

# Multi-fidelity prediction of molecular optical peaks with deep learning

Kevin P. Greenman,<sup>†</sup> William H. Green,<sup>†</sup> and Rafael Gómez-Bombarelli<sup>\*,‡</sup>

<sup>†</sup>*Department of Chemical Engineering, MIT, Cambridge, MA*

<sup>‡</sup>*Department of Materials Science and Engineering, MIT, Cambridge, MA*

E-mail: rafagb@mit.edu

## Abstract

Optical properties are central to molecular design for many applications, including solar cells and biomedical imaging. A variety of *ab initio* and statistical methods have been developed for their prediction, each with a trade-off between accuracy, generality, and cost. Existing theoretical methods such as time-dependent density functional theory (TD-DFT) are generalizable across chemical space because of their robust physics-based foundations but still exhibit random and systematic errors with respect to experiment despite their high computational cost. Statistical methods can achieve high accuracy at a lower cost, but data sparsity and unoptimized molecule and solvent representations often limit their ability to generalize. Here, we utilize directed message passing neural networks (D-MPNNs) to represent both dye molecules and solvents for predictions of molecular absorption peaks in solution. Additionally, we demonstrate a multi-fidelity approach based on an auxiliary model trained on over 28,000 TD-DFT calculations that further improves accuracy and generalizability, as shown through rigorous splitting strategies. Combining several openly-available experimental datasets, we benchmark these methods against a state-of-the-art regression tree algorithm and

compare the D-MPNN solvent representation to several alternatives. Finally, we explore the interpretability of the learned representations using dimensionality reduction and evaluate the use of ensemble variance as an estimator of the epistemic uncertainty in our predictions of molecular peak absorption in solution. The prediction methods proposed herein can be integrated with active learning, generative modeling, and experimental workflows to enable the more rapid design of molecules with targeted optical properties.

## Introduction

Dye molecules are used in many applications ranging from sensitizers for solar cells to biomedical imaging and diagnostics.<sup>1,2</sup> The optical properties of dyes, namely their absorption and emission characteristics, must be known to determine their suitability for particular applications. Although numerous theoretical and statistical methods exist to predict these properties, many of these methods are not sufficiently accurate or general, or require significant computational cost, all of which hinder their application to large and diverse sets of molecules. Herein, we propose new deep learning methods that use learned dye and solvent representations and multi-fidelity data to improve prediction accuracy and generalizability on rigorous splits of several of the largest open-source datasets. Our models are publicly available for making predictions with corresponding uncertainty estimates.

Many theoretical methods have been developed for predicting molecular optical properties, including empirical tables, semi-empirical methods, time-dependent density functional theory (TD-DFT), and wavefunction-based methods.<sup>3,4</sup> TD-DFT has been the most widely used method for at least the past decade because of its favorable accuracy/cost trade-off and its capacity to be coupled with continuum solvents approximations,<sup>5</sup> and it has been benchmarked and reviewed extensively.<sup>6,7</sup> In parallel to theoretical methods, researchers have also developed surrogate statistical models that predict UV/Vis spectra from molecular structure at a lower computational cost than TD-DFT. ML studies for predicting properties

related to the electronically excited states of molecules have been reviewed recently.<sup>8,9</sup>

Limitations in previous statistical modeling efforts can be classified into three categories: data sparsity, molecular representations, and solvent representations.<sup>10</sup> Many studies have focused on a narrow part of chemical space (e.g. a single dye family) because of the limited availability of large UV/Vis datasets. This data sparsity has been addressed recently with the publication of several experimental datasets,<sup>1,11–18</sup> described in Table 1. There are also several large computed datasets of excitation energies available (Table 2). However, studies are still lacking on how the chemical diversity of the training data impacts model performance on new, unrelated chemical space. Many prediction methods have created molecular representations based on generic structure-based fingerprints or human-selected descriptor features. Most previous studies did not consider solvent effects, but many leveraged descriptors derived from quantum chemical calculations. The issue of data sparsity is related to the shortcomings of solvent representations in previous models; with relatively few examples of dyes measured in more than one solvent, it was sometimes easier to train a model only on data in the most commonly reported solvent to remove this complexity from the model.

**Table 1: Existing Datasets of Experimental UV/Vis Spectroscopic Properties.** The properties listed for each dataset are not necessarily present for every measurement. "Full" refers to the full absorption/emission spectrum as xy-coordinate pairs,  $\lambda_{max}$  is the peak wavelength,  $\epsilon_{max}$  is the peak molar attenuation coefficient (also called the molar extinction coefficient or molar absorptivity),  $\sigma$  is the peak FWHM (bandwidth),  $\Phi$  is the quantum yield, and  $\tau$  is the fluorescence lifetime. A subset of the data in the ChemFluor<sup>16</sup> set was extracted from the Fluorophores<sup>12</sup> set. The number of entries for the UV/Vis+ dataset includes the count of the dye entries only, and the entries for NIST do not include ions.

Dataset	Entries	Dye	Solvent	Absorption	Emission	Other
ChemDataExtractor <sup>14</sup>	8,467	SMILES	Name	$\lambda_{max}, \epsilon_{max}$	-	-
ChemFluor <sup>16</sup>	4,386	SMILES	Name	$\lambda_{max}$	$\lambda_{max}$	$\Phi$
Deep4Chem <sup>17</sup>	20,236	SMILES	SMILES	$\lambda_{max}, \sigma_{abs}, \epsilon_{max}$	$\lambda_{max}, \sigma_{emi}$	$\Phi, \tau$
DSSCDB <sup>1</sup>	5,178	SMILES	Name	$\lambda_{max}$	$\lambda_{max}$	-
Dye Aggregation <sup>15</sup>	4,043	SMILES	Name	$\lambda_{max}$	-	-
Fluorophores.org <sup>12</sup>	955	Name	Name	Full, $\lambda_{max}$	Full, $\lambda_{max}$	$\Phi, \tau$
NIST <sup>11</sup>	2,306	MOL file	-	Full	-	-
PhotochemCAD <sup>13</sup>	552	Name	Name	$\lambda_{max}, \epsilon_{max}$	-	$\Phi$
UV/Vis+ <sup>18</sup>	112	Name	Name	Full	Full	-

Table 2: **Existing Large Datasets of Computed Excitation Energies.** Each dataset contains additional properties beyond excitation energies, such as oscillator strengths, highest occupied molecular orbital (HOMO), and lowest unoccupied molecular orbital (LUMO), but the set of reported properties is different for each dataset. Some datasets report results from multiple levels of theory (ranging from semi-empirical to coupled cluster), but all were calculated in vacuum. Many smaller datasets and datasets containing only ground-state properties (e.g. HOMO and LUMO) exist that are not referenced here.

Dataset	Entries	Dye
QM7b <sup>19,20</sup>	7,211	Coulomb matrix
QM8 <sup>21,22</sup>	21,786	XYZ file
QM-symex <sup>23</sup>	172,736	XYZ file
PubChemQC <sup>24</sup>	3,411,649	MOL file

Among the previous studies on predicting absorption peak wavelengths or excitation energies, the work of Ju et al.,<sup>16</sup> Kang et al.,<sup>25</sup> and Joung et al.<sup>26</sup> is particularly noteworthy because of the size of their training datasets and the accuracies this enabled them to achieve. Ju et al. trained a gradient boosted regression tree (GBRT) algorithm on composite fingerprints to predict the maximum absorption and emission wavelengths and photoluminescence quantum yield (PLQY) using a large set of experimental data they compiled from the literature. Kang et al. trained a random forest algorithm on a subset of the PubChemQC database<sup>24</sup> to predict B3LYP/6-31G\* excitation energies and oscillator strengths in a vacuum from molecular fingerprints. Joung et al. used their previously compiled experimental dataset<sup>17</sup> to train a model that uses graph-convolutional networks (GCN) to predict multiple molecular optical properties, including the absorption and emission peak wavelengths.

Although these recent works achieved impressive accuracies, their reported performance may be more representative of how they would perform in substituent-selection applications as opposed to *de novo* design tasks with unseen chemistries. Recent reviews of ML best practices in chemistry and materials science have warned against data leakage from the same compound or composition being present under multiple measurement conditions.<sup>27,28</sup> Random splitting into training/test or training/validation/test sets based on dye-solvent pair may not be sufficient for assessing model generalizability on this task because test error may be spuriously low if a dye appears in both training and test sets in different solvents or

even if the molecules in the test set are too chemically similar to the training data. We set out to explore how these decisions impact model performance.

Additionally, previous work has not attempted to leverage a combination of computed and experimental data in predictions of optical spectra. This multi-fidelity approach is desirable because of the lower cost and greater availability of calculations with respect to experiments. Multi-fidelity methods have been demonstrated in several other applications for integrating data from multiple levels of theory or from theory and experiments.<sup>29-31</sup> Furthermore, theoretical methods do not have a domain of applicability constrained by a training set, so they are more reliable for making predictions on chemistries with low similarities to existing data. These factors suggest that a multi-fidelity approach may improve model accuracy and generalizability on experimental predictions and may be more useful in active learning.

In this work, we leverage recently-compiled experimental datasets and directed message passing neural networks (D-MPNN)<sup>32</sup> to address previous limitations with molecular and solvent representations. The D-MPNN approach learns representations for the dye and solvent that are optimized for predicting absorption properties. We compare our optimized representations to a state-of-the-art fingerprint-based method and to alternative solvent descriptors. Our method produces interpretable representations and estimates the uncertainty in its predictions. We emphasize the importance of using rigorous splitting techniques for assessing the ability of a model to generalize to unseen chemistries, and we show that incorporating results from physics-based calculations into model training improves performance across several large datasets. The predictive capability of our techniques will enable more rapid design of molecules with target optical properties.

# Methods

## Data Sources and Preprocessing

We compiled experimental UV/Vis absorption data from several of the largest openly-available datasets: Deep4Chem,<sup>17</sup> ChemFluor,<sup>16</sup> Dye Aggregation (DyeAgg),<sup>15</sup> ChemDataExtractor (CDEx),<sup>14</sup> and the Dye-Sensitized Solar Cell Database (DSSCDB).<sup>1</sup> Among the datasets listed in Table 1, these were the largest and most easily machine-readable, and they included solvent information for each measurement. All of these data sources reported the dyes in the form of SMILES,<sup>33</sup> but only the Deep4Chem set reported the solvents in this form. For the other four sets, we converted the solvent names and abbreviations to SMILES through a manually-constructed dictionary mapping because automatic tools did not recognize the necessary variety of names and abbreviations for many solvents. We extracted all measurements that included a valid dye SMILES string, solvent SMILES string, and peak wavelength of maximum absorption. We determined the validity of the SMILES strings using RDKit<sup>34</sup> and dropped measurements with invalid dye or solvent SMILES (105 measurements) and those containing "." to represent clusters of molecules (373 measurements). The remaining dataset contained 28,734 measurements. We removed any dye-solvent pairs with duplicate measurements within the same dataset that disagreed by more than 5 nm. For those that agreed within 5 nm, the mean of the values was used. This resulted in datasets of size 1,825 (CDEx), 3,840 (ChemFluor), 14,771 (Deep4Chem), 1,720 (DSSCDB), 3,025 (DyeAgg), and 24,580 (a combined set of ChemFluor, Deep4Chem, DSSCDB, and DyeAgg).

## TD-DFT Calculations

For each dye molecule in the combined experimental dataset, as well as a set of molecules with dye-like substructures parsed from USPTO patents and commercial vendors, initial geometries were generated using RDKit to convert the SMILES strings into Cartesian coordinates.<sup>35</sup> These geometries were refined using semi-empirical tight-binding density functional

theory (GFN2-xTB)<sup>36</sup> in the ORCA software,<sup>37</sup> followed by geometry optimizations at the BP86<sup>38</sup>-D3<sup>39</sup>/def2-SVP<sup>40</sup> level of theory. Finally, TD-DFT calculations were performed with the Tamm-Dancoff approximation (TDA)<sup>41</sup> at the  $\omega$ B97X-D3<sup>42</sup>/def2-SVPD level of theory in Orca. This pipeline was completed for 28,772 molecules, of which 10,409 had a corresponding experimental measurement in at least one solvent from one of the aforementioned datasets. The total number of experiments with a corresponding TD-DFT calculation in vacuum was 19,409 (including measurements of the same molecule taken in more than one solvent).

For a subset of the complete dataset (only ChemFluor, DyeAgg, CDEx, and DSSCDB), we began with the optimized geometry calculated with BP86-D3/def2-SVP in ORCA and performed an additional TD-DFT calculation at the  $\omega$ B97XD/def2-SVP level with solvent corrections in Gaussian.<sup>43</sup> The solvent calculations were done using the integral equation formalism polarizable continuum model (IEFPCM) and Gaussian defaults for excited state solvation. This pipeline was completed for 6,707 dye-solvent pairs.

We extracted the peak vertical excitation energy from each of these calculations according to the following procedure: (1) if none of the energies were in the range of 1 - 4 eV, choose the lowest energy; (2) if only one energy is in the visible range, choose that one; (3) if multiple peaks are in the visible range, choose the one with the highest oscillator strength. While the vertical excitation energy is not exactly analogous to  $\lambda_{max,abs}$  because it does not account for nuclear vibronic effects, it is a relatively cheap computational surrogate that should improve the capability of a model to predict  $\lambda_{max,abs}$ .

## Dye and Solvent Representations

We compared three representation methods for the dye molecules and four for solvents. Two of the dye representations were derived from the open-source Chemprop D-MPNN framework,<sup>32</sup> and we compared these representations to the ChemFluor Functionalized Structure Descriptor (FSD) representation developed by Ju et al.<sup>16</sup> The FSD representation is a com-

posite fingerprint created by concatenating the E-state, CDK extended, and substructure presence and count fingerprints calculated by the PaDEL software<sup>44</sup> through PaDELPy.<sup>45</sup> Ju et al. found FSD to be superior in an extensive benchmark against other fixed fingerprint representations for predicting molecular absorption and emission peak energies. One Chemprop representation uses the D-MPNN framework "as is" to create a fingerprint embedding that is optimized for predicting absorption peak energies. The second Chemprop representation (which we call ChempropMultiFidelity) is similar to the first, but it uses a second Chemprop model trained on TD-DFT results to predict the TD-DFT peak vertical excitation energy and concatenates this predicted value onto the first Chemprop fingerprint embedding.

The four solvent representations compared herein are Morgan fingerprints, ChemFluor Comprehensive General Solvent Descriptors (CGSD), Minnesota solvent descriptors,<sup>46</sup> and Chemprop D-MPNN embeddings (SolventMPNN). We calculated the Morgan fingerprints with a radius of 4 and 256 bits. The five CGSD descriptors (developed by Ju et al.<sup>16</sup> in conjunction with the FSD dye representation) were extracted from the work of Reichardt<sup>47</sup> and Catalán<sup>48</sup> and represent the polarity ( $E_T(30)$ ), acidity (SA), basicity (SB), dipolarity (SdP), and polarizability (SP) of a solvent. We also matched solvents with their seven corresponding descriptors from the Minnesota Solvent Descriptor Database:<sup>46</sup> index of refraction ( $n$ ), Abraham's H-bond acidity ( $\alpha$ ), Abraham's H-bond basicity ( $\beta$ ), surface tension ( $\gamma$ ), dielectric constant ( $\epsilon$ ), aromaticity ( $\phi$ ), and electronegative halogenicity ( $\psi$ ). The solvent D-MPNN embeddings were optimized using a separate D-MPNN alongside that of the dye; this approach was previously shown to be successful in predicting solvation free energies.<sup>49</sup>

All dye-solvent pairs for which any of the above features could not be calculated were dropped from the dataset.

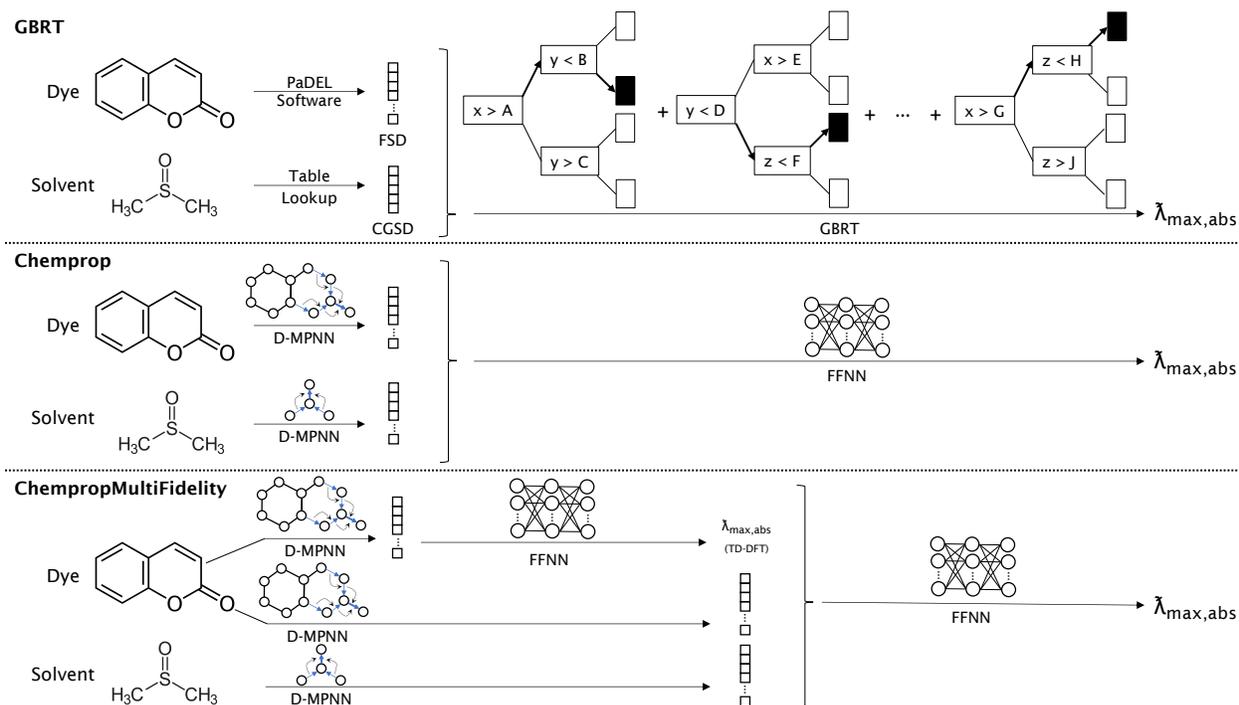
## Models

We compare the performance of three types of models, each corresponding to one of the dye representations described above. The FSD representation was used with the gradient boosted regression tree (GBRT) algorithm<sup>50</sup> as implemented by Ju et al.<sup>16</sup> For both D-MPNN representations described above, the resulting fingerprint embedding is passed to a feed-forward neural network (FFNN) to accomplish the regression task. The three types of models are illustrated in Figure 1. The ChempropMultiFidelity model is a variation of the hybrid physics-ML models reviewed by Jia et al.<sup>51</sup>

D-MPNN hyperparameters (including hidden sizes, numbers of layers, dropout, batch sizes, learning rates, and warm-up epochs) were tuned using SigOpt.<sup>52</sup> The GBRT used the hyperparameters reported by Ju et al., while all D-MPNN models used hyperparameters that were tuned on the same ChemFluor dataset used by Ju et al. The details of the model architectures, training, and predictions are given in the Supporting Information.

## Train-Validation-Test Splits

The type of splitting strategy used during the development of machine learning models is a crucial consideration when evaluating the accuracy and generalizability of a model.<sup>27,28</sup> We compare three splitting strategies to illustrate this principle and to encourage the use of rigorous splitting strategies in subsequent work. In our regression task, the dye molecule and solvent molecule are both inputs to predict the peak wavelength of maximum absorption. The most naive splitting strategy, therefore, is to split randomly by dye-solvent pairs. If there are no duplicate measurements in a dataset, this splitting strategy makes it trivial to ensure that no pair is present in more than one of the training, validation, and test sets. Although the solvent effect can sometimes cause a substantial shift in the peak wavelength, the peaks measurements of the same dye in different solvents will be correlated. In other words, knowing the peak absorption wavelength of a particular dye in one solvent will likely improve the predictions of that same dye in a different solvent. This suggests a more rigorous



**Figure 1: Model Architectures for Predicting Experimental Absorption Peak in Solvent.** GBRT uses the PaDEL software and a table lookup to arrive at the FSD and CGSD, which are concatenated and used as input for a gradient boosted regression tree algorithm. Chemprop uses separate D-MPNN networks to obtain fingerprint embeddings for the dye and solvent, then concatenates these for input to a FFNN. The ChempropMultiFidelity model is similar to Chemprop, except with the addition of a secondary D-MPNN and FFNN that are trained to predict the absorption peak from TD-DFT data. This value is concatenated with the two fingerprint embeddings before being passed to a FFNN. There is no weight sharing between any of the D-MPNN networks. The Chemprop and ChempropMultiFidelity models can alternatively use Morgan fingerprints, Minnesota descriptors, or CGSD to represent the solvent in place of a D-MPNN; GBRT can use Morgan fingerprints or Minnesota descriptors as alternatives to CGSD.

splitting strategy where measurements are split by dye molecules rather than by dye-solvent pairs. This method ensures that any given dye molecule is restricted to either the training, validation, or test set, regardless of how many different solvent measurements are available. The final and most rigorous strategy discussed in this work is a scaffold split using the Bemis-Murcko scaffold<sup>53</sup> implemented in RDKit through Chemprop. Scaffold splits ensure that any dye molecules that possess the same scaffold are restricted to a single set, which makes the regression task more challenging and provides a better evaluation of model generalizability. This splitting strategy is most reflective of performance on *de novo* design tasks with unseen chemistries. We used 80-10-10 training-validation-test proportions for all splits.

## Results and Discussion

We performed our analysis on a combination of five data sources, which comprised a total of 28,734 measurements (of which 26,623 were unique dye-solvent pairs). Of these 28,734, there were 1,870 included in more than one data source. The combined dataset contained 15,157 unique dyes and 364 unique solvents. Ten of the solvents were used in 1,000 or more measurements. The breakdown of our data by source and by solvent is represented in Tables 3 and S1.

Table 3: **Dataset Composition by Data Source.** The numbers for dye-solvent pairs correspond to the number of measurements after filtering. These numbers do not account for the aggregation of duplicate measurements either within or across datasets. There were 1,870 measurements present in more than one data source.

<b>Dataset</b>	<b>Measurements</b>
Deep4Chem	16,585
ChemFluor	4,170
DyeAgg	3,626
CDEx	1,915
DSSCDB	2,438
<b>Total</b>	<b>28,734</b>

The maximum absorption wavelengths of the dyes represented in our dataset cover the en-

tire visible spectrum and extend into the ultraviolet and near-infrared regions. The molecules in the dataset include many of the common dye substructures and families. Figure 2 illustrates the distribution of wavelengths and the prevalence of each substructure. The peak wavelength distributions of the individual datasets are different from that of the combined dataset and from each other, as shown in Figure S1.

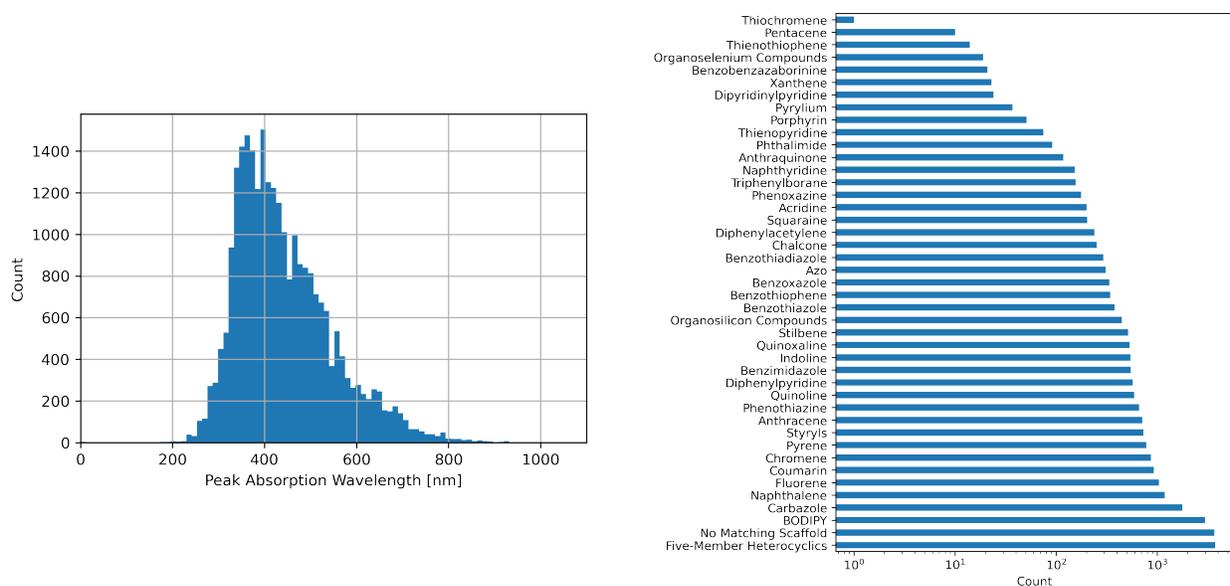


Figure 2: **Dataset Composition: Peak Location and Common Scaffold Matches.** (Left) The experimental peak wavelengths of maximum absorption from the combined dataset span the entire visible spectrum and extend into the infrared and ultraviolet, (Right) The combined dataset covers a wide variety of common dye scaffolds/families, as determined by SMARTS pattern matching.

The five datasets differ in their coverage of chemical dye space, as shown in Figure 3. The largest single data source (Deep4Chem) also has the most dense coverage of the chemical space. The smaller data sources, while covering a relatively large area of space, display more outliers that have few or zero close neighbors.

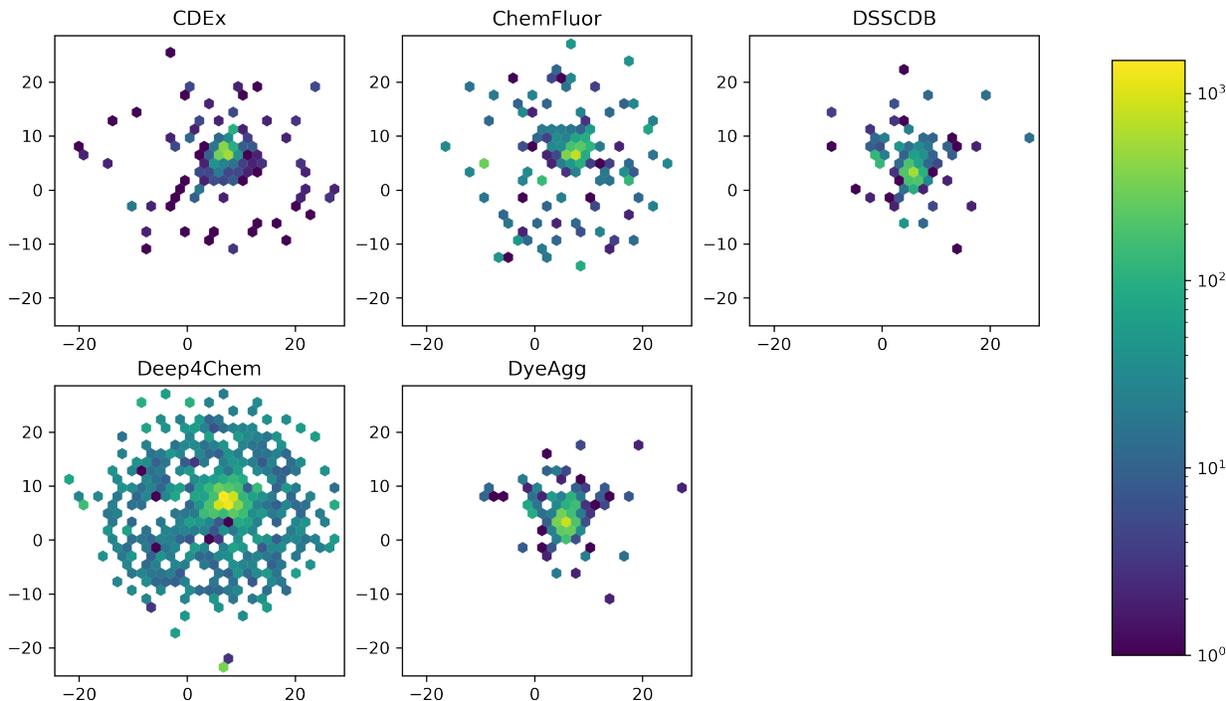


Figure 3: **UMAP of Morgan Fingerprints by Data Source.** A UMAP dimensionality reduction on the Morgan fingerprints shows a difference in the coverage and density of each dataset in chemical space.

We used TD-DFT to calculate the vertical excitation energies for 10,947 molecules in vacuum and 6,707 molecules with solvent corrections corresponding to the solvent measurements available in the experimental dataset. Since some molecules were measured in multiple solvents experimentally, the total number of experimental measurements with a corresponding vacuum TD-DFT calculation was 19,409. The results of these calculations are compared to experiments in Figure 4. Solvent corrections applied to TD-DFT yield results that have a smaller error than vacuum calculations with respect to the experimental ground truth. However, after fitting a linear regression to both sets of calculations, the error for the vacuum calculations is lower. Therefore, this systematic error of the vacuum calculations makes it suitable to use the results of the vacuum calculations as features for building our models.

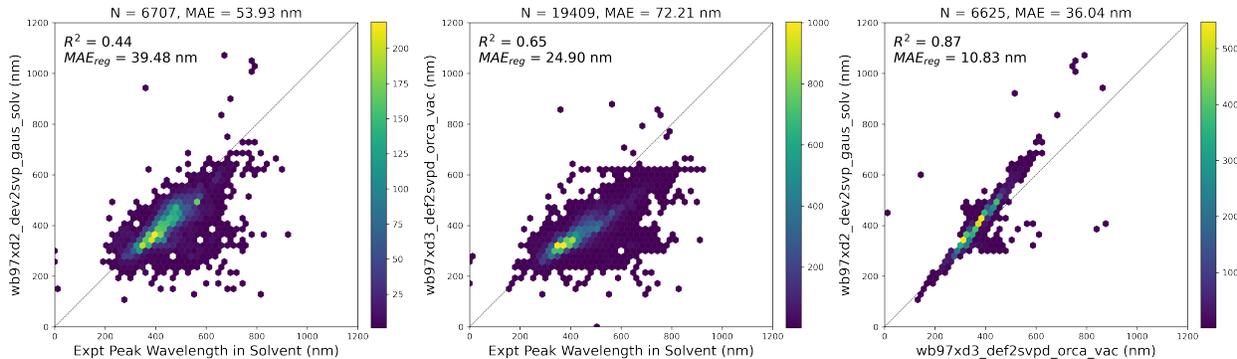


Figure 4: **TD-DFT Calculations in Vacuum and Solvent vs. Experiments.** (Left) Vertical excitation energy with maximum oscillator strength from solvent-corrected TD-DFT versus peak wavelength of maximum absorption from experiment, (Center) Vertical excitation energy with maximum oscillator strength from vacuum TD-DFT versus peak wavelength of maximum absorption from experiment, (Right) Vertical excitation energy with maximum oscillator strength from solvent-corrected TD-DFT versus Vertical excitation energy with maximum oscillator strength from vacuum TD-DFT. In each plot,  $MAE_{reg}$  refers to the adjusted MAE value after performing a simple linear regression. The linear regression equations for each plot above are as follows: (Left)  $\lambda_{expt,solv} = 1.69\lambda_{tddft,solv} - 238.24$ , (Center)  $\lambda_{expt,solv} = 1.82\lambda_{tddft,vac} - 226.82$ , and (Right)  $\lambda_{tddft,solv} = 1.18\lambda_{tddft,vac} - 32.86$ .

Although the combination of vacuum TD-DFT and a simple linear model performs well with respect to the experimental ground truth (MAE = 24.9 nm), the computational cost of TD-DFT may limit its applicability on large datasets. Nevertheless, this sets a baseline for the accuracy of computational methods in general when predicting this property.

We trained a D-MPNN and FFNN on 80-10-10 random splits of 28,772 vertical excitation energies from our full set of vacuum TD-DFT calculations. This model achieved a test MAE of 0.12 eV (14.99 nm), and the predictions are shown in Figure S9. This model became the auxiliary model used in the ChempropMultiFidelity approach for the remainder of this work.

The three ML methods we benchmarked are all able to match or exceed the accuracy of the linear regression on the vacuum TD-DFT result, but their ability to do so is heavily influenced by the strategy used to split the data into training, validation, and test sets (Figure 5). All three methods achieve a MAE under 10 nm when splitting by dye-solvent pair. For GBRT, this is similar to what Ju et al. observed in their predictions on the Chem-Fluor dataset (test set MAE of 10.46 nm). Performance worsens to MAEs of 13-21 nm when

splitting by dye molecule and to 18-27 nm when splitting by dye scaffold. ChempropMultiFidelity achieved a test RMSE of 27.47 nm using the SolventMPNN representation and scaffold splits. This outperforms the graph-convolutional networks (GCN) approach of Joung et al., who reported a test RMSE of 31.6 nm on random splits. Across all three splitting strategies, ChempropMultiFidelity and Chemprop perform better than GBRT. ChempropMultiFidelity performs better than Chemprop on the two more rigorous splitting strategies. This indicates that the Chemprop and ChempropMultiFidelity methods have generalizability superior to that of GBRT.

Ju et al. report that when they partitioned their data into training and test sets based on dye molecules rather than dye-solvent pairs, their test MAE in emission peak wavelength increased only slightly from 14.09 nm to 15.25 nm, but they did not report similar numbers for absorption predictions. Joung et al. do not report results for splitting by dye molecules, and neither report results when splitting by dye scaffold. Splitting more rigorously results in a wider error and ensemble variance distributions (Figure S6), but this is a better assessment of the ability of the model to generalize and is thus more reflective of performance for *de novo* design tasks. When using the random splitting strategy, dye molecules that appear in both the train and test set in different solvents have narrower error and ensemble variance distributions than those that only appear in the test set (Figure S7). As shown in Figure S8, we can also compare the splitting strategies by calculating the similarity (based on Morgan fingerprints or latent space coordinate) of each molecule in the test set to its nearest neighbor in the training set and plotting the test set error as a function of this similarity. When the similarity scores (which range from 0 to 1) are grouped into bins of size 0.1, this illustrates the error distributions as a function of similarity. The maximum prediction error should be lowest for the bin of test molecules that are most similar to the training set, but this was not true for the random splitting strategy. This indicates that the model may be relying too much on the training data and failing to learn the solvent effect. In contrast, the more rigorous splitting strategies exhibited the expected behavior.

The improved generalizability of the two D-MPNN approaches over the GBRT method may be the result of the automatically-learned dye representations. ChempropMultiFidelity outperforms Chemprop because in addition to this automatically-learned representation, it also incorporates additional physical knowledge through the inclusion of a predicted TD-DFT value in the learned embedding. It should be emphasized that this predicted TD-DFT value comes from an additional Chemprop model rather than an actual TD-DFT calculation, so there is no need to perform an additional calculation to predict on an unseen dye molecule. The ChempropMultiFidelity method achieves a MAE of 18.3 nm on scaffold splits, an improvement over the TD-DFT plus linear regression approach at a fraction of the cost.

We compared our ChempropMultiFidelity approach of incorporating a TD-DFT feature predicted by an auxiliary model to using true TD-DFT values. As shown in Figure S10, the detrimental effect of the noise introduced using the D-MPNN TD-DFT feature may be outweighed by the dramatic savings in computational cost and time from no longer needing to perform anew TD-DFT calculation for an unseen molecule.

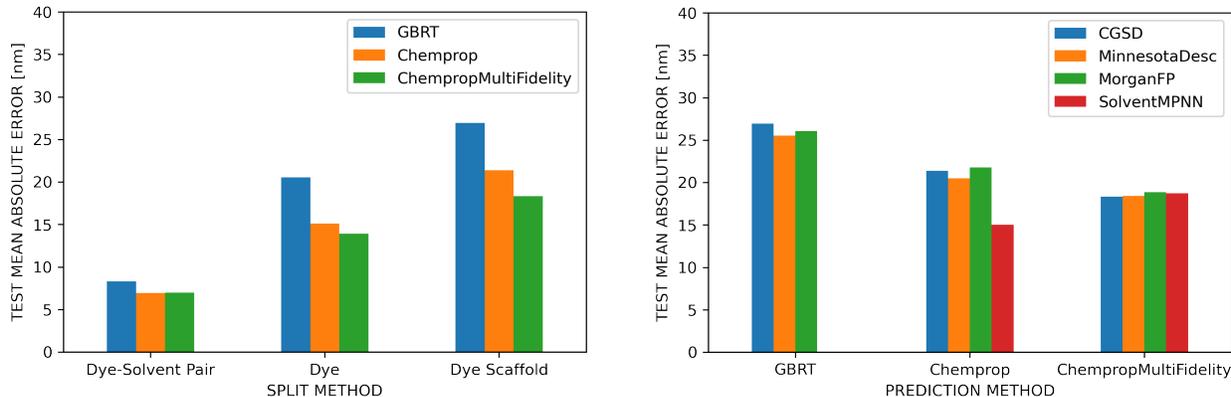


Figure 5: **Dye and Solvent Representations.** (Left) Comparison of three different molecular representations (GBRT, Chemprop, and ChempropMultiFidelity) across different methods for splitting into training, validation, and test sets on the Deep4Chem dataset. The CGSD is used to represent the solvent for all molecule representations. The ChempropMultiFidelity and Chemprop methods perform better than GBRT on all split types, and the difference in performance is more pronounced for the more rigorous split types. ChempropMultiFidelity performs best on the two more rigorous splits. (Right) Comparison of four different solvent representations (CGSD, Minnesota Solvent Descriptors, Morgan Fingerprint, and SolventMPNN) using scaffold splits of the Deep4Chem dataset. Regardless of which molecular representation is used, the CGSD and SolventMPNN representations outperform the others, albeit only slightly.

In our comparison of four different solvent representations, we found that none substantially and consistently outperform the others. It is necessary to represent the solvent in some way to achieve good predictions, but the CGSD, Minnesota descriptors, Morgan fingerprints, and SolventMPNN approaches all achieve similar results. The Morgan fingerprint and SolventMPNN approaches do, however, have the advantage that they are computable for any solvent since they are not restricted to look-up tables as are the CGSD and Minnesota methods.

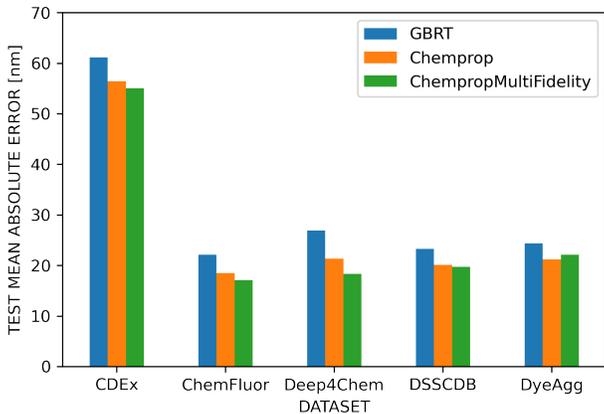


Figure 6: **Performance of Dye Representations Across Datasets.** Performance of three different molecular representations (GBRT, Chemprop, ChempropMultiFidelity) on several large, public datasets using scaffold splits and the CGSD solvent representation. Chemprop outperforms GBRT across all datasets. ChempropMultiFidelity is the best performer on all datasets except CDEx, for which all methods show substantially worse performance compared to the other datasets.

We used several additional datasets to further evaluate the performance of the three dye representations. ChempropMultiFidelity and Chemprop outperform GBRT on the ChemFluor, DSSCDB, and DyeAgg datasets and achieve MAEs of 17-23 nm, as shown in Figure 6. All models perform substantially worse (55 - 62 nm) on the CDEx dataset, and the inclusion of the TD-DFT feature degrades performance. This may be a result of errors introduced by the automatic extraction method used to construct this dataset.

We explored the effect of combining datasets together and performed 5-fold cross-validation to draw more rigorous conclusions. After observing the exceptionally poor performance of all models on the CDEx dataset, we excluded it from the combined dataset. The GBRT method was excluded from this analysis to compare the Chemprop and ChempropMultiFidelity models using the D-MPNN solvent representation (which cannot be integrated into the GBRT method). ChempropMultiFidelity achieved a mean MAE, RMSE, and  $R^2$  of 27.78 nm, 47.13 nm, and 0.80, respectively. These scores outperformed Chemprop in all metrics, but there was overlap in the standard errors of all three scores. The complete results are shown in Table 4. The prediction errors on this larger, combined dataset are larger than

those on the smaller datasets because of the different coverage of chemical space represented within each dataset. This is illustrated in Figures S11 - S15, which show the results of training and predicting on different datasets. These results indicate that the ChemFluor and Deep4Chem datasets are relatively similar to one another, as are the DSSCDB and DyeAgg. Combining all four datasets together results in an inhomogeneous chemical space and thus lowers performance.

Table 4: **Comparison of Dye Representations on a Combined Dataset.** 5-fold cross-validation on a dataset comprised of the union of the ChemFluor, Deep4Chem, DSSCDB, and DyeAgg dataset using scaffold splits and the D-MPNN solvent representation. The values represent the mean of the five cross-validation folds, while the error bars indicate the standard error.

Model	MAE (nm)	RMSE (nm)	R <sup>2</sup>
Chemprop	30.23 ± 5.62	52.08 ± 12.19	0.75 ± 0.08
ChempropMultiFidelity	27.78 ± 5.07	47.13 ± 11.10	0.80 ± 0.07

Having demonstrated the effectiveness of the D-MPNN models for modeling the peak wavelength of maximum absorption, we used dimensionality reduction to examine the interpretability of these models. We extracted the fingerprint embeddings from one of the Chemprop models and applied the t-SNE algorithm to reduce them to two dimensions. When these plots are colored by the experimental peak absorption wavelength and by the dye scaffolds present in the dataset, some patterns and clustering emerge, as shown in Figure 7. The number of dye scaffolds that are known to absorb light at lower visible (green - violet) and ultraviolet wavelengths is much greater than those known to absorb in the red - yellow range of the spectrum. This is apparent in the t-SNE plots, as the green points clustered in the lower right corner of the scaffold plot (corresponding to the BODIPY dye family) comprise nearly all of the red - yellow points on the wavelength plot.

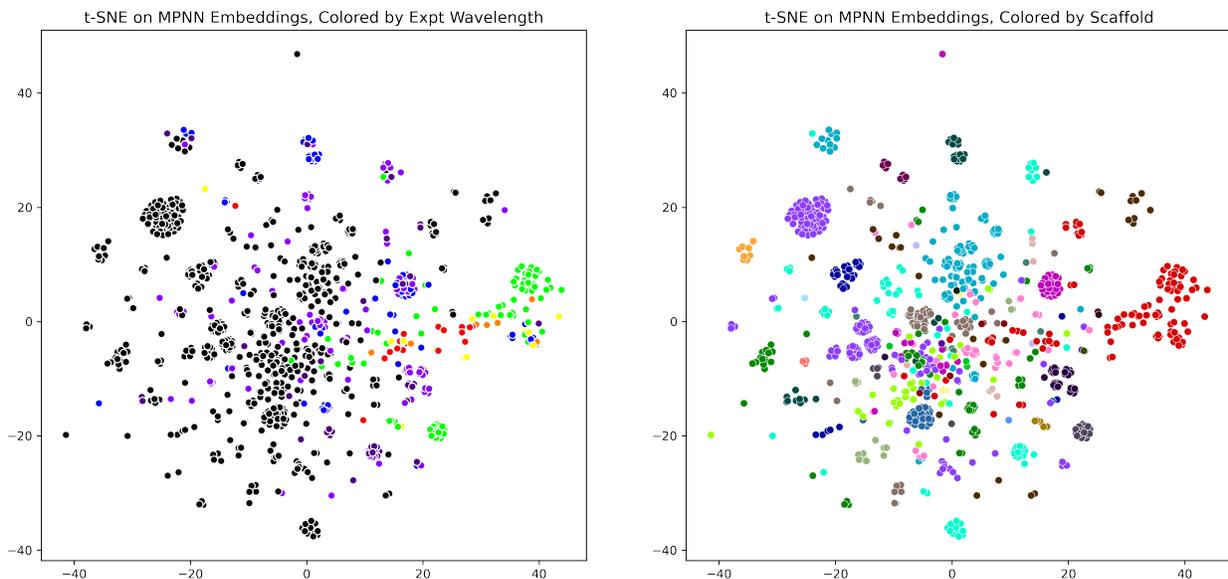


Figure 7: **MPNN Embedding Interpretability through t-SNE.** (Left) t-SNE plot of molecule D-MPNN embeddings from the scaffold-split test set of the Deep4Chem dataset, using the Chemprop molecule representation and SolventMPNN solvent representation, colored by the experimental peak wavelength of maximum absorption. Colors outside the visible spectrum are shown as black. (Right) Same as (Left), but colored by dye family scaffold.

While the accuracy of a model is very important, the ability to quantify uncertainty in predictions can greatly increase model utility. The level of confidence in a prediction or set of predictions can be used to motivate the selection of candidates that are most likely to succeed in experimental validation or to inform the choice of new measurements that will improve the model through active learning techniques. There are a plethora of methods available for estimating the uncertainty in NN models, and reviews of these methods have not found one method that consistently performs others across datasets and evaluation metrics.<sup>54–56</sup> However, one approach that is often used because of its ease of implementation is the ensemble variance as a measure of epistemic uncertainty. We evaluated the effectiveness of this method for quantifying the uncertainty in our Chemprop model. The parity plot in Figure 8 shows that the variance of an ensemble of five models is indeed high for many of the predictions that fall far away from the parity line. However, closer examination by plotting the square root of the ensemble variance versus the absolute prediction error gives a much

more sobering view of the effectiveness of this uncertainty metric. In fact, the Spearman rank correlation for this set of predictions and uncertainties is only 0.52, suggesting that one should not necessarily consider the rank order of the prediction uncertainties to be a good approximation of the rank ordering of the prediction errors.

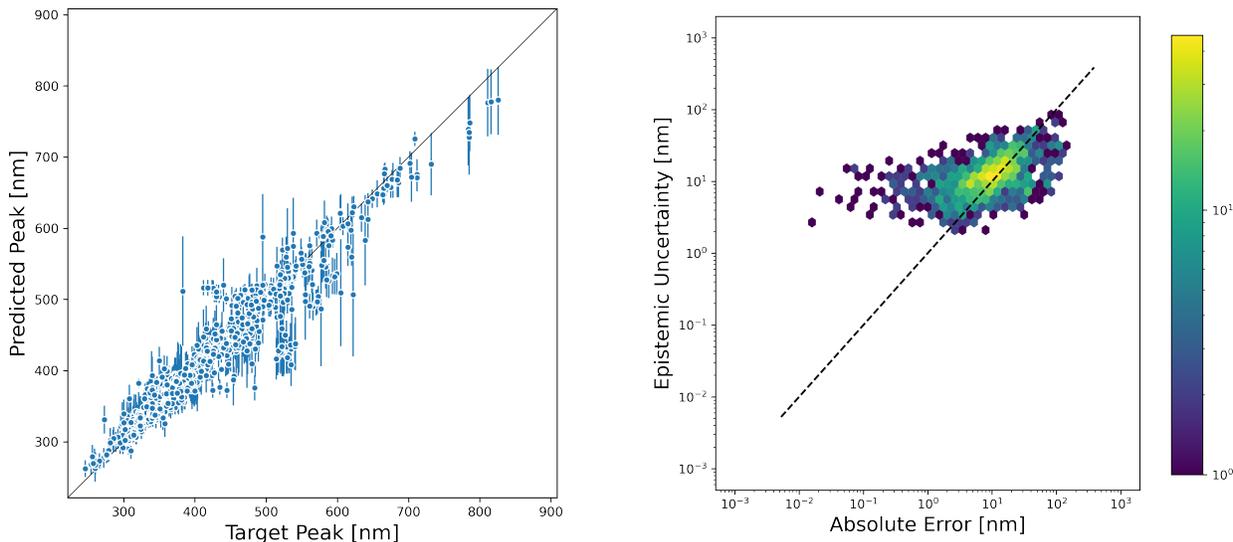


Figure 8: **Uncertainty in D-MPNN Models.** (Left) The epistemic uncertainty in test set predictions of an D-MPNN model estimated from the ensemble variance using an ensemble of five models. Predictions are for scaffold splits of the Deep4Chem dataset using the Chemprop molecule representation and SolventMPNN solvent representation. Error bars represent the square root of the ensemble variance for consistency in units. Many poorly-predicted points have a high ensemble variance. (Right) Epistemic uncertainty compared to absolute error between prediction and experiment. The Spearman rank correlation is 0.52.

## Conclusions

We have leveraged several recently-published datasets to benchmark models in their prediction of the peak wavelength of maximum absorption for dye molecules. Our results showed that D-MPNN models outperformed the best known fixed-fingerprint regression tree method, and the performance gain was more pronounced when we used more rigorous splitting strategies to evaluate the generalizability of the models to unseen chemistries. We also developed a multi-fidelity method for incorporating data from TD-DFT calculations to improve the

accuracy of experimental predictions. Vertical excitation energies from gas-phase TD-DFT calculations have a good linear correlation with the experimental peak positions (MAE = 24 nm) of dyes (measured in various solvents). TD-DFT PCM calculations in the solvents do not correlate as consistently with the experimental data.

Our best method (ChempropMultiFidelity) included a model trained on the results of previous TD-DFT calculations, and we used the predictions of this model as inputs to a second model that accounted for the solvent to predict the experimental peak wavelength. This multi-fidelity approach improved the model generalizability and improved performance for more rigorous splitting strategies. Our best model achieved a MAE of less than 7 nm on a held-out test set from a random split of dye-solvent pairs in the Deep4Chem dataset, and near 14 nm and 19 nm when splitting by dye molecule and dye scaffold, respectively. This is substantially better than the predictions of TD-DFT calculations alone, and at much lower cost. Our multi-fidelity approach has the advantage that the lower-fidelity data can cover a larger area of chemical space than the higher-fidelity data. Future work should compare this approach to additional methods for training ML models on multi-fidelity data, such as transfer learning (e.g. Figure S16), imputation,  $\Delta$ -ML, and multi-target weighted-loss-function approaches.

D-MPNN approaches perform well across many of the largest publicly-available datasets of the peak absorption wavelength. Additionally, the ChempropMultiFidelity model outperformed Chemprop on a union of the four largest datasets. It achieved a MAE of  $27.78 \pm 5.07$  nm with 5-fold cross validation on scaffold splits of this combined dataset.

We also compared several solvent representations and found that CGSD, Morgan fingerprints, Minnesota descriptors, and D-MPNN fingerprint representations all performed similarly. The Morgan fingerprint and D-MPNN approaches may be advantageous, however, because they can be applied to any solvent because they are not restricted to look-up tables. Future work could also use nearest-neighbor imputation techniques to estimate CGSD or Minnesota descriptors that are not present in the look-up tables.

We demonstrated the qualitative interpretability of our D-MPNN models using dimensionality reduction on their latent space fingerprints, which showed some clustering based on dye scaffold and observed peak wavelength. We also showed that although ensemble variance can be used as a measure of the epistemic uncertainty in our D-MPNN model predictions, and in this case the ensemble errors are comparable in magnitude to the prediction errors, these uncertainties are not necessarily well-correlated with true prediction error on these datasets.

This work is a step toward methods to predict full absorption and emission spectra, and it can enable more rapid design of dye molecules with targeted optical properties.

## **Author Contributions**

R.G.-B. conceived the project. K.P.G. performed the simulations, wrote the computer code, analyzed the data, and wrote the first manuscript draft. W.H.G. and R.G.-B. supervised the research and edited the manuscript.

## **Conflicts of Interest**

There are no conflicts to declare.

## **Acknowledgement**

K.P.G. was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1745302. This work was also supported by the DARPA Accelerated Molecular Discovery (AMD) program under contract HR00111920025. We acknowledge the MIT Engaging cluster and MIT Lincoln Laboratory Supercloud cluster<sup>57</sup> at the Massachusetts Green High Performance Computing Center (MGHPCC) for providing high-performance computing resources to run our TD-DFT calculations and train our deep

learning models. We also thank our colleagues Simon Axelrod, Camille Bilodeau, Michael Forsuelo, Esther Heid, Charles McGill, and Florence Vermeire for helpful scientific discussions.

## Supporting Information Available

The Supporting Information is available below and includes all data and code necessary to reproduce the results shown in this work.

## References

- (1) Venkatraman, V.; Raju, R.; Oikonomopoulos, S. P.; Alsberg, B. K. The dye-sensitized solar cell database. *Journal of Cheminformatics* **2018**, *10*, 1–9.
- (2) Zhu, S.; Tian, R.; Antaris, A. L.; Chen, X.; Dai, H. Near-Infrared-II Molecular Dyes for Cancer Imaging and Surgery. *Advanced Materials* **2019**, *31*, 1–25.
- (3) Labhart, H. Zur quantitativen Beschreibung des Einflusses von Substituenten auf das Absorptionsspektrum ebener Molekeln. Anwendung auf Anthrachinon. *Helvetica Chimica Acta* **1957**, *40*, 1410–1420.
- (4) Loos, P.-F. F.; Scemama, A.; Jacquemin, D. The Quest for Highly Accurate Excitation Energies: A Computational Perspective. *Journal of Physical Chemistry Letters* **2020**, *11*, 2374–2383.
- (5) Laurent, A. D.; Adamo, C.; Jacquemin, D. Dye chemistry with time-dependent density functional theory. *Physical Chemistry Chemical Physics* **2014**, *16*, 14334–14356.
- (6) Laurent, A. D.; Jacquemin, D. TD-DFT benchmarks: A review. *International Journal of Quantum Chemistry* **2013**, *113*, 2019–2039.

- (7) Adamo, C.; Jacquemin, D. The calculations of excited-state properties with time-dependent density functional theory. *Chemical Society Reviews* **2013**, *42*, 845–856.
- (8) Westermayr, J.; Marquetand, P. Machine Learning for Electronically Excited States of Molecules. *Chemical Reviews* **2020**, acs.chemrev.0c00749.
- (9) Dral, P. O.; Barbatti, M. Molecular excited states through a machine learning lens. *Nature Reviews Chemistry* **2021**, *5*, 388–405.
- (10) Chen, C. H.; Tanaka, K.; Funatsu, K. Random Forest Approach to QSPR Study of Fluorescence Properties Combining Quantum Chemical Descriptors and Solvent Conditions. *Journal of Fluorescence* **2018**, *28*, 695–706.
- (11) Talrose, V.; Yermakov, A. N.; Usov, A. A.; Goncharova, A. A.; Leskin, A. N.; Messineva, N. A.; Trusova, N. V.; Efimkina, M. V. In *NIST Chemistry WebBook, NIST Standard Reference Database Number 69*; Linstrom, P., Mallard, W., Eds.
- (12) Mayr, T. Fluorophores.org. <http://www.fluorophores.tugraz.at/substance/>.
- (13) Taniguchi, M.; Lindsey, J. S. Database of Absorption and Fluorescence Spectra of >300 Common Compounds for use in PhotochemCAD. *Photochemistry and Photobiology* **2018**, *94*, 290–327.
- (14) Beard, E. J.; Sivaraman, G.; Vázquez-Mayagoitia, Á.; Vishwanath, V.; Cole, J. M. Comparative dataset of experimental and computational attributes of UV/vis absorption spectra. *Scientific Data* **2019**, *6*, 1–11.
- (15) Venkatraman, V.; Chellappan, L. K. An open access data set highlighting aggregation of dyes on metal oxides. *Data* **2020**, *5*.
- (16) Ju, C. W.; Bai, H.; Li, B.; Liu, R. Machine Learning Enables Highly Accurate Predictions of Photophysical Properties of Organic Fluorescent Materials: Emission Wave-

- lengths and Quantum Yields. *Journal of Chemical Information and Modeling* **2021**, *61*, 1053–1065.
- (17) Joung, J. F.; Han, M.; Jeong, M.; Park, S. Experimental database of optical properties of organic compounds. *Scientific Data* **2020**, *7*, 1–6.
- (18) Noelle, A. et al. UV/Vis+ photochemistry database: Structure, content and applications. *Journal of Quantitative Spectroscopy and Radiative Transfer* **2020**, *253*, 107056.
- (19) Blum, L. C.; Raymond, J.-L. 970 Million Druglike Small Molecules for Virtual Screening in the Chemical Universe Database GDB-13. *J. Am. Chem. Soc.* **2009**, *131*, 8732.
- (20) Montavon, G.; Rupp, M.; Gobre, V.; Vazquez-Mayagoitia, A.; Hansen, K.; Tkatchenko, A.; Müller, K.-R.; von Lilienfeld, O. A. Machine learning of molecular electronic properties in chemical compound space. *New Journal of Physics* **2013**, *15*, 095003.
- (21) Ruddigkeit, L.; Van Deursen, R.; Blum, L. C.; Raymond, J.-L. Enumeration of 166 billion organic small molecules in the chemical universe database GDB-17. *Journal of chemical information and modeling* **2012**, *52*, 2864–2875.
- (22) Ramakrishnan, R.; Hartmann, M.; Tapavicza, E.; Von Lilienfeld, O. A. Electronic spectra from TDDFT and machine learning in chemical space. *The Journal of chemical physics* **2015**, *143*, 084111.
- (23) Liang, J.; Ye, S.; Dai, T.; Zha, Z.; Gao, Y.; Zhu, X. QM-symex, update of the QM-sym database with excited state information for 173 kilo molecules. *Scientific Data* **2020**, *7*, 1–6.
- (24) Nakata, M.; Shimazaki, T. PubChemQC project: a large-scale first-principles electronic structure database for data-driven chemistry. *Journal of chemical information and modeling* **2017**, *57*, 1300–1308.

- (25) Kang, B.; Seok, C.; Lee, J. Prediction of Molecular Electronic Transitions Using Random Forests. *Journal of Chemical Information and Modeling* **2020**,
- (26) Joung, J. F.; Han, M.; Hwang, J.; Jeong, M.; Choi, D. H.; Park, S. Deep Learning Optical Spectroscopy Based on Experimental Database: Potential Applications to Molecular Design. *JACS Au* **2021**, jacsau.1c00035.
- (27) Wang, A. Y.-T.; Murdock, R. J.; Kauwe, S. K.; Oliynyk, A. O.; Gurlo, A.; Brgoch, J.; Persson, K. A.; Sparks, T. D. Machine learning for materials scientists: An introductory guide toward best practices. *Chemistry of Materials* **2020**, *32*, 4954–4965.
- (28) Artrith, N.; Butler, K. T.; Coudert, F.-X.; Han, S.; Isayev, O.; Jain, A.; Walsh, A. Best practices in machine learning for chemistry. *Nature Chemistry* **2021**, *13*, 505–508.
- (29) Tran, A.; Tranchida, J.; Wildey, T.; Thompson, A. P. Multi-fidelity machine-learning with uncertainty quantification and Bayesian optimization for materials design: Application to ternary random alloys. *Journal of Chemical Physics* **2020**, *153*, 1–9.
- (30) Chen, C.; Zuo, Y.; Ye, W.; Li, X.; Ong, S. P. Learning properties of ordered and disordered materials from multi-fidelity data. *Nature Computational Science* **2021**, *1*, 46–53.
- (31) Huang, B.; von Lilienfeld, O. A. Ab Initio Machine Learning in Chemical Compound Space. *Chemical Reviews* **2021**, *121*, 10001–10036.
- (32) Yang, K.; Swanson, K.; Jin, W.; Coley, C. W.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M.; Palmer, A.; Settels, V.; Jaakkola, T. S.; Jensen, K. F.; Barzilay, R. Analyzing Learned Molecular Representations for Property Prediction. *Journal of Chemical Information and Modeling* **2019**, *59*, 3370–3388.
- (33) Weininger, D. SMILES, a chemical language and information system. 1. Introduction

- to methodology and encoding rules. *Journal of Chemical Information and Modeling* **1988**, *28*, 31–36.
- (34) Landrum, G. RDKit: Open-source cheminformatics. 2006; [www.rdkit.org](http://www.rdkit.org).
- (35) Riniker, S.; Landrum, G. A. Better informed distance geometry: using what we know to improve conformation generation. *Journal of chemical information and modeling* **2015**, *55*, 2562–2574.
- (36) Bannwarth, C.; Ehlert, S.; Grimme, S. GFN2-xTB - An Accurate and Broadly Parametrized Self-Consistent Tight-Binding Quantum Chemical Method with Multipole Electrostatics and Density-Dependent Dispersion Contributions. *Journal of Chemical Theory and Computation* **2019**, *15*, 1652–1671.
- (37) Neese, F.; Wennmohs, F.; Becker, U.; Riplinger, C. The ORCA quantum chemistry program package. *The Journal of chemical physics* **2020**, *152*, 224108.
- (38) Becke, A. D. Density-functional exchange-energy approximation with correct asymptotic behavior. *Physical Review A* **1988**, *38*, 3098–3100.
- (39) Grimme, S.; Ehrlich, S.; Goerigk, L. Effect of the damping function in dispersion corrected density functional theory. *Journal of Computational Chemistry* **2011**, *32*, 1456–1465.
- (40) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Physical Chemistry Chemical Physics* **2005**, *7*, 3297–3305.
- (41) Hirata, S.; Head-Gordon, M. Time-dependent density functional theory within the Tamm–Dancoff approximation. *Chemical Physics Letters* **1999**, *314*, 291–299.
- (42) Chai, J. D.; Head-Gordon, M. Long-range corrected double-hybrid density functionals. *Journal of Chemical Physics* **2009**, *131*.

- (43) Frisch, M. J. et al. Gaussian 09, Revision B.01. Gaussian Inc. Wallingford CT 2009.
- (44) Yap, C. W. PaDEL-descriptor: An open source software to calculate molecular descriptors and fingerprints. *Journal of computational chemistry* **2011**, *32*, 1466–1474.
- (45) Kessler, T.; Oliveira, L. PaDELPy: A Python wrapper for PaDEL-Descriptor software. 2021; <https://github.com/ecrl/padelpy>.
- (46) Winget, P.; Dolney, D. M.; Giesen, D. J.; Cramer, C. J.; Truhlar, D. G. Minnesota solvent descriptor database for DFT calculation. 2010; <https://comp.chem.umn.edu/solvation/mnsddb.pdf>.
- (47) Reichardt, C. Solvatochromic dyes as solvent polarity indicators. *Chemical Reviews* **1994**, *94*, 2319–2358.
- (48) Catalán, J. Toward a generalized treatment of the solvent effect based on four empirical scales: Dipolarity (SdP, a new scale), polarizability (SP), acidity (SA), and basicity (SB) of the medium. *Journal of Physical Chemistry B* **2009**, *113*, 5951–5960.
- (49) Vermeire, F. H.; Green, W. H. Transfer learning for solvation free energies: From quantum chemistry to experiments. *Chemical Engineering Journal* **2021**, *418*, 129307.
- (50) Friedman, J. H. Greedy function approximation: a gradient boosting machine. *Annals of statistics* **2001**, 1189–1232.
- (51) Jia, X.; Willard, J.; Karpatne, A.; Read, J. S.; Zwart, J. A.; Steinbach, M.; Kumar, V. Physics-guided machine learning for scientific discovery: An application in simulating lake temperature profiles. *ACM/IMS Transactions on Data Science* **2021**, *2*, 1–26.
- (52) Clark, S.; Hayes, P. SigOpt Web page. <https://sigopt.com>, 2019; <https://sigopt.com>.
- (53) Bemis, G. W.; Murcko, M. A. The properties of known drugs. 1. Molecular frameworks. *Journal of Medicinal Chemistry* **1996**, *39*, 2887–2893.

- (54) Scalia, G.; Grambow, C. A.; Pernici, B.; Li, Y.-P.; Green, W. H. Evaluating Scalable Uncertainty Estimation Methods for Deep Learning-Based Molecular Property Prediction. *Journal of Chemical Information and Modeling* **2020**, *60*, 2697–2717.
- (55) Hirschfeld, L.; Swanson, K.; Yang, K.; Barzilay, R.; Coley, C. W. Uncertainty Quantification Using Neural Networks for Molecular Property Prediction. *Journal of Chemical Information and Modeling* **2020**, *60*, 3770–3780.
- (56) Nigam, A.; Pollice, R.; Hurley, M. F. D.; Hickman, R. J.; Aldeghi, M.; Yoshikawa, N.; Chithrananda, S.; Voelz, V. A.; Aspuru-Guzik, A. Assigning confidence to molecular property prediction. *Expert Opinion on Drug Discovery* **2021**, 1–15.
- (57) Reuther, A.; Kepner, J.; Byun, C.; Samsi, S.; Arcand, W.; Bestor, D.; Bergeron, B.; Gadepally, V.; Houle, M.; Hubbell, M., et al. Interactive supercomputing on 40,000 cores for machine learning and data analysis. 2018 IEEE High Performance extreme Computing Conference (HPEC). 2018; pp 1–6.

# Supporting Information

## Code and Data Availability

All code to reproduce our workflow and figures and all data including TD-DFT calculation results is available at <https://doi.org/10.5281/zenodo.5500428>. To make predictions using Chemprop and ChempropMultiFidelity models, you can use the UVVisML tool at <https://github.com/learningmatter-mit/uvvisml>.

## Note on CGSD Solvent Descriptors

We identified two transcription errors by Ju et al. in the CGSD solvent descriptors (available from <https://figshare.com/articles/dataset/ChemFluor/12110619/3>) used to train their GBRT models.<sup>16</sup>

- In row 22 of `Solvent_Descriptors.xlsx`, the  $E_T(30)$  value for 1-methyl-2-pyrrolidinone should be 42.2 rather than 48 according to entry no. 284 of Table 2 of Reichardt's work.<sup>47</sup>
- In row 24 of `Solvent_Descriptors.xlsx`, the SP, SdP, SA, and SB values for N-methylformamide are taken from entry no. 29 (N,N-dimethylformamide) of Table 2 of Catalán's work.<sup>48</sup> No entry exists for N-methylformamide.

We corrected these errors in our file `data_solvents/chemfluor_cgsgd_solvent_db.csv`. This file contains more solvents than the original ChemFluor file because we parsed the entirety of the entries in Reichardt's and Catalán's works to have coverage over as many solvents as possible (rather than restricting ourselves to those in the ChemFluor dataset).

## Dataset Composition

Table S1: **Dataset Composition by Solvent.** Ten of the 364 unique solvents appeared in more than 1,000 measurements.

Solvent Name	Solvent SMILES	Number of Measurements
Other	-	5653
Dichloromethane	<chem>C1CCl</chem>	5621
Tetrahydrofuran	<chem>C1CCOC1</chem>	3328
Chloroform	<chem>ClC(Cl)Cl</chem>	2423
Acetonitrile	<chem>CC#N</chem>	2343
Ethanol	<chem>CCO</chem>	1824
Methylbenzene	<chem>Cc1ccccc1</chem>	1793
Methanol	<chem>CO</chem>	1758
DMSO	<chem>CS(C)=O</chem>	1637
Water	<chem>O</chem>	1181
Dimethylformamide	<chem>CN(C)C=O</chem>	1173

## Datasets Comparison

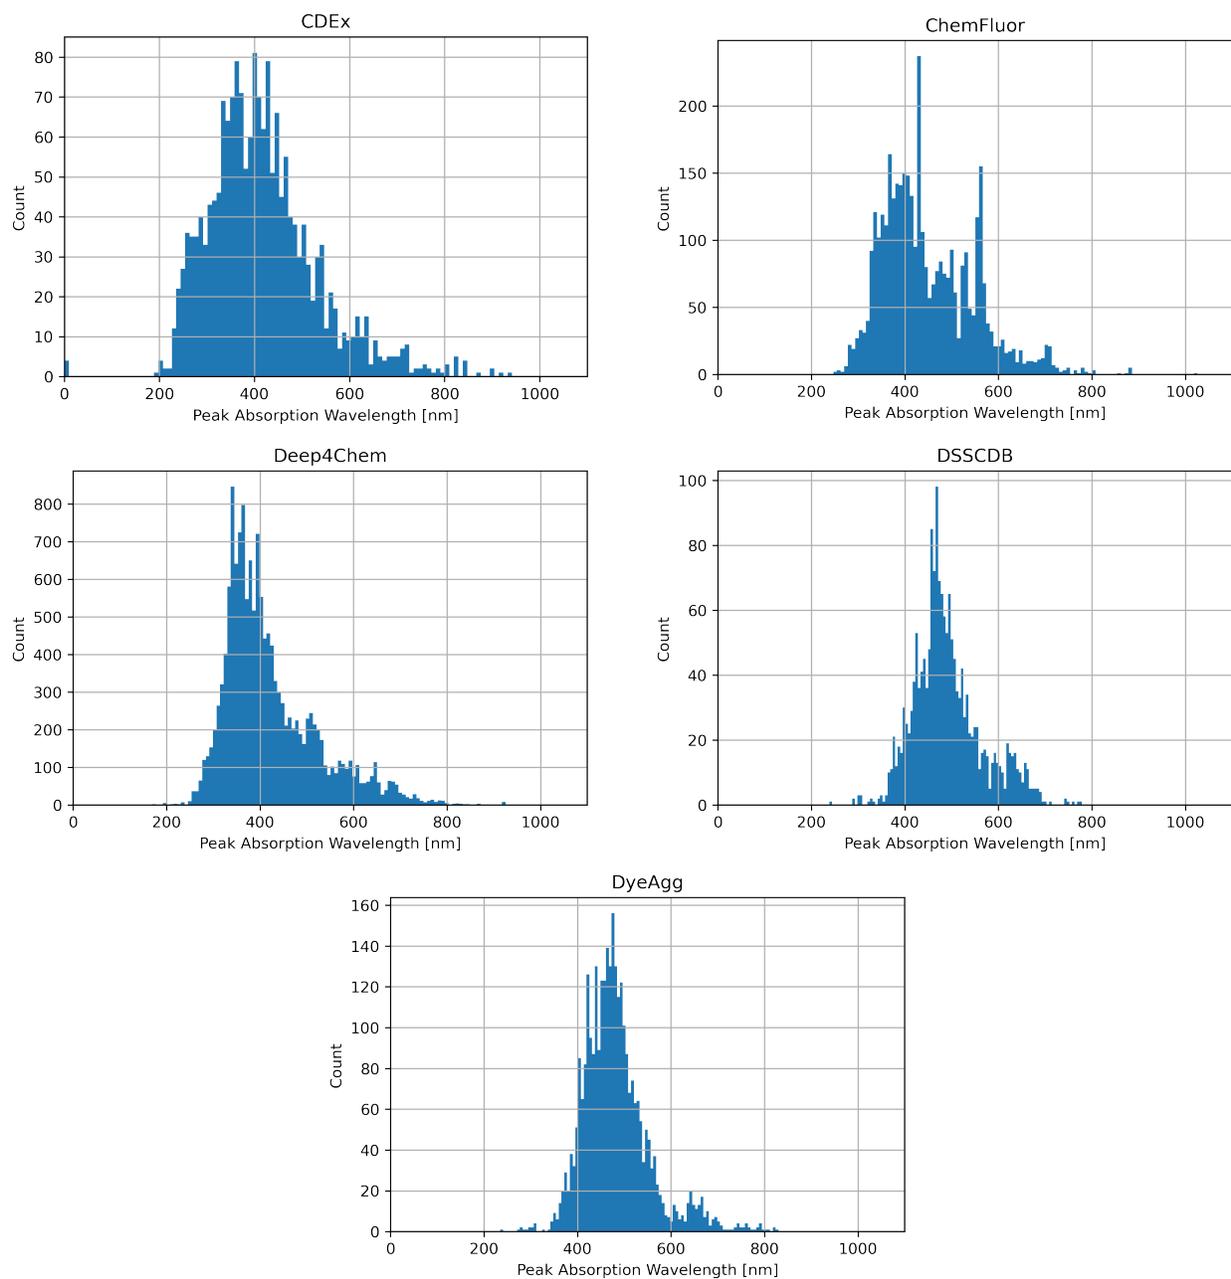


Figure S1: **Dataset Comparison.** The five datasets differ in their coverage of chemical space, which results in different distributions of peak absorption wavelength.

## RMSE Bar Plots

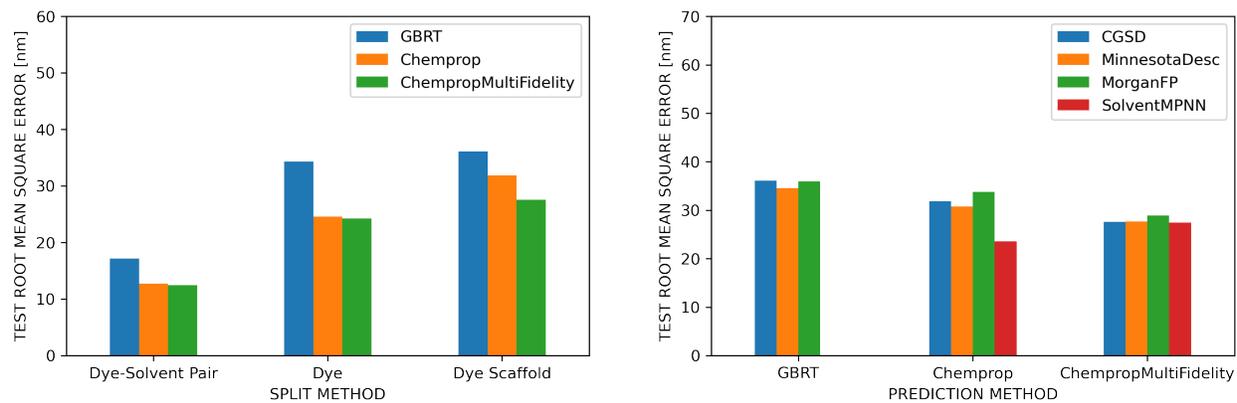


Figure S2

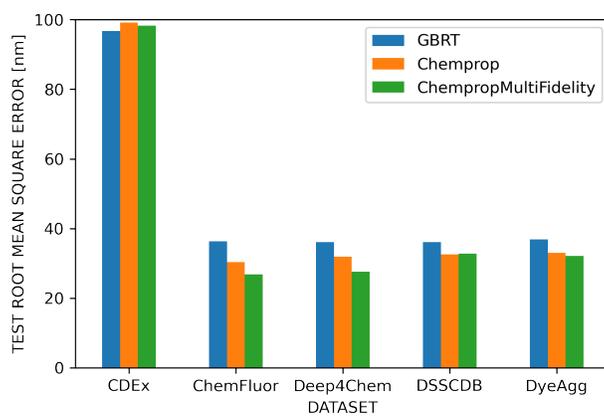


Figure S3

## R2 Bar Plots

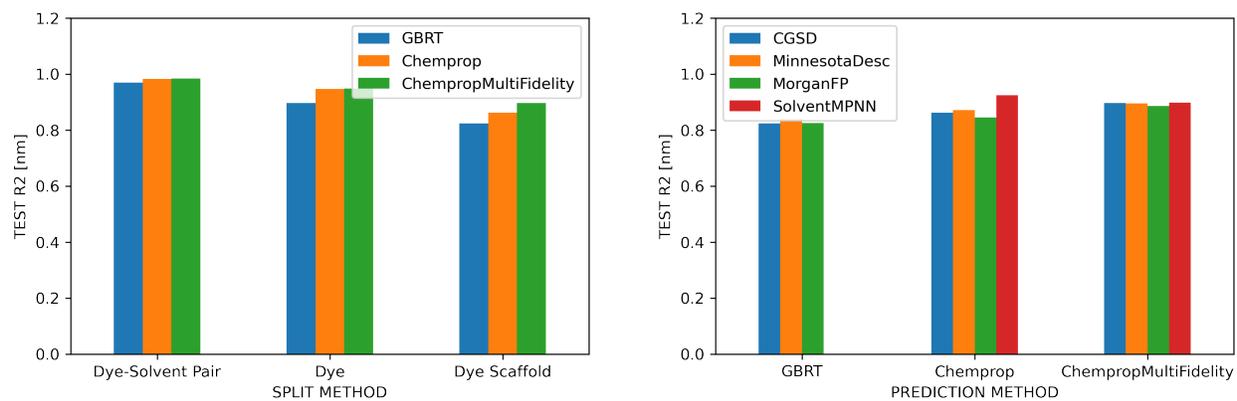


Figure S4

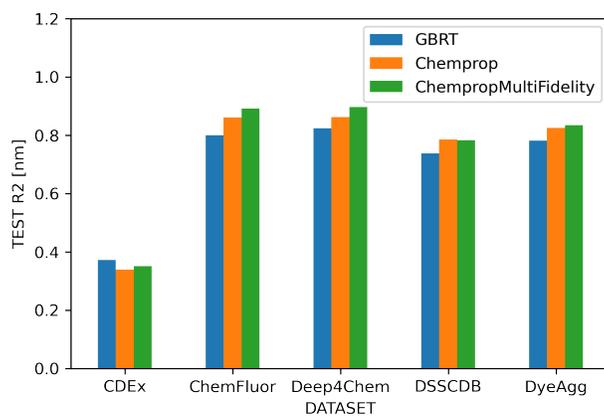


Figure S5

## MAE, RMSE, and $R^2$ Tables

Table S2: Molecule Representations - MAE (Deep4Chem Dataset)

GBRT	Chemprop	ChempropMultiFidelity	Solvent Rep.	Split Type
8.30	6.95	6.97	CGSD	random
8.31	7.32	7.08	Minnesota	random
8.98	8.27	7.89	Morgan	random
NaN	7.00	6.99	MPNN	random
20.55	15.12	13.90	CGSD	group_by_smiles
20.90	14.96	14.17	Minnesota	group_by_smiles
20.79	15.01	13.77	Morgan	group_by_smiles
NaN	14.29	14.16	MPNN	group_by_smiles
26.94	21.35	18.33	CGSD	scaffold
25.51	20.47	18.40	Minnesota	scaffold
26.05	21.77	18.84	Morgan	scaffold
NaN	15.03	18.72	MPNN	scaffold

Table S3: Molecule Representations - RMSE (Deep4Chem Dataset)

GBRT	Chemprop	ChempropMultiFidelity	Solvent Rep.	Split Type
17.18	12.74	12.49	CGSD	random
16.65	12.90	12.54	Minnesota	random
17.56	14.24	13.68	Morgan	random
NaN	12.93	12.70	MPNN	random
34.29	24.60	24.24	CGSD	group_by_smiles
34.36	23.94	25.24	Minnesota	group_by_smiles
34.11	25.31	23.37	Morgan	group_by_smiles
NaN	24.01	25.06	MPNN	group_by_smiles
36.08	31.87	27.57	CGSD	scaffold
34.58	30.74	27.70	Minnesota	scaffold
35.92	33.77	28.91	Morgan	scaffold
NaN	23.60	27.47	MPNN	scaffold

Table S4: Molecule Representations -  $R^2$  (Deep4Chem Dataset)

GBRT	Chemprop	ChempropMultiFidelity	Solvent Rep.	Split Type
0.97	0.98	0.98	CGSD	random
0.97	0.98	0.98	Minnesota	random
0.97	0.98	0.98	Morgan	random
NaN	0.98	0.98	MPNN	random
0.90	0.95	0.95	CGSD	group_by_smiles
0.90	0.95	0.94	Minnesota	group_by_smiles
0.90	0.94	0.95	Morgan	group_by_smiles
NaN	0.95	0.94	MPNN	group_by_smiles
0.82	0.86	0.90	CGSD	scaffold
0.84	0.87	0.90	Minnesota	scaffold
0.83	0.85	0.89	Morgan	scaffold
NaN	0.92	0.90	MPNN	scaffold

Table S5: Solvent Representations - MAE (Deep4Chem Dataset)

CGSD	MinnesotaDesc	MorganFP	SolventMPNN	Model	Split Type
8.30	8.31	8.98	NaN	GBRT	random
6.95	7.32	8.27	7.00	Chemprop	random
6.97	7.08	7.89	6.99	ChempropMultiFidelity	random
20.55	20.90	20.79	NaN	GBRT	group_by_smiles
15.12	14.96	15.01	14.29	Chemprop	group_by_smiles
13.90	14.17	13.77	14.16	ChempropMultiFidelity	group_by_smiles
26.94	25.51	26.05	NaN	GBRT	scaffold
21.35	20.47	21.77	15.03	Chemprop	scaffold
18.33	18.40	18.84	18.72	ChempropMultiFidelity	scaffold

Table S6: Solvent Representations - RMSE (Deep4Chem Dataset)

CGSD	MinnesotaDesc	MorganFP	SolventMPNN	Model	Split Type
17.18	16.65	17.56	NaN	GBRT	random
12.74	12.90	14.24	12.93	Chemprop	random
12.49	12.54	13.68	12.70	ChempropMultiFidelity	random
34.29	34.36	34.11	NaN	GBRT	group_by_smiles
24.60	23.94	25.31	24.01	Chemprop	group_by_smiles
24.24	25.24	23.37	25.06	ChempropMultiFidelity	group_by_smiles
36.08	34.58	35.92	NaN	GBRT	scaffold
31.87	30.74	33.77	23.60	Chemprop	scaffold
27.57	27.70	28.91	27.47	ChempropMultiFidelity	scaffold

Table S7: Solvent Representations -  $R^2$  (Deep4Chem Dataset)

CGSD	MinnesotaDesc	MorganFP	SolventMPNN	Model	Split Type
0.97	0.97	0.97	NaN	GBRT	random
0.98	0.98	0.98	0.98	Chemprop	random
0.98	0.98	0.98	0.98	ChempropMultiFidelity	random
0.90	0.90	0.90	NaN	GBRT	group_by_smiles
0.95	0.95	0.94	0.95	Chemprop	group_by_smiles
0.95	0.94	0.95	0.94	ChempropMultiFidelity	group_by_smiles
0.82	0.84	0.83	NaN	GBRT	scaffold
0.86	0.87	0.85	0.92	Chemprop	scaffold
0.90	0.90	0.89	0.90	ChempropMultiFidelity	scaffold

Table S8: Datasets - MAE (CGSD Solvent Representation)

GBRT	Chemprop	ChempropMultiFidelity	Dataset	Split Type
61.10	56.40	55.04	CDEx	scaffold
22.08	18.48	17.10	ChemFluor	scaffold
26.94	21.35	18.33	Deep4Chem	scaffold
23.27	20.08	19.69	DSSCDB	scaffold
24.36	21.18	22.14	DyeAgg	scaffold

Table S9: Datasets - RMSE (CGSD Solvent Representation)

GBRT	Chemprop	ChempropMultiFidelity	Dataset	Split Type
96.70	99.16	98.27	CDEx	scaffold
36.36	30.39	26.82	ChemFluor	scaffold
36.08	31.87	27.57	Deep4Chem	scaffold
36.07	32.59	32.79	DSSCDB	scaffold
36.93	33.01	32.11	DyeAgg	scaffold

Table S10: Datasets -  $R^2$  (CGSD Solvent Representation)

GBRT	Chemprop	ChempropMultiFidelity	Dataset	Split Type
0.37	0.34	0.35	CDEx	scaffold
0.80	0.86	0.89	ChemFluor	scaffold
0.82	0.86	0.90	Deep4Chem	scaffold
0.74	0.79	0.78	DSSCDB	scaffold
0.78	0.83	0.83	DyeAgg	scaffold

## Error and Uncertainty Distribution Plots

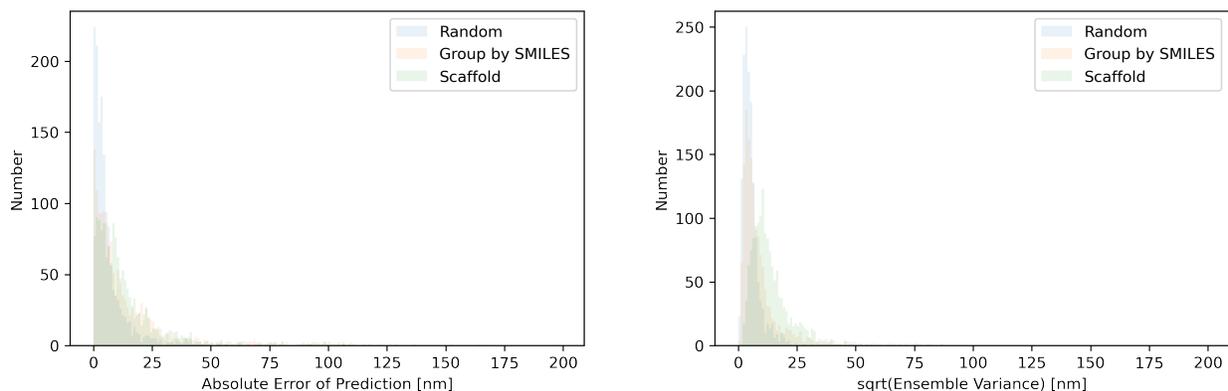


Figure S6: **Error and Uncertainty Distributions for 3 Split Types.** The error and uncertainty distributions both spread out as the splits become more rigorous from random to group by SMILES to scaffold.

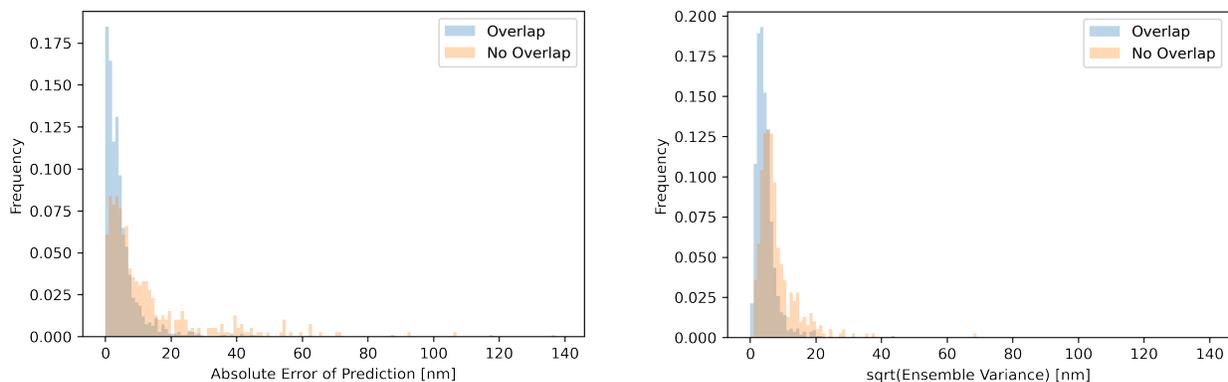


Figure S7: **Error and Uncertainty Distributions of Random Splits.** Using random splits, many dye molecules appear in both the train and test sets because they were measured in different solvents. For such molecules (“overlap”), the error and uncertainty distributions are tighter than for those molecules without overlap in train and test sets.

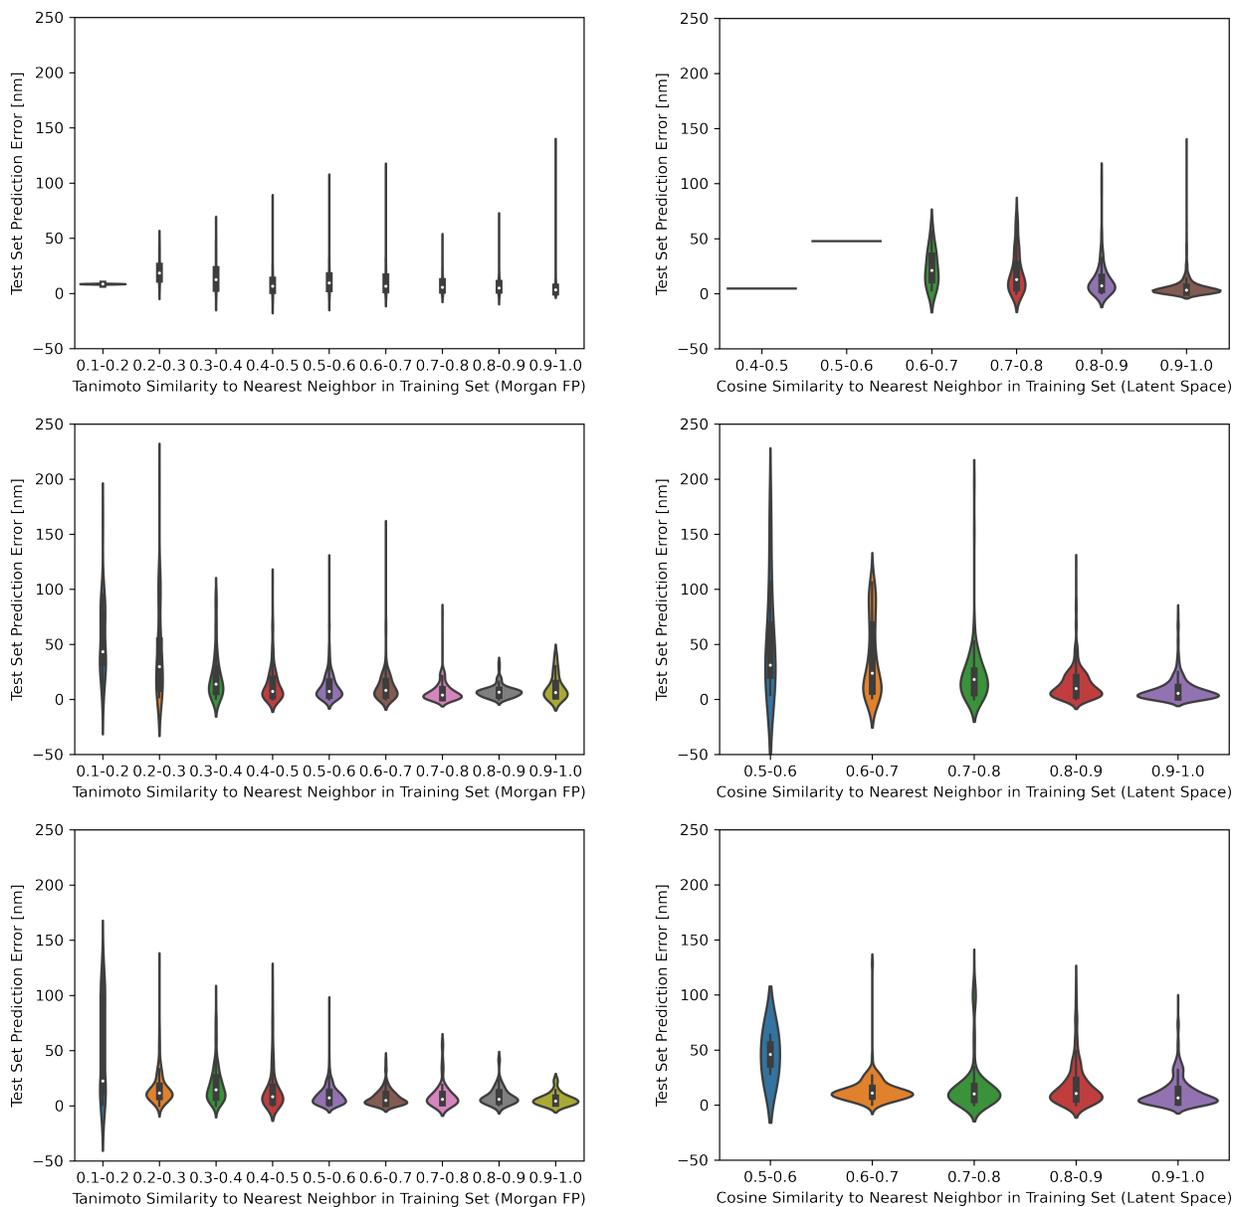


Figure S8: **Error distributions as a function of similarity to training set.** (Top) random splits; (Middle) group by SMILES splits; (Bottom) scaffold splits. As similarity to the nearest neighbor in the training set increases, we expect that the maximum of the prediction errors will decrease. This is mostly true for the more rigorous split types, but nearly the opposite is true for random splits.

## TD-DFT Model Performance

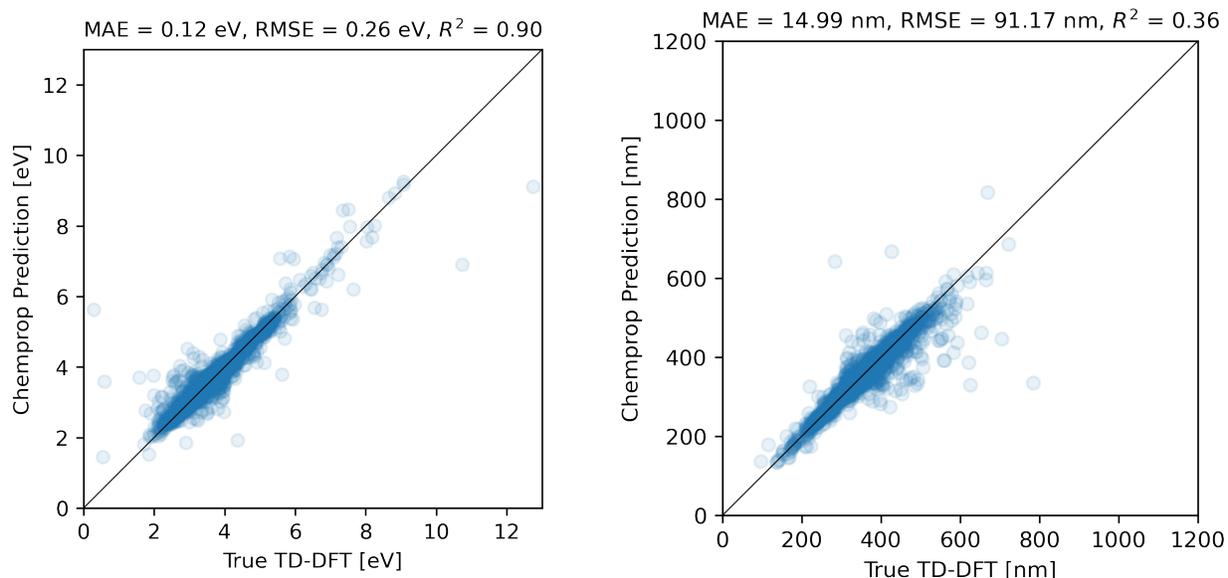


Figure S9: **Parity plot of Chemprop TD-DFT Predictions vs. True TD-DFT Results.** (Left) Units of eV; (Right) Units of nm. Both plots represent the exact same data (predictions of a model trained on data in eV). The TD-DFT model (a Chemprop D-MPNN and FFNN) used an ensemble size of 1 and was trained on 80-10-10 random splits of 28,772 vertical excitation energies from our full set of vacuum TD-DFT calculations. The large RMSE and low  $R^2$  in nanometers is a result of the few outliers with very low true TD-DFT values (these values are visible on the left-most part of the eV plot with true TD-DFT values less than 1 eV, but they exceed the upper limit on the x-axis in the nm plot).

## TD-DFT as a Feature

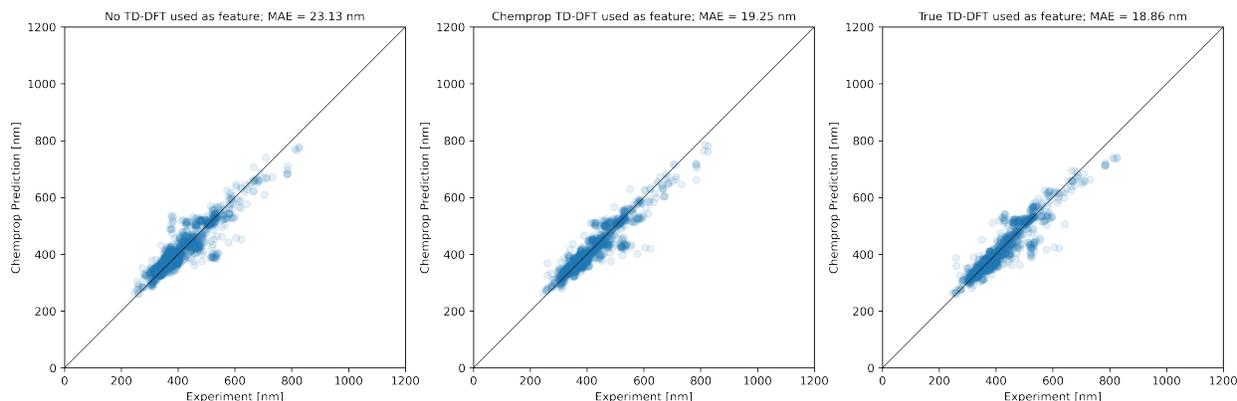


Figure S10: **Trade-Off Between True TD-DFT and Chemprop-Predicted TD-DFT as a Feature.** Prediction using (Left) no additional features, (Center) the TD-DFT value predicted by an auxiliary Chemprop D-MPNN model as a feature, or (Right) the true TD-DFT value as a feature. All models used a D-MPNN solvent representation, scaffold splits of the Deep4Chem dataset, and an ensemble size of 1. Using true TD-DFT values results in better performance (MAE = 18.86 nm) than using TD-DFT values predicted by the Chemprop D-MPNN (MAE = 19.25 nm), but the D-MPNN method retains most of the benefit over using no TD-DFT feature at all (23.13 nm).

## Comparing Predictions Across Data Sources

All models in this section used the CGSD solvent representation. Each data source (CDEx, ChemFluor, Deep4Chem, DSSCDB, and DyeAgg) was split by scaffolds, but the training set from one data source may have scaffolds that overlap with those in the test set of another data source

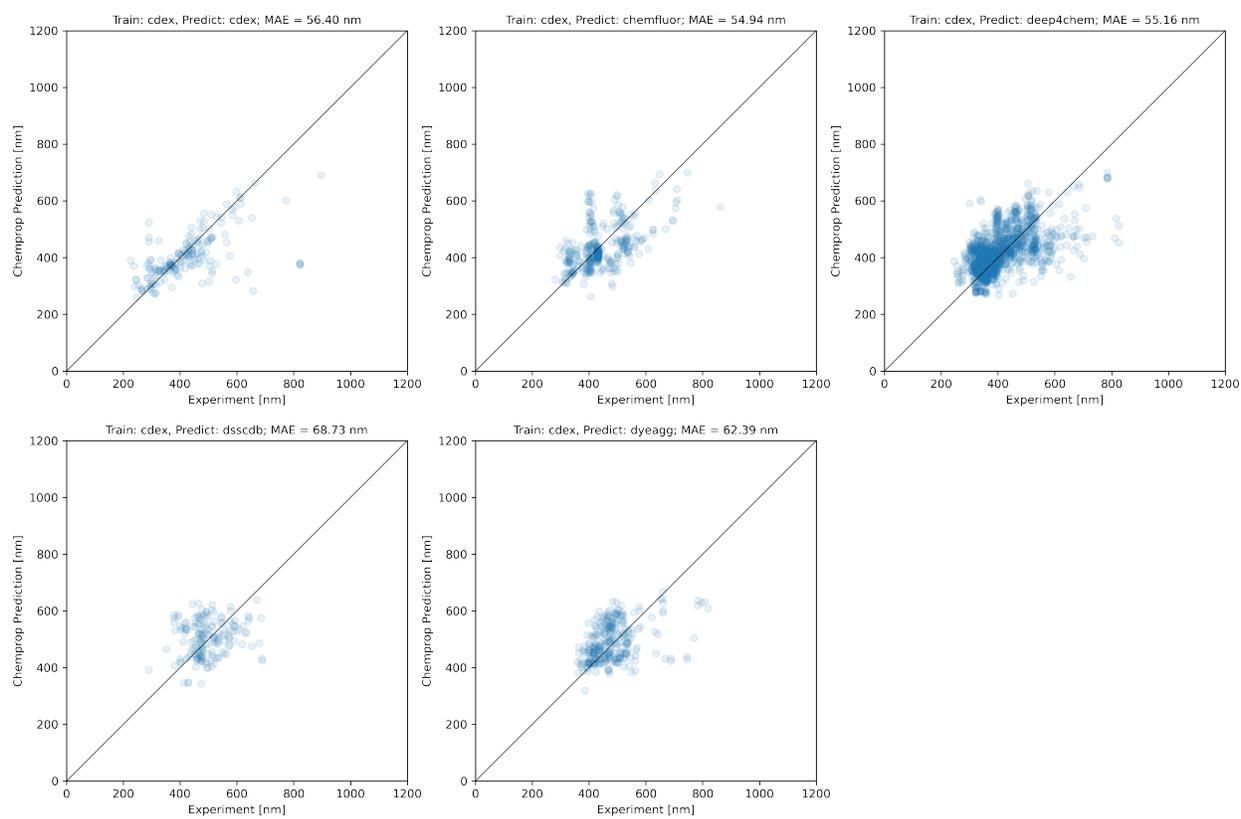


Figure S11: Predictions on all datasets after training on CDEX dataset.

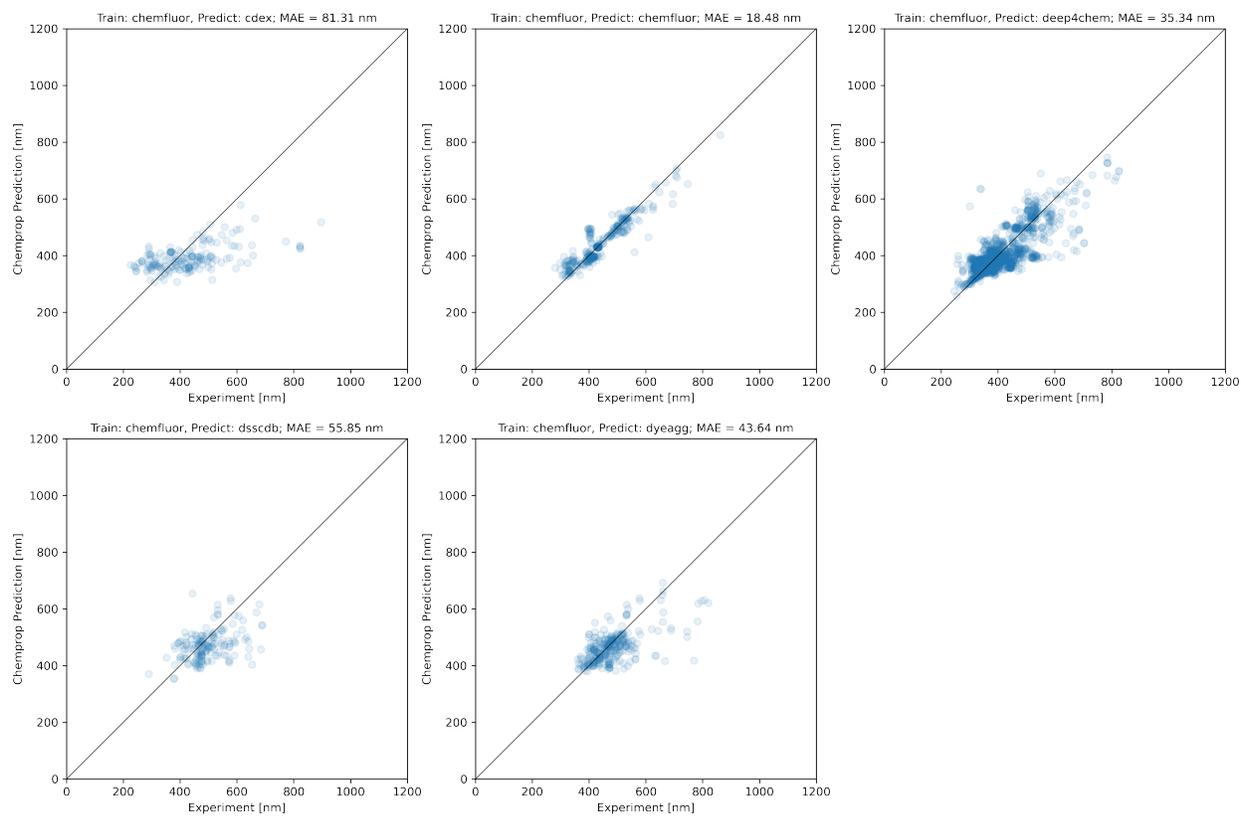


Figure S12: Predictions on all datasets after training on ChemFluor dataset.

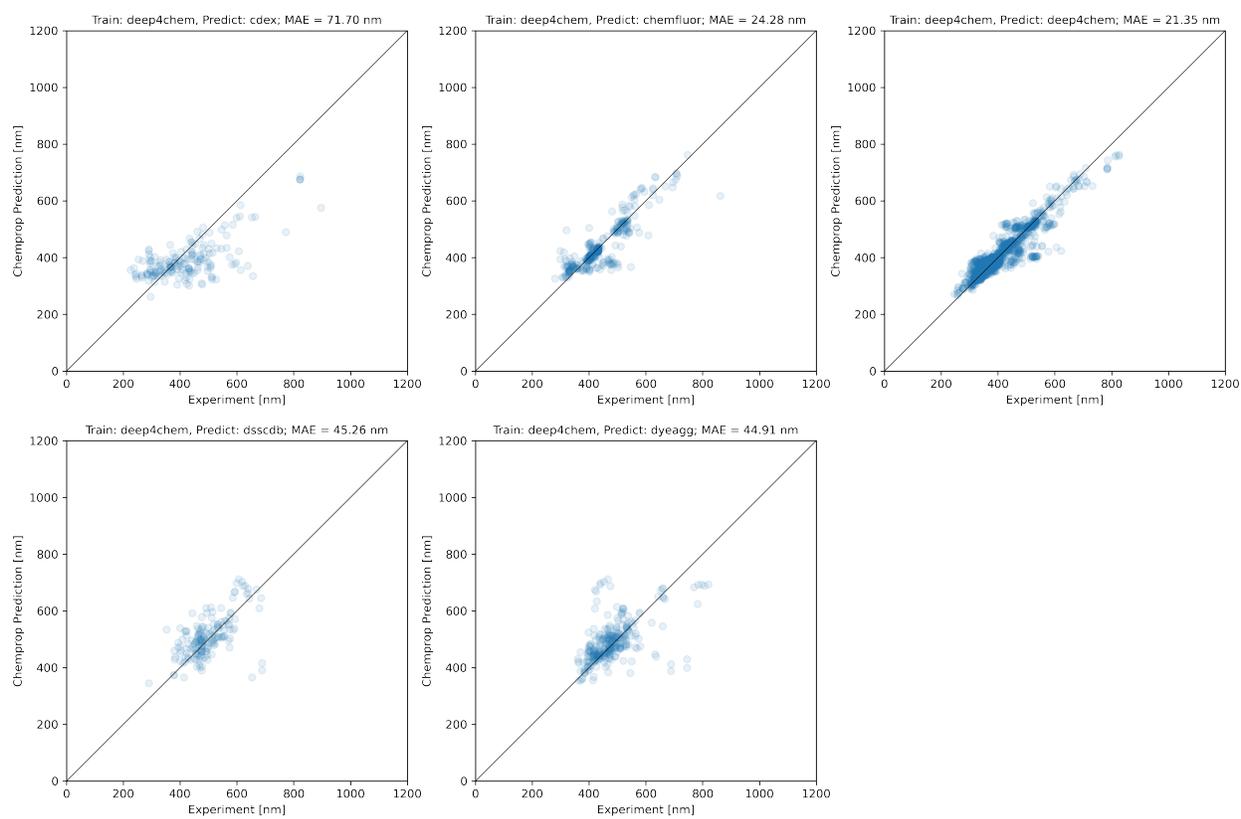


Figure S13: Predictions on all datasets after training on Deep4Chem dataset.

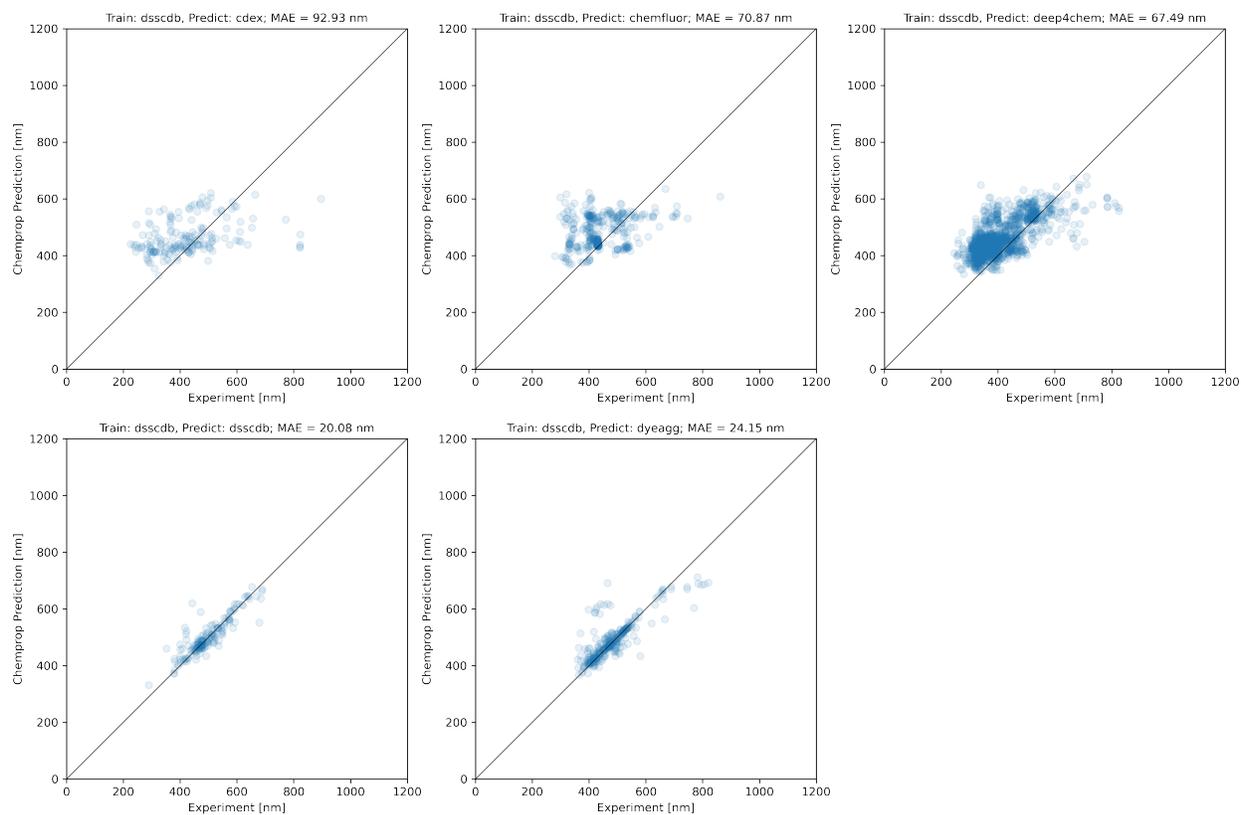


Figure S14: Predictions on all datasets after training on DSSCDB dataset.

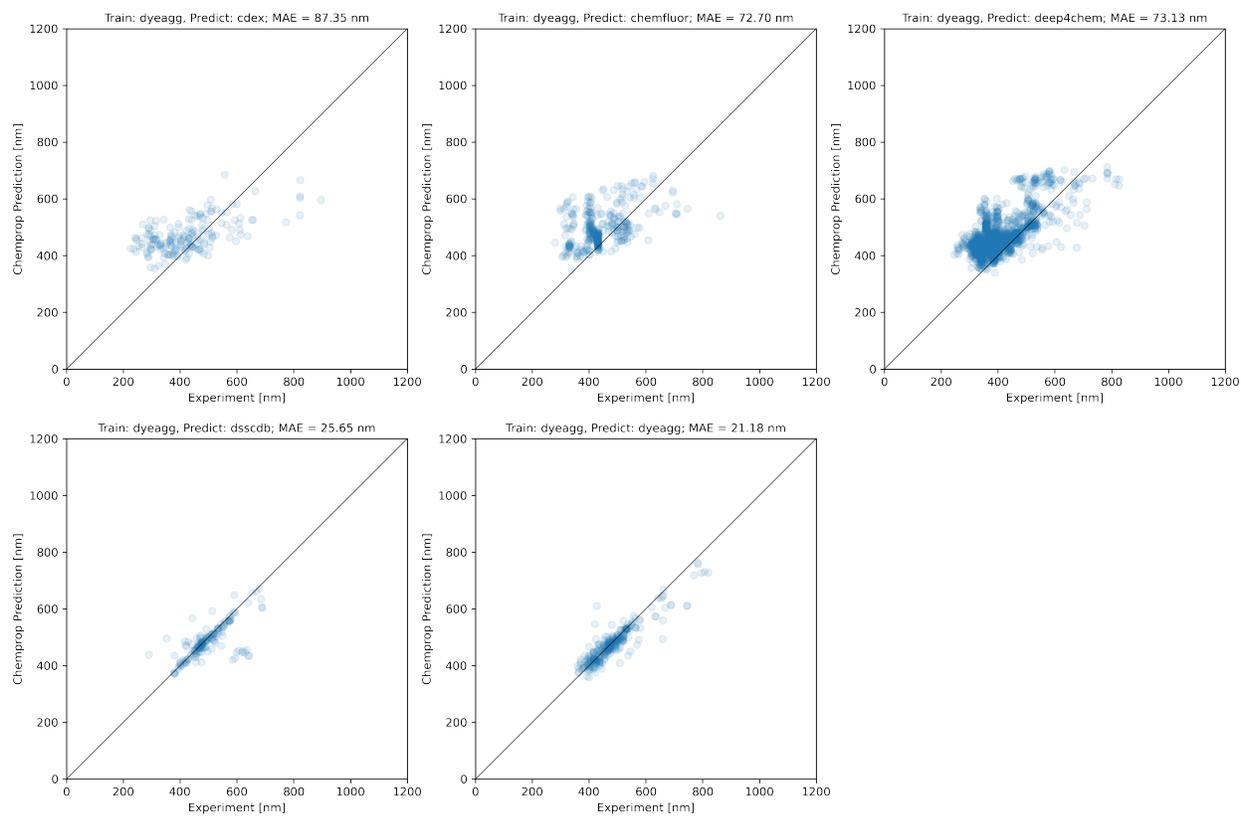


Figure S15: Predictions on all datasets after training on DyeAgg dataset.

## Transfer Learning Baseline

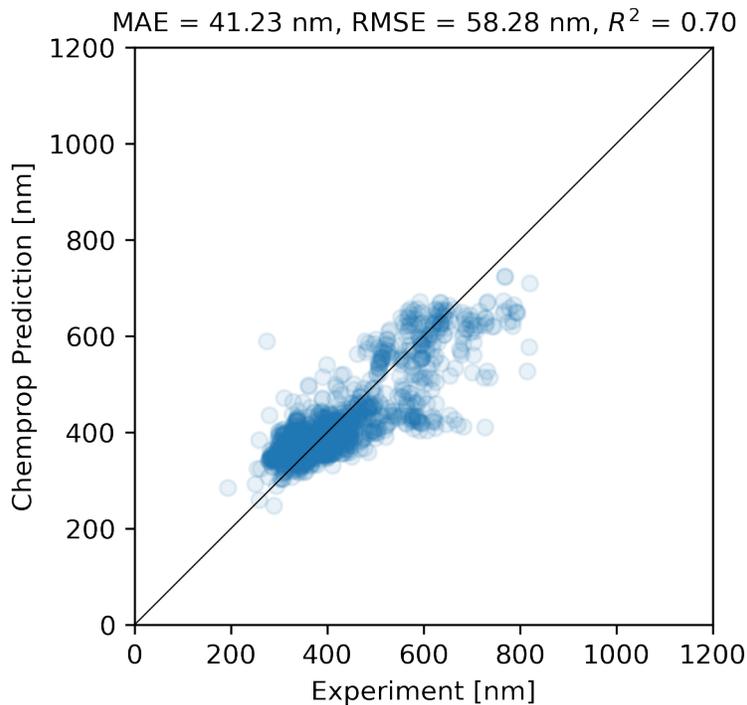


Figure S16: **Predicted Peak Wavelength vs. True Peak Wavelength Using Transfer Learning.** A model was pretrained on TD-DFT data and the D-MPNN parameters were frozen while training on experimental data. CGSD were used to represent the solvent for the experiments. This transfer learning approach performs much worse than our multi-fidelity approach.