MACE backbone, while the scalar energy readout preserves Maxwell reciprocity, the acoustic sum rule, and crystal tensor symmetries by construction. Because this coupling is a plug-in on top of standard MACE, existing energy/force foundation models can be upgraded to become field aware. Benchmarked against semilocal density-functional theory (DFT)/density-functional perturbation theory (DFPT) reference data, a directly trained cross-chemistry ferroelectric model reproduces the same-branch Berry phase and spontaneous polarizations across diverse inorganic crystals. Starting from the multihead foundation model mace-mp-mh-0 and its Open Materials 2024 (OMAT)-Perdew-Burke-Ernzerhof (PBE) head, joint fine-tuning on dielectric, ferroelectric, and replay data yields MACE-Field-MH-0 foundation models, which predict Z^* , α , derived dielectric constants, and cross-chemistry polarization trends with fidelity that captures branch-resolved polarization and spontaneous polarization, while retaining strong force-field accuracy. Further, single-material MACE-Field models and MACE-Field-MH-0 reproduce BaTiO₃ hysteresis loops and α -quartz infrared, Raman, and dielectric spectra from finite-field molecular dynamics, comparable to DFPT. These results show that a simple, physics-informed field coupling can endow atomistic foundation models with transferable dielectric and ferroelectric response, while targeted single-material training remains advantageous for the most quantitative spectroscopic predictions.

DOI: [10.1103/b116-xy8k](https://doi.org/10.1103/b116-xy8k)

Controlling how condensed matter responds to electric fields is central to chemistry, condensed-matter physics, and electronic engineering. It underpins technologies ranging from gate oxides and interconnect dielectrics in microelectronics to capacitors for power conversion, and to piezoelectrics and electro-optic crystals for sensing, imaging, and telecommunications [1–4]. Across these settings, performance is ultimately set by how a material redistributes charge and polarization under an applied field. Microscopically, this response is encoded in the electric polarization $\mathbf{P}$, Born effective charges Z^* , and polarizability tensors α , which connect local bonding to macroscopic observables such as dielectric permittivity, piezoelectric coefficients, and ferroelectric response [5–8]. The same quantities also gov-

ern optical and nonlinear-optical phenomena, including bulk and ferroelectric photovoltaic effects in noncentrosymmetric crystals and hybrid perovskites [9–13], and they shape field effects in catalysis and electrochemistry via screening and interfacial polarization [14–18]. A general and transferable way to predict $\mathbf{P}$, Z^* , and α across chemistries and structures would therefore enable data-driven screening, finite-temperature simulations, and targeted optimization of functional materials.

Achieving this goal is especially challenging in insulating crystals, where polarization is a geometric (Berry) phase and is multivalued on a lattice of polarization quanta [5,6]. Only *differences* in polarization are gauge invariant, so both dataset construction and learning objectives must contend with branch choice.

First-principles approaches based on Density-Functional Perturbation Theory (DFPT) and finite-field Berry-phase methods provide rigorous access to linear and nonlinear responses [7,19,20]. They yield Z^* , electronic polarizabilities, dielectric tensors, and electro-optic coefficients. Still, at high computational cost, DFPT typically requires solving linear-response equations for all occupied bands and perturbations and can be 1–2 orders of magnitude more expensive than

*Contact author: bradley.martin@ucl.ac.uk†Contact author: k.t.butler@ucl.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

a single total-energy calculation. This limits chemical diversity, system size, and accessible timescales, leaving a gap between accurate dielectric theory and the scale required for screening, finite-temperature dynamics, and device-relevant microstructures.

Machine learning (ML) offers a route to bridge this gap, and electric-response learning has developed along several complementary lines. For molecular systems, Christensen *et al.* introduced operator quantum ML (OQML), in which response properties are obtained by applying differential operators (e.g., $\partial/\partial\mathbf{E}$) to learned energy functionals, improving sample efficiency for forces and dipoles and enabling higher derivatives [21]. Symmetry-adapted kernel methods such as symmetry-adapted Gaussian process regression (SA-GPR) (AlphaML and LODE) with λ -smooth overlap of atomic positions (SOAP) kernels established how to learn covariant tensorial quantities, including molecular polarizabilities, condensed-phase dielectric-response tensors, and long-range Coulomb effects [22–24]. Neural models then learned charge and dipole degrees of freedom directly from structure: PhysNet predicts energies/forces together with dipoles and latent partial charges (e.g., on QM9 and MD17/ISO17) [25], while MuML combines local atomic dipoles with a nonlocal partial-charge model to capture both polarization and long-range charge separation (trained on QM7b with transfer tests) [26]. Field-explicit message passing (e.g., FieldSchNet) further showed that equivariance and field inputs enable direct access to response properties and vibrational/optical spectra [27]. In materials science, analogous structure-to-property models learn dielectric constants or tensors from DFPT-labeled crystal datasets [28,29].

Machine-learning interatomic potentials (MLIPs) extend these principles to scalable atomistic simulation by reproducing *ab initio* accuracy at much lower cost [30–32]. $O(3)$ -equivariant backbones such as NequIP/Allegro and MACE provide accurate, differentiable energies and forces across diverse chemistries [32,33]. Several strategies make these models electrically aware. One approach is to augment MLIPs with explicit electronic degrees of freedom with learned charges, dipoles, or Wannier-center proxies, to reconstruct long-range electrostatics and compute dielectric observables (e.g., DeepWannier/deep potential long-range (DPLR)-style pipelines) [34,35]. Another is to learn a field-coupled scalar functional (electric enthalpy) and obtain $\mathbf{P}$, Z^* , and α by exact differentiation, enforcing reciprocity by construction [36]. Complementary approaches recover polarization and Born-effective-charge-like quantities from latent long-range charge fields [37] or incorporate charge response into systematically improvable local bases such as charge-augmented Atomic Cluster Expansion (ACE) variants [38]. Together, these developments suggest that symmetry-aware, differentiable MLIPs can deliver derivative-consistent dielectric response at simulation cost, provided that models capture both nonlocal screening physics and the gauge structure of $\mathbf{P}$ in periodic crystals.

From the perspective of scalable computational materials science, two limitations remain. First, most unified electric-response MLIPs are trained as *single-material* models, restricting transfer across chemistries, space groups, and

bonding motifs. Second, many field-aware architectures diverge from widely used MLIP backbones, limiting the reuse of foundation models pretrained on large energy/force/stress datasets. This matters because high-quality dielectric labels (Berry-phase $\mathbf{P}$, DFPT Z^* , and α) are far scarcer than energetic data; reusing and fine-tuning existing foundation models is therefore particularly valuable.

To address these challenges, we introduce MACE-Field [39], a physics-informed extension of MACE that learns a single *electric-enthalpy* functional $\mathcal{F}(\{\mathbf{R}\}, \mathbf{E})$ across diverse inorganic crystals and obtains dielectric observables by *exact differentiation*:

$$\mathbf{P} = -\Omega^{-1} \frac{\partial \mathcal{F}}{\partial \mathbf{E}}, \quad Z_{\kappa,ij}^* = \Omega \frac{\partial P_i}{\partial u_{\kappa j}}, \quad \alpha_{ij} = \frac{\partial P_i}{\partial E_j}. \quad (1)$$

A uniform field couples *inside* each message-passing layer to latent equivariant features via Clebsch-Gordan tensor products with residual mixing. The readout remains scalar ($L = 0$), and a branch-invariant loss accommodates Berry-phase multivaluedness. Because all quantities derive from one scalar, Maxwell reciprocity and the acoustic sum rule (ASR) follow by construction, yielding a derivative-consistent description of dielectric response suitable for finite-field molecular dynamics (MD) and spectroscopy.

Our two key advances are as follows:

(1) *Foundation-model inheritance.* Field coupling is implemented as a plug-in at the latent irrep level that leaves the standard MACE backbone and readout unchanged, enabling reuse and fine-tuning of MACE foundation weights in a label-scarce regime.

(2) *Cross-chemistry training.* We train and fine-tune across broad chemistries from many Materials Project structures using MP-Ferroelectric ($\sim 2.5\text{k}$ structures, 61 elements) and MP-Dielectric ($\sim 5.3\text{k}$ DFPT entries, 81 elements), rather than a single material.

Together, these advances provide a derivative-consistent, foundation-compatible framework that scales across chemistry and predicts $\mathbf{P}$, Z^* , and α for diverse inorganic solids at a cost suitable for large-scale simulations and screening. Unless explicitly stated otherwise, the quantitative errors reported below are measured against the underlying generalized gradient approximation (GGA)-Perdew-Burke-Ernzerhof (PBE) or GGA-PBE(+U) DFPT/Berry-phase labels used to construct the datasets; they therefore quantify ML agreement with those reference workflows rather than absolute agreement with experiment.

We fine-tune a recent multihead MACE foundation, mace-mp-mh-0, using its OMAT-PBE head as the energy/force/stress prior and jointly supervising on MP-Ferroelectric, MP-Dielectric, and a 10 000-structure OMAT-PBE replay subset associated with the mace-mp-mh-0/1 family, to obtain the field-aware foundation model MACE-Field-MH-0. In Sec. IV, we show that this foundation route predicts BECs, polarizabilities, derived dielectric constants, same-branch Berry-phase polarizations, and spontaneous polarizations across chemistry. We contrast these cross-chemistry models with single-material BaTiO₃/ α -SiO₂ models trained on Allegro-pol trajectories [36] for finite-field hysteresis and IR/Raman/dielectric spectra. The remainder of the paper details the relevant background theory (Sec. I),

architecture and design choices (Sec. II), datasets and training protocols (Sec. III), quantitative results (Sec. IV), and discussion and outlook (Sec. V).

I. THEORY

A. Modern theory of polarization (Berry phase)

In periodic insulators, the macroscopic polarization is not the naive dipole per unit volume, because the position operator is ill-defined under periodic boundary conditions. The *modern theory of polarization* resolves this by expressing the electronic contribution as a Berry phase accumulated by the occupied Bloch bands along loops in the Brillouin zone [5,6]. We briefly summarize the key ingredients here.

We can decompose the total polarization into ionic and electronic parts,

$$\mathbf{P} = \mathbf{P}_{\text{ion}} + \mathbf{P}_{\text{el}}, \quad \mathbf{P}_{\text{ion}} = \frac{-e}{\Omega} \sum_{\kappa} Z_{\kappa} \mathbf{R}_{\kappa}, \quad (2)$$

where e is the electron charge, Ω the cell volume, Z_{κ} the nuclear charge for site κ , and $\mathbf{R}_{\kappa}$ the corresponding position vector. The electronic part of the polarization is

$$\mathbf{P}_{\text{el}} = -\frac{e}{(2\pi)^3} \sum_{n \in \text{occ}} \int_{\text{BZ}} d\mathbf{k} \langle u_{n\mathbf{k}} | i\nabla_{\mathbf{k}} | u_{n\mathbf{k}} \rangle, \quad (3)$$

where $|u_{n\mathbf{k}}\rangle$ are cell-periodic Bloch functions. Only *changes* of $\mathbf{P}$ between adiabatically connected states are gauge invariant; absolute values are defined modulo the *polarization lattice* generated by the quanta

$$\mathbf{Q}_k = \frac{e \mathbf{a}_k}{\Omega}, \quad k = 1, 2, 3, \quad (4)$$

with $\{\mathbf{a}_k\}$ the direct lattice vectors [5,6]. This multivaluedness must be handled explicitly in learning problems (Sec. II).

B. Finite electric fields and DFPT/finite-field formalisms

A uniform static electric field $\mathbf{E}$ couples to polarization through the electric enthalpy (fixed $\mathbf{E}$)

$$\mathcal{F}(\{\mathbf{R}\}, \mathbf{E}) = E_0(\{\mathbf{R}\}) - \Omega \mathbf{E} \cdot \mathbf{P}(\{\mathbf{R}\}), \quad (5)$$

where E_0 is the zero-field energy and Ω the cell volume [19,20]. This functional permits direct optimization of field-polarized states and yields forces and stresses at fixed $\mathbf{E}$.

Within density-functional perturbation theory, linear-response equations for first-order wavefunctions under atomic displacements or homogeneous fields provide access to the complete set of dielectric, piezoelectric, and vibrational properties at zero field without performing finite differences [7]. In this work, we use DFPT quantities as labels and as a reference for assessing our learned model.

C. Macroscopic response tensors as energy derivatives

In the International System of units, the primary response functions used here are defined by exact derivatives of the electric enthalpy:

$$\mathbf{P}(\{\mathbf{R}\}) = -\frac{1}{\Omega} \frac{\partial \mathcal{F}}{\partial \mathbf{E}}, \quad (6)$$

$$Z_{\kappa,ij}^*(\{\mathbf{R}\}) = \Omega \frac{\partial P_i}{\partial u_{\kappa j}} = \frac{\partial F_{\kappa j}}{\partial E_i}, \quad (7)$$

$$\alpha_{ij}(\{\mathbf{R}\}) = \frac{\partial P_i}{\partial E_j}. \quad (8)$$

Here, $u_{\kappa j}$ denotes the Cartesian displacement of sublattice κ , $F_{\kappa j} = -\partial \mathcal{F} / \partial u_{\kappa j}$ the force component, Z^* the Born effective charge tensor (dynamical charge), and α the electronic polarizability per unit volume [40]. Maxwell reciprocity implies symmetry of mixed second derivatives (e.g., $\alpha_{ij} = \alpha_{ji}$). Translational invariance yields the acoustic sum rule $\sum_{\kappa} Z_{\kappa,ij}^* = 0$ for all i, j . Higher-order responses (nonlinear optics, electro-optic) follow from higher derivatives of $\mathcal{F}$, but are not our primary focus here.

D. Ferroelectrics and spontaneous polarization

Ferroelectrics are insulating crystals with a switchable *spontaneous* polarization $\mathbf{P}_s \neq \mathbf{0}$ at zero field, arising from a structural instability that breaks inversion symmetry. Near the paraelectric-to-ferroelectric transition, a soft polar phonon condenses, producing a double-well free-energy landscape as a function of the order parameter (e.g., the unstable mode amplitude). Within a Landau-Devonshire expansion, one writes, schematically, $F(\mathbf{P}) = \frac{1}{2}a(T)\mathbf{P}^2 + \frac{1}{4}b(\mathbf{P}^2)^2 + \dots - \Omega \mathbf{E} \cdot \mathbf{P}$, with $a(T)$ changing sign at the Curie temperature T_C .

The coupling $-\Omega \mathbf{E} \cdot \mathbf{P}$ tilts the double well, leading to coercive fields and hysteresis loops under cyclic fields. Microscopically, the magnitude and direction of $\mathbf{P}_s$ follow from the Berry phase along an adiabatic distortion path connecting the paraelectric reference to the polar ground state, and its coupling to $\mathbf{E}$ is controlled by *mode effective charges*, projections of Z^* onto the polar eigenvectors. Anomalously large Z^* in perovskites signal strong cross-gap hybridization and underpin large dielectric and piezoelectric responses.

E. From linear dielectric response to spectra

The frequency-dependent dielectric tensor can be decomposed into electronic and ionic parts. At high frequency (well above phonon energies and below interband transitions), the response is purely electronic, $\epsilon_{\infty} = \epsilon_0(\mathbf{I} + \alpha)$, which we compute from the electronic polarizability α . At low frequency, the *ionic* contribution from infrared-active phonons adds a resonant term, schematically

$$\Delta\epsilon_{ij}(\omega) \propto \sum_{m \in \text{IR}} \frac{\tilde{Z}_{m,i} \tilde{Z}_{m,j}}{\omega_m^2 - \omega^2 - i\gamma_m \omega}, \quad (9)$$

where ω_m are transverse-optical phonon frequencies, γ_m the linewidths, and $\tilde{\mathbf{Z}}_m$ the mode effective charges obtained by contracting Z^* with phonon eigenvectors. In the static limit, this yields the lattice (ionic) permittivity and, together with ϵ_{∞} , the total static dielectric constant $\epsilon(0)$; the Lyddane-Sachs-Teller relation connects $\epsilon(0)/\epsilon_{\infty}$ to LO-TO splitting. The complex dielectric function is related to the optical conductivity by $\epsilon(\omega) = \epsilon_{\infty} + i\sigma(\omega)/(\epsilon_0\omega)$.

In molecular dynamics, infrared absorption follows from the dipole-dipole (or $\dot{\mathbf{P}}\text{-}\dot{\mathbf{P}}$) autocorrelation, Raman intensities

from polarizability-polarizability correlations [41,42], and the full frequency-dependent complex dielectric function from both. This links time-domain simulations under our learned $\mathcal{F}$ to frequency-domain spectra.

F. Multipole expansion and inspiration for equivariant field coupling

We introduce the multipole expansion as the inspiration for the field-aware MACE model. The interaction of a localized charge distribution $\rho(\mathbf{r})$ with a slowly varying external electrostatic potential $V_{\text{ext}}(\mathbf{r})$ can be expanded in multipoles. The electrostatic interaction energy is

$$E_{\text{int}} = \int d^3r \rho(\mathbf{r}) V_{\text{ext}}(\mathbf{r}). \quad (10)$$

Expanding V_{ext} about a reference point (e.g., the center of the charge distribution)

$$V_{\text{ext}}(\mathbf{r}) = V_0 + r_i \partial_i V_0 + \frac{1}{2} r_i r_j \partial_i \partial_j V_0 + \dots \quad (11)$$

and introducing the Cartesian multipole moments

$$q = \int d^3r \rho(\mathbf{r}), \quad p_i = \int d^3r r_i \rho(\mathbf{r}), \quad (12)$$

$$Q_{ij}^{(T)} = \int d^3r (3r_i r_j - r^2 \delta_{ij}) \rho(\mathbf{r}) \quad (13)$$

yields the familiar coupling to the electric field $\mathbf{E} = -\nabla V_{\text{ext}}$ and its gradients:

$$E_{\text{int}} = q V_0 - \mathbf{p} \cdot \mathbf{E}_0 - \frac{1}{6} Q_{ij}^{(T)} \partial_i \partial_j V_0 + \dots, \quad (14)$$

where subscripts 0 indicate evaluation at the expansion point. For a *uniform* field ($\nabla \mathbf{E} = 0$), only the dipole term survives. In charge-neutral insulating crystals, the total monopole q vanishes per cell; a nonzero dipole density (i.e., macroscopic polarization) arises when inversion symmetry is broken (ferroelectrics) [43].

This multipolar structure has a direct analog in the Atomic Cluster Expansion [44] and its neural realization MACE [33]. The local neighbor density around atom α , $\rho_\alpha(\mathbf{r}) = \sum_j f_{\text{cut}}(r_{\alpha j}) \delta(\mathbf{r} - \mathbf{r}_{\alpha j})$, admits a projection onto radial functions and spherical harmonics,

$$A_{n\ell m}^{(\alpha)} = \sum_j R_n(r_{\alpha j}) Y_{\ell m}(\hat{\mathbf{r}}_{\alpha j}), \quad (15)$$

which are *spherical tensors* of rank ℓ transforming under the $\text{SO}(3)$ irrep ℓ . These coefficients play the role of *generalized multipole moments* of the atomic environment. Scalar, rotation-invariant energies are then formed by contracting such tensors via Clebsch-Gordan coefficients (the spherical-tensor analog of combining multipoles into invariants).

A homogeneous electric field is itself a rank-1 spherical tensor. By the Wigner-Eckart theorem, the symmetry-allowed linear coupling between a latent rank- L feature $T^{(L)}$ (i.e., one irrep block of the learned node features, with multiple channels per irrep) and the field $E^{(1)}$ decomposes as

$$T^{(L)} \otimes E^{(1)} = \bigoplus_{J=|L-1|}^{L+1} [T^{(L)} \otimes E^{(1)}]^{(J)}, \quad (16)$$

with components

$$[T^{(L)} \otimes E^{(1)}]_M^{(J)} = \sum_{m,m'} C_{Lm,1m'}^{JM} T_m^{(L)} E_{m'}^{(1)}, \quad (17)$$

where C are Clebsch-Gordan coefficients. This is exactly the pattern we implement in the architecture through equivariant tensor products followed by irrep-wise mixing [see Fig. 1 and Eq. (26)]. The key idea is that the field perturbs learned latent spherical-tensor features according to angular-momentum selection rules, while the final energy remains a scalar invariant. Keeping the coupling *linear* in $\mathbf{E}$ realizes latent linear response; nonlinear field effects are nevertheless captured by stacking interaction layers and nonlinear readouts, which together generate higher-order dependence on $\mathbf{E}$ in the learned electric enthalpy $\mathcal{F}(\{\mathbf{R}\}, \mathbf{E})$.

G. Physical identities enforced by a single scalar enthalpy

All response tensors in this work are *exact derivatives* of the same twice-differentiable scalar electric-enthalpy functional $\mathcal{F}(\{\mathbf{R}\}, \mathbf{E})$ [Eqs. (6)–(8)]. This immediately enforces key symmetries and sum rules, without any extra penalties, because they are nothing more than properties of mixed partial derivatives and fundamental invariances of $\mathcal{F}$.

Maxwell/reciprocity (symmetry of mixed derivatives). From $\mathbf{P} = -\Omega^{-1} \partial \mathcal{F} / \partial \mathbf{E}$, one has

$$\alpha_{ij} = \frac{\partial P_i}{\partial E_j} = -\frac{1}{\Omega} \frac{\partial^2 \mathcal{F}}{\partial E_i \partial E_j} = -\frac{1}{\Omega} \frac{\partial^2 \mathcal{F}}{\partial E_j \partial E_i} = \alpha_{ji}. \quad (18)$$

That is, the electronic polarizability is *exactly symmetric* at the level of the model. Likewise, with $F_{\kappa j} = -\partial \mathcal{F} / \partial u_{\kappa j}$,

$$Z_{\kappa,ij}^* = \frac{\partial F_{\kappa j}}{\partial E_i} = -\frac{\partial^2 \mathcal{F}}{\partial E_i \partial u_{\kappa j}} = -\frac{\partial^2 \mathcal{F}}{\partial u_{\kappa j} \partial E_i} = \Omega \frac{\partial P_i}{\partial u_{\kappa j}}, \quad (19)$$

which is the standard Maxwell identity linking Z^* to the field derivative of the force and to the position derivative of the polarization.

Acoustic sum rule for Born effective charges. $\mathcal{F}$ is invariant under a rigid translation of all sublattices, $\{\mathbf{R}_\kappa\} \rightarrow \{\mathbf{R}_\kappa + \boldsymbol{\delta}\}$, for any $\boldsymbol{\delta}$ and any uniform $\mathbf{E}$; hence,

$$0 = \frac{\partial}{\partial \boldsymbol{\delta}_j} \mathcal{F}(\{\mathbf{R}_\kappa + \boldsymbol{\delta}\}, \mathbf{E}) = \sum_\kappa \frac{\partial \mathcal{F}}{\partial u_{\kappa j}} = -\sum_\kappa F_{\kappa j}. \quad (20)$$

Differentiating Eq. (20) with respect to E_i and using Eq. (19),

$$\begin{aligned} 0 &= \sum_\kappa \frac{\partial^2 \mathcal{F}}{\partial E_i \partial u_{\kappa j}} \\ &= -\sum_\kappa Z_{\kappa,ij}^* \Rightarrow \sum_\kappa Z_{\kappa,ij}^* = 0 \quad \forall i, j. \end{aligned} \quad (21)$$

That is, the ASR holds *by construction* for a uniform field and a translation-invariant $\mathcal{F}$.

Rotational covariance and crystal point-group constraints. Let $\mathcal{R} \in O(3)$ be a rigid rotation. Our construction ensures that

$$\mathcal{F}(\{\mathcal{R}\mathbf{R}_\kappa\}, \mathcal{R}\mathbf{E}) = \mathcal{F}(\{\mathbf{R}_\kappa\}, \mathbf{E}). \quad (22)$$

That is, the scalar output is rotationally invariant when the atomic configuration and the uniform field are corotated.




FIG. 1. MACE-field architecture. At message-passing layer t , MACE produces learned equivariant node features $h_{\alpha,k,L,M}^{(t)}$, where L denotes the $O(3)$ irrep order and k indexes channels. We occasionally refer to these latent tensors as “multipoles” only heuristically, because they transform like spherical tensors; they are not literal charge-density multipoles. A uniform external field $\mathbf{E}$ (irrep $l = 1$) couples to these features via a fully connected tensor product (Clebsch-Gordan contraction) to form $\Delta h_{\alpha,k,L,M}^{(t)}$; an irrep-wise linear map $W^{(t)}$ and a residual update yield field-aware features $\tilde{h}_{\alpha,k,L,M}^{(t)}$. Scalar components ($L = 0$) are read out at each layer and summed to give a rotationally invariant electric enthalpy $\mathcal{F}$. All dielectric observables are obtained as exact derivatives of this single scalar: polarization $\mathbf{P} = -\Omega^{-1} \partial \mathcal{F} / \partial \mathbf{E}$, Born effective charges $Z_{\alpha,ij}^* = \Omega \partial P_i / \partial R_{\alpha j} / e = \partial F_{\alpha j} / \partial E_i$, and electronic polarizability $\alpha_{ij} = \partial P_i / \partial E_j$. Hidden layers remain $O(3)$ equivariant; only the final readout in layer $t = T$ is strictly invariant.

Taking derivatives at $\mathbf{E} = \mathbf{0}$ yields the correct tensorial transformation laws:

$$\mathbf{P}(\{\mathcal{R}\mathbf{R}\}) = \mathcal{R} \mathbf{P}(\{\mathbf{R}\}), \quad (23)$$

$$\boldsymbol{\alpha}(\{\mathcal{R}\mathbf{R}\}) = \mathcal{R} \boldsymbol{\alpha}(\{\mathbf{R}\}) \mathcal{R}^{\top}, \quad (24)$$

$$Z_{\pi(\kappa)}^*(\{\mathcal{R}\mathbf{R}\}) = \mathcal{R} Z_{\kappa}^*(\{\mathbf{R}\}) \mathcal{R}^{\top}, \quad (25)$$

where π permutes sublattices according to the rotation.

Thermodynamic sign/definiteness. At fixed atomic positions, the second field derivative of $\mathcal{F}$ gives $-\Omega \boldsymbol{\alpha} = \partial^2 \mathcal{F} / \partial \mathbf{E} \partial \mathbf{E}$. For stable insulating states at zero field, thermodynamics implies $\mathbf{v}^{\top} \boldsymbol{\alpha} \mathbf{v} \geq 0$ for any vector $\mathbf{v}$ (passive linear response). While symmetry [Eq. (18)] is exact in our construction, strict positive semidefiniteness is a property of the true material response and is *approached* as the learned $\mathcal{F}$ converges; tiny negative eigenvalues may occur at inference.

II. METHODS

A. Overview

We learn a single scalar electric-enthalpy functional $\mathcal{F}(\{\mathbf{R}\}, \mathbf{E})$ that is differentiable with respect to both atomic positions and a *global, uniform* electric field. The MACE-Field network is a field-aware variant of MACE: It injects the field into latent $O(3)$ -equivariant features at each interaction layer while keeping the final readout scalar ($L = 0$). polarization, BECs, and polarizability are then obtained by *exact* automatic differentiation of $\mathcal{F}$ [Eqs. (6)–(8)], ensuring derivative consistency (Maxwell symmetries and acoustic sum rules) by construction. Conceptually, this makes MACE-Field a physics-informed, end-to-end differentiable model: All re-

sponse properties are mixed derivatives of a single learned potential, thereby avoiding finite-difference noise and naturally fitting into automatic adjoint-based optimization and downstream simulation workflows.

B. Architecture and field coupling

A schematic of the MACE-Field architecture is shown in Fig. 1. We start from a standard MACE backbone with T layers, K hidden irreps, and message interaction and product blocks as in Ref. [33]. Here, “latent features” refer to the learned per-atom node features in the equivariant network. The index L denotes the spherical-harmonic/ $O(3)$ irrep order, while K counts channels within each irrep.

Let $h_{\alpha,k,L,M}^{(t)}$ denote the latent features at layer t (atom α , channel k , irrep L , component M). In the current architecture, the field input is a single, spatially uniform global vector shared across the entire graph. It is represented as an $l = 1$ vector (odd parity $p = -1$) irrep feature E_{1m} and coupled at each layer $t < T$ via a FullyConnectedTensorProduct followed by a linear equivariant mixing (residual update):

$$\begin{aligned} \Delta h_{\alpha,k,L,M}^{(t)} &= \sum_{l_1 m_1, l_2 m_2} \mathcal{C}_{l_1 m_1, l_2 m_2}^{LM} h_{\alpha, k l_1 m_1}^{(t)} E_{l_2 m_2}, \\ \tilde{h}_{\alpha,k,L,M}^{(t)} &= h_{\alpha,k,L,M}^{(t)} - \sum_{\tilde{k}} W_{\tilde{k}kL}^{(t)} \Delta h_{\alpha, \tilde{k}, L, M}^{(t)}, \end{aligned} \quad (26)$$

with Clebsch-Gordan coefficients $\mathcal{C}$ and learned weights $W^{(t)}$. Message passing and products operate on $\tilde{h}^{(t)}$ thereafter. At the final layer $t = T$, following MACE, we retain only $L = 0$ scalars and employ a nonlinear readout to produce per-atom contributions, which are summed over atoms and layers to

yield $\mathcal{F}$. The field is identical for all atoms in a graph and does *not* depend on absolute positions, preserving translation invariance. At zero field $E = 0$, the residual update is identically zero, so we fall back to the normal MACE behavior. Because the field-coupling layers operate purely on latent features, the underlying ACE descriptors and scalar readout are unchanged, making MACE-Field a true plug-in: any existing MACE checkpoint can be upgraded to a field-aware model by inserting these layers and fine-tuning on response data. The same design choice also defines the scope of the method: it targets bulk-periodic response under spatially uniform bias and does not explicitly resolve site-dependent internal fields at defects, surfaces, interfaces, or domain walls.

C. Differentiable observables and units

1. Target properties

We compute observables by automatic differentiation (as implemented in PyTorch’s autograd functionality) on the *interaction* part of the energy (the atomic baselines are constant):

$$\mathbf{P} = -\frac{1}{\Omega} \frac{\partial \mathcal{F}}{\partial \mathbf{E}}, \quad Z_{\kappa,ij}^* = \frac{\partial(\Omega P_i)}{\partial u_{\kappa j}}, \quad \alpha_{ij} = \frac{\partial P_i}{\partial E_j}. \quad (27)$$

Internally, $\mathbf{P}$ is stored in $e/\text{\AA}^2$, $Z_{\kappa,ij}^*$ in units of e , and α_{ij} in ϵ_0 (vacuum permittivity). We report polarizations in units of $\mu\text{C cm}^{-2}$ where convenient.

All training losses use consistent internal units. To obtain Z^* and α , we differentiate $\Omega \mathbf{P}$ with respect to atomic positions and $\mathbf{P}$ with respect to $\mathbf{E}$, respectively, using nested automatic differentiation. During training, we retain and create graphs for higher-order derivatives only when those targets are present in the batch. We compute losses on the interaction energy (the atomic baselines are constant), so that derivatives reflect the learned $\mathcal{F}$.

2. Derived properties: Dielectric constants

We can relate the polarizability to the *relative* high-frequency dielectric tensor by

$$\epsilon_{\infty}^{(r)} = \mathbf{I} + \frac{\alpha}{\epsilon_0}. \quad (28)$$

Likewise, we obtain the polycrystalline dielectric constant $\epsilon_{\infty}^{(\text{poly})}$ at long wavelengths (ionic contributions vanish at high frequencies) from the eigenvalues (λ_i , $i \in \{1, 2, 3\}$) of the high-frequency dielectric tensor [45],

$$\epsilon_{\infty}^{(\text{poly})} \equiv \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}. \quad (29)$$

We can then obtain an estimate of the refractive index, n , at optical frequencies and far from resonance effects using

$$n = \sqrt{\epsilon_{\infty}^{(\text{poly})}}. \quad (30)$$

For a given relaxed structure, the real-space Hessian is

$$H_{\kappa\alpha,\kappa'\beta} = \frac{\partial^2 \mathcal{F}}{\partial u_{\kappa\alpha} \partial u_{\kappa'\beta}}, \quad (31)$$

where κ, κ' index atoms in the unit cell and $\alpha, \beta \in \{x, y, z\}$ are Cartesian components of the displacements $u_{\kappa\alpha}$. From this, we

construct the Γ -point dynamical matrix

$$D_{\kappa\alpha,\kappa'\beta} = \frac{1}{\sqrt{M_{\kappa} M_{\kappa'}}} H_{\kappa\alpha,\kappa'\beta}, \quad (32)$$

with M_{κ} the mass of atom κ . Diagonalizing D yields phonon eigenfrequencies and eigenvectors,

$$\sum_{\kappa'\beta} D_{\kappa\alpha,\kappa'\beta} e_{\kappa'\beta}^{(m)} = \omega_m^2 e_{\kappa\alpha}^{(m)}, \quad (33)$$

where m labels the (zone-center) normal modes, ω_m are the phonon frequencies, and the eigenvectors $e_{\kappa\alpha}^{(m)}$ are chosen orthonormal in the mass-weighted sense,

$$\sum_{\kappa\alpha} e_{\kappa\alpha}^{(m)} e_{\kappa\alpha}^{(m')} = \delta_{mm'}. \quad (34)$$

Given the Born effective charge tensors $Z_{\kappa,i\alpha}^*$ predicted by MACE-Field (with Cartesian indices $i, \alpha \in \{x, y, z\}$), we can form the mode effective charges by projecting Z^* onto the phonon eigenvectors,

$$\tilde{Z}_{m,i} = \sum_{\kappa\alpha} \frac{Z_{\kappa,i\alpha}^*}{\sqrt{M_{\kappa}}} e_{\kappa\alpha}^{(m)}. \quad (35)$$

In density-functional perturbation theory, the lattice (ionic) contribution to the dielectric tensor in the static limit $\omega \rightarrow 0$ [cf. Eq. (9)] is then given by

$$\epsilon_{ij}^{\text{ion}} = \frac{4\pi}{\Omega} \sum_{m \in \text{IR}} \frac{\tilde{Z}_{m,i} \tilde{Z}_{m,j}}{\omega_m^2}, \quad (36)$$

where Ω is the unit-cell volume and the sum runs over the IR-active, nonacoustic modes at Γ . In practice, we discard the three translational acoustic modes (for which $\omega_m \rightarrow 0$ and $\tilde{Z}_{m,i} \rightarrow 0$) and any modes with numerically ill-defined ω_m arising from slight residual stresses.

Finally, combining the electronic and ionic contributions yields the full static dielectric tensor

$$\epsilon_0 = \epsilon_{\infty} + \epsilon_{\text{ion}}. \quad (37)$$

D. Branch-invariant polarization supervision

polarization is multivalued modulo the polarization lattice, $\mathbf{Q}_k = e \mathbf{a}_k / \Omega$ [Eq. (4)]. As such, we follow the approach taken by Allegro-pol [36]: given a reference $\mathbf{P}_{\text{ref}}$ and prediction $\mathbf{P}_{\text{pred}}$, we compare the *folded* difference $\Delta \mathbf{P}_{\text{fold}}$:

$$\Delta \mathbf{P} = \mathbf{P}_{\text{pred}} - \mathbf{P}_{\text{ref}},$$

$$B = [\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3]^T, \quad Q_{\text{pol}} = B/|\Omega|,$$

$$Q_{\text{pol}}^T \mathbf{c}^T = \Delta \mathbf{P}^T, \quad \mathbf{n}^* = \arg \min_{\mathbf{n} \in \mathbb{Z}^3} \|(\mathbf{c} - \mathbf{n}) Q_{\text{pol}}\|_2,$$

$$\Delta \mathbf{P}_{\text{fold}} = \mathbf{c}_{\text{fold}} Q_{\text{pol}}, \quad \mathbf{c}_{\text{fold}} = \mathbf{c} - \mathbf{n}^*. \quad (38)$$

Here, B is the cell matrix with lattice vectors stored as rows, so in the lattice basis the polarization quanta become $Q_{\text{pol}} = B/|\Omega|$. The fractional coefficients $\mathbf{c}$ are therefore obtained by solving against Q_{pol}^T . The fold is then defined by the *exact* closest-vector problem on the polarization lattice rather than by a componentwise modulo. This closest-image construction works for general, nonorthogonal cells and makes the polarization loss branch invariant while leaving $\mathbf{P}$ itself defined by the conservative derivative in Eq. (27).

E. Loss function

Following the UniversalLoss used by MACE for energy, forces, and stress, our field-related per-task losses are

$$\mathcal{L}_P = \frac{1}{W_P} \sum_{g=1}^{N_b} w_g^P d_{\text{pol}}^2(\mathbf{c}_g^{\text{fold}}, Q_{\text{pol},g}), \quad (39)$$

$$\mathcal{L}_Z = \frac{1}{W_Z} \sum_{g=1}^{N_b} \frac{1}{n_g} \sum_{a \in g} w_{ga}^Z \overline{\mathcal{H}}_{\delta_Z}(\hat{Z}_{ga}^* - Z_{ga}^*), \quad (40)$$

$$\mathcal{L}_\alpha = \frac{1}{W_\alpha} \sum_{g=1}^{N_b} w_g^\alpha \overline{\mathcal{H}}_{\delta_\alpha}(\hat{\alpha}_g - \alpha_g), \quad (41)$$

where

$$d_{\text{pol}}^2(\mathbf{c}; Q) = \mathbf{c} M(Q) \mathbf{c}^T, \quad (42)$$

$$M(Q) = \frac{3 Q Q^T}{\text{tr}(Q Q^T)}$$

is the normalized, general cell-aware polarization loss, $\mathcal{H}_\delta$ is the Huber loss (with Huber delta δ), $\overline{\mathcal{H}}$ denotes the average of $\mathcal{H}$ over vector/tensor components, $Z_{ga}^* \in \mathbb{R}^{3 \times 3}$ is the BEC tensor on atom a , and α_g is the polarizability tensor. A batch g has size N_b and contains n_g atoms.

Normalization constants $W_T = \sum_g w_g^T$ (or $\sum_{g,a} w_{ga}^T$ for per-atom terms) prevent batch-size/weight drift. The w_g configuration weights can be used to mask per-task losses for heterogeneous labels by setting the contribution of a particular property to zero if its label is absent for a given data point. This is particularly useful for fine-tuning existing foundation models using a replay set of energies, forces, and stresses when field-related labels are absent.

The total loss is a non-negative weighted sum of per-task losses:

$$\mathcal{L}_{\text{tot}} = \lambda_E \mathcal{L}_E + \lambda_F \mathcal{L}_F + \lambda_\sigma \mathcal{L}_\sigma + \lambda_P \mathcal{L}_P + \lambda_Z \mathcal{L}_Z + \lambda_\alpha \mathcal{L}_\alpha, \quad (43)$$

where the energy, forces, and stress losses are the same as in MACE [33]. Unless otherwise stated, we use fixed weights λ during training (details for the different trained MACE-Field models are provided in the Supplemental Material [46]).

F. Implementation and training protocol

Models are implemented in PyTorch with cuEquivariance irreps (e3nn option is available) and tensor products, and built atop a public MACE codebase.

Exact hyperparameters, seeds, split definitions, and details on public data/code regeneration are listed in the Supplemental Material [46] and the ‘‘Data availability’’ statement.

We train with Adam (with decoupled weight decay), a cosine learning rate schedule with warm-up, and gradient clipping. We retain and create graphs for autograd of $\Omega \mathbf{P}$ with respect to $\mathbf{R}$ and $\mathbf{E}$ when BECs and α are requested (higher-order gradients). We use early stopping on validation $\mathbf{P}$ (ferroelectric split) and α/Z^* (dielectric split). All models are trained with distributed data parallel across graphics processing units (GPUs); representative dtypes, experiment-specific settings, and run-time numbers are listed in the Supplemental Material [46] and repository configuration files.

In practice, force-only machine-learning molecular dynamics (MLMD) with MACE-Field is expected to remain close to an architecture-matched ordinary MACE model, because the time-integration loop only queries energies and forces. The main extra overhead enters during training and whenever $\mathbf{P}$, Z^* , or α are explicitly requested, because these require additional nested-autograd evaluations; in the large-scale atomic/molecular massively parallel simulator (LAMMPS) machine-learning interatomic-potential (ML-IAP) workflows used here, these response quantities may therefore be back-filled in a separate postprocessing pass.

G. Finite-field MD protocols

We perform finite-field ML molecular dynamics using the learned $\mathcal{F}$, either directly in the atomic simulation environment (ASE) [47] or through LAMMPS/ML-IAP exports for the longer production trajectories.

BaTiO₃ hysteresis. We run field-driven dynamics and quasi-static ionic relaxations in a small bulk-periodic supercell under a spatially uniform field applied along the polar axis. In the LAMMPS/ML-IAP workflow, we first relax a 135-atom periodic cell at fixed DFT lattice vectors and then equilibrate athermally at 0 K; the driven segment is subsequently propagated at fixed cell with nve dynamics plus viscous damping. Because the cell is periodic and contains no explicit defects or domain walls, this setup probes intrinsic, spatially homogeneous switching rather than nucleation-and-growth kinetics.

α -SiO₂ IR/Raman and dielectric function. At 300 K, we run fixed-cell NVT trajectories in a 72-atom periodic α -quartz supercell, using 2 fs time steps and 200 ps production windows after equilibration. In the LAMMPS/ML-IAP workflow, energies and forces are evaluated during time integration and the derivative observables entering the spectra are reconstructed afterward on the stored frames. IR absorption is then obtained from the $\dot{\mathbf{P}} \cdot \dot{\mathbf{P}}$ autocorrelation, and Raman from the $\alpha \cdot \alpha$ autocorrelation, both Fourier transformed with a Hann window and modest Gaussian broadening. The complex dielectric $\epsilon(\omega)$ follows from $\epsilon_\infty^{(r)}$ and the same current-current (equivalently $\dot{\mathbf{P}} \cdot \dot{\mathbf{P}}$) response.

H. Comparison to other models

1. Allegro-pol (unified electric enthalpy)

Our work is closest in spirit to Allegro-pol, which also learns a single *electric enthalpy* and obtains $\mathbf{P}$, Z^* , and α by exact differentiation, enforcing reciprocity and the acoustic sum rule by construction [36]. Both approaches handle Berry-phase multivaluedness using a branch-invariant (minimum-image) loss on $\Delta \mathbf{P}$, though Allegro-pol’s implementation is limited to orthogonal cells. The main differences are architectural and data scientific. Allegro-pol injects the uniform field as a rank-1 input feature before the local equivariant layers, whereas MACE-Field couples $\mathbf{E}$ *inside* each message-passing layer as a perturbation to latent irreps via Clebsch-Gordan tensor products and residual mixing, preserving the standard MACE backbone and enabling foundation-weight reuse. Allegro-pol is demonstrated primarily in a single-material regime (e.g., α -quartz, BaTiO₃) with labels constructed from small finite-field calculations around

nominal zero field [36]; here, we focus on cross-chemistry training using DFPT/Berry-phase labels across thousands of Materials Project structures, enabling transfer across compositions and space groups.

2. Response-operator and symmetry-adapted learning

Complementary literature learns an electric response *without* a force field by enforcing derivative structure and covariance at the property level. OQML introduces response operators directly in the regression, so that quantities such as forces and dipoles arise as derivatives of a learned energy functional with improved data efficiency [21]. SA-GPR and λ -SOAP kernels provide a general framework to learn *tensorial* response with the correct $O(3)$ covariance, with demonstrations ranging from molecular and condensed-phase electrical response (e.g., water clusters/liquid snapshots) [22]. Similarly, the AlphaML work of Wilkins *et al.* uses symmetry-adapted linear-response coupled-cluster singles and doubles (LR-CCSD) *static polarizability tensors* for 7211 QM7b molecules and transfers to a chemically diverse showcase set [23]. These approaches are valuable baselines for screening and for building physical inductive biases (covariant tensors, operator consistency), but they do not provide a thermodynamically consistent electric-enthalpy functional for finite-field MD in periodic solids, nor do they address Berry-phase branch structure directly.

3. Long-range physics and electrostatics in atomistic ML

A second comparison axis concerns how models represent nonlocal electrostatics and screening. Grisafi and Ceriotti proposed a general scheme to incorporate *long-range physics* by building nonlocal representations (including forms with electrostatic-potential-like asymptotics) and remapping them into local, $O(3)$ -equivariant features; they demonstrated the approach on electrostatic energies, charged dimers, and the electronic dielectric response of liquid water [24]. More MLIP-specific routes (1) learn latent charge fields and apply Ewald-like summations [e.g., latent Ewald summation (LES)] [37], (2) promote charges to variational degrees of freedom within systematic bases (e.g., ACE+ Q/charge-augmented ACE) [38], or (3) add explicit long-range electrostatics to neural potentials (e.g., DPLR) [35]. Compared to these, MACE-Field targets dielectric response by *direct supervision* of $\mathbf{P}$, Z^* , and α while retaining strict derivative consistency through a single learned enthalpy; long-range/screening physics is then represented in the learned field-coupled functional rather than reconstructed via separate electrostatic postprocessing.

4. Empirical ferroelectric force fields and local-field limitations

Reactive and empirical ferroelectric force fields such as ReaxFF have already been used to study large-scale phenomena including surface chemistry, oxygen-vacancy effects, thickness dependence, and defect-assisted switching in BaTiO_3 and related systems [49,50]. These methods are attractive for mesoscale simulations because they can be deployed in large supercells and naturally access domain formation and defect physics. MACE-Field has a different emphasis: it learns a single differentiable electric enthalpy

and derives $\mathbf{P}$, Z^* , and α as exact mixed derivatives, which is advantageous for derivative consistency, finite-field MD, and dielectric analysis under spatially uniform bias. The present architecture, however, couples a single homogeneous external field to the entire graph and does not solve for self-consistent, spatially varying internal fields. It should therefore be viewed primarily as a model for bulk-periodic response, homogeneous switching in bulk cells, and spectroscopy, rather than a complete treatment of charged defects, surfaces, interfaces, or domain walls where explicit long-range electrostatics and local depolarizing fields are essential.

5. Molecular dipole models and charge-dipole decompositions

For finite systems, accurate dipoles can be obtained from learned decompositions into partial charges and local atomic dipoles. Veit *et al.* (MuML) combine a nonlocal charge model with a local atomic-dipole model to capture both long-range charge separation and short-range polarization, yielding accurate molecular dipoles on datasets such as QM7b and transfer tests beyond the training distribution [26]. Such models provide important points of comparison for the *representation* of polarization. Still, periodic insulating crystals introduce additional constraints (Berry-phase gauge/branch structure and cell-volume factors) that motivate our branch-invariant training objective and enthalpy-based formulation.

6. Direct tensor predictors (screening models; no force field)

Equivariant property predictors such as AnisoNet [29] and graph-based dielectric models that learn scalar ϵ or full dielectric tensors from crystal structure (e.g., the oxide-dielectric ML models of Morita *et al.* [28] and MACE- μ [51]) provide fast screening baselines for dielectric response. However, these approaches predict response *directly* without an interatomic potential: they do not yield forces, do not define a thermodynamically consistent electric-enthalpy functional, and therefore cannot be used directly for finite-field molecular dynamics, hysteresis simulations, or spectroscopic workflows that require end-to-end derivative consistency.

III. DATASETS

Our training data combine two cross-chemistry datasets that provide wide elemental coverage across the periodic table with two single-material MD trajectory sets. Figure 2 summarizes the elemental coverage of the two cross-chemistry datasets that underpin the field-aware foundation model. Unless explicitly stated otherwise, all errors reported below are with respect to the corresponding DFT/DFPT labels. Because the cross-chemistry datasets combine distinct semilocal workflows, the resulting error floor reflects both ML approximation and the intrinsic limitations and inconsistencies of the reference calculations.

A. MP-Dielectric: BECs and polarizabilities of Materials Project dielectrics

For dielectric response, we assemble a broad chemistry of ~ 6000 insulating materials from the Materials Project with DFPT-computed Born effective charges and electronic




FIG. 2. Elemental coverage of the cross-chemistry datasets used in this work. (a) Smidt *et al.* ferroelectric distortion-path set [48] covering 61 elements. (b) MP-Dielectric dataset (DFPT BECs and electronic polarizabilities) covering 81 elements. Color encodes the per-dataset normalized frequency, that is, the fraction of structures containing each element (see color bars); gray indicates no examples.

polarizabilities (equivalently, the electronic dielectric tensor ϵ_∞). The elemental coverage and per-element frequencies are shown in Fig. 2(b). This dataset spans 81 elements and includes oxides, halides, chalcogenides, and mixed-anion chemistries.

We assemble this dataset from the public Materials Project application programming interface (API) that queries dielectric-task provenance, extracts the DFPT tensors, and writes the filtered train/validation/test extxyz files used here. All DFPT calculations in this dataset were carried out within the GGA-PBE functional family [52]. As such, absolute values may differ from higher-level functionals and, as we discuss later, a small subset of entries even violate basic constraints such as the acoustic sum rule. We use $\alpha = \partial \mathbf{P} / \partial \mathbf{E}$ directly as labels and supervise Z^* componentwise. The DFPT forces in this dataset do not correlate to the dielectric response; therefore, we *do not* train on forces or stress for this set to avoid encoding artefacts. Splits are defined at the material level (unique MP identifiers), so polymorphs of a given material may appear in different splits. Within this head, the reference labels are therefore reasonably consistent across the exchange-correlation family and basis-set class (plane-wave VASP DFPT), even though convergence settings vary across MP workflows. Together with MP-Ferroelectric [Sec. III B and Fig. 2(a)], this dataset covers more than 80 elements and thousands of distinct chemistries, providing a realistic test of whether a single enthalpy model can generalize dielectric response beyond single-material manifolds.

B. MP-Ferroelectric: Polarizations and distortion paths of Materials Project ferroelectrics

To go beyond training on single-material examples, we use MACE-Field to predict the polarization of a wide set of ferroelectric materials and their corresponding polarization branches, so that we can calculate their spontaneous polarization. We use the automatically curated first-principles MP-Ferroelectric database of Smidt *et al.* [48], obtained using the MPContribs API [53–57]. The elemental coverage of this dataset, shown in Fig. 2(a), spans 61 elements and is enriched in oxide and perovskitelike chemistries that are typical ferroelectric candidates.

The MP-Ferroelectric workflow identifies symmetry-related nonpolar-to-polar pairs in the Materials Project. It computes Berry-phase polarizations along an adiabatic distortion path connecting the two end states. In this workflow, the electronic polarization is evaluated with the Berry-phase formalism (VASP implementation) and the ionic part is added from point charges; spin-polarized GGA-PBE(+U) calculations and path-based validation (smoothness and insulating branch) are used to recover a unique spontaneous $\mathbf{P}_s$ for each candidate [58].

From this database, we select ~ 250 materials for which the nonpolar and polar end points are available and insulating along the path, and include the eight evenly spaced interpolates (fixed cell, linear in fractional coordinates) between end points, yielding ~ 2500 structures in this ferroelectric dataset. For each structure in the training split, we train on the DFT total energy, forces, and the Berry-phase polarization. The Berry-phase polarizations are *prefolded* onto the polar lattice branch as in Ref. [48]. These folded values supervise the polarization term in our loss (Sec. II), while the model itself always predicts $\mathbf{P}$ as a derivative of the learned enthalpy. This dataset, therefore, probes a distinct semilocal workflow from MP-Dielectric: MP-Ferroelectric uses path-validated Berry-phase GGA-PBE(+U) calculations, whereas MP-Dielectric uses a separate GGA-PBE DFPT dataset assembled from the Materials Project API.

More broadly, we note that there is currently no public cross-chemistry dataset that provides energy, forces, stress, Berry-phase $\mathbf{P}$, Z^* , and α for all materials from a single workflow, so our multihead setting should be interpreted as learning across closely related but not identical semilocal references.

C. Finite-temperature MD trajectory sets: BaTiO₃ and α -SiO₂

To validate that MACE-Field works as intended, and to benchmark against the unified electric-enthalpy study Allegro-pol, we also train two MACE-Field models *exclusively* on finite-temperature MD trajectory datasets for tetragonal BaTiO₃ (300 K) and α -quartz (SiO₂, 300 K) [36].

Each dataset consists of time-ordered frames with per-configuration labels for total energy, forces, virial stress,

Berry-phase polarization, Born effective charges, and electronic polarizability. In contrast to the cross-chemistry setting above, these single-material datasets are internally fully consistent: all six labels are generated within one common electronic-structure workflow for a given material, following the unified-enthalpy dataset construction of Ref. [36].

These datasets serve a dual role: (1) as a demonstration of supervised training for multiproperty learning on single-material manifolds and (2) as the basis for our finite-field MLMD validations (Sec. IV). Specifically, for BaTiO₃ we reproduce polarization hysteresis at 0 K, extracting coercive fields and remanent $\mathbf{P}$; for α -SiO₂, we run long 300 K trajectories (~ 200 ps) and compute infrared absorption from the $\mathbf{P}$ - $\mathbf{P}$ autocorrelation and Raman intensities from the α - α autocorrelation, assembling the complex dielectric function $\epsilon(\omega)$ from time-domain response.

IV. RESULTS

We organize our results in three parts. First, we show that fine-tuning a MACE foundation model on the MP-Dielectric and MP-Ferroelectric datasets yields a field-aware foundation, MACE-Field-MH-0, that preserves its original energy/force accuracy while acquiring transferable predictions of Berry polarization, Born effective charges, polarizabilities, and derived refractive indices across chemistry. Second, we use the MP-Ferroelectric dataset to compare a directly trained cross-chemistry polarization model with MACE-Field-MH-0 for predicting Berry-phase branch structure and spontaneous polarization, and find comparable accuracy. Third, we demonstrate finite-field molecular dynamics for directly trained single-material models and MACE-Field-MH-0 to diagnose the limits of the foundation model. Unless stated otherwise, all quantitative errors in this section are reported against the corresponding semilocal DFT/DFPT labels rather than experiment or higher-level electronic-structure benchmarks.

A. Field-aware foundation model: MACE-Field-MH-0

1. Multihead fine-tuning and parities

We start from the recent multihead MACE foundation model `mace-mp-mh-0` from the broader MACE foundation-model family [59] and use its OMAT-PBE head as the energy/force/stress prior for our field-aware model. We then enable the MACE-Field coupling blocks while keeping the backbone and readout unchanged. For this “MACE-Field-MH-0” model, we perform joint multihead fine-tuning supervised on a “mp-dielectric” head with DFPT Born effective charges and electronic polarizabilities from the MP-Dielectric dataset (Sec. III A); a “mp-ferroelectric” head with DFPT polarizations from the MP-Ferroelectric dataset (Sec. III B); and a replay/pretrained head with energies, forces, and stresses from a 10 000-structure OMAT-PBE replay subset associated with the `mace-mp-mh-0/1` family. This joint dielectric + ferroelectric + replay protocol exploits multihead fine-tuning to combine partially heterogeneous datasets while preserving the accuracy of the underlying `mace-mp-mh-0` foundation model on its original training distribution.

Figure 3 shows parity plots on the OMAT-PBE replay subset, which remain tightly clustered around the diagonal




FIG. 3. OMAT-PBE replay-set parity for MACE-Field-MH-0. Parities on the 10 000-structure OMAT-PBE replay subset used for multihead fine-tuning. Color bars show the log number of data points. Top-left: Energy parity (eV). Top-right: Force parity with all components combined (eV/Å). Bottom: Stress parity for diagonal components (left) and off-diagonal components (right) (eV/Å³).

for energies, forces (all components combined), and both diagonal and off-diagonal stress components, with some increase in variance but no visible systematic bias. In other words, adding field-coupling layers and supervising on dielectric labels does not detrimentally degrade the core potential, so MACE-Field-MH-0 retains a high-quality general-purpose force field after fine-tuning.

Turning to dielectric response, Fig. 4 reports parities for individual BEC tensor components and electronic polarizabilities on the MP-Dielectric dataset. MACE-Field-MH-0 reproduces DFPT Z^* and α across chemistry with near-diagonal trends and narrow scatter, including both diagonal and off-diagonal entries. This indicates that a single electric-enthalpy functional $\mathcal{F}(\mathbf{R}, \mathbf{E})$, inherited from a foundation model and jointly fine-tuned on dielectric- and ferroelectric-response data, is sufficient to capture linear response over a broad chemical space.

To probe distributional behavior, Fig. 5 compares kernel-density estimates of the BEC and polarizability/electronic-dielectric components predicted on (1) the OMAT-PBE replay subset and (2) the MP-Dielectric structures, against the reference DFPT distributions. The three curves closely overlap for both BECs and electronic dielectrics/polarizabilities, with the MACE-Field-MH-0 model neither collapsing toward overly narrow peaks nor developing spurious heavy tails. This suggests that the inherited foundation prior, combined with moderate DFPT supervision, is enough to reconstruct realistic response distributions even on structures for which DFPT labels were never provided (the replay-only OMAT-PBE configurations). Notably, MACE-Field-MH-0 can capture the




FIG. 4. MP-Dielectric parity for MACE-Field-MH-0. Parities for Born effective charges and electronic polarizabilities from the MP-Dielectric dataset (train, validation, and test combined). Color bars show the log number of data points. Top: Born effective charge tensor (e) parity for diagonal (left) and off-diagonal (right) components. Bottom: polarizability tensor (ϵ_0) parity for diagonal (left) and off-diagonal (right) components.

small or zero BECs of many atoms in neutral, bulklike insulating crystal structures present in this OMAT-PBE replay set, which are largely absent in the MP-Dielectric dataset. In practical terms, this means that a single model, trained once, can be used as a general-purpose force field that also supplies consistent dielectric response across tens of thousands of inorganic structures, enabling workflows where dielectric properties are treated on the same footing as energies and forces.




FIG. 5. Density distribution of BECs and electronic dielectric constants. Kernel density estimate (KDE) curves for the Born effective charges Z^* (left) and the high-frequency dielectric constants ϵ_∞ (right). Solid blue: distribution predicted by the fine-tuned MACE-Field model on the MP-Dielectric dataset. Dashed orange: distribution from the reference DFPT MP-Dielectric dataset. Dotted green: distribution predicted by the fine-tuned MACE-Field model on the subselected OMAT-PBE replay subset.




FIG. 6. MP-dielectric parity for predicted ionic and electronic dielectric constants using MACE-Field-MH-0. Top: Parity for the diagonal (left) and off-diagonal (right) components of the ionic dielectric constant. Bottom: Parity for the diagonal (left) and off-diagonal (right) components of the electronic dielectric constant.

2. Direct prediction of dielectric constants

As MACE-Field-MH-0 can predict both Born effective charges and polarizability, we investigate whether we may combine this with analytical Hessians in MACE to evaluate the electronic and ionic dielectric constants directly, without supervising on dielectric tensors themselves.

Concretely, we first obtain the high-frequency (purely electronic) dielectric tensor directly from the predicted polarizability using Eq. (28); the electronic contribution is a straightforward algebraic postprocessing of a primary model output.

The ionic contribution requires vibrational information. Here, we exploit recent work on analytic Hessians for MACE-type models [60], which provides second derivatives of the energy with respect to atomic displacements at negligible additional cost compared to ordinary force evaluation. The ionic dielectric constant is obtained from the MACE-Field-MH-0 BECs and the corresponding mace-mp-mh-0 foundation-model Hessians following Sec. II C 2.

Figure 6 summarizes the parity between the DFPT reference dielectric tensors and the MACE-Field-MH-0 predictions for both ionic and electronic contributions (diagonal and off-diagonal components). The electronic dielectric constants closely follow the DFPT values, as expected from the strong agreement already observed for α . In contrast, the ionic dielectric constants, while still broadly following a near-diagonal trend, exhibit substantially larger scatter. This is not surprising: the ionic dielectric tensor depends on both Z^* and the phonon spectrum via Eq. (36), and is therefore sensitive to errors in both the BECs and the Hessian. Because the ionic dielectric response involves the *inverse* of the Hessian, any




FIG. 7. Matbench dielectric refractive-index benchmark for the fine-tuned MACE-Field foundation model. Refractive-index parity on the Matbench v0.1 test interface, obtained by deriving n from the learned MACE-Field-MH-0 α predictions. Because both Matbench and MP-Dielectric are derived from Materials Project data, this comparison should be interpreted with caution and is not a strict out-of-distribution benchmark. Still, it nevertheless provides a useful external point of comparison for derived optical properties. The color bar shows the logarithm of the number of data points. Inset: 99% percentile of data points.

residual bias in curvature inherited from the foundation prior is amplified in ϵ_{ion} .

3. Matbench dielectric benchmark

As an external benchmark, we test MACE-Field-MH-0 on the Matbench dielectric task [61]. Using the learned α to construct ϵ_{∞} and n [see Eqs. (28)–(30)], we obtain a reasonable refractive-index parity on the Matbench v0.1 test set (Fig. 7), with an root-mean-square error (RMSE) that remains competitive with strong direct-property predictors such as MODNet [62] and AnisoNet [29]. Importantly, this comparison is obtained without directly training on Matbench labels: the same interatomic model is used “as is” to predict α and then derive ϵ_{∞} and n . At the same time, it should not be interpreted as a strict out-of-distribution optical benchmark, because `matbench_dielectric` is itself derived from Materials Project dielectric data. Structural matching between our filtered MP-Dielectric set and the Matbench payload identifies 2842 common structures, corresponding to 53.8% of the MP-Dielectric entries used here and 60.7% of the 4764 Matbench records. We therefore interpret Fig. 7 as a useful external comparison point for derived refractive indices, while noting that the substantial overlap and the sensitivity of this derived quantity to dataset filtering and label-protocol differences warrant caution in overinterpreting the benchmark.




FIG. 8. Error analysis of outliers in BEC predictions. Left: Top 1% most erroneous BEC predictions and their corresponding chemical species. Right: Bar chart showing which elements most frequently appear in these outliers.

4. Error and outlier analysis

In Fig. 8, we summarize the error distribution for BEC components, highlighting that large deviations are confined to a small number of outliers, predominantly reactive nonmetallic elements such as oxygen and metalloids like tellurium, which typically exhibit negative BECs but are positive in the MP-Dielectric dataset. In Table I, we list the top ten most erroneous BEC predictions by MACE-Field-MH-0 compared to the DFPT reference. Notably, the majority of mismatches involve elements that would typically have negative BECs. In these cases, the DFPT BECs violate the acoustic sum rule and appear inconsistent. Despite these outliers, MACE-Field-MH-0 obeys the ASR by construction and so predicts sensible BECs for these structures, which explains the disparity. This illustrates a pragmatic benefit of the physics-informed setup: by enforcing charge neutrality and the ASR at the model level, we can, in some cases, “repair” inconsistent training labels and obtain more physically plausible predictions than those from the raw DFPT dataset.

B. Polarization across distortion paths and spontaneous P_s

In this experiment, we assess the capability of MACE-Field to learn ferroelectric polarization branches from the MP-Ferroelectric distortion-path dataset (Sec. III B).

Figure 9 compares DFPT Berry-phase labels with the *learned* polarizations obtained as $\mathbf{P} = -\Omega^{-1} \partial \mathcal{F} / \partial \mathbf{E}$ from (left) MACE-Field trained directly on polarization labels and (right) from MACE-Field-MH-0 (see Sec. IV A). per-panel (R^2), RMSE, and mean absolute error (MAE) are annotated for each model for the training (gray circles), validation (blue circles), and testing (orange squares) sets. Both models capture polarizations from all materials and distortion paths well, but the jointly fine-tuned OMAT-based MACE-Field-MH-0 model is slightly stronger on the folded branch-resolved task overall.

The spontaneous polarization is obtained as the difference in polarization between the polar and nonpolar reference structures at the end points of a given distortion path. Because the polarization is multivalued modulo a quantum, all polarizations along a path must first be mapped onto a common branch to define this difference unambiguously. The

TABLE I. Top ten BEC outliers by material. Largest erroneous BEC predictions (unit of e) by the fine-tuned MACE-Field model. The first listed element has the largest difference between the MP-Dielectric reference values and the MACE-Field prediction. The second and third elements' BECs in the material are then listed. Notably, all of these outliers correspond to clearly incorrect BECs in the MP-Dielectric dataset; this is most evident from the fact that they do not obey the acoustic sum rule. The MACE-Field predictions do obey the ASR and give sensible BECs.

Material	First element			Second element			Third element		
	El	Ref	Pred	El	Ref	Pred	El	Ref	Pred
$\text{Y}_2\text{Ag}_2\text{Te}_4$	Te	7.171	-2.366	Y	8.915	3.486	Ag	5.692	1.247
$\text{K}_2\text{Au}_2\text{Se}_4$	Se	7.119	-2.615	K	7.515	1.524	Au	7.082	3.707
$\text{Sm}_4\text{Ta}_4\text{O}_{16}$	O	5.362	-2.751	Sm	9.361	3.656	Ta	7.648	7.349
Ca_2SnS_4	S	6.824	-2.402	Ca	7.356	2.498	Sn	6.491	4.610
$\text{La}_4\text{Ta}_4\text{O}_{16}$	La	12.74	3.486	O	5.364	-2.937	Ta	7.648	8.261
CdPt_3O_6	O	5.392	-1.831	Pt	6.941	2.572	Cd	6.143	3.270
MgPt_3O_6	O	5.385	-2.032	Pt	6.947	3.148	Mg	3.821	2.748
$\text{Bi}_2\text{O}_2\text{F}_2$	O	5.389	-3.607	F	5.731	-1.786	Bi	6.588	5.393
$\text{Tm}_4\text{Ta}_4\text{O}_{16}$	O	5.359	-2.933	Tm	4.936	3.356	Ta	7.648	8.376
ZnPt_3O_6	O	5.392	-1.837	Pt	6.947	2.744	Zn	3.840	2.791

parities in Fig. 9 (one point per material; gray circles for the training set, blue circles for the validation set, and orange squares for the test set) compare DFPT reference values with MACE-Field predictions from (left) the directly trained model and (right) MACE-Field-MH-0. Consistent with the pointwise parities, the OMAT-based foundation model reproduces the spontaneous-polarization distribution at least as well as, and slightly better than, the directly trained cross-chemistry model.




FIG. 9. Polarization parity comparisons. Top: Componentwise parity between DFPT Berry-phase polarization and MACE-Field predictions over all materials and path frames; per-panel R^2 , RMSE, and MAE are annotated. (Left) Trained-from-scratch model. (Right) MACE-Field-MH-0. Bottom: Parity of spontaneous polarization P_s (folded end point difference) on the training (gray circles), validation (blue circles), and test set (orange squares); each point is one material. (Left) Trained-from-scratch model. (Right) MACE-Field-MH-0.

To illustrate the branch structure, Fig. 10 (left) shows the $\hat{z}$ Cartesian polarization branches for $\text{Cr}_4\text{Li}_4\text{O}_{16}\text{P}_4$ from the test set. The MACE-Field-MH-0 foundation model reproduces both the magnitude and slope of the Berry-phase branches and does so without any spurious branch hopping. Figure 10 (right) then provides a global view of branch behavior across the full dataset (train, validation, and test combined): it shows a combined violin and “fan” representation (shaded bands indicating the 95% percentile range) of the polarization, normalized by the corresponding polarization quantum, as a function of the distortion parameter. At each distortion step, the predicted polarization distributions closely overlap with the DFPT reference, demonstrating that the model captures realistic polarization branches even in the presence of $\mathbf{P}$'s multivalued nature.

From a methodological standpoint, the small but consistent advantage of MACE-Field-MH-0 on the folded branch-resolved and spontaneous-polarization parities is informative. Joint multihead fine-tuning on MP-Ferroelectric,




FIG. 10. Polarization across distortion paths. Left: Example polarization branches for $\text{Cr}_4\text{Li}_4\text{O}_{16}\text{P}_4$ from the test set; reference (solid) and MACE-Field-MH-0 (dashed) curves align across the path. Right: Branch-invariant “fan” plot: for each Cartesian component, polarization as a fraction of the corresponding polarization quantum vs distortion parameter. Reference DFPT distributions (blue) and MACE-Field-MH-0 (orange) show close agreement; shaded regions indicate 95% percentiles.




FIG. 11. BaTiO₃ ferroelectric hysteresis. BaTiO₃ finite-field MLMD at 0 K in a 135 atom supercell under a single 5 GHz cosine field cycle applied along the polar axis. Left: overlaid polarization-field hysteresis loops from MACE-Field-MH-0 (blue), the directly trained specialist (orange), and DFT branch points at zero field (black). Right: representative trajectory diagnostics for the MACE-Field-MH-0 loop, showing the evolution of $P_z(t)$, the applied field $E_z(t)$, and the average Ti off-centering along the tetragonal c axis during one cycle. The circular markers indicate the six representative frames shown in the Supplemental Material [46]. These simulations probe intrinsic homogeneous switching rather than domain nucleation and growth.

MP-Dielectric, and replay data does not merely preserve the ferroelectric signal from direct supervision; in these benchmarks, it slightly improves upon the directly trained cross-chemistry model. This suggests that foundation-model pretraining, together with simultaneous supervision on local dielectric response, provides a useful inductive bias for the global branch structure of $\mathbf{P}$. The transfer is therefore strong not only for local linear response but also for branch-consistent ferroelectric polarization across chemistry, which is encouraging for future universal multitask materials models.

Overall, these results demonstrate MACE-Field-MH-0 as a field-aware foundation model that (1) preserves its original accuracy on energies, forces, and stresses, and (2) acquires transferable predictions of Born effective charges, polarizabilities, derived refractive indices, and useful ferroelectric-response trends across diverse inorganic materials through joint multihead fine-tuning of a single enthalpy functional.

C. Finite-field MLMD for BaTiO₃ and α -SiO₂

To assess whether MACE-Field can drive field-dependent dynamics and reproduce time-domain observables, we follow the same broad validation strategy as Allegro-pol [36]. Using the same training signals (energy, forces, stress, Berry-phase $\mathbf{P}$, Z^* , and α), we perform fixed-cell MLMD under a spatially uniform field or at zero field, depending on the observable of interest (Sec. II G), comparing two models: a MACE-Field instance trained directly on MD trajectories and our fine-tuned MACE-Field-MH-0 foundation model.

For BaTiO₃, we drive a 135-atom periodic cell through a single 5 GHz cosine field cycle of amplitude $E_0 = 0.3636$ V/Å (about 36 MV/cm) along the polar axis, corresponding to a 200 ps trajectory at 2 fs time step. The trajectory is propagated at 0 K with fixed-cell nve dynamics plus viscous damping, and the polarization-field hysteresis is reconstructed afterward from the stored frames. These field scales are large compared with experimental coercive fields. Still, the

present setup is a small bulk-periodic cell under a spatially uniform bias. It therefore targets the intrinsic homogeneous-switching regime rather than defect- and domain-mediated switching in macroscopic samples [64]. Moreover, the supervised labels for these single-material models constrain the response near zero applied field, so the large-amplitude loop should be viewed as an extrapolative test of the learned enthalpy rather than an interpolation benchmark in field magnitude. Both models generate robust hysteresis loops (Fig. 11). The directly trained model produces a nearly square, symmetric loop with an average absolute remanent polarization $|P_r| \approx 54.5 \mu\text{C cm}^{-2}$ and a coercive field $|E_c| \approx 16.5$ MV/cm. The OMAT-based MACE-Field-MH-0 foundation model yields $|P_r| \approx 48.5 \mu\text{C cm}^{-2}$ and $|E_c| \approx 13.7$ MV/cm, giving a somewhat softer and narrower loop. Table II places these values alongside a DFT branch reference. In this intrinsic, homogeneous-switching regime, the foundation model gives a credible finite-field ferroelectric response. The right

TABLE II. BaTiO₃ intrinsic switching metrics from finite-field MLMD and reference DFT branches. For the dynamic MLMD loops (direct MACE-Field and OMAT foundation MACE-Field-MH-0), we report the average absolute remanent polarization $|P_r|$ and coercive field $|E_c|$ extracted from the loop. For the DFT reference from Fig. 3(a) of Ref. [36], we report the zero-field branch polarization and the terminal field magnitude of the stable upper/lower branches, so the listed $|E_c|$ should be interpreted as an intrinsic branch-limit field rather than a dynamic coercive field. In the present 0 K homogeneous-switching setup, $|P_r|$ and $|P_s|$ are numerically close, so the table provides a compact comparison of the intrinsic ferroelectric scale across models.

Model/reference	$ P_{s/r} $ ($\mu\text{C cm}^{-2}$)	$ E_c $ (MV/cm)
Direct	54.5	16.5
OMAT	48.5	13.7
DFT (PBE) [36]	54.8	18.1




FIG. 12. α -SiO₂ spectroscopy. α -quartz spectroscopy simulations at 300 K from MACE-Field-MH-0 (blue) and the directly trained specialist (orange). Top row: IR spectra (left; from $\mathbf{\hat{P}}\cdot\mathbf{\hat{P}}$), compared with the DFPT reference response (dashed black), and MLMD Raman spectra (right; from $\alpha\text{-}\alpha$) compared with experimental Raman [63] (dotted black). Middle row: real and imaginary parts of the dielectric function $\varepsilon(\omega)$ assembled from time-domain response, again with the DFPT reference shown as dashed black where available. Bottom row: mode-resolved Raman spectra reconstructed from finite-difference polarizability derivatives along optical normal modes for the direct model (orange), the foundation model in its own native mode basis (blue), the foundation response evaluated on the direct-model normal-mode basis (green dashed), and for experimental Raman (black dotted); black triangle markers denote experimental quartz Raman peak positions [63]. Spectra use Hann windows and Gaussian broadening $\sigma = 20 \text{ cm}^{-1}$; intensities are in arbitrary units.

panel of Fig. 11 further shows that the polarization reversal tracks both the applied field and the sign change of the average Ti off-centering along c , consistent with homogeneous tetragonal switching; the six circular markers pick out the representative atomistic snapshots shown in Fig. S1 of the Supplemental Material [46].

For α -quartz, we probe the linear response at 300 K. Spectra are obtained from derivative observables of the learned electric enthalpy evaluated on stored frames from 200 ps trajectories: IR from $\mathbf{\hat{P}}\cdot\mathbf{\hat{P}}$ and Raman from $\alpha\text{-}\alpha$. As in Allegro-pol, we employ a Hann window, Gaussian broadening with $\sigma = 20 \text{ cm}^{-1}$, and branch-invariant wrapping of the Berry-phase polarization $\mathbf{P}(t)$. Figure 12 combines the time-domain

IR/Raman and dielectric comparison with a mode-resolved Raman diagnostic, and Table III collects the main quantitative descriptors.

The directly trained MACE-Field model reproduces the three dominant IR peaks at $\omega_1 \approx 1042 \text{ cm}^{-1}$, $\omega_2 \approx 424 \text{ cm}^{-1}$, and $\omega_3 \approx 763 \text{ cm}^{-1}$, in good agreement (within $\sim 5 \text{ cm}^{-1}$) with DFPT and Allegro-pol values [36]. The corresponding high-frequency dielectric constants $\varepsilon_{\parallel,\perp}^\infty$ and static values $\varepsilon_{\parallel,\perp}^0$ also remain close to the Allegro-pol reference, and broadly consistent with representative PBE DFPT and room-temperature experimental quartz benchmarks (Table III). In the Raman fingerprint region, the directly trained model places most of its mode-resolved activity in the 350–500

TABLE III. Comparison of α -SiO₂ spectroscopic quantities at 300 K: main IR peak positions ω_i (cm⁻¹), high-frequency dielectric constants $\epsilon_{\parallel,\perp}^\infty$, and static dielectric constants $\epsilon_{\parallel,\perp}^0$, parallel ($\parallel$) and perpendicular ($\perp$) to the polar axis. “Allegro-pol” values are taken from Ref. [36]; “Direct” and “MFMH-0” denote, respectively, the directly trained MACE-Field model and the fine-tuned MACE-Field-MH-0 foundation model, all analyzed using the same time-domain protocol [Hann window, Gaussian broadening $\sigma = 20$ cm⁻¹, branch-invariant wrapping of $\mathbf{P}(t)$]. “DFPT (PBE)” and “Expt.” are representative quartz dielectric benchmarks compiled by He *et al.* [65] (experiment at 298 K). Peak positions are estimated from the plotted spectra with an uncertainty of ± 10 –20 cm⁻¹. The directly trained MACE-Field remains close to the Allegro-pol and DFPT references. At the same time, the foundation model is still somewhat softer and more polarizable, especially in the static dielectric response.

Quantity	Allegro-pol	Direct	MFMH-0	DFPT (PBE)	Expt.
ω_1 (cm ⁻¹)	1041	1042	1007	1044	1072
ω_2 (cm ⁻¹)	420	424	404	423	450
ω_3 (cm ⁻¹)	765	763	731	763	797
$\epsilon_{\parallel}^\infty$	2.41	2.43	2.64	2.41	2.42
$\epsilon_{\perp}^\infty$	2.37	2.38	2.58	2.38	2.39
$\epsilon_{\parallel}^0$	4.73	5.02	5.61	4.74	4.64
$\epsilon_{\perp}^0$	4.51	4.65	5.20	4.52	4.43

cm⁻¹ band (about 65% of the total common-basis activity), reproducing the dominant quartz feature around 430–475 cm⁻¹ and aligning well with the experimental [63] markers near 450 and 464 cm⁻¹ in Fig. 12. The direct model does retain weaker high-frequency Si–O stretching activity, but the mode-resolved panel shows that this activity remains relatively discrete; the broader high-frequency background in the MLMD Raman envelope is therefore likely to reflect finite-trajectory and spectral-estimator effects in addition to the underlying response model.

The MACE-Field-MH-0 foundation model retains not only the qualitative band structure but also a realistic quantitative IR/dielectric spectrum. The main IR peaks shift only moderately relative to the directly trained specialist, to $\omega_1 \approx 1007$ cm⁻¹, $\omega_2 \approx 404$ cm⁻¹, and $\omega_3 \approx 731$ cm⁻¹, while the dielectric constants move systematically high: $\epsilon_{\parallel,\perp}^\infty \approx 2.64$ and 2.58, and $\epsilon_{\parallel,\perp}^0 \approx 5.61$ and 5.20. MACE-Field-MH-0 is therefore somewhat softer and more polarizable than the directly trained specialist. The mode-resolved analysis shows that this Raman discrepancy is not primarily a phonon-basis failure: the native direct/foundation mode overlap remains high (mean overlap ≈ 0.83 , mean frequency ratio ≈ 0.96), yet the native Raman-activity correlation is only about 0.14 and remains low even when the foundation response is evaluated on the direct normal-mode basis (common-basis Raman-activity correlation ≈ 0.26). In that common-basis comparison, the foundation suppresses the dominant 350–500 cm⁻¹ quartz Raman band from about 65% of the total activity in the direct model to about 14%, while overweighting the 650–850 and 950–1250 cm⁻¹ bands to about 48% and 34%, respectively. The excess high-frequency Raman intensity is therefore primarily a response-head error, most plausibly

an overestimate of the relevant polarizability derivatives, rather than merely a consequence of slightly shifted phonon frequencies.

Overall, these BaTiO₃ and α -quartz case studies demonstrate that an explicitly field-aware MACE-Field model can be used to drive finite-field MD and to extract nontrivial time-domain observables. The directly trained single-material models provide the most quantitative agreement with DFPT and the Allegro-pol reference. The OMAT-based foundation model still produces a credible intrinsic BaTiO₃ hysteresis loop and a realistic quartz IR/dielectric response. Still, the mode-resolved Raman analysis indicates that its Raman-activity distribution is not yet quantitatively reliable. From a computational-science perspective, the framework is therefore not only a property regressor but also a stable dynamical engine for comparative finite-field simulations, with the Raman case providing a concrete example of where targeted response supervision remains important.

V. DISCUSSION

We introduced MACE-Field, a field-aware, $O(3)$ -equivariant interatomic potential that learns a *single* electric enthalpy $\mathcal{F}(\{\mathbf{R}\}, \mathbf{E})$ and obtains polarization $\mathbf{P}$, Born effective charge Z^* , and polarizability α by exact differentiation. A uniform field couples to learned latent spherical-tensor features via Clebsch-Gordan products, while the final readout remains a scalar. Benchmarked against semilocal DFT/DFPT reference data, a directly trained cross-chemistry model reproduces Berry-phase polarizations (including P_s) for ~ 250 ferroelectric materials spanning 61 elements. Separately, a fine-tuned foundation model, MACE-Field-MH-0, jointly trained on ~ 5.3 k DFPT dielectric/BEC entries, ~ 2.5 k ferroelectric distortion-path structures, and an OMAT-PBE replay set predicts Z^* and α with high fidelity (Figs. 4 and 7), recovers both electronic and ionic dielectric trends through analytic Hessians, and captures same-branch polarization and spontaneous-polarization trends across chemistry. Finally, single-material models for BaTiO₃ and α -SiO₂ demonstrate that the same framework can drive finite-field MD and reproduce hysteresis loops, IR/Raman spectra, and dielectric functions with DFPT accuracy, while the MACE-Field-MH-0 foundation model demonstrates comparative accuracy with a systematic softening of the resulting dynamics and a less reliable Raman-activity distribution (Figs. 11 and 12).

Conceptually, MACE-Field is therefore a prototype of a physics-informed, differentiable foundation model for materials [59,66]: it combines strong inductive biases (symmetries, enthalpy derivatives) with data-driven learning. It can be retrofitted onto existing energy/force models. Because all observables share one scalar origin, Maxwell reciprocity, the acoustic sum rule, point-group tensor forms, and branch-invariant learning of Berry-phase $\mathbf{P}$ hold *by construction* (Sec. 1G).

Key advances.

(1) *Plug-in field coupling that inherits MACE foundations.* The field enters only through irrep-wise latent tensor products and equivariant mixing; the ACE graph construction and scalar

readout are unchanged. Consequently, existing MACE foundation weights can be dropped in and fine-tuned to become field aware. Our MACE-Field-MH-0 model shows that this inheritance preserves the accuracy of the OMAT-based multihead `mace-mp-mh-0` prior on energies, forces, and stresses (Fig. 3), while adding transferable predictions of Z^* , α , and derived optical properties.

(2) *One model across diverse chemistry.* Unlike prior unified-enthalpy demonstrations restricted to single materials, MACE-Field is trained once across thousands of crystals covering $\gtrsim 80$ elements. The MACE-Field-MH-0 dielectric foundation model reproduces DFPT BECs and electronic polarizabilities across the MP-Dielectric set (Fig. 4) and yields realistic distributions of response quantities on unseen replay structures from the OMAT-PBE pool (Fig. 5). It also provides reasonable refractive-index parity on the Matbench dielectric task, with performance in the same broad regime as strong direct-property predictors. However, the substantial overlap with Materials-Project-derived data means that this comparison should be interpreted with caution rather than as a strict out-of-distribution benchmark (Fig. 7).

(3) *Time-domain finite-field validation.* Single-material MACE-Field models for BaTiO₃ and α -SiO₂ reproduce DFPT and Allegro-pol benchmarks for polarization hysteresis and IR/Raman/dielectric spectra, with MACE-Field-MH-0 maintaining comparative accuracy despite not being trained on the same MD trajectories. These MACE-Field models match IR peak positions and $\varepsilon_{\parallel,\perp}^{\infty,0}$ for α -quartz to within the uncertainty of the spectral analysis (Table III) and generate BaTiO₃ loops with the expected intrinsic, uniform-field coercive scales for a small periodic cell (Fig. 11).

What we learn from the fine-tuned foundation model. To our knowledge, MACE-Field-MH-0 offers a first demonstration of a field-aware *foundation* model built from a recent multihead OMAT-initialized prior: starting from `mace-mp-mh-0` and its OMAT-PBE head, we add field-coupling blocks and perform joint multiheaded fine-tuning on DFPT BECs, polarizabilities, Berry-phase polarizations, and a replay set of energies/forces/stresses. This succeeds remarkably well for cross-chemistry dielectric response and also transfers substantial ferroelectric information. The model captures Z^* and α across chemistry, reconstructs realistic refractive-index distributions, and (combined with analytic Hessians) yields reasonable electronic and ionic dielectric constants without direct supervision on ε (Fig. 6). On the folded same-branch and spontaneous-polarization benchmarks, it performs as well as a directly trained cross-chemistry ferroelectric model. Outlier analysis reveals that the largest discrepancies in Z^* often correspond to clear pathologies in the DFPT labels (violations of the acoustic sum rule). In contrast, the learned enthalpy automatically enforces ASR and produces physically sensible charges (Fig. 8 and Table I).

At the same time, MACE-Field-MH-0 exposes important limitations of fine-tuning a broadly transferable force field for field response. First, the ionic dielectric constants show significantly larger scatter than the electronic ones. This reflects the fact that ε_{ion} depends on both Z^* and the phonon spectrum via the inverse Hessian [Eq. (36)], and hence amplifies any residual bias in the underlying foundation-model Hessians. The softened phonons and overestimated static permittivities

in our quartz tests are consistent with a slightly oversmoothed potential-energy surface learned for broad transferability: total energies can remain usable while curvature-sensitive observables such as phonons, elastic constants, Raman derivatives, and ionic dielectric constants drift systematically. In our case, this manifests as mildly redshifted IR peaks and elevated permittivities, both in the analytic ε_0 and in the time-domain spectra of α -quartz (Fig. 12 and Table III).

Second, the cross-chemistry polarization benchmarks show that the fine-tuned MACE-Field-MH-0 foundation can do more than merely approach a directly trained specialist: in the present folded parity metrics, it slightly surpasses the direct model for both same-branch Berry-phase polarization and spontaneous P_s (Fig. 9). This is instructive. Learning polarization modulo the polarization quantum via a folded loss is intrinsically challenging, particularly when combined with heterogeneous semilocal workflows and multitask supervision. Joint training on BECs provides strong local information, but the branch structure of $\mathbf{P}$ is a global, path-dependent object; here, however, the combination of foundation pretraining, replay stabilization, and simultaneous dielectric supervision appears to provide a useful inductive bias rather than a detriment. In addition, MP-Ferroelectric and MP-Dielectric labels come from independent DFT calculations, so even subtle inconsistencies in geometry, functional, or convergence criteria effectively act as label noise when a single model is asked to reconcile both.

Third, Raman line shapes remain more model dependent than either IR peak positions or static dielectric constants. The mode-resolved quartz diagnostic sharpens this conclusion. The foundation model retains phonons broadly similar to those of the direct model (mean native mode overlap ≈ 0.83 and mean frequency ratio ≈ 0.96). Still, the Raman-activity parity remains poor both in the native comparison and when the foundation response is evaluated on the direct normal-mode basis. In other words, the excess foundation intensity in the 650–850 and 950–1250 cm⁻¹ bands survives even when the vibrational basis is held fixed. This indicates that the dominant error is in the learned polarizability response, especially the mode-resolved derivatives $\partial\alpha/\partial Q_m$, rather than primarily in the phonon frequencies themselves.

Taken together, these observations suggest that *field-aware foundations* such as MACE-Field-MH-0 can provide not only broadly accurate Z^* , α , and derived dielectric constants across chemistry, but also the strongest cross-chemistry ferroelectric parity among the models compared here, together with credible finite-field switching behavior. Targeted, single-material (or small-chemistry) fine-tuning is still beneficial when the goal is the most quantitative Raman spectra or material-specific finite-field dynamics.

Position within the broader ferroelectrics landscape. The MACE-Field architecture should be viewed as a derivative-consistent model for bulk-periodic response under a homogeneous external field. In particular, because the field input is global and no explicit self-consistent long-range Coulomb term is solved, spatially varying depolarizing fields, charged defects, and interface electrostatics are outside the present model class. Likewise, bulk phonon effects that depend on explicit long-range Coulomb physics, such as LO-TO splitting, are only learned implicitly from the training data rather than

enforced analytically. Addressing those regimes will likely require either site-dependent field conditioning or hybridization with explicit long-range electrostatics, and will be the subject of future work.

Outlook. There are several directions to extend and exploit the MACE-Field framework:

(1) *Better foundations and phonons.* The behavior of MACE-Field-MH-0 underscores the importance of high-quality forces and Hessians. Retraining foundations with tighter DFT settings, explicit phonon or elastic-constant targets, or incorporating DFPT force-constant or higher-order derivative information into the loss [60] could harden the phonon spectrum and improve ϵ_0 , hysteresis, and Raman response.

(2) *Mode-resolved Raman supervision.* The quartz diagnostics suggest that matching equilibrium α is not sufficient for quantitatively correct Raman spectra. A more targeted route is to generate DFPT or finite-field polarizabilities for $\pm Q_m$ displacements along Raman-relevant normal modes and to fine-tune directly on $d\alpha_{ij}/dQ_m$, together with isotropic, anisotropic, or activity-based Raman invariants. Mode-space validation should then become a standard check before trusting MLMD Raman envelopes.

(3) *Hybrid long-range, higher-order, and higher-fidelity response.* Combining MACE-Field with site-dependent auxiliary internal-field modules, latent/Ewald terms, or charge equilibration, while preserving derivative consistency, would extend its reach to highly ionic or low-dimensional systems. Nothing in the architecture ties it specifically to semilocal DFT: in principle, one could fine-tune on r^2 SCAN, hybrid-DFT, or multifidelity response datasets where such labels are available. Likewise, selected second- and third-order susceptibilities $\chi^{(2)}/\chi^{(3)}$ and hyperpolarizabilities are formally accessible as higher derivatives of $\mathcal{F}$. Still, quantitative non-linear responses will require dedicated training labels, as the current datasets primarily constrain the linear, near-zero-field regime.

(4) *From foundations to targeted specialists.* Field-aware foundations such as MACE-Field-MH-0 can serve as starting points for lightweight, material- or family-specific fine-tuning that sharpens polarization branches, P_s , and Raman intensities without retraining from scratch. This hierarchical strategy mirrors current practice in energy/force foundation models.

(5) *High-throughput dielectric and ferroelectric discovery.* A MACE-Field foundation fine-tuned on MP-Dielectric and MP-Ferroelectric can be deployed in large-scale screening campaigns for ferroelectrics, multiferroics, hyperferroelectrics, and high- κ dielectrics, including finite-field MD to probe switching pathways, leakage, and finite-temperature response.

(6) *Broader applications.* Beyond inorganic dielectrics, the same idea, a single enthalpy functional that is differen-

tiable with respect to fields and structural degrees of freedom, could be applied to a wide range of systems where field response matters: electrolytes and solid-electrolyte interfaces in batteries, polar liquids and biomolecules in strong fields, heterogeneous catalysis under bias, and geophysical materials in the deep Earth. In each case, a reusable, field-aware force field would enable computational scientists across communities to ask different questions about how materials behave under realistic operating conditions.

Methodologically, our work aligns with the broader push toward physics-informed machine learning in computational science, where enforcing physical structure in ML models improves robustness, interpretability, and transfer across tasks. To conclude, by learning a single electric enthalpy and differentiating it, MACE-Field unifies $\mathbf{P}$, Z^* , and α in a symmetry-consistent way that can inherit weights from existing MACE foundations. The present results show that simple coupling of a uniform field to learned equivariant latent features is already sufficient to endow a pretrained foundation with broadly accurate dielectric response, very strong cross-chemistry ferroelectric transfer capability, and credible finite-field dynamics. At the same time, the single-material studies show that the most quantitative Raman intensities and material-specific dynamical response still benefit from dedicated supervision. We therefore view MACE-Field as a practical route to field-aware atomistic modeling, whose natural next step is targeted fine-tuning on higher-fidelity or more locally field-rich data.

ACKNOWLEDGMENTS

K.T.B. and B.A.A.M. acknowledge support from the Engineering and Physical Sciences Research Council (EPSRC, grant EP/Y014405/1). Via our membership of the UK's HEC Materials Chemistry Consortium, funded by EPSRC (grant EP/L000202), this work used the UK Materials and Molecular Modelling Hub, which is partially funded by EPSRC (grants EP/T022213/1, EP/W032260/1, and EP/P020194/1).

DATA AVAILABILITY

The MACE-Field source code is available in Ref. [39]. The manuscript-specific repository, Ref. [67], contains processed train/validation/test split files, training shell scripts and configuration files, replay-set selection inputs, and the analysis code used for the figures reported hereâ including polarisation-branch folding, dielectric post-processing, Matbench-overlap analysis, and hysteresis/spectroscopy plotting.

-
- [1] J. F. Scott, Applications of modern ferroelectrics, *Science* **315**, 954 (2007).
 - [2] D. Damjanovic, Ferroelectric, dielectric and piezoelectric properties of ferroelectric thin films and ceramics, *Rep. Prog. Phys.* **61**, 1267 (1998).

- [3] R. Boyd, *Nonlinear Optics, Electronics & Electrical* (Academic Press, San Diego, CA, 2003).
- [4] J. Robertson, High dielectric constant gate oxides for metal oxide Si transistors, *Rep. Prog. Phys.* **69**, 327 (2006).

- [5] R. D. King-Smith and D. Vanderbilt, Theory of polarization of crystalline solids, *Phys. Rev. B* **47**, 1651 (1993).
- [6] R. Resta, Macroscopic polarization in crystalline dielectrics: The geometric phase approach, *Rev. Mod. Phys.* **66**, 899 (1994).
- [7] X. Gonze and C. Lee, Dynamical matrices, Born effective charges, dielectric permittivity tensors, and interatomic force constants from density-functional perturbation theory, *Phys. Rev. B* **55**, 10355 (1997).
- [8] K. M. Rabe, M. Dawber, C. Lichtensteiger, C. H. Ahn, and J.-M. Triscone, in *Physics of Ferroelectrics: A Modern Perspective, Topics in Applied Physics*, edited by K. M. Rabe, C. H. Ahn, and J.-M. Triscone (Springer, Berlin, 2007), Vol. 105, pp. 1–30.
- [9] S. M. Young and A. M. Rappe, First principles calculation of the shift current photovoltaic effect in ferroelectrics, *Phys. Rev. Lett.* **109**, 116601 (2012).
- [10] A. Cook, B. M. Fregoso, F. de Juan, S. Coh, and J. Moore, Design principles for shift current photovoltaics, *Nat. Commun.* **8**, 14176 (2017).
- [11] Z. Dai and A. M. Rappe, Recent progress in the theory of bulk photovoltaic effect, *Chem. Phys. Rev.* **4**, 011303 (2023).
- [12] K. T. Butler, J. M. Frost, and A. Walsh, Ferroelectric materials for solar energy conversion: Photoferroics revisited, *Energy Environ. Sci.* **8**, 838 (2015).
- [13] J. M. Frost, K. T. Butler, and A. Walsh, Molecular ferroelectric contributions to anomalous hysteresis in hybrid perovskite solar cells, *APL Mater.* **2**, 081506 (2014).
- [14] F. Che, J. T. Gray, S. Ha, N. Kruse, S. Scott, and J.-S. McEwen, Elucidating the roles of electric fields in catalysis: A perspective, *ACS Catal.* **8**, 5153 (2018).
- [15] X. Huang, C.-L. Tang, J. Li, L. Chen, J. Zheng, P. Zhang, J.-B. Le, R. Li, X. Li, J. Liu, *et al.*, Electric field-induced selective catalysis of single-molecule reaction, *Sci. Adv.* **5**, eaaw3072 (2019).
- [16] N. G. Léonard, R. Dhaoui, T. Chantarojsiri, and J. Y. Yang, Electric fields in catalysis: From enzymes to molecular catalysts, *ACS Catal.* **11**, 10923 (2021).
- [17] J. Yu, J. Yin, R. Li, Y. Ma, and Z. Fan, Interfacial electric field effect on electrochemical carbon dioxide reduction reaction, *Chem Catal.* **2**, 2229 (2022).
- [18] Z. Long, J. Meng, L. R. Weddle, P. E. Videla, J. P. Menzel, D. G. A. Cabral, J. Liu, T. Qiu, J. M. Palasz, D. Bhattacharyya, *et al.*, The impact of electric fields on processes at electrode interfaces, *Chem. Rev.* **125**, 1604 (2025).
- [19] R. W. Nunes and X. Gonze, Berry-phase treatment of the homogeneous electric field perturbation in insulators, *Phys. Rev. B* **63**, 155107 (2001).
- [20] I. Souza, J. Íñiguez, and D. Vanderbilt, First-principles approach to insulators in finite electric fields, *Phys. Rev. Lett.* **89**, 117602 (2002).
- [21] A. S. Christensen, F. A. Faber, and O. A. von Lilienfeld, Operators in quantum machine learning: Response properties in chemical space, *J. Chem. Phys.* **150**, 064105 (2019).
- [22] A. Grisafi, D. Wilkins, G. Csányi, and M. Ceriotti, Symmetry-adapted machine learning for tensorial properties of atomistic systems, *Phys. Rev. Lett.* **120**, 036002 (2018).
- [23] D. M. Wilkins, A. Grisafi, Y. Yang, K. U. Lao, R. A. DiStasio, and M. Ceriotti, Accurate molecular polarizabilities with coupled cluster theory and machine learning, *Proc. Natl. Acad. Sci. USA* **116**, 3401 (2019).
- [24] A. Grisafi and M. Ceriotti, Incorporating long-range physics in atomic-scale machine learning, *J. Chem. Phys.* **151**, 204105 (2019).
- [25] O. T. Unke and M. Meuwly, PhysNet: A neural network for predicting energies, forces, dipole moments, and partial charges, *J. Chem. Theory Comput.* **15**, 3678 (2019).
- [26] M. Veit, D. M. Wilkins, Y. Yang, R. A. DiStasio, and M. Ceriotti, Predicting molecular dipole moments by combining atomic partial charges and atomic dipoles, *J. Chem. Phys.* **153**, 024113 (2020).
- [27] M. Gastegger, K. T. Schütt, and K.-R. Müller, Machine learning of solvent effects on molecular spectra and reactions, *Chem. Sci.* **12**, 11473 (2021).
- [28] K. Morita, D. W. Davies, K. T. Butler, and A. Walsh, Modeling the dielectric constants of crystals using machine learning, *J. Chem. Phys.* **153**, 024503 (2020).
- [29] Y. Lou and A. M. Ganose, Discovery of highly anisotropic dielectric crystals with equivariant graph neural networks, *Faraday Discuss.* **256**, 255 (2025).
- [30] J. Behler and M. Parrinello, Generalized neural-network representation of high-dimensional potential-energy surfaces, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [31] K. Schütt, P.-J. Kindermans, H. E. S. Felix, S. Chmiela, A. Tkatchenko, and K. Müller, in *Advances in Neural Information Processing Systems*, edited by I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2017), Vol. 30, pp. 991–1001.
- [32] S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, and B. Kozinsky, E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials, *Nat. Commun.* **13**, 2453 (2022).
- [33] I. Batatia, D. Kovács, G. Simm, C. Ortner, and G. Csányi, in *Neural Information Processing Systems*, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Neural Information Processing Systems Foundation, Inc., La Jolla, CA, 2022), Vol. 35, pp. 11423–11436.
- [34] L. Zhang, M. Chen, X. Wu, H. Wang, W. E, and R. Car, Deep neural network for the dielectric response of insulators, *Phys. Rev. B* **102**, 041121(R) (2020).
- [35] L. Zhang, H. Wang, M. Muniz, A. Panagiotopoulos, R. Car, and W. E, A deep potential model with long-range electrostatic interactions, *J. Chem. Phys.* **156**, 124107 (2022).
- [36] S. Falletta, A. Cepellotti, A. Johansson, C. W. Tan, M. L. Descoteaux, A. Musaelian, C. J. Owen, and B. Kozinsky, Unified differentiable learning of electric response, *Nat. Commun.* **16**, 4031 (2025).
- [37] B. Cheng, Latent Ewald summation for machine learning of long-range interactions, *npj Comput. Mater.* **11**, 80 (2025).
- [38] M. Rinaldi, A. Bochkarev, Y. Lysogorskiy, and R. Drautz, Charge-constrained atomic cluster expansion, *Phys. Rev. Mater.* **9**, 033802 (2025).
- [39] B. A. A. Martin, A. M. Ganose, V. Kapil, T. Li, and K. T. Butler, MACE-Field: Electric-field-aware MACE models, GitHub repository, 2025, <https://github.com/mdi-group/mace-field>, accessed 5 June 2026.
- [40] With $\mathbf{D} = \epsilon_0 \mathbf{E} + \mathbf{P}$ and $\mathbf{P} = \epsilon_0 \chi^{(1)} \mathbf{E}$, one has $\alpha = \epsilon_0 \chi^{(1)}$ and $\epsilon_\infty = \epsilon_0 (\mathbf{I} + \chi^{(1)})$.

- [41] M. A. Thomas, M. Brehm, R. Fligg, P. Vöhringer, and B. Kirchner, Computing vibrational spectra from *ab initio* molecular dynamics, *Phys. Chem. Chem. Phys.* **15**, 6608 (2013).
- [42] S. Lubert, M. Iannuzzi, and J. Hutter, Raman spectra from *ab initio* molecular dynamics and its application to liquid S-methyloxirane, *J. Chem. Phys.* **141**, 094503 (2014).
- [43] Different quadrupole conventions exist (traced versus traceless tensors, differing numerical factors). Our use of the multipole series is solely as symmetry *guidance* for constructing equivariant couplings; numerical prefactors in Eq. (14) do not enter the learned architecture.
- [44] R. Drautz, Atomic cluster expansion for accurate and transferable interatomic potentials, *Phys. Rev. B* **99**, 014104 (2019).
- [45] I. Petousis, D. Mrdjenovich, E. Ballouz, M. Liu, D. Winston, W. Chen, T. Graf, T. D. Schladt, K. Persson, and F. Prinz, High-throughput screening of inorganic compounds for the discovery of novel dielectric and optical materials, *Sci. Data* **4**, 160134 (2017).
- [46] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/b116-xy8k> for training configurations and loss weights, run-time comparisons, dataset curation and splits, autograd and polarization-folding recipes, finite-field molecular-dynamics and spectroscopy protocols, representative BaTiO₃ hysteresis and SiO₂ trajectory snapshots, and additional parity/training-curve benchmarks.
- [47] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. Castelli, R. Christensen, M. Dulak, J. Friis, M. Groves, B. Hammer, C. Hargus, *et al.*, The atomic simulation environment—A Python library for working with atoms, *J. Phys.: Condens. Matter* **29**, 273002 (2017).
- [48] T. E. Smidt, S. A. Mack, S. E. Reyes-Lillo, A. Jain, and J. B. Neaton, An automatically curated first-principles database of ferroelectrics, *Sci. Data* **7**, 72 (2020).
- [49] D. Akbarian, D. Yilmaz, Y. Cao, Y. Cao, P. Ganesh, I. Dabo, J. M. Munro, R. Ginhoven, and A. Duin, Understanding the influence of defects and surface chemistry on ferroelectric switching: A ReaxFF investigation of BaTiO₃, *Phys. Chem. Chem. Phys.* **21**, 18240 (2019).
- [50] K. Kelley, A. Morozovska, E. Eliseev, V. Sharma, D. Yilmaz, A. C. T. van Duin, P. Ganesh, A. Borisevich, S. Jesse, P. Maksymovych, *et al.*, Oxygen vacancy injection as a pathway to enhancing electromechanical response in ferroelectrics, *Adv. Mater.* **34**, 2106426 (2022).
- [51] V. Kapil, D. P. Kovács, G. Csányi, and A. Michaelides, First-principles spectroscopy of aqueous interfaces using machine-learned electronic and quantum nuclear effects, *Faraday Discuss.* **249**, 50 (2024).
- [52] Public MP documentation describes DFPT workflows for dielectric properties (using VASP) within the semilocal GGA family; our API-assembled set consists of insulating GGA-PBE entries with dielectric tensors and Born effective charges.
- [53] M. K. Horton, P. Huck, R. X. Yang, J. M. Munro, S. Dwaraknath, A. Ganose, R. Kingsbury, M. Wen, J.-X. Shen, T. S. Mathis, *et al.*, Accelerated data-driven materials science with the Materials Project, *Nat. Mater.* **24**, 1522 (2025).
- [54] A. Jain, S. Ong, G. Hautier, W. Chen, W. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, *et al.*, Commentary: The Materials Project: A materials genome approach to accelerating materials innovation, *APL Mater.* **1**, 011002 (2013).
- [55] P. Huck, A. Jain, D. Gunter, D. Winston, and K. A. Persson, in *IEEE International Conference on e-Science* (IEEE, Piscataway, NJ, 2015), pp. 535–541.
- [56] P. Huck, D. Gunter, S. Cholia, D. Winston, A. N'Diaye, and K. Persson, User applications driven by the community contribution framework MPContribs in the Materials Project, *Concurrency Comput.* **28**, 1982 (2016).
- [57] *Handbook of Materials Modeling: Methods: Theory and Modeling*, edited by W. Andreoni and S. Yip (Springer, Cham, 2020), 2nd ed.
- [58] See the “Identifying ferroelectricity from first principles” and “Post-processing spontaneous polarisation values” sections in Ref. [48]; 255 structure pairs satisfy the “COMPLETED” workflow criteria there.
- [59] I. Batatia, P. Benner, C. Yuan, A. Elena, D. Kovács, J. Riebesell, X. R. Advincula, M. Asta, W. J. Baldwin, N. Bernstein, *et al.*, A foundation model for atomistic materials chemistry, *J. Chem. Phys.* **163**, 184110 (2025).
- [60] N. Gönheimer, K. Reuter, and J. T. Margraf, Beyond numerical Hessians: Higher-order derivatives for machine learning interatomic potentials via automatic differentiation, *J. Chem. Theory Comput.* **21**, 4742 (2025).
- [61] A. Dunn, Q. Wang, A. Ganose, D. Dopp, and A. Jain, Benchmarking materials property prediction methods: The Matbench test set and Automatminer reference algorithm, *npj Comput. Mater.* **6**, 138 (2020).
- [62] P.-P. D. Breuck, M. L. Evans, and G. Rignanese, Robust model benchmarking and bias-imbalance in data-driven materials science: A case study on MODNet, *J. Phys.: Condens. Matter* **33**, 404002 (2021).
- [63] Laboratoire de Sciences de la Terre ENS-Lyon, Quartz (powder) Raman spectrum, Handbook of Raman Spectra for Geology (2026), Raman spectrum of quartz powder collected with a 514.5 nm excitation line, accessed 5 May 2026, <https://www.geologie-lyon.fr/Raman/spectrum.php?nom=quartz+%28powder%29>.
- [64] M. Dawber, K. M. Rabe, and J. F. Scott, Physics of thin-film ferroelectric oxides, *Rev. Mod. Phys.* **77**, 1083 (2005).
- [65] L. He, F. Liu, G. Hautier, M. J. T. Oliveira, M. A. L. Marques, F. D. Vila, J. J. Rehr, G.-M. Rignanese, and A. Zhou, Accuracy of generalized gradient approximation functionals for density-functional perturbation theory calculations, *Phys. Rev. B* **89**, 064305 (2014).
- [66] C. Chen, W. Ye, Y. Zuo, C. Zheng, and S. P. Ong, Graph networks as a universal machine learning framework for molecules and crystals, *Chem. Mater.* **31**, 3564 (2019).
- [67] B. A. A. Martin, A. M. Ganose, V. Kapil, T. Li, and K. T. Butler, MACE-Field: General learning of the electric response of inorganic materials, GitHub, 2025, <https://github.com/mdigroup/2025-04-mace-field>.