Research Article | Open Access

Journal of Materials Informatics

Ellis *et al. J. Mater. Inf.* **2026**, *6*, 45 **DOI:**10.20517/jmi.2026.44

Prediction of electron localization functions from superposed atomic densities for accelerated superhydride discovery

Austin Ellis*, **Samantha Ellis**, **Maosheng Miao***

Keywords:
Electron localization function, materials informatics, deep learning, high-pressure materials, superhydrides, predictive modeling

Citation: Ellis, A.; Ellis, S.; Miao, M. Prediction of electron localization functions from superposed atomic densities for accelerated superhydride discovery. *J. Mater. Inf.* **2026**, *6*, 45. https://dx.doi.org/10.20517/jmi.2026.44

Received: 25 Jun 2026
First Decision: 9 Jul 2026
Revised: 30 Jul 2026
Accepted: 17 Aug 2026
Published: 16 Sep 2026

Academic Editor:
Hao Li
Copy Editor:
Pei-Yun Wang
Production Editor:
Pei-Yun Wang

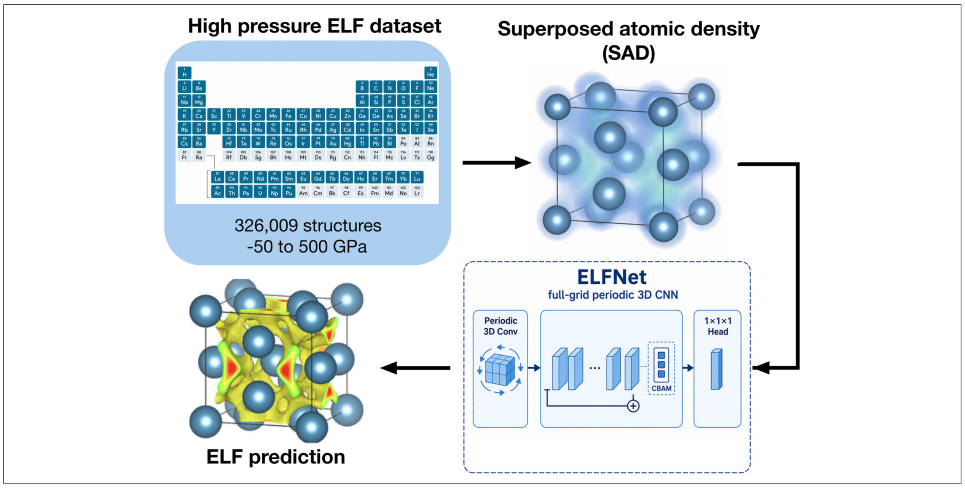

Abstract

The electron localization function (ELF) provides a chemically interpretable representation of bonding and interstitial electron localization, but generating volumetric ELF fields with density functional theory (DFT) is computationally expensive for large candidate libraries. We present ELFNet, a structure-to-field model that predicts DFT-derived ELF grids from computationally inexpensive superposed atomic density (SAD) inputs for high-pressure chemical-template screening. To train the model, we constructed a pressure-targeted Perdew–Burke–Ernzerhof (PBE) ELF dataset containing 326,009 converged unary and binary prototype structures spanning 89 elements and covering pressures from approximately -50 to 500 GPa. ELFNet uses periodic full-grid convolutions and a composite objective designed to recover sparse high-ELF features. We evaluated the model on 50,000 metal-only templates derived from known high-pressure superhydride prototypes. Predicted fields captured chemically meaningful differences in the spatial topology of interstitial localization, while global maximum ELF values and maximum-location errors provided complementary measures of agreement with DFT references. On the reported hardware, the full inference-through-write workflow reduced summed per-structure elapsed time by approximately 350-fold relative to fixed-geometry Vienna Ab initio Simulation Package (VASP) ELF generation. These results establish SAD-to-ELF

Department of Chemistry and Biochemistry, California State University, Northridge, CA 91330, USA.

***Correspondence to:** Austin Ellis, Prof. Maosheng Miao, Department of Chemistry and Biochemistry, California State University, Northridge, CA 91330, USA. E-mail: ellisal@ornl.gov; mmiao@csun.edu

prediction as a low-cost front-end screening method for prioritizing candidate metal templates before downstream first-principles validation.

INTRODUCTION

Electronic-structure calculations reveal bonding and localization patterns that energies alone do not capture. The electron localization function (ELF) represents these patterns as chemically interpretable three-dimensional maps. Introduced by Becke and Edgecombe, ELF is a dimensionless real-space descriptor of electron localization based on the local behavior of same-spin electron pairs[1]. ELF is now widely used to identify chemically meaningful features such as bonding and lone-pair basins, atomic shell structure, and interstitial electron localization in molecules and solids[2,3]. For each spin component σ, ELF may be written schematically as

$$\mathrm{ELF}_\sigma(\boldsymbol{r}) = \frac{1}{1+\chi_\sigma(\boldsymbol{r})^2}, \qquad \chi_\sigma(\boldsymbol{r}) = \frac{D_\sigma(\boldsymbol{r})}{D^0_\sigma(\boldsymbol{r})}, \tag{1}$$

where D_σ measures the excess local kinetic-energy density associated with Pauli repulsion and D_σ^0 is the corresponding reference for a homogeneous electron gas of the same spin density. ELF values range from 0 to 1, with values close to 1 indicating strongly localized electron-pair regions, values near 0.5 corresponding to behavior similar to a homogeneous electron gas, and values near 0 indicating suppressed localization. This representation is particularly valuable when conventional bonding descriptions are incomplete because it reveals localized regions that are not readily classified as two-center bonds, lone pairs, or formal valence assignments.

High-pressure chemistry is one such regime because compression can alter orbital ordering, bonding, reactivity, and the set of stable stoichiometries[4]. Hydrogen-rich compounds are a prominent example. Since Ashcroft's proposal that hydrogen-dominant metallic alloys could host high-temperature superconductivity[5], experimental and theoretical studies have identified hydrides with exceptionally high superconducting transition temperatures, including sulfur hydrides and lanthanum hydrides at megabar pressures[6-11]. Recent work has broadened this landscape through calcium clathrate and superhydrides, alloy superhydrides, and molecular hydrides that retain high transition temperatures or reduce stabilization pressures[12-21]. Dense hydrogen networks can provide the high phonon frequencies and strong electron–phonon coupling required for such transition temperatures. Nevertheless, discovering new high-pressure hydrides remains challenging because many metal frameworks, hydrogen arrangements, compositions, and pressures must be evaluated before promising candidates can be identified.

The high-pressure hydride search space can be reduced through chemical template theory. In metal superhydrides, stability is not governed exclusively by direct metal–hydrogen bonding. Chemical template theory instead emphasizes the metal sublattice, which can host localized interstitial electronic states described as quasi-atom orbitals. The associated ELF maxima coincide spatially with interstitial sites that are subsequently occupied by hydrogen networks. Previous work has shown that these localization patterns can template hydrogen-network assembly and rationalize structural trends across high-pressure hydride families[22]. ELF therefore provides a descriptor of template strength and of the electronic topology that favors hydrogen clathrates and related motifs. Related work has also shown that ELF-derived bonding-network descriptors correlate with superconducting critical temperatures in hydrogen-based superconductors[23].

Chemical template theory has recently been combined with machine learning for superhydride discovery. Previous work integrated chemical-template features into an ML-guided screening workflow and reported the batch discovery of complex metal superhydrides, including non-integer hydrogen-to-metal ratios, new

structural prototypes, and large primitive cells[24]. That study illustrates both the value and the computational bottleneck of ELF-based screening. Crystal-structure searches conventionally generate and relax large candidate pools[25,26], while recent machine-learning-guided, active-learning, and generative workflows can accelerate candidate generation, relaxation, and triage before first-principles validation[27-31]. As candidate cells grow, sometimes exceeding 50 atoms per primitive cell in non-stoichiometric superhydrides, direct density functional theory (DFT) calculations with ELF output remain expensive, and searches over thousands of hypothetical metal templates can devote substantial computational effort to generating ELFs for structures that are ultimately rejected.

These costs motivate a structure-to-field surrogate for ELF. The objective is to predict the ELF grid directly from structural information without performing a new electronic-structure calculation. In our workflow, the crystal structure and atomic species are converted into a computationally inexpensive superposed atomic density (SAD) grid, which ELFNet maps to the corresponding DFT-derived ELF grid. The model preserves periodicity, represents compression through the lattice geometry, and recovers the sparse high-value ELF maxima most relevant to template screening. Rather than replacing DFT, ELFNet supports a two-stage workflow in which predicted ELF fields rank many candidate metal templates before full DFT calculations are reserved for the most promising structures.

Recent machine-learning methods have predicted three-dimensional electron-density fields for molecules and periodic materials from atomic geometry or inexpensive initial electronic information, including grid-based and equivariant architectures that can accelerate or initialize electronic-structure calculations[32-36]. More directly related to ELF screening, machine-learning-predicted maximum ELF values have been used to discover ternary electrides from large hypothetical candidate libraries[37]. ELFNet differs by targeting complete periodic ELF fields across a broad, pressure-dependent 89-element dataset and evaluating them for high-pressure metal-template screening.

This study makes two contributions toward that goal. First, we introduce a Perdew–Burke–Ernzerhof (PBE) ELF dataset developed specifically for learning pressure-dependent localization fields. The dataset contains unary and binary prototype structures with broad element coverage, including noble gases and f-block elements that are often sparse or absent in general-purpose materials datasets. Although these elements may appear chemically inert under ambient conditions, high pressure can activate noble-gas chemistry and stabilize compounds with unusual stoichiometries and oxidation states[38,39]. Second, we present ELFNet, a full-grid residual three-dimensional convolutional network that maps SAD grids to ELF grids. The model uses a FlatResNet3D backbone, periodic circular-padding convolutions, and a peak-aware loss. In a case study involving 50,000 templates, the resulting workflow retained chemically meaningful localization features while reducing summed per-structure elapsed time by approximately 350-fold on the reported hardware relative to fixed-geometry VASP ELF generation.

MATERIALS AND METHODS

High-pressure ELF dataset

Design goals

The dataset was designed to learn the pressure-dependent mapping from structure-derived SAD fields to DFT-derived ELF fields for compact high-pressure prototype cells. Recent general-purpose potential-energy-surface resources, including MatPES, MP-ALOE, the Materials Project trajectory dataset introduced with CHGNet, and OMat24, provide broad energy, force, and stress labels for training machine-learning interatomic potentials[40-43]. Universal and foundational interatomic-potential architectures trained across broad chemical spaces include M3GNet, CHGNet, DPA-2, NequIP, Allegro, MACE, the MACE foundation

model, and GRACE[42,44-50]. In contrast, the dataset developed in this work was assembled around volumetric ELF output. Energies, forces, and stresses were retained for quality control and possible future use, but the learning task uses the SAD arrays as inputs and the ELF arrays as targets.

The dataset construction was guided by three requirements. First, the structures had to sample a broad range of pressures rather than only equilibrium or near-equilibrium volumes. Second, the enumeration had to include small prototype cells that could be systematically populated across many elements and element pairs. Third, the coverage had to include elements that may be underrepresented in common materials datasets but are chemically relevant under pressure, including noble gases and f-block elements.

Prototype families and element coverage

Unary structures were generated from six compact prototype families: simple cubic (SC, $Pm\bar{3}m$), body-centered cubic (BCC, $Im\bar{3}m$), face-centered cubic (FCC, $Fm\bar{3}m$), hexagonal close-packed (HCP, $P6_3/mmc$), diamond cubic (DC, $Fd\bar{3}m$), and a hexagonal family represented in the workflow by $P6/mmm$. Binary AB structures were generated from seven prototype families: CsCl ($Pm\bar{3}m$), HgS ($P3_221$), NaCl ($Fm\bar{3}m$), NiAs ($P6_3/mmc$), PbO ($P4/nmm$), zinc blende (ZB, $F4\bar{3}m$), and wurtzite (ZW, $P6_3mc$). These compact families were chosen to span common unary and binary coordination motifs while keeping the enumeration tractable.

The final high-pressure subset contains the following 89 elements: H through Bi, the lanthanides La–Lu, and the actinides Ac, Th, Pa, U, Np, and Pu. The AB subset covers all 3,916 unordered canonical element pairs overall, with nearly complete coverage within each binary prototype. This element coverage was chosen to support exploratory high-pressure screening beyond the chemical space that is most common in ambient-pressure materials datasets. Under compression, noble gases and rare-earth or actinide elements can participate in unusual bonding, insertion, or charge-transfer behavior. For example, high pressure stabilizes xenon oxides with mixed oxidation states and the helium-containing compound Na_2He[38,39]. Figure 1 summarizes the resulting element coverage and frequency.

The full 89-element inventory is listed in Supplementary Table 1.

Pressure-targeted lattice expansion

To sample compression more systematically, we used the converged fixed-lattice prototype calculations to construct a pressure-targeted lattice-multiplier grid. Within each structure/prototype/chemical-system group, final DFT pressures from the original lattice-scale calculations were summarized by the median at each scale and used to fit a monotonic non-increasing pressure–scale relation using isotonic regression[51]. The fitted relation was then inverted to select additional lattice-scale multipliers within the usable 0-500 GPa pressure interval. This procedure produced 164,133 converged pressure-targeted multiplier structures in the final dataset; details of the fitting, interpolation, target-pressure selection, and multiplier accounting are provided in Supplementary Note 1 and Supplementary Table 2.

DFT settings and final selection

All final structures were obtained from PBE calculations[52,53]. MatPES-compatible PBE input settings and PAW potentials were used throughout[54-56]. PBE was selected instead of r^{2}SCAN to maintain consistency with the MatPES-compatible high-throughput workflow and to make broad pressure sampling computationally tractable. Although r^{2}SCAN can improve semilocal descriptions across diverse bonding environments[57], the objective was broad pressure coverage with ELF output rather than a high-fidelity equilibrium-energy benchmark. Functional sensitivity was therefore evaluated empirically using the pressure-balanced PBE–r^{2}SCAN cross-check described below.

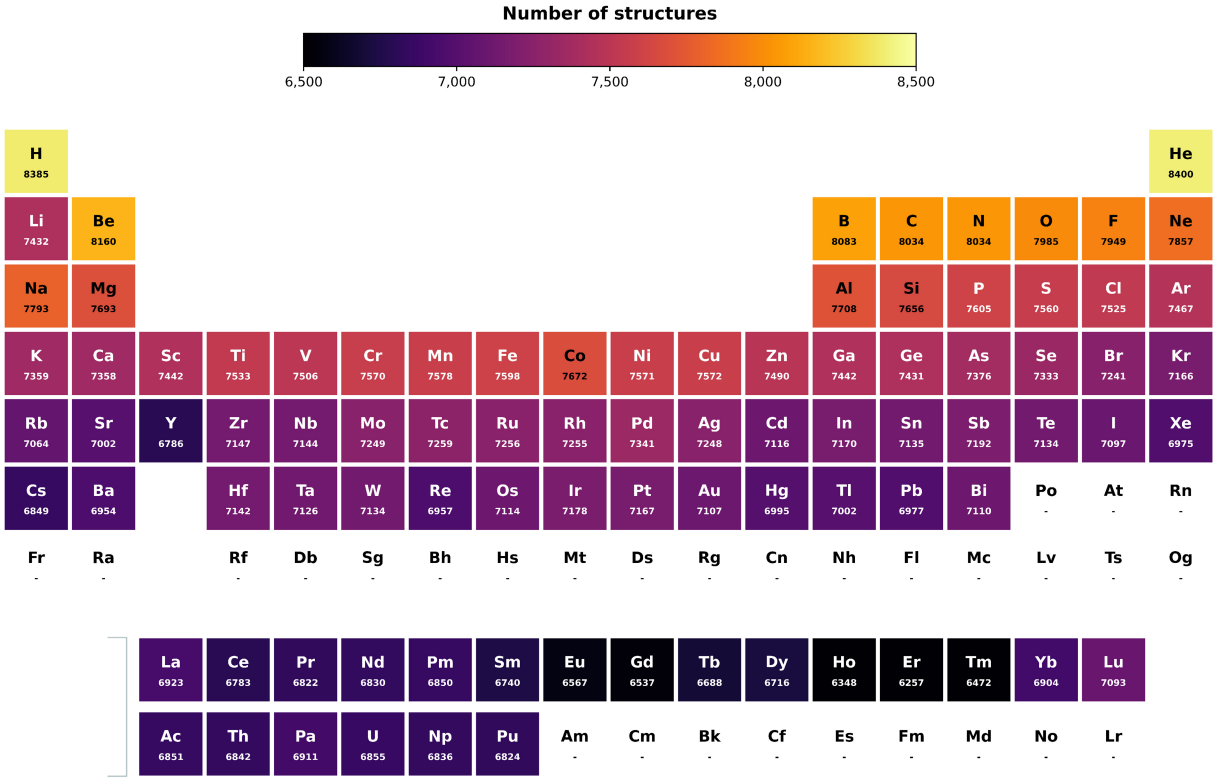

Figure 1. Element coverage in the final high-pressure PBE ELF dataset. Each periodic-table tile reports the number of structures containing that element; color encodes the same count on a logarithmic scale. The white color indicates zero coverage for that element. PBE: Perdew–Burke–Ernzerhof; ELF: electron localization function.

In a pressure-balanced 96-structure cross-check (16 per pressure bin), changing the reference from PBE to r^{2}SCAN increased the median voxel mean absolute error (MAE) only from 0.09871 to 0.10493 and the median maximum-location distance from 1.109 to 1.183 Å. ELFNet–r^{2}SCAN screening agreement was 92.71%, 78.13%, and 87.50% at ELF thresholds of 0.75, 0.80, and 0.90, respectively; PBE–r^{2}SCAN agreement was at least 95.83% at every threshold. These results indicate that ELFNet retains useful screening information across the change in functionality. Supplementary Table 3 provides additional summary error metrics for this cross-check.

The final dataset contains 326,009 structures, including 7,423 unary and 318,586 binary structures. It is nearly balanced between the fixed absolute-lattice and pressure-targeted multiplier generation routes. The achieved pressures span -49.983 to 499.997 GPa, with a median of 83.26 GPa and a mean of 149.12 GPa. The median force magnitude is zero because many one-site and two-site high-symmetry prototypes have symmetry-cancelled forces; this does not imply zero stress or zero electronic response. Additional dataset statistics are provided in Supplementary Table 4, and the distributions of cohesive energy, force magnitude, and achieved pressure are summarized in Figure 2.

Negative achieved pressures arise only from expanded fixed-lattice configurations and represent DFT tensile-stress states.

SAD representation and learning target

The model input is an SAD field. SAD provides a computationally inexpensive, structure-derived representation of the crystal by summing tabulated neutral spherical atomic densities over all atoms in the periodic unit cell[58]. Consider a periodic crystal represented as $C = (\{Z_i\}_{i=1}^{N}, \{\boldsymbol{x}_i\}_{i=1}^{N}, L)$ where Z_i is the atomic

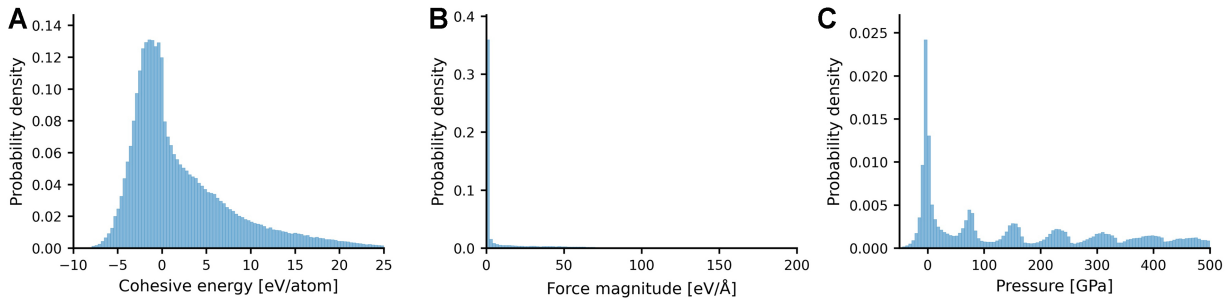

Figure 2. Distribution summary for the final high-pressure PBE ELF dataset. The panels report (A) cohesive energy per atom, (B) force magnitude, and (C) achieved density-functional-theory pressure. PBE: Perdew–Burke–Ernzerhof; ELF: electron localization function.

number, $\boldsymbol{x}_i \in [0,1)^3$ is the fractional coordinate, and $L \in \mathbb{R}^{3\times3}$ is the lattice matrix. Grid operations are performed with periodic boundary conditions, so distances are evaluated using minimum-image wrapping across the unit-cell boundaries.

In continuous notation, the SAD field is

$$\rho_{\mathrm{SAD}}(\boldsymbol{r}) = \sum_{i=1}^{N} q_{Z_i}\, \rho^{0}_{Z_i}\left(d_{\mathrm{per}}(\boldsymbol{r}, L\boldsymbol{x}_i)\right), \tag{2}$$

where $\rho_{Z_i}^{0}$ is a tabulated neutral density for element Z_i, q_{Z_i} is the configured valence scaling used in this work, and d_{per} is the minimum-image periodic distance. In the inference implementation, each atom contributes to the grid through radial interpolation of the neutral density table, with contributions truncated once the radial density falls below a fixed threshold.

Pressure enters this representation implicitly through the lattice geometry. No pressure scalar is supplied to the neural network. Instead, the same fractional coordinates evaluated in a compressed lattice produce different real-space distances, different density overlaps, and different SAD values on the grid. This distinction matters for high-pressure screening because a candidate template at 200 GPa and the same idealized template at 0 GPa produce distinct inputs with different lattice geometries and SAD fields.

The supervised target is the DFT-derived ELF field on the same unit-cell grid. During training and inference, SAD and ELF arrays are handled as paired full-grid tensors with identical shapes. During inference from a starting structure, the model estimates the grid dimensions from the lattice, evaluates the tabulated neutral densities on that periodic grid, and predicts the corresponding ELF tensor.

ELFNet architecture

Full-grid FlatResNet3D

ELFNet is a full-grid neural operator

$$\begin{aligned} f_\theta : \mathbb{R}^{1\times D\times H\times W} &\to [0,1]^{1\times D\times H\times W}, \\ \rho_{\mathrm{SAD}} &\longmapsto \widehat{y} = f_\theta(\rho_{\mathrm{SAD}}), \end{aligned} \tag{3}$$

where $\widehat{y}$ is the predicted ELF field. It processes the complete unit-cell SAD grid in one same-resolution forward pass and returns one sigmoid-bounded ELF grid.

The production model uses one input SAD channel, one output ELF channel, a base width of 32, 16 same-resolution flat residual blocks, a kernel size of five, and three-dimensional convolutional block attention module (CBAM) after every four residual blocks. The stem applies a circularly padded 3 × 3 × 3 convolution,

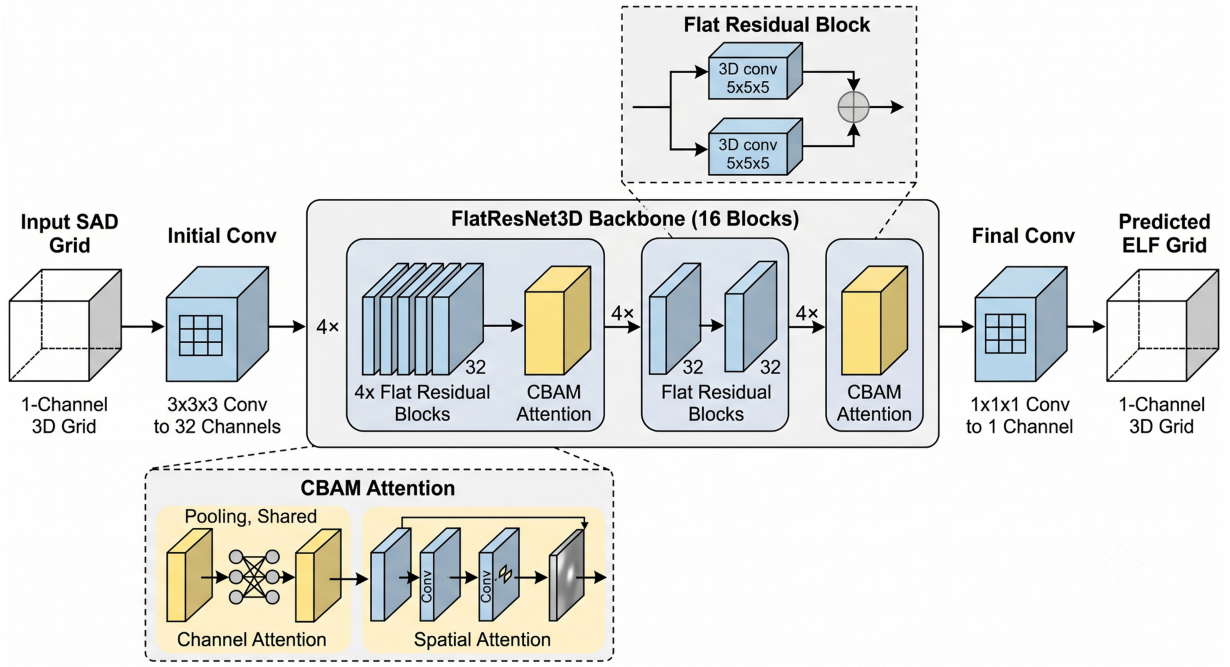

Figure 3. ELFNet architecture. A one-channel SAD grid is transformed at its native spatial resolution. The stem uses a periodic 3 × 3 × 3 convolution to produce 32 channels. Four stages, each containing four flat residual blocks followed by a three-dimensional CBAM, form the 16-block body. A post-body periodic 5 × 5 × 5 convolution is combined with a long stem skip connection, and a 1 × 1 × 1 head with sigmoid activation produces the ELF grid. Each residual block uses GroupNorm, GELU activation, and SE. SAD: Superposed atomic density; CBAM: convolutional block attention module; ELF: electron localization function; GELU: Gaussian error linear unit; SE: squeeze-excitation.

GroupNorm, and Gaussian error linear unit (GELU) activation. Each residual block contains two circularly padded 5 × 5 × 5 convolutions with GroupNorm, GELU, and squeeze-excitation (SE). A CBAM module follows blocks 4, 8, 12, and 16[59-63].

After the residual body, a further circularly padded 5 × 5 × 5 convolution and GroupNorm are combined with a long skip connection from the stem, followed by GELU activation. A 1 × 1 × 1 convolution and sigmoid activation then produce the one-channel ELF field. The bundled checkpoint contains 4,232,989 trainable parameters. This same-grid design preserves the image-to-image character of the SAD-to-ELF mapping while avoiding the memory overhead of multi-resolution feature maps and patch reconstruction. Figure 3 summarizes the implemented data flow.

Training data loading

The training loader yields complete unit-cell SAD/ELF pairs. Samples are bucketed by grid shape to avoid unnecessary padding, so the production model trains on native full-grid arrays rather than extracted patches.

Training procedure

The 326,009 paired full-grid SAD/ELF samples were divided deterministically at the sample level into 309,709 training samples and 16,300 validation samples using a validation fraction of 0.05 and a random seed of 42. The independent 50,000-structure hydride-derived template set was not used for training, validation, checkpoint selection, or hyperparameter tuning.

The production model was trained for 80 epochs using the AdamW optimizer with a maximum learning rate of 1×10^{-4}, a weight decay of 1×10^{-4}, and a OneCycleLR learning-rate schedule. The batch size was 16 per distributed process, with two gradient-accumulation steps and gradient-norm clipping at 1.0. Training used

mixed bfloat16 precision on eight distributed ranks across two compute nodes, each equipped with four NVIDIA A100 80 GB GPUs.

The full bfloat16 production run completed all epochs and optimizer steps with finite training and validation losses, no NaN or Inf events, and finite model and optimizer tensors in the final checkpoint.

Why explicit symmetry is not enforced

ELFNet does not explicitly enforce point-group or space-group symmetry beyond the periodic boundary conditions built into the convolutional layers. Recent equivariant atomistic and electronic-field models demonstrate how rotational covariance can be incorporated directly into learned three-dimensional representations[32,34,35,46]. In principle, crystallographic symmetry could also be incorporated by averaging predictions or intermediate feature fields over space-group operations written in Seitz form, $g = \{R|\boldsymbol{t}\}$, where R is a rotation and $\boldsymbol{t}$ is a fractional translation. For a grid-based convolutional neural network (CNN), however, this would require resampling feature tensors under multiple symmetry operations and averaging the transformed copies. This would increase memory use and wall time, especially for full-grid ELF learning in high-symmetry cells.

This choice is consistent with the intended screening task. ELFNet is designed to identify candidate templates before more expensive DFT calculations, not to replace final electronic-structure analysis. Circular padding ensures that local neighborhoods wrap across unit-cell boundaries, preserving periodicity in the learned representation. For template screening, the most important requirement is to recover the presence, approximate magnitude, and spatial location of high-ELF maxima. A prediction with small symmetry-breaking artifacts can still be useful if it identifies the same interstitial localization basins that would motivate follow-up DFT calculations. Explicit symmetry averaging remains a useful future extension, but it is omitted here to prioritize memory efficiency and full-grid throughput.

Across 117 $Fm\bar{3}m$ structures, symmetry averaging produced a raw-to-averaged field MAE of 0.01672 and reduced the median DFT-referenced maximum-location error by 0.185 Å. The 0.640 Å single-voxel argmax change reflects switching among near-degenerate maxima rather than physical displacement of an ELF feature. Small field and peak-amplitude changes indicate that raw predictions are already close to their symmetrized counterparts, but averaging is advisable when exact space-group consistency is required. Results are reported in Supplementary Table 5.

Training objective

ELF prediction is a strongly imbalanced field-regression problem. Most voxels in a unit cell have small or moderate ELF values, while a small fraction of localized basins or interstitial maxima carries much of the chemical information needed for template screening. A uniform voxel-reconstruction objective can therefore produce smooth fields with low average error while suppressing rare high-ELF features.

ELFNet uses a composite objective designed to balance global reconstruction quality with recovery of chemically relevant maxima. The voxel term, $\mathcal{L}_{\mathrm{vox}}$, is an L1 reconstruction loss between the predicted and reference ELF values. The gradient term, $\mathcal{L}_{\mathrm{grad}}$, compares periodic finite differences along the three grid axes using an L1 loss, which helps preserve sharp features and reduce displaced basins. The distribution term, $\mathcal{L}_{\mathrm{cdf}}$, compares sorted predicted and reference ELF values using an L1 loss. The adaptive peak term, $\mathcal{L}_{\mathrm{peak}}$, compares soft peak-location distributions, matches ELF values at the strongest target voxels, and penalizes spurious high predicted peaks.

The four components are combined with learned Kendall-style uncertainty weights[64].

For $k \in \{\text{vox, grad, cdf, peak}\}$, the model learns a scalar log-variance parameter η_k and minimizes

$$\mathcal{L} = \sum_{k\in\{\text{vox,grad,cdf,peak}\}} [\exp(-\eta_k)\mathcal{L}_k + \eta_k] . \quad (4)$$

This objective is aligned with the template-screening metrics because it targets both the value and real-space location of high-ELF features. Training-loss curves for the production checkpoint are shown in Supplementary Figure 1.

RESULTS AND DISCUSSION

A plain 3D U-Net trained on the same split and compute budget contained 4.94% more parameters than ELFNet. Under the same non-peak objective, it achieved a 71.840% joint pass rate, whereas the FlatResNet3D backbone achieved 76.432% with lower median maximum-value, maximum-location, and voxel errors. The complete ELFNet model yielded the strongest peak-sensitive performance. Results are reported in Supplementary Table 6.

In otherwise matched ablations, the full composite loss increased the joint pass rate by 15.99 percentage points and reduced the median maximum-location error by 0.270 Å relative to the voxel plus gradient plus CDF objective. Individual voxel and gradient objectives could favor their corresponding fieldwise metrics but did not reproduce peak-sensitive acceptance. Removing CBAM degraded all primary metrics, and replacing circular padding with zero padding produced the largest architecture-related degradation. Results are reported in Supplementary Tables 6 and 7.

Qualitative template classification from predicted ELF fields

We first evaluated whether ELFNet preserves the spatial interstitial-localization patterns required for chemical-template screening before applying scalar metrics. Figure 4 compares predicted and DFT reference ELF fields for two representative templates. Although the reported global maximum for AuGa is numerically larger, that maximum does not form a pronounced basin in an open interstitial region and the field is instead fragmented or concentrated near atomic regions. Ca exhibits a clear ELF basin within an open interstitial region. The relevant screening signal is therefore the location and spatial topology of localization, not the scalar global maximum alone. Templates lacking a pronounced open-interstitial basin can be deprioritized, while those exhibiting such a basin can be retained for explicit hydrogen insertion and first-principles validation[22,24].

Hydride-derived template-screening case study

The screened structures were metal-only sublattices derived from known high-pressure superhydride prototypes. Construction began with parent hydride scaffolds. Hydrogen atoms were removed while the metal-site geometry was retained, and the metal sites were reassigned to mono-element or balanced binary element combinations. This procedure implements a chemical-template screening strategy in which the metal sublattice is evaluated for interstitial localization features that could template hydrogen networks before explicit hydrogen insertion.

Figure 5 summarizes 50,000 evaluated structures spanning 50 parent hydride formulas, 28 space groups, and pressure labels from 50 to 300 GPa. The 200 GPa parent-pressure bin contributes the largest number of structures. Additional information for this evaluation set is provided in Supplementary Table 8.

The mapping between the 50 parent hydride formulas and 55 prototype labels is listed in Supplementary Table 9.

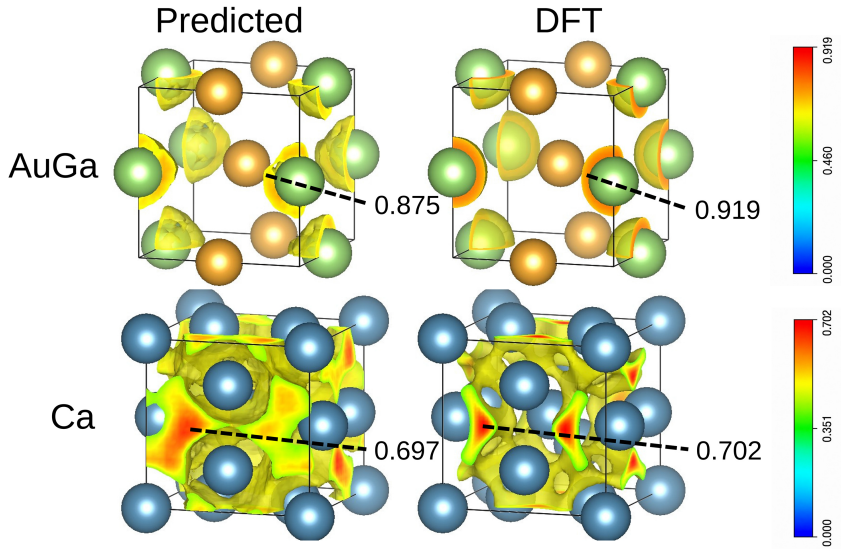

Figure 4. Predicted and DFT reference ELF fields for representative AuGa and Ca templates. Isosurfaces are rendered at ELF = 0.70 for AuGa and ELF = 0.55 for Ca. Surface color represents ELF values on the row-specific scales shown at right. Au, Ga, and Ca atoms are shown in gold, green, and blue, respectively. Although AuGa has the larger scalar global maximum, it does not form a pronounced basin in an open interstitial region; Ca exhibits a clear open-interstitial ELF basin. The dashed leaders identify the global ELF maxima and report their values. DFT: Density functional theory; ELF: electron localization function.

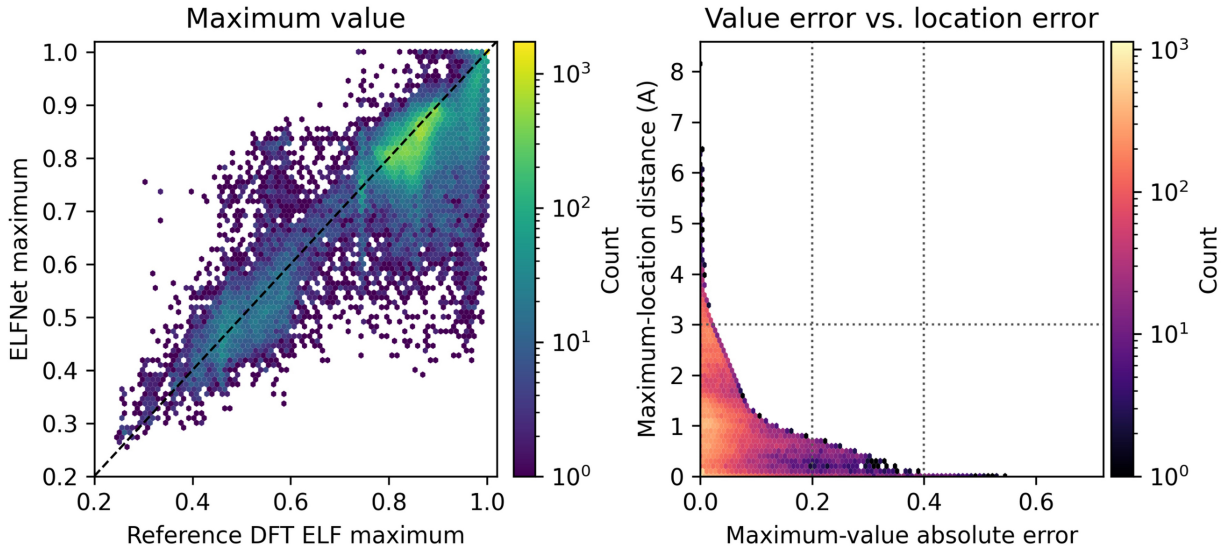

Figure 5. Maximum-value and maximum-location performance over the 50,000-structure hydride-derived template set. Left: parity plot comparing predicted and DFT reference global ELF maxima; the dashed line indicates perfect agreement. Right: absolute error in the maximum ELF value versus the real-space distance between predicted and DFT maximum locations. Both panels use logarithmic color density to indicate the number of structures in each bin. DFT: Density functional theory; ELF: electron localization function.

Maximum-value and maximum-location performance

For each template, the DFT and ELFNet ELF grids were reduced to two screening metrics: the value of the global ELF maximum and the real-space distance between the predicted and DFT maximum locations. The maximum-value metric tests whether ELFNet recovers the amplitude of the strongest localization feature. The maximum-location metric tests whether that feature appears in the same region of the unit cell as the DFT reference. These metrics are complementary because a prediction can closely match the maximum ELF

value while placing it in a different local basin, or it can place the maximum near the DFT reference location while under- or overestimating its value.

Figure 5 summarizes the template screening evaluation. The left panel compares the predicted and reference global ELF maxima, and the right panel compares maximum-value absolute error with maximum-location distance. The high density of structures with small value errors and small location distances shows that many templates are suitable for first-stage ranking, while the broader tails identify lower-confidence cases that should be prioritized for direct DFT validation if they appear chemically promising.

Across the 50,000-template set, the median absolute error in the global ELF maximum was 0.0268, and the median maximum-location distance was 0.895 Å. The corresponding 95th-percentile values were 0.1577 and 2.845 Å, respectively. At the thresholds shown in Figure 5, 93.57% of structures had both a maximum-value error no greater than 0.20 and a maximum-location distance no greater than 3 Å. These results show that the selected set is concentrated in the low-error region relevant to first-stage screening, while the remaining tail identifies lower-confidence cases for direct DFT validation. Large location errors should not necessarily be interpreted as complete field failures because the global maximum can shift between nearly equivalent localization basins. Additional percentile statistics and screening-threshold region counts are provided in Supplementary Table 10.

Pressure-stratified medians remained within the screening thresholds in every bin, with maximum-value errors no greater than 0.0361 and maximum-location distances no greater than 1.296 Å. The 160 GPa bin had the largest median location distance and the 300 GPa bin the largest median maximum-value error, but both remained within threshold; the aggregate result is therefore not solely a 200 GPa effect. Results are reported in Supplementary Table 11.

Among the 3,214 threshold failures, 1,515 exceeded only the maximum-value threshold, and 1,699 exceeded only the maximum-location threshold. No structure exceeded both thresholds. Pressure-stratified failure rates ranged from 2.36% at 100 GPa to 13.56% at 160 GPa. Failure rates also rose with structural and spatial complexity: from approximately 4% for one- or two-metal-site structures to 13.17% for eight-site structures and 26.57% for twelve-site structures, and from 3%-5% in the two smallest grid-volume quartiles to 9%-10% in the two largest. This pattern is consistent with increasing peak-location ambiguity. Element enrichments were confounded by pressure, scaffold, space group, and co-occurring elements and should not be interpreted as causal. Supplementary Tables 10 and 11 provide additional threshold-region counts and pressure-stratified summaries.

Cost of DFT ELF generation and model inference

To test whether ELFNet provides a practical timing advantage for screening, we compared fixed-geometry DFT ELF generation and ELFNet prediction for the same 50,000-structure evaluation set. For VASP, the final Total CPU time used (sec) and Elapsed time (sec) fields were parsed from each OUTCAR. The VASP calculations excluded ionic relaxation. Each calculation used eight MPI ranks, with up to six calculations running concurrently within a 48-task allocation on Intel Xeon Gold 6342 CPU nodes.

The ELFNet benchmark measured per-structure elapsed time from reading the benchmark input/reference ELFCAR through SAD construction, model inference, and writing the predicted ELFCAR. The input read supplied structure and grid metadata for the benchmark; a deployment workflow can instead begin from a structure file. Inference used one node containing four NVIDIA A100-SXM4-80GB GPUs and 32 CPU cores, with the evaluation set divided among four single-GPU processes. Because both workflows used

concurrent execution, the reported totals are sums of per-structure timer values rather than the elapsed wall time of either complete campaign.

Summed over the evaluation set, the VASP Total CPU time used fields were 922.6 h and the VASP Elapsed time fields were 935.3 h. The 922.6 h quantity is a sum of VASP timer fields and is not scheduler-accounted core-hours. The full ELFNet inference-through-write pipeline required 2.67 h of summed per-structure elapsed time. Using elapsed time consistently, the mean per-structure values were 67.34 s for VASP ELF generation and 0.192 s for the complete ELFNet pipeline. Their ratio corresponds to an approximately 350-fold reduction in summed per-structure elapsed time on the specified CPU and GPU systems. This is a hardware-specific timing ratio rather than a hardware-normalized compute-efficiency metric or a measured full-campaign wall-clock speedup.

The model forward pass accounted for 18.5 min of aggregate processing time; much of the end-to-end pipeline time arose from input/output and preprocessing. In a production screening workflow, maxima could be evaluated directly from predicted tensors, with full ELFCAR files written only for selected candidates. Additional timing definitions and values are provided in Supplementary Table 12.

Discussion, limitations, and outlook

This work frames ELF generation as a structure-to-field learning problem for high-pressure chemical-template screening. In the chemical-template picture, interstitial ELF basins in the metal sublattice indicate where the lattice may support the formation of dense hydrogen networks. The results show that a periodic full-grid residual network can learn a useful SAD-to-ELF mapping and apply it to large template libraries. This shifts ELF analysis from a post-processing step after DFT toward a front-end screen for prioritizing candidate metal templates.

The dataset design helps explain this screening behavior. It was built for volumetric ELF learning, not only for energy and force prediction. Its pressure-targeted structures expose the model to changes in localization fields under compression, making the dataset more relevant than one limited to equilibrium or near-equilibrium cells. The broad element coverage also supports exploratory searches beyond conventional ambient-pressure chemistry, including elements such as noble gases, lanthanides, and actinides, whose bonding behavior can change substantially under compression.

The hydride-derived template-screening case study tests the intended use of the model directly. By applying ELFNet to metal-only sublattices derived from known superhydride prototypes, the workflow evaluates whether a candidate metal framework produces interstitial localization features before explicit hydrogen insertion. The qualitative AuGa and Ca examples show that the predicted fields can distinguish whether localization forms a pronounced basin in an open interstitial region. The maximum-value and maximum-location metrics provide complementary checks, but the scalar maximum alone does not determine template character.

This evaluation is aligned with the composite training objective, which combines full-field reconstruction with explicit sensitivity to sparse high-ELF basins through the adaptive peak term. Cases with large value or location errors should therefore be treated as lower-confidence candidates and prioritized for direct DFT validation if they remain chemically promising.

Several methodological choices should be interpreted in the context of this screening goal. The reference ELFs were generated with PBE, so the predicted fields represent a consistent PBE-level ELF topology rather than a functional-independent electronic descriptor. The 96-structure r^{2}SCAN cross-check showed only modest changes in error and retained useful screening agreement across the change in functional. Templates

near screening thresholds should nevertheless receive direct higher-level validation. In addition, ELFNet enforces periodicity through circular padding but does not explicitly enforce crystallographic point-group or space-group symmetry. For a front-end screening model, this is acceptable because the central requirement is to recover the presence, approximate magnitude, and location of high-ELF basins that motivate follow-up DFT calculations.

Other limitations are related to the scope of the training data and the target property. The model was trained on compact unary and binary high-pressure prototypes, while the screening case study applies it to hydride-derived metal scaffolds with larger and more varied metal sublattices. Further validation on complete hydrides and larger multicomponent systems is therefore needed. ELFNet also predicts only the ELF field. It does not determine thermodynamic stability, dynamical stability, metallicity, electron–phonon coupling, or superconducting critical temperature, so promising templates still require downstream first-principles validation.

ELFNet is intended as a first-stage electronic screening tool rather than a substitute for stability or superconductivity calculations. It ranks hydrogen-free metal templates by their predicted interstitial localization. For selected templates, candidate hydrogen arrangements can then be generated through targeted or active-learning-accelerated crystal-structure searches, generative crystal models, or conditional structure-inpainting methods[28,29,31,65]. The resulting complete hydrides require structural relaxation and evaluation of thermodynamic stability, dynamical stability, electronic structure, electron–phonon coupling, and superconducting critical temperature.

The timing comparison supports the practical value of ELFNet as a screening step. On the reported hardware, the complete workflow reduced summed per-structure elapsed time by approximately 350-fold relative to fixed-geometry VASP ELF generation while including benchmark-input reading, SAD construction, model inference, and predicted ELFCAR writing. This ratio is hardware-specific; the elapsed time of a complete screening campaign also depends on concurrency, load balancing, storage performance, and scheduler conditions. In a production workflow, ranking metrics could be evaluated directly from predicted tensors, with full ELFCAR files saved only for selected structures.

Future work should focus on improving candidate ranking and identifying cases outside the training domain. Uncertainty-aware screening and active learning could prioritize structures whose predicted ELF fields are most likely to require additional DFT data[29,30]. Optional symmetry averaging at inference time could reduce small symmetry-breaking artifacts when higher-fidelity fields are needed. Joint prediction of related electronic descriptors, such as charge density, could also provide a broader picture of the electronic structure used in high-pressure template screening. Future extensions could pair ELFNet with recent equivariant and grid-based electron-density models to combine ELF-based localization screening with complementary charge-density information[32-36].

CONCLUSIONS

We developed ELFNet, a structure-to-field framework for predicting DFT-derived ELF fields from computationally inexpensive SAD inputs. The work combines a pressure-targeted PBE ELF dataset containing 326,009 converged unary and binary prototype structures across 89 elements with a periodic full-grid residual network trained with a composite objective to recover high-ELF localization features relevant to chemical-template screening under pressure.

The hydride-derived template-screening case study demonstrates the practical value of this approach. Applied to 50,000 metal-only templates derived from known high-pressure superhydride prototypes, ELFNet retained chemically meaningful differences in the spatial topology of interstitial localization, including

whether a pronounced basin forms in an open interstitial region. Performance was evaluated using global maximum ELF values and maximum-location errors against DFT reference ELFs. On the reported hardware, the full inference-through-write workflow reduced summed per-structure elapsed time by approximately 350-fold relative to fixed-geometry VASP ELF generation. These results establish ELF-based chemical-template analysis as a low-cost front-end screening method for prioritizing candidate metal templates before downstream first-principles validation.

DECLARATIONS

Authors' contributions

Contributed to the conception and design of the study, data curation, formal analysis, investigation, methodology, software, validation, visualization, and manuscript drafting and revision: Ellis, A.
Contributed to data curation, formal analysis, investigation, and manuscript drafting and revision: Ellis, S.
Contributed to the conception and design of the study, funding acquisition, project administration, resources, supervision, and manuscript revision: Miao, M.
All authors reviewed and approved the final manuscript.

Availability of data and materials

The ELFNet source code, bundled checkpoint, model documentation, training and inference scripts, configuration files, and example POSCAR files are available at https://github.com/Austin243/ELFNet. The code version used for this study is the public repository state at the time of submission. Data supporting this article, including the pressure-targeted high-pressure PBE ELF training dataset, training-triplet files, and hydride-derived inference and screening data with provenance, are available from the public Zenodo record at https://zenodo.org/records/20481629. Supplementary Materials provide pressure-targeting, dataset, training-dynamics, template-set, Figure 5 performance, and timing summaries.

AI and AI-assisted tools statement

During preparation of this manuscript, generative AI tools, including ChatGPT (OpenAI), were used solely to assist with language editing, organization, and formatting. Figure 3 was created with the assistance of Gemini Nano Banana Pro. The authors reviewed and edited all AI-assisted content, verified Figure 3 for scientific accuracy, and accept full responsibility for the manuscript. No AI tool was used to generate primary scientific data, perform the reported analyses, or draw the scientific conclusions.

Financial support and sponsorship

This work was supported by the U.S. Department of Defense Research and Education Program for Historically Black Colleges and Universities and Minority-Serving Institutions Basic Research Funding (No. W911NF2310232); the National Science Foundation (DMR-1848141 and OAC-2117956); the Camille and Henry Dreyfus Foundation; and California State University Research, Scholarship and Creative Activity awards. The funding bodies had no role in the study design, data collection, analysis, interpretation, preparation of the manuscript, or decision to submit the work for publication.

Conflicts of interest

All authors declared that there are no conflicts of interest.

Ethical approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Supplementary Materials

Supplementary Materials

REFERENCES

1. Becke, A. D.; Edgecombe, K. E. A simple measure of electron localization in atomic and molecular systems. *J. Chem. Phys.* **1990**, *92*, 5397-403. DOI
2. Silvi, B.; Savin, A. Classification of chemical bonds based on topological analysis of electron localization functions. *Nature* **1994**, *371*, 683-6. DOI
3. Savin, A.; Nesper, R.; Wengert, S.; Fässler, T. F. ELF: the electron localization function. *Angew. Chem. Int. Ed. Engl.* **1997**, *36*, 1808-32. DOI
4. Hilleke, K. P.; Bi, T.; Zurek, E. Materials under high pressure: a chemical perspective. *Appl. Phys. A.* **2022**, *128*, 441. DOI
5. Ashcroft, N. W. Hydrogen dominant metallic alloys: high temperature superconductors? *Phys. Rev. Lett.* **2004**, *92*, 187002. DOI
6. Drozdov, A. P.; Eremets, M. I.; Troyan, I. A.; Ksenofontov, V.; Shylin, S. I. Conventional superconductivity at 203 kelvin at high pressures in the sulfur hydride system. *Nature* **2015**, *525*, 73-6. DOI PubMed
7. Duan, D.; Liu, Y.; Tian, F.; et al. Pressure-induced metallization of dense $(H_2S)_2H_2$ with high-T_c superconductivity. *Sci. Rep.* **2014**, *4*, 6968. DOI PubMed PMC
8. Li, Y.; Hao, J.; Liu, H.; Li, Y.; Ma, Y. The metallization and superconductivity of dense hydrogen sulfide. *J. Chem. Phys.* **2014**, *140*, 174712. DOI
9. Drozdov, A. P.; Kong, P. P.; Minkov, V. S.; et al. Superconductivity at 250 K in lanthanum hydride under high pressures. *Nature* **2019**, *569*, 528-31. DOI
10. Somayazulu, M.; Ahart, M.; Mishra, A. K.; et al. Evidence for superconductivity above 260 K in lanthanum superhydride at megabar pressures. *Phys. Rev. Lett.* **2019**, *122*, 027001. DOI
11. Errea, I.; Belli, F.; Monacelli, L.; et al. Quantum crystal structure in the 250-kelvin superconducting lanthanum hydride. *Nature* **2020**, *578*, 66-9. DOI
12. Wang, H.; Tse, J. S.; Tanaka, K.; Iitaka, T.; Ma, Y. Superconductive sodalite-like clathrate calcium hydride at high pressures. *Proc. Natl. Acad. Sci. U. S. A.* **2012**, *109*, 6463-6. DOI PubMed PMC
13. Peng, F.; Sun, Y.; Pickard, C. J.; Needs, R. J.; Wu, Q.; Ma, Y. Hydrogen clathrate structures in rare earth hydrides at high pressures: possible route to room-temperature superconductivity. *Phys. Rev. Lett.* **2017**, *119*, 107001. DOI
14. Ma, L.; Wang, K.; Xie, Y.; et al. High-temperature superconducting phase in clathrate calcium hydride CaH_6 up to 215 K at a Pressure of 172 GPa. *Phys. Rev. Lett.* **2022**, *128*, 167001. DOI
15. Li, Z.; He, X.; Zhang, C.; et al. Superconductivity above 200 K discovered in superhydrides of calcium. *Nat. Commun.* **2022**, *13*, 2863. DOI PubMed PMC
16. Bi, J.; Nakamoto, Y.; Zhang, P.; et al. Giant enhancement of superconducting critical temperature in substitutional alloy $(La,Ce)H_9$. *Nat. Commun.* **2022**, *13*, 5952. DOI PubMed PMC
17. Chen, W.; Huang, X.; Semenok, D. V.; et al. Enhancement of superconducting properties in the La-Ce-H system at moderate pressures. *Nat. Commun.* **2023**, *14*, 2660. DOI PubMed PMC
18. Chen, L. C.; Luo, T.; Cao, Z. Y.; et al. Synthesis and superconductivity in yttrium-cerium hydrides at high pressures. *Nat. Commun.* **2024**, *15*, 1809. DOI PubMed PMC
19. Shan, P.; Ma, L.; Yang, X.; et al. Molecular hydride superconductor BiH_4 with T_c up to 91 K at 170 GPa. *J. Am. Chem. Soc.* **2025**, *147*, 4375-81. DOI
20. Liu, P.; Zhuang, Q.; Xu, Q.; Cui, T.; Liu, Z. Mechanism of high-temperature superconductivity in compressed H_2-molecular-type hydride. *Sci. Adv.* **2025**, *11*, eadt9411. DOI PubMed PMC
21. Zhao, W.; Ellis, A.; Duan, D.; et al. Unlocking the origin of high-temperature superconductivity in molecular hydrides at moderate pressures. *Adv. Funct. Mater.* **2025**, *35*, 2415910. DOI
22. Sun, Y.; Miao, M. Chemical templates that assemble the metal superhydrides. *Chem* **2023**, *9*, 443-59. DOI
23. Belli, F.; Novoa, T.; Contreras-García, J.; Errea, I. Strong correlation between electronic bonding network and critical temperature in hydrogen-based superconductors. *Nat. Commun.* **2021**, *12*, 5381. DOI PubMed PMC
24. Sun, Y.; Ellis, A.; Chen, X.; Miao, M. Batch discovery of complex metal superhydrides via an effective machine learning method structured by chemical template concept. *J. Am. Chem. Soc.* **2025**, *147*, 40407-19. DOI PubMed
25. Glass, C. W.; Oganov, A. R.; Hansen, N. USPEX - evolutionary crystal structure prediction. *Comput. Phys. Commun.* **2006**, *175*, 713-20. DOI
26. Wang, Y.; Lv, J.; Zhu, L.; Ma, Y. Crystal structure prediction via particle-swarm optimization. *Phys. Rev. B.* **2010**, *82*, 094116. DOI

27. Merchant, A.; Batzner, S.; Schoenholz, S. S.; Aykol, M.; Cheon, G.; Cubuk, E. D. Scaling deep learning for materials discovery. *Nature* **2023**, *624*, 80-5. DOI PubMed PMC
28. Zeni, C.; Pinsler, R.; Zügner, D.; et al. A generative model for inorganic materials design. *Nature* **2025**, *639*, 624-32. DOI PubMed PMC
29. Hessmann, S. S. P.; Schütt, K. T.; Gebauer, N. W. A.; Gastegger, M.; Oguchi, T.; Yamashita, T. Accelerating crystal structure search through active learning with neural networks for rapid relaxations. *npj. Comput. Mater.* **2025**, *11*, 44. DOI
30. Kulichenko, M.; Barros, K.; Lubbers, N.; et al. Uncertainty-driven dynamics for active learning of interatomic potentials. *Nat. Comput. Sci.* **2023**, *3*, 230-9. DOI PubMed PMC
31. Conway, L. J.; Pickard, C. J. Accelerating crystal structure prediction using data-derived potentials: high-pressure binary hydrides. *Ann. Phys.* **2026**, *538*, e00608. DOI
32. Jørgensen, P. B.; Bhowmik, A. Equivariant graph neural networks for fast electron density estimation of molecules, liquids, and solids. *npj. Comput. Mater.* **2022**, *8*, 183. DOI
33. Lv, T.; Zhong, Z.; Liang, Y.; Li, F.; Huang, J.; Zheng, R. Deep charge: deep learning model of electron density from a one-shot density functional theory calculation. *Phys. Rev. B.* **2023**, *108*, 235159. DOI
34. Koker, T.; Quigley, K.; Taw, E.; Tibbetts, K.; Li, L. Higher-order equivariant neural networks for charge density prediction in materials. *npj. Comput. Mater.* **2024**, *10*, 161. DOI
35. Qin, X.; Lv, T.; Zhong, Z. EAC-Net: predicting real-space charge density via equivariant atomic contributions. *J. Chem. Theory. Comput.* **2026**, *22*, 4813-21. DOI PubMed
36. Li, C.; Sharir, O.; Yuan, S.; Chan, G. K. Image super-resolution inspired electron density prediction. *Nat. Commun.* **2025**, *16*, 4811. DOI PubMed PMC
37. Wang, Z.; Gong, Y.; Evans, M. L.; et al. Machine learning-accelerated discovery of A_2BC_2 ternary electrides with diverse anionic electron densities. *J. Am. Chem. Soc.* **2023**, *145*, 26412-24. DOI
38. Dewaele, A.; Worth, N.; Pickard, C. J.; et al. Synthesis and stability of xenon oxides Xe_2O_5 and Xe_3O_2 under pressure. *Nat. Chem.* **2016**, *8*, 784-90. DOI
39. Dong, X.; Oganov, A. R.; Goncharov, A. F.; et al. A stable compound of helium and sodium at high pressure. *Nat. Chem.* **2017**, *9*, 440-5. DOI
40. Kaplan, A. D.; Liu, R.; Qi, J.; et al. A foundational potential energy surface dataset for materials. *arXiv* **2025**, arXiv:2503.04070. Available online: https://doi.org/10.48550/arXiv.2503.04070. (accessed 2026-08-20).
41. Kuner, M. C.; Kaplan, A. D.; Persson, K. A.; Asta, M.; Chrzan, D. C. MP-ALOE: an r^{2}SCAN dataset for universal machine learning interatomic potentials. *npj. Comput. Mater.* **2025**, *11*, 352. DOI
42. Deng, B.; Zhong, P.; Jun, K.; et al. CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nat. Mach. Intell.* **2023**, *5*, 1031-41. DOI
43. Barros-Luque, L.; Shuaibi, M.; Fu, X.; et al. The Open Materials 2024 (OMat24) inorganic materials dataset and models. *Nat. Comput. Sci.* **2026**, *6*, 642-52. DOI
44. Chen, C.; Ong, S. P. A universal graph deep learning interatomic potential for the periodic table. *Nat. Comput. Sci.* **2022**, *2*, 718-28. DOI PubMed
45. Zhang, D.; Liu, X.; Zhang, X.; et al. DPA-2: a large atomic model as a multi-task learner. *NPJ. Comput. Mater.* **2024**, *10*, 293. DOI PubMed PMC
46. Batzner, S.; Musaelian, A.; Sun, L.; et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **2022**, *13*, 2453. DOI PubMed PMC
47. Musaelian, A.; Batzner, S.; Johansson, A.; et al. Learning local equivariant representations for large-scale atomistic dynamics. *Nat. Commun.* **2023**, *14*, 579. DOI PubMed PMC
48. Batatia, I.; Kovács, D. P.; Simm, G. N. C.; Ortner, C.; Csányi, G. MACE: higher order equivariant message passing neural networks for fast and accurate force fields. *arXiv* **2022**, arXiv:2206.07697. Available online: https://doi.org/10.48550/arXiv.2206.07697. (accessed 2026-08-20).
49. Batatia, I.; Benner, P.; Chiang, Y.; et al. A foundation model for atomistic materials chemistry. *J. Chem. Phys.* **2025**, *163*, 184110. DOI
50. Lysogorskiy, Y.; Bochkarev, A.; Drautz, R. Graph atomic cluster expansion for foundational machine learning interatomic potentials. *npj. Comput. Mater.* **2026**, *12*, 114. DOI
51. Barlow, R. E.; Bartholomew, D. J.; Bremner, J. M.; Brunk, H. D. Statistical inference under order restrictions. The theory and application of isotonic regression. John Wiley & Sons: London and New York, 1972. https://books.google.com/books/about/Statistical_Inference_Under_Order_Restri.html?id=DEamySUDBWcC&utm_source. (accessed 2026-08-20).
52. Kohn, W.; Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **1965**, *140*, A1133. DOI
53. Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **1996**, *77*, 3865. DOI

54. Kresse, G.; Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B.* **1996**, *54*, 11169. DOI
55. Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B.* **1994**, *50*, 17953. DOI
56. Kresse, G.; Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B.* **1999**, *59*, 1758. DOI
57. Furness, J. W.; Kaplan, A. D.; Ning, J.; Perdew, J. P.; Sun, J. Accurate and numerically efficient r²SCAN meta-generalized gradient approximation. *J. Phys. Chem. Lett.* **2020**, *11*, 8208-15. DOI
58. Hirshfeld, F. L. Bonded-atom fragments for describing molecular charge densities. *Theoret. Chim. Acta.* **1977**, *44*, 129-38. DOI
59. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, USA. Jun 27-30, 2016. IEEE; 2016. pp 770-8. DOI
60. Wu, Y.; He, K. Group normalization. In Ferrari, V.; Hebert, M.; Sminchisescu, C.; Weiss, Y.; editors. *Computer Vision - ECCV 2018*. Cham: Springer International Publishing; 2018. pp. 3-19. DOI
61. Hendrycks, D.; Gimpel, K. Gaussian error linear units (GELUs). *arXiv* **2016**, arXiv:1606.08415. Available online: https://doi.org/10.48550/arXiv.1606.08415. (accessed 2026-08-20).
62. Hu, J.; Shen, L.; Sun, G. Squeeze-and-excitation networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, USA. Jun 18-23, 2018. IEEE; 2018. pp 7132-41. DOI
63. Woo, S.; Park, J.; Lee, J. Y.; Kweon, I. S. CBAM: convolutional block attention Module. In: Ferrari, V.; Hebert, M.; Sminchisescu, C.; Weiss, Y.; editors. *Computer Vision - ECCV 2018*. Cham: Springer International Publishing; 2018. pp. 3-19. DOI
64. Cipolla, R.; Gal, Y.; Kendall, A. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, USA. Jun 18-23, 2018. IEEE; 2018. pp. 7482-91. DOI
65. Reents, T.; Cantarella, A.; Bercx, M.; Bonfà, P.; Pizzi, G. Score-based diffusion models for accurate crystal-structure inpainting and reconstruction of hydrogen positions. *npj. Comput. Mater.* **2026**, *12*, 203. DOI