

# Known Unknowns: Out-of-Distribution Property Prediction in Materials and Molecules

Nofit Segal<sup>1\*</sup>, Aviv Netanyahu<sup>2\*</sup>, Kevin P. Greenman<sup>3,4</sup>,  
Pulkit Agrawal<sup>2†</sup>, Rafael Gómez-Bombarelli<sup>1†</sup>

<sup>1</sup>Materials Science and Engineering, MIT, Memorial Dr, Cambridge, 02139, MA, USA.

<sup>2</sup>Electrical Engineering and Computer Science, MIT, Vassar St, Cambridge, 02139, MA, USA.

<sup>3</sup>Chemical Engineering, Catholic Institute of Technology, Broadway, Cambridge, 02142, MA, USA.

<sup>4</sup>Chemistry, Catholic Institute of Technology, Broadway, Cambridge, 02142, MA, USA.

\*These authors contributed equally.

†Equal Advising.

Corresponding authors: [pulkitag@mit.edu](mailto:pulkitag@mit.edu); [rafagb@mit.edu](mailto:rafagb@mit.edu);

## Abstract

Discovery of high-performance materials and molecules requires identifying extremes with property values that fall outside the known distribution. Therefore, the ability to extrapolate to out-of-distribution (ood) property values is critical for both solid-state materials and molecular design. Our objective is to train predictor models that extrapolate zero-shot to higher ranges than in the training data, given the chemical compositions of solids or molecular graphs and their property values. We propose using a transductive approach to ood property prediction, achieving improvements in prediction accuracy. In particular, the True Positive Rate (TPR) of ood classification of materials and molecules improved by 3x and 2.5x, respectively, and precision improved by 2x and 1.5x compared to non-transductive baselines. Our method leverages analogical input-target relations in the training and test sets, enabling generalization beyond the training target support, and can be applied to any other material and molecular tasks.

**Keywords:** machine learning, materials property prediction, extrapolation, out-of-distribution, transduction

## 1 Introduction

Designing new materials and molecules is essential for the development of new technologies. Traditionally, this design process involves extensive experimental iteration or high-throughput methods to screen databases, which are time-consuming and resource-intensive [1,2]. As a result, there is increasing interest in applying machine learning (ML) techniques to accelerate the discovery of materials and molecules with desired properties [1–5]. There is particular interest in property values that are outside the known property value distribution as these will most likely lead to discovering new materials that will, in turn, unlock new capabilities and technologies.

One strategy for finding materials and molecules with desired properties is inverse design through conditional generation where materials with out-of-distribution (ood) property values are unavailable, and the goal is to generate them.[2, 4–8]. A complementary approach is screening a large database of candidates based on their predicted properties. [9–14]. In this setting, the objective is to identify ood materials and molecules from a set of known candidates with unknown property values. However, both approaches typically struggle when property values fall outside the training distribution. [4, 11, 13, 15–17]. Enhancing extrapolative capabilities in property prediction would improve the screening of large candidate spaces in terms of precision by identifying promising compounds and molecules with exceptional properties. This approach could help guide further synthesis and computational efforts, ultimately advancing materials and molecular design.

Extrapolation in materials science can refer to both the domain (materials space) and range (property values) of the predictive function. It is often used to refer *generalization* to unseen classes of materials structures and chemical spaces — for example, training on metals and predicting ceramics or training on artificial molecules and predicting natural products. Here, we address extrapolation in material property values.

When ood generalization is defined with respect to the materials space, extrapolation often reduces to interpolation. This occurs because test sets tend to remain within the same distribution as the training data representation space [18]. This includes predictive models employing leave-one-cluster-out extrapolation strategies [17, 19, 20], as well as generative approaches designed to achieve ood generalization to structures with different atomic compositions or larger numbers of atoms [21–23].

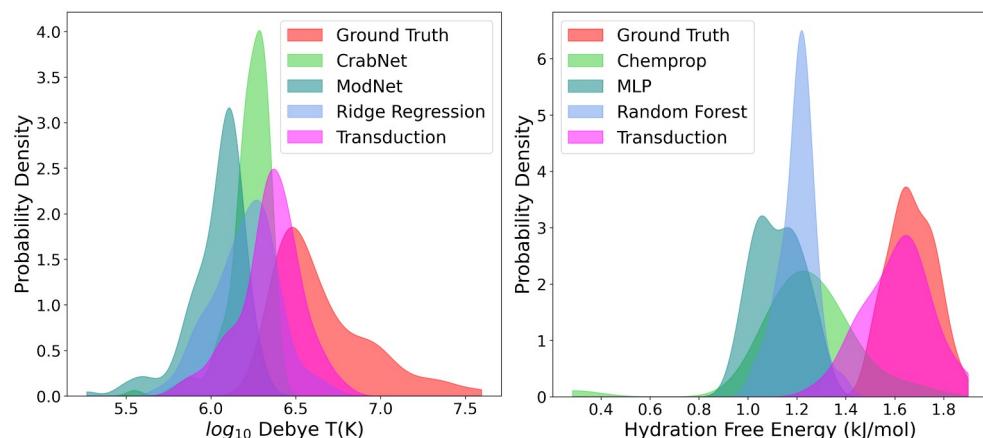
When ood generalization is defined with respect to the range of the predictive function, classical machine learning models face significant challenges in extrapolating property predictions through regression. As a result, several approaches have shifted toward classifying ood materials instead [15, 16].

In this work, we focus on predicting property values from materials representations and exploring zero-shot extrapolation to property value ranges beyond the training distribution. We adapt the Bilinear Transduction method [24], which has shown success in extrapolation tasks in other domains, and investigate its effectiveness for property value extrapolation in materials science (Figure 1). The core idea of Bilinear Transduction is reparameterizing the prediction problem. Rather than making property value predictions on a new candidate material, predictions are made based on a known training example and the difference in representation space between the two

materials. During inference, property values are predicted similarly – based on a chosen training example and the difference between it and the new sample. This method has the potential to extrapolate by learning *how* property values change as a function of material differences rather than predicting these values from new materials.

Our work advances the field by enhancing extrapolation capabilities, enabling more accurate predictions of material and molecular properties beyond known value ranges. Specifically, our method significantly improves ood precision, ensuring a higher percentage of high-potential candidates with desirable properties during the screening of large databases. This approach has the potential to streamline the identification process by reducing time and resource expenditure on low-potential candidates, thereby accelerating the discovery of materials and molecules with high synthesis viability.

We present a comprehensive evaluation of our approach using common benchmarks, spanning high-throughput computational and experimental datasets for solid-state materials and molecules.



**Fig. 1: Bilinear Transduction prediction distribution on out-of-distribution values is closer to the ground truth distribution compared with other machine learning methods.** **(Left)** probability density functions for out-of-distribution **materials** Debye temperature. Compared with CrabNet, ModNet, and Ridge Regression, Bilinear Transduction (purple) has the most overlap with the ground truth (red). **(Right)** probability density functions for out-of-distribution **molecular** hydration free energy. Ground truth (red) and machine learning method predictions estimated with kernel density estimation – Chemprop, MLP, Random Forest, and Bilinear Transduction (purple) which has the most overlap with the ground truth.

**Table 1:** Solids ood mean average prediction error and standard error of the mean.

Dataset	Property	#Samples	Ridge Reg. [15]	MODNet [13]	CrabNet [11]	Ours
AFLOW [25]	Band Gap [eV]	14123	$2.59 \pm 0.03$	$2.65 \pm 0.04$	<b><math>1.47 \pm 0.03</math></b>	$1.51 \pm 0.04$
	Bulk Modulus [GPa]	2740	$74.0 \pm 3.8$	$93.06 \pm 3.7$	$59.25 \pm 3.2$	<b><math>47.4 \pm 3.4</math></b>
	Debye Temperature [K]	2740	$0.45 \pm 0.03$	$0.62 \pm 0.03$	$0.38 \pm 0.02$	<b><math>0.31 \pm 0.02</math></b>
	Shear Modulus [GPa]	2740	$0.69 \pm 0.03$	$0.78 \pm 0.04$	$0.55 \pm 0.02$	<b><math>0.42 \pm 0.02</math></b>
	Thermal Conductivity [ $\frac{W}{mK}$ ]	2734	$1.07 \pm 0.05$	$1.5 \pm 0.05$	$0.97 \pm 0.03$	<b><math>0.83 \pm 0.04</math></b>
Matbench [10]	Thermal Expansion [K $^{-1}$ ]	2733	$0.44 \pm 0.02$	$0.47 \pm 0.02$	<b><math>0.37 \pm 0.02</math></b>	$0.39 \pm 0.02$
	Band Gap [eV]	2154	$6.37 \pm 0.28$	$3.26 \pm 0.13$	$2.70 \pm 0.13$	<b><math>2.54 \pm 0.16</math></b>
	Refractive Index	4764	$14.4 \pm 2.0$	$4.24 \pm 0.48$	$3.92 \pm 0.5$	<b><math>3.81 \pm 0.49</math></b>
MP [26]	Yield Strength [MPa]	312	$972 \pm 34$	$731 \pm 82$	$740 \pm 49$	<b><math>591 \pm 62</math></b>
	Bulk Modulus [GPa]	6307	$151 \pm 14$	$60.1 \pm 3.9$	$57.8 \pm 4.2$	<b><math>45.8 \pm 3.9</math></b>
	Elastic Anisotropy	6331	$165 \pm 17$	$60.0 \pm 4.5$	$61.4 \pm 4.6$	<b><math>59.8 \pm 4.5</math></b>
	Shear Modulus [GPa]	6184	$134.5 \pm 7.2$	$65.6 \pm 2.5$	$65.3 \pm 2.8$	<b><math>63.2 \pm 2.6</math></b>

**Table 2:** Solids 30% extrapolative precision

Dataset	Property	Ridge Reg. [15]	MODNet [13]	CrabNet [11]	Ours
AFLOW [25]	Band Gap	0.16	0.15	0.14	<b>0.22</b>
	Bulk Modulus	0.22	0.30	0.17	<b>0.40</b>
	Debye Temperature	0.19	0.06	0.07	<b>0.20</b>
	Shear Modulus	<b>0.07</b>	<b>0.07</b>	0.06	<b>0.07</b>
	Thermal Conductivity	<b>0.15</b>	0.05	0.06	0.09
Matbench [10]	Thermal Expansion	0.26	0.28	0.15	<b>0.36</b>
	Band Gap	0.05	0.11	0.17	<b>0.20</b>
	Refractive Index	0.17	0.32	0.20	<b>0.40</b>
MP [26]	Yield Strength	0.00	0.00	0.00	<b>0.67</b>
	Bulk Modulus	0.22	0.21	0.14	<b>0.60</b>
	Elastic Anisotropy	0.01	<b>0.09</b>	0.03	0.03
	Shear Modulus	0.28	0.27	0.20	<b>0.53</b>

## 2 Results

### 2.1 Solids

We demonstrate the extrapolation capability of bilinear transduction on three common benchmarks for solid materials property prediction tasks involving 12 prediction tasks, with dataset sizes ranging between  $\sim 300$  and  $\sim 14000$ , and compared against three baselines.

The solids datasets include material compositions and their property values. **AFLOW** contains material property values obtained from high-throughput calculations [25]. Following Kauwe et al. [15], who evaluated classical ML algorithms on AFLOW, we curate a subset of six properties: band gap, bulk modulus, Debye temperature, shear modulus, thermal conductivity, and thermal expansion, out of which the last four are scaled by applying a base 10 logarithm. **Matbench** is an automated leaderboard for benchmarking ML algorithms predicting solid material properties [10]. Matbench contains three composition-based regression tasks: experimentally measured band gap [12], experimentally measured yield strength of steels [27], and calculated

refractive index [28]. **Materials Project (MP)** provides materials and their property values derived from high-throughput calculations [26]. Following Wang et al. [11], we focus on bulk modulus, shear modulus, and ratio of elastic anisotropy. In cases of duplicate compositions, we retain the entry with the target value corresponding to the lowest formation enthalpy.

We compare with **Ridge Regression**, the strongest method in Kauwe et al. [15], who evaluate classical ML algorithms on ood property values, and also **MODNet** [13] and **CrabNet** [11], all of which use stoichiometry as input and are leading models in composition-based property prediction.

Table 1 presents the mean average error (MAE) for ood predictions on solids and molecules. Bilinear transduction consistently outperforms or performs comparably to the baseline methods.

In addition, we report model performance on a purely extrapolative task of discovering high-performance extremes: identifying the 30% of test samples with the highest property values, which we refer to as *the top ood candidates* (Table 2). Specifically, it measures the fraction of true top ood candidates correctly identified among the top predicted ood candidates. The score is calculated as the ratio of true top ood candidates included in the top predictions to the total number of predicted top candidates. The held-out set is composed of both in-distribution and ood samples in equal parts (see Section 3), so the 30% of the test dataset corresponds to 60% of the ood-sourced part of the test set. Furthermore, this metric re-weights misranked in-distribution errors 19-fold, reflecting the 95:5 split between in-distribution validation and ood test sets.

Figure 1 demonstrates that Bilinear Transduction produces a prediction for band gap distribution that is closer to the ood ground truth distribution. Figure 2 shows that Bilinear Transduction extrapolates to some extent, whereas the other baselines do not exceed the training support threshold.

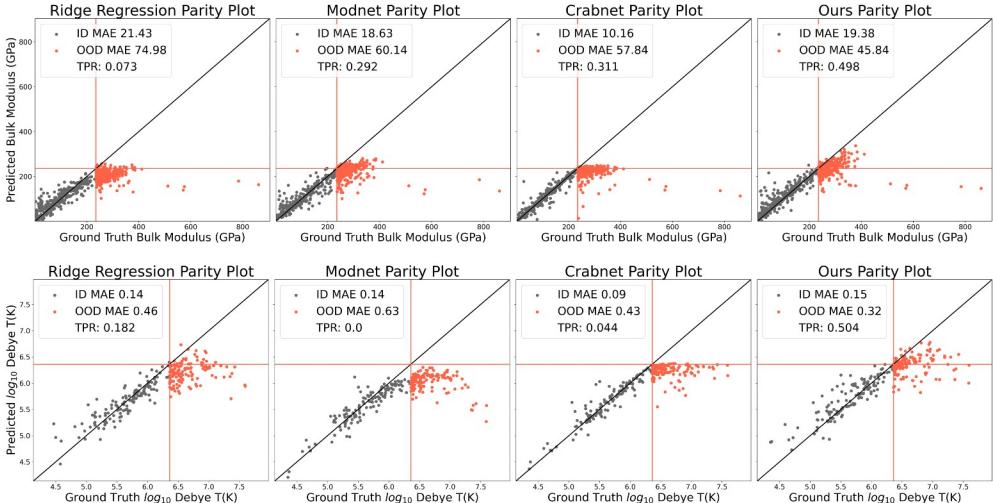
## 2.2 Molecules

For molecules, we sourced one dataset involving 4 graph-to-property prediction tasks with dataset sizes ranging from  $\sim 600$  to 4200 and compared against three baselines.

**Table 3:** Molecules ood mean average prediction error and standard error of the mean.

Dataset	Property	#Samples	Chemprop [29]	Random Forest [30]	MLP [31]	Ours
MoleculeNet [32]	ESOL [ $\frac{mol}{L \cdot kj}$ ]	1128	$0.47 \pm 0.04$	$0.67 \pm 0.04$	$0.5 \pm 0.04$	<b><math>0.42 \pm 0.04</math></b>
	Freesolv [ $\frac{mol}{L}$ ]	643	$0.44 \pm 0.03$	$0.42 \pm 0.02$	$0.5 \pm 0.02$	<b><math>0.08 \pm 0.01</math></b>
	Lipophilicity [ $\log D$ ]	4200	$0.75 \pm 0.02$	$1.02 \pm 0.02$	$0.9 \pm 0.02$	<b><math>0.7 \pm 0.02</math></b>
	BACE binding [ $IC50$ ]	1513	$1.03 \pm 0.06$	$0.93 \pm 0.05$	$0.95 \pm 0.07$	<b><math>0.73 \pm 0.05</math></b>

**MoleculeNet** includes molecular graphs encoded as SMILES representations [33] and their property values derived from high-throughput calculations and experimental trials [32]. We focus on physical chemistry and biophysics properties suitable for regression – ESOL, Freesolv, Lipophilicity and BACE binding affinity.



**Fig. 2: In-distribution and out-of-distribution Bulk Modulus (top) and Debye Temperature (bottom) predictions vs ground truth values.** While Ridge Regression [15], MODNet [13], CrabNet [11], and Bilinear Transduction (ours), perform well within the training distribution (gray dots bounded by the red horizontal line), Bilinear Transduction extends predictions beyond this range on ood data (red dots) closest to the ground truth, achieving a lower ood MAE and higher TPR.

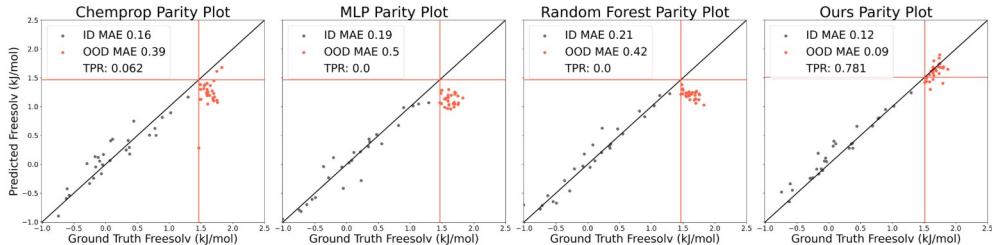
**Table 4:** Molecules 30% extrapolative precision

Dataset	Property	Chemprop [29]	Random Forest [30]	MLP [31]	Ours
MoleculeNet [32]	ESOL	<b>0.17</b>	0.11	0.15	0.16
	Freesolv	0.50	0.22	0.44	<b>0.83</b>
	Lipophilicity	0.05	0.05	<b>0.07</b>	<b>0.07</b>
	BACE binding	0.04	<b>0.07</b>	0.05	0.06

We compare with **Random Forest (RF)** [30], a classical ML tree-based method, and **Multi Layer Perceptron (MLP)** which use RDKit descriptors as input [34]. These serve as ablations of our method, using the same representation with partial structural information as we do. We also compare with **Chemprop** [29], a leading method for property prediction from molecular graphs via message-passing. Chemprop learns representations from structural information that is not explicitly available in the descriptor-based representation.

Table 3 presents the mean average error (MAE) for ood predictions on solids and molecules. Bilinear transduction consistently outperforms or performs comparably to the baseline methods. Figure 1 demonstrates that Bilinear Transduction produces a prediction for the Freesolv data distribution that is closer to the ood ground truth distribution and Figure 3 shows that Bilinear Transduction extrapolates, whereas the other baselines rarely surpass the boundary marking the beginning of the test range.

The fraction of true top ood candidates correctly identified among the top predicted ood candidates is presented in Table 4.



**Fig. 3: In-distribution, and out-of-distribution Freesolv predictions vs.**

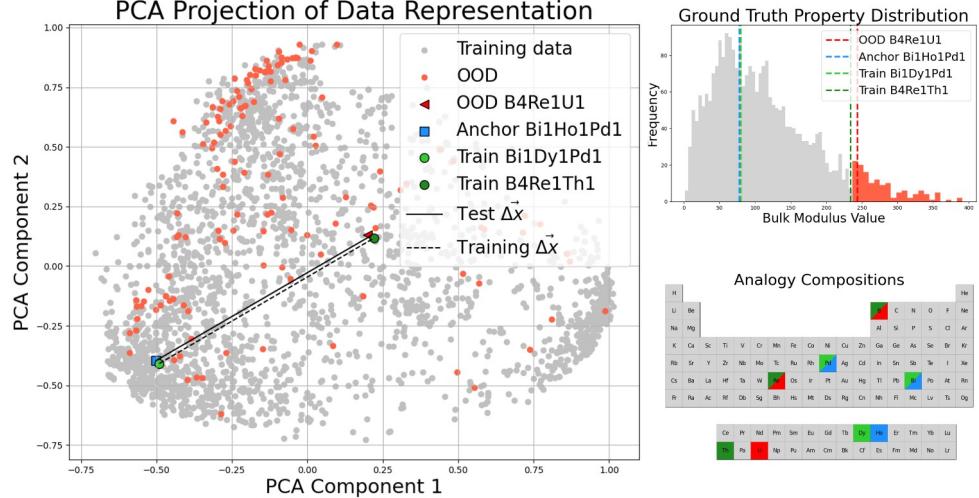
**ground truth values.** While Chemprop [29], RF, MLP and Bilinear Transduction (ours), perform well within the training distribution (gray dots bounded by the red horizontal line), only Bilinear Transduction performs well beyond this range on ood data (red dots).

## 2.3 Learning Using Analogies

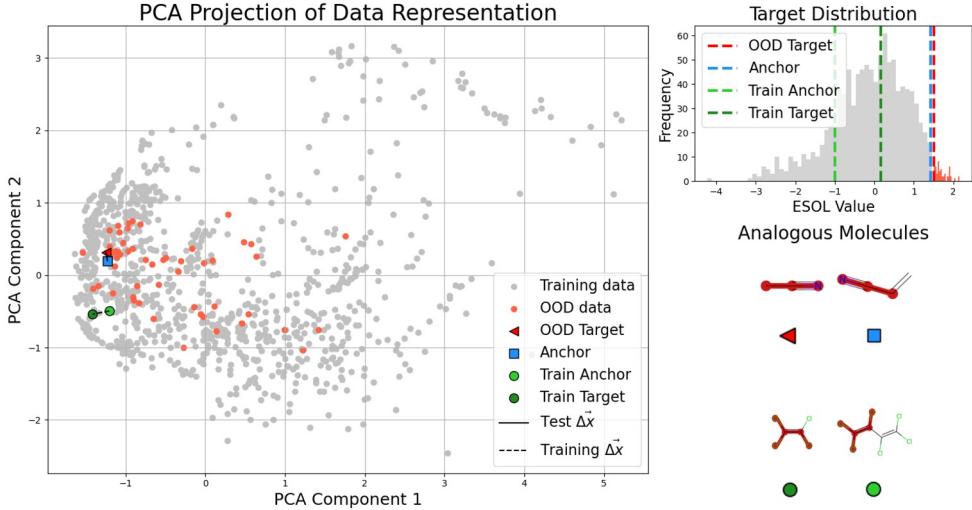
The success of transduction in these tasks is related to one of the tenets of chemistry: *similar materials have similar properties*. The transductive approach effectively builds on this idea, by proposing that *similar changes in chemical compounds or molecular structure imply similar changes in properties*. At inference, our approach selects the anchor by minimizing the difference between  $\Delta x_{te,an}$  and  $\Delta x_{tr}$ . We demonstrate how these algebraic operations in the embedding space relate to chemical changes as measured in the domain.

For solids, these operations are expressed as element changes. Figure 4 demonstrates this in bulk modulus inference of an ood sample with stoichiometry  $B_4ReU$ . The model selects in-distribution  $BiHoPd$  as the anchor, analogous to training anchor  $BiDyPd$  and training target  $B_4ReTh$ . Specifically, the compositions of these anchors ( $BiDyPd$  and  $BiHoPd$ ), and the targets ( $B_4ReTh$  and  $B_4ReU$ ), differ by only one neighboring f-block element: Dy ( $Z=66$ ) to Ho ( $Z=67$ ), and Th ( $Z=90$ ) to U ( $Z=92$ ). See Appendix 6.1.3 for more examples.

For molecules, these operations manifest as structural similarity. This is notable given that the RDKit descriptor vector used as input lacks detailed structural and connectivity information from the SMILES representation. Figure 5 illustrates this in ESOL inference, with the Maximum Common Structure (MCS) highlighted (Figure 5c) between the anchor-ood pair, and the training anchor-target pair. Each pair’s structures show high similarity, differing by the addition of a conjugated double bond that extends the molecular backbone. See Appendix 6.1.3 for additional examples.



**Fig. 4: Analogy Visualization Solids** AFLOW bulk modulus ood predictions are based on in-distribution anchors, that paired with ood points, form analogies to training pairs. (a) PCA plot of all samples in the dataset. **ood-anchor** difference is similar to **training-anchor** difference. (b) Ground truth bulk modulus training and **test** distributions and **ood**, **anchor**, and analogous **training** point and **anchor** values. (c) Analogy compositional visualization. **ood** and **training** points differ by one neighboring f-block element. So do **ood** and **training** anchors.



**Fig. 5: Analogy Visualization Molecules** MoleculeNet ESOL ood predictions are based on differences between in-distribution anchors and ood points, that form analogies to training pairs. (a) **ood-anchor** and training **target-anchor** differences. (b) Ground truth **training** and **test** distributions and **ood**, **anchor**, and analogous training **target**, **anchor** values. (c) Analogous molecule pairs. **ood-anchor** and training **target-anchor** similarities measured with MCS highlighted in red.

### 3 Methods

Let  $X \subseteq \mathbb{R}^n$  denote a representation space of data points, such as RDKit descriptors derived from Simplified Molecular Input Line Entry System (SMILES) for molecules, and composition-based descriptors derived from elemental properties for solids. We choose fixed descriptor-based representations as they offer interpretable features, useful for smaller datasets that are common in materials science. Additionally, for solids, a composition-based approach enables more robust predictions, as it implies weaker assumptions about the material [35]. Let  $Y \subseteq \mathbb{R}$  denote the space of labels – in our case, material property values. We split the data into training, validation and test sets such that ood test samples  $D^{\text{te}} \subseteq Y$  are 5% of the data with the highest property values, an in-distribution validation set is a random sample of 5% of the remaining data, and the training set  $D^{\text{tr}} = \{(x, y)\}_{n \in X \times Y}$  is the rest.

We adapt Bilinear Transduction [24], a regression method, which offers the following formulation (see description in Appendix 6.3 and adapted version summary in Algorithm 1). Let  $\Delta X = \{x_i - x_j \mid x_i, x_j \in X\}$  denote the differences distribution. During training, rather than predicting  $h_{\theta} : X \rightarrow Y$ , predictor  $h_{\theta} : \Delta X \times X \rightarrow Y$  predicts value  $y_2 \in D^{\text{tr}}_Y$  for data point  $x_2 \in D^{\text{tr}}_X$  given anchor point  $x_1 \in D^{\text{tr}}_X$  and the difference between them  $\Delta x_{21} = x_2 - x_1 \in D^{\text{tr}}_{\Delta X}$ . The predictor  $h_{\theta}(\Delta x, x) = f_{\theta}(\Delta x)g_{\theta}(x)$  is implemented as a bilinear function in non-linear embeddings of  $\Delta x$  and  $x$ . I.e. during

training we predict the property value  $y_i$  of material  $x_i$  from material  $x_j$  and their difference  $x_i - x_j$ , where material  $x_j$  has a lower property value.

Given test point  $x_{te} \in D_X^{te}$ , its predicted value is  $h_\theta(\Delta x_{te,an}, x_{an})$ . Anchor  $x_{an} \in D_X^{tr}$  is the training point that minimizes the distance between its difference with the test point  $\Delta x_{te,an} = x_{te} - x_{an}$ , and differences within the training distribution  $D_{\Delta X}^{tr}$ . Formally,  $x_{an} = \operatorname{argmin}_{x_i \in D_X^{tr}} \{ \|\Delta x_{te,i} - \Delta x\|_2 \mid \Delta x \in D_{\Delta X}^{tr} \}$ .

Descriptor-based features encapsulate fundamental chemical and physical information that directly influences materials and molecular characteristics. Therefore, the difference between feature vectors is related, possibly intricately, to the change in property value. Bilinear Transduction has the potential to extrapolate by learning *how* property values change as a function of compositional differences instead of predicting these values directly.

---

### Algorithm 1 Bilinear Transduction for Molecules and Materials

---

- 1: **Input:** Training set  $(x_1, y_1), \dots, (x_n, y_n)$
- 2: **Train:** Train  $\theta$  on loss  $L(\theta) = \sum_{i=1}^n \sum_{j: y_j < y_i} \ell(h_\theta(x_i - x_j, x_j), y_i)$
- 3: **Test:** For each new  $x_{te}$ , let  $x_{an} = \operatorname{argmin}_{x_{an} \in D_X^{tr}} \{ \|x_{te} - x_{an} - \Delta x_{tr}\|_2 \mid \Delta x_{tr} \in D_{\Delta X}^{tr} \}$ , and predict

$$y = h_\theta(x_{te} - x_{an}, x_{an})$$


---

## 4 Data availability

The raw data, along with processing scripts, are available at <https://github.com/learningmatter-mit/matex>.

## 5 Code availability

An open-source implementation of Matex is available at <https://github.com/learningmatter-mit/matex>.

## 6 Supplementary Information

### 6.1 Additional Results

#### 6.1.1 In-Distribution Predictions

In Table 5 we demonstrate that all methods perform well, making Bilinear Transduction a strong predictor both in and out of distribution alike.

**Table 5:** Solids (top) and molecules (bottom) **In-distribution** mean average prediction error and standard error of the mean.

Dataset	Property	Ridge Reg. [15]	MODNet [13]	CrabNet [11]	Ours
AFLOW [25]	Band Gap [eV]	0.87 ± 0.04	0.56 ± 0.02	0.35 ± 0.02	0.61 ± 0.02
	Bulk Modulus [GPa]	15.41 ± 1.21	15.1 ± 1.3	8.01 ± 1.05	13.06 ± 1.6
	Debye Temperature [K]	0.13 ± 0.01	0.13 ± 0.01	0.09 ± 0.01	0.14 ± 0.01
	Shear Modulus [GPa]	0.31 ± 0.03	0.27 ± 0.02	0.19 ± 0.02	0.31 ± 0.03
	Thermal Conductivity [ $\frac{W}{mK}$ ]	0.47 ± 0.03	0.4 ± 0.02	0.31 ± 0.03	0.43 ± 0.04
Matbench [10]	Thermal Expansion [ $K^{-1}$ ]	0.11 ± 0.01	0.07 ± 0.01	0.04 ± 0.0	0.11 ± 0.01
	Band Gap [eV]	1.75 ± 0.07	0.32 ± 0.03	0.24 ± 0.03	0.49 ± 0.05
	Refractive Index	1.00 ± 0.05	0.15 ± 0.01	0.13 ± 0.02	0.16 ± 0.01
MP [26]	Yield Strength [MPa]	411 ± 75	62.5 ± 11.8	52.4 ± 18.1	156 ± 33
	Bulk Modulus [GPa]	36.9 ± 1.21	18.63 ± 1.16	10.2 ± 0.8	19.4 ± 1.3
	Elastic Anisotropy	22.00 ± 2.01	2.12 ± 0.34	1.24 ± 0.06	2.4 ± 0.3
MoleculeNet [32]	Shear Modulus [GPa]	35.7 ± 1.2	12.8 ± 0.7	8.75 ± 0.63	13.6 ± 0.7
	ESOL [ $\frac{mol}{\frac{mol}{kj}}$ ]	0.28 ± 0.03	0.25 ± 0.03	0.28 ± 0.03	0.29 ± 0.04
	Freesolv [ $\frac{mol}{mol}$ ]	0.16 ± 0.02	0.20 ± 0.06	0.18 ± 0.06	0.12 ± 0.02
	Lipophilicity [ $\log D$ ]	0.36 ± 0.02	0.40 ± 0.02	0.38 ± 0.03	0.46 ± 0.03
	BACE binding [IC50]	0.45 ± 0.04	0.37 ± 0.04	0.43 ± 0.05	0.51 ± 0.05
		Chemprop [29]	Random Forest [30]	MLP [31]	

#### 6.1.2 True Positive Rate (TPR)

In Table 6, we compare the TPR of ood detection, showcasing our ability to produce candidate materials with outstanding property values, with an improvement of 3x compared with the strongest baseline.

#### 6.1.3 Analogy Analysis

##### Solids

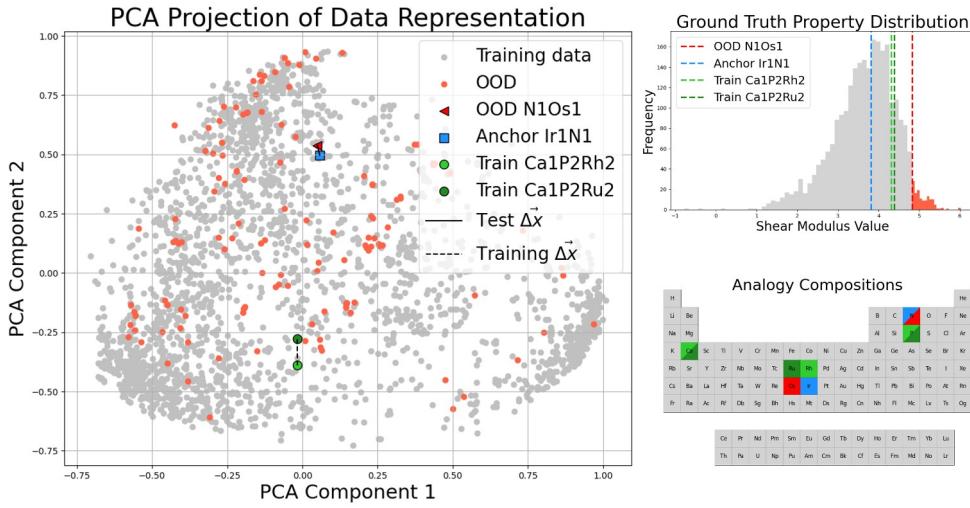
Figure 6 describes shear modulus inference of ood NOs via anchor Nlr, analogous to training anchor  $CaP_2Rh_2$  and training target  $CaP_2Ru_2$ . In this case, the training anchor and target ( $CaP_2Rh_2$  and  $CaP_2Ru_2$ ) and test anchor and target (Nlr and NOs) differ by one d-block element: Rh ( $Z=45$ ) to Ru ( $Z=44$ ) and Ir ( $Z=77$ ) to Os ( $Z=76$ ). Figure 7 describes bulk modulus inference of ood  $NiNb_4$  via anchor  $HfNbP$ , analogous to training anchor  $HfMoP$  and training target  $Mo_8P_5$ . Additional examples include shear modulus prediction for ood  $NbSiIr$  via anchor  $NbSiPt$  analogous to training anchor  $GeNbIr$  and target  $GeNbPt$ , and ood  $ReSiNb$  via anchor  $OsSiZr$  analogous to training anchor  $GeIrNb$  and target  $GeIrPt$ .

**Table 6:** Solids (top) and molecules (bottom) ood **True Positive Rate (TPR)**.

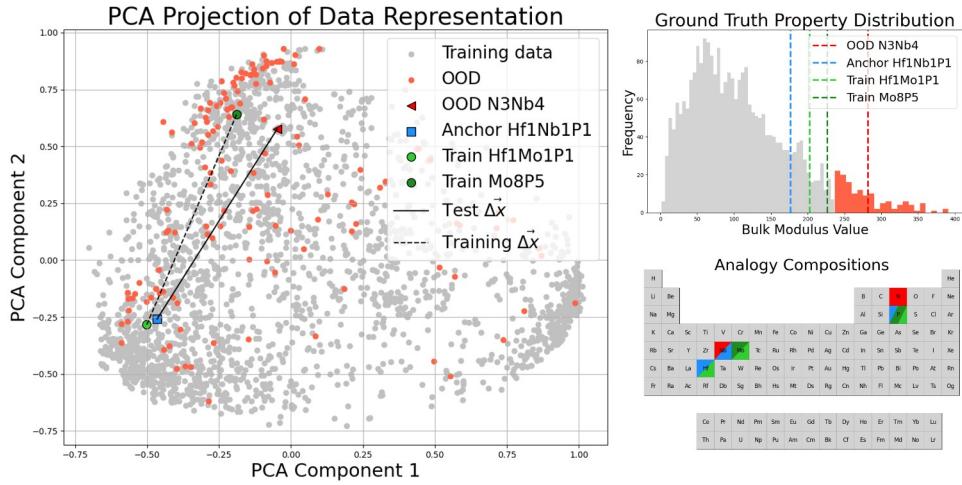
Dataset	Property	Ridge Reg. [15]	MODNet [13]	CrabNet [11]	Ours
AFLOW [25]	Band Gap [eV]	0.0	0.0	0.032	<b>0.132</b>
	Bulk Modulus [GPa]	0.124	0.007	0.199	<b>0.336</b>
	Debye Temperature [K]	0.182	0.0	0.0	<b>0.504</b>
	Shear Modulus [GPa]	0.058	0.088	0.044	<b>0.182</b>
	Thermal Conductivity [ $\frac{W}{mK}$ ]	0.088	0.0	0.0	<b>0.132</b>
Matbench [10]	Thermal Expansion [ $K^{-1}$ ]	0.154	0.154	0.132	<b>0.191</b>
	Band Gap [eV]	0.0	0.004	0.0	<b>0.009</b>
	Refractive Index	0.009	0.101	0.004	<b>0.122</b>
MP [26]	Yield Strength [MPa]	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>	<b>0.0</b>
	Bulk Modulus [GPa]	0.073	0.003	0.311	<b>0.498</b>
	Elastic Anisotropy	0.0	0.0	0.0	<b>0.006</b>
MoleculeNet [32]	Shear Modulus [GPa]	0.0	0.0	0.003	<b>0.084</b>
	ESOL [ $\frac{mol}{kJ}$ ]	<b>0.357</b>	0.0	0.196	0.268
	Freesolv [ $\frac{mol}{kJ}$ ]	0.062	0.0	0.0	<b>0.781</b>
	Lipophilicity [ $\log D$ ]	0.024	0.0	0.014	<b>0.057</b>
	BACE binding [IC50]	0.0	0.0	0.0	<b>0.013</b>
		Chemprop [29]	Random Forest [30]	MLP [31]	

### Molecules

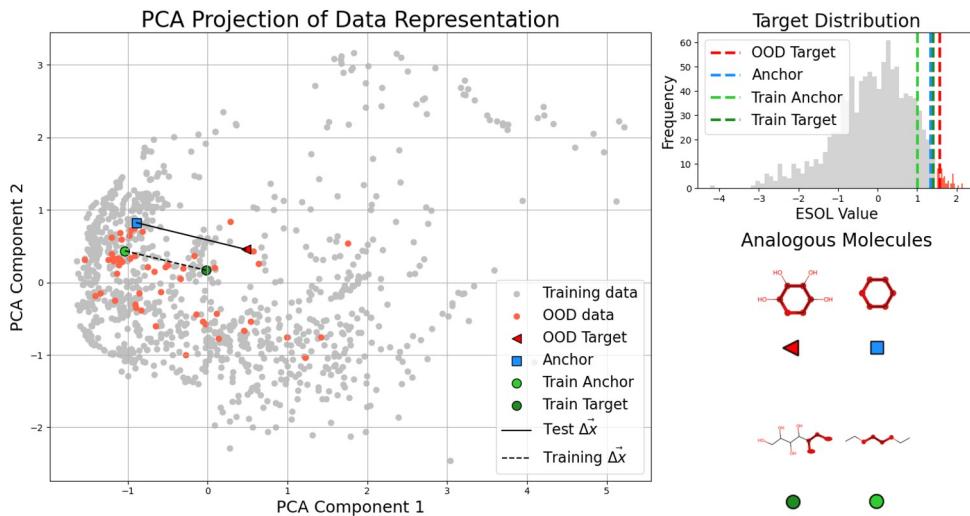
Figures 5 and 9 display additional analogical molecule pairs. Two distinct modes of similarity can be identified: one is between the training anchor and target, and between the test anchor and target Figures 9(a,b,c), and the other is between the anchors, and between the targets Figure 9(d,e,f). In the anchor selection process, the model can converge to an anchor that is either very similar to the ood target, in a way that two training samples are similar (in  $X$  space), or that is different from the ood target and similar to the training anchor and the training target will be similar to the ood. For both modes, we can spot analogous differences between the molecules.



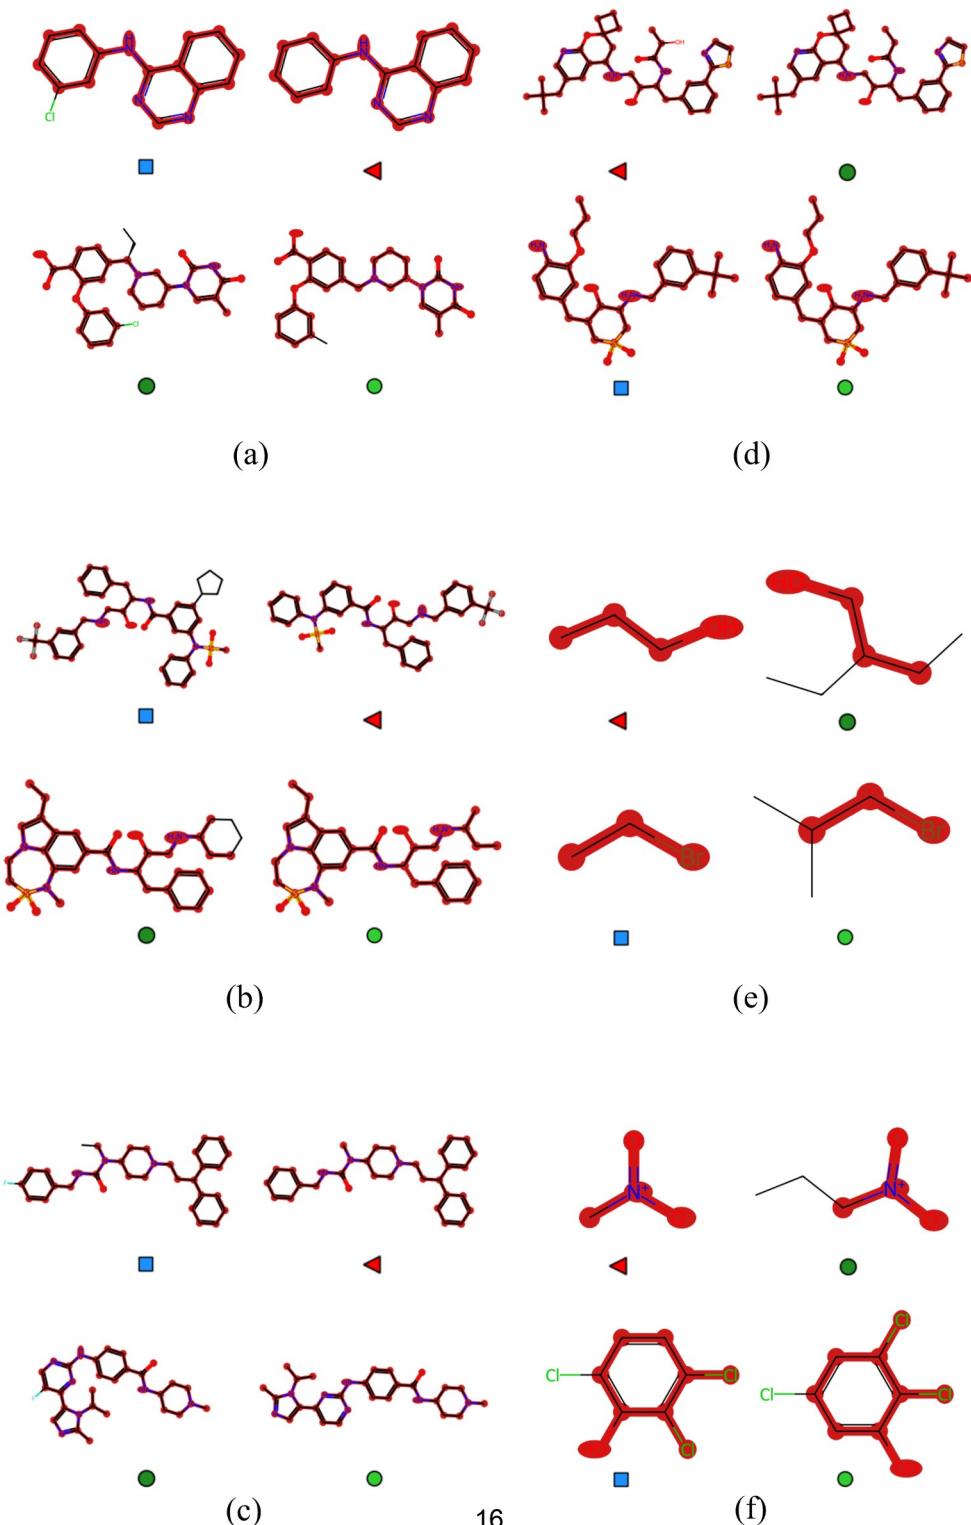
**Fig. 6: Visualizing Analogies in Bilinear Transduction.** AFLOW shear modulus ood predictions are based on in-distribution anchors, that paired with ood points, form analogies to training pairs. (a) PCA plot of all samples in the dataset. The difference between the ood point and its anchor is similar to the difference between the training point and anchor. (b) Ground truth shear modulus training (gray) and test (red) distributions and ood, anchor, and analogous training pair values. (c) Analogy compositional visualization. **anchor** and **ood** differ by one neighboring d-block element. So do training **anchor** and **target**.



**Fig. 7: Visualizing Analogies in Bilinear Transduction.** AFLOW bulk modulus ood predictions are based on in-distribution anchors, that paired with ood points, form analogies to training pairs. (a) PCA plot of all samples in the dataset. The difference between the ood point and its anchor is similar to the difference between the training point and anchor. (b) Ground truth shear modulus training (gray) and test (red) distributions and ood, anchor, and analogous training pair values. (c) test anchor and training anchor differ by one neighboring d-block element (Mo-Nb). ood and training target differ by the same elements, and in a group 15 element (N-P).



**Fig. 8: Visualizing Analogies in Bilinear Transduction.** MoleculeNet ESOL ood predictions are based on in-distribution anchors, that paired with ood points, form analogies to training pairs. (a) PCA plot of all samples in the dataset. The difference between the ood point and its anchor is similar to the difference between the training point and anchor. (b) Ground truth ESOL training (gray) and test (red) distributions and ood, anchor, and analogous training pair values. (c) Analogical molecules pairs. The similarity between **ood-anchor** and training **anchor** and **target** is highlighted in red, using the MCS metric.



**Fig. 9: Analogical Molecules pairs.** Molecules are paired according to maximum similarity. **(a-c) top** pair: **ood** are similar to their **anchors**, **bottom** pair: training **targets** are similar to their **anchors**. **(d-f) top** pair: **ood** are similar to training **targets**, **bottom** pair: **anchors** are similar to training **anchors**. For each pair, we denote the relevant benchmark and chemical operator differentiating samples within each pair. (a) lipophilicity, Cl addition. (b) BACE, addition or completion of a ring. (c) lipophilicity, F addition. (d) BACE, targets differ in OH functional group and anchors are the same. (e) ESOL, additional C. (f) ESOL, targets differ in additional C and anchors differ in functional group position in the ring.

## 6.2 Implementation Details

### 6.2.1 Data Representation

**Chemprop.** Chemprop derives a molecular graph representation based on SMILES, featurizes the atom nodes and bond edges using RDKit, then performs message-passing in order to learn an intermediate representation for property prediction via a feed-forward network [29, 33, 34].

**RF and MLP.** We derive RDKit descriptors from the SMILES representation using the deepchem library [33, 34, 36].

**Ridge Regression.** We adhere to the data preprocessing scheme outlined in Kauwe et al. [15] on AFLOW and featurize data using element-based Oliynyk descriptors [37]. The data is scaled using the StandardScaler and normalized using the Normalizer from the sklearn library based on the training data statistics.

**MODNet.** Following De Breuck et al. [13] on Matbench, we create element-based feature vectors using Matminer [38], and then select features based on the normalized mutual information [39]. The data is scaled using the MinMaxScaler based on the training data statistics and NaN values are replaced with a constant, -1.

**CrabNet.** Following Wang et al. [11] on MP, we leverage mat2vec, learned via self-supervised natural language processing techniques, trained on a large corpus of scientific literature.

**Bilinear Transduction.** For MoleculeNet, we derive RDKit descriptors from the SMILES representation using the deepchem library [33, 34, 36]. We use the Ridge Regression representation and processing for AFLOW and MP, and the MODNet representation and processing for Matbench.

### 6.2.2 Bilinear Transduction Hyperparameter Search

In Table 1, we report the best ood MAE score for Bilinear Transduction on AFLOW over a hyperparameter search on the number of predictor network layers (3, 4), layer size (256, 512, 1024) and embedding size (32, 42, 48, 64). The hyperparameter search revealed little sensitivity to changes in hyperparameter values, indicating the robustness of our evaluation.

## 6.3 Methods

### Bilinear Transduction

Bilinear Transduction [24] is a regression method for out-of-support (oos) generalization, a type of distribution shift where data becomes oos at test time. By reparameterizing the problem, oos generalization is possible via a bilinear predictor. Let  $X \subseteq \mathbb{R}^n$  denote a representation space of data points and let  $Y \subseteq \mathbb{R}$  denote the space of labels. Given a training set  $D^{\text{tr}} = \{(x, y)\}_n \subseteq X \times Y$ , Bilinear Transduction learns a transductive predictor that extrapolates in a zero-shot manner to oos test points. Naively, one can learn predictor  $h_\theta : X \rightarrow Y$ . However, ML methods often fail under covariate shift [40, 41]. Instead, Bilinear Transduction offers the following formulation.

Let  $\Delta X = \{x_{i-1} - x_i \mid x_i, x_{i-1} \in X\}$  denote the differences distribution. During training, predictor  $h_\theta : \Delta X \times X \rightarrow Y$  predicts value  $y_{i-2} \in D^{\text{tr}}_X$  for data point  $x_{i-2} \in D^{\text{tr}}_X$  given anchor point  $x_1 \in D^{\text{tr}}_X$  and the difference between them  $\Delta x_{21} = x_2 - x_1 \in D^{\text{tr}}_{\Delta X}$ . The predictor  $h_\theta(\Delta x, x) = f_\theta(\Delta x)g_\theta(x)$  is implemented as a bilinear function in non-linear embeddings of  $\Delta x$  and  $x$ .

The test set  $D^{\text{te}} \subseteq X$  includes oos data points. Given test point  $x_{\text{te}} \in D^{\text{te}}$ , its predicted value is  $h_\theta(\Delta x_{\text{te,an}}, x_{\text{an}})$ . Anchor  $x_{\text{an}} \in D^{\text{tr}}_X$  is the training point that minimizes the distance between its difference with the test point  $\Delta x_{\text{te,an}} = x_{\text{te}} - x_{\text{an}}$ , and differences within the training distribution  $D^{\text{tr}}_{\Delta X}$ . Formally,  $x_{\text{an}} = \operatorname{argmin}_{x_i \in D^{\text{tr}}} \{ \|\Delta x_{\text{te},i} - \Delta x\|_2 \mid \Delta x \in D^{\text{tr}}_{\Delta X} \}$ . This reparameterization converts the problem to within support, as  $\Delta x_{\text{te,an}}$  and  $x_{\text{an}}$  are within the training distribution. Under some coverage and rank assumptions Bilinear Transduction has theoretical convergence guarantees [24].

We adopt Bilinear Transduction for out setting where the output  $Y$  is ood rather than the input being oos (described in Section 3). Under our current formulation where  $Y$  is ood, the theoretical guarantees for Bilinear Transduction may not fully apply. Future work will focus on learning data representations that align with the assumptions of Bilinear Transduction for ood  $X$ .

## References

- [1] Axelrod, S., Schwalbe-Koda, D., Mohapatra, S., Damewood, J., Greenman, K.P., Gómez-Bombarelli, R.: Learning matter: Materials design with machine learning and atomistic simulations. *Accounts of Materials Research* **3**(3), 343–357 (2022)
- [2] Sanchez-Lengeling, B., Aspuru-Guzik, A.: Inverse molecular design using machine learning: Generative models for matter engineering. *Science* **361**(6400), 360–365 (2018)
- [3] Bilodeau, C., Jin, W., Jaakkola, T., Barzilay, R., Jensen, K.F.: Generative models for molecular discovery: Recent advances and challenges. *Wiley Interdisciplinary Reviews: Computational Molecular Science* **12**(5), 1608 (2022)
- [4] Noh, J., Gu, G.H., Kim, S., Jung, Y.: Machine-enabled inverse design of inorganic solid materials: promises and challenges. *Chemical Science* **11**(19), 4871–4881 (2020)
- [5] Kim, Y., Kim, Y., Yang, C., Park, K., Gu, G.X., Ryu, S.: Deep learning framework for material design space exploration using active transfer learning and data augmentation. *npj Computational Materials* **7**(1), 140 (2021)
- [6] Zeni, C., Pinsler, R., Zügner, D., Fowler, A., Horton, M., Fu, X., Shysheya, S., Crabbé, J., Sun, L., Smith, J., et al.: Mattergen: a generative model for inorganic materials design. *arXiv preprint arXiv:2312.03687* (2023)
- [7] Yang, M., Cho, K., Merchant, A., Abbeel, P., Schuurmans, D., Mordatch, I., Cubuk, E.D.: Scalable diffusion for materials generation. *arXiv preprint arXiv:2311.09235* (2023)
- [8] Xie, T., Fu, X., Ganea, O.-E., Barzilay, R., Jaakkola, T.: Crystal diffusion variational autoencoder for periodic material generation. *arXiv preprint arXiv:2110.06197* (2021)
- [9] Walters, W.P., Barzilay, R.: Applications of deep learning in molecule generation and molecular property prediction. *Accounts of chemical research* **54**(2), 263–270 (2020)
- [10] Dunn, A., Wang, Q., Ganose, A., Dopp, D., Jain, A.: Benchmarking materials property prediction methods: the matbench test set and automatminer reference algorithm. *npj Computational Materials* **6**(1), 138 (2020)
- [11] Wang, A.Y.-T., Kauwe, S.K., Murdock, R.J., Sparks, T.D.: Compositionally restricted attention-based network for materials property predictions. *Npj Computational Materials* **7**(1), 77 (2021)
- [12] Zhuo, Y., Mansouri Tehrani, A., Brgoch, J.: Predicting the band gaps of inorganic

solids by machine learning. *The journal of physical chemistry letters* **9**(7), 1668–1673 (2018)

[13] De Breuck, P.-P., Hautier, G., Rignanese, G.-M.: Materials property prediction for limited datasets enabled by feature selection and joint learning with modnet. *npj computational materials* **7**(1), 83 (2021)

[14] Ward, L., Agrawal, A., Choudhary, A., Wolverton, C.: A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Computational Materials* **2**(1), 1–7 (2016)

[15] Kauwe, S.K., Graser, J., Murdock, R., Sparks, T.D.: Can machine learning find extraordinary materials? *Computational Materials Science* **174**, 109498 (2020)

[16] Zhao, Z.-W., Del Cueto, M., Troisi, A.: Limitations of machine learning models when predicting compounds with completely new chemistries: possible improvements applied to the discovery of new non-fullerene acceptors. *Digital Discovery* **1**(3), 266–276 (2022)

[17] O mee, S.S., Fu, N., Dong, R., Hu, M., Hu, J.: Structure-based out-of-distribution (ood) materials property prediction: a benchmark study. *npj Computational Materials* **10**(1), 144 (2024)

[18] Li, K., Rubungo, A.N., Lei, X., Persaud, D., Choudhary, K., DeCost, B., Dieng, A.B., Hattrick-Simpers, J.: Probing out-of-distribution generalization in machine learning for materials. *arXiv preprint arXiv:2406.06489* (2024)

[19] Meredig, B., Antono, E., Church, C., Hutchinson, M., Ling, J., Paradiso, S., Blaiszik, B., Foster, I., Gibbons, B., Hattrick-Simpers, J., *et al.*: Can machine learning identify the next high-temperature superconductor? examining extrapolation performance for materials discovery. *Molecular Systems Design & Engineering* **3**(5), 819–825 (2018)

[20] Muckley, E.S., Saal, J.E., Meredig, B., Roper, C.S., Martin, J.H.: Interpretable models for extrapolation in scientific machine learning. *Digital Discovery* **2**(5), 1425–1435 (2023)

[21] Yang, H., Hu, C., Zhou, Y., Liu, X., Shi, Y., Li, J., Li, G., Chen, Z., Chen, S., Zeni, C., *et al.*: Mattersim: A deep learning atomistic model across elements, temperatures and pressures. *arXiv preprint arXiv:2405.04967* (2024)

[22] Merchant, A., Batzner, S., Schoenholz, S.S., Aykol, M., Cheon, G., Cubuk, E.D.: Scaling deep learning for materials discovery. *Nature* **624**(7990), 80–85 (2023)

[23] Batatia, I., Benner, P., Chiang, Y., Elena, A.M., Kovacs, D.P., Riebesell, J., Advincula, X.R., Asta, M., Baldwin, W.J., Bernstein, N., *et al.*: A foundation model for atomistic materials chemistry. *arXiv preprint arXiv:2401.00096* (2023)

[24] Netanyahu, A., Gupta, A., Simchowitz, M., Zhang, K., Agrawal, P.: Learning to extrapolate: A transductive approach. In: International Conference on Learning Representations (2023)

[25] Curtarolo, S., Setyawan, W., Hart, G.L., Jahnatek, M., Chepulskii, R.V., Taylor, R.H., Wang, S., Xue, J., Yang, K., Levy, O., *et al.*: Aflow: An automatic framework for high-throughput materials discovery. *Computational Materials Science* **58**, 218–226 (2012)

[26] Jain, A., Ong, S.P., Hautier, G., Chen, W., Richards, W.D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., *et al.*: Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials* **1**(1) (2013)

[27] G., C., S., B.: Mechanical properties of some steels. [https://citrination.com/datasets/153092/show\\_files/](https://citrination.com/datasets/153092/show_files/) (2017)

[28] Petousis, I., Mrdjenovich, D., Ballouz, E., Liu, M., Winston, D., Chen, W., Graf, T., Schladt, T.D., Persson, K.A., Prinz, F.B.: High-throughput screening of inorganic compounds for the discovery of novel dielectric and optical materials. *Scientific data* **4**(1), 1–12 (2017)

[29] Heid, E., Greenman, K.P., Chung, Y., Li, S.-C., Graff, D.E., Vermeire, F.H., Wu, H., Green, W.H., McGill, C.J.: Chemprop: a machine learning package for chemical property prediction. *Journal of Chemical Information and Modeling* **64**(1), 9–17 (2023)

[30] Breiman, L.: Random forests. *Machine learning* **45**, 5–32 (2001)

[31] Gardner, M.W., Dorling, S.: Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. *Atmospheric environment* **32**(14-15), 2627–2636 (1998)

[32] Ramsundar, B., Eastman, P., Walters, P., Pande, V., Leswing, K., Wu, Z.: Deep Learning for the Life Sciences. O'Reilly Media, ??? (2019)

[33] Weininger, D.: Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences* **28**(1), 31–36 (1988)

[34] RDKit: Open-source cheminformatics. <https://www.rdkit.org>

[35] Damewood, J., Karaguesian, J., Lunger, J.R., Tan, A.R., Xie, M., Peng, J., Gómez-Bombarelli, R.: Representations of materials for machine learning. *Annual Review of Materials Research* **53**(1), 399–426 (2023)

[36] Ramsundar, B.: Molecular machine learning with deepchem. PhD thesis, Stanford

University (2018)

- [37] Oliynyk, A.O., Antono, E., Sparks, T.D., Ghadbeigi, L., Gaulois, M.W., Meredig, B., Mar, A.: High-throughput machine-learning-driven synthesis of full-heusler compounds. *Chemistry of Materials* **28**(20), 7324–7331 (2016)
- [38] Ward, L., Dunn, A., Faghaninia, A., Zimmermann, N.E., Bajaj, S., Wang, Q., Montoya, J., Chen, J., Bystrom, K., Dylla, M., *et al.*: Matminer: An open source toolkit for materials data mining. *Computational Materials Science* **152**, 60–69 (2018)
- [39] Kraskov, A., Stögbauer, H., Grassberger, P.: Estimating mutual information. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* **69**(6), 066138 (2004)
- [40] Shimodaira, H.: Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference* **90**(2), 227–244 (2000)
- [41] Shimodaira, H., Sugiyama, M., Storkey, A., Gretton, A., David, S.-B., Quinone-roCandela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.: Dataset Shift in Machine Learning. *Neural Information Processing Series*, pp. 201–205. Yale University Press in association with the Museum of London, ??? (2009)