Symmetry-Aware Prediction of Electron Localization Functions from Superposed Atomic Densities

Anonymous Author(s)
Affiliation
Address
email

Abstract

The electron localization function (ELF) is a powerful diagnostic of bonding and electronic structure across materials conditions, including the extreme regimes relevant to high-pressure chemistry. However, its direct generation from the chemical formula and crystal structure is very challenging due to its highly non-linear nature. We propose developing a supervised deep learning method that can transform the 3D superposition of atomic densities (SAD) and yield the ELF. The method can naturally incorporate pressure-implicit structural representations and can be used to rapidly score candidate metal sublattices (templates) for compounds under compression. Our approach combines a periodic 3D U-Net with circular padding, an explicit symmetry-pooling layer built from space-group Seitz operators in the local patch frame, and memory-aware training on periodic patches with epoch-wise origin jitter. The model has been trained on 50,000 metal-only structures drawn from a curated subset of Alexandria-MP20, using a 90/10 train/test split. Reproducible and comparable results have been achieved after detailing the representation, symmetry handling, patching strategy, and learning objectives. Our implementation is symmetry-aware at the data and network levels and is designed to scale to large unit cells without significant memory use.

1 Introduction

The electron localization function (ELF), valued for its bounded range and interpretability, provides a powerful assessment of the chemical nature of molecules and compounds [1,2]: regions with ELF close to 1 indicate strong localization (e.g., lone pairs, covalent basins), to 0.5 indicates delocalization similar to a uniform electron gas, and to 0 tends toward nodes [3,4]. A fast, symmetry-respecting predictor of ELF conditioned on crystal structure would thus be a useful "inner loop" for exploring candidate materials across a large chemical space [5,6].

ELF is particularly useful for screening high-pressure compounds, a task hindered by limited data since fewer materials are known under extreme conditions. A prominent example is the family of metal superhydrides, which have been extensively investigated over the past decade for their promise of achieving room-temperature superconductivity [7–10]. The ELF of these superhydrides has been demonstrated to correlate strongly with their superconducting behavior. Moreover, the ELF associated with the metal sublattices provides a measure of the so-called chemical template strength, the essential driving force behind the stability of metal superhydrides. This makes ELF a valuable descriptor for identifying candidate superhydrides, especially those stable at lower pressures with higher critical temperatures. However, ELF is a highly non-linear function whose values vary unpredictably across compounds, making its direct prediction from chemical composition and crystal structure a major challenge.

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

We propose to learn ELF directly from a pressure-implicit representation of the crystal: the superposition of atomic densities (SAD) rasterized on the unit-cell grid [11-13]. Because the SAD is evaluated in the actual lattice (compressed or expanded), pressure enters implicitly via interatomic distances and unit cell volumes, without requiring an explicit pressure scalar as input. The task is then a 3D image transform on a fixed grid size per sample: given SAD, predict ELF on the same grid. To include crystal symmetries and periodicity, we build periodicity into both the data pipeline and the network [14]. At data time, we train on periodic patches with epoch-wise origin jitter and pass patch-local Seitz operators; at model time, we use circular padding everywhere and average features over the patch's symmetry operations via a batched geometric warp [15-18]. This combination enforces invariances that are standard in generative models for crystals such as E(3), while keeping compute tractable [19,20].

2 Background

2.1 Crystal structure representation

Let a crystal unit cell be $M = (A, X, L)$ with $A = \{a_i\}_{i=1}^{N}$ the atom types, $X = [x_1, \dots, x_N]^T \in [0,1)^{N\times3}$ the fractional coordinates, and $L \in R^{3\times3}$ the lattice columns. The infinite crystal is the periodic set $\hat{X} = \{x_i + \sum_{j=1}^{3} k_j e_j \mid k_j \in Z\}$ in fractional space and $L\hat{X}$ in Cartesian space [21]. A superposition of atomic density is a one-channel scalar field defined on the unit-cell volume that approximates the in-cell electron density using an independent-atom ansatz,

$$\rho_{SAD}(r) = \sum_{i=1}^{N} \rho_{Z_i}(\|r - Lx_i\|),$$

where ρ_{Z_i} is a spherically symmetric density associated with atomic number Z_i [22,23]. When evaluated on the actual lattice L, ρ_{SAD} is implicitly pressure-aware: compression changes L and hence interatomic separations, reshaping the superposed density without providing pressure explicitly as a feature. We rasterize ρ_{SAD} on a regular $N_x \times N_y \times N_z$ grid covering the unit cell; the supervised target is the ELF field on the same grid. Although ELF depends on kinetic-energy density and Pauli effects, not just on ρ, the SAD field acts as a physically informed, translation– and rotation-covariant summary of local environments that a sufficiently expressive network can map to the bounded ELF.

2.2 Seitz invariance

Space-group operations are represented in Seitz form $\{R \mid t\}$, where $R \in \mathrm{GL}(3,\mathbb{Z}) \cap \mathrm{O}(3)$ is an integer rotation (orthogonal; the inverse equals the transpose) acting on fractional coordinates and $t \in [0,1)^3$ is a fractional translation [16]. The group action on fractional coordinates is $x \mapsto Rx + t \pmod 1$; the lattice L transforms as $L \mapsto LR^{-1}$ to preserve Cartesian geometry [24]. Periodic $E(3)$ invariance for a crystal-conditioned predictor requires invariance under (i) permutation of atom indices, (ii) translation of all atomic positions, (iii) rigid rotations coupled with the induced lattice transform, and (iv) periodic choices of the unit cell [25-27]. Generative crystal models enforce these symmetries through equivariant backbones; we realize the same principle by averaging features over Seitz operations supplied with each sample and by making all convolutions periodic so the domain is a 3-torus rather than a bounded box [28].

2.3 3D convolutional networks on periodic domains

A 3D convolution with circular padding computes the discrete convolution on the quotient domain $\mathbb{Z}_{N_x} \times \mathbb{Z}_{N_y} \times \mathbb{Z}_{N_z}$, i.e., a torus, which exactly matches unit-cell periodicity [29]. U-Nets, encoder–decoder architectures with skip connections, provide an effective inductive bias for dense field prediction because they aggregate multiscale context while preserving spatial detail through skips [30,31]. In a symmetry-aware variant, intermediate feature maps can be warped by group actions and averaged, yielding features that are invariant (or equivariant, depending on where the averaging is inserted) [32-34]. Our backbone combines these ideas with explicit space-group averaging at input and optionally after each stage, ensuring that the predicted ELF respects the crystal's symmetry class up to the numerical tolerance of interpolation.

3 Model Architecture

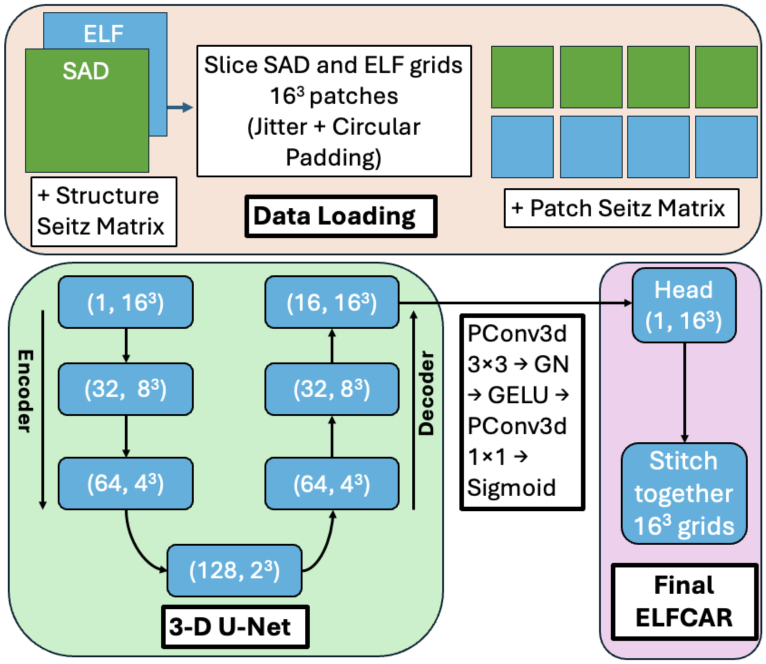

Figure 1: Symmetry-aware pipeline and model for ELF prediction. SAD and ELF volumetric grids are sliced into 16^3 patches with jitter and circular padding; structure- and patch-level Seitz matrices apply crystallographic symmetry for augmentation. Patches are processed by a 3-D U-Net (box labels give #channels and spatial size), and a head PConv3d(3×3) $\rightarrow$ GroupNorm $\rightarrow$ GELU $\rightarrow$ PConv3d(1×1) $\rightarrow$ Sigmoid produces a 1×16^3 prediction. Patch outputs are stitched to reconstruct the full ELFCAR volume.

3.1 Data loading: patches, Seitz operators, and jitter

We train on periodic $p \times p \times p$ patches extracted from SAD/ELF volumes defined on a unit-cell grid of size (N_x, N_y, N_z). Patches are sampled with stride $s \leq p$ using wrap-around indexing so that periodic boundaries are exactly respected. To decorrelate the patch lattice from the crystal grid, each epoch applies an *origin jitter*: a single offset $(o_x, o_y, o_z) \in \{0, \ldots, s-1\}^3$ is added to all patch starts modulo (N_x, N_y, N_z); jitter is disabled for validation to ensure determinism.

Each structure provides space-group symmetry as Seitz operators $\{R \mid t\}$ in fractional coordinates. Because training uses patches, translations are expressed in the *patch frame* whose fractional origin is $o = (i_x/N_x, i_y/N_y, i_z/N_z)$, yielding the transformed translation

$$t' = (R\,o + t - o) \bmod 1, \qquad \{R \mid t\} \mapsto \{R \mid t'\}.$$

The data loader returns per-patch operators $\{R \mid t'\}$ together with the corresponding patch tensors, enabling symmetry-consistent warping and averaging during training.

3.2 3D U-Net backbone with periodic convolutions and symmetry pooling

We use a 3D U-Net whose convolutions are *periodic* so that feature extraction respects lattice periodicity. Encoder stages downsample and decoder stages upsample in the usual U-Net fashion, and the head maps to a single channel with a `Sigmoid` to bound the ELF in $[0, 1]$. To enforce space-group invariance, we interleave feature extraction with a symmetry-averaging layer that uses the per-item Seitz operators (rotations and fractional translations) to sample features under all symmetry-equivalent coordinate transforms and average them. This can be applied to the input and, optionally, after each encoder/decoder stage, yielding representations invariant to the structure's space group while retaining equivariance at the sampling-grid level. Training minimizes a composite objective comprising voxelwise fidelity, periodic-gradient agreement, and value-distribution alignment. Relative weights are learned via uncertainty weighting, and optimization uses AdamW with cosine annealing.

4 Discussion

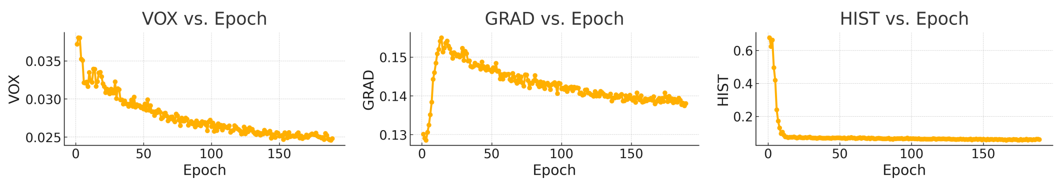

Figure 2: Validation loss decomposed into its three terms across training epochs. Each point is a per-epoch mean computed from step-level logs. The HIST term drops sharply in the early epochs and then stabilizes, VOX declines steadily throughout training, and GRAD shows a modest early rise as the loss weights are learned, followed by a gradual decay to a plateau.

We have demonstrated that a symmetry-aware, periodic 3D U-Net can learn a supervised transformation from a pressure-implicit structural field (SAD) to the electron localization function (ELF) on the unit-cell torus. The model's invariances (periodicity and space-group handling by Seitz pooling), its memory-aware patch training with origin jitter, and a composite loss that couples voxel fidelity, periodic gradients, and value-distribution alignment together yield an operator that is fast enough for inner-loop screening of crystal templates. At the same time, our present training set—metal-only structures from a curated Alexandria-MP20 subset—places deliberate constraints on composition and pressure coverage. Below we outline the near-term steps and longer-term program needed to turn this prototype into a practical high-pressure ELF engine for superhydride discovery and broader materials design.

The next step is to expose the model to compression. Although pressure enters implicitly through the lattice L in SAD, reliable generalization to the extreme compressions of interest requires that the learned mapping see such regimes during training. The most direct next step is to assemble a high-pressure ELF training set by (i) generating families of structures at multiple compressions for each composition/topology and (ii) computing reference ELFs at those volumes. Two practical curricula are natural: (a) isotropic volume sweeps $V/V_0 \in \{1.0, 0.9, 0.8, \ldots\}$ with fixed fractional coordinates, followed by (b) relaxed high-pressure structures including lower symmetry cells to capture the complexity of the chemistry at higher pressures.

The same symmetry-aware pipeline can learn other field-valued operators on the 3-torus by changing the target (e.g. charge density, charge density difference) and optionally augmenting inputs with multi-channel SADs or gradients. Coupled with differentiable stitching and gradient-aware objectives, this suggests a route to fast, lightweight, physically aligned models that can steer closed-loop discovery in high-pressure chemistry and general crystal design, while retaining exact periodicity and explicit space-group handling.

References

[1] Becke, A. D. Edgecombe, K. E. (1990) A simple measure of electron localization in atomic and molecular systems. *The Journal of Chemical Physics* **92**:5397–5403.

[2] Silvi, B. Savin, A. (1994) Classification of chemical bonds based on topological analysis of electron localization functions. *Nature* **371**:683–686.

[3] Savin, A., Nesper, R., Wengert, S. Fässler, T. F. (1997) ELF: The Electron Localization Function. *Angewandte Chemie International Edition* **36**:1808–1832.

[4] Clements, R. J. et al. (n.d.) Electron localisation descriptors in ONETEP. Notes/review, ePrints Soton.

[5] Jørgensen, P. B. Bhowmik, A. (2022) Equivariant graph neural networks for fast electron density estimation of molecules, liquids, and solids. *npj Computational Materials* **8**:183.

[6] Achar, S. K., Bernasconi, L. Johnson, J. K. (2023) Machine Learning Electron Density Prediction Using Weighted Smooth Overlap of Atomic Positions. *Nanomaterials* **13**:1853.

[7] Drozdov, A. P. et al. (2015) Conventional superconductivity at 203,K at high pressures in the sulfur hydride system. *Nature* **525**:73–76.

[8] Drozdov, A. P. et al. (2019) Superconductivity at 250,K in lanthanum hydride under high pressures. *Nature* **569**:528–531.

[9] Sun, Y. Miao, M. (2023) Chemical Templates That Assemble the Metal Superhydrides. *Chem* (Cell Press).

[10] Belli, F. et al. (2021) Strong correlation between electronic bonding network and critical temperature in hydrogen-based superconductors. *Nature Communications* **12**:5381.

[11] Hirshfeld, F. L. (1977) Bonded-atom fragments for describing molecular charge densities. *Theoretica Chimica Acta* **44**(2):129–138.

[12] Lehtola, S. (2019) Assessment of initial guesses for self-consistent field calculations. *Journal of Chemical Theory and Computation* **15**:1593–1604.

[13] Q-Chem 5.x Manual (n.d.) Section 4.4.2: “SCF$_{G}UESS = SAD(Superposition of AtomicDensities).''Softwaredocumentation.$

[14] Kaba, S.-O. Ravanbakhsh, S. (2022) Equivariant Networks for Crystal Structures. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*.

[15] Litvin, D. B. Kopský, V. (2011) Seitz notation for symmetry operations of space groups. *Acta Crystallographica A* **67**:415–418.

[16] Bilbao Crystallographic Server (n.d.) Seitz symbols for Space Groups $R, |, t$. Online resource.

[17] Cohen, T. S. Welling, M. (2016) Group Equivariant Convolutional Networks. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*. PMLR.

[18] Fu, C., Sun, J. Xu, R. (2022) Lattice Convolutional Networks for Learning Ground States of Quantum Many-Body Systems. *arXiv*:2206.07370.

[19] Xie, T. et al. (2021) Crystal Diffusion Variational Autoencoder for Periodic Material Generation. *arXiv*:2110.06197.

[20] Jiao, R. et al. (2023) Crystal Structure Prediction by Joint Equivariant Diffusion. *arXiv*:2309.04475; in *Advances in Neural Information Processing Systems (NeurIPS 2023)*.

[21] Shmueli, U. (ed.) (2008) *International Tables for Crystallography, Volume B: Reciprocal Space* (3rd ed.). Chichester: Wiley / IUCr.

[22] Koritsanszky, T. S. Coppens, P. (2001) Chemical Applications of X-ray Charge-Density Analysis. *Chemical Reviews* **101**:1583–1627. doi:10.1021/cr990112c.

[23] Hirshfeld, F. L. (1977) Bonded-atom fragments for describing molecular charge densities. *Theoretica Chimica Acta* **44**(2):129–138. doi:10.1007/BF00549096.

[24] Wondratschek, H. (2002) *Matrices, Mappings and Crystallographic Symmetry*. IUCr Educational Pamphlet.

[25] Aroyo, M. I. (n.d.) Transformations. Bilbao Crystallographic Server (Crystallography Online).

[26] Zaheer, M. et al. (2017) Deep Sets. In *Advances in Neural Information Processing Systems (NIPS 2017)*.

[27] Batzner, S. et al. (2022) E(3)-equivariant graph neural networks for data-efficient interatomic potentials. *Nature Communications* **13**:2453.

[28] Achour, E. M. et al. (2022) Existence, Stability and Scalability of Orthogonal Convolutional Layers. *Journal of Machine Learning Research* **23**.

[29] Baraniuk, R. G. et al. (2022) 7.05: Discrete Time Circular Convolution and the DTFS. *LibreTexts: Signals Systems*.

[30] Ronneberger, O., Fischer, P. Brox, T. (2015) U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)*; see also *arXiv*:1505.04597.

[31] Drozdzal, M. et al. (2016) The Importance of Skip Connections in Biomedical Image Segmentation. In *MICCAI 2016 Workshops*; see also *arXiv*:1608.04117.

[32] Cohen, T. S. Welling, M. (2016) Group Equivariant Convolutional Networks. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*. PMLR.

[33] Dieleman, S., De Fauw, J. Kavukcuoglu, K. (2016) Exploiting Cyclic Symmetry in Convolutional Neural Networks. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*; see also *arXiv*:1602.02660.

[34] Benton, G., Finzi, M., Izmailov, P. Wilson, A. G. (2020) Learning Invariances in Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS 2020)*.

5 Appendix

File discovery and I/O. The loader discovers triplets of `*.npy` files with matching stems: `stem_sad`, `stem_elf`, `stem_sym`. Arrays are opened via memory-mapped NumPy for low-overhead header reads and I/O. The grid shape (N_x, N_y, N_z) is read from the ELF file header and stored per sample to support consistent patch extraction and rescaling across structures.

Patch extraction and channels. The core dataset class emits periodic $p \times p \times p$ patches with optional overlap controlled by a stride $s \leq p$. Given integer starts (i_x, i_y, i_z), a patch is cut out using `np.take(..., mode="wrap")` along each axis, which is equivalent to modular indexing and exactly enforces periodic boundaries. The input tensor stacks channels in the specified order; by default, SAD and ELF are concatenated to form $\mathbf{X} \in \mathbb{R}^{C \times p \times p \times p}$ with $C \in \{1, 2\}$ depending on whether the target (ELF) is carried through the loader for supervised training.

Epoch-wise jitter and samplers. To avoid pathological alignment between the patch lattice and the crystal grid, the dataset implements epoch-wise origin jitter. Let $(o_x, o_y, o_z) \in \{0, \ldots, s-1\}^3$ be a random offset drawn once per epoch from a deterministic seed; patch starts are then

$$(i_x s + o_x,\ i_y s + o_y,\ i_z s + o_z) \mod (N_x, N_y, N_z).$$

A dataset hook `set_epoch(e)` sets this offset deterministically from the seed and epoch number. Training samplers call this hook once per epoch in both single-GPU and distributed settings: `_RandomSamplerWithEpoch` subclasses `RandomSampler` to increment an internal epoch counter on each `__iter__`, and `_DistributedSamplerWithJitter` forwards the framework's `set_epoch(e)` to the dataset, preserving DDP semantics and data partitioning. Validation uses the same stride with jitter disabled.

Symmetry bookkeeping and batching. Space-group symmetries are provided per structure as Seitz operators $\{R \mid t\}$ in fractional coordinates, stored as an $(R, 4, 4)$ array with integer rotation blocks and fractional translations. Because training uses patches, these global operators must be expressed in the patch frame whose fractional origin is $o = (i_x/N_x, i_y/N_y, i_z/N_z)$. The translational part is shifted via

$$t' = (R\,o + t - o) \mod 1, \qquad \{R \mid t\} \mapsto \{R \mid t'\}.$$

For each item, the dataset returns both the global operators and the patch-frame operators. A dedicated `collate_patches` routine pads ragged symmetry lists in a batch to the maximum group size and returns

$$\big(X,\ \texttt{sym\_batch},\ \texttt{mask},\ \texttt{origin\_frac},\ \texttt{orig\_shape},\ \texttt{stems}\big),$$

where `sym_batch`$\in \mathbb{R}^{B \times R_{\max} \times 4 \times 4}$ and `mask`$\in \{0,1\}^{B \times R_{\max}}$ preserve per-item group sizes for downstream weighting and symmetry-averaged computation.

Periodic convolutions and U-Net wiring. The predictor f_θ is a 3D U-Net whose every convolution is periodic. A `PeriodicConv3d` layer performs circular padding of width $(k-1)/2$ for an odd kernel k and then applies a standard `Conv3d` with zero explicit padding. Residual blocks comprise

`PeriodicConv3d` $\to$ `GroupNorm` $\to$ `GELU` $\to$ (`Dropout3d`) $\to$ `PeriodicConv3d` $\to$ `GroupNorm`, with a residual connection and final `GELU`, providing stable training at depth. Downsampling uses a stride-2 periodic convolution followed by `GroupNorm` and `GELU`; upsampling uses trilinear interpolation followed by periodic convolution, concatenation with the corresponding encoder skip, further residual processing, and a periodic "merge" convolution. The head comprises a periodic $3 \times 3 \times 3$ convolution, normalization, `GELU`, a $1 \times 1 \times 1$ periodic projection to one channel, and a `Sigmoid` to bound the ELF prediction in $[0, 1]$.

Symmetry pooling (`SymmAvg3D`). Let $f \in \mathbb{R}^{B\times C\times D\times H\times W}$ be a feature map on the patch grid and $S = \{(R, t')\}$ the per-item Seitz operators in the patch frame returned by the loader. The layer constructs base sampling coordinates $\rho_{out} \in [-1, 1]^3$ (and the associated fractional grid $\zeta_{out} = 12(\rho_{out} + 1) \in [0, 1]^3$) and, for each operator, computes the input sampling grid by

$$\zeta_{in} = \mathrm{wrap}\big(S^{-1}R^{-1}S\,\zeta_{out} - S^{-1}R^{-1}t'\big)\,, \quad \rho_{in} = 2\zeta_{in} - 1,$$

where $S = \mathrm{diag}(W/N_x,\, H/N_y,\, D/N_z)$ rescales between patch-voxel and fractional coordinates and wrap projects component-wise to $[0, 1)$. Because R is orthogonal with integer entries, $R^{-1} = R^\top$ is implemented by transpose for numerical stability. The network samples f at ρ_{in} using `grid_sample` in trilinear mode and averages across valid operations with a per-item mask, yielding a symmetry-averaged feature map with the same shape as f. The layer is applied to the input SAD and, optionally, after every encoder and decoder stage.

Training objective and optimization. Let $\hat{y} = f_\theta(x)$ and y be predicted and true ELF patches. We use a voxelwise Smooth-L1 loss L_{vox}, a periodic-gradient loss

$$L_\nabla = \|\Delta_x\hat{y} - \Delta_x y\|_1 + \|\Delta_y\hat{y} - \Delta_y y\|_1 + \|\Delta_z\hat{y} - \Delta_z y\|_1,$$

where forward differences use circular shifts along each axis, and a soft-histogram KL that compares Gaussian-smoothed marginal histograms of ELF values in $[0, 1]$. Learnable log-precisions η_k balance the terms,

$$L(\theta) = \sum_{k\in\{vox,\nabla,hist\}} e^{-\eta_k}L_k + \eta_k,$$

removing the need for manual loss-weight tuning during training. Optimization uses AdamW with cosine annealing; the module is implemented with PyTorch Lightning for reproducibility.