Supplementary Information for: "Long-Range Electrostatics and Born Effective Charges from Invariant Moment Tensor Potentials with Latent Ewald Summation (MTP-LES"

This document collects supporting material for the main paper. Sections are referenced as "Supplementary Note SX" and figures, tables, and equations carry an "S" prefix throughout.

How the SI maps to the main text.

- Supplementary Note S1 gives the full mathematical specification of the rotationally invariant MTP descriptors summarized in the main-text Methods.
- Supplementary Note S2 provides extended quantitative detail supporting the net-charge comparison in Fig. 1 of the main text.
- Supplementary Note S3 gives the full VASP input specification for the HfO_2 dataset.
- Supplementary Note S4 provides quantitative analysis of the liquid-water infrared spectrum.
- Supplementary Note S5 reports per-species BEC RMSE values for water and dipeptides.
- Supplementary Note S6 reports the dipeptide dipole-loss ablation.
- Supplementary Note S7 provides the raw timing data behind the runtime benchmark.
- Supplementary Note S8 documents the periodic-image neighbour-list implementation detail required to reproduce periodic MTP–LES training.

Contents

S1 Full specification of the invariant MTP descriptors

This note specifies the rotationally invariant moment-tensor descriptors used in this work, following the Moment Tensor Potential construction of Shapeev [1, 2]. We employ a compact set of invariant contractions of moment tensors, given explicitly below together with its body-order and level truncation; this set is sufficient to reach the energy and force accuracy reported for the systems studied.

S1.1 Moment tensors

For each atom i, moment tensors of rank $\nu \in \{0, 1, 2\}$ are constructed from the local neighborhood within a cutoff radius r_{cut}:

$$M^{(0)}_{\mu,i} = \sum_{j\in\mathcal{N}(i)} R_\mu(r_{ij}, s_i, s_j), \tag{S1}$$

$$M^{(1)}_{\mu,i,\alpha} = \sum_{j\in\mathcal{N}(i)} R_\mu(r_{ij}, s_i, s_j)\, r_{ij,\alpha}, \tag{S2}$$

$$M^{(2)}_{\mu,i,\alpha\beta} = \sum_{j\in\mathcal{N}(i)} R_\mu(r_{ij}, s_i, s_j)\, r_{ij,\alpha}\, r_{ij,\beta}, \tag{S3}$$

where $r_{ij} = |\mathbf{r}_j - \mathbf{r}_i|$, $r_{ij,\alpha}$ is the α-th Cartesian component of $\mathbf{r}_j - \mathbf{r}_i$, s_i, s_j are species indices, and $\mathcal{N}(i)$ is the set of neighbours of i within r_{cut}.

S1.2 Chebyshev radial basis

The radial functions are expanded using the Chebyshev polynomial basis employed in the MLIP-2 framework [2]. For each pair of atomic species (s_i, s_j), the radial basis is written as

$$R_\mu(r, s_i, s_j) = \sum_{n=1}^{n_{\mathrm{rad}}} c_{\mu,s_i,s_j,n}\, T_n(\xi(r))\, f_{\mathrm{cut}}(r), \tag{S4}$$

where $c_{\mu,s_i,s_j,n}$ are learnable species-pair-dependent coefficients and T_n denotes the nth Chebyshev polynomial of the first kind, generated recursively as

$$T_0 = 1, \qquad T_1 = \xi, \qquad T_n = 2\xi T_{n-1} - T_{n-2}. \tag{S5}$$

The scaled radial coordinate

$$\xi(r) = \frac{2r - r_{\mathrm{min}} - r_{\mathrm{cut}}}{r_{\mathrm{cut}} - r_{\mathrm{min}}} \tag{S6}$$

maps the interval $[r_{\mathrm{min}}, r_{\mathrm{cut}}]$ onto $[-1, 1]$. Each radial basis function is multiplied by the smooth cutoff envelope

$$f_{\mathrm{cut}}(r) = \begin{cases} \left(\dfrac{r_{\mathrm{cut}} - r}{r_{\mathrm{cut}} - r_{\mathrm{min}}}\right)^2, & r < r_{\mathrm{cut}}, \\ 0, & r \geq r_{\mathrm{cut}}. \end{cases} \tag{S7}$$

This formulation is equivalent to the Chebyshev radial basis used in MLIP-2, differing only by the constant normalization factor $(r_{\mathrm{cut}} - r_{\mathrm{min}})^{-2}$, which is absorbed into the learnable coefficients $c_{\mu,s_i,s_j,n}$.

S1.3 Rotational invariants by tensor contraction

Following the Moment Tensor Potential (MTP) formalism, each moment tensor is assigned a level

$$\mathrm{lev}\left(M^{(\nu)}_\mu\right) = 2\mu + 4\nu, \tag{S8}$$

which determines the complexity of the corresponding basis function. The level of a product of moment tensors is defined as the sum of the levels of the individual tensors. Rotationally invariant basis functions are constructed by contracting moment tensors whose total level does not exceed ℓ_{max} and whose tensor rank is limited by ν_{max}. Throughout this work, we use $\ell_{\mathrm{max}} = 24$ and $\nu_{\mathrm{max}} = 2$.

Rather than employing the complete MTP invariant basis, we retain a reduced subset that provides sufficient expressiveness while keeping the descriptor compact. In particular, invariants such as $\operatorname{tr} M_\mu^{(2)}$ and higher-order contractions involving $M^{(2)}$ are omitted. The rotational invariants used in this work are

$$\text{zeroth-order moment:} \qquad B = M_\mu^{(0)}, \tag{S9}$$

$$\text{scalar products:} \qquad B = M_{\mu_0}^{(0)} M_{\mu_1}^{(0)}, \qquad \left(M_{\mu_0}^{(0)}\right)^2, \tag{S10}$$

$$\text{vector contraction:} \qquad B = \sum_\alpha M_{\mu_1,\alpha}^{(1)} M_{\mu_2,\alpha}^{(1)}, \tag{S11}$$

$$\text{Frobenius contraction:} \qquad B = \sum_{\alpha\beta} M_{\mu_1,\alpha\beta}^{(2)} M_{\mu_2,\alpha\beta}^{(2)}, \tag{S12}$$

$$\text{scalar-weighted vector contraction:} \qquad B = M_{\mu_0}^{(0)} \sum_\alpha M_{\mu_1,\alpha}^{(1)} M_{\mu_2,\alpha}^{(1)}, \tag{S13}$$

$$\text{vector–matrix–vector contraction:} \qquad B = \sum_{\alpha\beta} M_{\mu_1,\alpha}^{(1)} M_{\mu_2,\alpha\beta}^{(2)} M_{\mu_3,\beta}^{(1)}. \tag{S14}$$

The scalar-weighted vector-contraction family includes the special case $\mu_1 = \mu_2$, namely

$$M_{\mu_0}^{(0)} \left|\mathbf{M}_{\mu_1}^{(1)}\right|^2. \tag{S15}$$

These contractions generate rotationally invariant combinations of the scalar, vector, and second-order tensor moments while retaining information about the local atomic environment.

Finally, a one-hot encoding of the central atomic species is concatenated with the invariant vector to form the per-atom descriptor

$$\mathbf{d}_i \in R^{n_\text{feat}}. \tag{S16}$$

The resulting descriptor is used as the input to both the short-range energy network and the independent latent-charge prediction network introduced below.

S2 Charge-head construction: supplementary detail

The main text adopts a decoupled and neutrality-enforced charge head over the native LES construction (see Results and Fig. 1 of the main text). This note provides supporting details for that choice on $PbTiO_3$: the two constructions compared, the per-configuration net cell charge that distinguishes them, and the energy, force, and BEC metrics that explain why the residual charge is invisible to standard MLIP error checks.

The two constructions. Both heads sit on the same short-range MTP backbone and are trained on identical data with identical descriptors, optimizer, and loss weights; they differ only in how the latent charges are produced. The native head predicts the latent charges directly from the shared MTP descriptors, as in existing LES integrations. The decoupled head uses an independent charge branch with a per-species bias and projects the mean charge to zero per configuration, enforcing $\sum_i q_i = 0$ exactly.

Net cell charge. Across the 45 BEC configurations, the native head develops a large residual cell charge, $|\sum_i q_i| \approx 12\,e$ per cell, whereas the decoupled head is neutral to within $10^{-5}\,e$. The latent charges in LES are fixed only by the requirement that they reproduce energies and forces, which for a periodic system leaves the total cell charge unconstrained; the reciprocal-space (tinfoil) treatment then absorbs the residual as a uniform compensating background. The result is energetically admissible but inconsistent with the charge-neutral physical system, which is why an explicit neutrality constraint is required rather than optional.

Energies, forces, and BECs are not enough. The residual charge leaves no signature in the standard metrics. Both heads reach essentially identical Born-effective-charge correlation against the PBE-DFPT references ($R = 0.995$ native, $R = 0.994$ decoupled), and neither energy nor force accuracy is degraded by the residual charge. The non-neutrality is therefore exposed only by direct evaluation of the latent charges, not by the parity metrics that would normally be used to validate the model—which is the central reason we report $\sum_i q_i$ explicitly and adopt the decoupled head as the production model.

S3 Full VASP input specification for HfO_2

Table S1: VASP input parameters for the HfO_2 dataset.

Parameter	Value	Description
VASP version	6.4.3	Code version
Functional	PBE	GGA exchange–correlation
PAW potentials	`PAW_PBE Hf` (20 Jan 2003)	Released with VASP 6.4.3
	`PAW_PBE O` (08 Apr 2002)	
`ENCUT`	600 eV	Plane-wave cutoff
`EDIFF`	10^{-8} eV	Electronic convergence
`ALGO`	Normal	Blocked Davidson
`ISMEAR`	0	Gaussian smearing
`SIGMA`	0.1 eV	Smearing width, extrapolated to 0
`ISPIN`	2	Spin polarization on
`LASPH`	.TRUE.	Non-spherical PAW contributions
`ISYM`	0	Symmetry disabled
k-point density	0.040 Å^{-1} per recip. vector	$3 \times 2 \times 2$ for 96-atom cell
k-point grid	Γ-centred Monkhorst–Pack	
`NSW`	0	Static evaluations
`IBRION`	-1	No ionic relaxation
`ISIF`	2	Positions & cell fixed
`LWAVE`	.FALSE.	Wavefunctions not retained
`LCHARG`	.FALSE.	Charge density not retained

S4 Quantitative IR spectrum analysis

Peak positions are determined from the smoothed spectrum maxima; integrated intensities are the area within ± 200 cm^{-1} of each peak after area-normalisation over 0–4000 cm^{-1}.

Table S2: IR peak positions and band intensities for liquid water: MTP–LES vs experiment [3].

Band	MTP–LES position (cm^{-1})	Exp. position (cm^{-1})	MTP–LES intensity	Exp. intensity
Libration	601	676	0.196	0.173
H–O–H bend	1624	1648	0.080	0.046
O–H stretch	3437	3414	0.541	0.632

Intensities are band areas (dimensionless after area normalisation over 0–4000 cm^{-1}). Model spectrum from the dipole-current autocorrelation without quantum correction, consistent with the main-text figure.

The librational position is within 11%, the bending mode within 1.5%, and the O–H stretching position within 1% of experiment. The lower-frequency band intensities are overestimated and the O–H stretching intensity underestimated relative to experiment, consistent with the absence of nuclear quantum effects in classical MD [4, 5].

S5 Per-species BEC RMSE

Table S3: Per-species diagonal and off-diagonal BEC RMSE for water and dipeptides, in elementary charge units.

System and species	Diagonal RMSE (e)	Off-diagonal RMSE (e)
Water — O	0.103	0.081
Water — H	0.087	0.064
Dipeptides — H	0.092	0.087
Dipeptides — C	0.187	0.198
Dipeptides — N	0.256	0.258
Dipeptides — O	0.147	0.212

S6 Dipeptide dipole-loss ablation

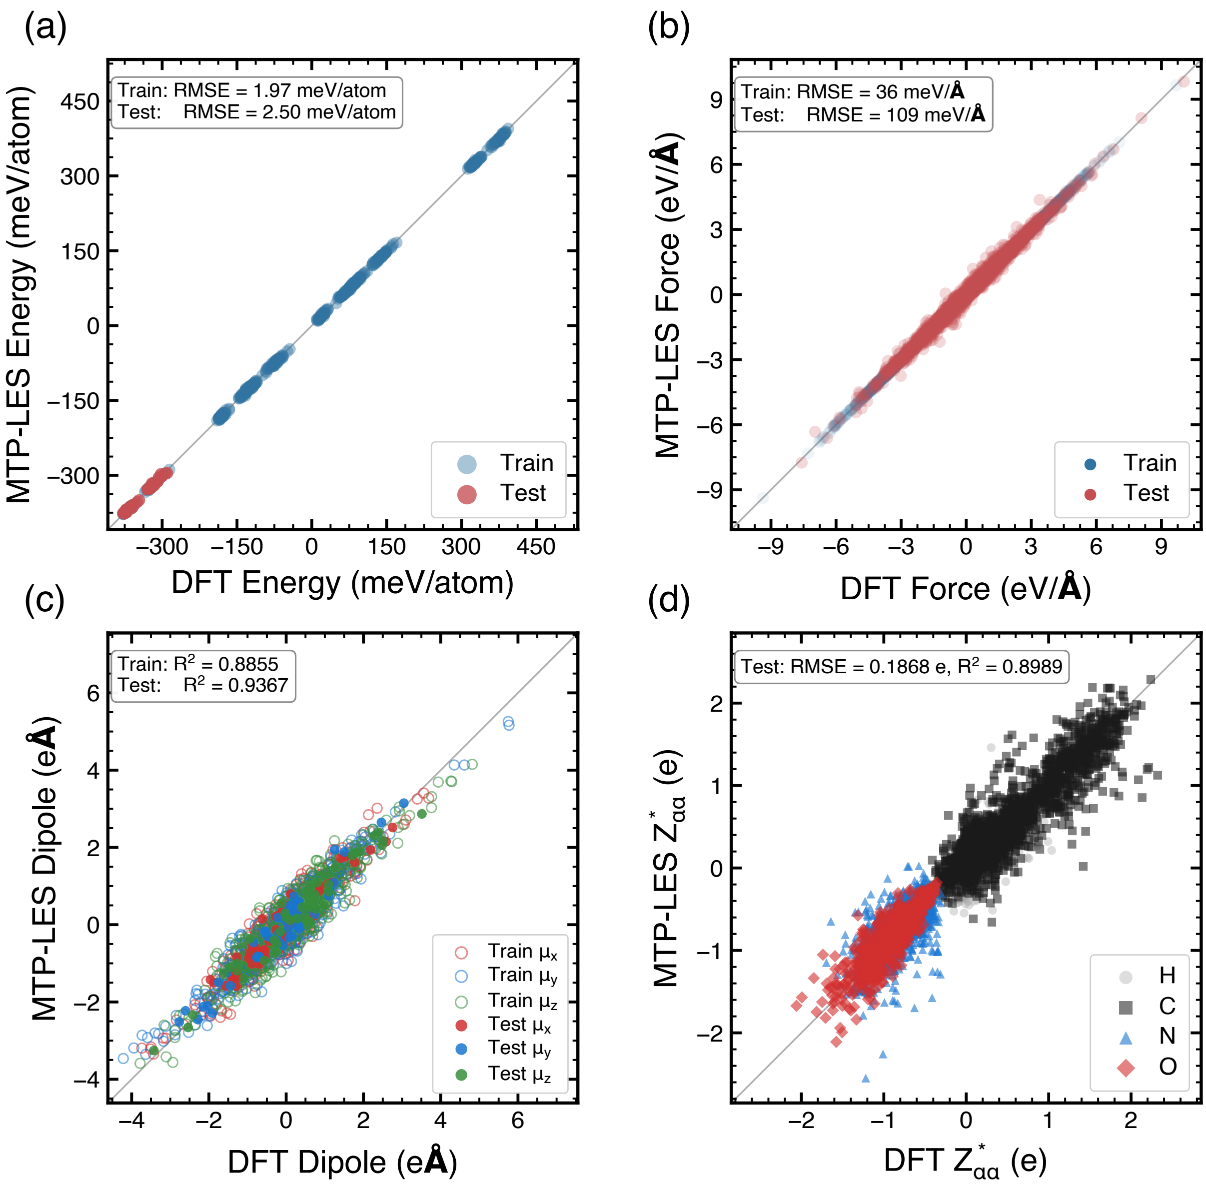

Figure S1: Validation of MTP–LES on dipeptides *without* dipole loss ($w_\mu = 0$). Panels and color coding follow main-text Fig. 7. The increased scatter in the dipole parity (**c**) illustrates the benefit of the optional dipole loss for this finite-molecule dataset.

Table S4: Effect of the optional dipole loss on the dipeptide test set.

Metric	$w_\mu = 0$	$w_\mu = 0.1$
Dipole RMSE (eÅ)	0.307	0.075
Diagonal BEC R^2	0.899	0.931
Diagonal BEC RMSE (e)	0.187	0.154
Off-diagonal BEC RMSE (e)	0.181	0.164

S7 Wall-time benchmark

Table S5: Wall time per MD step for MTP-only and MTP–LES on liquid water, single A100-PCIE-40GB GPU, float32, mean over 1 000 force evaluations.

N (atoms)	MTP-only (ms)	MTP–LES (ms)	Overhead (%)
192	76	83	9.0
384	104	108	3.8
768	120	122	1.3
1 536	168	169	0.5
2 304	217	222	2.5
3 456	282	289	2.6
5 184	397	411	3.6
6 912	484	501	3.4
9 216	613	635	3.5
12 288	777	826	6.3
15 360	996	1 059	6.3

S8 Implementation note: neighbour-list offsets in periodic training

One implementation detail is important for reproducing periodic MTP–LES force training. For periodic systems, the periodic-image offsets in the neighbour list must be preserved as non-differentiable constants during backpropagation. If these offsets are recomputed from positions at each step, the resulting forces can become inconsistent between training and validation because the offsets carry gradients that are not part of the intended force model. We therefore precompute the periodic-image offsets from the reference positions and attach them to the neighbour list as non-differentiable constants. With this treatment, training and validation forces agree to within numerical precision for the periodic benchmarks reported in the main text.

References

[1] Alexander V. Shapeev. Moment tensor potentials: A class of systematically improvable interatomic potentials. *Multiscale Modeling & Simulation*, 14(3):1153–1173, 2016.

[2] Ivan S Novikov, Konstantin Gubaev, Evgeny V Podryabinkin, and Alexander V Shapeev. The mlip package: moment tensor potentials with mpi and active learning. *Machine Learning: Science and Technology*, 2(2):025002, dec 2020.

[3] John E. Bertie and Zhida Lan. Infrared intensities of liquids xx: The intensity of the oh stretching band of liquid water revisited, and the best current values of the optical constants of h2o(l) at 25°c between 15,000 and 1 cm-1. *Appl. Spectrosc.*, 50(8):1047–1057, Aug 1996.

[4] Joseph A. Morrone and Roberto Car. Nuclear quantum effects in water. *Phys. Rev. Lett.*, 101:017801, Jul 2008.

[5] Ondrej Marsalek and Thomas E. Markland. Quantum dynamics and spectroscopy of *Ab Initio* liquid water: The interplay of nuclear and electronic quantum effects. *The Journal of Physical Chemistry Letters*, 8(7):1545–1551, 2017.