
High Performance of Gradient Boosting in Binding Affinity Prediction

Dmitrii Gavrilev¹ Nurlybek Amangeldiuly¹ Sergei Ivanov^{2,1} Evgeny Burnaev¹

Abstract

Prediction of protein-ligand (PL) binding affinity remains the key to drug discovery. Popular approaches in recent years involve graph neural networks (GNNs), which are used to learn the topology and geometry of PL complexes. However, GNNs are computationally heavy and have poor scalability to graph sizes. On the other hand, traditional machine learning (ML) approaches, such as gradient-boosted decision trees (GBDTs), are lightweight yet extremely efficient for tabular data. We propose to use PL interaction features along with PL graph-level features in GBDT. We show that this combination outperforms the existing solutions.

1. Introduction

Binding affinity is one of the major determinants to identify whether a molecule binds to a desired macromolecular target. It helps speculate if the molecule can induce a biological response in an organism, which is the central mechanism in drug action. Reliable and fast prediction of such parameter is an important and challenging task in cheminformatics (Wan et al., 2020). To build an accurate binding affinity prediction model PL complex 3D structures are needed (Son & Kim, 2021), and are readily accessible from steadily growing open source databases (Jumper et al., 2021). While docking is an established approach to predict affinity (Kitchen et al., 2004), recent ML-based algorithms consistently outperform existing classical methods (Meli et al., 2021). Considered a natural way of processing chemical structures, recently GNNs are proposed to predict binding affinity. In a typical prediction pipeline for cheminformatics, a GNN generalizes to a labeled and weighted graph, nodes of which represent molecular atoms, and edges may represent their bonds by including various interaction features such as bond angles and distances (Jones et al.,

2021; Son & Kim, 2021). This application domain of GNN can lead to a promising results (Kearnes et al., 2016). The main bottleneck of GNNs is, however, their computationally complex iterative convergence (Gupta et al., 2020; Tiezzi et al., 2020).

In contrast, GBDTs provide faster convergence and perform out-of-the-box, which makes them a popular tool for ML contests and industrial applications (Gorishniy et al., 2021). Notwithstanding the advantages, powerful model of GBDTs often overlooked (Qin et al., 2021).

In this work, we leverage the advantages of GBDTs to graph-level structural and interaction feature representations of PL complexes for binding affinity prediction. Assessment of the model performance on external test sets showed that the GBDTs can perform equally well or outperform the recent state-of-the-art GNN-based approach, SIGN (Li et al., 2021).

2. Related Work

The early ML methods of binding affinity prediction rely on careful feature engineering. Namely, RF-Score proposed training a Random Forest model on intermolecular interaction features (Ballester & Mitchell, 2010). For sets of atom types P in protein and L in ligand, the authors count the number of atom type pairs in the complex. More formally, the resulting interaction features vector x is $|P| \times |L|$ dimensional and

$$x_{ji} = \sum_{a \in \text{protein}} \sum_{b \in \text{ligand}} \delta_{aj} \delta_{bi} \Theta(d_{\text{cutoff}} - d_{ab}),$$

where δ and Θ are the Kronecker and Heavyside functions respectively, d_{cutoff} is the interaction cutoff distance and d_{ab} is the distance between atoms a and b . Particularly, the authors set $d_{\text{cutoff}} = 12\text{\AA}$. As for P and L , the following sets are recommended:

$$P = \{C, N, O, S\}, \\ L = \{C, N, O, F, P, S, Cl, Br, I\}.$$

Therefore, the interaction features can be described by a 36-dimensional vector. Designing representations for the molecules has also received a great deal of attention in the context of graph kernels (Kriege et al., 2020; Ivanov

^{*}Equal contribution ¹Skolkovo Institute of Science and Technology ²Criteo AI. Correspondence to: Dmitrii Gavrilev <Dmitrii.Gavrilev@skoltech.edu>.

Table 1. Metrics on PDBbind core and CSAR-HiQ sets.

| Models | PDBbind core set | | | | CSAR-HiQ set | | | |
|-----------------------|------------------|-------|-------|-------|--------------|-------|-------|-------|
| | RMSE ↓ | MAE ↓ | SD ↓ | R ↑ | RMSE ↓ | MAE ↓ | SD ↓ | R ↑ |
| LR | 1.608 | 1.285 | 1.576 | 0.690 | 1.994 | 1.555 | 1.909 | 0.650 |
| MLP | 1.456 | 1.158 | 1.454 | 0.744 | 2.195 | 1.727 | 2.046 | 0.575 |
| SIGN | 1.316 | 1.027 | 1.312 | 0.797 | 1.735 | 1.327 | 1.709 | 0.754 |
| Ours: CatBoost | 1.321 | 1.045 | 1.270 | 0.812 | 1.798 | 1.391 | 1.679 | 0.744 |
| Ours: LightGBM | 1.316 | 1.040 | 1.279 | 0.809 | 1.725 | 1.305 | 1.660 | 0.751 |

& Burnaev, 2018) which have been applied successfully for various applications on graphs (Sharaev et al., 2018). On the other hand, recent trends in binding affinity prediction involve extensive usage of GNNs to obtain spatial representations. Such approach also allows characterizing intermolecular and intramolecular interactions. To enrich the representation, the efforts are being directed towards integrating atomic distances (Lim et al., 2019), bond angles and physics-derived descriptors (Moon et al., 2022). In such work (Jiang et al., 2021), authors construct protein graphs, ligand graphs and PL graphs separately to represent PL structure and interaction by also integrating distance, angle, and various 2D node and edge features. Another approach, graphDelta (Karlov et al., 2020) characterizes binding site of the complex by computing Behler-Parrinello symmetry functions (Behler & Parrinello, 2007) per each common PL heavy atom types in the complex and obtains complex interaction features to train a message-passing neural network.

Very recent work, SIGN exploits the interaction features similar to RF-Score. Moreover, the authors of SIGN propose to use not only topological information but also angles and distances. Despite outperforming ML and NN-based models, following drawbacks inherent to GNN models are not solved in SIGN:

- Time-consuming construction of graphs from raw 3D coordinates (Wang & Xia, 2012);
- Requires large computational resources for training (Liu et al., 2020).

We provide time and memory consumption details of SIGN in Section 4.

3. Our Approach

In this work, we incorporate PL interaction features, as introduced in RF-Score, alongside graph-level features. To construct graph-level features, the atomic features are pooled together. For each atomic feature, we sum the values over all atoms in the complex. Additionally, we calculate the standard deviation of atomic features to characterize atomic distribution statistics. Categorical features are encoded to a one-hot numeric array: for example, atom types are trans-

formed into a binary array, where each bit encodes a type. Then, the graph-level feature constructed from these arrays represents the distribution of atom types in a complex. Below we list the features that we used to train our models:

- RF-Score PL interaction features;
- Summarized atomic features;
- Standard deviation of atomic features.

We follow the same feature extraction procedure as the authors of SIGN and Pafnucy with the help of Open Babel toolbox (Stepniewska-Dziubinska et al., 2018; O’Boyle et al., 2011). Each atom is represented by a 36-dimensional vector, 18 of which describe protein features and the rest describe ligand features. Those features include

- Binary: atom type, aromaticity, hydrophobicity, acceptor, donor, ring;
- Integer: hybridization, number of non-hydrogen atoms attached, number of heteroatoms attached;
- Float: partial charge.

As a result, we obtain 108 features for each complex — 36 from the interaction matrix, 36 by summing atomic features, and the other 36 by calculating standard deviation.

After obtaining features, we pass them to a GBDT model. GBDT (Friedman, 2001) is a widely used algorithm that builds ensemble of decision trees. The idea of GBDT is to iteratively add weak decision trees and aggregate their predictions to obtain an overall strong model. At each iteration the model is updated by building a new tree that minimizes the loss function w.r.t. previous predictions of the ensemble. In particular, at each iteration t the model $f(\mathbf{x})$ is updated in an additive manner:

$$f^t(\mathbf{x}) = f^{t-1}(\mathbf{x}) + \epsilon h^t(\mathbf{x}), \tag{1}$$

where f^{t-1} is previously built trees, h^t is a weak learner that is chosen from some family of functions \mathcal{H} (e.g. decision trees), and ϵ is a learning rate. The weak learner $h^t \in \mathcal{H}$ approximates the direction of the negative gradient of a loss function L w.r.t. the current model’s predictions:

$$h^t = \arg \min_{h \in \mathcal{H}} \sum_i \left(-\frac{\partial L(f^{t-1}(\mathbf{x}_i), y_i)}{\partial f^{t-1}(\mathbf{x}_i)} - h(\mathbf{x}_i) \right)^2. \tag{2}$$

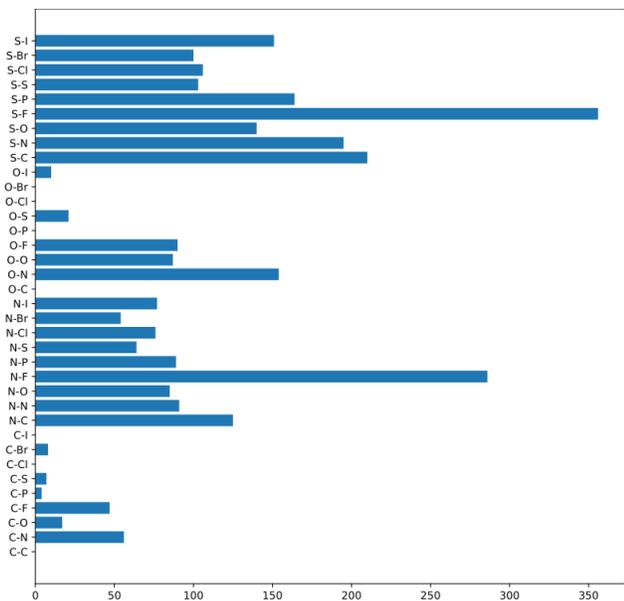


Figure 1. Feature importance (interaction features).

As such one can think of GBDT’s training as functional gradient descent, where at each iteration we move in the direction that minimizes the loss. This allows flexibly training GBDT on various tabular continuous/discrete, noisy and heterogeneous data.

4. Experiments

Data. For evaluating the performance of the models, we use the same dataset and splits as in (Li et al., 2021). Specifically, all experiments were conducted on PDBbind v2016 (Liu et al., 2015) and CSAR-HiQ (Dunbar Jr et al., 2011), which contain experimental PL binding affinity data. The core set of PDBbind (290 complexes) was used as a benchmark for model comparison (Su et al., 2018). The models were trained on the difference between the refined and core sets (3767 complexes). For tuning hyperparameters, 377 out of 3767 complexes were used as a validation set. After training and tuning, the models were additionally evaluated on CSAR-HiQ dataset. Since CSAR-HiQ and the refined set of PDBbind overlap, we filter out these overlapping complexes for evaluation. The filtered dataset consists of 135 complexes.

Models. Our work focuses mainly on GBDT models since they have shown great performance for tabular data in both time and quality. We compare two popular GBDT models, CatBoost (Dorogush et al., 2018) and LightGBM (Ke et al., 2017), with SIGN. First, we train the models on PDBbind refined set and tune their hyperparameters on validation set.

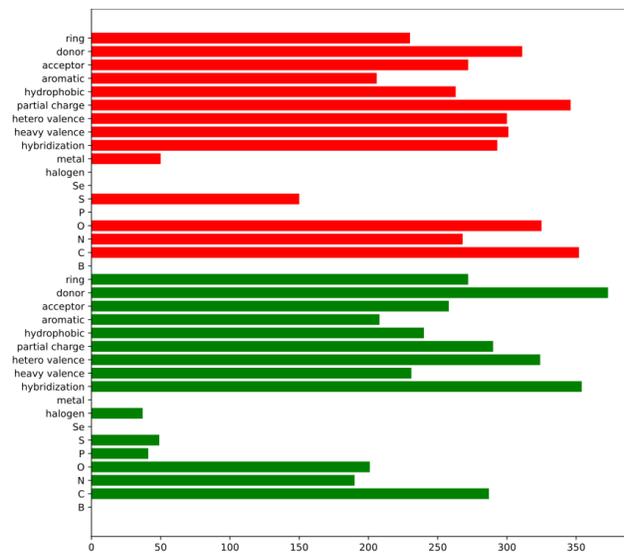


Figure 2. Feature importance (sum of atomic features).

For tuning, we utilize Optuna (Akiba et al., 2019), an automatic hyperparameter search framework. While validating, we aim to minimize root mean squared error (RMSE). The best hyperparameters for CatBoost are following: depth = 10, grow policy = symmetric tree, L2 leaf regularization = 73.47. As for LightGBM, the optimal hyperparameters are: bagging fraction = 1, bagging frequency = 0, feature fraction = 0.8, L1 regularization = 0, L2 regularization = 0, minimal child samples = 20, number of leaves = 244.

Next, we evaluate the models on the PDBbind core set. Table 1 summarizes our main findings. The following metrics are reported: RMSE (root mean squared error), MAE (mean absolute error), SD (standard deviation), R (Pearson’s correlation). The metrics are averaged across 5 different runs. On the core set, LightGBM has the same RMSE as SIGN on average, while CatBoost performs better in terms of SD and R. The only metric on which SIGN outperforms GBDT model is MAE. Besides GBDT models, we also provide the results for linear regression and three-layer perceptron.

To assess the generalization ability, the models are further evaluated on CSAR-HiQ set. As can be seen from Table 1, classical GBDT models perform on par with SIGN method. LightGBM demonstrates slight improvement over SIGN in terms of RMSE and is able to outperform all presented models. Nevertheless, SIGN shows better R score. It should be noted that GBDTs achieved these results easily, owing to efficiency and scalability of their architecture.

Feature importance. Finally, we obtain the feature importance to analyze their contribution. Figures 2 and 3 show the degree of contribution of atomic features (their sum and

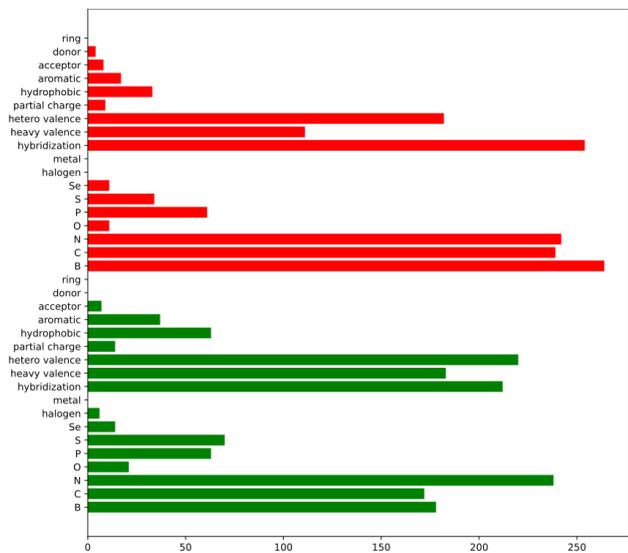


Figure 3. Feature importance (standard deviation of atomic features).

standard deviation). Green and red bars correspond to ligand and protein features respectively. Figure 1 illustrates the importance of interaction features. It is notable that mainly uncommon interactions in the complexes, such as $S - F$, $N - F$ have the highest importance in decision making of the model. It is also important to note that the sum of atomic features has a more even distribution of importance than those of interaction features while having the same peak in magnitude. This might suggest that pooling atomic-level features into graph-level features can significantly boost the performance of RF-Score.

Training time and memory consumption. We also report the benefits of using GBDTs in terms of speed and memory consumption. All experiments were conducted on 22 CPU cores and 1 NVIDIA A10 GPU (GPU was used for training SIGN). While training SIGN took 11 hours (300 epochs with batch size set to 16), training CatBoost took only 3 minutes. Moreover, LightGBM has been trained only under 4 seconds in total. As for memory consumption, SIGN allocates 19 GB of RAM and 13 GB of VRAM, whereas GBDM models required less than 1 GB of RAM.

5. Discussion and Future Work

In this work, we looked back at traditional ML methods for predicting PL binding affinity and proposed a way to improve their performance. We showed that the inclusion of graph-level features with their standard deviation results in better metrics. To our surprise, GBDM models can be comparable to heavy GNN models without much effort. We

hope that our findings would bring attention back to simpler models, as there might be more undiscovered potential. Next, we plan to integrate GBDM and GNN into one end-to-end architecture to utilize the advantages of heterogeneous learning and representation learning, as recently reported in our work, BGNN (Ivanov & Prokhorenkova, 2021). The approach demonstrated state-of-the-art performance and efficiency results on node-level prediction tasks, and adapting the architecture to process the graph-level molecular features is a promising direction.

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.
- Ballester, P. J. and Mitchell, J. B. A machine learning approach to predicting protein–ligand binding affinity with applications to molecular docking. *Bioinformatics*, 26(9), 2010.
- Behler, J. and Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98, 2007.
- Dorogush, A. V., Ershov, V., and Gulin, A. Catboost: gradient boosting with categorical features support. *arXiv preprint arXiv:1810.11363*, 2018.
- Dunbar Jr, J. B., Smith, R. D., Yang, C.-Y., Ung, P. M.-U., Lexa, K. W., Khazanov, N. A., Stuckey, J. A., Wang, S., and Carlson, H. A. Csr benchmark exercise of 2010: selection of the protein–ligand complexes. *Journal of chemical information and modeling*, 51(9), 2011.
- Friedman, J. H. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 2001.
- Gorishniy, Y., Rubachev, I., Khrulkov, V., and Babenko, A. Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems*, 34, 2021.
- Gupta, P., Gasse, M., Khalil, E., Mudigonda, P., Lodi, A., and Bengio, Y. Hybrid models for learning to branch. *Advances in neural information processing systems*, 33, 2020.
- Ivanov, S. and Burnaev, E. Anonymous walk embeddings. In *International conference on machine learning*. PMLR, 2018.
- Ivanov, S. and Prokhorenkova, L. Boost then convolve: Gradient boosting meets graph neural networks. *arXiv preprint arXiv:2101.08543*, 2021.

- Jiang, D., Hsieh, C.-Y., Wu, Z., Kang, Y., Wang, J., Wang, E., Liao, B., Shen, C., Xu, L., Wu, J., et al. Interaction-graphnet: A novel and efficient deep graph representation learning framework for accurate protein–ligand interaction predictions. *Journal of medicinal chemistry*, 64(24), 2021.
- Jones, D., Kim, H., Zhang, X., Zemla, A., Stevenson, G., Bennett, W. D., Kirshner, D., Wong, S. E., Lightstone, F. C., and Allen, J. E. Improved protein–ligand binding affinity prediction with structure-based deep fusion inference. *Journal of chemical information and modeling*, 61(4), 2021.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnoy, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873), 2021.
- Karlov, D. S., Sosnin, S., Fedorov, M. V., and Popov, P. graphdelta: Mpnn scoring function for the affinity prediction of protein–ligand complexes. *ACS omega*, 5(10), 2020.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30, 2017.
- Kearnes, S., McCloskey, K., Berndl, M., Pande, V., and Riley, P. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8), 2016.
- Kitchen, D. B., Decornez, H., Furr, J. R., and Bajorath, J. Docking and scoring in virtual screening for drug discovery: methods and applications. *Nature Reviews Drug Discovery*, 3(11), 2004.
- Kriege, N. M., Johansson, F. D., and Morris, C. A survey on graph kernels. *Applied Network Science*, 5(1), 2020.
- Li, S., Zhou, J., Xu, T., Huang, L., Wang, F., Xiong, H., Huang, W., Dou, D., and Xiong, H. Structure-aware interactive graph neural networks for the prediction of protein–ligand binding affinity. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021.
- Lim, J., Ryu, S., Park, K., Choe, Y. J., Ham, J., and Kim, W. Y. Predicting drug–target interaction using a novel graph neural network with 3d structure-embedded graph representation. *Journal of Chemical Information and Modeling*, 59(9), 2019.
- Liu, S., Park, J. H., and Yoo, S. Efficient and effective graph convolution networks. In *Proceedings of the 2020 SIAM International Conference on Data Mining*. SIAM, 2020.
- Liu, Z., Li, Y., Han, L., Li, J., Liu, J., Zhao, Z., Nie, W., Liu, Y., and Wang, R. Pdb-wide collection of binding data: current status of the pdbname database. *Bioinformatics*, 31(3), 2015.
- Meli, R., Anighoro, A., Bodkin, M. J., Morris, G. M., and Biggin, P. C. Learning protein–ligand binding affinity with atomic environment vectors. *Journal of Cheminformatics*, 13(1), 2021.
- Moon, S., Zhung, W., Yang, S., Lim, J., and Kim, W. Y. Pignet: a physics-informed deep learning model toward generalized drug–target interaction predictions. *Chemical Science*, 13(13), 2022.
- O’Boyle, N. M., Banck, M., James, C. A., Morley, C., Van dermeersch, T., and Hutchison, G. R. Open babel: An open chemical toolbox. *Journal of cheminformatics*, 3(1): 1–14, 2011.
- Qin, Z., Yan, L., Zhuang, H., Tay, Y., Pasumarthi, R. K., Wang, X., Bendersky, M., and Najork, M. Are neural rankers still outperformed by gradient boosted decision trees? In *International Conference on Learning Representations (ICLR)*, 2021.
- Sharaev, M., Artemov, A., Kondrateva, E., Ivanov, S., Sushchinskaya, S., Bernstein, A., Cichocki, A., and Burnaev, E. Learning connectivity patterns via graph kernels for fmri-based depression diagnostics. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, 2018.
- Son, J. and Kim, D. Development of a graph convolutional neural network model for efficient prediction of protein–ligand binding affinities. *PLoS ONE*, 16(4), 2021.
- Stepniewska-Dziubinska, M. M., Zielenkiewicz, P., and Siedlecki, P. Development and evaluation of a deep learning model for protein–ligand binding affinity prediction. *Bioinformatics*, 34(21):3666–3674, 2018.
- Su, M., Yang, Q., Du, Y., Feng, G., Liu, Z., Li, Y., and Wang, R. Comparative assessment of scoring functions: the casf-2016 update. *Journal of chemical information and modeling*, 59(2), 2018.
- Tiezzi, M., Marra, G., Melacci, S., Maggini, M., and Gori, M. A lagrangian approach to information propagation in graph neural networks. *arXiv preprint arXiv:2002.07684*, 2020.
- Wan, S., Bhati, A. P., Zasada, S. J., and Coveney, P. V. Rapid, accurate, precise and reproducible ligand–protein binding free energy prediction. *Interface Focus*, 10(20200007), 2020.
- Wang, J. and Xia, Y. Fast graph construction using auction algorithm. *arXiv preprint arXiv:1210.4917*, 2012.