

Enhancing Dimensionality Prediction in Hybrid Metal Halides via Feature Engineering and Class-Imbalance Mitigation

Mariia Karabin¹, Isaac Armstrong², Leo Beck², Paulina Apanel², Markus Eisenbach⁴, David B. Mitzi³, Hanna Terletska¹, Hendrik Heinz²

¹ Department of Physics and Astronomy, Middle Tennessee State University, Murfreesboro, TN, 37132, USA

² Department of Chemical and Biological Engineering, University of Colorado, Boulder, CO, 80309, USA

³ Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC, 27708, USA

⁴ National Center for Computational Sciences Division, Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA

Abstract. We present a machine learning framework for predicting the structural dimensionality of hybrid metal halides (HMHs), including organic-inorganic perovskites, using a combination of chemically-informed feature engineering and advanced class-imbalance handling techniques. The dataset, consisting of 494 HMH structures, is highly imbalanced across dimensionality classes (0D, 1D, 2D, 3D), posing significant challenges to predictive modeling. This dataset was later augmented to 1336 via the Synthetic Minority Oversampling Technique (SMOTE) to mitigate the effects of the class imbalance. We developed interaction-based descriptors and integrated them into a multi-stage workflow that combines feature selection, model stacking, and performance optimization to improve dimensionality prediction accuracy. Our approach significantly improves F1-scores for underrepresented classes, achieving robust cross-validation performance across all dimensionalities.

Keywords: Hybrid metal halides, hybrid organic-inorganic perovskites, data imbalance, feature engineering, predictive modeling, interaction descriptors, data augmentation

. Introduction

Hybrid metal halides (HMHs) are a promising class of materials for optoelectronic and energy technologies [1–3]. Their unique crystal structures, tunable band gaps, and defect tolerance have contributed to the progress in photovoltaics, light-emitting diodes, and photodetectors [4–8]. From a computational standpoint, HMHs are particularly challenging because their hybrid organic-inorganic character combines organic molecules with inorganic frameworks. This computational challenge sometimes requires advanced simulation tools such as the INTERFACE force field (IFF) [9] and motivates data-driven approaches [10] that can complement first-principles calculations. While 3D perovskites are a well-known subset of HMHs due to their high efficiencies in photovoltaic devices, lower-dimensional structures (0D, 1D, and 2D) may offer improved stability and distinct optoelectronic properties [4, 11–13]. Identifying structural dimensionality within the HMH family thus represents a highly important step toward rational materials discovery and design. However, experimental synthesis and characterization of HMHs are resource-intensive, thus motivating the development of computational tools that can accelerate screening and classification.

Machine learning (ML) shows great potential for HMH dimensionality prediction, though faces significant challenges [14]. First, available HMH datasets are relatively small and far below the volumes

typically required for robust ML training [15]. This deficiency increases the risk of overfitting and requires models that can generalize from scarce information. Second, datasets are highly imbalanced, with 2D structures dominating and 0D and 1D being highly underrepresented [16]. This imbalance can bias learning algorithms towards majority classes, while maintaining low per-class accuracy for underrepresented classes. Third, the complex interplay of steric, electronic, and structural factors in HMHs requires careful feature design to capture relevant relationships [17, 18].

To address these challenges, we propose an ML framework combining interaction-based feature engineering, class imbalance mitigation, and advanced ensemble learning [19–21]. First, we design chemically grounded descriptors that include elemental properties, structural characteristics, and interactions relevant to dimensionality. These descriptors are then used to design interaction-based features to capture more subtle, non-linear relationships. We implement the Synthetic Minority Oversampling Technique (SMOTE) [22, 23] to augment minority classes and reduce representational bias. Lastly, we evaluate the predictive accuracy and robustness of ML models, which also include stacking ensembles. This cohesive pipeline enhances predictive per-class accuracy while minimizing potential overfitting.

Our study contributes to the broader discussion of small data machine learning strategies in materials science [14]. More specifically, how to extract reliable insights from small and imbalanced datasets. Many classes of functional materials suffer from similar data scarcity. This limitation motivates the development of ML strategies that involve data-level, algorithm-level, and strategy-level techniques [24, 25]. Approaches that succeed in these settings are therefore highly transferable and offer generalizable frameworks for tackling diverse small-data challenges in computational materials research. Similar to recent efforts in algorithmic scalability and transfer learning [26], our approach demonstrates that thoughtful feature design and data augmentation can provide performance gains comparable to larger datasets.

This work is organized as follows: Section 2 provides a review of the relevant work on machine learning for hybrid metal halide property prediction and small-data strategies. Section 3 describes our methodology, including dataset preparation, model development, feature engineering, and class imbalance mitigation. Section 4 presents the results of our framework, the effects of feature engineering, the impact of SMOTE-based augmentation, and model evaluation. Section 5 discusses the broader implications of our work for HMH discovery and small-data ML strategies, while section 6 concludes the paper and outlines future research directions.

2. Related Work

Machine learning for perovskites has emerged as a powerful tool to complement first-principles methods and enable rapid screening of chemical space and property prediction [27–30]. Early studies often focused on band gap and stability prediction using composition-based features, geometric or density functional theory (DFT)-derived descriptors [31–34]. Other efforts have applied ML to accelerate the discovery of Pb-free perovskite candidates with favorable optoelectronic properties [2, 35]. These works highlight the ability of ML to handle complex, nonlinear relationships between composition, structure, and functional properties.

Despite these advances, dimensionality prediction in hybrid metal halides (including perovskites) remains relatively unexplored. While several studies have demonstrated the relationship between dimensionality and optoelectronic performance, few ML frameworks explicitly target dimensionality prediction [36, 37]. Yuan et al. [36] trained ML models on 195 amines to classify low-dimensional organic-inorganic halide perovskites. Lyu et al. [37] developed ML model to identify 2D perovskites and report steric effect index, eccentricity, largest ring size, and hydrogen-bond donor as key controlling

features. One challenge with perovskite dimensionality prediction is that available HMH datasets are typically small and imbalanced. This imbalance can bias model training and limit predictive reliability. Addressing these challenges requires data-level, algorithm-level, and strategy-level techniques tailored for small datasets [14, 38]. Recent progress within small-data strategies also explores the importance of domain-informed descriptors and balanced sampling [39]. Techniques such as shotgun transfer learning [38] and active learning have been demonstrated in broader materials contexts, showing that careful use of limited data can rival results from larger datasets.

Ensemble-based models [20, 21, 40–43] have been shown to be effective in capturing nonlinear feature interactions in small, noisy datasets while maintaining interpretability. Feature importance analysis in these models further provides chemical insight, linking descriptors to structural outcomes [17, 44–46]. These methods complement deep learning approaches [47] such as graph neural networks, but with the added advantage of interpretability. Such interpretability is important for advancing HMH dimensionality prediction, ensuring that ML not only improves classification accuracy but also contributes to understanding the underlying chemistry [18]. Ensemble models such as random forests and gradient boosted trees remain workhorses of small materials datasets, frequently outperforming deep learning models in limited data regimes [48, 49]. The Synthetic Minority Oversampling Technique (SMOTE) has proven to be effective in balancing materials datasets [50, 51]. Prior studies show that imbalance skews recall, motivating data-level augmentation approaches like those adopted here. The model tends to overfit the dominant class in the dataset, leading to disproportionately low recall for underrepresented classes and a bias in overall predictive performance.

Our study addresses the underexplored dimensionality prediction in HMHs by designing interaction-based features that encode chemical and structural relationships, mitigating dataset imbalance with SMOTE, and using ensemble models to ensure robustness and interpretability. This bridges the perovskite ML research with the broader small-data ML frameworks, positioning HMHs dimensionality prediction as both a materials challenge and a methodological testbed.

3. Methods

In this section, we outline the framework, beginning with dataset construction and preprocessing (section 3.1), followed by model development (section 3.2) and feature engineering (section 3.3), and conclude with techniques for mitigating class imbalance (section 3.4). Together these steps define a coherent pipeline to enable robust ML predictions on small, imbalanced datasets of hybrid metal halides.

3.1. Dataset and feature description

The dataset used in this study describes 494 hybrid metal halide (HMH) structures, obtained from the HybriD³ Perovskite Database produced by Qin et al. [16]. Each dataset entry is correspondingly labeled according to its structural dimensionality. HMHs are a large materials family composed of 'A' organic cations that stabilize 'BX' metal-halide polyhedra which consist of group 14 metals (B) and halide anions (X) [52]. Conventional ABX₃ hybrid perovskites, of recent solar research interest, are a subset of HMHs (Figure 1a). Dependent on the underlying chemistry, these metal-halide octahedra may share corners, edges, or faces, and generally arrange into one or more different dimensionalities: 0D - lacking an extended inorganic framework, 1D - chains, 2D - sheets, and 3D - bulk (Figure 1b). Structural dimensionality of HMHs has to-date been difficult to predict a priori, with similar compounds exhibiting different dimensionalities with only minor changes in composition or synthesis environment [3, 37, 53].




Figure 1: (a) ABX_3 Classic perovskite crystal structure illustrating general arrangement; (b) Examples of structural differences between hybrid metal halide dimensionality classes.




Figure 2: Data distribution of structural dimensionality of hybrid metal halides (HMHs)

Figure 2 reflects high data imbalance and prevalence of layered 2D HMHs, which comprise 67% of the dataset. In contrast, 0D and 1D structures correspond to only 5% and 10%, respectively. In addition to the assigned dimensionality, each entry in the HybriD³ dataset contains the complete chemical composition of the HMH. From these compositions, we generated atomic-level, molecular-level, as well as geometric features. These features, together with the interaction-based features described in section 3.3, serve as the input for the machine learning model described in section 3.2.

To start, we derived our 11 original features (Table 1) from our dataset in an effort to incorporate chemical effects that could influence dimensionality. This effort was guided by an analysis of our dataset that identified 17 instances of two or more HMHs that changed dimensionality upon inclusion of a different

organic cation for the same inorganic lattice composition (i.e. $B = \text{Pb}$, $X = \text{I}$). Due to the higher prevalence of examples of dimensionality change upon replacement of cations, as opposed to swapping out inorganic compositions, we focused our features on traits of organic cations. Geometric dimensions of the cation were an immediate starting point, leading us to define the presence of alkyl chains, aromatic rings, and their relative sizes, generating 6 descriptors. A vast majority of our listed cations consisted of C, H, and N atoms, with N being the most electronegative. As such, we described relative N placements within the cations by simply counting how many 'terminal' N atoms were present (N atoms bonded to only a single non-hydrogen atom). Some entries contained multiple 'A' cations per unit cell of the HMH, the number of which we included. Additionally, some of the HybriD³ data had been reported with water present, potentially displacing cations, so the presence of water was also tracked. These features, coupled with the chemical formulas for the organic and inorganic components, gave 11 unique descriptors from which to initialize ML training. Further details on these descriptors, and why they were chosen, can be found in the SI.

3.2. ML model development and validation

To predict the dimensionality of hybrid metal halides, we employed both single model learners and an ensemble strategy to leverage complementary strengths. The implemented models were selected for their robustness in handling heterogeneous features, non-linear relationships, and imbalanced datasets [21, 54].

Gradient Boosted Trees (GBT) were used as a core predictive model due to their ability to capture complex non-linear decision boundaries and interactions between features [55]. GBT iteratively fits weak learners (decision trees) to the residual errors of previous trees to reduce bias and variance. This iterative boosting framework is particularly well-suited for tabular scientific datasets, where descriptors exhibit varying degrees of freedom. Hyperparameters such as learning rate, maximum tree depth, and number of estimators were tuned to balance predictive accuracy with generalization.

In addition to GBT, we have also used several classifiers as base learners in the ensemble. Logistic regression is a linear model that provides interpretable baseline decision boundaries, while Random Forest is a bagging-based ensemble of decision trees that mitigates overfitting. Support Vector Machine (SVM) is effective in high-dimensional feature spaces, particularly for separating overlapping classes. These models capture distinct structural patterns and provide complementary perspectives on the dataset [56].

We implemented a stacking ensemble to integrate these predictive strengths. The base learners generate first-level predictions, which are then used as input features for the meta-learner, GBT. This two-layer architecture makes use of diverse error structures and learns optimized decision rules that outperform an individual model. Stacking has been shown to improve robustness against overfitting and provide better generalization on small, imbalanced datasets [14, 24]. Our model was trained and validated using stratified k-fold cross-validation to preserve class balance across folds. Performance was evaluated using precision, recall, F1-score, and receiver operating characteristic (ROC) curves and corresponding area under the curve (AUC) scores (section 4). The ROC curve shows the model's ability to distinguish between different classes by plotting the true positive rate against the false positive rate. A higher AUC value indicates a stronger capability of the classifier to distinguish between different structural dimensionality classes, with an AUC of 1.0 representing perfect separation.

3.3. Feature engineering

Our workflow involves interaction-based feature engineering [19, 57], as the original descriptors alone were insufficient to fully capture the structure-dimensionality relationships of HMH. Given the relatively

small size of our dataset (494 compounds), careful feature design was required to maximize the prediction accuracy while minimizing the risk of overfitting. This approach combines chemically motivated descriptors with interaction-based features that encode more subtle relationships, and therefore improve separability among minority classes.

The original 11 descriptors were derived from our dataset (section 3.1) and reflect steric hindrance, hydrophilicity, and molecular complexity, all of which highly affect HMH dimensionality. To capture nonlinear effects and coupled contributions, we introduced 11 interaction-based features (Table 1). Unlike statistical feature construction, this approach reduces the risk of spurious correlations, which is particularly important in small datasets where ML models are more prone to overfitting [14]. Moreover, interpretable features allow ML predictions to be validated against existing chemical intuition.

The pairwise interaction features were generated using pairwise combinations of the original features (e.g., differences, products, and ratios). This encodes competitive or synergistic effects between features that influence the dimensionality of HMHs. Other types of interaction features were generated using domain-informed combinations and non-linear transformations (Table 1).

Feature Type	Feature	Description
Original features	Organic and inorganic formula units, number of cation rings, cation ring C count, cation ring non-C count, longest alkyl length chain, number of alkyl chain, water presence (boolean), number of terminal nitrogens, longest chain C count, number of same type cations	Captures compositional effects and structural factors that govern steric packaging, hydrogen bonding, and polarity effects influencing dimensionality.
Interaction-based features	Total alkyl chain weight, ring length interaction, normalized count of terminal nitrogens per chain length, molecular weight per chain length, cation complexity, nitrogen/ weight ratio, compactness, hydrophilicity index, size complexity	Encode nonlinear or coupled effects between the original features reflecting how chemical interactions and packing affect HMH dimensionality.

Table 1: Summary of original features and engineered interaction-based features used for HMH dimensionality prediction, along with their description.

As shown in sections 4.1 and 4.2, the engineered interaction features improved F1-scores across all dimensionality classes, particularly in underrepresented categories (0D and 1D).

3.4. Class imbalance

The original HMH dataset (section 3.1) exhibits high class imbalance, with 2D structures being prevalent and 0D and 1D structures highly underrepresented (Figure 2). This imbalance presents a significant challenge for developing a machine learning model with high generalizability and prediction accuracy [58–61]. The models trained on the original distribution tend to overfit the majority class (2D) and perform poorly on the minority classes (0D and 1D), as reflected in the low F1-scores (section 4.1).

To address this issue, we applied a Synthetic Minority Oversampling Technique (SMOTE) [22, 23] to generate additional samples for minority classes. Producing a new synthetic sample requires interpolation between a randomly chosen minority class sample and one of its k-nearest neighbors in feature space:

$$x_{\text{new}} = x_i + \lambda(x_{nn} - x_i), \quad \lambda \in [0, 1]$$

SMOTE constructs a synthetic data point by taking a weighted combination of two real samples from the same minority class in the descriptor space. This means that each generated sample does not correspond to a physically realized compound, but to an interpolated representation with feature values lying between two chemically plausible points. The algorithm selects k nearest neighbors for each class, ensuring that the synthetic points populate regions of the feature space that were originally sparse. This feature-space interpolation effectively creates pseudo-structures that extend the distribution of the existing samples. The augmented samples maintain chemical realism since they were derived from physically valid feature relationships.

This approach improves model generalization and mitigates overfitting by increasing the diversity of the minority class samples without simply duplicating the existing ones. In our workflow, the augmentation balanced the dataset from its original 494 HMMs samples to a total of 1336 samples, with each dimensionality class containing an equal number of 334 samples per class. This reduces the bias towards the 2D majority class and improves decision boundaries between classes in the multi-class setting, as reflected in the ROC curves and per-class AUC values in section 4.2.

4. Results

In this section, we present the results of our dimensionality prediction framework for HMMs. Section 4.1 highlights the effect of feature engineering, particularly the introduction of interaction descriptors on classification performance. Section 4.2 evaluates the impact of class imbalance mitigation using SMOTE and its role in improving prediction for minority classes. Section 4.3 includes a detailed analysis of per-class F1-scores as well as cross-validation analysis.

4.1. Feature engineering

To improve the prediction accuracy of the structural dimensionality of HMMs, we developed chemically-informed interaction-based descriptors alongside original compositional features. The initial feature set includes basic atomic and molecular descriptors as well as features derived from them. While these features provided some separation between dimensionality classes, they were not sufficient to capture the subtleties of structural dimensionality patterns corresponding to lower-dimensional HMMs.

As shown in Table 2, the model trained on the original set of descriptors provides high accuracy for the majority class (2D, F1-score = 0.89), but underperforms significantly on the minority classes. Specifically, F1-score for 0D and 1D structures are 0.29 and 0.36, respectively, showing poor generalization to underrepresented dimensionalities. The support values in Table 2 show the number of test samples for each dimensionality class. These values highlight the dataset imbalance, with 2D structures being the majority class, and provide context for interpreting the class-specific F1-scores.

To address these limitations, we designed an additional set of descriptors based on key interaction components, such as metal-halide coordination geometries, organic cation steric effects, and measures of structural compactness. These descriptors were designed to better reflect the spatial and chemical connectivity relevant to structural dimensionality.

Dimensionality	F-1 score	Support
0	0.29	5
1	0.36	5
2	0.89	67
3	0.71	22

Table 2: Per-class F1-score and support for ML model trained on original descriptors.

To gain insight into the contribution of the individual descriptors we performed feature importance analysis using a gradient boosting classifier trained on the original descriptors as well as interaction-based descriptors. As shown in Figure 3, the most influential descriptors include 7 original features and 8 interaction-based features, most of which relate to the size and chemical composition of the cations. Structural descriptors such as the longest carbon chain, number of terminal nitrogen atoms, and hydrogen bonding also contribute significantly to prediction accuracy. These results highlight the relevance of both steric and electronic features in capturing the nuances relevant to dimensionality prediction.

The interaction descriptors enriched the feature space with the information that is highly important to distinguishing structural dimensionality of HMMs. This refinement was a necessary first step before applying additional class imbalance mitigation techniques described in the next section.




Figure 3: Feature importance analysis of all of the descriptors for HMMs dimensionality prediction.

4.2. Effect of class imbalance mitigation

The original dataset shows a significant class imbalance, with 2D structures constituting the majority of samples, while 0D and 1D classes are highly underrepresented (Figure 2). As demonstrated in Figure 4, this imbalance was addressed using the Synthetic Minority Oversampling Technique (SMOTE), which synthetically generates new samples of the minority classes by interpolating between the existing data




Figure 4: Per-class data distribution before and after SMOTE-based oversampling. The original dataset (blue) shows a strong imbalance favoring 2D structures. After the data augmentation (orange), all four dimensionality classes contain an equal number of samples, enabling better model generalization.

Dimensionality	F-1 score	Support
0	0.74	69
1	0.83	61
2	0.75	73
3	0.79	65

Table 3: F1-scores and support after applying interaction-based descriptors and SMOTE oversampling. Minority class performance improves substantially, while maintaining high performance for majority classes.

points. After applying SMOTE, the dataset was expanded from 494 samples to 1336 samples, with all four dimensionality classes being represented equally.

This resampling strategy led to substantial improvements in classification performance for minority classes. As shown in Table 3, the F1-score for 0D structures increased from 0.29 (see Table 2) to 0.74, and for 1D structures from 0.36 (Table 2) to 0.83. The F1-scores for the 2D and 3D majority classes remained high at 0.75 and 0.79, respectively. Support values in Table 3 show the number of test samples for each dimensionality class after applying SMOTE oversampling. The balanced support across 0D - 3D classes reflects the effectiveness of data-level augmentation in mitigating class imbalance.

The improved balance in class representation enabled the model to learn more generalizable patterns across all dimensionalities, rather than overfitting to the dominant 2D class. These results highlight the importance of data-level augmentation techniques in small and imbalanced materials datasets.

4.3. Model evaluation

To further evaluate the model performance and prediction accuracy across individual classes, we calculated one-vs-rest receiver operating characteristic (ROC) curves and corresponding area under the curve (AUC) scores. Figure 5 shows the ROC curves generated using the original (imbalanced) dataset. The classifier performs well on the majority classes (AUC = 0.89 for class 2 (2D HMs) and 0.94 for class 3 (3D HMs)), but fails to accurately distinguish minority classes, with AUC for class 0 (0D HMs) of only 0.54.




Figure 5: Multi-class ROC curves (one-vs-rest) before SMOTE augmentation. Minority class 0 (0D HMHs) performance is low (AUC=0.54), indicating poor classification due to class imbalance.




Figure 6: Multi-class ROC curves after SMOTE augmentation. AUC values improve for all classes, particularly for class 0 (AUC=0.95), demonstrating that oversampling significantly enhances classification performance for minority classes.

After applying SMOTE, performance significantly improves across all classes, as shown in Figure 6. The AUC for class 0 (0D HMHs) increases from 0.54 to 0.95, and classes 1, 2, and 3 also show improvements to 0.94, 0.93, and 0.96, respectively. These results indicate that SMOTE-based augmentation significantly improved class representation as well as the model’s ability to accurately distinguish between the dimensionality classes.

We conducted a 5-fold cross-validation using the final model trained on the interaction-based descriptors and the SMOTE-augmented dataset to assess the robustness and generalization capabilities of our model. As shown in Figure 7, the model exhibits consistently high performance across all folds, with F1-scores ranging from 0.947 to 0.98 and a mean F1-score of 0.964 (red dashed line). The resulting




Figure 7: Cross-validation F1-scores across five folds for the model trained on interaction-based descriptors with SMOTE. The model achieves a mean F1-score of 0.964 (red dashed line), with low variance across folds, demonstrating the model’s stability (blue line).

variation between folds corresponds to an uncertainty margin of $\pm 1.3\%$, confirming the robustness and reproducibility of the model’s predictive performance. This stability across folds suggests that the model is not overfitting to the specific subset of the data and is generalizable across a range of HMH compositions and structural dimensionalities. To evaluate model performance consistently, the dataset was split into training and test subsets using a stratified random split, ensuring that each dimensionality class was proportionally represented in both sets. The standard split ratio of 80% training and 20% testing was applied. To reduce the influence of random partitioning on model performance, this process was repeated across five independent random seeds, and the reported metrics were averaged across all runs.

Together, the cross-validation and ROC analysis confirm that our model achieves high predictive accuracy, is generalizable, and effectively mitigates the class imbalance effects that hinder minority class performance in the small HMH dataset.

5. Discussion

Accurately predicting the structural dimensionality of hybrid metal halides (HMHs) is highly important for the rational design of functional materials [62, 63]. Especially for applications where low-dimensionality is associated with enhanced optoelectronic properties [64]. One of the challenges is that the existing HMH datasets are often small and highly imbalanced. This makes the development of machine learning models much more complex. Our study shows that feature engineering and class imbalance mitigation can substantially improve the predictive accuracy while preserving interpretability.

Models trained on original descriptors show limitations in generalization and especially in capturing the features of 0D and 1D HMHs, resulting in low F1-scores. To circumvent these challenges we have introduced interaction-based descriptors, which encode steric hindrance, bonding geometry, and coordination effects to better reflect the structure-property relationships that influence HMHs’ structural dimensionality. Feature importance analysis confirms that these engineered descriptors were among the top contributors to model prediction accuracy. Especially the descriptors related to organic cation structure and hydrogen bonding, supporting observations from prior studies [36, 37].

Beyond the feature design, our results highlight that addressing data imbalance is essential for model

generalization and reliable prediction. Consistent with earlier findings that imbalance reduces minority class recall [22]. We have improved the minority class performance by applying SMOTE oversampling, which led to a substantial increase in the F1-score for 0D structures from 0.29 to 0.74. The model successfully integrated the synthetic data without overfitting. ROC analysis further confirmed this trend and shows significant improvements in AUC for minority classes after augmentation.

Ensemble-based models such as random forests and gradient boosted trees proved to be highly effective in this small-data regime, aligning with relevant studies [20, 21, 40, 48]. The cross-validation scores are fairly high and stable (mean F1 = 0.964), which suggests that the final model generalizes well and is not overly sensitive to data splits. This is especially important given the variability in small datasets. Our framework, including the integration of interaction descriptors, class imbalance mitigation, and ensemble learning, provides the possibility of applying ML models to other underexplored HMH properties.

While SMOTE oversampling helps to mitigate the effects of class imbalance, it relies on interpolation and may not capture the full diversity of HMH structures. Future work will explore generative modeling approaches to incorporate structures for minority classes. In addition, expanding the dataset with experimentally or computationally validated structures would further improve the model’s robustness. Coupling this framework with models that learn directly from compositional strings or text-based representations could reduce the reliance on handcrafted descriptors and accelerate discovery.

Overall, these results demonstrate that domain-informed feature design, combined with data-level augmentation and ensemble learning, provides a transferable workflow for small, imbalanced materials datasets. Although demonstrated here for HMH dimensionality prediction, this approach is broadly applicable to other structure-property prediction tasks in functional materials.

6. Conclusion and future work

In this study, we present a machine learning framework for the structural dimensionality prediction of hybrid metal halides (HMHs). We address two main challenges, namely small dataset size and significant class imbalance. By introducing the interaction-based descriptors and applying SMOTE-based oversampling, we significantly improved the prediction accuracy for minority classes (0D and 1D) without sacrificing the performance. Feature importance analysis highlighted the relevance of interaction descriptors, while ROC and cross-validation results confirm the model’s robustness and generalizability.

Our results show that combining domain-specific feature engineering with class imbalance mitigation leads to more reliable ML models for underrepresented regions of materials space. The methods presented here can be applied to broader HMH discovery efforts that go beyond the well-studied 2D structures. In addition to predictive gains, the model provided interpretable insights into the structural chemistry of HMHs. Feature importance analysis revealed that interaction features emphasizing the role of hydrogen bonding and steric interactions were strongly weighted. Collectively, these findings show that dimensionality is strongly governed by a balance between steric bulk, polarity, and hydrogen bonding, which were encoded in the interaction features.

Future work will involve several directions to extend this study. First, we will expand our dataset by generating more computational and experimental data to get a more in-depth understanding of the property-structure relationship. Additionally, large language models (LLMs) represent a rapidly emerging tool for materials research. Future work will involve training and fine-tuning LLMs to assist in structural dimensionality analysis by leveraging textual descriptors of chemical formulas. Together, these strategies can advance the development of data-driven tools for novel and underexplored materials.

Descriptor Name	Example Cation #1	Descriptor Value for Example Cation #1	Example Cation #2	Descriptor Value for Example Cation #2	Example Cation #3	Descriptor Value for Example Cation #3
Nitrogen Atoms bonded to 1 non-Hydrogen Atom (Terminal Nitrogen Count)	Tropylium 	0	Cyclohexylammonium 	1	Bis(aminoethyl)-thiophene 	2
Cation Ring Count	Guanidinium 	0	Cyclohexylammonium 	1	Benzodiimidazolium 	3
Largest Cation Ring Carbon Count	Imidazolium 	3	Cyclohexylammonium 	6	Tropylium 	7
Largest Cation Ring Non-carbon Count	Imidazolium 	2	Cyclohexylammonium 	0	Bis(aminoethyl)-thiophene 	1
Number of Alkyl Chains (2+ carbons) (Alkyl = attached to ring)	Phenylethylammonium 	1	Cyclohexylammonium 	0	Bis(aminoethyl)-thiophene 	2
Longest Alkyl Chain Carbon Count (Alkyl = attached to ring)	Phenylethylammonium 	2	Cyclohexylammonium 	0	Butylammonium 	0
Longest Chain Carbon Count (do not count ring)	Bis(aminoethyl)-thiophene 	2	Cyclohexylammonium 	0	Butylammonium 	4
Number of Organic Cations per Unit Cell	Methylammonium 	1	Bis(2-chloroethyl)ammonium 	2		
Water Present?	Cation + dihydrate 	1	Cyclohexylammonium 	0		

Figure 8: Original features. These chemical descriptors and the formulas for the organic and inorganic components provided the initial features for training our ML model. These features sought to capture key cation structural traits in a simple description.

7. Supplementary Information

This section provides an in-depth description of the original features used in this work (Figure 8). These features were derived from the HybriD³ dataset and capture structural and chemical properties of HMHs. Each descriptor encodes a specific aspect of the organic cation or inorganic framework that influences the resulting structural dimensionality. The first two descriptors (not shown in Figure 8) are text strings of the organic and inorganic formula units, respectively. For example, for HMH formula C₈H₈N₄PbI₄, the organic text string descriptor would be C₈H₈N₄, and the inorganic would be PbI₄. These formula descriptors contain compositional information such as atomic identities and similarities in available lattice

volume. While HMH total formulas do contain the number of duplicate cations within them (i.e. "2" in $(\text{C}_6\text{H}_{14}\text{N})_2\text{PbI}_4$), this information was not always included in organic formulas in HybriD³, so the number of duplicate cations was included as another descriptor. The terminal nitrogen count descriptor determines the number of nitrogen atoms bonded to a single non-hydrogen atom. N was chosen for this descriptor due to its prevalence in the dataset and importance to overall charge distributions within the organic cations. For instance, a cation with three terminal N will apply stronger forces on the surrounding inorganic lattice when compared to a cation of similar size with only a single terminal N, possibly distorting the inorganic structure and changing dimensionality of the HMH. The number of cation rings descriptor corresponds to the number of rings within the organic cation, which can affect steric packing and $\pi - \pi$ interactions that may influence dimensionality. The number of carbon vs. non-carbon atoms within the largest rings were also tracked as separate descriptors to provide more information on local charge. Number of alkyl chains captures the aspects of molecular shape and steric hindrance introduced by side chains. Relative length of the alkyl chains, and carbon number within the longest chain, also sought to simply capture alkyl chain contributions. Finally, water presence may affect stability and is a Boolean type descriptor. In total, these simple descriptors captured major geometric details of cation structures and proved to be a successful basis for the creation of interaction features for our ML model.

Acknowledgements

This work has been supported by the National Science Foundation grants DMREF 2323546, DMREF 2323547, and DMREF 2323548, and used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Data Access Statement

The HMHs dataset used in this work is available at the HybriD³ database: <https://materials.hybrid3.duke.edu>

Conflict of Interest Statement

The authors do not have conflicts of interest to declare.

Ethics Statement

No confidential information about human subjects is used in this manuscript.

References

- [1] Pikee Priya and N. R. Aluru. Accelerated design and discovery of perovskites with high conductivity for energy applications through machine learning. *npj Computational Materials*, 7, 06 2021.
- [2] Jino Im, Seongwon Lee, Tae-Wook Ko, Hyun Woo Kim, YunKyong Hyon, and Hyunju Chang. Identifying Pb-free perovskites for solar cells by machine learning. *npj Computational Materials*, 5, 03 2019.
- [3] D B Mitzi, K. Chondroudis, and C R Kagan. Organic-inorganic electronics. *IBM Journal of Research and Development*, 45:29–45, 01 2001.

- [4] Akihiro Kojima, Kenjiro Teshima, Yasuo Shirai, and Tsutomu Miyasaka. Organometal halide perovskites as visible-light sensitizers for photovoltaic cells. *Journal of the American Chemical Society*, 131:6050–6051, 04 2009.
- [5] Samuel D Stranks and Henry J Snaith. Metal-halide perovskites for photovoltaic and light-emitting devices. *Nature Nanotechnology*, 10:391–402, 05 2015.
- [6] Brandon R Sutherland and Edward H Sargent. Perovskite photonic sources. *Nature Photonics*, 10:295–302, 04 2016.
- [7] Kebin Lin, Jun Xing, Li Na Quan, F Pelayo, Xiwen Gong, Jianxun Lu, Liqiang Xie, Weijie Zhao, Di Zhang, Chuankong Yan, Wenqiang Li, Xinyi Liu, Yan Lu, Jeffrey Kirman, Edward H Sargent, Qihua Xiong, and Zhanhua Wei. Perovskite light-emitting diodes with external quantum efficiency exceeding 20 per cent. *Nature*, 562:245–248, 10 2018.
- [8] Li Zhang, Changjiu Sun, Tingwei He, Yuanzhi Jiang, Junli Wei, Yanmin Huang, and Mingjian Yuan. High-performance quasi-2d perovskite light-emitting diodes: from materials to devices. *Light Science & Applications*, 10, 03 2021.
- [9] Hendrik Heinz, Tzu-Jen Lin, Ratan Kishore Mishra, and Fateme S. Emami. Thermodynamically consistent force fields for the assembly of inorganic, organic, and biological nanostructures: The interface force field. *Langmuir*, 29:1754–1765, 01 2013.
- [10] Zhuo Wang, Zhehao Sun, Hang Yin, Xinghui Liu, Jinlan Wang, Haitao Zhao, Cheng Heng Pang, Tao Wu, Shuzhou Li, Zongyou Yin, and Xue-Feng Yu. Data-driven materials innovation and applications. *Advanced Materials*, 34, 07 2022.
- [11] Best research-cell efficiency chart — photovoltaic research — nrel, 2025. nrel.gov.
- [12] Bayrammurad Saparov and David B Mitzi. Organic–inorganic perovskites: Structural versatility for functional materials design. *Chemical Reviews*, 116:4558–4596, 04 2016.
- [13] Chenkun Zhou, Haoran Lin, Qingquan He, Liangjin Xu, Michael Worku, Maya Chaaban, Sujin Lee, Xiaoqin Shi, Mao-Hua Du, and Biwu Ma. Low dimensional metal halide perovskites and hybrids. *Materials Science and Engineering R Reports*, 137:38–65, 12 2018.
- [14] Pengcheng Xu, Xiaobo Ji, Minjie Li, and Wencong Lu. Small data machine learning in materials science. *npj Computational Materials*, 9, 03 2023.
- [15] Yuxuan Yao, Dan Han, Kieran B. Spooner, Xiaoyu Jia, Hubert Ebert, David O. Scanlon, and Harald Oberhofer. Adapting explainable machine learning to study mechanical properties of 2d hybrid halide perovskites. *Advanced Functional Materials*, 35, 08 2024.
- [16] Qin. Hybrid³: Materials property database for hybrid organic-inorganic perovskites. *APS March Meeting Abstracts*, 2021:X54.007, 2021.
- [17] Junya Wang, Pengcheng Xu, Xiaobo Ji, Minjie Li, and Wencong Lu. Feature selection in machine learning for perovskite materials design and discovery. *Materials*, 16:3134–3134, 04 2023.
- [18] Yilei Wu, Chang-Feng Wang, Ming-Gang Ju, Qiangqiang Jia, Qionghua Zhou, Shuaihua Lu, Xinying Gao, Yi Zhang, and Jinlan Wang. Universal machine learning aided synthesis approach of two-dimensional perovskites in a typical laboratory. *Nature Communications*, 15, 01 2024.
- [19] Tom Overman, Diego Klabjan, and Jean Utke. Ife: Interaction information based automated feature engineering. *2021 IEEE International Conference on Data Mining (ICDM)*, pages 803–808, 12 2024.
- [20] R. Suguna, J. Suriya Prakash, H. Aditya Pai, T. R. Mahesh, Venkatesan Vinoth Kumar, and Temesgen Engida Yimer. Mitigating class imbalance in churn prediction with ensemble methods and smote. *Scientific Reports*, 15, 05 2025.

- [21] Ibomoiye Domor Mienye and Yanxia Sun. A survey of ensemble learning: Concepts, algorithms, applications, and prospects. *IEEE Access*, 10:99129–99149, 2022.
- [22] Alberto Fernandez, Salvador Garcia, Francisco Herrera, and Nitesh V Chawla. Smote for learning from imbalanced data: Progress and challenges, marking the 15-year anniversary. *Journal of Artificial Intelligence Research*, 61:863–905, 04 2018.
- [23] N V Chawla, K W Bowyer, L O Hall, and W P Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 06 2002.
- [24] Ying Zhang and Chen Ling. A strategy to apply machine learning to small datasets in materials science. *npj Computational Materials*, 4, 05 2018.
- [25] Md Hosne Mobarak, Mariam Akter Mimona, Md. Aminul Islam, Nayem Hossain, Fatema Tuz Zohura, Ibnul Imtiaz, and Md Israfil Hossain Rimon. Scope of machine learning in materials research—a review. *Applied Surface Science Advances*, 18:100523, 12 2023.
- [26] Vishu Gupta, Kamal Choudhary, Francesca Tavazza, Carelyn Campbell, Wei-keng Liao, Alok Choudhary, and Ankit Agrawal. Cross-property deep transfer learning framework for enhanced predictive analytics on small materials data. *Nature Communications*, 12, 11 2021.
- [27] Qiuling Tao, Pengcheng Xu, Minjie Li, and Wencong Lu. Machine learning for perovskite materials design and discovery. *npj Computational Materials*, 7, 01 2021.
- [28] Jin-Soo Kim, Juhwan Noh, and Jino Im. Machine learning-enabled chemical space exploration of all-inorganic perovskites for photovoltaics. *npj Computational Materials*, 10, 05 2024.
- [29] Ming Chen, Zhenhua Yin, Zhicheng Shan, Xiaokai Zheng, Lei Liu, Zhonghua Dai, Jun Zhang, Shengzhong Frank Liu, and Zhuo Xu. Application of machine learning in perovskite materials and devices: A review. *Journal of Energy Chemistry*, 94:254–272, 03 2024.
- [30] Yiming Liu, Xinyu Tan, Jie Liang, Hongwei Han, Peng Xiang, and Wensheng Yan. Machine learning for perovskite solar cells and component materials: Key technologies and prospects. *Advanced Functional Materials*, 33, 02 2023.
- [31] Christopher J Bartel, Amalie Trewartha, Qi Wang, Alexander Dunn, Anubhav Jain, and Gerbrand Ceder. A critical examination of compound stability predictions from machine-learned formation energies. *npj Computational Materials*, 6, 07 2020.
- [32] G. Pilia, A. Mannodi-Kanakkithodi, B P Uberuaga, R Ramprasad, J E Gubernatis, and T. Lookman. Machine learning bandgaps of double perovskites. *Scientific Reports*, 6, 01 2016.
- [33] Marina R Filip and Feliciano Giustino. The geometric blueprint of perovskites. *Proceedings of the National Academy of Sciences*, 115:5397–5402, 05 2018.
- [34] Guocai Lv, Jie Cheng, Changxin Wang, and Mei Yang. Thermodynamic stability prediction of organic-inorganic hybrid perovskites using machine learning. *Journal of physics. Condensed matter : an Institute of Physics journal*, 37:10.1088/1361-648X/ae06e6, 13 2025.
- [35] Shuaihua Lu, Qionghua Zhou, Yixin Ouyang, Yilv Guo, Qiang Li, and Jinlan Wang. Accelerated discovery of stable lead-free hybrid organic-inorganic perovskites via machine learning. *Nature Communications*, 9, 08 2018.
- [36] Songyang Yuan, Yulin Liu, Jing Lan, Wenjian Yang, Hui Xiong, Wenzhe Li, and Jiandong Fan. Accurate dimension prediction for low-dimensional organic-inorganic halide perovskites via a self-established machine learning strategy. *The Journal of Physical Chemistry Letters*, 14:7323–7330, 08 2023.

- [37] Ruiyang Lyu, Curtis E. Moore, Tianyu Liu, Yongze Yu, and Yiyang Wu. Predictive design model for low-dimensional organic–inorganic halide perovskites assisted by machine learning. *Journal of the American Chemical Society*, 143:12766–12776, 08 2021.
- [38] Hironao Yamada, Chang Liu, Stephen Wu, Yukinori Koyama, Shenghong Ju, Junichiro Shiomi, Junko Morikawa, and Ryo Yoshida. Predicting materials properties with little data using shotgun transfer learning. *ACS Central Science*, 5:1717–1730, 09 2019.
- [39] Bangtan Zong, Jinshan Li, Tinghuan Yuan, Jun Wang, and Ruihao Yuan. Recent progress on machine learning with limited materials data: Using tools from data science and domain knowledge. *Journal of Materiomics*, 11:100916, 05 2025.
- [40] Omer Sagi and Lior Rokach. Ensemble learning: A survey. *Wiley Interdisciplinary Reviews Data Mining and Knowledge Discovery*, 8, 02 2018.
- [41] Lior Rokach. Ensemble-based classifiers. *Artificial Intelligence Review*, 33:1–39, 11 2009.
- [42] R. Polikar. Ensemble based systems in decision making. *IEEE Circuits and Systems Magazine*, 6:21–45, 01 2006.
- [43] Shahid Ali, Sreenivas Sremath Tirumala, and Abdolhossein Sarrafzadeh. Ensemble learning methods for decision making: Status and future prospects. *Unitec Research Bank (Unitec Institute of Technology)*, pages 211–216, 07 2015.
- [44] Basir Akbar, Kil To Chong, and Hilal Tayara. Analysis of ruddlesden-popper and dion-jacobson 2d lead halide perovskites through integrated experimental and computational analysis. *Battery energy*, 4, 02 2025.
- [45] Changjiao Li, Hua Hao, Ben Xu, Zhonghui Shen, Enhao Zhou, Dongbing Jiang, and Hanxing Liu. Improved physics-based structural descriptors of perovskite materials enable higher accuracy of machine learning. *Computational Materials Science*, 198:110714–110714, 07 2021.
- [46] D Vijay Anand, Qiang Xu, JunJie Wee, Keli Xia, and Tze Chien Sum. Topological feature engineering for machine learning based halide perovskite materials design. *npj Computational Materials*, 8, 09 2022.
- [47] M.A. Ganaie, Minghui Hu, A.K Malik, M Tanveer, and P.N. Suganthan. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115:105151–105151, 07 2022.
- [48] Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do tree-based models still outperform deep learning on tabular data?, 2022. arXiv:2207.08815.
- [49] Witold Wydmański, Oleksii Bulenok, and Marek Śmieja. Hypertab: Hypernetwork approach for deep learning on small tabular datasets, 2023. arxiv:2304.03543.
- [50] Dina Elreedy, Amir F Atiya, and Firuz Kamalov. A theoretical distribution analysis of synthetic minority oversampling technique (smote) for imbalanced learning. *Machine Learning*, 113:4903–4923, 01 2023.
- [51] Ying Li, Yali Yang, Peihua Song, Lian Duan, and Rui Ren. An improved smote algorithm for enhanced imbalanced data classification by expanding sample generation space. *Scientific Reports*, 15, 07 2025.
- [52] Thomas M Brenner, David A Egger, Leor Kronik, Gary Hodes, and David Cahen. Hybrid organic–inorganic perovskites: low-cost semiconductors with intriguing charge-transport properties. *Nature Reviews Materials*, 1, 01 2016.
- [53] Manoj K Jana, Ruyi Song, Haoliang Liu, Dipak Raj Khanal, Svenja M Janke, Rundong Zhao, Chi Liu, Z. Valy Vardeny, Volker Blum, and David B Mitzi. Organic-to-inorganic structural chirality

- transfer in a 2d hybrid perovskite and impact on rashba-dresselhaus spin-orbit coupling. *Nature Communications*, 11, 09 2020.
- [54] Assaf Shmuel, Oren Glickman, and Teddy Lazebnik. A comprehensive benchmark of machine and deep learning models on structured data for regression and classification. *Neurocomputing*, 655:131337, 11 2025.
- [55] Candice Bentéjac, Anna Csörgő, and Gonzalo Martínez-Muñoz. A comparative analysis of gradient boosting algorithms. *Artificial Intelligence Review*, 54:1937–1967, 08 2020.
- [56] Ruoyu Li, Qin Deng, Dong Tian, Daoye Zhu, and Bin Lin. Predicting perovskite performance with multiple machine-learning algorithms. *Crystals*, 11:818, 07 2021.
- [57] Tim Verdonck, Bart Baesens, María Óskarsdóttir, and Seppe vanden Broucke. Special issue on feature engineering editorial. *Machine Learning*, 113:3917–3928, 08 2021.
- [58] Jian Jiang, Chunhuan Zhang, Lu Ke, Nicole Hayes, Yueying Zhu, Huahai Qiu, Bengong Zhang, Tianshou Zhou, and Guo-Wei Wei. A review of machine learning methods for imbalanced data challenges in chemistry. *Chemical Science*, 16, 01 2025.
- [59] Bashar Hamad Aubaidan, Rabiah Abdul Kadir, Mohamed Taha Lajb, Muhammad Anwar, Kashif Naseer Qureshi, Bakr Ahmed Taha, and Kayhan Ghafoor. A review of intelligent data analysis: Machine learning approaches for addressing class imbalance in healthcare - challenges and perspectives. *Intelligent Data Analysis*, 29, 02 2025.
- [60] Wuxing Chen, Kaixiang Yang, Zhiwen Yu, Yifan Shi, and C. L. Philip Chen. A survey on imbalanced learning: latest research, applications and future directions. *Artificial Intelligence Review*, 57, 05 2024.
- [61] Shuxian Zhao, Jie Gui, Minjing Dong, Baosheng Yu, Zhipeng Gui, Lu Dong, Yuan Yan Tang, and James Tin-Yau Kwok. A survey on small sample imbalance problem: Metrics, feature analysis, and solutions, 2025. arXiv:2504.14800v1.
- [62] Ling Mao and Changying Xiang. A comprehensive review of machine learning applications in perovskite solar cells: Materials discovery, device performance, process optimization and systems integration. *Materials Today Energy*, 47:101742, 01 2025.
- [63] Bonghyun Jo, Wenning Chen, and Hyun Suk Jung. Comprehensive review of advances in machine-learning-driven optimization and characterization of perovskite materials for photovoltaic devices. *Journal of Energy Chemistry*, 101:298–323, 10 2024.
- [64] Feiyue Lu, Yanyan Liang, Nana Wang, Lin Zhu, and Jianpu Wang. Machine learning for perovskite optoelectronics: a review. *Advanced Photonics*, 6, 08 2024.