

A CODE AND DATA AVAILABILITY

Anonymized code and datasets available at:

<https://drive.google.com/drive/folders/1UKqVvz9Ae35PThyqB7un1ZNb2bT8d740?usp=sharing>

Final versions will be released on GitHub and made available on a dedicated website, respectively.

B TIPS AND TRICKS FOR WORKING WITH 3D ATOMIC DATA

In order to facilitate the entry of new practitioners to the nascent field of 3D molecular learning, we provide some high-level guidelines for working with the datasets we provide and for curating new ones. For many of the tasks mentioned, we provide computational tools on our GitHub repository.

B.1 ASSEMBLING NEW DATASETS

Data sources and repositories. The success of deep learning methods strongly depends on the availability of sufficient training data. Unless they have a laboratory that can produce the necessary data, most scientists will probably have to use public databases. The go-to repository for protein structures is the Protein Data Bank (PDB).¹ RNA structures can be found at the RNA 3D hub of Bowling Green State University.² An exhaustive repository for small molecules is ChEMBL.³ The 3D structures of small molecules are mostly not directly deposited. They can be generated by quantum-chemical methods (expensive) or in good approximation by cheminformatics tools such as RDKit. Many more specific databases are out there and worth being explored. We provide methods to mine and convert data from many common formats in the field on our GitHub repository.

Scope and limitations of the data. Even the most extensive databases cannot capture the large diversity of biological macromolecules or the space of potential drug molecules. It is thus necessary to think about the scientific problem at hand and whether the available data adequately represents the range of structures that are responsible for the studied effects. An important general limitation of structural data is that molecules change conformation fluidly in real life due to thermal fluctuations. Additionally, interactions with other molecules, disordered regions, or environmental factors like pH can result in significant differences from their experimentally determined forms.

Incomplete or corrupted data. Structural data is rarely perfect. Experimental uncertainties are mostly caused by limited resolution of the involved techniques such as X-ray crystallography or electron cryo-microscopy. Computationally generated structures are also prone to flaws in molecular force fields or basis sets for wave functions. These limitations can lead to problems such as unrealistic conformations, missing or duplicate atoms, non-resolved amino acid side chains, and more. One has to decide whether to keep those structures or to sanitize them using computational tools. Additionally, hydrogen atoms are often not included in the data and, if needed for the task, have to be added when assembling the dataset. The most important guideline here is to be consistent and clear in the way these issues are treated. Sometimes it can be necessary to assemble two different datasets with different treatments of missing data.

B.2 DEVELOPING AND BENCHMARKING NEW ALGORITHMS

Reading and preprocessing. Algorithms represent data in various ways and a given dataset is not always compatible with the representation needed for the algorithm. For example, certain structures, residues, or atoms may need to be filtered out. Ideally, these steps are considered as dataset preparation and are separated from the algorithm itself, i.e. not hard-coded into the dataloader. This has two main advantages: (1) it saves time upon multiple reruns of the algorithm as structural data can be large and expensive to process, and (2) saving the preprocessed input dataset separately increases reproducibility, because small differences in preprocessing are often not recorded. We provide our

¹<https://www.rcsb.org/>

²<http://rna.bgsu.edu/rna3dhub/nrlist/>

³<https://www.ebi.ac.uk/chembl/>

benchmarking datasets in a format that is easy to read for many Python-based algorithms and provide the necessary tools on GitHub.

Comparing algorithms. Predictions can be tested with various metrics. Depending on the prediction problem, some of the metrics grasp the scientific aims of the training better than others. It is usually recommended to stick to the metrics that are common in the field and are given in the benchmarks. As science develops, new metrics for a specific problem might come up. These should be well justified and the old metrics should still be reported alongside them to allow for a comparison. Ideally, the new metrics are calculated for older models, too. To facilitate this in advance, when benchmarking an algorithm, specific predictions should be stored and not only metrics.

Interpretation of results. When judging the performance of an algorithm, one should take into account the experimental uncertainties both in the structures but also in the label data. While small molecules can be investigated in much detail, it will rarely be possible to get near perfect performance for tasks involving complex biological macromolecules. Over time, even held-out test sets become part of the selection process for new methods as only those methods that perform better on the test set will prevail. A measured improvement can thus be caused by minor specifics of the test set. As the field matures and performance becomes saturated, the benchmark sets will still be valid as sanity checks for new methods, but harder tasks will be the ones driving new development.

C METHODOLOGICAL DETAILS

C.1 DATASET PREPARATION

We present a set of methods to mine task-specific atomic datasets from several large databases (e.g. PDB) as well as to filter them, split them, and convert them to a format suitable for standard machine learning libraries (esp. PyTorch and TensorFlow). We store these datasets in HDF5 format, where each atom is stored as a row in a standardized data frame. This data format accurately captures the natural hierarchy of atom subgroups in biomolecules, especially proteins, and enables data loading and processing to be consistent across datasets, tasks, and computational environments.

To capture hierarchical information in a way that is task-specific but standardized, we define an "ensemble" to be the highest-level of structure for each example, e.g. the PDB entry for the protein. Within each ensemble, we define "subunits", which represent the specific units of structure used for that task. For example, for the paired tasks (PIP, LEP, MSP), there is one subunit corresponding to each structure in the pair; for RES, there is one subunit for each residue microenvironment, and for structure ranking (PSR, RSR), there is one subunit for each candidate 3D structure. In this way, it is simple to iterate over each dataset and extract each atomistic structure, which can then be augmented and processed into any desired format (e.g. voxelized for the 3DCNN, converted to graphs for the GNN).

In the following sections, we describe the specific steps used to mine and process each dataset.

Small Molecule Properties The QM9 dataset (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) contains the results of quantum-chemical calculations for 134,000 stable small organic molecules made up of maximally nine atoms of C, O, N, and F. For each molecule, it contains the geometry of a molecule’s conformation in its ground state as well as calculated energetic, electronic, and thermodynamic properties. In particular, these properties are:

- μ - Dipole moment (unit: D)
- α - Isotropic polarizability (unit: bohr^3)
- ϵ_{HOMO} - Highest occupied molecular orbital energy (unit: Ha, reported in eV)
- ϵ_{LUMO} - Lowest unoccupied molecular orbital energy (unit: Ha, reported in eV)
- ϵ_{gap} - Gap between HOMO and LUMO (unit: Ha, reported in eV)
- R^2 - Electronic spatial extent (unit: bohr^2)
- ZPVE - Zero point vibrational energy (unit: Ha, reported in meV)

- U_0 - Internal energy at 0 K (unit: Ha)
- U_{298} - Internal energy at 298.15 K (unit: Ha)
- H_{298} - Enthalpy at 298.15 K (unit: Ha)
- G_{298} - Free energy at 298.15 K (unit: Ha)
- C_v - Heat capacity at 298.15 K (unit: $\frac{\text{cal}}{\text{mol K}}$)

It is common to subtract the reference thermochemical energy from $U_0, U_{298}, H_{298}, G_{298}$ to obtain:

- U_0^{at} - Atomization energy at 0K (unit: $\frac{\text{kcal}}{\text{mol}}$, reported in eV)
- U_{298}^{at} - Atomization energy at 298.15K (unit: $\frac{\text{kcal}}{\text{mol}}$, reported in eV)
- H_{298}^{at} - Atomization enthalpy at 298.15K (unit: $\frac{\text{kcal}}{\text{mol}}$, reported in eV)
- G_{298}^{at} - Atomization free energy at 298.15K (unit: $\frac{\text{kcal}}{\text{mol}}$, reported in eV)

We report metrics for these quantities in the benchmark.

The QM9 dataset is processed from the one provided in MoleculeNet (Wu et al., 2018). As recommended by the authors of the original dataset, we exclude 3,054 molecules that do not pass a geometrical consistency test (Ramakrishnan et al., 2014). Additionally, we excluded all 1,398 molecules that RDKit is unable to process - as in former GNN work (Fey & Lenssen, 2019). In this way, we ensure that all models in this work can be trained on the same data. Following previous work (Wu et al., 2018; Gilmer et al., 2017; Schütt et al., 2017; Anderson et al., 2019), we split the remaining dataset randomly in training, validation, and test set - containing 103547, 12943, and 12943 molecules, respectively.

Protein Interface Prediction For our test set, we download the cleaned PDB files from the DB5 dataset as provided in (Townshend et al., 2019), and convert to our standardized format. Each complex is an ensemble, with the bound/unbound ligand/receptor structures forming 4 distinct subunits of said ensemble. We use the bound forms of each complex to define neighboring amino acids (those with any heavy atoms within 6 Å of one another), and then map those onto the corresponding amino acids in the unbound forms of the complex (removing those that do not map). These neighbors are then included as the positive examples, with all other pairs being defined as negatives. At prediction time, we attempt to re-predict which possible pairings are positive or negative, downsampling negatives to achieve a 1:1 positive to negative split. We use the unbound subunits as our pair of input structures for testing. We use AUROC of these predictions as our metric to evaluate performance.

For our training set, we reproduce the Database of Interacting Protein Structures (DIPS) (Townshend et al., 2019). Specifically, we take the snapshot of all structures in the PDB from November 20, 2015. We apply a number of filtering operations, removing structures with no protein present, structures with less than 50 amino acids, structures with worse than 3.5 Å resolution, and structures not solved by X-ray crystallography or Cryo-EM. We then split the dataset into all pairs of interacting chains. These pairs form our ensembles, with each of the two chains being one subunit. We then remove pairs with less than 500 Å² buried surface area as measured by the FreeSASA Python library (Mitternacht, 2016) (using total area computed the naccess classifier, including hydrogens and skipping unknown residues). Furthermore, to ensure there is no train/test contamination, we prune this set against the DB5 set defined above, removing any pairs that have a chain with more than 30% sequence identity, using the software BLASTP (Altschul et al., 1990)). We also prune the set based on structural similarity, removing any pairs in DIPS that map to corresponding SCOP (Andreeva et al., 2014) pairs of superfamilies that are also present across a pair in DB5 (i.e., we remove a DIPS pair if the first subunit in that pair has a chain with a SCOP superfamily that is present in an unbound subunit of a DB5 pair, and the second subunit in that DIPS pair also has a SCOP superfamily that is present in the other unbound subunit of that same DB5 pair). Once this pruning is done, we split the DIPS set into a training, validation, and (internal) testing set based on PDB sequence clustering at a 30% identity level, to ensure little contamination between them. We perform a 80%, 10%, 10% split for training, validation, and testing, respectively. Note this internal testing set is not used for performance reporting.

Residue Identity Environments are extracted from a non-redundant subset of high-resolution structures from the PDB. Specifically, we use only X-ray structures with resolution < 3.0 Å, and enforce a 60% sequence identity threshold. We then split the dataset by structure based on domain-level CATH 4.2 topology classes (Dawson et al., 2017), as described in (Anand et al., 2020). This resulted in a total of 21147, 964, and 3319 PDB structures for the train, validation, and test sets, respectively. Rather than train on every residue for each of these structures, we balance the classes in the train set by downsampling to the frequency of the least-common amino acid (cysteine). The original class balance is preserved in the test set. In total, the train, validation, and test sets comprised 3733710, 188530, and 1261342 environments, respectively. We ignore all non-standard residues. We represent the physico-chemical environment around each residue using all C, O, N, S, and P atoms in the protein and any co-crystallized ligands or ions. All non-backbone atoms of the target residue are removed, and each environment is centered around a “virtual” $C\beta$ position of the target residue defined using the average $C\beta$ position over the training set.

Mutation Stability Prediction Mutation data are extracted from the SKEMPI 2.0 database (Jankauskaitė et al., 2019). Non-point mutations or mutants that cause non-binding of the complex are screened out. Additionally, mutations involving a disulfide bond and mutants from the PDBs 1KBH or 1JCK are ignored due to processing difficulties. A label of 1 is assigned to a mutant if the K_d of the mutant protein is less than that of the wild-type protein, indicating better binding, and 0 otherwise. Atoms from the twenty canonical amino acids were extracted from the PDBs provided in SKEMPI using PyMOL (Schrödinger, LLC, 2015), and in silico mutagenesis is accomplished using PyRosetta (Chaudhury et al., 2010), where dihedrals within 10 Å of the mutated residue are repacked. This protocol produces 893 positive examples and 3255 negative examples. For ENN training, we use structures that are reduced to a size that is tractable for Cormorant. To this end, we only selected the regions within a radius of 6 Å around the $C\alpha$ -atom of the mutated residue. For 3DCNNs, we analogously used a radius of 25 Å. GNNs are trained on complete structures. This dataset is split by sequence identity at 30%.

Ligand Binding Affinity PDDBind contains X-ray structures of proteins bound to small molecule and peptide ligands. We use the “refined set” (v.2019) consisting of 4,852 complexes filtered for various quality metrics, including resolution ≤ 2.5 Å, R-factor ≤ 0.25 , lack of steric clashes or covalent bonding, and more (Li et al., 2014). We further exclude complexes with invalid ligand bonding information. The binding affinity provided in PDDBind is experimentally determined and expressed in terms of inhibition constant (K_i) or dissociation constant (K_d), both in Molar units. As in previous works (Ballester & Mitchell, 2010; Zilian & Sotriffer, 2013; Ragoza et al., 2017; Jiménez et al., 2018), we do not make the distinction between K_i and K_d , and instead predict the negative log-transformed binding affinity, or pK . The majority of prior scoring functions have used the “core set” provided by the Critical Assessment of Scoring Functions (CASF) (Su et al., 2019) as a test set for evaluating prediction performance. However, by construction every protein in this test set is at least 90% identical to several proteins in the training set. Thus, performance on this test set does not accurately represent the generalizability of a scoring function, and has been shown to overestimate the performance of machine learning models in particular (Kramer & Gedeck, 2010; Gabel et al., 2014; Li & Yang, 2017). Therefore, to prevent overfitting to specific protein families, we create a new split in training, validation and test set, based on a 30% sequence identity threshold to limit homologous proteins appearing in both train and test sets. Specifically, for the creation of each of validation and test set, we extract a single protein at random and all other proteins with which it shares $> 30\%$ sequence identity, as calculated by BLASTP. We perform this procedure iteratively until each of the training and test sets are 10% as large as the overall dataset. We additionally enforce that no single sequence identity cluster represents more than 20% of the overall split to prevent overrepresentation of any single protein family. Splitting using this procedure resulted in training, validation, and test sets of size 3507, 466, and 490, respectively.

For comparison, we provide an additional, less restrictive, split based on a 60% sequence identity threshold (results in Table 8). This leads to training, validation, and test sets of size 3678, 460, and 460, respectively.

For Cormorant, we use a reduced dataset without hydrogens and only the most abundant heavy elements in the full dataset (C, N, O, S, Zn, Cl, F, P, Mg) as well as only structures with 500 or fewer atoms. The Cormorant implementation requires that training, validation, and test set contain

the same set of atom types. To achieve this, we removed all structures containing elements that are very rare in PDBBind (Br, I, Ca, Na, K, Mn, Fe, Co, Ni, Cu, Cs, Sr, Cd, Ni). The limitation of atom numbers is purely technical. The Kronecker products involved in the covariant neurons are very memory intensive and training Cormorant on larger structures was limited by the memory of the GPUs available to us. For the 30% identity split, the reduced training, validation, and test datasets contain 3274, 421, and 457 structures, respectively, and for the 60% identity split, they contain 3314, 415, and 423 structures, respectively.

Ligand Efficacy Prediction Each input consists of a ligand bound to both the active and inactive conformation of a specific protein. The goal is to predict the label for this drug/ligand, either an “activator” or “inactivator” of the protein function. Why include these protein conformations in the input? From a biochemical perspective, if the drug binds much more favorably to the active protein conformation, it will act as an activator of the protein function. The model may then learn this differential binding strength to improve predictions of ligand function.

Pairs of structures for 27 proteins are obtained through manual curation of the Protein Data Bank structures where “active” and “inactive” conformational states are both available. For example, for ion channels, this means a channel in an open vs. closed state. 527 ligands with known protein binding and labeled function are selected from the IUPHAR database. We label ligands as activators if they are listed as “agonists” or “activators” and label ligands as inactivators if they are listed as “inhibitors” or “antagonists”. We select up to 15 of both activating and inactivating ligands for each protein.

We model the drugs bound to the relevant protein. To prepare protein structures for use in docking, we first prepare structures using the Schrödinger suite. All waters are removed, the tautomeric state of the ligand present in the experimentally determined structure is assigned using Epik at pH 7.0 \pm 2.0, hydrogen bonds are optimized, and energy minimization is performed with non-hydrogen atoms constrained to an RMSD of less than 0.3 Å from the initial structure. For ligands to be docked, the tautomeric state is assigned using Epik tool at target pH 7.0. Ligands are docked using default Glide SP. This results in 527 pairs of complexes. These are split into training, validation, and tests sets by protein target to ensure generalizability across proteins.

For ENN training, we have to reduce the structures to a size that is tractable for Cormorant. To this end, we only use the regions within a radius of 5.5 Å around the ligand. For 3D-CNNs, we analogously use a radius of 25 Å. GNNs were trained on complete structures.

Protein Structure Ranking The Critical Assessment of Structure Prediction (CASP) (Kryshtafovych et al., 2019) is a long-running international competition held biennially, of which CASP13 is the most recent, that addresses the protein structure prediction problem by withholding newly solved experimental structures (called *targets*) and allowing computational groups to make predictions (called *decoys*), which are then evaluated for their closeness to their targets after submission. Those submissions are then carefully curated and released as decoy sets in two stages (20 decoys per target for Stage 1, 150 decoys per target for Stage 2) for the Model Quality Assessment (MQA), one of the categories in CASP which aims to score a set of decoys of a target based on how closely they are to the target. For the PSR dataset, we compiled those decoys sets released in CASP5-13, then relaxed those structures with the SCWRL4 software (G. G. Krivov & Dunbrack, 2009) to improve side-chain conformations. For all decoys in the dataset, we computed the RMSD, TM-score, GDT_TS, and GDT_HA scores using the TM-score software (Zhang & Skolnick, 2007).

Mirroring the setup of the competition, we split the decoy sets based on target and released year. More specifically, we randomly split the targets in CASP5-10 and randomly sample 50 decoys for each target to generate the training and validation sets (508 targets for training, 56 targets for validation), and use the whole CASP11 Stage 2 as test set (85 targets total, with 150 decoys for each target). We chose CASP11 as test set, as the targets in CASP12-13 are not fully released yet.

RNA Structure Ranking The RNA Puzzles competition (Cruz et al., 2012) is a rolling international competition dealing with the RNA structure prediction problem. Similarly to CASP, newly solved experimental structures, referred to as natives, are withheld until computational groups make prediction, referred to as candidates. These candidates are then evaluated by their RMSD from the native. For this task, we use candidate structures created by the state-of-the-art structure generation

method, FARFAR2 (Watkins & Das, 2019), for each of the 21 first RNA Puzzles. These are made available as part of the FARFAR2 publication. There are an average of 21303 (standard deviation of 13973) candidates generated per puzzles, with a large range of RMSDs. For the RSR dataset we randomly sample 1000 candidates per puzzle. We split temporally, by puzzle, using RNA Puzzles 1-13 for training, 14-17 for validation, and 18-21 for testing.

C.2 TASK-SPECIFIC EXPERIMENTAL DETAILS

3DCNNs Our base 3DCNN architecture consists of four 3D-convolutional layers with increasing filter size (32, 64, 128, and 256) — each followed by max-pooling, dropout, and ReLU activation — and one fully-connected layer of size 512, followed by dropout and ReLU activation. For single model task (PSR, RSR, LBA, SMP), we add an additional fully-connected layer to transform to the required output dimension size. For paired tasks (PIP, LEP, MSP), we adapt this base architecture into a Siamese network, add an additional fully-connected layer followed by dropout and ReLU activation to combine the output of each member of the pair, and finally add a final fully-connected layer to transform to the required output dimension size, as in (Townshend et al., 2019).

For input to the 3DCNNs, we represent our data as cube in 3D space of certain radius (50 Å for PSR, RSR; 17 Å for PIP; 20 Å for LBA; 7.5 Å for SMP; 25 Å for LEP, MSP; 10 Å for RES) that are discretized into voxels with resolution of 1 Å to form a grid. For paired tasks (PSR, RSR, and PIP), we form a separate voxel grid for each member of the pair. For most tasks, we use the centroid of each input structure as center of the grid, excluding LBA where we use the centroid of the ligand as center and MSP where we use the mutation point as center. Each grid voxel is associated with a binary feature vector which encodes the presence or absence of each specified atom type in that voxel. For PSR, RSR, PIP, and RES, we encode the presence of heavy atoms C, O, N, and S (P for RSR since S does not exist in RNA structures). For other tasks where hydrogen bonds might play an important role, we encode the hydrogen atom (H) in addition to C, O, N, and few other abundant atoms (F for LBA and SMP; S, Cl, F for LEP; S for MSP). To encode rotational symmetries, we apply a data augmentation strategy in which we apply 20 random rotations to the input grid, as in (Townshend et al., 2019), except for RES, where we instead apply the canonicalization procedure described in (Anand et al., 2020). During testing of LEP, we applied 35 random rotations to the each sample, and took majority voting among those 35 predictions as the final predicted class.

For binary classification tasks, we use binary cross-entropy weighted by the class distribution (i.e. rarer class is weighted more heavily on mistakes). To address issues with imbalanced datasets, we randomly oversample/undersample the less/more frequent class respectively during training. For regression tasks, we use mean-squared error loss for training. All models were trained with Adam optimizer with default beta parameters and learning rate 0.001 for SMP; 0.0001 for PSR, RSR, PIP, LBA, RES; and 0.00001 for LEP, MSP. We monitor the loss on the validation set at every epoch. The weights of the best-performing are then used to evaluate on the held-out test set. The models were all trained on 1 Titan X(p) GPU for 4–24 hours depending on the task.

GNNs Our base GNN architecture consists of five layers of graph convolutions as defined by Kipf and Welling (Kipf & Welling, 2016), with increasing hidden dimension (64, 128, 128, 256, 256) each followed by batch normalization and ReLU activation. For tasks with paired input structures (PIP, LEP, MSP), we apply this convolutional architecture to each input structure separately in a Siamese architecture with tied weights, and then concatenate the outputs before passing through two fully-connected layers of size 256 to transform to an output dimension of one neuron for binary classification. We regularize using dropout with a probability of 0.25 after the first fully-connected layer. Some tasks require classification of an entire structure, and thus are well-suited to graph-level outputs (PSR, RSR, LBA). Here, we apply global addition pooling across all nodes before applying the final two layers. For PIP, RES, and MSP, instead of pooling we instead extract the embedding of the node corresponding to the $C\alpha$ atom of the residue in question (interacting residue, deleted residue, and mutated residue, respectively) after the final convolutional layer. For SMP, we use the previously-developed architecture presented in Gilmer et al. (2017), which is publicly available.

We use a very simple featurization scheme for atomic systems. We define edges between all atoms separated by less than 4.5 Å. Edges are weighted by the distance between the atoms, and nodes are featurized by one-hot-encoding all heavy (non-hydrogen) atoms. The only exceptions are LBA, where we include hydrogens due to the importance of polar interactions and hydrogen bonding in

determining binding affinity, and SMP, where we adopt the established featurization scheme used in MoleculeNet (Wu et al., 2018). All GNNs were implemented in PyTorch Geometric (Fey & Lenssen, 2019).

For binary tasks, we use a binary cross-entropy loss criterion weighted by the class distribution (e.g. a 1:4 positive:negative ratio would result in positive examples being up-weighted four-fold). For regression tasks, we use a mean-squared error criterion. For all models, we train with the Adam optimizer with learning rate 0.0001 (except for PIP, which uses a learning rate of 0.001) and monitor the relevant metrics (see Table 8) on the validation set after every epoch. The weights of the best-performing are then used to evaluate on the held-out test set. Models were all trained using 1 Tesla V100 GPU for 4–48 hours depending on the task.

ENNs For the core of all Cormorant architectures in this work, we use a network of four layers of covariant neurons that use the Clebsch–Gordan transform as nonlinearity, with $L = 3$ as the largest index in the $SO(3)$ representation and 16 channels, followed by a single $SO(3)$ -vector layer with $L = 0$. An input featurization network encodes the atom types as one-hot vectors. For SMP, input and output are passed through multi-layer perceptrons (MLP) as in Anderson et al. (2019). For the input, a weighted adjacency matrix with a learnable cut-off radius is constructed. This mask is passed alongside the input vector through a MLP with a single hidden layer with 256 neurons and ReLU activation. The output network is constructed from a set of scalar invariants that are passed through a network of two MLPs. Each of these MLPs has a single hidden layer of size 256, and the intermediate representation has 96 neurons. For all other tasks, input and output layers are a single learnable mixing matrix, as used in the original Cormorant implementation for MD-17 (Anderson et al., 2019). The Siamese architecture required for MSP and LEP was constructed by training two ENNs that are then connected by passing on the squared difference of the single-network outputs: For MSP, the two structures corresponded to the wild-type structure and the mutated one; for LEP to the active and inactive one. We extend the original Cormorant implementation to handle classification problems (binary and multi-class) and the Siamese network architecture.

We use MSE loss for regression tasks and cross-entropy loss for classification tasks. For all tasks, we used the AMSgrad optimizer with an initial learning rate of 0.001 and a final learning rate of 0.00001, decaying in a cosine function over the training process. We trained SMP for 255 epochs, LBA for 150 epochs, RES for 30 epochs, and MSP for 50 epochs. We monitor the loss for the validation set after every epoch. The weights of the best-performing are then used to evaluate on the held-out test set. The models were all trained on 1 Titan X(p) GPU for 1–5 days depending on the task.

C.3 1D AND 2D BASELINES

For each task, we select a baseline that fulfills the following criteria: (1) represents the current state-of-the-art (SOTA) for that task—or as close as possible—using only 1D (sequence only) or 2D (sequence and/or bond connectivity) molecular representations, and (2) either has a publicly available implementation or has reported results for the same task and splitting criteria. For PSR and RSR, which are inherently 3D tasks and have no appropriate 1D or 2D representation, we compare to the state-of-the-art 3D methods instead. Below we describe the choice and implementation of baseline models for each task.

SMP As a 2D method for predicting molecular properties, we choose molecular GNNs (Tsubaki et al., 2019) which are based on learning representations of subgraphs in molecules. We use an implementation that only uses the SMILES representation of the molecular graph.⁴ As an additional 2D baseline, we compare to N-Gram Graph XGB (Liu et al., 2019). This method is based on an unsupervised representation called N-gram graph which first embeds the vertices in the molecule graph and then assembles the vertex embeddings in short walks in the graph. This representation is combined with the XGBoost learning method (Chen & Guestrin, 2016).⁵

⁴https://github.com/masashitsubaki/molecularGNN_smiles

⁵https://github.com/chaol224/n_gram_graph

PIP For the PIP task, our non-3D method is the BIPSPI (Sanchez-Garcia et al., 2018) model, a gradient-boosted decision tree. We compare to their model that uses only sequence and sequence conservation features and is evaluated on DB5.

RES As a 1D sequence-based model for predicting residue identity, we choose the transformer architecture TAPE, introduced by Rao et al. (2019). We use their reported accuracy on heldout families for language modeling, as that corresponds to a sequence-only version of our RES tasks, with similar stringency in terms of splitting criteria.

MSP We use the publicly provided implementation of TAPE (Rao et al., 2019)⁶. We modify their sequence-to-sequence head to predict the effect of mutations at specific positions, using the original unmutated protein as the input sequence and writing the output sequence as a one-hot-encoded 20-dimensional vector, indicating if a given mutation would be beneficial or detrimental. Note that the vast majority of positions would be unlabeled and therefore not included in the learning task.

LBA As a 1D method for predicting ligand binding affinity, we choose DeepDTA (Öztürk et al., 2018)⁷, a 1DCNN based model that takes in pairs of ligand SMILES string and protein sequence as input. We use the same hyperparameters as in the original paper for the baseline.

We also compare our results against DeepAffinity (Karimi et al., 2019)⁸. We compare to their unified RNN/GCNN-CNN model that takes in pairs of ligand SMILES string and their novel representations of structurally-annotated protein sequences (SPS/Structure Property-annotated Sequence) as input. Per the authors’ recommendation, we use the DSSP software (Joosten et al., 2010; Kabsch & Sander, 1983) to generate the protein secondary structure and the protein relative solvent accessibility used in the SPS representation directly from the protein 3D structure, rather than the predicted ones by the SSpro/ACCpro software (Magnan & Baldi, 2014; Cheng et al., 2005) as done in the DeepAffinity paper. We use the same hyperparameters as in the original paper for the baseline, except for the maximum compound and protein sizes which we increase to 160 and 168, respectively, to accommodate for larger compounds/proteins in the PDDBind dataset. We trained the DeepAffinity models for 1000 epochs.

LEP We also train DeepDTA (Öztürk et al., 2018)⁷ (with the same hyperparameters as in the original paper) on the LEP dataset as baseline. As the inherent protein sequences and ligand SMILES strings are the same for both the inactive and active structures, the problem is reduced to binary classification task given a pair of protein sequence and the ligand SMILES string, and does not require modifying the DeepDTA architecture to make the Siamese network as in the GNN, ENN, or 3DCNN case.

PSR We compare our results against the state-of-the-art single-model methods as reported in Pagès et al. (2019). These include 3DCNN (Hou et al., 2019) and Ornate (Pagès et al., 2019), 3DCNN voxel-based methods trained on structural information, and Proq3D (Uziela et al., 2017), a deep-learning based method which employs structural information, Rosetta energy terms (Leaver-Fay et al., 2011), and evolutionary information derived from the amino acid sequence. We exclude ProteinGCN (Sanyal et al., 2020), a recent GNN-based method, from comparison as they do not provide results on CASP11 dataset.

RSR For RNA structure ranking, we compare our results against the Rosetta scoring function (Alford et al., 2017). In past RNA Puzzles competitions, methods using the Rosetta scoring function have been found to most consistently produce the lowest RMSD candidates. This is a physical- and knowledge-based potential specifically tuned for biomolecular structure.

C.4 STATE-OF-THE-ART METHODS

When possible, for tasks in Table D, we choose 3D methods that fulfill the following criteria: (1) they represent the current state-of-the-art for that task, or as close as possible, and (2) they either have

⁶<https://github.com/songlab-cal/tape>

⁷<https://github.com/hkmztrk/DeepDTA>

⁸<https://github.com/Shen-Lab/DeepAffinity>

a publicly available implementation or has reported results for the same task and splitting criteria. Here our choice of methods is described in more detail if not already discussed in the section above.

SMP We compare to the state of the art, i.e., the best achieved prediction on each task, as reported in Anderson et al. (2019). Many methods have been tested on QM9 and have reached excellent performance which makes them comparably hard to beat for new methods. In general, the best methods for QM9 so far are message passing neural networks Gilmer et al. (2017), continuous-filter convolutional neural networks Schütt et al. (2017), and Cormorant Anderson et al. (2019). Differences in performance between earlier Cormorant studies and this work can be attributed to the different (random) split.

PIP We compare our results against the BIPSPI Sanchez-Garcia et al. (2018) model, a gradient-boosted decision tree. In contrast to the 1D/2D baseline comparison to BIPSPI, in this case we compare against their model that employs both structural- and sequence-based amino acid features.

RES Since there have been no standardized datasets for this task to date, it is difficult to perform a direct comparison of methods. The closest comparison for a CNN trained on a balanced dataset of residue environments is 0.425, as reported in Torng & Altman (2017). While higher performance was reported by Anand et al. (2020) (accuracy 0.572), this model was trained on an unbalanced dataset comprising every standard residue environment in all training set PDBs. Similar performance has also been reported with other deep learning architectures Weiler et al. (2018); Boomsma & Frellsen (2017), but these do not describe their training/evaluation data or splitting criteria. In contrast, we restrict our training and evaluation to a balanced subset, downsampled to the frequency of the rarest class, which limits performance slightly. Additionally, to enable fair comparison over three replicates between 3DCNN and GNN, we then trained on only half of these down-sampled environments. The discrepancy in performance we observe on this subset is indicative of the fact that the differences in residue environment are subtle and complex, so simply increasing training data can result in higher performance. This is especially true for common classes such as leucine and glycine, which are over five times as frequent than the least common class, cysteine. Within these common classes, accuracy exceeds 80%, which increases the average accuracy when classes are imbalanced.

LEP Because this was a novel dataset, we computed initial results a non-deep learning method, Schrödinger’s Glide, to score each protein-ligand complex. Glide is state-of-the-art for scoring protein-ligand complexes and determining how “good” a pose is. This resulted in 2 scores per ligand; the score to the inactive protein structure and the score to the active protein structure. We then performed a binary classification by training an SVM on these two features to predict the ligand activity class. This approach is reasonable from a physical basis: if the ligand binds much better to the active protein structure than the inactive protein structure, then it will be an activator of the protein’s function.

LBA Many methods have been developed for the prediction of ligand binding affinity using the PDBBind dataset. However, the standard has been to evaluate performance on the so-called “core set”, as described in Section C.1, after training and validating on the remainder of the refined set. The state-of-the-art reported on this core set has been achieved by the 3DCNN-based KDEEP Jiménez et al. (2018), followed closely by the popular random forest-based method RF-score Ballester & Mitchell (2010). However, because the core set contains only proteins that are also present in the training set, this only measures *in-distribution* performance, not generalizable scoring ability. Thus, the most comparable baseline for our dataset, which was split at 30% sequence identity, is the performance of the empirical linear regression-based scoring function X-score fitted to complexes with less than 30% identity to the core set, as reported in Li & Yang (2017). We note that this is not a perfect comparison, since the procedure used in Li & Yang (2017) reduces the size of the training set significantly; however, as an empirical scoring function the performance of X-score is not very sensitive to training set size, compared to RF-score, which was significantly affected.

D COMPARISON TO STATE-OF-THE-ART METHODS

Table 7: Comparison of performance against state-of-the-art methods, where available. The 3DCNN, GNN, and ENN networks achieve state-of-the-art in several tasks; for those where they do not (SMP, PIP, LBA), we note that the competing methods also use the 3D geometry of molecules. Asterisks (*) indicate that the exact training data differed (though splitting criteria were the same).

Task	Metric	3D			SOTA
		3DCNN	GNN	ENN	
SMP	μ [D]	0.572	0.068	0.046	*0.030 (Gilmer et al., 2017)
	ε_{gap} [eV]	0.589	0.091	0.065	*0.061 (Anderson et al., 2019)
	U_0^{at} [eV]	1.615	0.070	0.023	*0.014 (Schütt et al., 2017)
PIP	AUROC	0.844	*0.669	—	0.919 (Sanchez-Garcia et al., 2018)
RES	accuracy	0.451	0.082	*0.072	*0.425 (Torng & Altman, 2017)
MSP	AUROC	0.520	0.637	0.678	—
LBA	RMSD	1.520	1.936	*1.429	*1.838 (Li & Yang, 2017)
	glob. R_P	0.558	0.581	*0.541	*0.645 (Li & Yang, 2017)
	glob. R_S	0.556	0.647	*0.532	*0.697 (Li & Yang, 2017)
LEP	AUROC	0.824	0.678	0.569	0.770 (Friesner et al., 2004)
PSR	mean R_S	0.177	0.327	—	0.432 (Pagès et al., 2019)
	glob. R_S	0.837	0.716	—	0.796 (Pagès et al., 2019)
RSR	mean R_S	0.414	0.195	—	0.173 (Alford et al., 2017)
	glob. R_S	0.656	0.309	—	0.304 (Alford et al., 2017)

E COMPLETE BENCHMARKING RESULTS

Table 8: Complete benchmarking results from Tables 3–6, with additional metrics and standard deviations reported over three replicates. R_K is Kendall correlation and AUPRC is area under the precision-recall curve. SMP metrics are all mean absolute error (MAE). Asterisks (*) indicate that the exact training data differed (though splitting criteria were the same).

Task	Metric	3DCNN	GNN	ENN	SOTA Baseline (Pagès et al., 2019)
PSR	mean R_P	0.191 \pm 0.017	0.388 \pm 0.022	—	0.444
	mean R_K	0.122 \pm 0.008	0.227 \pm 0.005	—	0.304
	mean R_S	0.177 \pm 0.012	0.327 \pm 0.327	—	0.432
	global R_P	0.836 \pm 0.019	0.703 \pm 0.020	—	0.772
	global R_K	0.668 \pm 0.024	0.515 \pm 0.015	—	0.594
	global R_S	0.837 \pm 0.023	0.716 \pm 0.013	—	0.796
SOTA Baseline (Alford et al., 2017)					
RSR	mean R_P	0.428 \pm 0.018	0.219 \pm 0.284	—	0.129
	mean R_K	0.287 \pm 0.014	0.136 \pm 0.183	—	0.119
	mean R_S	0.414 \pm 0.018	0.195 \pm 0.265	—	0.173
	global R_P	0.613 \pm 0.007	0.203 \pm 0.023	—	0.161
	global R_K	0.462 \pm 0.006	0.195 \pm 0.022	—	0.206
	global R_S	0.656 \pm 0.007	0.309 \pm 0.041	—	0.304
Non-3D Baseline (Sanchez-Garcia et al., 2018)					
PIP	AUROC	0.844 \pm 0.002	*0.669 \pm 0.001	—	0.841
Non-3D Baseline (Rao et al., 2019)					
RES	accuracy	0.451 \pm 0.002	0.082 \pm 0.002	*0.072 \pm 0.005	*0.30
MSP	AUROC	0.520 \pm 0.015	0.637 \pm 0.020	0.678 \pm 0.092	0.554
	AUPRC	0.269 \pm 0.009	0.258 \pm 0.017	0.304 \pm 0.137	
Non-3D Baseline (Tsubaki et al., 2019)					
SMP	μ [D]	0.572 \pm 0.005	0.068 \pm 0.001	0.046 \pm 0.008	0.496 \pm 0.002
	ε_{gap} [eV]	0.589 \pm 0.012	0.091 \pm 0.003	0.065 \pm 0.001	0.154 \pm 0.001
	U_0^{at} [eV]	1.615 \pm 0.040	0.070 \pm 0.009	0.023 \pm 0.002	0.182 \pm 0.004
	α [bohr ³]	2.090 \pm 0.240	0.207 \pm 0.010	0.183 \pm 0.090	0.392 \pm 0.004
	$\varepsilon_{\text{HOMO}}$ [eV]	0.295 \pm 0.009	0.060 \pm 0.002	0.038 \pm 0.002	0.107 \pm 0.001
	$\varepsilon_{\text{LUMO}}$ [eV]	0.499 \pm 0.008	0.056 \pm 0.001	0.081 \pm 0.014	0.115 \pm 0.001
	R^2 [bohr ²]	65.653 \pm 2.064	0.629 \pm 0.113	1.139 \pm 0.071	27.976 \pm 0.212
	ZPVE [meV]	71.410 \pm 4.328	4.677 \pm 0.097	1.727 \pm 0.029	10.614 \pm 0.270
	U_{298}^{at} [eV]	1.700 \pm 0.123	0.070 \pm 0.005	0.023 \pm 0.000	0.181 \pm 0.001
	H_{298}^{at} [eV]	1.744 \pm 0.049	0.070 \pm 0.005	0.023 \pm 0.001	0.180 \pm 0.004
	G_{298}^{at} [eV]	1.633 \pm 0.053	0.067 \pm 0.009	0.022 \pm 0.000	0.173 \pm 0.000
	C_v [$\frac{\text{cal}}{\text{molK}}$]	1.218 \pm 0.145	0.074 \pm 0.003	0.040 \pm 0.010	0.187 \pm 0.004
Non-3D Baseline (Öztürk et al., 2018)					
LBA	RMSD	1.520 \pm 0.009	1.936 \pm 0.120	* 1.429 \pm 0.042	1.565 \pm 0.018
(30%)	global R_P	0.558 \pm 0.005	0.581 \pm 0.039	*0.541 \pm 0.029	0.573 \pm 0.022
	global R_S	0.556 \pm 0.004	0.647 \pm 0.071	*0.532 \pm 0.033	0.574 \pm 0.010
LBA	RMSD	1.517 \pm 0.032	1.493 \pm 0.010	* 1.450 \pm 0.024	1.760 \pm 0.415
(60%)	global R_P	0.669 \pm 0.015	0.669 \pm 0.013	* 0.716 \pm 0.008	0.713 \pm 0.013
	global R_S	0.656 \pm 0.012	0.691 \pm 0.010	* 0.714 \pm 0.009	0.702 \pm 0.013
LEP	AUROC	0.824 \pm 0.034	0.678 \pm 0.034	0.569 \pm 0.030	0.696 \pm 0.021
	AUPRC	0.714 \pm 0.031	0.525 \pm 0.051	0.519 \pm 0.034	0.550 \pm 0.024