npj | computational materials

Article

Published in partnership with the Shanghai Institute of Ceramics of the Chinese Academy of Sciences

https://doi.org/10.1038/s41524-025-01863-4

Hierarchical transfer learning: an agile and equitable strategy for machine-learning interatomic models

Rebecca K. Lindsey[1] ✉, Awwal D. Oladipupo[1], Sorin Bastea[2], Bradley A. Steele[2], I-Feng W. Kuo[2] & Nir Goldman[2,3]

Machine-learned interatomic models are growing in popularity due to their ability to afford near quantum-accurate predictions for complex phenomena with orders-of-magnitude greater computational efficiency. However, these models struggle when applied to systems of many element types due to the approximately exponential increase in number of parameters that must be determined. To mitigate this challenge, we present a new hierarchical transfer learning approach that allows the fitting problem to be decomposed into smaller independent and reusable parameter blocks that enable development of explicitly chemically extensible ML-IAM. Application of this strategy is demonstrated for C and N mixtures under conditions ranging from nominally ambient to ~10,000 K and 200 GPa for compositions from 0 to 100% N. Ultimately, this strategy makes model generation for chemically complex systems more tractable and efficient, facilitates comprehensive model validation, and makes ML-IAM development for problems of this nature more accessible to users with limited access to extreme computing infrastructure.

Machine learning (ML) has emerged as a powerful tool for developing interatomic models (IAM) capable of bridging the computational efficiency of classical molecular mechanics approaches and the predictive power of first-principles-based methods. In essence, this is achieved by directly learning the target potential energy surface (PES) topography onto a flexible set of basis functions, rather than trying to reconstruct the PES through analytical expressions as is typically done in classical "force field' strategies. This model generation approach is particularly useful for simulating complex systems (e.g., condensed phase reacting systems and materials under extremely high temperature and pressure conditions) for which suitable molecular mechanics descriptions are not known a priori.

ML-IAM development is typically accomplished through five high-level steps: (1) an initial training data is generated, comprising a series of system configurations with corresponding forces, energies, and/or stresses assigned via, e.g., Kohn–Sham Density Functional Theory[1] (DFT); (2) system configurations are then recast as a series of descriptors that encode the local chemical environments; (3) a model architecture is selected, e.g., a smooth set of basis functions that ingest these descriptors and can be used to predict corresponding energies and by consequence, forces, and stresses; (4) model parameters are determined via optimization; (5) the model is iteratively refined until desired accuracy is achieved through strategies such as active learning.

Many open source tools for each of these steps are now available. For example, for the ML portion (steps 2–5), a multitude of descriptor approaches[2–6], ML model architectures[7–16] and even active learning tools[6,12–14,17–20], have been published, each of which excels in different problem spaces (e.g., data-rich vs data-poor fitting, large-scale vs relatively small-scale simulation, materials vs chemistry applications). Ultimately, this has enabled what were once viewed as challenging fitting problems to serve as basic model and method benchmarks[21], as well as generation of highly transferable general-purpose ML-IAM parameter sets[22–26]. However, there remain a number of application spaces necessitating the balance of accuracy and efficiency afforded by ML-IAMs for which model development remains far from trivial. These problems tend to exist at the confluence of high chemical and configurational complexity due to the need for large models, large training sets, and exhaustive validation. In particular, efforts to develop ML-IAMs for systems with highly complex potential energy surfaces – such as those involving both bonded and non-bonded interactions that span widely separated energy scales, as seen in condensed-phase molecular reactions and covalent phase transformations – have largely been restricted to systems with fewer than three element types, due to the roughly exponential growth in the number of required parameters with increasing chemical complexity. For example, modeling evolution in

[1]Department of Chemical Engineering, University of Michigan, Ann Arbor, MI, USA. [2]Physical and Life Sciences Directorate, Lawrence Livermore National Laboratory, Livermore, CA, USA. [3]Department of Chemical Engineering, University of California, Davis, CA, USA. ✉e-mail: rklinds@umich.edu

carbon-containing systems under reactive conditions is particularly difficult due to the disparate energy scales relevant for conformational change, bond formation/breaking, and non-bonded interactions and the fact that these materials tend to contain three or more element types (e.g., C, N, O, and H)[6,26–28].

Active learning, in which the fitting framework autonomously attempts to identify maximally informative candidate (i.e., unlabeled) training data for labeling and subsequent addition to the training set can make these fitting problems more tractable by reducing training data volume requirements[6,13,18,19]. However, it does not address the remaining practical challenges of large models, large training sets, and the need for exhaustive validation. Therefore, in this work, we describe a new transfer learning strategy to (1) reduce fit complexity, (2) enable the physicochemical space over which models are suited to be better defined, and (3) enable generation of models that are more transferable and efficiently generated than those fit using traditional "direct-learned" methods. Our transfer learning approach is distinct from that used for neural network-type ML-IAMs[29], and is designed for use with parametrically linear ML-IAMs that employ a descriptor that provides unique and fully separable representations of n-body interaction clusters based on both order and chemical composition[7,26]. Our strategy draws inspiration from classical transferable force fields[30–32], allowing the fitting problem to be decomposed into small reusable parameter blocks and thereby enables far greater chemically extensibly than currently available strategies. Our hierarchical strategy also enables immediate application to targeted subsets of chemical space while concurrently refining and expanding model chemical scope. This allows for real-time problem-solving within the initial domain as fitting efforts progressively scale to accommodate increasingly complex chemistry.

In the following sections, we provide an overview of our approach within the context the Chebyshev Interaction Model for Efficient Simulation (ChIMES) ML-IAM, followed by discussion of our target application and initial training data generation strategy. We apply our approach to a testbed system comprising mixtures of carbon (C) an nitrogen (N) under conditions ranging from nominally ambient temperature (T) and pressure (P) up to 10,000 K and 200 GPa. We note that this testbed was chosen due to the relatively low atomic complexity (i.e., only two species are present) but high configurational complexity (e.g., including multiple phases, compositions, and chemistry). Performance of transfer- and direct-learned models are compared with DFT, and results are discussed within the context of the following guiding questions:

- Given models for pure-C and -N systems, can we build a high-quality C/N model without having to refit any parameters?
- How does performance of hierarchally-transfer-learned models compare with a model fit via the standard approach?
- Beyond agility and parameter reuse, do any other advantages emerge from this strategy?

Results

ML-IAMs are characterized by two main features: model architecture and descriptor. For the majority of ML-IAMs, these features are structured such that parameters for interactions between atoms of various types are inseparable; introduction of an additional atom type requires generating a new model with all parameters updated[2,3,5,11,13,15,22]. As a consequence, ML-IAMs are generally fit *a-la-carte* for a given target system, atomic composition, and set of conditions. Specifically, use of an atom-centered descriptor *or* use of complex architecture (e.g., neural network, graph based, etc) generally precludes generating models that are explicitly compositionally-extensible. In the section below, we will show that the cluster-based descriptor and parametrically linear form used by the ChIMES ML-IAM overcomes this limitation, enabling a unique opportunity for developing chemically extensible models through a hierarchical transfer learning strategy.

Model overview

In this section, we provide a brief overview of the ChIMES ML-IAM, emphasizing features salient to the presently described transfer learning strategy. For a more detailed discussion of the model and its underlying form, we direct the reader to refs. 6 and 33.

ChIMES describes system energy through an explicit many-body cluster expansion, i.e:

$$E = \sum_{i}^{n_a} E_i + \sum_{i>j}^{n_a}\sum E_{ij} + \sum_{i>j>k}^{n_a}\sum\sum E_{ijk} + \cdots, \tag{1}$$

where E is the total ChIMES energy for a system of n_a atoms, E_i is the energy for a single atom, i, E_{ij}, and E_{ijk} are the energy for a cluster of two or three atoms (i.e., ij or ijk), respectively. This expansion can extend to arbitrary bodiedness, though all models produced to date contain a maximum of 4-body terms[6,7,26,27,33–38]. Interactions between pairs of atoms are described through Chebyshev polynomial series that take as input interatomic pair distances, i.e., for a pair of two atoms ij:

$$E_{ij} \propto \sum_{\alpha}^{\mathcal{O}_{2B}} c_{\alpha}^{e_i e_j} T_{\alpha}\left(s_{ij}^{e_i e_j}\right), \tag{2}$$

where T_α is a Chebyshev polynomial of order α, $s_{ij}^{e_i e_j}$ is a transformed pair distance between atoms ij of element type $e_i e_j$, $\mathcal{O}_{2B}$ is the user-defined two-body order for the polynomial series, and $c_{\alpha}^{e_i e_j}$ are the Chebyshev polynomial coefficients that comprise the fitting parameters of the model. Note that we use the "$\propto$" symbol to indicate that smoothing and penalty functions have been excluded from these equations for simplicity. Many-body interactions are treated as the product of interactions for constituent atom pairs. For example, a three-body interaction is given by:

$$E_{ijk} \propto \sum_{\beta}^{\mathcal{O}_{3B}}\sum_{\gamma}\sum_{\delta}{}' c_{\beta\gamma\delta}^{e_i e_j, e_i e_k, e_j e_k} T_{\beta}\left(s_{ij}^{e_i e_j}\right) T_{\gamma}\left(s_{ik}^{e_i e_k}\right) T_{\delta}\left(s_{jk}^{e_j e_k}\right), \tag{3}$$

Where the "′" indicates that the sum only considers terms for which at least two of β, γ, and δ are non-zero, ensuring a true three-body interaction. Previous work has shown that this functional form is well suited for describing variety of systems, including inorganic materials, covalent materials, condensed phase reacting systems, and molecular systems and materials[23,26,33,34,36,38].

Models are fit by force-, energy-, and/or stress matching to gas- and/or condensed-phase atomic configurations labeled by a ground-truth method such as DFT; this is achieved by minimizing an objective function of the form:

$$F_{\text{obj}} \propto \sqrt{\sum_{i=1}^{n_f}\left(w_{\text{E}_i}^2 \Delta E_i^2 + \sum_{j=1}^{n_a}\sum_{k=1}^{3} w_{\text{F}_{ijk}}^2 \Delta F_{ijk}^2 + \sum_{j=1}^{9} w_{\sigma_{ij}}^2 \Delta\sigma_{ij}^2\right)}, \tag{4}$$

where $\Delta X = X^{\text{DFT}} - X^{\text{ChIMES}\{c\}}$ and X is a force, energy or stress predicted by the superscripted method. F_{obj} and $\{c\}$ are the weighted root-mean-squared error and model coefficients, respectively. The number of frames and atoms are given by n_f and n_a, respectively. F_{ijk} indicates the k^{th} Cartesian component of the force on atom j in configuration i while σ_{ij} and E_i indicates the j component of the stress tensor and the energy for configuration i, respectively. Weights for each force, energy, and stress are given by w.

Since ChIMES is entirely linear in its fitted parameters $\{c\}$, the model optimization problem can be recast as the following over-determined matrix equation:

$$\boldsymbol{wMc} = \boldsymbol{wX}_{\text{DFT}}, \tag{5}$$

where $\boldsymbol{X}_{\mathrm{DFT}}$ is the vector of F_{ijk}^{DFT}, $\sigma_{ij}^{\mathrm{DFT}}$, and E^{DFT} values, $\boldsymbol{w}$ is a diagonal matrix of weights to be applied to the elements of $\boldsymbol{X}_{\mathrm{DFT}}$ and rows of $\boldsymbol{M}$, and the elements of design matrix $\boldsymbol{M}$ are given by:

$$M_{ab} = \frac{\partial X_{a,\mathrm{ChIMES}\{c\}}}{\partial c_b}. \tag{6}$$

In the above, a represents a combined index over force and energy components, and b is the index over permutationally invariant model parameters. This allows model parameters to be rapidly generated by application of advanced linear solvers[38–41] and makes the model well suited for iterative and/or active-learning training strategies. For additional details, the reader is directed to refs. 6 and 26.

Parameter hierarchy and transfer learning overview

ChIMES models view system configurations as a collection of fully connected graphs between atoms in n-body clusters, where atoms form nodes and the transformed distances between those atoms form the edges. We refer to these cluster graphs as the ChIMES descriptor. The parametric linearity characteristic to ChIMES models coupled with use of a cluster-centered many-body descriptor gives rise to an inherently hierarchical parameter structure that can be leveraged for transfer learning. Specifically, the atom cluster energy terms given in Eq. (1) can be further decomposed based on constituent atom types. For example, the two-body energy contributions for a system comprised entirely of C and N can be written as:

$$\sum_{i>j}^{n_a}\sum E_{ij} = \sum_{i>j}^{n_C}\sum E_{ij}^{\mathrm{CC}} + \sum_{i>j}^{n_N}\sum E_{ij}^{\mathrm{NN}} + \sum_{i}^{n_C}\sum_{j}^{n_N} E_{ij}^{\mathrm{CN}}, \tag{7}$$

where n_{C} and n_{N} are the number of C and N atoms in the system, respectively i.e., $n_{\mathrm{C}} + n_{\mathrm{N}} = n_{\mathrm{a}}$ and $E_{ij}^{e_i e_j}$ is the two-body energy for a set of atoms ij of element types $e_i e_j$. Similarly, for a three-body interaction:

$$\begin{aligned}\sum_{i>j>k}^{n_a}\sum E_{ijk} = {} & \sum_{i>j>k}^{n_C}\sum E_{ijk}^{\mathrm{CCC}} \\ & + \sum_{i>j>k}^{n_N}\sum E_{ijk}^{\mathrm{NNN}} + \sum_{i>j}^{n_C}\sum\sum_{k}^{n_N} E_{ijk}^{\mathrm{CCN}} \\ & + \sum_{i}^{n_C}\sum_{j>k}^{n_N}\sum E_{ijk}^{\mathrm{CNN}}.\end{aligned} \tag{8}$$

This logic can be extended for construction of higher-body interactions. Notably, pure component terms (and thus, parameters) for C and N interactions are confined to the first two terms of Eqs. (7) and (8), and are non-interacting. Two-atom-type cross-interactions are contained in the remaining terms. This point is illustrated in Fig. 1. For a C, H, O, and N ChIMES model, all terms for interactions between only C atoms and between only N atoms are contained in the orange C and N blocks, respectively, while all cross-interaction terms are contained within the yellow CN block. Critically, this means parameters for a given block in the same column can be fit in parallel, completely independently of one another. For example, a ChIMES model containing up to four body interactions for a C/N system is comprised of two "building blocks" containing {C, CC, CCC, CCCC} and {N, NN, NNN, NNNN} parameters, and a CN cross-term building block containing {CN, CCN, CNN, CCCN, CCNN, CNNN}. We note that, throughout this text, 'C/N' refers to systems or datasets containing both carbon and nitrogen atoms (e.g., C/N systems), as well as the corresponding models that fully describe them, including pure-C, pure-N, and C–N cross interactions. In contrast, 'CN' is used specifically to denote parameters associated with 'C–N` cross-interactions.

Following the precedent set by previous ChIMES model development endeavors, a ChIMES-CN model would traditionally be generated by fitting all of these parameters at once. However, the unique model structure also allows a "hierarchical transfer learning" approach wherein C- and N-

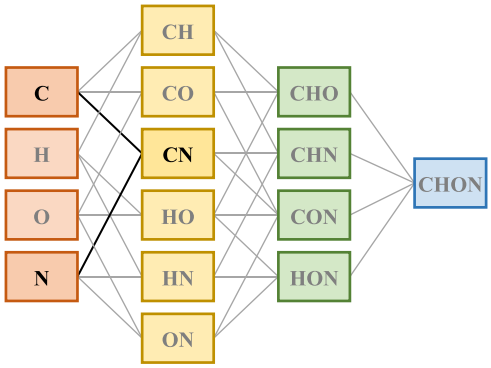

Fig. 1 | Schematic of a ChIMES parameter hierarchy for a model describing C, H, O, and N containing systems. Parameters in a given column can be fit completely independent of one another. Parameters blocks with two or more atoms represent cross-interactions between the indicated atom types.

building blocks are fit independently to pure-C and pure-N training data, respectively. CN-block parameters can then be fit to two-element system training data by replacing the definition of ΔX used in Eq. (4) with: $\Delta X = X^{\mathrm{DFT}'} - X^{\mathrm{ChIMES}\{c\}}$, where $X^{\mathrm{DFT}'} = X^{\mathrm{DFT}} - X^{\mathrm{ChIMES-C}} - X^{\mathrm{ChIMES-N}}$ and ChIMES–C and ChIMES–N indicate X computed using the C and N block parameters, respectively. This same logic can be extended to trinary and quaternary systems e.g., the resulting CN parameter block along with the previously fit C and N parameter blocks could be used in development of, e.g. CHN, CON, and CHON models.

Prototypical system overview and training strategy

Like other ML-IAMs, these building blocks have historically been fit all at once, yielding ChIMES models for which applications are confined to specific set of atom types and limited to certain compositional ranges[27,28,36,38,43]. Here, we explore efficacy of the hierarchical transfer learning strategy described above to develop a C/N model valid from near-ambient conditions to extreme conditions of ~10,000 K and 200 GPa that is suitable for describing any range of compositions from 0 to 100% N, by building upon previously generated ChIMES models for C (i.e., the 2024-Large model[26]) and N[33]. We note that this testbed C/N system is also interesting within the context of synthesis of N-doped graphictic materials for applications including catalysis, energy storage, and sensing[44–49]. For example, shock-compression of C/N-rich precursor materials has been shown capable of producing nitrogen-containing graphitic nanoonions on sub μs timescales, which holds incredible promise as a high-throughput strategy for tailored synthesis of exotic and technologically relevant carbon nanomaterials[50,51]. However, governing phenomena and associated kinetics remain poorly understood due to the extreme associated T and P and far-from equilibrium events that occur. Hence the present pure C/N systems can serve as a reductionist model for understanding this process.

Training data were generated through a combination of Kohn–Sham DFT molecular dynamics (MD) simulations and single point calculations, details of which can be found in Section "Methods". The C/N binary phase diagram is unknown under our conditions of interst; thus, to generate training data, simulations were launched DFT-MD for three different C/N compositions, at a variety of temperatures and densities spanning 300 K/1 g cm^3 to 9000 K/4 g cm^3 as shown in the plot in Fig. 2. Systems at densities below 2.9 g cm^3 were initialized with a graphitic structure, while higher density initial configurations were initialized with a diamond-like structure; N-atoms were then introduced by random substitution. Simulations were run for at least 5 ps; 20 training configurations were taken from each of the 10 simulations to build the initial 298 configuration training set. As is shown in Fig. 2, resulting configurations span graphitic, compressed gas, and high-density liquid, containing both small molecules and larger, polymeric structures. This training set was supplemented with 98 configurations for 3 solid C/N materials, mp-1985, mp-571653, and mp-563, found in the

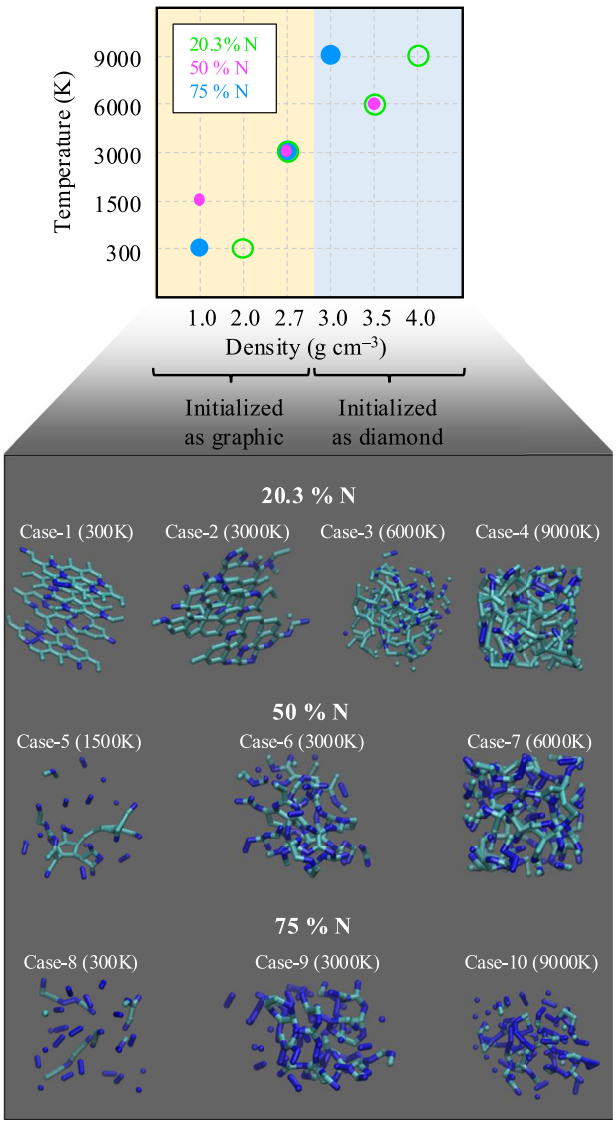

Fig. 2 | Overview of the training data used for model development in this work. All training data for mixed C/N systems with nitrogen fraction, temperature, and pressure as given in the plot inset. Simulations used to generate training data points were initialized as N-doped graphite or diamond, as indicated below the plot. Representative snapshots of training configurations at each composition state points are provided below the plot, with N atoms in blue and C atoms in cyan and the corresponding composition, temperature, and label ("case") given in the figure. Connections are drawn between atoms within 1.8 Å of one another.

materials project database[52], comprising cell optimization trajectories under 0, 5, 10, 20 and 40 GPa. Note that the former two structures have been observed experimentally[53–55].

Models were generated using version the 2.0.0 of the `ChIMES-LSQ` package[8]. Fitting was automated via version 2.0.0 of the ChIMES Active Learning Driver, `ALDriver`[6,17] an open source python workflow tool for automated ChIMES model generation via iterative fitting. ChIMES simulations were conducted using `ChIMES_MD`, which is a part of the ChIMES-LSQ package.The Standard iterative learning strategy coupled with the newly implemented hierarchical learning capability was used for the present work. Briefly, in the iterative learning strategy, an initial (It-1) model is trained to the DFT-MD-generated training set. The model is then deployed in parallel simulations at a selection of the training compositions, temperatures, and densities. Generally, early models are not adequately informed by the available training data and can give rise to unstable simulations that frequently sample poorly informed regions

Table 1 | Model hyperparameters including the inner cutoff ($r_{cut,in}$), 2-, 3-, and 4-body outer cutoffs ($r_{cut,out,2b}$, $r_{cut,out,3b}$, and $r_{cut,out,4b}$, repectively), and Morse variable for distance transformation (λ)

	CC	NN	CN
$r_{cut,in}$	0.98	0.86	0.90
$r_{cut,out,2b}$	5.00	8.00	5.00
$r_{cut,out,3b}$	5.00	5.00	5.00
$r_{cut,out,4b}$	4.50	4.00	4.50
λ	1.40	1.09	1.34

All values are given in Å.

of the model (e.g., near the inner cutoff) and do not conserve the appropriate quantity. Our active learning strategy[33] attempts to select up to 20 such configurations from each simulation, as well as 20 configurations from otherwise stable portions of the simulation. These ChIMES-generated configurations are then assigned labels (forces, energies, and stresses) via single-point DFT calculation, and then added to the training set, from which the next iteration model is generated; this is repeated for a user-specified number of iterative learning cycles. A total of 10 cycles were used in the present work. A weighting factor of $w = n_I/I$ is applied to each training point, where n_I is the total number of requested iterative learning cycles and I is the current cycle, counting from 1. This has the effect giving DFT-MD generated configurations highest priority weights, which prevents the unphysical configurations generated by early ChIMES models (e.g., that may have extremely small interatomic distances) from driving the fit away from relevant physicochemical space. We note that the need for this weighting strategy arises from our desire to generate maximally efficient models, which means that at our target level of model complexity we may not be able to simultaneously describe near and extremely far from optimal structures equally well. Instead, by applying this decaying weighting scheme, we maintain importance of "ground truth" DFT configurations and ensures models converge with subsequent iterations while still adding the benefit of "rare event sampling" and longer-time scales accessible to ChIMES-based MD.

As in previous work[33], initial weights were set to $w_F = 1.0$ kcal mol^{-1} Å, $w_E = 0.3$ kcal mol^{-1}, and $w_\sigma = 100.0$ kcal mol^{-1} Å^{-3}. These weights account for differences in the relative abundance and magnitude of the force, energy, and stress tensor training data. Models contained 1- through 4-body interactions, with corresponding polynomial orders of $\mathcal{O}_{2b} = 25$, $\mathcal{O}_{3b} = 10$, and $\mathcal{O}_{4b} = 4$. Remaining hyperparameters for the fits were selected using previously described ChIMES heuristic approaches[33,36] and are given in Table 1. All ChIMES simulations were run using either the `ChIMES_MD` code available in the `CHIMES_LSQ` repository[8], or via `LAMMPS`[56] through version 2.0.0 of the `ChIMES_Calculator` Library[57]. Simulations used a 0.2 fs time step and were run for up to 100 ps, however all analysis was performed on only the first 5 ps to enable consistent comparison with DFT. Additional details on use of these tools is provided in Supplementary Section I.

To assess efficacy of the proposed hierarchical transfer learning capability, three models were generated, henceforth referred to as "Standard", "Hierarch," and "Partial" for models fit using the standard a-la-carte approach, the new hierarchical transfer learning strategy, or a partial hierarchical strategy where two parameter blocks are learned simultaneously, respectively. We begin by comparing model performance relative to DFT, for C/N systems, and then extend our study to pure C and pure N.

Performance for C/N systems

Though one might intuitively expect the Standard and Partial models to outperform the Hierarchical model for the C/N system, we find that in general, all models perform equally well for this system. Thus, in this subsection, we will only present data from the Standard model when it shows significant deviations from the Hierarchical model.

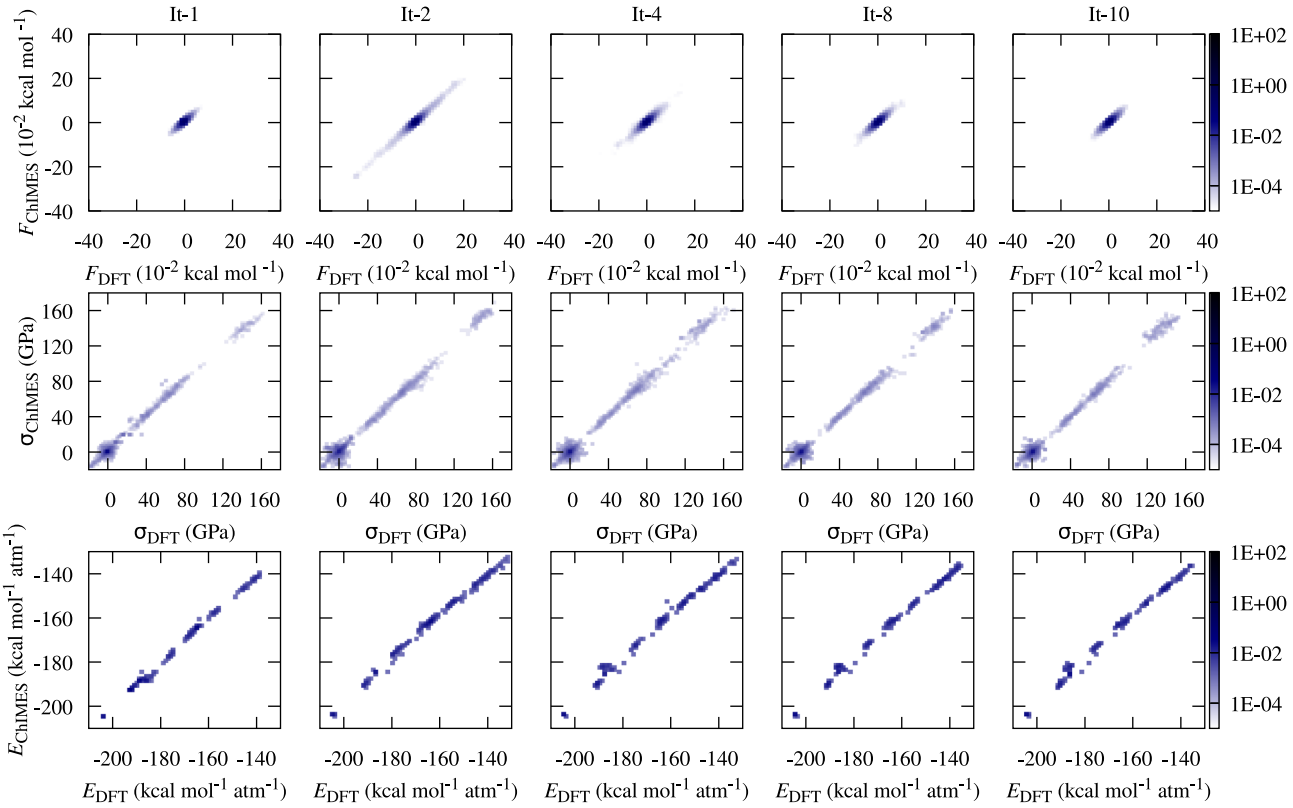

Fig. 3 | Force, stress, and energy parity plots for successive active learning iterations of the Hierarch model. Data in each plot represent only new training data added at each cycle and are given in terms of point density as indicated in the color bar.

We begin with discussion of numerical metrics. The complete Hierarch C/N model contains a total of 3026 parameters. Of those parameters, the 442 for C and 462 for N were fixed, taken from earlier work, and only the remaining 2122 corresponding to CN cross-interactions were fit. Figure 3 provides force, energy, and stress parity plots for fitting iterations 1, 2, 4, 8, and 10 of 10. Note that each plot only shows *new* training data introduced at that iteration. Overall, we find excellent agreement with DFT. Distributions of stress and energy remain relatively unchanged between iterations but there is a clear increase in the spread of forces at It-2, arising from configurations generated by the poorly-constrained It-1 ChIMES model. By It-4, the DFT generated range and distribution of forces (i.e., in It-1) are recovered by the ChIMES model; by It-10, models yield simulations for which the relevant quantity is conserved (see Supplementary Information Fig. 1).

Moving on to physical property metrics, we find that pressure predictions at each training state point are within error of the DFT-predicted value (see Supplementary Information Fig. 2), while C/N crystal cold compression curves exhibits absolute percent differences ranging from 0.1 to 1.7% (see Supplementary Information Table II). Diffusion coefficients are also in good agreement with DFT (See Supplementary Fig. 3), though ChIMES models underpredict these values for the two graphitic structures comprising cases 1 and 2. This disagreement is due to use of $r_{cut,out} \leq 5$, which precludes recovery of the low-lying dispersion forces that modulate inter-sheet interactions, but greatly reduces the model's computational expense. Ongoing work is exploring overcoming this limitation by explicitly including D2 corrections in the ChIMES model[26]. Radial pair distribution functions (RDFs) and vibrational power spectra for simulations using the Hierarch model are provided in Fig. 4. In general, we find excellent agreement with DFT, despite the diversity of structure, chemistry, and bonding across the investigated state points.

To further quantify chemical evolution in each system, we determine mole fractions and corresponding lifetimes for atomic species and small molecules (C, N, N_2, N_3, and C_2N_2) observed in each of the 8 reactive simulations (i.e., with $T > 300$ K). A comprehensive overview of data is given in Supplementary Information Figs. 4, 5. As shown in Fig. 5, we once again find excellent agreement with DFT, consistent with the accuracy typical for an a-la-carte model[28,33,38,42]. Species lifetimes are also in good agreement with DFT. Notably, the ChIMES simulations indicate a large spread in lifetimes for case-5 (1500 K, 1 g cm^{-3}, 50 %N). This is due to coupling between low [N], low density, and short timescales, which makes ensuing chemistry sensitive to simulation initialization (e.g., structure and initial velocities).

Performance for pure C and N systems

In the previous section, performance of models trained *on* C/N data was evaluated for *predicting* CN data. In this section, we ask how well these models perform when predicting properties of pure C and pure N systems. While reading this section, recall the following:

- The Standard model attempts to learn pure-C and pure-N, and C-N cross interaction parameters simultaneously from C/N training data.
- The Partial model uses C parameters that were trained on C data, and attempts to learn pure-N and C-N cross interaction parameters from C/N training data.
- The Hierarch model uses C and N parameters that were trained on just C and just N data, respectively, and attempts to learn only C-N cross interaction parameters from C/N training data.

Hence, the Hierarch model will represent best possible performance for both pure-C and pure-N and the Partial will yield the exact same performance as the Hierarch model for pure-C, since it uses the exact same C parameters. We note that results for the Standard and Partial fit models are taken from a single independent simulation, and that results are only presented when they are found to vary significantly between model development strategies.

Beginning with analysis of performance for Pure C systems, comparing the Standard and Hierarch models, we find that predicted pressures, RDFs, and vibrational power spectra are of similar accuracy (see Supplementary Information Table III and Supplementary Fig. 6). Particularly notable

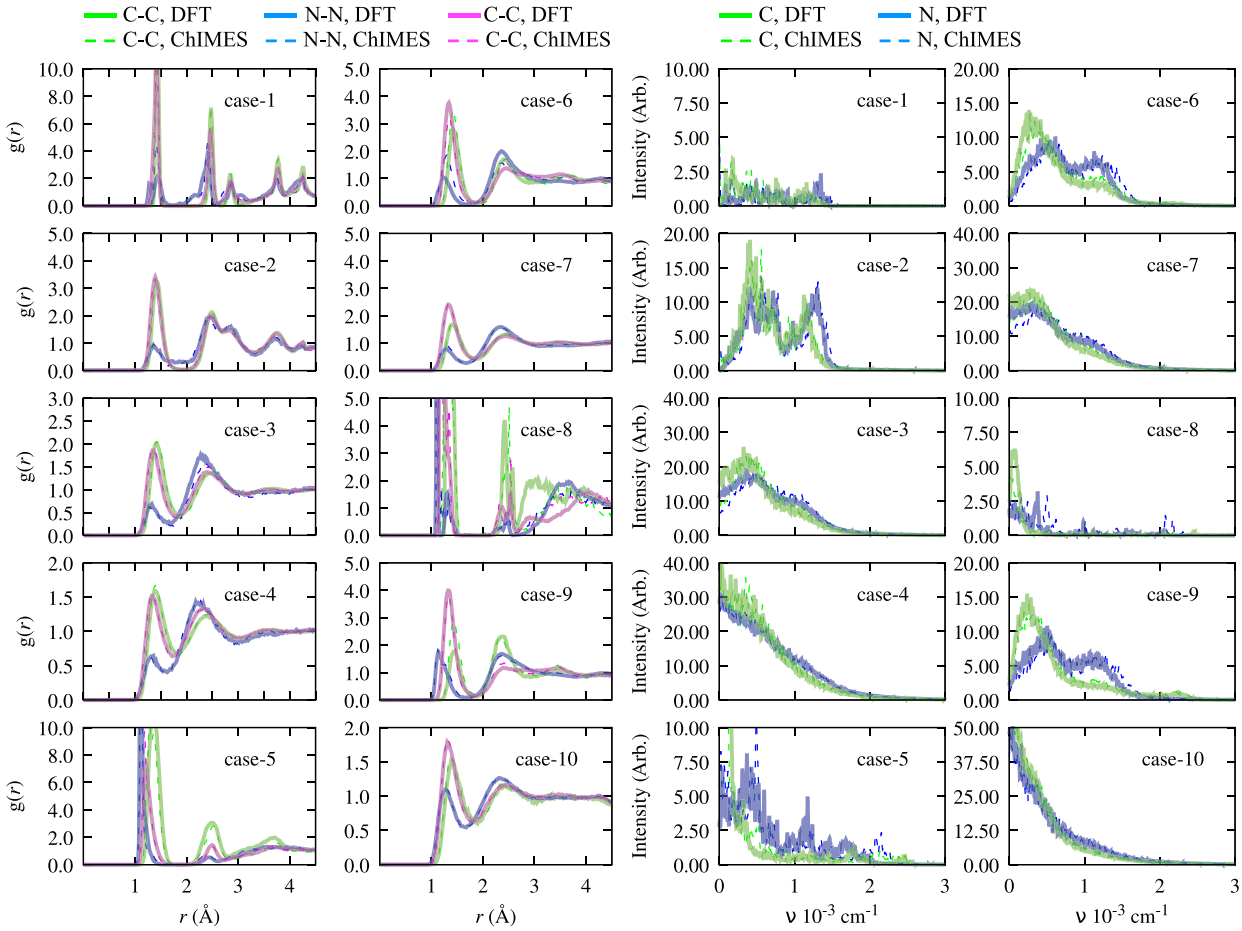

Fig. 4 | RDFs (left) and corresponding vibrational power spectra (right) predicted by DFT (thick solid lines) and the Hierarchically learned ChIMES model (thin dashed lines). In the RDF figures, C–C, N–N, and C–N are given in green, blue, and magenta, respectively. In the power spectra, data for C and N are given in green and blue, respectively.

exceptions to this are shown in Fig. 6, where for low density, high temperature carbyne-like state points, the Standard model yields poor recovery of the corresponding RDFs. This result is unsurprising, since the system structures are dissimilar from anything in the C/N training set (see Fig. 2). Additionally, we find that diffusion coefficients for C at 6000 K and 2.5 g cm^{-3} and 7000 K and 2 g cm^{-3} are significantly over predicted relative to DFT (see Fig. 7). As shown in Table 2, we find the Standard model performs notably worse when predicting diamond and graphite lattice parameters, which is surprising since the C/N training data contains both graphite-like configurations and high density liquid. In particular, our 2017 ChIMES-Carbon model[7] was trained to only a single liquid carbon state point, yet produced a significantly improved diamond lattice parameter (i.e., a = 3.565 Å).

Whereas in the paragraph above focused on pure C systems performance of the Hierarch and Partial models are expected to be identical (i.e., since they use the same underlying C-block parameters), performance for pure N systems will vary. Hence here, we compare performance of the Standard, Hierarch, and Partial models against DFT. In general, we find that all models yield good recovery of the DFT equation of state (Supplementary Table IV), RDFs, vibrational power spectra (Fig. 8), diffusion coefficients, and mole fractions and lifetimes for species formed (Fig. 9), with the Hierarch model yielding slightly better results, just as was seen for the pure C in the previous section. Notable exceptions to this performance include predicted pressure at 8000 K, 4.5 g cm^{-3}, where DFT and the Hierarch model are in good agreement with P = 204.67 and 202.4 GPa respectively, versus the Standard and Partial, which over predict this value by approximately 60 GPa (i.e., with P = 257.20 and 259.4, respectively). The 300 K diffusion coefficient is also significantly underestimated by the Standard and Partial models. Small deficiencies are also observed in RDF and vibrational power spectra for the Standard and Partial Hierarch models. Namely, the 300 K, 1 g cm^{-3} RDF peak at $r \approx 3.5$ Å is sharper shifted to larger r relative to DFT, and the and 6000 K, 2.5 g cm^{-3} power spectrum, which is missing the N_2 vibration peak near 2250 cm^{-1}, and exhibits non-zero vibrations between 1000 and 1500 cm^{-1} unseen in the DFT data.

Discussion

In this work, a new hierarchical transfer learning strategy for development of explicitly chemically extensible ML-IAM was demonstrated. Strategy efficacy was evaluated by comparing performance of models fit using the new hierarchical, partial-herarchical, and standard strategies. We provide key findings from this work below, noting that these insights extend beyond the C/N system studied here. In particular, we find that multielement models fit through the standard ChIMES strategy can perform well across a broad range of T, P, and composition. These models can also extrapolate reasonably well to, e.g., single component properties so long as T and P are not at the training limits. The hierarchical training strategy produces models for multielement systems that are as good as those developed using the standard strategy, yet outperform those models for, e.g., single element properties. At the same time, partial hierarchical learning is also effective for cases where pre-existing models exist for some portion of the fitting problem, but for which the user does not wish to fit remaining parameters in multiple separate steps (e.g., like the present Partial model). We note that efficacy of this transfer learning strategy depends on suitability of the

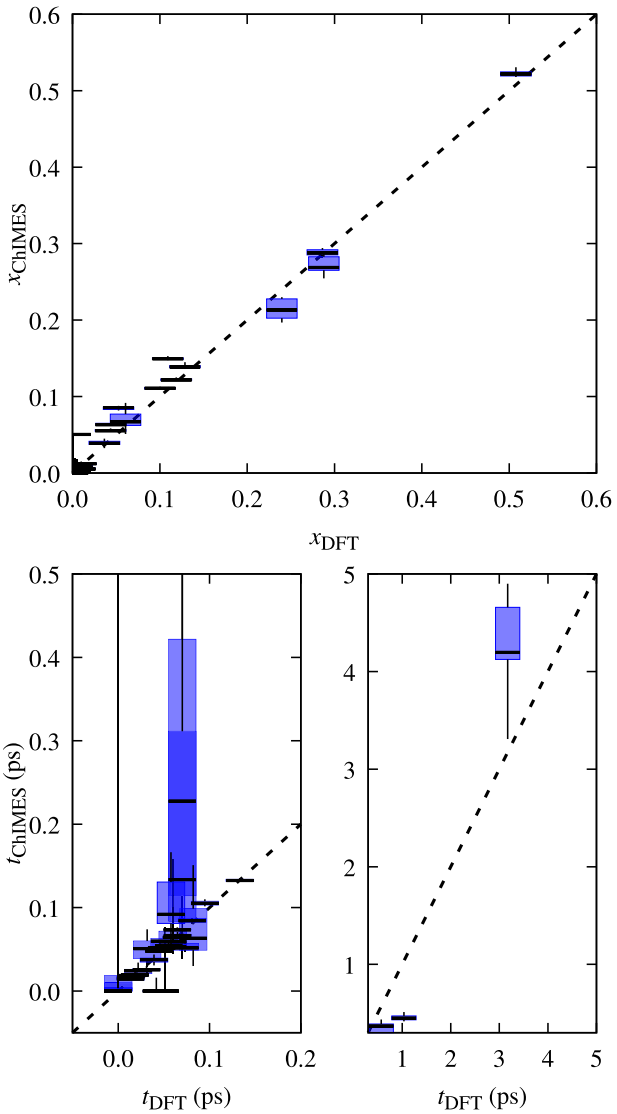

Fig. 5 | Parity plot comparing molfractions (top) and corresponding lifetimes (bottom) for small molecules and atomic species found in each case simulation, predicted by DFT and the Hierarchically learned ChIMES model. The whisker plot gives the maximum, minimum, and 1st through 3rd quartiles based on predictions from 8 independent ChIMES simulations at each case.

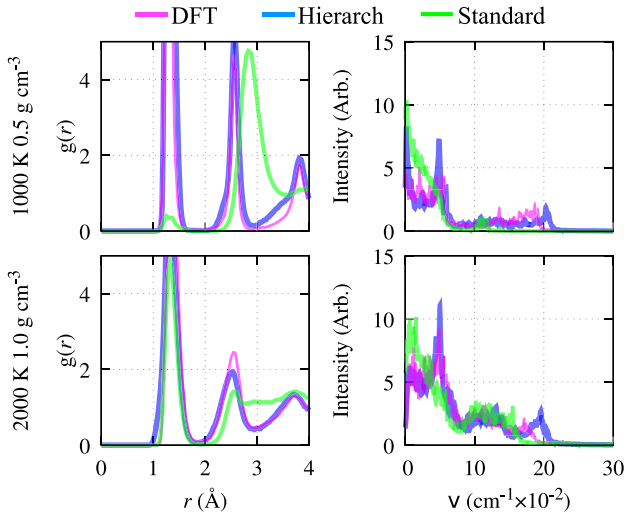

Fig. 6 | Radial pair distribution functions and corresponding vibrational power spectra for the pure carbon system at selected temperatures and pressures. Data from DFT, the 2024 ChIMES carbon model (used by the Hierarchically and Partially-Hierarchically learned models), and the traditionally fit (Standard) models are given in blue, magenta, and green, respectively. DFT and 2024 ChIMES carbon model data are adapted from ref. 26.

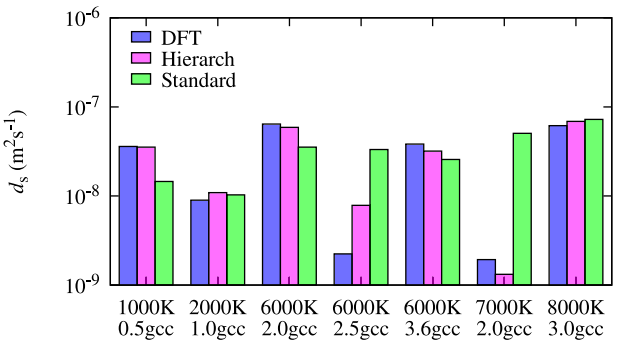

Fig. 7 | Diffusion coefficients for liquid carbon at selected temperatures and pressures. Data for DFT and 2024 ChIMES carbon model (used by the Hierarchically and Partially-Hierarchically learned models) are adapted from ref. 26.

Table 2 | Carbon lattice parameters (Å) for carbon predicted by DFT, the 2024 ChIMES carbon model (used by the Hierarchically and Partially-Hierarchically learned models), and the traditionally fit (Standard) models

		DFT	Hierarch	Standard
Diamond	*a*	3.565	3.569	3.528
Graphite	*a*	2.466	2.465	2.445
Graphite	*c*	6.391	6.521	6.173

DFT and 2024 ChIMES carbon model data are adapted from ref. 26.

training sets used to generate the transferred parameters for the target application.

Ultimately, our new hierarchical transfer learning approach provides a means of breaking down high complexity fitting problems into smaller, more manageable tasks. Critically, this strategy can make ML-IAM development for high complexity systems more accessible to users with limited access to extreme computing infrastructure. Moreover, it also facilitates comprehensive validation, i.e., ensuring that the model will yield quality predictions for all compositions realizable within the hierarchically assembled model space. Future work will investigate whether this strategy can also reduce training data requirements.

Methods

All DFT calculations were performed using version 5.4.4 of the VASP software package[58–61]. Interactions were described through the Perdew–Burke–Ernzerhof generalized gradient approximation functional[62,63], projector-augmented wave pseudopotentials[64,65] (PAW), and the DFT-D2 method[66] for description of dispersion interactions, which has previously been shown to be well suited for describing C and N containing materials under extreme conditions[27,50,67]. The plane wave cutoff was set to 1000 eV, which is required for accurate stress tensor calculation. All reported calculations were spin-restricted; we note that, consistent with other studies of C, H, O, and N-containing materials shocked to comparable conditions[43,68] spin-restricted and -unrestricted calculations yielded similar results. DFT simulations were run in the canonical ensemble with a 0.5 fs timestep, for 5 ps. Electronic eigenstates were occupied according to the Fermi–Dirac distribution with the electronic temperature set equivalent to the target ionic temperature, enabling treatment of ionization and

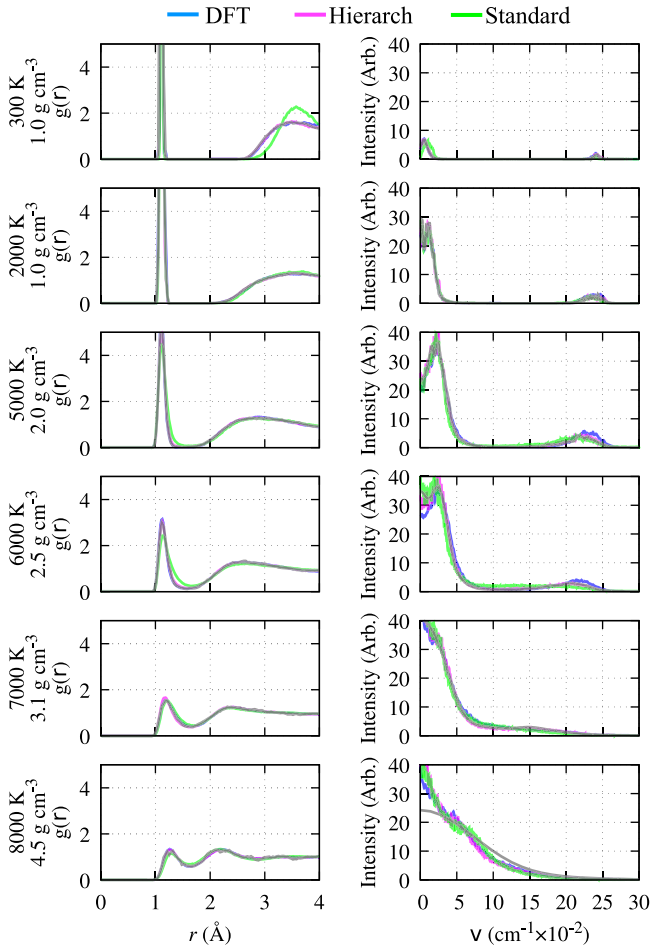

Fig. 8 | RDFs (left) and corresponding vibrational power spectra (right) for nitrogen at state points indicated in the RDF *y*-axis. The plots provide predictions from DFT (blue) as well as ChIMES models generated using the full Hierarchical learning (dashed magenta) and Standard strategy (green). Data for DFT and the Standard fit ChIMES model predictions are adapted from ref. 33.

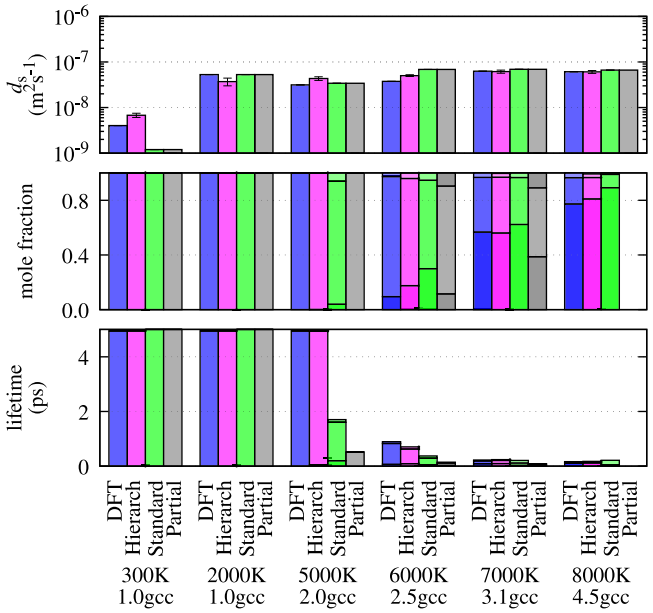

Fig. 9 | Diffusion coefficients, mole fractions, and corresponding lifetimes for nitrogen at the conditions indicated under the *x*-axis. Mole fractions and lifetimes are given for N_1, N_2, and N_3 within the same bar, in progressively more transparent colors with N_1 at the bottom and N_3 at the top. Data for DFT and the Standard fit ChIMES model predictions are adapted from ref. 33. Note that lifetimes predicted by the Standard and Partial model at 5000 K, 2.0 g cm^{-3} is within the spread of values predicted by the Hierarch model.

excitation. Simulation cells were selected to be greater than twice the models' outer cutoffs, i.e., large enough to be sampled accurately at the gamma point. We note that all state points with $T > 300$ K exhibited reactivity; hence, speciation of C, N, N_2, N_3, and C_2N_2 were tracked for those state points.

Data availability

The following files are available free of charge:

- Supplementary_Information.pdf: Additional validation for the models developed in this work
- Supplementary_Information.tar.gz: Training set and active learning files needed to generate the Hierarchically transfer-learned model.

We note that parameters for the Hierarchically transfer-learned model[68] are available in the `ChIMES_Calculator` GitHub repository, under `serial_interface/tests/force_fields/published_params.CN-hierarch.2+3+4b.Tersoff.txt`.

Code availability

All work reported herein was conducted using version 2.0.0 of `CHIMES_LSQ`, the `ChIMES_Calculator`, and the `ChIMES_Active_Learning_Driver`, available at https://github.com/rk-lindsey/.

Received: 10 November 2024; Accepted: 3 November 2025;
Published online: 09 January 2026

References

1. Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133 (1965).
2. Dusson, G. et al. Atomic cluster expansion: Completeness, efficiency and stability. *J. Comput. Phys.* **454**, 110946 (2022).
3. Cusentino, M. A., Wood, M. A. & Thompson, A. P. Explicit multielement extension of the spectral neighbor analysis potential for chemically complex systems. *J. Phys. Chem. A* **124**, 5456–5464 (2020).
4. Plimpton, S. Recent developments in lammps. https://doi.org/10.2172/1889354 (2021).
5. Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134**, 074106 (2011).
6. Lindsey, R., Fried, L. E., Goldman, N. & Bastea, S. Active learning for robust, high-complexity reactive atomistic simulations. *J. Chem. Phys.* **153**, 134117 (2020).
7. Lindsey, R. K., Fried, L. E. & Goldman, N. ChIMES: A force matched potential with explicit three-body interactions for molten carbon. *J. Chem. Theory Comput.* **13**, 6222–6229 (2017).
8. https://github.com/rk-lindsey/chimes_lsq.
9. Wood, M. A. & Thompson, A. P. Extending the accuracy of the snap interatomic potential form. *J. Chem. Phys.* **148**, 241721 (2018).
10. https://github.com/FitSNAP/FitSNAP.
11. Linfeng, Z., Jiequn, H., Han, W., Roberto, C. & Weinan, E. Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics. *Phys. Rev. Lett.* **120**, 143001 (2018).
12. https://github.com/deepmodeling/deepmd-kit.
13. Vandermause, J. et al. On-the-fly active learning of interpretable bayesian force fields for atomistic rare events. *npj Comput. Mater.* **6**, 1–11 (2020).
14. https://github.com/mir-group/flare.

15. Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **104**, 136403 (2010).
16. https://github.com/libAtoms/QUIP.
17. https://github.com/rk-lindsey/al_driver.
18. Wang, H., Zhang, L., Han, J. & Weinan, E. Deepmd-kit: A deep learning package for many-body potential energy representation and molecular dynamics. *Comput. Phys. Commun.* **228**, 178–184 (2018).
19. Novikov, I. S., Gubaev, K., Podryabinkin, E. V. & Shapeev, A. V. The mlip package: moment tensor potentials with mpi and active learning. *Mach. Learn.* **2**, 025002 (2020).
20. https://mlip.skoltech.ru/download/.
21. Zuo, Y. et al. Performance and cost assessment of machine learning interatomic potentials. *J. Phys. Chem. A* **124**, 731–745 (2020).
22. Smith, J. S., Isayev, O. & Roitberg, A. E. Ani-1: an extensible neural network potential with dft accuracy at force field computational cost. *Chem. Sci.* **8**, 3192–3203 (2017).
23. Pham, C. H., Lindsey, R. K., Fried, L. E. & Goldman, N. High-accuracy semiempirical quantum models based on a minimal training set. *J. Phys. Chem. Lett.* **13**, 2934–2942 (2022).
24. Smith, J. S. et al. Automated discovery of a robust interatomic potential for aluminum. *Nat. Commun.* **12**, 1–13 (2021).
25. https://doi.org/10.26434/chemrxiv-2022-15ct6-v3.
26. Lindsey, R.K. et al. ChIMES carbon 2.0: A transferable machine-learned interatomic model harnessing multifidelity training data. *npj. Comput. Mater.* **11**, 3332 (2025).
27. Lindsey, R. K., Bastea, S., Goldman, N. & Fried, L. E. Investigating 3, 4-bis (3-nitrofurazan-4-yl) furoxan detonation with a rapidly tuned density functional tight binding model. *J. Chem. Phys.* **154**, 164115 (2021).
28. Lindsey, R. K., Huy Pham, C., Goldman, N., Bastea, S. & Fried, L. E. Machine-learning a solution for reactive atomistic simulations of energetic materials. *Propellants, Explos. Pyrotech.* **47**, e202200001 (2022).
29. Smith, J. S. et al. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Commun.* **10**, 1–8 (2019).
30. Martin, M. G. & Siepmann, J. I. Transferable potentials for phase equilibria. 1. united-atom description of n-alkanes. *J. Phys. Chem. B* **102**, 2569–2577 (1998).
31. Jorgensen, W. L., Maxwell, D. S. & Tirado-Rives, J. Development and testing of the opls all-atom force field on conformational energetics and properties of organic liquids. *J. Am. Chem. Soc.* **118**, 11225–11236 (1996).
32. MacKerell Jr, A. D., Wiorkiewicz-Kuczera, J. & Karplus, M. An all-atom empirical energy function for the simulation of nucleic acids. *J. Am. Chem. Soc.* **117**, 11946–11975 (1995).
33. Lindsey, R. K. et al. Chemical evolution in nitrogen shocked beyond the molecular stability limit. *J. Chem. Phys.* **159**, 084502 (2023).
34. Lindsey, R. K., Goldman, N., Fried, L. E. & Bastea, S. Many-body reactive force field development for carbon condensation in c/o systems under extreme conditions. *J. Chem. Phys.* **153**, 054103 (2020).
35. Goldman, N., Aradi, B., Lindsey, R. K. & Fried, L. E. Development of a multicenter density functional tight binding model for plutonium surface hydriding. *J. Chem. Theory Comput.* **14**, 2652–2660 (2018).
36. Lindsey, R. K., Fried, L. E. & Goldman, N. Application of the ChIMES force field to nonreactive molecular systems: Water at ambient conditions. *J. Chem. Theory Comput.* **15**, 436–447 (2019).
37. Goldman, N. et al. Semi-automated creation of density functional tight binding models through leveraging chebyshev polynomial-based force fields. *J. Chem. Theory Comput.* **17**, 4435–4448 (2021).
38. Pham, C. H., Lindsey, R. K., Fried, L. E. & Goldman, N. Calculation of the detonation state of hn3 with quantum accuracy. *J. Chem. Phys.* **153**, 224102 (2020).
39. Efron, B., Hastie, T., Johnstone, I. & Tibshirani, R. Least angle regression. *Ann. Stat.* **32**, 407–499 (2004).
40. Friedman, J., Hastie, T. & Tibshirani, R. Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**, 1 (2010).
41. Tibshirani, R. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B* **58**, 267–288 (1996).
42. Lindsey, R. K., Goldman, N., Fried, L. E. & Bastea, S. Many-body reactive force field development for carbon condensation in C/O systems under extreme conditions. *J. Chem. Phys.* **153**, 054103 (2020).
43. Steele, B. A., Bastea, S. & Kuo, I.-F. W. Ab initio structural dynamics of pure and nitrogen-containing amorphous carbon. *Sci. Rep.* **13**, 19657 (2023).
44. Bagge-Hansen, M. et al. Detonation synthesis of carbon nano-onions via liquid carbon condensation. *Nat. Commun.* **10**, 3819 (2019).
45. Jia, R. et al. Synthesis of highly nitrogen-doped hollow carbon nanoparticles and their excellent electrocatalytic properties in dye-sensitized solar cells. *J. Mater. Chem.* **20**, 10829–10834 (2010).
46. Ma, G. et al. Nitrogen-doped hollow carbon nanoparticles with excellent oxygen reduction performances and their electrocatalytic kinetics. *J. Phys. Chem. C.* **115**, 25148–25154 (2011).
47. Bhattacharjya, D. et al. Nitrogen-doped carbon nanoparticles by flame synthesis as anode material for rechargeable lithium-ion batteries. *Langmuir* **30**, 318–324 (2014).
48. Huang, H. et al. One-pot green synthesis of nitrogen-doped carbon nanoparticles as fluorescent probes for mercury ions. *RSC Adv.* **3**, 21691–21696 (2013).
49. Armstrong, M. R. et al. Ultrafast shock synthesis of nanocarbon from a liquid precursor. *Nat. Commun.* **11**, 1–7 (2020).
50. Dattelbaum, D. et al. Carbon clusters formed from shocked benzene. *Nat. Commun.* **12**, 1–9 (2021).
51. Jain, A. et al. The Materials Project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002 (2013).
52. Teter, D. M. & Hemley, R. J. Low-compressibility carbon nitrides. *Science* **271**, 53–55 (1996).
53. Liu, A. Y. & Wentzcovitch, R. M. Stability of carbon nitride solids. *Phys. Rev. B* **50**, 10362 (1994).
54. Marqués, M., Osorio, J., Ahuja, R., Flórez, M. & Recio, J. Pressure effects on the structure and vibrations of β-and γ- c 3 n 4. *Phys. Rev. B* **70**, 104114 (2004).
55. Plimpton, S. Fast parallel algorithms for short-range molecular dynamics. *J. Phys. Chem. C.* **117**, 1–19 (1995).
56. https://github.com/rk-lindsey/chimes_calculator.
57. Kresse, G. & Hafner, J. Ab initio molecular dynamics for liquid metals. *Phys. Rev. B Condens. Matter Mater. Phys.* **47**, 558 (1993).
58. Kresse, G. & Hafner, J. Ab initio molecular-dynamics simulation of the liquid-metal–amorphous-semiconductor transition in germanium. *Phys. Rev. B Condens. Matter Mater. Phys.* **49**, 14251 (1994).
59. Kresse, G. & Furthmüller, J. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **6**, 15–50 (1996).
60. Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B Condens. Matter Mater. Phys.* **54**, 11169 (1996).
61. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**, 3865 (1996).
62. Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized gradient approximation made simple [erratum to phys. rev. lett. 77, 3865 (1996)]. *Phys. Rev. Lett.* **78**, 1396–1396 (1997).

63. Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B Condens. Matter Mater. Phys.* **50**, 17953 (1994).
64. Kresse, G. & Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B Condens. Matter Mater. Phys.* **59**, 1758 (1999).
65. Grimme, S. Semiempirical GGA-type density functional constructed with a long-range dispersion correction. *J. Comput. Chem.* **27**, 1787–1799 (2006).
66. Lindsey, R. K., Goldman, N., Fried, L. E. & Bastea, S. Chemistry-mediated ostwald ripening in carbon-rich c/o systems at extreme conditions. *Nat. Commun.* **13**, 1–7 (2022).
67. Mundy, C. J. et al. Ultrafast transformation of graphite to diamond: An ab initio study of graphite under shock compression. *J. Chem. Phys.* **128**, 184701 (2008).
68. https://github.com/rk-lindsey/chimes_calculator/tree/main/serial_interface/tests/force_fields.

Acknowledgements

This work has been primarily supported by the DOE, Office of Science, BES, under Award #DE-SC0022305 (Formulation engineering of energy materials via multiscale learning spirals). This work was performed in part under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. The project is assigned release number LNL-JRNL-864545.

Author contributions

R.K.L. conceived of and implemented the transfer learning framework method, developed the models, conducted the ChIMES simulations, and contributed to DFT simulations. A.D.O. contributed to model development and simulations. I.F.W.K. and B.A.S. contributed to generation of initial DFT training set. I.F.W.K. and S.B. contributed to identification of the application space. S.B. and N.G. contributed to interpretation of results. All authors contributed to scientific discussions and manuscript development.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41524-025-01863-4.

Correspondence and requests for materials should be addressed to Rebecca K. Lindsey.