

Chemprop: A Machine Learning Package for Chemical Property Prediction

Esther Heid,^{†,‡} Kevin P. Greenman,[†] Yunsie Chung,[†] Shih-Cheng Li,^{†,¶} David E. Graff,^{†,§} Florence H. Vermeire,^{†,||} Haoyang Wu,[†] William H. Green,[†] and Charles J. McGill^{*,†,⊥}

[†]*Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States*

[‡]*Institute of Materials Chemistry, TU Wien, 1060 Vienna, Austria*

[¶]*Department of Chemical Engineering, National Taiwan University, Taipei 10617, Taiwan*

[§]*Department of Chemistry and Chemical Biology, Harvard University, Cambridge, Massachusetts 02138, United States*

^{||}*Department of Chemical Engineering, KU Leuven, Celestijnenlaan 200F, B-3001 Leuven, Belgium*

[⊥]*Department of Chemical and Life Science Engineering, Virginia Commonwealth University, Richmond, Virginia 23284, United States*

E-mail: mcgillc2@vcu.edu

Abstract

Deep learning has become a powerful and frequently employed tool for the prediction of molecular properties, thus creating a need for open-source and versatile software solutions that can be operated by non-experts. Among current approaches, directed message-passing neural networks (D-MPNNs) have proven to perform well on a variety of property prediction tasks. The software package Chemprop implements

the D-MPNN architecture, and offers simple, easy, and fast access to machine-learned molecular properties. Compared to its initial version, we present a multitude of new Chemprop functionalities such as the support of multi-molecule properties, reactions, atom/bond-level properties, and spectra. Further, we incorporate various uncertainty quantification and calibration methods along with related metrics, as well as pretraining and transfer learning workflows, improved hyperparameter optimization, and other customization options concerning loss functions or atom/bond features. We benchmark D-MPNN models trained using Chemprop with the new reaction, atom-level and spectra functionality on a variety of property prediction datasets, including MoleculeNet and SAMPL, and observe state-of-the-art performance on the prediction of water-octanol partition coefficients, reaction barrier heights, atomic partial charges, and absorption spectra. Chemprop enables out-of-the-box training of D-MPNN models for a variety of problem settings in a fast, user-friendly, and open-source software.

1 Introduction

Machine learning in general, and especially deep learning, has become a powerful tool in various fields of chemistry. Applications range from the prediction of physico-chemical¹⁻⁹ and pharmacological¹⁰ properties of molecules to the design of molecules or materials with certain properties,¹¹⁻¹³ the exploration of chemical synthesis pathways,¹⁴⁻²⁷ or the prediction of properties important for chemical analysis like IR,²⁸ UV/VIS,²⁹ or mass spectra.³⁰⁻³³

Many combinations of molecular representation and model architecture have been developed to extract features from molecules and predict molecular properties. Molecules can be represented as graphs, strings, pre-computed feature vectors, or sets of atomic coordinates and processed using graph-convolutional neural networks, transformers, or feed-forward neural networks to train predictive models. While early works focused on hand-made features or simple fingerprinting methods combined with kernel regression or neural networks,³⁴ the current state-of-the-art has shifted to end-to-end trainable models which directly learn to

extract their own features.³⁵ Here, the models can achieve extreme complexity based on the mechanisms of information exchange between parts of the molecule. For example, graph convolutional neural networks (GCNNs) extract local information from the molecular graph for single or small groups of atoms, and use that information to update the immediate neighborhood.^{1-3,36} They offer robust performance for properties dependent on the local structure, and if the three-dimensional conformation of a molecule is not known or not relevant for a prediction task. Graph attention transformers allow for a less local information exchange via attention layers, which learn to accumulate the features of atoms both close and far away in the graph.^{37,38} Another important line of research comprises the prediction of properties dependent on the three-dimensional conformation of a molecule, such as the prediction of properties obtained from quantum mechanics.^{2,39-41} Finally, transformer models from natural language processing can be trained on string representations such as SMILES or SELFIES, also leading to promising results.⁴²⁻⁴⁵ In this work, we discuss our application of GCNNs, namely Chemprop,³⁶ a directed-message passing algorithm derived from the seminal work of Gilmer et al.¹

An early version of Chemprop was published in Ref. 36. Since then, the software has substantially evolved, and now includes a vast collection of new features. For example, Chemprop is now able to predict properties for systems containing multiple molecules, such as solute/solvent combinations, or reactions with and without solvent. It can train on molecular targets, spectra, or atom/bond-level targets, and output the latent representation for analysis of the learned feature embedding. Available uncertainty metrics include popular approaches such as ensembling, mean-variance estimation, and evidential learning. Chemprop is thus a general and versatile deep learning toolbox, and enjoys a wide user base.

The remainder of the article is structured as follows: First we summarize the architecture of Chemprop. We discuss a selection of Chemprop features, with a focus on features introduced after the initial release of Chemprop. We then conclude the main body of the article, and provide details on the data and software, which we have open-sourced includ-

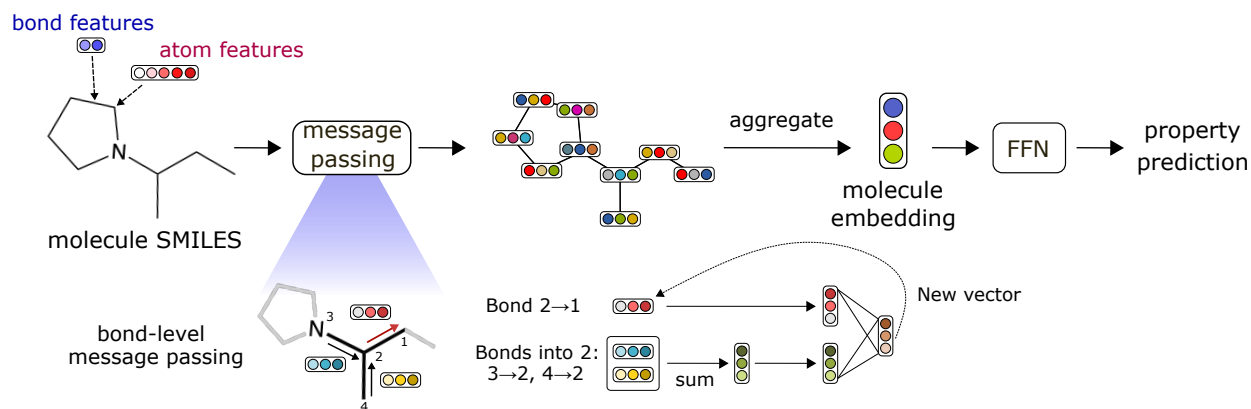


Figure 1: Overview of the architecture of Chemprop. The message passing update of the hidden vector for directed edge $2 \rightarrow 1$ is expanded for demonstration.

ing all scripts to allow for full reproducibility. Alongside the main body of this article, we provide *Supporting Information* that contains: further model design details; descriptions of the data acquisition, preprocessing, and splitting of all datasets used in benchmarking; and the results of Chemprop benchmarks on a variety of datasets showcasing its performance on both simple and advanced prediction tasks.

2 Model structure

Chemprop consists of four modules: 1) a local features encoding function, 2) a directed message passing neural network (D-MPNN) to learn atomic embeddings from the local features, 3) an aggregation function to join atomic embeddings into molecular embeddings, and 4) a standard feed-forward neural network (FFN) for the transformation of molecular embeddings to target properties, summarized in Fig. 1. The D-MPNN is a class of graph-convolutional neural network (GCNN), which updates hidden representations of the vertices V and edges E of a graph G based on the local environment. In the following, we will use bold lower case to denote vectors, bold upper case to denote matrices, and italic light font for scalars and objects.

For a molecule, the molecule SMILES string is used as input, which is then transformed to a molecular graph using RDKit⁴⁶, where atoms correspond to vertices, and bonds to edges.

Initial features are constructed based on the identity and topology of each atom and bond. For each vertex v , initial feature vectors $\{\mathbf{x}_v | v \in V\}$ are obtained from a one-hot encoding of the atomic number, number of bonds linked to each atom, formal charge, chirality (if encoded in the SMILES), number of hydrogens, hybridization and aromaticity of the atom, as well as the atomic mass (divided by 100 for scaling). For each edge e , initial feature vectors $\{\mathbf{e}_{vw} | \{v, w\} \in E\}$ arise from the bond type, whether the bond is conjugated or in a ring, and whether it contains stereochemical information such as a cis/trans double bond. The D-MPNN uses directed edges in a graph to pass information, where each undirected edge (bond) has two corresponding directed edges, one in each direction. Initial directed edge features \mathbf{e}_{vw}^d are obtained via simple concatenation of the atom features of the first atom of a bond \mathbf{x}_v to the respective undirected bond features \mathbf{e}_{vw}

$$\mathbf{e}_{vw}^d = \text{cat}(\mathbf{x}_v, \mathbf{e}_{vw}) \quad (1)$$

where $\text{cat}()$ denotes simple concatenation. The directed edges \mathbf{e}_{vw}^d and \mathbf{e}_{wv}^d are only distinguished by the choice of which atom to use in Eq. 1. Chemprop also offers the option to read in custom atom and bond features in addition to or as a replacement for the default features, as described in *SI* Section S1.2, and thus offers full control of the initial features. In summary, Module 1 of Chemprop constructs atom and directed bond feature vectors \mathbf{x}_v and \mathbf{e}_{vw}^d from the input molecules.

The initial atom and bond features are then passed to a D-MPNN. In a D-MPNN structure, messages are passed between directed edges rather than between nodes as would be done in a traditional MPNN. To construct the hidden directed edge features \mathbf{h}_{vw}^0 of hidden size h , the initial directed edge features \mathbf{e}_{vw}^d are passed through a single neural network layer with learnable weights $\mathbf{W}_i \in \mathbb{R}^{h \times h_i}$

$$\mathbf{h}_{vw}^0 = \tau(\mathbf{W}_i \mathbf{e}_{vw}^d) \quad (2)$$

and a nonlinear activation function τ which can be chosen by the user (default ReLU). The size h of \mathbf{h}_{vw}^0 can be chosen by the user (default 300). The size of \mathbf{e}_{vw}^d , which we term h_i , is set by the lengths of initial feature encodings, per Eq. 1. The directed edge features are then iteratively updated based on the local environment via T (default 3) message passing steps

$$\mathbf{h}_{vw}^{t+1} = \tau(\mathbf{h}_{vw}^0 + \mathbf{W}_h \sum_{k \in \{N(v) \setminus w\}} \mathbf{h}_{kv}^t) \quad (3)$$

until $t+1 = T$, where $\mathbf{W}_h \in \mathbb{R}^{h \times h}$ and $N(v) \setminus w$ denotes the neighbors of node v excluding w . The opposite facing directed edge is excluded from the message passing update for increased numerical stability (see Mahé et al.⁴⁷). Finally, the updated hidden states \mathbf{h}_{vw}^T are aggregated into atomic embeddings via

$$\mathbf{h}_v = \tau(\mathbf{W}_o \mathbf{q}) \quad (4)$$

where \mathbf{q} is a concatenation of the initial atom features \mathbf{x}_v and the sum of all incoming directed edge hidden states

$$\mathbf{q} = \text{cat}(\mathbf{x}_v, \sum_{w \in N(v)} \mathbf{h}_{vw}^T) \quad (5)$$

with $\mathbf{W}_o \in \mathbb{R}^{h \times h_o}$. Here, h_o is the size of \mathbf{q} , i.e. the sum of the hidden size h and the size of \mathbf{x}_v . In summary, in Module 2, the D-MPNN weights \mathbf{W}_i , \mathbf{W}_h , and \mathbf{W}_o are learned from the training data, outputting learnable atomic embeddings \mathbf{h}_v . Customizations and hyperparameter tuning include the choice of activation function τ , the hidden size h , and the number of message passing steps T . Chemprop offers the option to add bias terms to all neural network layers (defaults to False).

The atomic embeddings \mathbf{h}_v of all atoms in a molecule are then aggregated into a single molecular embedding \mathbf{h}_m via

$$\mathbf{h}_m = \text{cat}(\mathbf{h}'_m, \mathbf{x}_m) \quad (6)$$

with

$$\mathbf{h}'_m = \sum_{v \in V} \mathbf{h}_v \quad (7)$$

where \mathbf{x}_m is an optional vector of additional molecular features. Chemprop offers three aggregation options: summation (as shown in Eq. (7)), a scaled sum (divided by a user specified scaler, called norm within Chemprop), or an average (the default aggregation). Schweidtmann et al.⁴⁸ compared the performance of such aggregation functions for different datasets. The optional additional molecular features, \mathbf{x}_m , may be provided features from outside sources (*SI* Section S1.2) or generated engineered fingerprints (Morgan circular fingerprints⁴⁹ and RDKit 2D fingerprints⁴⁶ are implemented in Chemprop). By default, \mathbf{x}_m is empty such that the molecular embedding is simply an aggregation over atomic embeddings, i.e. $\mathbf{h}_m = \mathbf{h}'_m$. In summary, Module 3 produces molecular embeddings \mathbf{h}_m of length h plus the size of \mathbf{x}_m . Chemprop offers the option to circumvent Modules 1-3 and only using \mathbf{x}_m as fixed molecular embedding, so that $\mathbf{h}_m = \mathbf{x}_m$.

Finally, in the last module, molecular target properties are learned from the molecular embeddings \mathbf{h}_m via a feed-forward neural network, where the number of layers (default 2) and the number of hidden neurons (default 300) can be chosen by the user. The number of input neurons is set by the length of \mathbf{h}_m , and the number of output neurons by the number of targets. The activation function between linear layers is set to be the same as in the D-MPNN, and bias is turned on per default. For binary classification tasks, the final model output is passed through a sigmoid function to constrain values to the range (0,1). For multiclass classification, the final model output is transformed with a softmax function, such that the classification scores sum to 1 across classes.

Chemprop is fully end-to-end trainable, so that the weights for the D-MPNN and FFN are updated simultaneously. Users have the option to train models using cross-validation and ensembles of submodels. By default, a single model is trained on a random data split for 30 epochs. We note that small datasets need a much larger number of epochs to train, and advise to check for convergence of the learning curve. Chemprop uses the Adam optimizer.⁵⁰ The default learning rate schedule increases the learning rate linearly from 10^{-4} to 10^{-3} for the first two warmup epochs and then decreases learning rate exponentially from 10^{-3} to

Table 1: Selected published studies based on Chemprop.

Ref.	Year	Prediction	Ref.	Year	Prediction
10	2020	Growth inhibitory activity against <i>E. coli</i> ; led to an identification of a potential new drug	51	2022	Absorption, distribution, metabolism, excretion (ADME) properties for drug discovery
52	2021	Chemical synergy against SARS-CoV-2; identified two drug combinations with antiviral synergy in vitro	53	2023	Growth inhibitory activity against <i>A. baumannii</i> ; led to an identification of a potential new drug
28	2021	IR spectra of molecules	54	2023	Fuel properties
55	2021	Atomic charges, Fukui indices, NMR constants, bond lengths, and bond orders	56	2023	Critical properties, acentric factor, and phase change properties
57	2021	Lipophilicity	58	2023	Molecular optical peaks
59	2022	Reaction rates and barrier heights	60	2023	Lipophilicity
61	2022	Barrier heights of reactions	62	2023	Solvent effects on reaction rate constants
29	2022	Molecular optical peaks	63	2023	Vapor pressure in the low volatility regime
64	2022	Solvation free energy, solvation enthalpy, and Abraham solute descriptors	65	2023	Molecular optical peaks and partition coefficients for closed-loop active learning
66	2022	Solid solubility of organic solutes in water and organic solvents	67	2023	Senolytic activity of compounds to selectively target senescent cells
68	2022	Activity coefficients	69	2023	Toxicity measurements using 12 nuclear receptor signaling and stress response pathways

10^{-4} for the remaining epochs. By default, a batch size of 50 data points is used for each optimizer step. Early stopping and dropout are available as means of regularization. The PyTorch backend of Chemprop enables seamless GPU acceleration of both model training and inference. The acceleration of training and inference processes when used with a GPU can be significant, as shown in *SI* Section S3.3.

3 Discussion of Features

Table 1 lists a non-exhaustive selection of studies based on Chemprop, showcasing its versatility and applicability for the prediction of a large variety of chemical properties, but also its ease of use. Models can be trained and tested with a single line on the command line (or a few lines of python code) and a user-supplied CSV file (See *SI* for some examples). In the following, we discuss specialty options introduced since its first release.

3.1 Additional features

Chemprop can take additional features at the molecule-, atom-, or bond-level as input. While Chemprop often generates accurate models without requiring any input beyond the SMILES,

it has been shown that outside information added as additional features can further improve performance.^{36,64} Users can provide their custom additional features by adding keywords and paths to the data files containing the features. See *SI* for command-line arguments and details.

3.2 Multi-molecule models

Chemprop can also train on dataset containing more than one molecule as input. For example, when properties related to solvation need to be predicted both a solute and a solvent are required as input to the model. Users can provide multiple molecules as input to Chemprop. When multiple molecules are used, by default Chemprop trains a separate D-MPNN for each molecule (Figure S1a). If the option `--mpn_shared` is specified, the same D-MPNN is used for all molecules (Figure S1b). The embeddings of the different molecules are then concatenated prior to the FFN. Note that the current implementation of multiple molecules in Chemprop does not ensure permutational invariance towards the input molecules. This is suited to situations where the input molecules have different roles, e.g. molecule 1 = solute, molecule 2 = solvent. For additional input information and a figure depicting the multi-molecule model structure, see *SI*.

3.3 Reaction support

Chemprop supports the input of atom-mapped reactions, i.e. pairs of reactants and products SMILES connected via the “>” symbol by using the keyword `--reaction`. The pair of reactants and products is transformed into a single pseudomolecule, namely the condensed graph of reaction (CGR), and then passed to a regular D-MPNN block. The construction of a CGR within Chemprop is described in detail in Ref. 59, and summarized in the following. In general, the input of a reaction vs. a molecule only affects the setup of the graph object and its initial features, but not any other part of the architecture. The graph of a reaction has a different set of edges E as the graph of a molecule, as shown in Fig. 2 for an example

Reaction SMILES: [CH2:1]=[CH1:2][CH1:3]=[CH2:4].[CH2:5]=[CH1:6][CH3:7]
>>[CH2:1]1[CH1:2]=[CH1:3][CH2:4][CH2:5][CH1:6]1[CH3:7]




Figure 2: Construction of the condensed graph of reaction (CGR) of an example reaction. The vertices and edges are obtained as the union of the respective reactant and product vertices and edges. The features are obtained as a combination of the reactant (white background) and product (gray background) features for atoms and bonds.

Diels-Alder reaction. To build the CGR pseudomolecule, the set of atoms is obtained as the union of the sets of atoms in the reactants and products. Similarly, the set of bonds is obtained as the union of the sets of bonds in the reactants and products. Once constructed, the CGR is passed through the D-MPNN and other model architecture components in the same way as a molecular graph would. Optionally, Chemprop accepts an additional molecule object as input, such as a solvent, or a reagent, etc. which is passed to its own D-MPNN similar to the multi-molecule model. The output of the reaction D-MPNN and molecule D-MPNN is concatenated after atomic aggregation, before the FFN. This option is available via the `reaction_solvent` keyword. See *SI* for further commandline options and details.

3.4 Spectra data support

Chemprop supports prediction of whole-spectrum properties for molecules. An initial version of this capability was discussed in Ref. 28 for use with IR absorbance spectra. Targets for the spectra dataset type are composed of an array of intensity values, set at fixed bin locations typically specified in terms of wavelengths, frequencies, or wavenumbers. Spectral Information Divergence was originally developed as a method of comparing spectra to reference databases⁷⁰ and is adapted in Chemprop to be used as a loss function, considering the deviation of the spectrum as a whole rather than independently at each bin location. The treatment of spectra can handle targets with gaps or missing values within a dataset. With the expectation that spectra will often be collected in systems where a portion of the range will be obscured or invalid (e.g., from solvent absorbance), Chemprop can create exclusion regions in specified spectra where no predictions are provided and targets are ignored for training purposes.

3.5 Latent representations

Graph neural networks enable learning both molecular representation and property end-to-end directly from the molecular graph. As detailed above in Section 2, the learned node representations are aggregated into a molecule-level representation after the message-passing phase, which we refer to as the “learned fingerprint.” This embedding is then further fed into the FFN network. Within the FFN, we consider the final hidden representation, which we refer to as the “ffn embedding”. Both of these vectors are latent representations of a molecule, as it relates to a particular trained model. Molecule latent representations can be useful for data clustering or use as additional features in other models. Chemprop supports the calculation of either from a trained model for a given set of molecules.

3.6 Loss function options

Chemprop can train models according to many common loss functions. The loss functions available for a given task are determined by the dataset type (regression, classification, multi-class, or spectra). Regression models can be trained with mean squared error (MSE), bounded MSE (which allows inequalities as targets), or negative log-likelihood (NLL) based on prediction uncertainty distributions consistent with mean-variance estimation (MVE)⁷¹ or evidential uncertainty⁷². Classification tasks default to the binary cross entropy loss and have additional options of Matthews correlation coefficient (MCC) and Dirichlet (evidential classification)⁷³. Cross entropy and MCC are also available for multi-class problems. There are two options available for training on spectra: spectral information divergence (SID)⁷⁰ and first-order Wasserstein distance (a.k.a. earthmover’s distance).⁷⁴ Loss functions must be differentiable since they are used to calculate gradients that update the model parameters, but Chemprop also provides the option to use several non-differentiable metrics for model evaluation.

3.7 Transfer learning

Transfer learning is a general strategy of using information gained through the training of one model to inform and improve the training of a related model. Often, this strategy is used to transfer information from a previously trained model of a large dataset to a model of a small dataset in order to improve the performance of the model of the small dataset. The simplest method of transfer learning would be taking predictions or latent representations from one model and supplying them as additional features to another model (Sections 3.1 and 3.5).

In Chemprop, different strategies are available to transfer learned model parameters from a previously trained model to a new model as a form of transfer learning, as shown in Figure 3. A pretrained model may be used to initialize a new model, with normal updating of the transferred weights in training. Alternatively, parameters from the transferred model




Figure 3: Options to transfer model parameters from a pre-trained model (squared) to a new model, by (a) initializing the new model parameters or by freezing the (b) D-MPNN layers and (c) n FFN layers

can be frozen, holding it constant during training. Freezing parameters always includes the D-MPNN weights but can be specified to include some FFN layers as well. See *SI* for the related corresponding arguments.

3.8 Hyperparameter optimization

Chemprop provides a command-line utility, allowing for a simple initiation of hyperparameter optimization jobs. Options for the hyperparameter job such as how many trials to carry out and which hyperparameters to include in the search (options in Table S3) can be specified with simple command-line arguments. The optimization initially uses randomly sampled trials, followed by targeted sampling using the Tree-structured Parzen Estimator

algorithm.^{75,76}

Hyperparameter optimization is often the most resource-intensive step in model training. In order to search a large parameter space adequately, a large number of trials is needed. Chemprop allows for parallel operation of multiple hyperparameter optimization instances, removing the need to carry out all trials in series and reducing the wall time needed to perform the optimization significantly. See *SI* for the available hyperparameter options and other details.

3.9 Uncertainty tools

Chemprop includes a variety of popular uncertainty estimation, calibration, and evaluation tools. The estimation methods include deep ensembles⁷⁷, dropout⁷⁸, mean-variance estimation (MVE)⁷¹, and evidential⁷², as well as a special version of ensemble variance for spectral predictions²⁸, and the inherently probabilistic outputs of classification models. After estimating the uncertainty in a model's predictions, it is often helpful to calibrate these uncertainties to improve their performance on new predictions. We provide four such methods for regression tasks (z -scaling⁷⁹, t -scaling, Zelikman's CRUDE,⁸⁰ and MVE weighting⁸¹) and two for classification (single-parameter Platt scaling⁸² and isotonic regression⁸³). In addition to standard metrics such as RMSE, MAE, etc. for evaluating predictions, Chemprop also includes several metrics specifically for evaluating the quality of uncertainty estimates. These include negative log likelihood, Spearman rank correlation, expected normalized calibration error (ENCE),⁸⁴ and miscalibration area.⁸⁴ Any valid classification or multiclass metric used to assess predictions can also be used to assess uncertainties.

3.10 Atom/bond-level targets

Chemprop supports a multitask constrained D-MPNN architecture for predicting atom- and bond-level properties, such as charge density or bond length. This model enables a D-MPNN to be trained on multiple atomic and bond properties simultaneously, though unlike

molecular property targets they do not share a single FFN. Optionally, an attention-based constraining method may be used to enforce that predicted atomic or bond properties sum to a specified molecular net value, such as the overall charge of a molecule. An initial, more limited version of this capability was developed in Ref. 55. For details on the input formats to be used for atom/bond targets, both for training and for inference, see *SI*.

4 Benchmarking

See *SI* for general performance benchmarks, benchmarks using specific Chemprop features (atom/bond-level targets, reaction support, multi-molecule models, spectra prediction, and uncertainty estimation), and system timing benchmarks.

5 Conclusion

We have presented the software package Chemprop, a powerful toolbox for machine learning of chemical properties of molecules and reactions. Significant improvements have been made to the software since its initial release and study,³⁶ including the support of multi-molecule properties, reactions, atom/bond-level properties, and spectra. Additionally, several state-of-the-art approaches to estimate the uncertainty of predictions have been incorporated, as well as pretraining and transfer learning procedures. Furthermore, the code now offers a variety of customization options, such as custom atom and bond features, a large variety of loss functions, and saving the learned feature embeddings for subsequent use with different algorithms. We have showcased and benchmarked Chemprop on a variety of example tasks and datasets, and have found competitive performances for molecular property prediction compared to other approaches available on public leaderboards. In summary, Chemprop is a powerful, fast and convenient tool to learn conformation-independent properties of molecules, sets of molecules, or reactions.

Data and Software Availability

Chemprop, including all features described in this manuscript, is available under the open-source MIT License on GitHub, github.com/chemprop/chemprop. An extensive documentation including tutorials is available online,⁸⁵ including a workshop on YouTube.⁸⁶ Scripts and data splits to fully reproduce this study are available on GitHub, github.com/chemprop/chemprop_benchmark and on Zenodo, doi.org/10.5281/zenodo.8174267 respectively.

Acknowledgement

E.H., Y.C., F.H.V., and W.H.G. acknowledge support from the Machine Learning for Pharmaceutical Discovery and Synthesis Consortium (MLPDS). K.P.G., S.L., F.H.V., H.W., W.H.G., and C.J.M. acknowledge support from the DARPA Accelerated Molecular Discovery (AMD) program (DARPA HR00111920025). E.H. acknowledges support from the Austrian Science Fund (FWF), project J-4415. K.P.G. was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1745302. D.E.G. acknowledges support from the MIT IBM Watson AI Lab. F.H.V. would like to acknowledge the KU Leuven Internal Starting Grant (STG/22/032). C.J.M. would like to acknowledge support from VCU Startup Funding. Parts of the data reported within this paper were generated with resources from the MIT SuperCloud Lincoln Laboratory Supercomputing Center.⁸⁷

Supporting Information

Supporting Information Available: additional software details and usage examples, dataset and data handling details for benchmarking, results of software benchmarks (general performance, feature demonstrations, timing).

References

- (1) Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; Dahl, G. E. Neural Message Passing for Quantum Chemistry. *Proceedings of the International Conference on Machine Learning* **2017**, 1263–1272.
- (2) Gasteiger, J.; Groß, J.; Günnemann, S. Directional Message Passing for Molecular Graphs. *Proceedings of the International Conference on Learning Representations* **2020**, arXiv:2003.03123.
- (3) Zhang, S.; Liu, Y.; Xie, L. Molecular Mechanics-Driven Graph Neural Network with Multiplex Graph for Molecular Structures. *Machine Learning for Molecules Workshop at NeurIPS* **2020**, arXiv:2011.07457.
- (4) Vermeire, F. H.; Chung, Y.; Green, W. H. Predicting Solubility Limits of Organic Solutes for a Wide Range of Solvents and Temperatures. *J. Am. Chem. Soc.* **2022**, *144*, 10785–10797.
- (5) Dobbelaere, M. R.; Ureel, Y.; Vermeire, F. H.; Tomme, L.; Stevens, C. V.; Van Geem, K. M. Machine Learning for Physicochemical Property Prediction of Complex Hydrocarbon Mixtures. *Ind. Eng. Chem. Res.* **2022**, *61*, 8581–8594.
- (6) Dobbelaere, M. R.; Plehiers, P. P.; Van de Vijver, R.; Stevens, C. V.; Van Geem, K. M. Learning Molecular Representations for Thermochemistry Prediction of Cyclic Hydrocarbons and Oxygenates. *J. Phys. Chem. A* **2021**, *125*, 5166–5179.
- (7) Rittig, J. G.; Ben Hicham, K.; Schweidtmann, A. M.; Dahmen, M.; Mitsos, A. Graph Neural Networks for Temperature-Dependent Activity Coefficient Prediction of Solutes in Ionic Liquids. *Comput. Chem. Eng.* **2023**, *171*, 108153.
- (8) Rittig, J. G.; Ritzert, M.; Schweidtmann, A. M.; Winkler, S.; Weber, J. M.; Morsch, P.;

- Heufer, K. A.; Grohe, M.; Mitsos, A.; Dahmen, M. Graph Machine Learning for Design of High-Octane Fuels. *AIChE J.* **2023**, *69*, e17971.
- (9) Fleitmann, L.; Ackermann, P.; Schilling, J.; Kleinekorte, J.; Rittig, J. G.; vom Lehn, F.; Schweidtmann, A. M.; Pitsch, H.; Leonhard, K.; Mitsos, A.; Bardow, A.; Dahmen, M. Molecular Design of Fuels for Maximum Spark-Ignition Engine Efficiency by Combining Predictive Thermodynamics and Machine Learning. *Energ. Fuel.* **2023**, *37*, 2213–2229.
- (10) Stokes, J. M.; Yang, K.; Swanson, K.; Jin, W.; Cubillos-Ruiz, A.; Donghia, N. M.; MacNair, C. R.; French, S.; Carfrae, L. A.; Bloom-Ackermann, Z., et al. A Deep Learning Approach to Antibiotic Discovery. *Cell* **2020**, *180*, 688–702.
- (11) De Cao, N.; Kipf, T. MolGAN: An Implicit Generative Model for Small Molecular Graphs. *arXiv Preprint* **2022**, arXiv:1805.11973.
- (12) Dan, Y.; Zhao, Y.; Li, X.; Li, S.; Hu, M.; Hu, J. Generative Adversarial Networks (GAN) Based Efficient Sampling of Chemical Composition Space for Inverse Design of Inorganic Materials. *Npj Comput. Mater.* **2020**, *6*, 1–7.
- (13) Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; Aspuru-Guzik, A. Automatic Chemical Design using a Data-Driven Continuous Representation of Molecules. *ACS Cent. Sci.* **2018**, *4*, 268–276.
- (14) Wei, J. N.; Duvenaud, D.; Aspuru-Guzik, A. Neural Networks for the Prediction of Organic Chemistry Reactions. *ACS Cent. Sci.* **2016**, *2*, 725–732.
- (15) Segler, M. H.; Waller, M. P. Neural-Symbolic Machine Learning for Retrosynthesis and Reaction Prediction. *Chem. Eur. J.* **2017**, *23*, 5966–5971.
- (16) Coley, C. W.; Barzilay, R.; Jaakkola, T. S.; Green, W. H.; Jensen, K. F. Prediction of Organic Reaction Outcomes using Machine Learning. *ACS Cent. Sci.* **2017**, *3*, 434–443.

- (17) Coley, C. W.; Green, W. H.; Jensen, K. F. Machine Learning in Computer-Aided Synthesis Planning. *Acc. Chem. Res.* **2018**, *51*, 1281–1289.
- (18) Szymkuć, S.; Gajewska, E. P.; Klucznik, T.; Molga, K.; Dittwald, P.; Startek, M.; Bajczyk, M.; Grzybowski, B. A. Computer-Assisted Synthetic Planning: The End of the Beginning. *Angew. Chem. Int. Ed.* **2016**, *55*, 5904–5937.
- (19) Segler, M. H.; Preuss, M.; Waller, M. P. Planning Chemical Syntheses with Deep Neural Networks and Symbolic AI. *Nature* **2018**, *555*, 604–610.
- (20) Kayala, M. A.; Azencott, C.-A.; Chen, J. H.; Baldi, P. Learning to Predict Chemical Reactions. *J. Chem. Inf. Model.* **2011**, *51*, 2209–2222.
- (21) Kayala, M. A.; Baldi, P. ReactionPredictor: Prediction of Complex Chemical Reactions at the Mechanistic Level using Machine Learning. *J. Chem. Inf. Model.* **2012**, *52*, 2526–2540.
- (22) Fooshee, D.; Mood, A.; Gutman, E.; Tavakoli, M.; Urban, G.; Liu, F.; Huynh, N.; Van Vranken, D.; Baldi, P. Deep Learning for Chemical Reaction Prediction. *Mol. Syst. Des. Eng.* **2018**, *3*, 442–452.
- (23) Bradshaw, J.; Kusner, M. J.; Paige, B.; Segler, M. H.; Hernández-Lobato, J. M. A Generative Model for Electron Paths. *Proceedings of the International Conference on Learning Representations* **2019**, arXiv:1805.10970.
- (24) Bi, H.; Wang, H.; Shi, C.; Coley, C.; Tang, J.; Guo, H. Non-Autoregressive Electron Redistribution Modeling for Reaction Prediction. *Proceedings of the International Conference on Machine Learning* **2021**, 904–913.
- (25) Coley, C. W.; Jin, W.; Rogers, L.; Jamison, T. F.; Jaakkola, T. S.; Green, W. H.; Barzilay, R.; Jensen, K. F. A Graph-Convolutional Neural Network Model for the Prediction of Chemical Reactivity. *Chem. Sci.* **2019**, *10*, 370–377.

- (26) Sacha, M.; Błaz, M.; Byrski, P.; Dabrowski-Tumanski, P.; Chrominski, M.; Loska, R.; Włodarczyk-Pruszyński, P.; Jastrzebski, S. Molecule Edit Graph Attention Network: Modeling Chemical Reactions as Sequences of Graph Edits. *J. Chem. Inf. Model.* **2021**, *61*, 3273–3284.
- (27) Schwaller, P.; Gaudin, T.; Lanyi, D.; Bekas, C.; Laino, T. “Found in Translation”: Predicting Outcomes of Complex Organic Chemistry Reactions using Neural Sequence-to-Sequence Models. *Chem. Sci.* **2018**, *9*, 6091–6098.
- (28) McGill, C.; Forsuelo, M.; Guan, Y.; Green, W. H. Predicting Infrared Spectra with Message Passing Neural Networks. *J. Chem. Inf. Model.* **2021**, *61*, 2594–2609.
- (29) Greenman, K. P.; Green, W. H.; Gómez-Bombarelli, R. Multi-Fidelity Prediction of Molecular Optical Peaks with Deep Learning. *Chem. Sci.* **2022**, *13*, 1152–1162.
- (30) Dührkop, K. Deep Kernel Learning Improves Molecular Fingerprint Prediction from Tandem Mass Spectra. *Bioinf.* **2022**, *38*, i342–i349.
- (31) Nguyen, D. H.; Nguyen, C. H.; Mamitsuka, H. ADAPTIVE: leArning DAta-dePendenT, concIse molecular VEctors for Fast, Accurate Metabolite Identification from Tandem Mass Spectra. *Bioinf.* **2019**, *35*, i164–i172.
- (32) Nguyen, D. H.; Nguyen, C. H.; Mamitsuka, H. Recent Advances and Prospects of Computational Methods for Metabolite Identification: A Review with Emphasis on Machine Learning Approaches. *Brief. Bioinf.* **2019**, *20*, 2028–2043.
- (33) Stravs, M. A.; Dührkop, K.; Böcker, S.; Zamboni, N. MSNovelist: De Novo Structure Generation from Mass Spectra. *Nat. Methods* **2022**, *19*, 865–870.
- (34) Muratov, E. N.; Bajorath, J.; Sheridan, R. P.; Tetko, I. V.; Filimonov, D.; Poroikov, V.; Oprea, T. I.; Baskin, I. I.; Varnek, A.; Roitberg, A., et al. QSAR without Borders. *Chem. Soc. Rev.* **2020**, *49*, 3525–3564.

- (35) Kearnes, S.; McCloskey, K.; Berndl, M.; Pande, V.; Riley, P. Molecular Graph Convolutions: Moving Beyond Fingerprints. *J. Comput. Aided Mol. Des.* **2016**, *30*, 595–608.
- (36) Yang, K.; Swanson, K.; Jin, W.; Coley, C.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M.; Palmer, A.; Settels, V.; Jaakkola, T.; Jensen, K.; Barzilay, R. Analyzing Learned Molecular Representations for Property Prediction. *J. Chem. Inf. Model.* **2019**, *59*, 3370–3388.
- (37) Maziarka, Ł.; Danel, T.; Mucha, S.; Rataj, K.; Tabor, J.; Jastrzebski, S. Molecule Attention Transformer. *arXiv Preprint* **2020**, arXiv:2002.08264.
- (38) Kreuzer, D.; Beaini, D.; Hamilton, W.; Létourneau, V.; Tossou, P. Rethinking Graph Transformers with Spectral Attention. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 21618–21629.
- (39) Schütt, K. T.; Sauceda, H. E.; Kindermans, P.-J.; Tkatchenko, A.; Müller, K.-R. SchNet—A Deep Learning Architecture for Molecules and Materials. *J. Chem. Phys.* **2018**, *148*, 241722.
- (40) Unke, O. T.; Meuwly, M. PhysNet: A Neural Network for Predicting Energies, Forces, Dipole Moments, and Partial Charges. *J. Chem. Theory Comput.* **2019**, *15*, 3678–3693.
- (41) Bigi, F.; Pozdnyakov, S. N.; Ceriotti, M. Wigner Kernels: Body-Ordered Equivariant Machine Learning without a Basis. *arXiv Preprint* **2023**, arXiv:2303.04124.
- (42) Winter, B.; Winter, C.; Schilling, J.; Bardow, A. A Smile is All you Need: Predicting Limiting Activity Coefficients from SMILES with Natural Language Processing. *Digital Discovery* **2022**, *1*, 859–869.
- (43) Bagal, V.; Aggarwal, R.; Vinod, P.; Priyakumar, U. D. MolGPT: Molecular Generation using a Transformer-Decoder Model. *J. Chem. Inf. Model.* **2021**, *62*, 2064–2076.

- (44) Honda, S.; Shi, S.; Ueda, H. R. Smiles Transformer: Pre-trained Molecular Fingerprint for Low Data Drug Discovery. *arXiv Preprint* **2019**, arXiv:1911.04738.
- (45) Chithrananda, S.; Grand, G.; Ramsundar, B. Chemberta: Large-Scale Self-Supervised Pretraining for Molecular Property Prediction. *Machine Learning for Molecules Workshop at NeurIPS* **2020**, arXiv:2010.09885.
- (46) Landrum, G. RDKit: Open-Source Cheminformatics. 2006; <https://www.rdkit.org/>.
- (47) Mahé, P.; Ueda, N.; Akutsu, T.; Perret, J.-L.; Vert, J.-P. Extensions of marginalized graph kernels. *Proceedings of the International Conference on Machine Learning* **2004**, 70.
- (48) Schweidtmann, A. M.; Rittig, J. G.; Weber, J. M.; Grohe, M.; Dahmen, M.; Leonhard, K.; Mitsos, A. Physical Pooling Functions in Graph Neural Networks for Molecular Property Prediction. *Comput. Chem. Eng.* **2023**, *172*, 108202.
- (49) Morgan, H. L. The Generation of a Unique Machine Description for Chemical Structures – A Technique Developed at Chemical Abstracts Service. *J. Chem. Doc.* **1965**, *5*, 107–113.
- (50) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. *Proceedings of the International Conference on Learning Representations* **2017**, arXiv:1412.6980.
- (51) Lim, M. A.; Yang, S.; Mai, H.; Cheng, A. C. Exploring Deep Learning of Quantum Chemical Properties for Absorption, Distribution, Metabolism, and Excretion Predictions. *J. Chem. Inf. Model.* **2022**, *62*, 6336–6341.
- (52) Jin, W.; Stokes, J. M.; Eastman, R. T.; Itkin, Z.; Zakharov, A. V.; Collins, J. J.; Jaakkola, T. S.; Barzilay, R. Deep Learning Identifies Synergistic Drug Combinations for Treating COVID-19. *Proc. Natl. Acad. Sci. U.S.A.* **2021**, *118*.

- (53) Liu, G.; Catacutan, D. B.; Rathod, K.; Swanson, K.; Jin, W.; Mohammed, J. C.; Chiappino-Pepe, A.; Syed, S. A.; Fragis, M.; Rachwalski, K., et al. Deep Learning-Guided Discovery of an Antibiotic Targeting *Acinetobacter Baumannii*. *Nat. Chem. Biol.* **2023**, *19*, 1342–1350.
- (54) Larsson, T.; Vermeire, F.; Verhelst, S. Machine Learning for Fuel Property Predictions: A Multi-Task and Transfer Learning Approach. *SAE Tech. Pap.* **2023**, 2023-01-0337.
- (55) Guan, Y.; Coley, C. W.; Wu, H.; Ranasinghe, D.; Heid, E.; Struble, T. J.; Pattanaik, L.; Green, W. H.; Jensen, K. F. Regio-Selectivity Prediction with a Machine-Learned Reaction Representation and On-the-Fly Quantum Mechanical Descriptors. *Chem. Sci.* **2021**, *12*, 2198–2208.
- (56) Biswas, S.; Chung, Y.; Ramirez, J.; Wu, H.; Green, W. H. Predicting Critical Properties and Acentric Factor of Fluids using Multi-Task Machine Learning. *J. Chem. Inf. Model.* **2023**, *63*, 4574–4588.
- (57) Lenselink, E. B.; Stouten, P. F. W. Multitask machine learning models for predicting lipophilicity (logP) in the SAMPL7 challenge. *Journal of Computer-Aided Molecular Design* **2021**, *35*, 901–909.
- (58) McNaughton, A. D.; Joshi, R. P.; Knutson, C. R.; Fnu, A.; Luebke, K. J.; Mallerich, J. P.; Madrid, P. B.; Kumar, N. Machine Learning Models for Predicting Molecular UV–Vis Spectra with Quantum Mechanical Properties. *J. Chem. Inf. Model.* **2023**, *63*, 1462–1471.
- (59) Heid, E.; Green, W. H. Machine Learning of Reaction Properties via Learned Representations of the Condensed Graph of Reaction. *J. Chem. Inf. Model.* **2022**, *62*, 2101–2110.
- (60) Isert, C.; Kromann, J. C.; Stiefl, N.; Schneider, G.; Lewis, R. A. Machine Learning for

- Fast, Quantum Mechanics-Based Approximation of Drug Lipophilicity. *ACS Omega* **2023**, *8*, 2046–2056.
- (61) Spiekermann, K. A.; Pattanaik, L.; Green, W. H. Fast Predictions of Reaction Barrier Heights: Toward Coupled-Cluster Accuracy. *J. Phys. Chem. A* **2022**, *126*, 3976–3986.
- (62) Chung, Y.; Green, W. H. Machine Learning from Quantum Chemistry to Predict Experimental Solvent Effects on Reaction Rates. *ChemRxiv Preprint* **2023**, ChemRxiv.
- (63) Lansford, J. L.; Jensen, K. F.; Barnes, B. C. Physics-informed Transfer Learning for Out-of-sample Vapor Pressure Predictions. *Propellants, Explosives, Pyrotechnics* **2023**, *48*, e202200265.
- (64) Chung, Y.; Vermeire, F. H.; Wu, H.; Walker, P. J.; Abraham, M. H.; Green, W. H. Group Contribution and Machine Learning Approaches to Predict Abraham Solute Parameters, Solvation Free Energy, and Solvation Enthalpy. *J. Chem. Inf. Model.* **2022**, *62*, 433–446.
- (65) Koscher, B.; Canty, R. B.; McDonald, M. A.; Greenman, K. P.; McGill, C. J.; Bilodeau, C. L.; Jin, W.; Wu, H.; Vermeire, F. H.; Jin, B., et al. Autonomous, Multi-Property-Driven Molecular Discovery: From Predictions to Measurements and Back. *ChemRxiv Preprint* **2023**, ChemRxiv.
- (66) Vermeire, F. H.; Chung, Y.; Green, W. H. Predicting Solubility Limits of Organic Solutes for a Wide Range of Solvents and Temperatures. *J. Am. Chem. Soc.* **2022**, *144*, 10785–10797.
- (67) Wong, F.; Omori, S.; Donghia, N. M.; Zheng, E. J.; Collins, J. J. Discovering Small-Molecule Senolytics with Deep Neural Networks. *Nature Aging* **2023**, *3*, 734–750.
- (68) Felton, K. C.; Ben-Safar, H.; Alexei, A. DeepGamma: A Deep Learning Model for

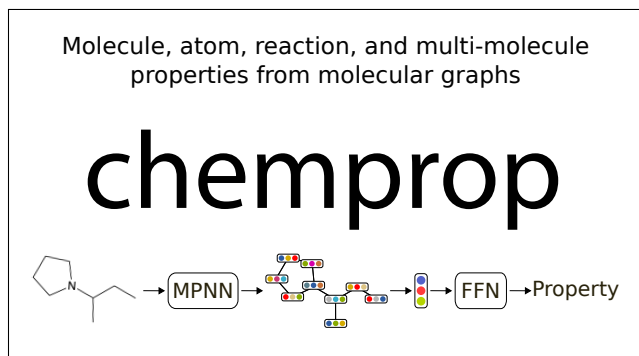
Activity Coefficient Prediction. *1st Annual AAAI Workshop on AI to Accelerate Science and Engineering (AI2ASE)* **2022**, AI2ASE.

- (69) Colomba, M.; Benedetti, S.; Fraternali, D.; Guidarelli, A.; Coppari, S.; Freschi, V.; Crinelli, R.; Kass, G. E. N.; Gorassini, A.; Verardo, G.; Roselli, C.; Meli, M. A.; Di Giacomo, B.; Albertini, M. C. Nrf2-Mediated Pathway Activated by Prunus spinosa L. (Rosaceae) Fruit Extract: Bioinformatics Analyses and Experimental Validation. *Nutrients* **2023**, *15*, 2132.
- (70) Chang, C.-I. An Information-Theoretic Approach to Spectral Variability, Similarity, and Discrimination for Hyperspectral Image Analysis. *IEEE Transactions on Information Theory* **2000**, *46*, 1927–1932.
- (71) Nix, D. A.; Weigend, A. S. Estimating the Mean and Variance of the Target Probability Distribution. *Proceedings of IEEE International Conference on Neural Networks* **1994**, *1*, 55–60.
- (72) Amini, A.; Schwarting, W.; Soleimany, A.; Rus, D. Deep Evidential Regression. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 14927–14937.
- (73) Sensoy, M.; Kaplan, L.; Kandemir, M. Evidential Deep Learning to Quantify Classification Uncertainty. *Adv. Neural Inf. Process. Syst.* **2018**, *31*.
- (74) Villani, C. *Optimal transport: old and new*; Springer, 2009; Vol. 338.
- (75) Bergstra, J.; Bardenet, R.; Bengio, Y.; Kégl, B. Algorithms for Hyper-Parameter Optimization. *Adv. Neural Inf. Process. Syst.* **2011**, *24*.
- (76) Bergstra, J.; Yamins, D.; Cox, D. Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures. *Proceedings of the International Conference on Machine Learning* **2013**, 115–123.

- (77) Lakshminarayanan, B.; Pritzel, A.; Blundell, C. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. *Adv. Neural Inf. Process. Syst.* **2017**, *30*.
- (78) Gal, Y.; Ghahramani, Z. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. *Proceedings of the International Conference on Machine Learning* **2016**, *48*, 1050–1059.
- (79) Levi, D.; Gispan, L.; Giladi, N.; Fetaya, E. Evaluating and calibrating uncertainty prediction in regression tasks. *Sensors* **2022**, *22*, 5540.
- (80) Zelikman, E.; Healy, C.; Zhou, S.; Avati, A. CRUDE: Calibrating Regression Uncertainty Distributions Empirically. *arXiv Preprint* **2020**, arXiv:2005.12496.
- (81) Wang, D.; Yu, J.; Chen, L.; Li, X.; Jiang, H.; Chen, K.; Zheng, M.; Luo, X. A Hybrid Framework for Improving Uncertainty Quantification in Deep Learning-Based QSAR Regression Modeling. *J. Cheminf.* **2021**, *13*, 1–17.
- (82) Guo, C.; Pleiss, G.; Sun, Y.; Weinberger, K. Q. On Calibration of Modern Neural Networks. *Proceedings of the International Conference on Machine Learning* **2017**, 1321–1330.
- (83) Zadrozny, B.; Elkan, C. Transforming Classifier Scores into Accurate Multiclass Probability Estimates. *Proceedings of the International Conference on Knowledge Discovery and Data Mining* **2002**, 694–699.
- (84) 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. *J. Chem. Inf. Model.* **2020**, *60*, 2697–2717.
- (85) Chemprop. <https://chemprop.readthedocs.io/en/latest/>, (accessed April 6 2023).

- (86) Chemprop Workshop. <https://www.youtube.com/watch?v=Te015E8Wo2M>, (accessed April 6 2023).
- (87) 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. *Proceedings of the IEEE High Performance Extreme Computing Conference* **2018**, 1–6.

TOC Graphic



Supporting Information

Chemprop: A Machine Learning Package for Chemical Property Prediction

Esther Heid,^{†,‡} Kevin P. Greenman,[†] Yunsie Chung,[†] Shih-Cheng Li,^{†,¶} David E.
Graff,^{†,§} Florence H. Vermeire,^{†,||} Haoyang Wu,[†] William H. Green,[†] and Charles J.
McGill^{*,†,⊥}

[†]*Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge,
Massachusetts 02139, United States*

[‡]*Institute of Materials Chemistry, TU Wien, 1060 Vienna, Austria*

[¶]*Department of Chemical Engineering, National Taiwan University, Taipei 10617, Taiwan*

[§]*Department of Chemistry and Chemical Biology, Harvard University, Cambridge,
Massachusetts 02138, United States*

^{||}*Department of Chemical Engineering, KU Leuven, Celestijnenlaan 200F, B-3001 Leuven,
Belgium*

[⊥]*Department of Chemical and Life Science Engineering, Virginia Commonwealth
University, Richmond, Virginia 23284, United States*

E-mail: mcgillc2@vcu.edu

Contents

1	Additional Model Details	4
1.1	Example commands	4
1.2	Additional features	4
1.3	GPU support	8
1.4	Regularization	8
1.5	Multi-molecule models	9
1.6	Reaction support	10
1.7	Hyperparameter optimization	11
1.8	Atom/bond-level targets	13
2	Benchmark methods	14
2.1	Hyperparameter tuning	14
2.2	Datasets	15
2.2.1	MoleculeNet & OGB	15
2.2.2	SAMPL	16
2.2.3	Atom/bond-level targets	17
2.2.4	Reaction barrier heights	19
2.2.5	UV/Vis absorption	20
2.2.6	IR spectra	20
2.2.7	HOMO-LUMO gaps	20
3	Benchmark results	21
3.1	Model Performance: General benchmarking	21
3.1.1	MoleculeNet & OGB	21
3.1.2	PCQM4Mv2	23
3.1.3	SAMPL	23

3.2	Model Performance: Specific feature demonstrations	24
3.2.1	Atom/bond-level targets	24
3.2.2	Reaction barrier heights	26
3.2.3	UV/Vis absorption	27
3.2.4	IR spectra	28
3.2.5	Uncertainty estimation for QM9 gap	29
3.3	Timing	29

1 Additional Model Details

1.1 Example commands

To train a default model on the ESOL solubility dataset¹ which is distributed with Chemprop as CSV file, and save the results to the folder “checkpoint”, run

```
chemprop_train --data_path data/delaney.csv  
--dataset_type regression --save_dir checkpoint  
--save_smiles_splits
```

on the command line after installation of Chemprop following the instructions on Github.² This splits the data randomly into training, validation and test sets in the ratio 80/10/10, trains a default model and computes the performance on the test set. To compute predictions using an already trained model, run

```
chemprop_predict --checkpoint_dir checkpoint  
--test_path checkpoint/fold_0/test_smiles.csv  
--preds_path checkpoint/test_preds.csv
```

which takes the previously generated test set, computes predictions using all models in the checkpoint folder and saves them to the indicated path. For the use of Chemprop within a Python script or a graphical web interface, as well as many options to customize the model, data splits, and performance metrics please consult the instructions on Github² or the Chemprop documentation.³ Hyperparameter optimization can be performed with similar commands, as detailed in the Discussion of Features section.

1.2 Additional features

Users can provide their custom additional features by adding keywords and paths to the data files containing the features.

For molecule-level features x_m , a path to the features can be specified using the keyword `--features_path PATH/TO/FEATURES`. The provided molecular features are concatenated to the learned molecular embedding prior to the FFN network. The features can be provided as a numpy `.npy` file or CSV file. For both file formats, the features must be in the same order as the SMILES strings in the data file. The features file should not contain the SMILES strings, since features will be associated with the corresponding molecule based on the ordering in the file. The features file should contain numerical values, with columns corresponding to different features and rows corresponding to molecule data points. By default, provided features are normalized unless the flag `--no_features_scaling` is used.

For additional atomic features x_v , the path to the features can be provided using the keyword `--atom_descriptors_path PATH/TO/FEATURES`. The supported file formats include `.npz`, `.pkl`, and `.sdf`. Two options are available to select in which way atom descriptors are used. The option `--atom_descriptors descriptor` concatenates the additional features to the embedded atomic features after the D-MPNN. On the other hand, the option `--atom_descriptors feature` concatenates the features to the initial atomic feature vectors prior to the D-MPNN, such that they can be used during message-passing. Additional bond-level features can be provided via `--bond_descriptors_path PATH/TO/FEATURES` in the same format as the atom-level features. Similarly, users must choose in which way bond descriptors are used. The option `--bond_descriptors descriptor` concatenates the new bond-level features to the embedded bond features after the D-MPNN, which can only be used for bond-level property prediction, while the option `--bond_descriptors feature` concatenates the new features with the default bond feature vectors before the D-MPNN.

Users must ensure that the order of additional atom and bond features match the atom and bond ordering in the RDKit molecule object. If users wish to only use their custom features instead of the default features, the keywords `--overwrite_default_atom_features` and `--overwrite_default_bond_features` can be used to overwrite the default atom and bond features, respectively. The overwrite option is only available when the additional fea-

tures are used as **feature**. Similar to the molecular-level features, the atom- and bond-level features will be normalized automatically by default. This can be disabled with the options `--no_atom_descriptor_scaling` and `--no_bond_descriptor_scaling`.

The inputs of atom and bond features can be provided via three file formats:

- **.npz** format

Atomic descriptors are saved as 2D array ([number of atoms x number of descriptors]) for each molecule in the exact same order as the SMILES strings in the data file. Similarly, bond descriptors are saved as 2D array ([number of bonds x number of descriptors]). For example:

```
np.savez('descriptors.npz', *descriptors)
```

where `descriptors` is a list of atomic or bond descriptors in 2D array in the order of molecules in the training/predicting datafile.

- **.pkl/.pckl/.pickle** format

It contains a pandas dataframe with SMILES as index and a numpy array of descriptors as columns. For example:

Table S1: Custom atomic features for each atom provided in 1D arrays.

smiles	descriptors
<chem>CCOc1ccc2nc(S(N)(=O)=O)sc2c1</chem>	[0.6377, 0.7075...]
<chem>CCN1C(=O)NC(c2ccccc2)C1=O</chem>	[0.0958, 0.6521...]

Table S2: Multiple atomic features for each atom provided in multiple 1D arrays.

smiles	desc1	desc2
<chem>CCOc1ccc2nc(S(N)(=O)=O)sc2c1</chem>	[0.6377, 0.7075...]	[0.8266, 0.8964...]
<chem>CCN1C(=O)NC(c2ccccc2)C1=O</chem>	[0.0958, 0.6521...]	[0.2847, 0.8410...]

- **.sdf** format

Each molecule is presented as a mol block in the file. Descriptors should be saved as entries for each mol block in the format of comma separated values. Each molecule must have an entry named SMILES that stores the SMILES string. For example:

CHEMBL1308_loner5

```
RDKit          3D

6 6 0 0 1 0 0 0 0 0999 V2000
-0.7579 -0.5337 -2.8744 C 0 0 0 0 0 0 0 0 0 0 0 0
0 0
-0.2229 -1.3763 -1.7558 C 0 0 0 0 0 0 0 0 0 0 0 0
0 0
-0.0046 -1.0089 -0.4029 C 0 0 0 0 0 0 0 0 0 0 0 0
0 0
0.4824 -2.0104 0.3280 N 0 0 0 0 0 0 0 0 0 0 0 0
0 0
0.5806 -3.0317 -0.5484 N 0 0 0 0 0 0 0 0 0 0 0 0
0 0
0.1735 -2.6999 -1.8031 C 0 0 0 0 0 0 0 0 0 0 0 0
0 0
1 2 1 0
2 6 2 0
2 3 1 0
3 4 2 0
4 5 1 0
5 6 1 0
M END
> <desc1> (1)
-8.568031e-05,0.0001865207,-0.0002012379,-5.054658e
-05,0.0002148434,-0.0003503839,1.970448e-05,3.081137e
```

```

-05 , 2.997883e-05 , 9.446278e-05 , -7.194711e-05 , 0.0001527364

> <desc2> (1)
5.462954e-05 , -2.415399e-06 , 0.0001044788 , -2.274438e
-05 , 0.0001698836 , 5.206409e-06 , 4.5825e-06 , -8.882181e
-07 , -1.08787e-05 , 2.993307e-05 , -4.069051e-06 , 1.338413e-05

> <SMILES> (1)
Cc1cnnHc1

$$$$

```

1.3 GPU support

The PyTorch backend of Chemprop enables seamless GPU acceleration of both model training and inference. The acceleration of training and inference processes when used with a GPU can be significant, as shown later in Section 3.3. This feature is enabled by default on machines possessing a CUDA-enabled GPU, but it may be turned off via the `--no-cuda` argument on the command line. For machines with multiple GPUs, Chemprop will use the default GPU as PyTorch (typically the GPU at index 0). It is possible to specify the k^{th} GPU by either passing `--gpu <k>` on the command or setting the `CUDA_VISIBLE_DEVICES` environment to show only device `k` (i.e., executing `CUDA_VISIBLE_DEVICES=k` prior to running Chemprop).

1.4 Regularization

Chemprop has two builtin forms of regularization, intended to help reduce overfitting in trained models. These two regularization techniques were present in the initial release of Chemprop and remain an important contributor to model quality. The first form of regu-

larization is called early stopping. With early stopping, the performance of the model on the validation set is calculated at the end of each epoch. The version of the model that is stored at the end of training is the one saved at the end of the best scoring epoch. This has the effect of discarding later epochs of training where the model would be overfitting to the training data, continuing to improve the training loss at the cost of hurting performance on the validation and test sets. Contrary to what the name implies, early stopping as implemented in Chemprop does not shorten the amount of time needed for training.

The second form of regularization is called dropout. During training, dropout regularization will randomly zero out a fraction of the latent variables for that forward pass. This practice has been shown to reduce overfitting and lead to higher quality latent variables.⁴ The level of dropout regularization can be specified using the option `--dropout <p>` where p is the dropout probability. By default, dropout is inactive. We have observed dropout to be a helpful addition to models in a variety of contexts and recommend that users include it in their choices of hyperparameters.

1.5 Multi-molecule models

The number of molecules N is specified with the keyword `--number_of_molecules`. For the example of a solute-solvent pair $N = 2$. To train a new model using multiple molecules as an input, the SMILES string of each molecule must be provided as a separate column in the input CSV file. If N molecules are used, Chemprop assumes that the SMILES strings are located in the first N columns by default. Alternatively, the names of the specific columns containing the SMILES of the different molecules can be specified using the `--smiles_columns <column_1> ...` option. The embedding of multiple molecules in Chemprop can be done in two different ways as schematically represented in Figure S1 for $N = 2$. When multiple molecules are used, by default Chemprop trains a separate D-MPNN for each molecule (Figure S1a). If the option `--mpn_shared` is specified, the same D-MPNN is used for all molecules (Figure S1b). In both cases, the D-MPNN of each molecule uses the




Figure S1: Example of how two molecules ($N = 2$) can be embedded in Chemprop. a) A separate D-MPNN is used for each molecule or b) the same D-MPNN is used. After embedding, the different molecular vectors are concatenated (CAT) and used as input to the feed forward network (FFN) for property prediction.

same hyperparameters. The embeddings of the different molecules are then concatenated prior to the FFN. Note that the current implementation of multiple molecules in Chemprop does not ensure permutational invariance towards the input molecules. This is suited to situations where the input molecules have different roles, e.g. molecule 1 = solute, molecule 2 = solvent. For additional input information and a figure depicting the multi-molecule model structure, see the *SI*.

1.6 Reaction support

The initial atom and bond feature vectors in the CGR contain information on both the reactant and product features. Whenever information is not available, e.g. because a bond did not exist in either the reactants or products, the features are set to zero. A simple concatenation of reactant and product features can be used to obtain the pseudomolecule features

(keyword `--reaction_mode reac_prod`). Since the atomic number does not change upon reaction, its one-hot encoding is not repeated in the second part of the feature vector. For many reaction properties the change in the local structure upon reaction, i.e. the difference between reactants and products, is very informative. Since neural networks are known to not perform well for adding and subtraction operations, we also provide options to include the difference in properties directly. Namely, one can concatenate the difference in atom and bond features with the reactant properties (keyword `--reaction_mode reac_diff`, default) or with the product properties (keyword `--reaction_mode prod_diff`). Further, we provide options to automatically balance imbalanced reactions (via setting `--reaction_mode` to `reac_prod_balance`, `reac_diff_balance` or `prod_diff_balance`), which is useful if a dataset contains both balanced and imbalanced reactions. For datasets containing only balanced reactions, or only imbalanced reactions, e.g. where leaving groups are never specified, this is usually not necessary. Optionally, Chemprop accepts an additional molecule object as input, such as a solvent, or a reagent, etc. which is passed to its own D-MPNN similar to the multi-molecule model. The output of the reaction D-MPNN and molecule D-MPNN is concatenated after atomic aggregation, before the FFN. This option is available via the `reaction_solvent` keyword. The size of the molecule D-MPNN can be varied with the keywords `hidden_size_solvent` and `depth_solvent`. We note that this additional solvent or reagent needs to be passed as a separate SMILES column in the input CSV file. Reagents passed within the reaction SMILES string (between the `>` symbols) are disregarded. An example of the reaction-solvent model can be found in Ref. 5, in which a Chemprop model is trained to predict kinetic solvent effects for a diverse range of reactions and solvents.

1.7 Hyperparameter optimization

Hyperparameter optimization is a key step when building ML models that can lead to significant gains in model performance. Given the sheer size of the search space with deep learning models, this step can often be operationally challenging. Chemprop provides a

Table S3: Searchable hyperparameters using `chemprop_hyperopt`.

Keyword	Description
<code>activation</code>	the activation function used after each linear layer, when necessary
<code>aggregation</code>	the aggregation function used when constructing a molecule-level representation from node-level representations
<code>aggregation_norm</code>	the normalization factor if using <code>norm</code> aggregation
<code>batch_size</code>	the minibatch size
<code>depth</code>	the number of message-passing iterations
<code>dropout</code>	the dropout probability after each layer in both the D-MPNN encoder and FFN
<code>ffn_hidden_size</code>	the size of each hidden layer in the FFN
<code>ffn_num_layers</code>	the number of layers in the FFN
<code>hidden_size</code>	the message size in the D-MPNN encoder
<code>linked_hidden_size</code>	the size of <i>both</i> the messages in the D-MPNN encoder and the hidden layers in the FFN. This argument is overridden by either <code>hidden_size</code> or <code>ffn_hidden_size</code>
<code>max_lr</code>	the maximum learning rate used in the learning rate scheduler
<code>init_lr</code>	the initial learning rate expressed as the ratio of <code>init_lr</code> to <code>max_lr</code>
<code>final_lr</code>	the final learning rate expressed as the ratio of <code>final_lr</code> to <code>max_lr</code>
<code>warmup_epochs</code>	the number of epochs over which to ramp up the learning rate up from <code>init_lr</code> to <code>max_lr</code> expressed as a fraction of the total training epochs
<code>basic</code>	search over <code>depth</code> , <code>ffn_num_layers</code> , <code>dropout</code> , and <code>linked_hidden_layers</code>
<code>learning_rate</code>	search over <code>init_lr</code> , <code>max_lr</code> , <code>final_lr</code> , and <code>warmup_epochs</code>
<code>all</code>	all of the above hyperparameters

command line utility, `chemprop_hyperopt`, that automates this process by removing the need to manually define the search space of hyperparameters. Users can simply supply a list of keywords from which to build a hyperparameter search space (Table S3). The number of trials of hyperparameter combinations to be tested can be set using the `--num_iters` argument. By default, the search space will first be randomly sampled for `num_iters/2` trials before switching to targeted sampling via the tree-structured Parzen estimator algorithm^{6,7} for the remaining trials. The number of random trials to be used can be changed by setting `--startup_random_iters` to a value less than `num_iters`.

Hyperparameter optimization can be the most resource-intensive step in model training. In order to search a large parameter space adequately, a large number of trials would be needed. Chemprop allows for parallel operation of multiple hyperparameter optimization instances, so that the entire set of trials does not need to be run in series. Parallel operation can be achieved by setting the location of trial checkpoint files with `--hyperopt_checkpoint_dir` to be a single shared location for multiple hyperparameter optimization instances. This allows for multiple instances of the program to share and contribute to the same trial history, reducing the wall time needed to perform hyperparameter optimization significantly.

1.8 Atom/bond-level targets

The input is provided as a CSV file. The targets of atomic properties must be a 1D list in the same order as the atoms in the RDKit⁸ molecule object. The bond properties can either be a 2D list of shape $n \times n$, where n is the number of atoms, or a 1D list in the same order as the bonds in the RDKit molecule object. An example file with both atomic and bond targets is shown in Table S4. It is also important to note that Chemprop can autodetect whether a target should be an atomic or bond target.

Table S4: Example input file for atom- and bond-level property prediction. The value of `hirshfeld_charges` is presented as a 1D list, while the value of `bond_index_matrix` is presented as a 2D list.

smiles	hirshfeld_charges	bond_index_matrix
<chem>CNC(=S)N/N=C/c1c(O)ccc2ccccc12</chem>	[-0.0266, -0.0755...]	[[[0.0000, 0.9595...] ...[0.0000, 0.0000...]]]
<chem>O=C(NCCn1cccc1)c1cccc2ccccc12</chem>	[-0.2924, 0.1702...]	[[[0.0000, 1.6334...] ...[0.0059, 0.0009...]]]
<chem>C=C(C)[C@H]1C[C@@H]2OO[C@H]1C=C2C</chem>	[-0.1017, 0.0123...]	[[[0.0000, 1.9083...] ...[0.0000, 0.0000...]]]
<chem>OCCc1cc[nH]n1</chem>	[-0.2684, 0.0276...]	[[[0.0000, 0.9446...] ...[0.0000, 0.0000...]]]
<chem>CC(=N)NCc1cccc(CNCc2ccncc2)c1</chem>	[-0.0832, 0.1150...]	[[[0.0000, 1.0036...] ...[0.0005, 0.0009...]]]

Alternatively, the `--keeping_atom_map` option can be used if users wish to use atom-mapped SMILES. To apply the summation constraint to properties for each molecule, a path to the constraints can be specified using the keyword `--constraints_path PATH/TO/CONSTRAINTS` in the same order as the SMILES strings in the data file. Different constraints should be separated into different columns with a header row and one row per molecule, and the file should not contain the SMILES string. Which targets will be constrained is controlled by the names of the tasks in the constraint file header. For properties without constraints, the atomic or bond embeddings will be linked with FFN layers. Conversely, for properties with sum constraints, attention-based layers will also be constructed for each target.⁹ By default, the atom tasks share FFN weights, and bond tasks share FFN weights so that the FFN

weights might benefit from multitask training. The argument `--no_shared_atom_bond_ffn` can be used if users want to train the FFN weights for each task independently. The argument `--no_adding_bond_types` will let the bond types of each bond determined by RDKit molecules not be added to the output of bond targets. For attention-based constraining, the argument `--weights_ffn_num_layers` can be used to change the number of layers in the FFN for determining weights used to correct the constrained targets (default 2).

2 Benchmark methods

In the following, we describe the hyperparameter tuning procedure for our benchmark studies as well as the source, splitting routines, and further information on all benchmark datasets employed in this study.

2.1 Hyperparameter tuning

Training of benchmark models was carried out using hyperparameters optimized for each task. Throughout the remainder of this study, we classify datasets as small/large if they contain less/more than 10k data points in total. Models trained on small datasets were optimized for hyperparameters using 100 search iterations, whereas models trained on large datasets were optimized for only 30 iterations. During hyperparameter tuning and final model training, we trained for 200/50 epochs for small/large datasets. All models were trained on a single data split with an ensemble size of 5 during the final training, and without ensembling for hyperparameter tuning. During hyperparameter tuning, we optimized for the number of message passing steps, the hidden size during message passing, the number of layers of the feed forward neural network, as well as its hidden size, and the dropout ratio. For small datasets, furthermore the learning rate (initial, final and maximum), warm-up period and batch size were optimized. For both hyperparameter tuning and model production, scaled sums were used to aggregate atomic into molecular feature vectors. All other

parameters were left at their default values.

2.2 Datasets

The benchmarking datasets used in this study are listed in Table S5. All datasets are publicly available from the literature as described in the following. Various evaluation metrics are used to assess the performance of the Chemprop models on each dataset and against other models previously reported in the literature:

- ROC-AUC: area under the receiver operating characteristic curve
- PRC-AUC: area under the precision-recall curve
- AP: average precision
- MAE: mean absolute error
- RMSE: root-mean-square error
- R^2 : coefficient of determination
- SID: spectral information divergence

2.2.1 MoleculeNet & OGB

The HIV and PCBA datasets from MoleculeNet¹⁰ and Open Graph Benchmark (OGB)¹¹ were selected for classification tasks. Both MoleculeNet and OGB provide a diverse set of benchmark datasets that have been widely used to compare the performance of various machine learning models. They also host public leaderboards that allow us to directly compare our results to other public models. The HIV dataset contains results from an assay designed to detect HIV inhibition for 41,127 compounds. It has been observed that many of the species in the HIV dataset are at risk for assay result artifacts¹², so in the narrowest sense dataset performance should be viewed as a test for the assay result rather

than strictly as a predictor for HIV inhibition. The PCBA dataset includes the 128 biological activities selected from PubChem BioAssay¹³ for 437,929 compounds. The datasets were evaluated using the random and scaffold splits that were provided by MoleculeNet and OGB. We adopted the training, validation, and test sets of the scaffold-split HIV data and the random-split PCBA data from MoleculeNet. The scaffold-split PCBA data were adopted from OGB as MoleculeNet did not evaluate the PCBA model on the scaffold split. In all splits, the datasets were split into 80% training, 10% validation, and 10% test sets. For the random-split PCBA, MoleculeNet sets all missing targets to zero (in contrast to OGB), so we report performances for either case, i.e. filled-in zeros for comparability to the MoleculeNet leaderboard, and without filled-in values, to showcase how the observed performance drops when adopting a scaffold split versus a random split.

QM9 is a dataset of DFT calculation values commonly used for chemical model benchmarking. The calculations for this dataset were originally carried out by Ramakrishnan et al.¹⁴ and later distributed as part of the MoleculeNet benchmarks.¹⁰ The dataset is made up of 133,885 molecules with properties and structures calculated at the B3LYP/6-31G(2df,p) level of theory. The molecules were chosen as the set of possible molecules containing up to nine heavy atoms of the types C, N, O, and F. Data sources for QM9 provide 3D coordinates for the atoms in the optimized structures, but we only use molecule SMILES as inputs for model training in this work. QM9 provides 12 target values for each molecule, provided in Table S6. In the MoleculeNet presentation of the properties, atomized versions of the thermochemical properties U0, U298, H298, and G298 are provided alongside the original versions of the properties. In this work, we will use the atomized thermochemical properties. This dataset was randomly divided into 80% training, 10% validation, and 10% test data.

2.2.2 SAMPL

Experimental logP data of the SAMPL6, SAMPL7 and SAMPL9 challenges was downloaded from the SAMPL GitHub repository.¹⁵ SAMPL runs a series of blind challenges for compu-

Table S5: Summary of the benchmarking datasets.

Dataset/Category	Property/Data type	Type	N tasks	N data	Metric(s)	Ref. ^a
MoleculeNet & OGB	HIV (HIV replication inhibition)	Class.	1	41,127	ROC-AUC	10
	PCBA (biological activities)	Class.	128	437,929	PRC-AUC, AP	10,11
	QM9 (DFT calculated properties)	Regr.	12	133,885	MAE, RMSE	10,14
SAMPL	logP	Regr.	1	23,469 ^b	RMSE	15,16
Atom/bond-level targets	Quantum mechanical descriptors	Regr.	6	136,219	MAE, RMSE	9
	Bond dissociation enthalpy	Regr.	1	42,577	MAE	17,18
	Partial charge	Regr.	1	130,267	MAE, RMSE	19
Reaction barrier heights	E2	Regr.	1	1264	MAE	20–23
	S _N 2	Regr.	1	2361	MAE	20–23
	Cycloaddition	Regr.	1	5269	MAE	24
	RDB7	Regr.	1	23,852 ^c	MAE	25
	RGD1-CNHO	Regr.	1	353,984 ^c	MAE	26
UV/Vis Absorption	UV/Vis peak absorption wavelength	Regr.	1	26,395	MAE, RMSE, R^2	27–30
IR Spectra	Spectra between 400 and 4000 cm ⁻¹	Spectra	1	8,754	SID	31
PCQM4MV2	HOMO-LUMO gap	Regr.	1	3,452,151	MAE, RMSE	32

^a References for the data and data splits.

^b The size of the training set. The SAMPL6, SAMPL7, and SAMPL9 data are used as a test set.

^c Including reverse reactions.

tational chemistry, providing the identity of test molecules for which predictions of physico-chemical properties, among them the water-octanol partition coefficients, can be submitted using quantum-mechanics, molecular mechanics, or empirical models. In this work, we build an empirical model based on Chemprop, which we train on a publicly available dataset of logP measurements from Ref. 16. Molecules present in the SAMPL challenges were removed from the logP training dataset. The remaining 23,469 data points were randomly split into 80% training, 10% validation, and 10% test data. The test data was used to obtain a measure of performance for the literature dataset. We then retrained a production model on the full dataset (no validation or test data) using the best hyperparameters and number of epochs identified earlier, with which we made predictions for the three SAMPL challenges.

2.2.3 Atom/bond-level targets

To predict atom-level and bond-level targets, we selected three benchmark datasets. The framework we used to predict atomic and bond properties in Chemprop was based on modifications made to the approach developed by Guan et al.⁹ They published a dataset of

Table S6: Target values present in the QM9 dataset, presented with the target labels used in the dataset.

Label	Units	Description
mu	D	Dipole moment
alpha	α_0^3	Polarizability
HOMO	Ha	Highest occupied molecular orbital energy
LUMO	Ha	Lowest unoccupied molecular orbital energy
gap	Ha	Homo-Lumo gap
r2	α_0^2	Electronic spacial extent
ZPVE	Ha	Zero point vibrational energy
U0	kcal/mol	Internal energy at 0 K
U298	kcal/mol	Internal energy at 298.15 K
H298	kcal/mol	Enthalpy at 298.15 K
G298	kcal/mol	Free energy at 298.15 K
Cv	cal/(mol K)	Heat capacity at 298.15 K

quantum mechanical (QM) descriptors for 136,219 organic molecules with atom types H, C, O, N, F, S, Cl, Br, B, I, P, and Si. This dataset included atomic charges, Fukui indices, NMR shielding constants, bond length, and bond orders. Those molecules were optimized using GFN2-xTB and subjected to population analysis using the B3LYP/def2-SVP level of theory. We used this dataset to evaluate the performance of different implementations, randomly splitting it into 80% training, 10% validation, and 10% test data.

For benchmarking, we also used the BDE-db dataset from St. John et al.¹⁷ This dataset contains bond dissociation enthalpies (BDEs) for 42,577 closed-shell organic molecules with up to 9 heavy atoms of types C, H, O, and N, resulting in 290,664 BDEs. BDEs were calculated using the M06-2X/def2-TZVP level of theory. We used the same data splits as their study,¹⁸ with 40,577 data points as training set and 1000 molecules each in the validation and test sets.

Lastly, we included a dataset of DDEC partial charges, which includes partial charges calculated with different dielectric constants ($\epsilon = 4$ for charges in protein and $\epsilon = 78$ for charges in water).¹⁹ The dataset comprises 130,267 moderate size organic molecules with elements of types C, H, N, O, S, P, F, Cl, Br, and I, curated from ZINC and ChEMBL databases. A small fraction of data in the $\epsilon = 78$ dataset was dropped due to issues with SMILES conversion. We then randomly split the datasets into 80% training, 10% validation, and 10% test data. Two external test sets of 146 organic liquids and 1081 FDA-approved

drugs were used to test the transferability of the models.

2.2.4 Reaction barrier heights

To benchmark Chemprop's reaction functionality, four datasets of computational barrier heights were selected to cover a broad range of dataset size, diversity and quality. Since some of the original publications only report model mean absolute errors, we also report mean absolute errors, although we train on mean squared errors similar to all other benchmarks in this study.

First, E2 and S_N2 reactions originally published in Ref. 20 were used with reaction SMILES as provided in Ref. 21. We used the same split sizes as Ref. 22 and Ref. 23, which are 1440 training and 360 validation data points for S_N2, as well as 800 training and 200 validation data points for E2 with random splits. All other data points were used for testing. We then compared the performance of Chemprop to those reported in Ref. 22 and Ref. 23.

Second, cycloaddition reactions from Ref. 24 were used, with random splits into 80% training, 10% validation and 10% test data. We then compared Chemprop against all models reported in Ref. 33.

Third, the RDB7 dataset, which contains 11,926 high-accuracy reaction barrier heights and enthalpies calculated at CCSD(T)-F12/cc-pVDZ-F12 as provided in Ref. 25. In contrast to the E2, S_N2, and cycloaddition datasets that focus on one specific reaction class, this dataset spans a large range of barrier heights and is used to assess Chemprop's performance on substantially more reaction diversity. We randomly split the data into 80% training, 10% validation and 10% test data and then added reverse reactions to each set.

Fourth, the RGD1-CNHO dataset²⁶ was used, which comprises the largest and most diverse dataset out of the four, and also the most difficult to learn. We again randomly split the data into 80% training, 10% validation and 10% test data and then added reverse reactions to each set.

2.2.5 UV/Vis absorption

Multi-molecule models are demonstrated using prediction of UV/Vis peak absorption wavelength, a prediction model that involves both the absorbing molecule and the solvent. Our dataset of the peak wavelength of maximum absorption ($\lambda_{max,abs}$) is a combination of several databases²⁷⁻³⁰ that were extracted from the experimental literature. There are 26,395 samples across a variety of dye molecule families and solvents. Each sample consists of a dye molecule SMILES, solvent molecule SMILES, and a peak wavelength value. There are no multi-component species of either dyes or solvents. The train-validation-test splits are in 80/10/10 proportions and are constrained to avoid data leakage of highly-correlated measurements of the same dye in multiple solvents.

2.2.6 IR spectra

The dataset used for whole-spectra predictions was collected from infrared absorption spectra made public by NIST.³¹ This dataset comprises 8,754 gas-phase spectra, with absorbance magnitudes indicated at 2 cm^{-1} intervals between 400 and 4000 cm^{-1} . The spectra for different molecules have different ranges of collected absorbance and may have regions of missing or excluded values. We randomly split this dataset into 80% training, 10% validation, and 10% test data.

2.2.7 HOMO-LUMO gaps

The PCQM4MV2 dataset is a collection of DFT-calculated molecular HOMO-LUMO gaps, originally collected as part of the PubChemQC project³² and now curated as part of the Open Graph Benchmark.¹¹ This dataset contains HOMO-LUMO gaps measured in units of eV for 3,452,151 molecules. A further 294,470 molecules have targets privately held by the Open Graph Benchmark for blinded testing purposes and are not included in the benchmarks performed in this work. For benchmark training, the data we had available was randomly divided into 80% training, 10% validation, and 10% test data. The Open Graph Benchmark

Table S7: Test set metrics for the different targets of QM9. The top grouping of tasks was trained together in a single multitask model. The bottom grouping of results for U0 and gap shows the results for single-task models. The atomized basis of the thermochemical properties U0, U298, H298, and G298 were used for training.

Model	Target	MAE	RMSE
multitask	mu	0.326	0.581
	alpha	0.231	0.523
	HOMO	0.002 40	0.004 17
	LUMO	0.002 38	0.004 07
	gap	0.003 28	0.005 80
	r2	17.4	33.4
	ZPVE	0.000 262	0.000 366
	Cv	0.111	0.222
	U0	1.90	3.21
	U298	1.91	3.22
	H298	1.92	3.22
	G298	1.87	3.19
individual	U0	1.11	2.45
	gap	0.003 14	0.005 86

provides 3D coordinates for training data used in the dataset, but we only use molecule SMILES as inputs for model training in this work.

3 Benchmark results

3.1 Model Performance: General benchmarking

In the following, we present benchmarking results on predicting molecular targets on single-molecule datasets.

3.1.1 MoleculeNet & OGB

In the original publication of the algorithm behind Chemprop,³⁴ the MoleculeNet datasets were used as benchmark to compare against other non-deep-learning algorithms, such as Morgan Fingerprints used with random forest regression.³⁴ In this work, we do not fully repeat the original coverage of the MoleculeNet datasets. We revisit three of the datasets that continue to be of interest: QM9, HIV, and PCBA. The most significant differences in the benchmark models presented here and those presented in Ref. 34 are the improved

Table S8: Test set results for HIV and PCBA classification tasks compared with MoleculeNet (MolNet) and OGB leaderboards (higher = better). For the PCBA random split, we also report the performance with missing targets set to None (in brackets). For the PCBA scaffold split, the test set only had a single class for 'PCBA-493208' task, and therefore 'PCBA-493208' was omitted from the training, validation and test set.

Dataset	Metric	Split type	MolNet best	MolNet average	OGB best	OGB average	This work
HIV	ROC-AUC	Scaffold	0.841	0.788	0.8420	0.7936	0.8028
PCBA	PRC-AUC	Random	0.136	0.119	-	-	0.2089 (0.3840)
PCBA	AP	Random	-	-	-	-	0.2143 (0.3904)
PCBA	PRC-AUC	Scaffold	-	-	-	-	0.3012
PCBA	AP	Scaffold	-	-	0.3167	0.2716	0.3056

hyperparameter search and the adoption of the “norm” aggregation setting. For PCBA, we additionally test the performance of the model on the scaffold split that is obtained from OGB.¹¹

First, we trained a multitask model on all 12 targets in the QM9 dataset, which produced an average MAE of 2.14 and RMSE of 3.96 across all targets. Though reporting averaged metrics is common, the differing orders of magnitude among the target properties biases the averaged result heavily toward targets of larger magnitudes. In Table S7, we report the test set metrics individually by task. We also trained benchmark single-task models on the U0 and HOMO-LUMO gap targets, reported in Table S7. In this benchmark, the performance observed on the single-task treatment of U0 is significantly better than the multitask version with RMSE of 2.45 and 3.21 Ha, respectively. The single-task model did not have a clear improvement on HOMO-LUMO gap performance.

The results of the benchmark Chemprop models trained on the HIV and PCBA datasets are presented in Table S8. For the PCBA dataset which has 128 classification tasks, the test scores are averaged over all tasks. The Chemprop model achieves a ROC-AUC of 0.8028 on the scaffold split for the HIV prediction. While our model underperforms compared to the best models from MoleculeNet and OGB leaderboards, our model provides better predictions than the average models from both leaderboards on the HIV scaffold split. For the PCBA random split, the Chemprop model has a PRC-AUC of 0.2089, outperforming the best model from MoleculeNet, the DeepChem graph convolutional model³⁵ with a PRC-AUC of 0.136.

Table S9: RMSEs for predicting logP (dimensionless) for the SAMPL6, SAMPL7, and SAMPL9 challenges.

	SAMPL6	SAMPL7	SAMPL9
Number of submissions	105	36	18
Average RMSE of submissions	1.51	1.44	2.11
Best RMSE	0.38	0.49	1.12
RMSE of this work	0.27	0.49	1.11

Compared to the best model from OGB, which uses the heterogeneous interpolation on graph³⁶, our model has a lower AP of 0.3028 on the PCBA scaffold split, but we are able to achieve better performance than the average models.

3.1.2 PCQM4Mv2

A benchmark model was trained on the PCQM4Mv2 dataset curated by the Open Graph Benchmark.¹¹ This dataset contains the molecular HOMO-LUMO gap calculated by DFT in units of eV. The benchmark Chemprop model achieved a test set MAE of 0.0956 eV and RMSE of 0.154 eV. The OGB hosts a leaderboard for performance for this dataset, based on a blinded test set. The test set used for this model was part of the open data and is expected to be similar but not the same as the test set reported on the leaderboard.

3.1.3 SAMPL

When training Chemprop to predict water-octanol partition coefficients (logP), we obtain an RMSE of 0.53 on a random test set with our conventional data splits of 80% test, 10% validation and 10% test. This corresponds to an RMSE of 0.72 kcal mol⁻¹ for the transfer free energy ΔG from water to octanol at 298 K, which is related to logP by

$$\Delta G = -RT \ln 10 \cdot \log P \quad (1)$$

where R is 8.314 J mol⁻¹ K⁻¹ and T is the temperature. We then retrained a production model on the full logP dataset without any validation or test splits using the same hyperparameters to predict logP of the molecules in the SAMPL6, SAMPL7, and SAMPL9

blind prediction challenges. The performance of Chemprop is shown in Table S9, where Chemprop outperforms all other submissions from the SAMPL6, SAMPL7 and SAMPL9 challenges. We note that submissions range from quantum mechanics (QM) models and molecular mechanics (MM) models to empirical models relying on heuristic rules or machine learning, as well as mixtures thereof. Our work therefore does not only outperform other empirical models, but also a large variety of QM and MM models. In general, logP is often used in drug development, where it serves as an indicator of lipophilicity, which is known to impact the absorption, distribution, metabolism, excretion, and toxicity of drug candidates.³⁷ We thus demonstrate the ability of Chemprop to aid in important tasks such as drug discovery. Moreover, we note that the best performing submission in SAMPL7 was made by a biotechnology company independent of our group, using a Chemprop model trained on a different database, further highlighting the usefulness and impact of our software.

3.2 Model Performance: Specific feature demonstrations

In the following, we present benchmarking results for speciality features of Chemprop, namely the training on reactions or multiple molecules, the prediction of atom/bond-level targets or spectra, and the use of uncertainty quantification methods.

3.2.1 Atom/bond-level targets

As shown in Fig. S2, the performance of a multitask constrained D-MPNN was evaluated on a dataset containing six atomic and bond QM descriptors, with testing errors agreeing well with previous findings.⁹

BDE prediction was also examined using a single-task model for BDE and a multitask model for both BDE and partial charge. The single-task model achieved an MAE of $0.60 \text{ kcal mol}^{-1}$, which is comparable to the testing error of 0.58 reported by the GNN model in ALFABET.¹⁸ However, the GNN model in ALFABET was exclusively engineered for the purpose of BDE prediction, whereas the multitask model in Chemprop is capable of training

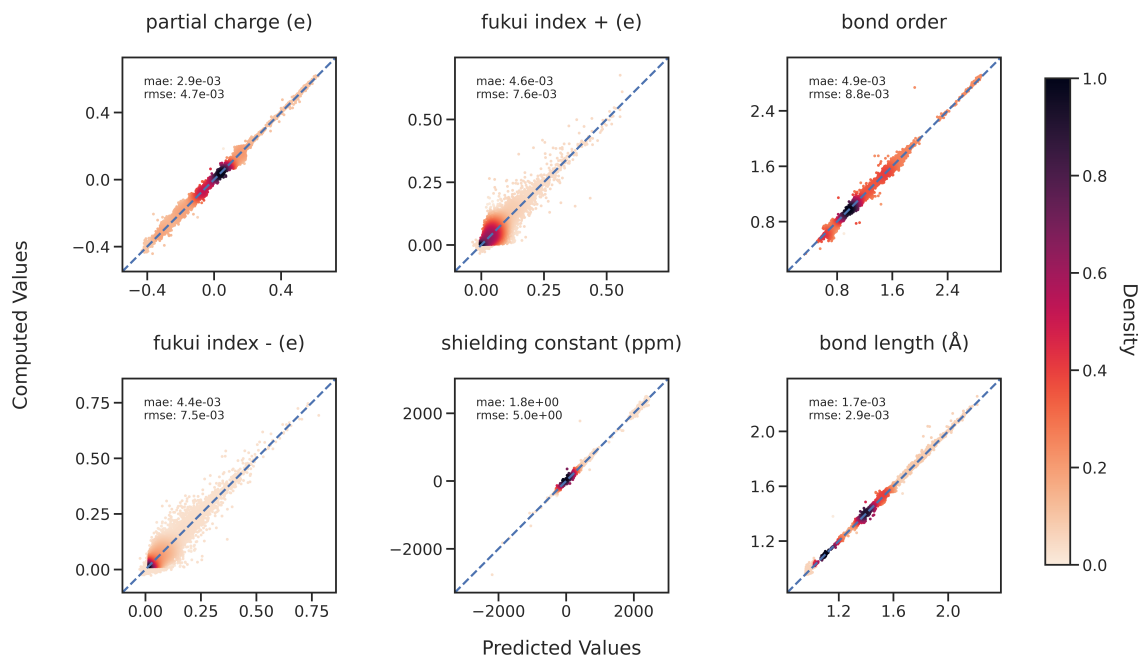


Figure S2: Comparing QM computed descriptors with multitask constrained model predictions on a held-out testing set.

on diverse atom- and bond-level properties concurrently. The multitask model, trained on both atom-level (i.e., partial charges) and bond-level (i.e., BDE) descriptors, achieved better performance with a testing MAE of $0.56 \text{ kcal mol}^{-1}$ on BDE predictions, outperforming both the GNN model and the single-task model. This was due to the synergistic effect of co-training.

The Chemprop prediction of DDEC partial charge on small organic molecules gave a low MAE (RMSE) of 0.0053 (0.0079) e and 0.0047 (0.0097) e for dielectric constants of $\epsilon = 4$ (in protein) and $\epsilon = 78$ (in solvent), respectively. Furthermore, the models were tested on external test sets to assess their transferability to real-world applications. The MAE (RMSE) for external test set 1 (organic liquids) and external test set 2 (FDA-approved drugs) was 0.0083 (0.013) e and 0.0065 (0.012) e, respectively. The results indicated that Chemprop significantly outperformed random forest regression model, with an RMSE of 0.023 e and 0.018 e for the two external test sets, as reported by Bleiziffer et al.¹⁹

These findings suggest that Chemprop is promising for predicting various atom- and

Table S10: MAEs for predicting the barrier heights of organic reactions in kcal mol⁻¹ for this work (top), other graph-convolutional approaches (middle, taken from Ref. 22 and 33), and simple machine learning approaches (bottom, taken from Ref. 22, 23 and 33).

	S _N 2	E2	Cycloadd	RDB7	RGD1
This work	2.53	2.65	2.58	4.05	5.85
WL GNN ²²	8.49	8.04	2.76	-	-
WL ml-QM-GNN ^{22,33}	2.76	2.65	2.64	-	-
Chemprop reactant ²²	2.85	2.75	-	-	-
Regression QM desc. ^{22,33}	6.43	6.07	5.94	-	-
KRR + BoB ²³	3.49	3.27	-	-	-
KRR + SLATM ²³	2.92	2.92	-	-	-
KRR + FCHL19 ²³	2.87	2.75	-	-	-
KRR + one-hot ²³	2.14	2.40	-	-	-
KNN + Morgan ³³	-	-	4.33	-	-
RF + QM desc. ³³	-	-	3.73	-	-
RF + Morgan ³³	-	-	3.50	-	-
XGBoost + QM desc. ³³	-	-	3.78	-	-
XGBoost + Morgan ³³	-	-	3.84	-	-

bond-level properties of molecules, with potential applications in drug discovery and materials science.

3.2.2 Reaction barrier heights

Table S10 summarizes the MAEs obtained for different reaction barrier height datasets.

For E2 and S_N2 reactions we can directly compare or work against the models by Stuyver et al.²² and Heinen et al.²³ We find that Chemprop significantly outperforms the Weisfeiler-Lehman (WL) architecture from Stuyver et al.^{22,33} even when adding quantum-mechanical (QM) descriptors (termed “ml-QM-GNN” in Table S10) to the WL network. We furthermore note that Ref. 22 also reports the performance of a Chemprop model, but trained it only on the reactants, not the full reactions. Here, we can directly observe the advantage offered by using the full reaction to construct the input graph representations. Chemprop furthermore outperforms the multivariate regression of quantum-mechanical descriptors of Ref. 22. Compared to the kernel ridge regression (KRR) models of Ref. 23, Chemprop outperforms the models based on BoB, SLATM, and FCHL19 representations by a large margin. The KRR model on a simple one-hot encoding of the nucleophile, electrophile and substituents close to the reactive center offers a slight performance benefit at the disadvantage of not being able

to generalize at all to new reactants or reactions.

For [3 + 2] dipolar cycloadditions, we compare Chemprop in reaction mode to WL type models with and without QM features, regressions on QM descriptors, as well as different machine learning model (k -nearest neighbor (KNN), random forests (RF) and XGBoost) on either QM descriptors or Morgan fingerprints, and find that Chemprop outperforms all other models.

A large benefit of Chemprop in reaction mode over all other architectures in Table S10 is furthermore its generality and versatility. It is straightforward to train any machine learning model on a single type of reactions (like S_N2 , E2, or cycloadditions), but finding a representation and architecture that can predict reaction properties of a large variety of reactions is much more difficult. Here, we showcase the ability of Chemprop to learn from diverse reaction datasets using the RDB7²⁵ and the RGD1-CNHO²⁶ datasets. We find larger MAEs compared to the simpler single-type datasets. Albeit not reaching chemical accuracy, our models still produce state-of-the-art performances given the diversity of reactions and range of barrier heights in both datasets. In Ref. 38, a refined Chemprop model with customized atom features and pretrained on DFT data of lower level of theory yields an MAE of 2.6 kcal mol⁻¹. Importantly, Chemprop does not make use of the three-dimensional structure of the reactants, products and transition states but estimates barrier heights solely from the change in bonds, thus requiring minimal information to predict a new reaction. We furthermore note that simpler approaches such as models trained only on reactant structures or descriptors are not applicable to diverse reaction datasets.

3.2.3 UV/Vis absorption

Chemprop achieved an MAE of 15.5 nm, RMSE of 29.7 nm, and R^2 of 0.920 on our dataset of experimental absorption peak wavelengths across a diverse set of dye molecules in a variety of solvents. This was previously demonstrated to outperform state-of-the-art fingerprint-based methods³⁹. Our train-validation-test splitting for this task was constrained such that all

Table S11: Uncertainty evaluation metrics (dimensionless) for QM9 gap predictions. *NLL*: negative log likelihood; ρ : Spearman rank correlation; *ENCE*: expected normalized calibration error; *MA*: miscalibration area. The arrows indicate if smaller or larger values indicate better performance.

Method	NLL (\downarrow)	ρ (\uparrow)	ENCE (\downarrow)	MA (\downarrow)
Ensemble	-3.94	0.33	0.21	0.14
Evidential	-4.06	0.47	0.27	0.15
MVE	-4.13	0.47	0.17	0.10

measurements of the same molecule in different solvents would be assigned to the same split to avoid data leakage. Previous work has shown that data leakage can lead to overly-optimistic estimates of generalization ability on similar datasets.³⁹ Any work on multi-molecule tasks should carefully consider the implications of the choice of splitting technique to avoid data leakage from highly correlated samples (in cases such as measurements of the same property in different solvents) or from duplicated samples with flipped molecule columns (in cases of symmetric multi-molecule properties).

3.2.4 IR spectra

Chemprop spectra prediction was benchmarked using gas-phase IR absorbance data provided publicly by NIST.³¹ Similarity between the predicted spectra and the target spectra are assessed using Spectral Information Divergence, SID. The benchmark average SID for predictions on the test set was 0.27. Qualitatively, a SID value of 0.27 is a good prediction which generally tends to match the location and magnitude of all major peaks and the location of most minor peaks, while smoothing some of the details of peak shape. To give some context to this value, we also provide some simple baselines for comparison. The average SID of a uniform distribution against the dataset was 2.52. The average SID of a roundrobin pairing of every spectrum in the dataset with every other spectrum in the dataset was 2.89. The average SID of a normalized sum of all the spectra against each individual member of the dataset was 1.45.

3.2.5 Uncertainty estimation for QM9 gap

Table S11 summarizes the performance of three uncertainty quantification (UQ) methods (ensemble, evidential, and mean-variance estimation (MVE)) selected from Chemprop’s available UQ options. For the evidential uncertainty, we used the total uncertainty (the sum of aleatoric and epistemic components). We trained all models on the gap values from the QM9 dataset and calibrated the predictions using the z-scaling method⁴⁰ with the standard deviation as the regression calibrator metric. We then evaluated the methods based on four metrics: negative log likelihood (NLL), Spearman rank correlation (ρ), expected normalized calibration error (ENCE), and miscalibration area (MA). On this task, we observe that MVE performs the best across all four metrics, while ensemble performs the worst across all metrics, with evidential in between. However, we emphasize that UQ performance can vary depending on task, dataset size, representation, and other factors. The results presented here simply serve as an illustration of Chemprop UQ capabilities; we refer readers to substantial UQ benchmarking work done previously for further discussion of the merits and pitfalls of various methods and metrics on chemical and materials data.^{41–46}

3.3 Timing

Training and inference timing benchmarks for Chemprop can be found in Tables S12, S13, and S14. These benchmarks were measured on three systems: a compute cluster node with CPU only, a compute cluster node with a GPU resource, and a laptop. We used an Intel Xeon Platinum 8260 processor (2.4 GHz, 48 CPU cores) for cluster CPU benchmarks and an Intel Xeon Gold 6248 (2.5 GHz, 40 CPU cores) processor with an Nvidia Volta V100 GPU for the cluster GPU benchmark timing. Both devices are part of the MIT Supercloud.⁴⁷ For both systems, we restricted the maximum numbers of CPU cores accessible to Chemprop to 8. For laptop timing, we used a Thinkpad X1 Carbon with an Intel Core i7-1280P (1.8 GHz, 14 CPU cores) processor and no enabled GPU. Our benchmark datasets were randomly

Table S12: Train times in hours:minutes:seconds for subsets of the QM9 dataset

Device	1k	10k	100k
Laptop	0:01:53	0:19:15	3:24:57
Cluster CPU	0:00:46	0:07:06	1:15:18
Cluster GPU	0:00:19	0:02:48	0:31:06

sampled subsets of the QM9 HOMO-LUMO gap targets with sizes of 100,000, 10,000, and 1,000.

The training times for Chemprop models found in Table S12 includes all training processes, including time for data preprocessing and model evaluation. This training was carried out with a 80/10/10 training-validation-test split of the data. The hyperparameters were chosen to be in the typical range used for datasets of this size: hidden size of 1000, feed forward hidden size of 1000, 4 message passing layers, 2 feed forward layers, and 50 epochs. The training times found in Table S13 are average training times on a per epoch basis. This average time includes the feed forward, loss function calculation, model update, and model evaluation steps. The average time excludes the time taken in the first training epoch, data preprocessing, and final model evaluation. Inference timing benchmarks can be found in Table S14. These times include all inference processes, including postprocessing of the predictions.

Training time shows significant speed improvement when moving from the laptop platform to the cluster CPU system and further improvement moving from the cluster CPU system to the cluster GPU system. This trend is followed across the tested dataset sizes. In each case, the speedup is greater than a factor of 2. Training for single models on moderately sized datasets can be carried out reasonably even on a laptop. For large datasets, hyperparameter optimization, and model structures involving many submodels, training on cluster resources or using a GPU is recommended. Inference times in the 10,000 and 100,000 dataset sizes are also improved when moving from laptop to cluster CPU to cluster GPU, but the progressive improvement is smaller than for training. Inference using any of the system levels tested is relatively fast for these dataset sizes.

Table S13: Average training times for an epoch in seconds for subsets of the QM9 dataset, excluding the first epoch.

Device	1k	10k	100k
Laptop	2.2	23.1	245
Cluster CPU	0.9	8.4	90
Cluster GPU	0.3	3.2	36

Table S14: Inference times in hours:minutes:seconds for subsets of the QM9 dataset.

Device	1k	10k	100k
Laptop	0:00:01	0:00:11	0:01:34
Cluster CPU	0:00:02	0:00:08	0:01:12
Cluster GPU	0:00:03	0:00:07	0:00:46

References

- (1) Delaney, J. S. ESOL: Estimating Aqueous Solubility Directly from Molecular Structure. *J. Chem. Inf. Comput.* **2004**, *44*, 1000–1005.
- (2) Chemprop: Molecular Property Prediction. <https://www.github.com/chemprop/chemprop>, (accessed April 6 2023).
- (3) Chemprop. <https://chemprop.readthedocs.io/en/latest/>, (accessed April 6 2023).
- (4) Hinton, G. E.; Srivastava, N.; Krizhevsky, A.; Sutskever, I.; Salakhutdinov, R. R. Improving Neural Networks by Preventing Co-adaptation of Feature Detectors. *arXiv preprint arXiv:1207.0580* **2012**,
- (5) Chung, Y.; Green, W. H. Machine Learning from Quantum Chemistry to Predict Experimental Solvent Effects on Reaction Rates. *ChemRxiv Preprint* **2023**, ChemRxiv.
- (6) Bergstra, J.; Bardenet, R.; Bengio, Y.; Kégl, B. Algorithms for Hyper-Parameter Optimization. *Adv. Neural Inf. Process. Syst.* **2011**, *24*.
- (7) Bergstra, J.; Yamins, D.; Cox, D. Making a Science of Model Search: Hyperparameter

Optimization in Hundreds of Dimensions for Vision Architectures. *Proceedings of the International Conference on Machine Learning* **2013**, 115–123.

- (8) Landrum, G. RDKit: Open-Source Cheminformatics. 2006; <https://www.rdkit.org/>.
- (9) Guan, Y.; Coley, C. W.; Wu, H.; Ranasinghe, D.; Heid, E.; Struble, T. J.; Pattanaik, L.; Green, W. H.; Jensen, K. F. Regio-Selectivity Prediction with a Machine-Learned Reaction Representation and On-the-Fly Quantum Mechanical Descriptors. *Chem. Sci.* **2021**, *12*, 2198–2208.
- (10) Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; Pande, V. MoleculeNet: A Benchmark for Molecular Machine Learning. *Chem. Sci.* **2018**, *9*, 513–530.
- (11) Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; Leskovec, J. Open Graph Benchmark: Datasets for Machine Learning on Graphs. 2021.
- (12) Filtering Chemical Libraries. <https://practicalcheminformatics.blogspot.com/2018/08/filtering-chemical-libraries.html>, (accessed November 14 2023).
- (13) Wang, Y.; Xiao, J.; Suzek, T. O.; Zhang, J.; Wang, J.; Zhou, Z.; Han, L.; Karapetyan, K.; Dracheva, S.; Shoemaker, B. A.; Bolton, E.; Gindulyte, A.; Bryant, S. H. PubChem’s BioAssay Database. *Nucleic Acids Res.* **2012**, *40*, D400–D412.
- (14) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; Von Lilienfeld, O. A. Quantum Chemistry Structures and Properties of 134 kilo Molecules. *Sci. Dat.* **2014**, *1*, 140022.
- (15) The SAMPL Challenges. <https://github.com/samplchallenges>, (accessed March 15, 2023).
- (16) Koscher, B.; Canty, R. B.; McDonald, M. A.; Greenman, K. P.; McGill, C. J.; Bilodeau, C. L.; Jin, W.; Wu, H.; Vermeire, F. H.; Jin, B., et al. Autonomous, Multi-

Property-Driven Molecular Discovery: From Predictions to Measurements and Back. *ChemRxiv Preprint* **2023**, ChemRxiv.

- (17) St. John, P. C.; Guan, Y.; Kim, Y.; Kim, S.; Paton, R. S. BDE-db: A Collection of 290,664 Homolytic Bond Dissociation Enthalpies for Small Organic Molecules. <https://doi.org/10.6084/m9.figshare.10248932.v1>, 2019; (accessed March 22, 2023).
- (18) St. John, P. C.; Guan, Y.; Kim, Y.; Kim, S.; Paton, R. S. Prediction of Organic Homolytic Bond Dissociation Enthalpies at Near Chemical Accuracy with Sub-Second Computational Cost. *Nat. Commun.* **2020**, *11*, 2328.
- (19) Bleiziffer, P.; Schaller, K.; Riniker, S. Machine Learning of Partial Charges Derived from High-Quality Quantum-Mechanical Calculations. *J. Chem. Inf. Model.* **2018**, *58*, 579–590.
- (20) von Rudorff, G. F.; Heinen, S. N.; Bragato, M.; von Lilienfeld, O. A. Thousands of Reactants and Transition States for Competing E2 and S2 Reactions. *Mach. Learn.: Sci. Technol.* **2020**, *1*, 045026.
- (21) Heid, E.; Green, W. H. Machine Learning of Reaction Properties via Learned Representations of the Condensed Graph of Reaction. *J. Chem. Inf. Model.* **2022**, *62*, 2101–2110.
- (22) Stuyver, T.; Coley, C. W. Quantum Chemistry-Augmented Neural Networks for Reactivity Prediction: Performance, Generalizability, and Explainability. *J. Chem. Phys.* **2022**, *156*, 084104.
- (23) Heinen, S.; von Rudorff, G. F.; von Lilienfeld, O. A. Toward the Design of Chemical Reactions: Machine Learning Barriers of Competing Mechanisms in Reactant Space. *J. Chem. Phys.* **2021**, *155*, 064105.

- (24) Stuyver, T.; Jorner, K.; Coley, C. W. Reaction Profiles for Quantum Chemistry-Computed [3+ 2] Cycloaddition Reactions. *Sci. Dat.* **2023**, *10*, 66.
- (25) Spiekermann, K.; Pattanaik, L.; Green, W. H. High Accuracy Barrier Heights, Enthalpies, and Rate Coefficients for Chemical Reactions. *Sci. Dat.* **2022**, *9*, 417.
- (26) Zhao, Q.; Vaddadi, S. M.; Woulfe, M.; Ogunfowora, L. A.; Garimella, S. S.; Isayev, O.; Savoie, B. M. Comprehensive Exploration of Graphically Defined Reaction Spaces. *Sci. Dat.* **2023**, *10*, 145.
- (27) Joung, J. F.; Han, M.; Jeong, M.; Park, S. Experimental Database of Optical Properties of Organic Compounds. *Sci. Dat.* **2020**, *7*, 1–6.
- (28) Ju, C.-W.; Bai, H.; Li, B.; Liu, R. Machine Learning Enables Highly Accurate Predictions of Photophysical Properties of Organic Fluorescent Materials: Emission Wavelengths and Quantum Yields. *J. Chem. Inf. Model.* **2021**, *61*, 1053–1065.
- (29) Venkatraman, V.; Raju, R.; Oikonomopoulos, S. P.; Alsberg, B. K. The Dye-Sensitized Solar Cell Database. *J. Cheminf.* **2018**, *10*, 1–9.
- (30) Venkatraman, V.; Kallidanthiyil Chellappan, L. An Open Access Data Set Highlighting Aggregation of Dyes on Metal Oxides. *Data* **2020**, *5*, 45.
- (31) NIST Mass Spectrometry Data Center, In *"Infrared Spectra" in NIST Chemistry WebBook, NIST Standard Reference Database Number 69*; Lindstrom, P. J., Mallard, W. G., Eds.; National Institute of Standards and Testing: Gaithersburg, MD, <http://webbook.nist.gov> (Accessed 2019-10-3).
- (32) Nakata, M.; Shimazaki, T. PubChemQC Project: A Large-Scale First-Principles Electronic Structure Database for Data-Driven Chemistry. *J. Chem. Inf. Model.* **2017**, *57*, 1300–1308.

- (33) Stuyver, T.; Coley, C. Machine Learning-Guided Computational Screening of New Candidate Reactions with High Bioorthogonal Click Potential. *Chem. Eur. J.* **2023**, e202300387.
- (34) Yang, K.; Swanson, K.; Jin, W.; Coley, C.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M.; Palmer, A.; Settels, V.; Jaakkola, T.; Jensen, K.; Barzilay, R. Analyzing Learned Molecular Representations for Property Prediction. *J. Chem. Inf. Model.* **2019**, *59*, 3370–3388.
- (35) Altae-Tran, H.; Ramsundar, B.; Pappu, A. S.; Pande, V. S. Low Data Drug Discovery with One-shot Learning. *arXiv preprint arXiv:1611.03199* **2016**,
- (36) Heterogeneous Interpolation on Graph. <https://github.com/TencentYoutuResearch/HIG-GraphClassification>, 2021; (accessed May 28 2023).
- (37) Sun, Y.; Hou, T.; He, X.; Man, V. H.; Wang, J. Development and Test of Highly Accurate Endpoint Free Energy Methods. 2: Prediction of Logarithm of n-Octanol–Water Partition Coefficient (logP) for Druglike Molecules using MM-PBSA Method. *J. Comp. Chem.* **2023**, *44*, 1300–1311.
- (38) Spiekermann, K. A.; Pattanaik, L.; Green, W. H. Fast Predictions of Reaction Barrier Heights: Toward Coupled-Cluster Accuracy. *J. Phys. Chem. A* **2022**, *126*, 3976–3986.
- (39) Greenman, K. P.; Green, W. H.; Gómez-Bombarelli, R. Multi-Fidelity Prediction of Molecular Optical Peaks with Deep Learning. *Chem. Sci.* **2022**, *13*, 1152–1162.
- (40) Levi, D.; Gispan, L.; Giladi, N.; Fetaya, E. Evaluating and calibrating uncertainty prediction in regression tasks. *Sensors* **2022**, *22*, 5540.
- (41) 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. *J. Chem. Inf. Model.* **2020**, *60*, 2697–2717.

- (42) Tran, K.; Neiswanger, W.; Yoon, J.; Zhang, Q.; Xing, E.; Ulissi, Z. W. Methods for Comparing Uncertainty Quantifications for Material Property Predictions. *Machine Learning: Science and Technology* **2020**, *1*, 025006.
- (43) Hirschfeld, L.; Swanson, K.; Yang, K.; Barzilay, R.; Coley, C. W. Uncertainty Quantification using Neural Networks for Molecular Property Prediction. *J. Chem. Inf. Model.* **2020**, *60*, 3770–3780.
- (44) Nigam, A.; Pollice, R.; Hurley, M. F.; Hickman, R. J.; Aldeghi, M.; Yoshikawa, N.; Chithrananda, S.; Voelz, V. A.; Aspuru-Guzik, A. Assigning Confidence to Molecular Property Prediction. *Expert Opin. Drug Discov.* **2021**, *16*, 1009–1023.
- (45) Soleimany, A. P.; Amini, A.; Goldman, S.; Rus, D.; Bhatia, S. N.; Coley, C. W. Evidential Deep Learning for Guided Molecular Property Prediction and Discovery. *ACS Cent. Sci.* **2021**, *7*, 1356–1367.
- (46) Heid, E.; McGill, C. J.; Vermeire, F. H.; Green, W. H. Characterizing Uncertainty in Machine Learning for Chemistry. **2023**,
- (47) 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. *Proceedings of the IEEE High Performance Extreme Computing Conference* **2018**, 1–6.