

Gradient-based optimization of complex nanoparticle heterostructures enabled by deep learning on heterogeneous graphs

Received: 22 December 2024

Accepted: 30 October 2025

Published online: 8 December 2025

 Check for updates

Eric Sivonxay¹, Lucas Attia^{1,2}, Evan Walter Clark Spotte-Smith³, Benjamin Sanchez-Lengeling^{4,6,7}, Xiaojing Xia⁵, Daniel Barter¹, Emory M. Chan⁵✉ & Samuel M. Blau¹✉

Applications of deep learning (DL) to design nanomaterials are hampered by a lack of suitable data representations and training data. Here we report efforts to overcome these limitations and leverage DL to optimize the nonlinear optical properties of core–shell upconverting nanoparticles (UCNPs). UCNPs, which have applications in fields such as biosensing, super-resolution microscopy and three-dimensional printing, can emit visible and ultraviolet light from near-infrared excitations. We report a large-scale dataset of UCNP emission spectra based on accurate but expensive kinetic Monte Carlo simulations ($N > 6,000$) and use these data to train a heterogeneous graph neural network using a physically motivated representation of UCNP nanostructure. Applying gradient-based optimization on the trained graph neural network, we identify structures with $6.5\times$ higher predicted emission under 800-nm illumination than any UCNP in our training set. Our work reveals design principles for UCNP heterostructures and presents a roadmap for DL-based inverse design of nanomaterials.

Applications in clean energy, advanced manufacturing, biomedicine, photonics and microelectronics increasingly demand new materials with complex structures and heterogeneous composition¹. In principle, machine learning (ML) offers a strategy to accelerate the discovery of such materials, as it has emerged as a transformational tool for the design of small molecules, bulk inorganic crystals and even single-component nanomaterials^{2–4}. Deep learning (DL) approaches are particularly well suited to model the behavior of systems with large numbers of parameters, but several obstacles hinder DL from being used to guide the discovery of the complex materials often needed for real-world applications, including nanostructures and composites.

First, state-of-the-art approaches for representing materials often fail to capture the structural complexity of nanomaterials, such as multishell nanoparticles^{5–12} (Fig. 1a), in which nanostructure controls energy transport¹ (Fig. 1b). Nanomaterials exhibit distinct and often superior properties compared with their bulk counterparts^{13–15}; however, the large number of features required to adequately describe a nanoscale material (including surface ligands and the morphology, dimensions, composition and defects of each domain) make training on naive tabular representations computationally inefficient, in part because they neglect physical relationships between features¹⁶. While bulk crystals can be represented by their unit cell coordinates and small organic molecules by atomic coordinates, graphs or strings¹⁷ such as

¹Energy Technologies Area, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ²Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. ³Department of Materials Science and Engineering, Carnegie Mellon University, Pittsburgh, PA, USA. ⁴Google DeepMind, Cambridge, MA, USA. ⁵The Molecular Foundry, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ⁶Present address: Department of Chemical Engineering and Applied Chemistry, University of Toronto, Toronto, Ontario, Canada. ⁷Present address: Vector Institute for Artificial Intelligence, Toronto, Ontario, Canada. ✉e-mail: emchan@lbl.gov; smblau@lbl.gov

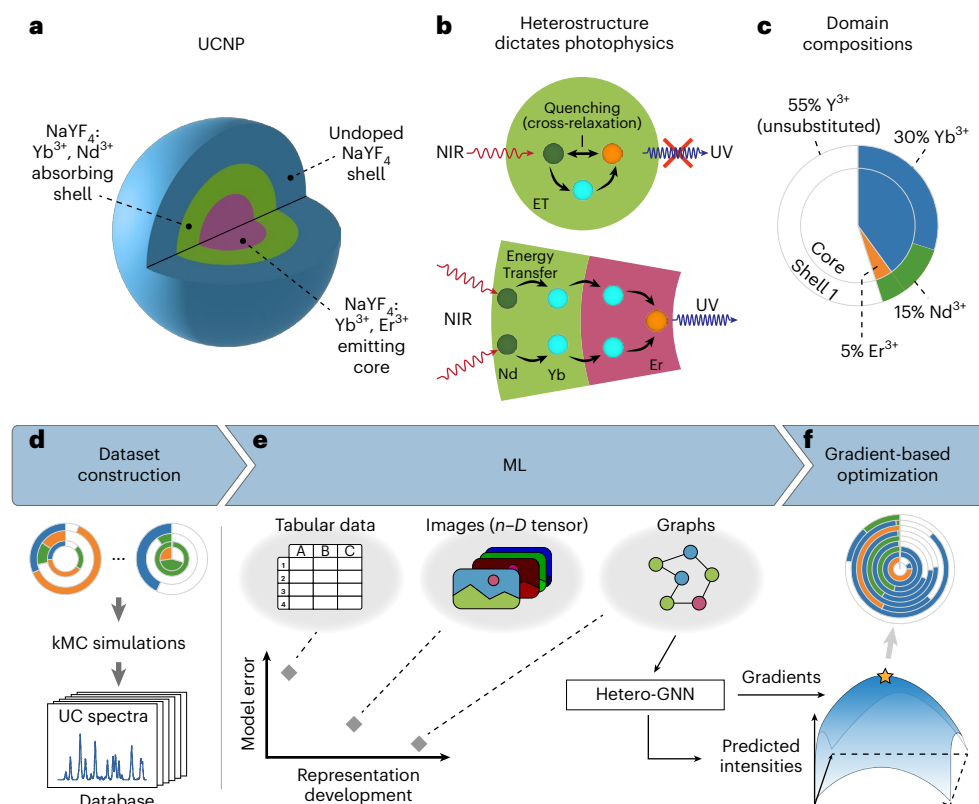


Fig. 1 | DL approach to enable UCNP heterostructure optimization.

a, Schematic of doping heterostructured UCNPs. **b**, Illustration of the importance of heterostructuring to physically separate dopants. In the core-only particle, all dopants interact to a high degree, leading to cross-relaxation and consequently low upconversion efficiency, while heterostructuring with Er (orange) and Nd (green) separated into different regions doped with Yb (blue) prevents cross-relaxation and facilitates UV emission. **c**, Two-dimensional visualization of a heterostructured UCNP with core and shell doping levels for Yb³⁺ (blue), Nd³⁺ (green) and Er³⁺ (orange) represented as colored slices in a concentric pie chart. White regions represent the remaining fraction of Y³⁺ ions in

the NaYF₄ host matrix that have not been substituted by dopant ions. **d**, To enable physics-informed DL for the inverse design of UCNPs, we first construct a large dataset of simulated UCNPs and their upconversion (UC) luminescence spectra using high-throughput kMC simulations. **e**, Then, we investigate different representations of heterostructured UCNPs and find the best performance with a heterogeneous graph representation in a GNN. **f**, The resulting trained model can predict the gradient of emission intensity with respect to each structural parameter, enabling inverse design via gradient-based optimization of UCNP heterostructure to maximize UV emission. $n-D = n$ -dimensional.

the Simplified Molecular Input Line Entry System (SMILES)¹⁸, such atomistic representations are impractical for complex nanomaterials because their critical features often span length scales of one to >10⁶ atoms and cannot necessarily be reduced to periodic subunits^{19,20}. Recent DL approaches encode spatial information as pixels or voxels, but these fixed-resolution representations cannot efficiently capture the structural hierarchy of a wide range of nanomaterials, including those of different sizes.

Beyond the challenge of representing nanomaterials, it is also challenging to generate datasets sufficiently large to train DL models that can accurately predict the properties of heterogeneous, multi-component nanostructures²¹. Although high-throughput experimental and computational approaches are growing in their availability and utility^{22–24}, the synthesis and simulation of complex heterostructures is often time-consuming, limiting the scale of available datasets²⁵ and constraining campaigns to the ‘small data’ regime where DL techniques often struggle. Modern DL models can also have difficulty extrapolating outside of the envelope of their training data, which is necessary for the discovery of novel materials with enhanced properties. Finally, the discovery of fundamentally new materials is complicated by the rough response surfaces of material properties with respect to their composition, necessitating tedious ‘needle-in-a-haystack’ searches across a parameter space. The prediction of materials with targeted properties would be substantially accelerated by surrogate models that are differentiable so that gradient-based optimization techniques can

be used to direct efficient searches³. Thus, DL-guided inverse design of complex nanomaterials would benefit from the development of large, relevant structure–property datasets, new methods to represent them across length scales and differentiable models that are accurate beyond their training distribution.

In this work, we develop a heterogeneous graph representation for nanomaterials with a variable number of spatial domains, each containing multiple components that can interact within the same domain and across interfaces. We demonstrate that graph neural networks (GNNs) built atop such representations can accurately predict properties of nanostructures that are far more complex than any contained in the training dataset. As a model system, we center our investigation on lanthanide-doped upconverting nanoparticle (UCNP) heterostructures (Fig. 1a), whose unique nonlinear optical properties have applications in biological and super-resolution imaging²⁶, optogenetics, sensing, photonics²⁷, three-dimensional printing²⁸ and photovoltaics²⁹. These applications leverage the ability of UCNPs to absorb multiple near-infrared (NIR) photons and convert them into higher energy radiation, often visible and ultraviolet (UV) light. Such nonlinear processes are the result of complex networks of energy transfer (ET) interactions between different lanthanide ions (including, Yb³⁺, Er³⁺ and Nd³⁺, as in Fig. 1b). To promote advantageous ET interactions and inhibit those that quench emission, nearly all practical implementations of UCNPs use doped heterostructures in which a spherical core is surrounded by one to four concentric shells, with each domain having a distinct

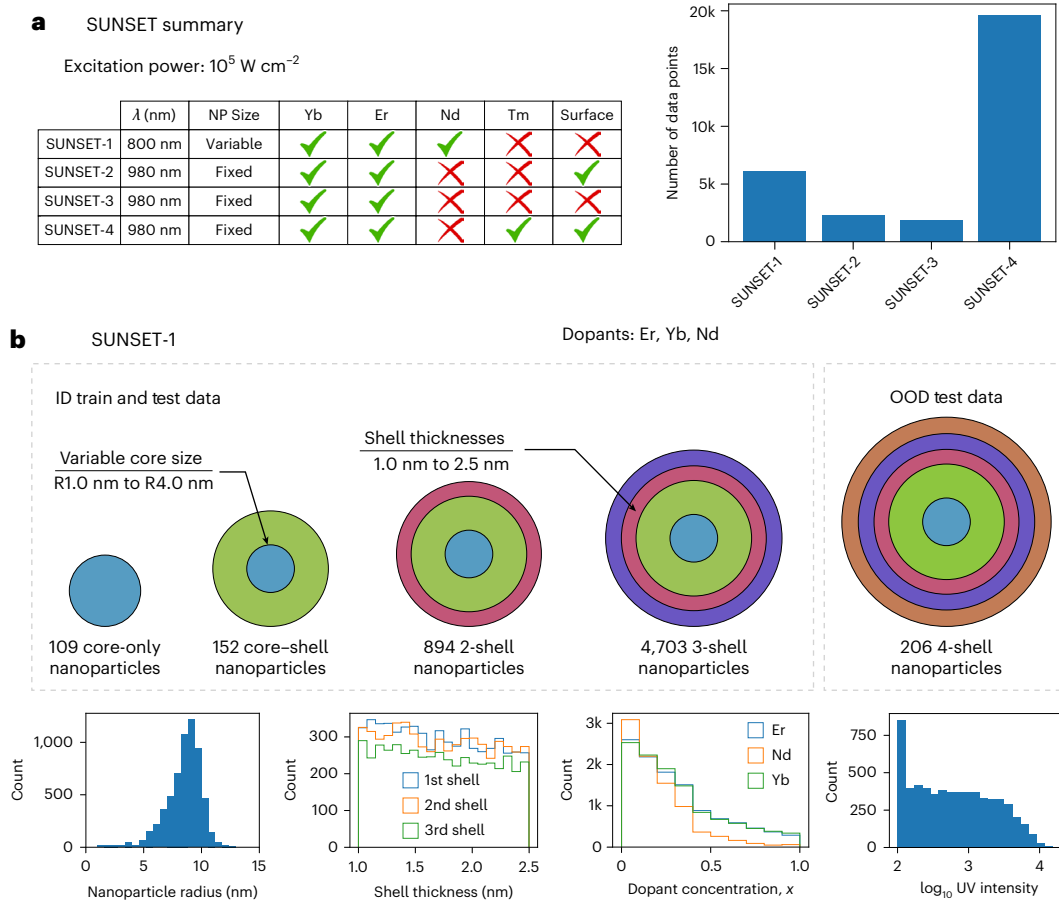


Fig. 2 | Description of SUNSET datasets. **a**, Summary of the SUNSET datasets, including excitation wavelength (λ , nm), fixed versus variable nanoparticle size, dopants used, presence or absence of luminescence-quenching surface ligands, and scale in terms of number of data points. **b**, Top: graphical depiction of SUNSET-1 heterostructures and ID versus OOD split. The OOD test data contain 206 4-shell nanoparticles. Bottom: statistical distributions of structural features

in the SUNSET-1 dataset. From left to right, distributions show nanoparticle radius (R , nm), shell thickness (nm) with colored series for each region (first shell, blue; second shell, orange; third shell, green), dopant concentration (x) with colored series for each dopant type (Er, blue; Nd, orange; Yb, green) and \log_{10} of the UV intensity.

combination and composition of lanthanide ions (Fig. 1c). Due to the large numbers of tunable structural and compositional parameters and the complex network of ET interactions between dopants, optimizing the intensity and wavelength for such complex heterostructures is extremely challenging³⁰. Thus, multishell UCNP heterostructures present a stringent test for any new DL model and representation.

To train these DL models, we generated a dataset of Simulated Upconverting Nanoparticle Spectra for Emission Tuning (SUNSET), consisting of results from ~6,000 kinetic Monte Carlo (kMC) simulations of nanoparticle photophysics (Fig. 1d). Models trained on SUNSET aim to predict photon emission within a specified wavelength band as a function of UCNP heterostructure. By training on simulations of UCNPs with up to three shells and evaluating on simulations with four shells, we can quantify the capacity for models to extrapolate to larger and more complex heterostructures. We find that our heterogeneous graph representation, informed by UCNP physics and geometry, allows DL models to achieve far higher prediction accuracy than tabular, image and homogeneous graph representations (Fig. 1e), especially when extrapolating beyond the training data. The differentiability of our heterogeneous GNN (hetero-GNN) also yields gradients of emission intensity with respect to layer thicknesses and dopant concentrations (Fig. 1e), which are not accessible from kMC. Our trained model thus facilitates inverse design of UCNP heterostructure via gradient-based optimization (Fig. 1f), identifying superior UCNPs with a range of sizes and up to ten shells. When excited at 980 or 800 nm, these optimized

UCNP heterostructures exhibit exceptionally high emission between 300 nm and 450 nm, a spectral range useful for inducing photochemistry for optogenetic, catalytic, therapeutic and three-dimensional printing applications. To validate these predictions, we perform additional kMC simulations, which indicate that our model possesses considerable ability to extrapolate far out of distribution and can suggest never-before-seen structures with accurate predictions of their emission intensity, further revealing valuable design principles. These findings demonstrate a path forward for the optimization and discovery of technologically useful UCNPs and offer inspiration for the development of DL representations and models that enable inverse design for a broad range of optical nanomaterials.

Results

Dataset construction

To develop and train DL models that can predict core-shell UCNP photophysics and ultimately enable inverse design of UCNPs with complex heterostructures that exhibit efficient UV and blue emission (300–450 nm), we generated SUNSET, a dataset of over 30,000 multi-shell UCNP spectra calculated with a high-performance kMC simulation package (RNMC)³¹ optimized for chemical reaction networks and UCNP photophysics⁹ (Fig. 2). SUNSET consists of four subcollections (SUNSET-[1, 2, 3, 4]) that include different dopant ion combinations and surface effects (Fig. 2a). While each of the subcollections provides utility for model development and testing, we focus exclusively on

Table 1 | Model performance on the SUNSET-1 dataset

	RFR	FCNN	CNN	GNN	Hetero-GNN
ID test	57.2 (7.3%)	55.8 (6.6%)	17.6 (2.1%)	84.3 (9.9%)	13.8 (1.6%)
OOD test	365.8 (36.6%)	526.6 (55.5%)	49.0 (5.2%)	89.6 (9.4%)	22.1 (2.3%)

An overview of the performance of different models on the SUNSET-1 dataset, assessed using the MSEs (equation (5)) of predicted \log_{10} of $I_{\text{vis-UV}}$ derived from tenfold cross-validation both for the SUNSET-1 ID test, containing structures with up to three shells, and for the SUNSET-1 OOD test, containing structures with four shells. The first value in each cell is the unitless MSE multiplied by 10^3 for easier interpretation. The second value in parentheses is the NMSE, which is the MSE normalized by the sum of squares to yield a relative error (equation (6)). Bold text indicates the lowest errors for each category, both of which are achieved by the hetero-GNN.

SUNSET-1 in our main narrative; discussion of the SUNSET-[2–4] collections can be found in the Methods and Supplementary Information sections 8 and 12. We focus on SUNSET-1 because this subcollection includes nanoparticles with variable sizes and numbers of shells, which are essential for training models capable of extrapolating to more complex heterostructures and thereby facilitating impactful inverse design. Further, the nanoparticle structures in SUNSET-1 are substantially larger than in SUNSET-[2–4] (given that they contain multiple layers, and each layer must be at minimum 1 nm thick to be synthesizable³²), and thus the vast majority of the computational cost of SUNSET went toward the ~6,000 simulations in SUNSET-1. We note that individual kMC trajectories often took weeks to complete, necessitating the use of high-throughput self-checkpointing workflows (Methods; Supplementary Fig. 1).

SUNSET-1 utilizes a dopant set of Er^{3+} , Nd^{3+} and Yb^{3+} because this combination of dopants has been used to sensitize upconversion⁶ and optogenetic activity³³ with 800-nm excitation, a wavelength that lies in the NIR-I biological imaging window. Segregation of these dopants into different shells of UCNP heterostructures has been shown to dramatically enhance emission⁶. In these systems, Nd is typically included to sensitize the absorption of the 800-nm excitation, Er to upconvert absorbed energy and emit UV or visible light, and Yb to act as a conduit to transfer energy between Nd and Er dopants that would otherwise quench each other via cross-relaxation ET. The nanoparticle heterostructures sampled in SUNSET-1 are variable, with core radii ranging from 1 to 4 nm and up to three shells, each measuring between 1 nm and 2.5 nm in thickness, as depicted in Fig. 2b. We sample the core radius and shell thickness from uniform distributions. Dopant concentrations are also drawn from uniform distributions, with the constraint that their total sum must be less than 1. To probe the extrapolatory power of developed models, we simulate explicit four-shell nanoparticles and hold them out of training data to use as an out-of-distribution (OOD) test set. Thus, SUNSET-1 as a whole has nanoparticle radii spanning from 1 to 13.6 nm, and the brightest particle has an intensity of ~20,000 photon counts per second (cps). $I_{\text{vis-UV}}$ most closely follows an exponential distribution, so we use the unitless $\log_{10}(I_{\text{vis-UV}})$ as the target label for model training.

Representing nanoparticle structure for ML

To train ML models on the SUNSET data, we initially investigated several existing representations for encoding the compositional and dimensional features of each UCNP. As we summarize in Table 1 and discuss in more detail in the ‘Model performance’ section, we found that standard ML models (for example, random forest regressors (RFRs) and convolutional neural networks (CNNs) utilizing tabular and image-based representations exhibited poor ability to extrapolate, with three- to tenfold lower accuracy during OOD testing than during in-distribution (ID) testing.

Seeking representations and models with greater ability to extrapolate to more complex nanostructures, we explored the use of graph-structured representations, which have recently gained prominence due to their ability to effectively capture complex relationships (edges) between entities (nodes)³⁴. The simplest graph representation of a UCNP is a homogeneous graph in which each node is labeled with

the identity and concentrations of a single type of dopant (for instance, Er) in a specific domain of the UCNP (for instance, the core), while edges encode interactions (ET processes) between dopants represented by those nodes. However, we found that GNNs utilizing these homogeneous graph representations exhibited equally poor accuracy for both ID and OOD testing (Table 1).

Reasoning that the poor performance of homogeneous GNNs (homo-GNNs) was related to inadequate representation of the physical interactions between dopants, we developed a UCNP representation based on a directed heterogeneous graph (Fig. 3a). Unlike the homogeneous graphs, dopant–dopant interactions in our heterogeneous graphs are represented by interaction nodes that connect dopant nodes (via edges), allowing the encoding of additional physical features of the interactions. Two different types of interaction nodes are used, intralayer and translayer, to delineate interactions between dopants within the same geometric region (either core or shell domains) and those in different regions, respectively. It is important to note that dopant nodes are never connected to other dopant nodes, and every interaction node connects exactly two dopant nodes. A self-interaction node, describing interactions between different dopant ions of the same type in the same region, has edges both from and to the same dopant node. The use of a directed graph introduces asymmetry in ET between two dopants (for example, $\text{Yb} \rightarrow \text{Er}$ as compared with $\text{Er} \rightarrow \text{Yb}$). This is important for ET processes that are not reversible, such as nonresonant, phonon-assisted energy ET that results in irreversible heat dissipation³⁵.

When establishing features encoded in the nodes, we chose a minimal set of descriptors that are most relevant for UCNPs. Dopant node features include dopant type, dopant concentration (within the respective region) and geometric bounds (inner and outer radii of the core/shell domain they reside in). The interaction nodes contain the interaction type (Yb–Er, Er–Yb, Er–Er and so on) and features derived from the pair of connected dopant nodes—dopant concentrations and geometric bounds. Because we explore only spherical nanoparticles with multiple concentric shells, layer radii fully specify heterostructure geometry. To further account for the effect of distance and region geometry on dopant ETs, we introduce a quantity that we call the ‘integrated interaction’ (see the Methods for an explanation).

The proposed heterogeneous graph structure lends naturally to the use of GNNs for DL. In our hetero-GNN (Fig. 3c), we first embed information from each node into a continuous vector space. We construct the dopant node embeddings by passing the dopant type (Z_i) through an embedding layer, then contextualizing the initial embedding on the dopant concentration (x_i) and the radii (r_i^{inner} and r_i^{outer}) using feature-wise linear modulation (FiLM) layers, as shown in Fig. 3d. Likewise, to obtain the embedding vector for the interaction nodes (Fig. 3e), we pass the interaction type (τ_{ij}) through an embedding layer and then condition on the integrated interaction using a FiLM layer. A batch normalization is applied before the FiLM layer to shift the distribution of integrated interaction values. Note that ‘contextualizing’ or ‘conditioning’ a vector (for instance, an initial embedding) on another value (for example, dopant concentration, layer radii, or integrated interaction) with a FiLM layer is a way of combining the information contained in each via an operation that is controlled by many learnable




Fig. 3 | UCNPs graph representation and hetero-GNN model architecture.

a, Heterogeneous graph structure of an example two-region (core plus one shell) nanoparticle with Yb and Er dopants. Node types: dopant nodes (green circles), translayer interaction nodes (blue triangles) and intralayer interaction nodes (red triangles). Node features: dopant type Z (atomic number), dopant concentration x (molar fraction), inner radius r^{inner} , outer radius r^{outer} and interaction type τ (atomic number of outgoing dopant, atomic number of incoming dopant). **b**, Depiction of the integrated interaction Φ between two regions V_i and V_j over all pairwise distances, s , parameterized by learnable parameters, σ . Subscripts i and j denote different dopant-region indices. For instance, in **a**, Yb in the core, Er in the core, Yb in the shell and Er in the shell would all have distinct dopant-region indices. **c**, Architecture of the heterogeneous

graph model, showing dopant node embeddings h_i^0 , intralayer interaction node embeddings e_{ij}^0 , translayer interaction node embeddings e_{ij}^0 and their updates through heterogeneous convolution layers. The embeddings are combined through mean aggregation into an overall UCNPs embedding which is regressed to the \log_{10} of the UV intensity through a FCNN. **d**, Dopant node embedding process, where the dopant embedding (Z_i) is conditioned on r_i^{inner} , r_i^{outer} and x_i using feature-wise linear modulation (FiLM) generators (Gen.) to form the initial dopant node embedding for dopant-region i , h_i^0 . **e**, Interaction node embedding process, where features r_i^{inner} , r_i^{outer} and x_i for dopant-region i , r_j^{inner} , r_j^{outer} and x_j for dopant-region j , and the interaction type, τ_{ij} , form the initial interaction node embedding between dopant-regions i and j , e_{ij}^0 . Batch normalization (Batch Norm) layers are used in the embedding process to stabilize training.

parameters, which often provides better expressivity and performance than a simple concatenation or addition³⁶. The resulting dopant and interaction embeddings are then used as the inputs for three iterations of message passing (MP) based on the heterogeneous graph's directional edges, where each MP iteration employs graph attention via the graph attention operator GATv2³⁷, after which we use mean aggregation to obtain the global latent representation. Finally, a fully connected neural network (FCNN) is used for predicting the label, the \log_{10} of the emission intensity over the specified wavelength band, from the global latent representation

Model performance

To assess the performance of the hetero-GNN described above with respect to other models and representations, we train each on the SUNSET-1 dataset, where our target label is the unitless \log_{10} of the emission intensity from the UV-blue (300–450 nm) wavelength band, and training data include 800-nm excited UCNPs with zero to three shells. We evaluate the mean squared errors (MSEs) of each model when predicting the intensities of held-out ID samples as well as for

OOD nanoparticles with four shells (Table 1). We arbitrarily multiply our unitless MSE values by 10^3 for easier interpretation. We compare the hetero-GNN with four well-established supervised learning models: two models (a random forest regressor and a FCNN multilayer perceptron) using a tabular representation; a CNN using an image representation; and the homo-GNN described above. Model hyperparameters are provided in Supplementary Information section 9. Additional model, representation and feature details are provided in Supplementary Information sections 4–6.

Of the five models tested, the hetero-GNN exhibited the lowest error for both ID and OOD testing, with MSE values of 13.9 and 22.2, respectively (Table 1). The ID loss is fourfold lower than that of models utilizing tabular representations (RFR and FCNN) and 21.5% lower than the CNN utilizing an image representation. The fact that the image- and heterogeneous graph-based models have the best ID accuracies demonstrates how their representations allow them to leverage spatial information to learn relationships between heterostructure and properties and to connect the common behavior of dopant ions of the same type but located in different regions. It is notable, however, that the least

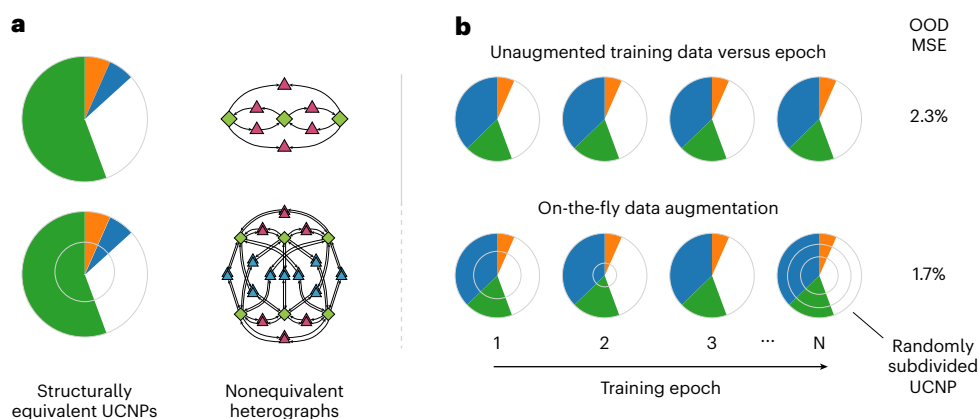


Fig. 4 | Increasing hetero-GNN accuracy using on-the-fly data augmentation to promote subdivision invariance. **a**, Arbitrarily subdividing a structural region in a heterostructured nanoparticle leaves the nanoparticle physically unchanged. Such ‘subdivision invariance’ is not reflected in nanoparticle graph representations. Two-dimensional visualizations of physically equivalent core and arbitrarily subdivided core–shell nanoparticles, along with their nonequivalent heterogeneous graph representations. **b**, Subdivision invariance

is promoted in hetero-GNNs using on-the-fly data augmentation. In contrast to the nonaugmented training procedure (top), on-the-fly data augmentation is performed during each training epoch by inserting a random number of subdivisions into UCNPs at random radial coordinates (bottom). This process implicitly multiplies the number of unique heterograph representations of UCNPs seen by the model during training by the number of training epochs, and reduces OOD MSE by 25%.

accurate model for ID testing was also based on a graph representation. The homo-GNN exhibited sixfold higher MSE than the hetero-GNN, highlighting that the enhanced accuracy of the hetero-GNN is the result of its incorporation of interactions as nodes in the graphs.

We believe that the promotion of dopant–dopant interactions to a node-level property specifically improves hetero-GNN performance by elevating the prominence of interaction features during message passing. When included as explicit nodes, the interactions are able to alter the content of the passed messages, allowing for the transmission of richer and more physically relevant information. This parallels the photophysics of UCNPs, in which ET interactions between dopants critically determine the excited state populations of donors, acceptors and their neighbors^{30,38}, driving nonlinear processes such as upconversion, photon avalanching^{5,39} and quantum cutting⁴⁰. By contrast, interaction features in homo-GNNs have less influence on predictions because edge properties can only contribute to the attention score, influencing the weighting of the messages being passed between dopant nodes rather than the information contained therein. The heterogeneous graph structure also enables the embedding of interactions between lanthanide dopants. This embedding introduces valuable inductive bias concerning the distinctness of lanthanide interactions, constraining the model to treat dopant pairs such as Yb–Er equally, agnostic to which layers they reside in, albeit with varying strengths.

The most striking benefit of the hetero-GNN is its ability to extrapolate, in this case to four-shelled nanostructures not included in its training set. When switching from ID to OOD testing, the MSE for the hetero-GNN model increased by 8.3- or 1.6-fold. This modest increase in loss is in stark contrast to the tabular-representation-based RFR and FCNN models, for which extrapolation resulted in six- to tenfold increases in the MSE, respectively. The tabular models lack the geometric and relational information of the graph models and therefore must learn the influence of the dopants in each layer independently. This increases the data demand of these models, making them prone to overfitting and reducing their ability to predict the properties of unseen heterostructures. Even the image-based CNN, which had high ID accuracy, exhibited 2.8-fold greater loss during OOD testing. We attribute the superior generalization ability of the hetero-GNN to its graph representation. This conclusion is supported by the fact that the homo-GNN also exhibits very little change in loss (+6.5%) when shifting from interpolation to extrapolation.

In summary, representing UCNPs heterostructures as heterogeneous graphs results in a hetero-GNN that exhibits both high accuracy and extrapolative capacity, in contrast to existing models that are inaccurate (homo-GNN), poor at extrapolation (image CNN) or both (tabular models). We have improved the OOD accuracy of these hetero-GNN models by an additional 25% using an on-the-fly data augmentation approach (Fig. 4) that randomly subdivides UCNPs into physically equivalent heterostructures during training (see ‘Data augmentation for training subdivision invariance’ section in the Methods and Supplementary Information section II).

Nanoparticle optimization

The hetero-GNN is fully differentiable and takes the features that define the UCNPs heterostructure (layer radii and dopant concentrations) as explicit inputs. For optimization, the hetero-GNN acts as a surrogate model for kMC that is not only orders of magnitude faster, but also provides derivatives of a predicted label with respect to structural features (which are inaccessible with kMC), enabling the use of more powerful gradient-based optimizers to identify UCNPs structures that minimize or maximize one or multiple properties.

To explore the utility of our differentiable model for inverse design, we use the hetero-GNN trained on SUNSET-1 with data augmentation to search for Yb/Er/Nd-codoped UCNPs with the highest UV/blue intensities under 800-nm excitation. To facilitate the discovery of complex UCNPs heterostructures, we conduct optimizations far beyond the structural distribution spanned by our training data. While the SUNSET-1 training set contains UCNPs structures with up to four regions (a core and three shells) and with a maximum radius of 11.5 nm, our optimizations explore UCNPs with up to ten regions and a maximum radius of 15 nm. Furthermore, during optimization we relax the constraints on region thicknesses used to generate training data (see ‘Optimization’ section in the Methods for more details).

Given randomly initialized nanoparticle structures, we use a combination of trust region constrained local optimization^{41,42} and basin hopping global optimization to identify UCNPs with maximized UV/blue emission as a function of maximum allowed nanoparticle radius and the number of distinct regions in the heterostructure. The ‘Optimization’ section in the Methods contains additional details of our optimization approach. The maximum intensities identified for UCNPs of different sizes and different numbers of regions (core + shells) are illustrated in Fig. 5a. The optimal structure for each distinct radius +

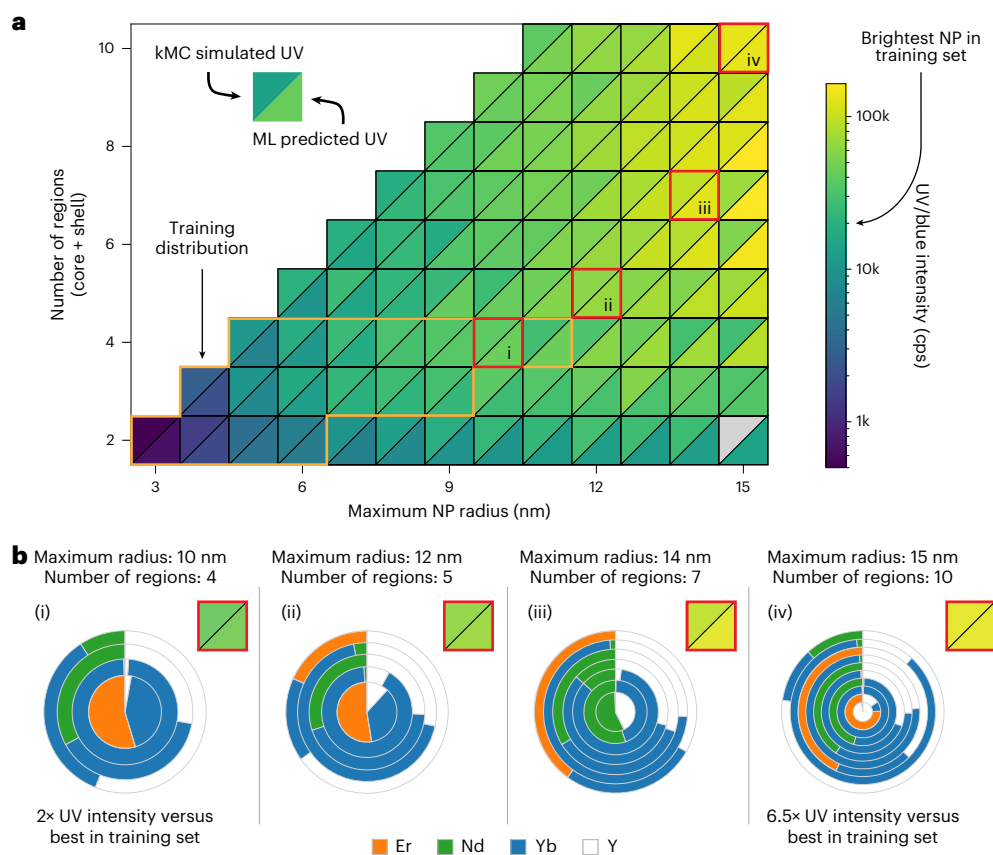


Fig. 5 | Gradient-based optimization of UCNP heterostructure. **a**, The optimization matrix illustrates $I_{\text{vis-UV}}$ intensity (cps) predictions for different particle sizes and numbers of layers. Each square represents the total UV and near-UV intensity obtained from validation (top left triangle) and predicted by the ensemble of models (bottom right triangle). The orange outline encapsulates the training distribution. Red boxes indicate optimized UCNPs that were selected for visualization. See Supplementary Fig. 10 for an alternate color scheme provided

for accessibility. **b**, Two-dimensional (2D) visualizations of four representative optimized UCNPs from the optimization matrix. The brightest structure, UCNP (iv), has a UV intensity 6.5× brighter than the brightest UCNP in the SUNSET-1 training set. Er (orange), Nd (green), Yb (blue) and Y (white) are represented by their respective colors in the 2D visualizations. Note: kMC simulations of the optimized two-layer 15-nm particle were not possible due to the large number of Er^{3+} dopants.

region number was then simulated with kMC, for which larger radius particles often required months-long simulations. The colors of the two triangles in each square indicate the kMC-simulated and ML-predicted intensities for the optimized structure. Remarkably, the model demonstrates accurate predictions for the upconversion luminescence of particles with substantially OOD radii and numbers of regions, including when the UV emissions approach an order of magnitude higher than those in the training set. On the total set of optimized UCNP structures ($N = 80$), the model achieves an MSE of 33.8, which is slightly worse than the SUNSET-1 OOD test set (22.1), but is impressive given the optimized UCNPs are far further OOD in terms of emission intensity and heterostructure complexity.

Optimization results display several trends that are well established in the experimental literature on UCNP heterostructures. Optimized UCNPs generally achieve higher absolute brightness at larger diameters, presumably because they are able to host a greater number of absorbing and emitting dopants⁴³. The optimized heterostructures for several representative sizes (Fig. 5b) show that the domains of these champion UCNPs are in fact heavily doped—or rather, alloyed⁴⁴, with up to 100% lanthanide substitution—to maximize absorption and emission throughput. Rather than spreading dopants homogeneously through UCNPs, the brightest structures partition Er and Nd dopants into separate shells, reflecting the established knowledge that Er and Nd are prone to quench each other via cross-relaxation⁴⁵. Because the energy absorbed by Nd must be transferred to Er for upconversion,

the optimizer produced structures that separate Nd- and Er-rich domains by a thin shell heavily doped with only Yb (ref. 6). Such layers transmit the energy absorbed by Nd dopants to the upconverting Er dopants via rapid energy migration through the Yb sublattice. To maximize Nd–Yb and Er–Yb ET, the Nd- and Er-containing shells are also heavily doped with Yb. Many of the GNN-optimized structures, particularly ones with fewer layers, are reminiscent of the three-layered heterostructure refined by Zhong et al.⁶ and others³⁸. Nd-rich domains are located in outer shells to maximize absorption by a larger number (volume) of the sensitizers. Meanwhile, an Er-rich core is used to promote upconversion by concentrating absorbed energy into a smaller volume and smaller number of Er activators. The fact that gradient-based optimization of the hetero-GNN surrogate model can rapidly learn design rules that have been developed over decades of UCNP research suggests its potential to discover even more complex and functional nanostructures.

In addition to validating established domain knowledge, the extrapolated results from hetero-GNN optimization provide the opportunity to understand the behavior of complex UCNPs with a greater number of shells than can be readily synthesized or simulated. For example, a major unanswered question is the optimal number of layers for a UCNP; for instance, are more layers better? The optimization matrix in Fig. 5a suggests that, for smaller particles (below 10 nm radius), moving beyond two to three layers does not substantially improve the brightness, probably because the shell thicknesses would

be thinner than characteristic ET distances. However, larger UCNP heterostructures ($r > 10$ nm) do see benefit from complex many-shell architectures, with the brightest 14- and 15-nm UCNPs having seven and ten shells, respectively. These many-shelled structures also suggest strategies to enhance upconversion efficiency. The most striking characteristic of the optimized 12-, 14- and 15-nm UCNPs (Fig. 5b(ii–iv)) is their interleaving of multiple layers of Nd- and Er-rich shells. Rather than converging on one large layer of Nd sensitizer, the optimized 12-nm UCNP sandwiches a layer of Nd sensitizer between two layers of Er activator (with the appropriate Yb buffer layers, as in Fig. 5b(ii)), while the brightest 15-nm UCNP exhibits the inverse arrangement (Fig. 5b(iv)). This sandwich shell arrangement allows ET to occur from two sides, maximizing the number of donors or acceptors within a given distance while minimizing concentration quenching in those outer shells. Curiously, the 14- and 15-nm-radius UCNPs also exhibit motifs in which two Nd-rich shells are separated by an intermediate shell of less concentrated Nd. It is unclear what advantage this motif provides. It is possible that UCNPs may be relatively insensitive to variations near their core (where this motif is observed) because the fraction of dopants is relatively small compared with those in outer shells. This argument may also explain the curious dearth of dopants in the core of the 15-nm UCNP. Microscopic analysis of the trajectories and energy flows inside the corresponding kMC simulations³⁰ can be used in the future to elucidate the mechanistic origins of these structural motifs (see Supplementary Fig. 11 for the mechanism extracted from the simulation of the optimized four-layer, 10-nm UCNP heterostructure shown in Fig. 5b(i)). Thus, while hetero-GNN models do not provide direct physical insight, they can accelerate understanding by quickly identifying optimized structures for deeper mechanistic investigation with physical models.

In a particle utilizing all of the photophysical strategies observed in Fig. 5b, we find a $6.5\times$ increase in $J_{\text{vis-UV}}$, as compared with the brightest nanoparticle in the training set. Supplementary Video 1 visualizes a local optimization trajectory (where many local optimizations compose an overall global optimization) for a particle with eight total regions, showing the modification of layer geometry and dopant concentrations as the emission intensity is being maximized. Even within the training feature distribution, the optimization identifies a particle that utilizes these design rules to achieve a $2\times$ increase in emission intensity as compared with the brightest particle in the training set. These results illustrate that, both within the training distribution but especially far OOD, optimization with a differentiable hetero-GNN can rapidly discover structures with properties that exceed historical training data and identify heterostructure design rules.

Ultimately, the optimized structures in Fig. 5 and their intriguing structural design motifs must be validated and investigated more thoroughly through future experiments. The synthesis of the more complex four- to ten-layered structures, while tedious, should be feasible because the layer-by-layer synthesis of multishell UCNP heterostructures is well established^{5,6,46} and can be automated using precise robotic workflows^{5,40}. However, true comparisons of these synthesized structures with model predictions will require future advances in nanoscale characterization (for example, atomic tomography), because imaging tools such as transmission electron microscopy and electron energy loss spectroscopy currently cannot resolve the small lanthanide ion concentrations (as low as 0.1%) inside the thin shells (as small as 1 nm) of the optimized heterostructures, especially when buried within many other shells and projected into a two-dimensional image.

Discussion

To assess the accuracy of our model prediction during optimization, particularly in the far OOD region, we performed explicit kMC simulations. For the largest particles, these simulations were extremely expensive, and we terminated many simulations early (after 20–80% of the requested kMC steps had run) to reduce cost. Overall, the validating

kMC simulations took $>120,000$ central processing unit (CPU)-hours on AMD EPYC 7763 and Intel Xeon Gold 6330 CPUs, and individual simulations could take dozens of weeks. All optimizations using our trained hetero-GNN took approximately 2,000 graphics processing unit (GPU)-hours on NVIDIA A100 GPUs. Because optimal particles often emerged early during the optimization process, this GPU-hour figure could probably be reduced by improving our optimization procedure. Coupled with the fact that the kMC simulations could not be directly used for gradient-based optimization, as they are not inherently differentiable, this indicates the massive acceleration in nanomaterial design that can be achieved using DL.

The models that we developed here were trained to predict a single property, namely emission intensity. An alternative direction would be to directly learn the UCNP photodynamics, training models to predict the state of a UCNP system at a particular point in time. This approach, in line with recent, related work in the area of neural network solutions to differential equations^{47,48}, could enable more facile and automated mechanistic reasoning and is a worthy direction for future study.

While we have here focused on UCNPs, we believe that the heterogeneous graph representation that we have described and implemented could be suitable to predict heterostructure-dependent properties in other multilayered nanomaterials. Possible applications include engineering the nanophotonic properties of plasmonic^{7,10} and dielectric nanoparticles⁸, the catalytic properties of polyelemental heterostructures⁴⁹, the optoelectronic properties of complex semiconductor nanoparticle heterostructures¹¹, the layer-by-layer assembly of nanoparticles for drug delivery⁵⁰, multilayered magnetic nanospheres¹² and multilayer graphene sheets for diverse energy and mechanical applications⁵¹. Our framework could be applied to such materials by encoding the composition of the layers, dimensional information and physical interactions as features in the nodes of heterogeneous graphs. A hetero-GNN could be trained to output desired properties such as scattering spectra. The application-specific physics of these systems can be customized simply by substituting the physics-informed features in the interaction nodes, highlighting a major advantage of our heterogeneous graph approach. We also note that, given a DL model that can predict multiple properties controlled by heterostructure, our approach could enable structural optimization while maximizing or minimizing multiple properties simultaneously. Overall, our approach has the potential to considerably improve the rate at which we discover functional nanomaterials and could provide inspiration for applications of DL to underexplored areas of chemistry and nanoscience.

Methods

Simulation

We use the high-performance C++ kMC implementation in the RNMC software package³¹ to simulate the optical response of lanthanide-doped nanoparticles. Input generation is handled by NanoParticleTools. The effects of parasitic surface ligands are incorporated by including dopant species that are acceptors, which effectively act as energy sinks.

Our kMC method uses the Gillespie rejection-free kMC approach⁵² without additional optimization for efficiency. As such, we observe a large fraction of computation consumed by rapid fluctuations between configurations connected by resonant ET between nearby ions or by thermal equilibration between nearby energy levels⁵³. Simulation efficiency could be increased through established solutions to this ‘low barrier problem’ of temporally stiff kMC simulations, including by addressing nearby energy levels and ions as a single entity⁵³ or by reducing rate constants for the most rapid fluctuations⁵⁴. Such methods must be adapted and validated for the complex interactions of UCNP photophysics. Beyond kMC, there are alternative methods to simulate the photophysics of upconversion, including solving systems of differential rate equations (DREs) across spatially discretized UCNP

heterostructures⁵⁵. While a comparison between these methods is outside the scope of our current work, they suggest that it may be possible to study UCNP as we have here without relying on kMC. Although the complexity of ET networks in complex UCNP heterostructures negates many of the advantages of DREs, recent developments using neural network solutions to differential equations could accelerate such DRE-based approaches.

Workflow and Infrastructure

To scale UCNP simulations from single, hand-tailored calculations to high-throughput simulations on the scale of tens of thousands, we built an automated workflow to orchestrate the simulations. The workflow is implemented in Python, utilizing JobFlow and FireWorks^{56,57}.

To begin the process, we randomly generate a large set of nanoparticle structures using code from NanoParticleTools. A ‘firework’ (calculation unit) is created for each candidate nanoparticle structure and inserted into a MongoDB collection (Supplementary Fig. 1). For maximum throughput, we execute simulations with multiple tiers of concurrency, including across various supercomputers, multiple jobs per system and multiple workers per job. Each worker runs simulations for a single nanoparticle heterostructure, corresponding to four individual simulations with different seeds for the random number generator in event sampling.

When a resource is provisioned and a worker is activated, it queries a job specification from the MongoDB database. If the simulation is new, the input files are generated, and the directory and machine information are recorded in the database. If the job is a continuation of an existing run, the worker navigates to the previously logged directory. We note that resumed simulations are allocated only to workers on the same machine as the initial run, as supercomputers often do not have direct access to each other’s file systems. Simulations are executed until completion or until the resource allocation is revoked. SLURM (Simple Linux Utility for Resource Management) jobs are configured to send a signal to the worker 2 min before termination, allowing the worker to save the current state and checkpoint the job. Once a simulation is completed, the trajectory is analyzed, and a summary of the results is written to the database. The worker then repeats the cycle of querying a job definition, and so on. This automated checkpointing functionality allowed us to use low-priority preemptible queues for the majority of SUNSET data generation.

SUNSET dataset

We present the dataset titled SUNSET. SUNSET consists of five sub-datasets that encompass a range of nanoparticles. Evolution of the ET network via kMC is carried out to 10 ms.

Within each subdataset, we provide a training (and validation) split, an ID test split and an OOD test split. In SUNSET-2, SUNSET-3 and SUNSET-4, the OOD test split also contains the particles with the top 5% and bottom 5% of core sizes. In the case of SUNSET-2 and SUNSET-3, where many of the structures are obtained from Bayesian optimization⁹, we also partition the highest-emitting particles into the OOD test splits. We set these particles aside to ensure the model performance, and therefore the optimization is not biased by these preoptimized configurations. SUNSET-1 comprises a wider parameter space and is sampled in an unbiased manner; thus, we do not follow the same splitting scheme. Instead, we use the four-shell nanoparticles as the OOD, rather than splitting based on emissions.

- SUNSET-1 targets the Er–Nd–Yb system, commonly used to achieve upconversion with 800-nm light^{6,45}. It is composed of multilayered UCNP with a core and up to four shells, illustrated in Fig. 2b. Each nanoparticle consists of a doped core with variable radius r_{core} ranging from 1 to 4 nm and doped shells with variable thickness r_{shell} ranging between 1 nm and 2.5 nm. The total UCNP size ranges between 1 nm to 13.6 nm.

- SUNSET-2 targets the Yb–Er system, a pair that is known to absorb at 980 nm and emit UV light. It is composed of core–shell UCNP characterized by a total radius of 4 nm. Each UCNP consists of a doped core with variable radii, r_{core} , ranging from 0 to 3.4 nm, enclosed within a fixed outer shell radius of 3.4 nm. In addition, the UCNP features a cap shell containing surface species, which mimic the parasitic nature of surface ligands in real UCNP systems. The total UCNP size remains constant at 4 nm. Notably, certain particles lack a shell when $r_{\text{core}} = 0$ or 3.4 nm.
- SUNSET-3 also focuses on the Yb–Er system but excludes the presence of surface species. Removing parasitic surface ligands mirrors the effect of growing an inert external shell of undoped NaYF₄ over the synthesized UCNP. With the external shell excluded, the UCNP size is limited to 3.4 nm. Parameters, including the core and shell sizes, dopants and incident wavelength, remain identical to SUNSET-2.
- SUNSET-4 expands on SUNSET-2 by introducing Tm as a possible dopant. Like Er, Tm can also facilitate UV emission. All other parameters in SUNSET-4 remain identical to SUNSET-2. This dataset is intended to be used in conjunction with SUNSET-2.

Averaging kMC trajectories

Due to the stochastic nature of kMC simulations, where slight differences in sampling lead to drastically different trajectories, even with the same starting conditions, we average across multiple replicate trajectories when reporting photophysical properties. Specifically, the simulations are influenced by the dopant placement on the NaYF₄ host and the pseudorandom number generator used for event sampling. In the reported datasets, we generate four uniquely doped nanoparticles and simulate each with four different sampling seeds, for a total of 16 replicates. In some cases where simulations are very costly and run for weeks, such as in large or heavily doped nanoparticles, we use fewer dopant configurations resulting in averages of 4, 8 or 12 replicates. This is the case for ~6% of the reported data.

To examine the effect of the simulation averaging, we examine the mean statistics using Delete- d jackknife⁵⁸, treating the average of 16 replicates as the true value. The error introduced by averaging 4–12 trajectories is between 60 cps and 140 cps (Supplementary Fig. 2). While this may be orders of magnitude for low-emitting particles, for the target high-emitting particles, this is an error of <1%. We find training with more data, by inclusion of the lower fidelity points, outweighs the negatives of less accurate data. The lower fidelity points probably include a higher proportion of more heavily doped particles, which are more likely to run into limitations with runtime, and thus their inclusion in training is important for dataset diversity and to avoid an unintentional biasing against heavily doped particles.

Aggregation and data transformation

The output and model target, $I_{\text{vis-UV}}$, aggregated across the wavelength range of 300–450 nm (equation (1)), presents two major challenges. First, the summed intensities exhibit a broad range spanning multiple orders of magnitude. Regression using metrics such as mean absolute error or MSE may emphasize fitting higher-magnitude values, leading to neural network parameters becoming biased to more accurately predict these large values. Second, the intensities follow an exponential distribution, posing challenges for ML models that might prioritize learning patterns in the majority class (lower emissions) and struggle to generalize to the minority class (higher emissions).

$$I_{\text{vis-UV}} = \sum_{\lambda=300}^{450} I(\lambda). \quad (1)$$

To address these challenges, a logarithmic transformation is applied to compress the labels into a more manageable range (Equation

(2)). To prevent undefined values, a constant is added before the log transform for nanoparticles with no emissions. In addition to avoiding undefined values, this constant impacts the spacing of low-valued data. To ensure no gap in the label distribution, this constant is set to 100, considering that the lowest nonzero value in the dataset is 25 cps based on an average of a minimum of four 10-ms simulations.

$$\text{target} = \log_{10}(I_{\text{vis-UV}} + 100). \quad (2)$$

Integrated interaction

The integrated interaction is derived by integrating a Gaussian function, denoted as $\mathcal{N}(s; 0, \sigma)$, over all pairwise distances, s , between interacting regions, V_i and V_j , as illustrated in Fig. 3b:

$$I = x_i \times x_j \int \int \mathcal{N}(s; 0, \sigma) dV_i dV_j, \quad (3)$$

$$s(r_1, \theta_1, \phi_1, r_2, \theta_2, \phi_2) = \sqrt{r_1^2 + r_2^2 - 2r_1r_2(\sin(\theta_1)\sin(\theta_2)\cos(\phi_1 - \phi_2) + \cos(\theta_1)\cos(\theta_2))}, \quad (4)$$

where x_i and x_j are the doping concentrations in regions i and j , respectively, σ are the neural network parameters, and s is written in spherical coordinates. While the probability of ET between two dopants is actually proportional to distance⁻⁶, we chose to represent this probability as a sum of Gaussian functions because they are continuous at $x = 0$ and can be integrated multiple times while still capturing the decaying nature of ET with increasing distance between two ions. Adjusting the σ parameters within the Gaussians enables the modulation of the effective interaction distance of dopants. When used in the ML model, the integrated interaction module is parameterized by n learnable weights (here $n = 5$), each corresponding to a σ value of one of the Gaussians in the sum.

Data augmentation for learning subdivision invariance

While the hetero-GNN exhibits superior performance to the CNN, the image (CNN) representation has the physically intuitive property that arbitrarily subdividing a given nanoparticle region (for example, dividing a shell into two smaller shells, where both have the same dopant concentrations as the originally undivided region) has no impact on the model's structural representation or subsequent label prediction. This property, which we call subdivision invariance, is physically motivated by the fact that region subdivision is arbitrary and leaves the nanoparticle being described completely unchanged. However, neither our heterogeneous graph nor any of the other nonimage representations are inherently subdivision invariant. For example, as shown in Fig. 4a, subdividing an originally core-only particle into a core and a shell dramatically changes the heterogeneous graph, and thus our hetero-GNN model may predict very different latent representations for physically identical nanoparticles. This is clearly undesirable and may be detrimental to both the learning process and subsequent structural optimization. Meanwhile, the voxelization that makes the image representation subdivision invariant simultaneously causes layer dimensions to only be present in the model implicitly, preventing UCNP emission from being differentiated with respect to layer thicknesses and precluding gradient-based optimization of UCNP heterostructure. Thus, any DL model that aims to enable inverse design of nanomaterial heterostructure via gradient-based optimization will need to reckon with the problem of subdivision invariance.

Even when an input representation is not inherently subdivision invariant, it is possible to design the DL model built atop the representation to explicitly enforce subdivision invariance such that physically identical structures yield identical latent representations. However, in the context of a graph representation, such explicit enforcement is

only possible by avoiding the use of any nonlinear operations, which dramatically limits model expressivity and performance.

An alternate strategy is to train models to approximate subdivision invariance via data augmentation. Using data augmentation to train approximate invariances in different DL contexts (for example, image rotation, reflection in CNNs⁵⁹, molecular rotation and translation in interatomic potentials⁶⁰) is well established and can enhance model prediction accuracy and robustness. We apply this augmentation strategy to our hetero-GNN model by artificially subdividing the UCNP input with the same labels (emission intensities) but with structural representations modified with random subdivisions. More specifically, for each UCNP input, we randomly subdivide each parent layer into up to three child layers in the augmented UCNP. The subdivision is inserted between 5% and 90% of the parent layer radii. This data augmentation is meant to guide the learned latent representation to exhibit approximate subdivision invariance, which should improve model performance. This strategy is implemented on-the-fly (Fig. 4b), so that data in each batch are augmented as they appear during training, rather than augmenting the entire dataset before training. Thus, this does not explicitly increase the size of the training dataset, but implicitly multiplies the number of unique heterograph representations of UCNP seen by the model during training by the number of training epochs. Random subdivisions result in a node with $r_{\text{inner}}, r_{\text{outer}}$ being split into two nodes with $r_{\text{inner}}, r_{\text{subdivision}}$ and $r_{\text{subdivision}}, r_{\text{outer}}$, where $r_{\text{inner}} < r_{\text{subdivision}} < r_{\text{outer}}$.

When the hetero-GNN is trained using on-the-fly augmentation, its performance improves by 23% on the ID test set (with error falling from 13.8 to 10.6) and by 25% on the OOD test set (with error falling from 22.1 (2.3%) to 16.5 (1.7%)). In addition, we validated that this augmentation scheme actually trains the model to learn subdivision invariance by evaluating the vector distance between the representation of nanoparticles and their subdivided analogs, observing that the augmented hetero-GNN model more closely represented the subdivided UCNP in the ID and OOD test sets in embedding space than the nonaugmented model across a range of subdivisions (Supplementary Fig. 6). These results underscore the importance of considering subdivision invariance in model training.

Model training

A learning rate of 1×10^{-3} is used, with a warm-up period of ten epochs, where the learning rate linearly increases from 1×10^{-4} to 1×10^{-3} . During training, the MSE of the validation set is monitored and the learning rate is reduced on plateau, with a patience of 50 epochs. Early stopping is triggered when the validation MSE has not decreased for 200 epochs. Model performance is reported using the MSE and normalized MSE (NMSE), as shown in equations (5) and (6) below:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (\log_{10}(I_{\text{vis-UV}}^{\hat{}} + 100) - \log_{10}(I_{\text{vis-UV}} + 100))^2 \quad (5)$$

$$\text{NMSE} = \frac{\sum_{i=1}^N (\log_{10}(I_{\text{vis-UV}}^{\hat{}} + 100) - \log_{10}(I_{\text{vis-UV}} + 100))^2}{\sum_{i=1}^N \log_{10}(I_{\text{vis-UV}} + 100)^2}, \quad (6)$$

where N is the number of UCNP, $I_{\text{vis-UV}}^{\hat{}}$ is the predicted UV emission intensity and $I_{\text{vis-UV}}$ is the actual emission intensity.

Optimization

For a nanoparticle with l control volumes and z dopant elements, we can define the following bounds.

$$0 \leq x_n^i \leq 1, \quad \text{for } i \in [0..l] \text{ and } n \in [0..z] \quad (7)$$

$$0 < r_i^{\text{frac}} \leq 1, \quad \text{for } i \in [0..l], \quad (8)$$

where the outer radius of one control volume is explicitly the inner radius of the subsequent control volume. All dopant concentrations are within the closed interval $[0, 1]$. The fractional radius of region i , r_i^{frac} , is defined on an interval of $[0, 1]$ as a fraction of the maximum nanoparticle size, r_{max} (which we identify a priori). This is necessary to keep on the same interval as the concentration, because the trust region optimizer defines the same trust region for all independent variables $r_i^{\text{frac}} = r_i/r_{\text{max}}$.

In addition, we define linear constraints that bound the total concentration within each region $[0, 1]$ and restrict the thickness of each region.

$$0 \leq \sum_{j=z \times n+1}^{z \times n+z} x_j \leq 1, \quad \text{for } n \in [0, l-1] \quad (9)$$

$$c_{\text{min}} \leq r_0 \leq c_{\text{max}} \quad (10)$$

$$t_{\text{min}} \leq r_{n+1} - r_n \leq t_{\text{max}}, \quad \text{for } n \in [0, l-1], \quad (11)$$

where j is a dopant-region index, the minimum core radius $c_{\text{min}} = 1$ nm, the maximum core radius $c_{\text{max}} = 5$ nm, the minimum layer thickness $t_{\text{min}} = 0.5$ nm and the maximum layer thickness $t_{\text{max}} = 5$ nm. We initialize random starting configuration within the distributions outlined for each dataset as a starting point for optimization. We perform local optimization using the trust region constrained optimization as implemented in SciPy⁴¹ with an initial trust radius of 1.0. An initial constraint penalty of 1×10^3 was applied to strongly penalize constraint violation, ensuring that concentrations stayed within the range of $[0, 1.0]$ and total radius within specification. The criterion used for termination of local optimization is when the trust radius is less than 1×10^{-8} . To search for a globally optimal particle, we repeatedly perturb the local minima and reoptimize the nanoparticle heterostructure using the basin-hopping functionality of SciPy with up to 500 steps, a step size of 0.15 and a temperature value of 0.25. Following global optimization, the best identified candidate structures are validated via kMC simulations. For the larger particles in the optimization matrix, these kMC simulations often took months to run.

Data availability

The SUNSET dataset is freely available via Figshare at <https://doi.org/10.6084/m9.figshare.25130921> (ref. 61). Each subset (for instance, SUNSET-1) is presented in Javascript Object Notation (JSON) format; separate JSON files are also provided for ID and OOD collections for each subset. The hetero-GNN models used for UCNP optimization are freely available via Figshare at <https://doi.org/10.6084/m9.figshare.27941694.v1> (ref. 62). The nanoparticle structures discovered by optimizing the hetero-GNN model are available via Figshare at <https://doi.org/10.6084/m9.figshare.27973206> (ref. 63). Data for Figs. 2 and 5 are available via Figshare at <https://doi.org/10.6084/m9.figshare.29916992> (ref. 64).

Code availability

The RNMC program, which contains the NPMC kMC tool, is available via GitHub at <https://github.com/BlauGroup/RNMC> and via Zenodo at <https://doi.org/10.5281/zenodo.14360064> (ref. 65). Code defining the ML representations, data featurization and model training is available via GitHub at <https://github.com/BlauGroup/NanoParticleTools> and via Zenodo at <https://doi.org/10.5281/zenodo.16878169> (ref. 66).

References

- Voznyy, O., Sutherland, B. R., Ip, A. H., Zhitomirsky, D. & Sargent, E. H. Engineering charge transport by heterostructuring solution-processed semiconductors. *Nat. Rev. Mater.* **2**, 17026 (2017).
- Sun, W. et al. Machine learning–assisted molecular design and efficiency prediction for high-performance organic photovoltaic materials. *Sci. Adv.* **5**, 4275 (2019).
- Sanchez-Lengeling, B. & Aspuru-Guzik, A. Inverse molecular design using machine learning: generative models for matter engineering. *Science* **361**, 360–365 (2018).
- Wu, S. et al. Machine-learning-assisted discovery of polymers with high thermal conductivity using a molecular design algorithm. *npj Comput. Mater.* **5**, 66 (2019).
- Skripka, A. et al. A generalized approach to photon avalanche upconversion in luminescent nanocrystals. *Nano Lett.* **23**, 7100–7106 (2023).
- Zhong, Y. et al. Elimination of photon quenching by a transition layer to fabricate a quenching-shield sandwich structure for 800 nm excited upconversion luminescence of Nd³⁺-sensitized nanoparticles. *Adv. Mater.* **26**, 2831–2837 (2014).
- So, S., Mun, J. & Rho, J. Simultaneous inverse design of materials and structures via deep learning: demonstration of dipole resonance engineering using core-shell nanoparticles. *ACS Appl. Mater. Interfaces* **11**, 24264–24268 (2019).
- Peurifoy, J. et al. Nanophotonic particle simulation and inverse design using artificial neural networks. *Sci. Adv.* **4**, 4206 (2018).
- Xia, X., Sivonxay, E., Helms, B. A., Blau, S. M. & Chan, E. M. Accelerating the design of multishell upconverting nanoparticles through Bayesian optimization. *Nano Lett.* **23**, 11129–11136 (2023).
- Liu, G.-X. et al. Inverse design in quantum nanophotonics: combining local-density-of-states and deep learning. *Nanophotonics* **12**, 1943–1955 (2023).
- Hamachi, L. S. et al. Precursor reaction kinetics control compositional grading and size of cdse1-xsx nanocrystal heterostructures. *Chem. Sci.* **10**, 6539–6552 (2019).
- Albrecht, M. et al. Magnetic multilayers on nanospheres. *Nat. Mater.* **4**, 203–206 (2005).
- Xu, M., Liang, T., Shi, M. & Chen, H. Graphene-like two-dimensional materials. *Chem. Rev.* **113**, 3766–3798 (2013).
- Chen, Y. et al. Two-dimensional metal nanomaterials: synthesis, properties, and applications. *Chem. Rev.* **118**, 6409–6455 (2018).
- Xie, C., Niu, Z., Kim, D., Li, M. & Yang, P. Surface and interface control in nanoparticle catalysis. *Chem. Rev.* **120**, 1184–1249 (2019).
- Haghighatlari, M. et al. Learning to make chemical predictions: the interplay of feature representation, data, and machine learning methods. *Chem* **6**, 1527–1542 (2020).
- Krenn, M., Häse, F., Nigam, A., Friederich, P. & Aspuru-Guzik, A. Self-referencing embedded strings (SELFIES): a 100% robust molecular string representation. *Mach. Learn. Sci. Technol.* **1**, 045024 (2020).
- Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.* **28**, 31–36 (1988).
- O'Brien, M. N., Jones, M. R. & Mirkin, C. A. The nature and implications of uniformity in the hierarchical organization of nanomaterials. *Proc. Natl Acad. Sci. USA* **113**, 11717–11725 (2016).
- Liang, R. et al. Hierarchically engineered nanostructures from compositionally anisotropic molecular building blocks. *Nat. Mater.* **21**, 1434–1440 (2022).
- Brown, K. A., Brittman, S., Maccaferri, N., Jariwala, D. & Celano, U. Machine learning in nanoscience: big data at small scales. *Nano Lett.* **20**, 2–10 (2019).
- Zhou, J. et al. 2Dmatpedia, an open computational database of two-dimensional materials from top-down and bottom-up approaches. *Sci. Data* **6**, 86 (2019).
- Reker, D. et al. Computationally guided high-throughput design of self-assembling drug nanoparticles. *Nat. Nanotechnol.* **16**, 725–733 (2021).

24. McCandler, C. A., Dahl, J. C. & Persson, K. A. Phosphine-stabilized hidden ground states in gold clusters investigated via a $Au_n(PH_3)_m$ database. *ACS Nano* **17**, 1012–1021 (2022).
25. Barnard, A. et al. Nanoinformatics, and the big challenges for the science of small things. *Nanoscale* **11**, 19190–19201 (2019).
26. Hlaváček, A. et al. Bioconjugates of photon-upconversion nanoparticles for cancer biomarker detection and imaging. *Nat. Protoc.* **17**, 1028–1072 (2022).
27. Pan, J.-A., Qi, X. & Chan, E. M. Enhanced upconversion and photoconductive nanocomposites of lanthanide-doped nanoparticles functionalized with low-vibrational-energy inorganic ligands. *Nanoscale Horiz.* **10**, 596–604 (2025).
28. Sanders, S. N. et al. Triplet fusion upconversion nanocapsules for volumetric 3d printing. *Nature* **604**, 474–478 (2022).
29. He, M. et al. Monodisperse dual-functional upconversion nanoparticles enabled near-infrared organolead halide perovskite solar cells. *Angew. Chem. Int. Ed.* **55**, 4280–4284 (2016).
30. Teitelboim, A. et al. Energy transfer networks within upconverting nanoparticles are complex systems with collective, robust, and history-dependent dynamics. *J. Phys. Chem. C* **123**, 2678–2689 (2019).
31. Zichi, L. et al. RNMC: kinetic Monte Carlo implementations for complex reaction networks. *J. Open Source Softw.* **9**, 7244 (2024).
32. Xu, X. et al. Optimising passivation shell thickness of single upconversion nanoparticles using a time-resolved spectrometer. *APL Photonics* **4**, 026104 (2019).
33. Wu, X. et al. Dye-sensitized core/active shell upconversion nanoparticles for optogenetics and bioimaging applications. *ACS Nano* **10**, 1060–1066 (2016).
34. Corso, G., Stark, H., Jegelka, S., Jaakkola, T. & Barzilay, R. Graph neural networks. *Nature Rev. Methods Primers* **4**, 17 (2024).
35. Shi, R. & Mudring, A.-V. Phonon-mediated nonradiative relaxation in In^{3+} -doped luminescent nanocrystals. *ACS Mater. Lett.* **4**, 1882–1903 (2022).
36. Perez, E., Strub, F., De Vries, H., Dumoulin, V. & Courville, A. Film: visual reasoning with a general conditioning layer. *Proc. AAAI Conf. Artif. Intell.* **32**, 39412–3951 (2018).
37. Brody, S., Alon, U. & Yahav, E. How attentive are graph attention networks. In *International Conference on Learning Representations (ICLR, 2022)*.
38. Chan, E. M., Levy, E. S. & Cohen, B. E. Rationally designed energy transfer in upconverting nanoparticles. *Adv. Mater.* **27**, 5753–5761 (2015).
39. Skripka, A. & Chan, E.M. Unraveling the myths and mysteries of photon avalanching nanoparticles. *Mater. Horiz.* <https://doi.org/10.1039/D4MH01798F> (2025).
40. Chan, E. M. Combinatorial approaches for developing upconverting nanomaterials: high-throughput screening, modeling, and applications. *Chem. Soc. Rev.* **44**, 1653–1679 (2015).
41. Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).
42. Conn, A. R., Gould, N. I. & Toint, P. L. *Trust Region Methods* (SIAM, 2000).
43. Gargas, D. J. et al. Engineering bright sub-10-nm upconverting nanocrystals for single-molecule imaging. *Nat. Nanotechnol.* **9**, 300–305 (2014).
44. Tian, B. et al. Low irradiance multiphoton imaging with alloyed lanthanide nanocrystals. *Nat. Commun.* **9**, 3082 (2018).
45. Xie, X. et al. Mechanistic investigation of photon upconversion in Nd^{3+} -sensitized core-shell nanoparticles. *J. Am. Chem. Soc.* **135**, 12608–12611 (2013).
46. Li, X., Wang, R., Zhang, F. & Zhao, D. Engineering homogeneous doping in single nanoparticle to enhance upconversion efficiency. *Nano Lett.* **14**, 3634–3639 (2014).
47. Du, Y., Chalapathi, N. & Krishnapriyan, A. Neural spectral methods: self-supervised learning in the spectral domain. In *Proc. Twelfth International Conference on Learning Representations (2024)*.
48. Chalapathi, N., Du, Y. & Krishnapriyan, A. Scaling physics-informed hard constraints with mixture-of-experts. In *Proc. Twelfth International Conference on Learning Representations (2024)*.
49. Wahl, C. B. et al. Machine learning-accelerated design and synthesis of polyelemental heterostructures. *Sci. Adv.* **7**, 5505 (2021).
50. Yan, Y., Such, G. K., Johnston, A. P., Lomas, H. & Caruso, F. Toward therapeutic delivery with layer-by-layer engineered particles. *ACS Nano* **5**, 4252–4257 (2011).
51. Lee, J.-H., Loya, P. E., Lou, J. & Thomas, E. L. Dynamic mechanical behavior of multilayer graphene via supersonic projectile penetration. *Science* **346**, 1092–1096 (2014).
52. Gillespie, D. T. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.* **22**, 403–434 (1976).
53. Andersen, M., Panosetti, C. & Reuter, K. A practical guide to surface kinetic monte carlo simulations. *Front. Chem.* <https://doi.org/10.3389/fchem.2019.00202> (2019).
54. Chatterjee, A. & Voter, A. F. Accurate acceleration of kinetic monte carlo simulations through the modification of rate constants. *J. Chem. Phys.* **132**, 194101 (2010).
55. Pini, F. et al. Spatial and temporal resolution of luminescence quenching in small upconversion nanocrystals. *ACS Appl. Mater. Interfaces* **14**, 11883–11894 (2022).
56. Jain, A. et al. Fireworks: a dynamic workflow system designed for high-throughput applications. *Concurr. Comput. Pract. Exp.* **27**, 5037–5059 (2015).
57. Rosen, A. S. et al. Jobflow: computational workflows made simple. *J. Open Source Softw.* **9**, 5995 (2024).
58. Wu, C.-F. J. Jackknife, bootstrap and other resampling methods in regression analysis. *Ann. Stat.* **14**, 1261–1295 (1986).
59. Shorten, C. & Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *J. Big Data* **6**, 1–48 (2019).
60. Batzner, S. et al. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **13**, 2453 (2022).
61. Sivonxay, E., Attia, L., Chan, E. M. & Blau, S. M. SUNSET Dataset (For 'Gradient-based optimization of complex nanoparticle heterostructures enabled by deep learning on heterogeneous graphs'). Version v1. *Figshare* <https://doi.org/10.6084/m9.figshare.25130921.v1> (2025).
62. Attia, L. Hetero-GNN checkpoints. Version v1. *Figshare* <https://doi.org/10.6084/m9.figshare.27941694.v1> (2024).
63. Attia, L. Optimized nanoparticle structures. Version v1. *Figshare* <https://doi.org/10.6084/m9.figshare.27973206.v1> (2024).
64. Sivonxay, E., Attia, L., Blau, S. & Chan, E. M. Source data (For 'Gradient-based optimization of complex nanoparticle heterostructures enabled by deep learning on heterogeneous graphs'). Version v1. *Figshare* <https://doi.org/10.6084/m9.figshare.29916992.v1> (2025).
65. Zichi, L. et al. RNMC: kinetic Monte Carlo implementations for complex reaction networks. Version v1.1.1. *Zenodo* <https://doi.org/10.5281/zenodo.14360064> (2024).
66. Sivonxay, E., Attia, L., Barter, D., Xia, X. & Blau, S. M. BlauGroup/NanoParticleTools. Version Publication. *Zenodo* <https://doi.org/10.5281/zenodo.16878169> (2025).

Acknowledgements

This work was primarily funded by the Charter Hill Laboratory Directed Research and Development program of Lawrence Berkeley National Laboratory (LBNL), supported by the Office of Science, Office of Basic Energy Sciences (BES), of the US Department of Energy (DOE)

under contract no. DE-AC02-05CH11231. This DOE-BES contract also supported work at the Molecular Foundry as well as computational resources at the National Energy Research Scientific Computing Center (NERSC, award no. BES-ERCAPO023292) and the Lawrence Livermore computational cluster provided by the LBNL IT Division. We are especially grateful for the NERSC and Lawrence Livermore low-priority queues, without which this work would not have been possible. L.A. was supported by the DOE Computational Science Graduate Fellowship under award no. DE-SC0022158. E.W.C.S.-S. was supported by the Carnegie Bosch Institute Postdoctoral Fellowship. We thank K. Chiao for helpful discussions related to optimizing C++ code.

Author contributions

E.S. and D.B. implemented nanoparticle kMC in C++, and X.X. provided validation against a previous implementation. E.S. and X.X. built simulation analysis capabilities in Python. E.S. implemented the high-throughput workflow infrastructure. E.S. and X.X. ran kMC simulations to construct SUNSET. E.S. implemented RFR, FCNN, CNN and homo-GNN models and trained on SUNSET data. E.S. and L.A. implemented the hetero-GNN model and trained on SUNSET data. L.A. and E.S. implemented the on-the-fly data augmentation scheme. E.S. implemented and carried out global optimization, and then ran kMC on the optimized particles to validate predicted intensities. B.S.-L. provided important input on components of the GNN architecture and global optimization. E.S., L.A., E.W.C.S.-S., E.M.C. and S.M.B. wrote the manuscript. E.M.C. and S.M.B. conceived and supervised the project.

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/s43588-025-00917-3>.

Correspondence and requests for materials should be addressed to Emory M. Chan or Samuel M. Blau.

Peer review information *Nature Computational Science* thanks Kuan Sun and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Primary Handling Editor: Katilin McCardle, in collaboration with the *Nature Computational Science* team. Peer reviewer reports are available.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2025