

---

# DualBind: A Dual-Loss Framework for Protein-Ligand Binding Affinity Prediction

---

Meng Liu<sup>1</sup> Saeed Gopal Paliwal<sup>1</sup>

## Abstract

Accurate prediction of protein-ligand binding affinities is crucial for drug development. Recent advances in machine learning show promising results on this task. However, these methods typically rely heavily on labeled data, which can be scarce or unreliable, or they rely on assumptions like Boltzmann-distributed data that may not hold true in practice. Here, we present DualBind, a novel framework that integrates supervised mean squared error (MSE) with unsupervised denoising score matching (DSM) to accurately learn the binding energy function. DualBind not only addresses the limitations of DSM-only models by providing more accurate absolute affinity predictions but also improves generalizability and reduces reliance on labeled data compared to MSE-only models. Our experimental results demonstrate that DualBind excels in predicting binding affinities and can effectively utilize both labeled and unlabeled data to enhance performance.

## 1. Introduction

Accurate prediction of protein-ligand binding energies is a fundamental task in drug discovery workflow, as it directly influences the prioritization and optimization of potential drug candidates by evaluating their interaction strength with the target. This task presents significant challenges due to the complex nature of protein-ligand interactions and the intricate potential energy landscape they create. Traditionally, various scoring functions designed by experts, such as those in Glide (Friesner et al., 2004) and AutoDock Vina (Trott & Olson, 2010; Eberhardt et al., 2021), have been employed to estimate binding affinities. In recent years, data-driven machine learning methods have also emerged, offering new avenues for enhancing prediction accuracy through flexible and learnable neural networks (Jiménez et al., 2018; Öztürk

et al., 2018; Chen et al., 2020; Li et al., 2020; Nguyen et al., 2021; Lu et al., 2022; Jiang et al., 2021; Moon et al., 2022; Zhang et al., 2023a; Jin et al., 2023b; Lin et al., 2024; Zhang et al., 2023b).

The most straightforward approach in machine learning is to train a model that minimizes prediction errors directly, requiring reliable binding affinity labels for protein-ligand complexes. Notably, models relying solely on a supervised loss might overfit to noisy or sparse data, reducing their robustness and generalizability. Recently, DSMBind (Jin et al., 2023b;a) adopted a generative modeling strategy by training an energy-based model (EBM) with a denoising score matching (DSM) objective. This DSM-driven approach aimed to maximize the log-likelihood of crystal structures within the training set. While this approach did not yield direct comparisons to actual binding affinity values, it suggests that the learned energy function correlates with experimental binding energies. DSMBind only needs the protein-ligand complex structures for training, without requiring binding affinity labels.

In this work, we first identify and analyze an unrealistic assumption in DSM-trained model; that is, the distribution of training data follows a Boltzmann distribution. Then we introduce DualBind, a simple and novel framework which synergistically combines the supervised mean squared error (MSE) objective with the denoising score matching (DSM) objective to learn the energy function of protein-ligand interactions. Compared to models relying solely on MSE or DSM, DualBind offers several notable advantages. DSM-only models often struggle with the Boltzmann distribution assumption, that does not reliably hold within the training set. In contrast, DualBind aligns closely with actual training affinity data, which significantly enhances the performance of energy predictions. Crucially, this alignment allows DualBind to predict absolute affinity values, rather than merely comparative values. Compared to MSE-only methods that depend heavily on labeled data and may suffer from overfitting, DualBind achieves better generalization by incorporating the DSM objective. In addition, DualBind is not strictly dependent on the availability of labeled data. It can effectively utilize additional unlabeled data by learning the broader energy distribution, which is especially valuable

---

<sup>1</sup>NVIDIA. Correspondence to: Meng Liu <men-liu@nvidia.com>, Saeed Gopal Paliwal <saeep@nvidia.com>.

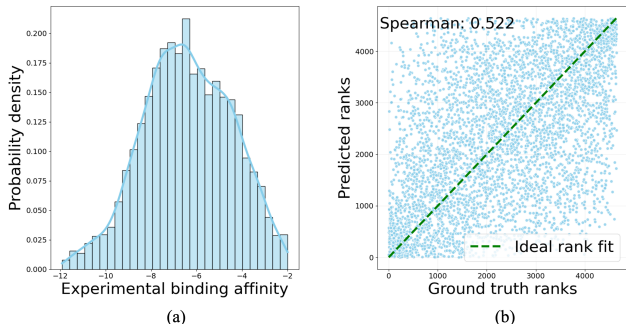


Figure 1. (a) Distribution of binding affinity in the PDBbind v2020 refined dataset. It departs from the Boltzmann distribution. (b) Rank fit of a DSM-only model on training complexes.

in scenarios where experimental data are scarce.

To validate the benefits of DualBind, we conducted extensive experiments on the task of protein-ligand binding affinity prediction. The results consistently show that DualBind outperforms both DSM-only and MSE-only models across a variety of metrics. Furthermore, our empirical evidence demonstrates that DualBind’s capability to simultaneously train on both labeled and unlabeled data significantly enhances model performance.

## 2. Method

In this section, we first point out the unrealistic nature of the underlying Boltzmann distribution assumption typically employed in DSM-only Models in Section 2.1. Then, in Section 2.2, we present our proposed DualBind approach and describe its unique advantages.

### 2.1. Boltzmann Distribution Assumption in DSM-only Models

Energy-based models (EBMs) are a class of models that represent complex systems through an energy function. These models assign a scalar energy value to each configuration of the system (LeCun et al., 2006; Song & Kingma, 2021). A denoising score matching (DSM) objective has been widely used to train EBMs (Vincent, 2011; Song & Ermon, 2019; Ho et al., 2020; Song et al., 2021; Jin et al., 2023a). The efficacy of the DSM objective, which focuses on accurately learning the energy function by maximizing the likelihood of the training data, inherently relies on the assumption that the distribution of training samples aligns with the Boltzmann distribution. Specifically, in the context of protein-ligand binding, this assumption implies that the probability  $p(C)$  of a complex  $C$  being observed is proportional to  $\exp(-E(C))$ , where  $E(C)$  denotes its corresponding binding energy. Note that a lower binding energy indicates a stronger interaction between a ligand and a protein.

However, the actual distribution of protein-ligand complexes in experimental datasets often diverges from this assumption due to experimental biases, selective data reporting, and other confounding factors (Li et al.; Hernández-Garrido & Sánchez-Cruz, 2023). These discrepancies typically result in a dataset that does not exhibit the expected exponential decay of energy states. To validate this observation, we analyze the distribution of binding affinity in the PDBbind v2020 refined dataset, a widely used training dataset where the affinity data are collected from the scientific literature (Su et al., 2018). As illustrated in Figure 1 (a), the binding affinity distribution significantly departs from the expected Boltzmann distribution. Consequently, although the learned EBM can effectively assign local minima to observed protein-ligand complexes, as demonstrated in Jin et al. (2023b), we conjecture it struggles to accurately differentiate between their relative binding affinities.

To verify this, we reproduce an EBM learned by a DSM objective based on Gaussian noise. The learned model achieves a Pearson correlation coefficient of 64.2 on the CASF-2016 test benchmark (Su et al., 2018), which is comparable to the performance of DSMBind (Jin et al., 2023b). To evaluate the model’s capability in learning relative binding affinities from the training dataset, we analyzed the rank fit results *on the training set* using the learned EBM model, as illustrated in Figure 1 (b). The plot shows the predicted versus ground truth ranks of training protein-ligand complexes. Ideally, these ranks should align closely along the dashed green line, which represents perfect correlation. However, there is a significant scatter with a Spearman correlation coefficient of only 0.522. This divergence indicates that DSM-only models struggle to accurately rank the binding affinities of protein-ligand complexes. This finding supports our hypothesis that the DSM’s effectiveness is compromised by its dependence on a theoretical Boltzmann distribution that does not hold in the widely used protein-ligand complex dataset.

### 2.2. DualBind: A Dual-Loss Approach

Here, we propose DualBind, which employs a dual-loss framework to effectively learn the energy function. This integration is not only simple and straightforward but also provides crucial advantages over models that rely solely on DSM or MSE loss. In the following, we first describe the overall dual-loss framework, then introduce the architecture of the energy function model, and finally discuss the significant benefits of DualBind compared to DSM-only and MSE-only models.

**The Dual-Loss Framework.** We represent a protein-ligand complex structure as  $C = (\mathbf{A}, \mathbf{X})$ , where  $\mathbf{A} \in \mathbb{R}^{n \times d}$  denotes the features of the  $n$  atoms in the complex, including both protein and ligand atoms, and  $\mathbf{X} \in \mathbb{R}^{n \times 3}$  represents




Figure 2. An illustration of the DualBind methodology. DualBind employs a dual-loss framework that combines the MSE loss  $\mathcal{L}_{\text{MSE}}$  and the DSM loss  $\mathcal{L}_{\text{DSM}}$ . Specifically,  $\mathcal{L}_{\text{MSE}}$  anchors the predicted binding affinity of the crystal structure to its experimentally determined binding affinity. Concurrently,  $\mathcal{L}_{\text{DSM}}$  shapes the gradient at the perturbed structure. Details are described in Section 2.

the coordinates of these atoms. We define a binding affinity prediction model as  $E_{\theta}(\mathbf{A}, \mathbf{X}) : \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times 3} \rightarrow \mathbb{R}$ , which produces a scalar energy value for any given complex.  $\theta$  denotes the learnable parameters in the model.

Intuitively, our dual-loss framework combines the DSM loss  $\mathcal{L}_{\text{DSM}}$ , which learns the energy landscape by shaping the gradient of the energy function, with the MSE loss  $\mathcal{L}_{\text{MSE}}$ , which directly ties the predictions to known binding affinity values. The overview of our DualBind approach is illustrated in Figure 2.

*MSE Loss.* As shown in the upper branch of Figure 2, the use of the MSE loss is straightforward. We feed a complex crystal structure  $(\mathbf{A}, \mathbf{X})$  into the model and then compute the error between the predicted and the ground truth binding affinity. Formally, for one data sample,

$$\mathcal{L}_{\text{MSE}} = (E_{\theta}(\mathbf{A}, \mathbf{X}) - y)^2, \quad (1)$$

where  $y$  denotes the the ground truth binding affinity. Intuitively, this MSE loss encourages that the points, which represent observed crystal structures within the energy landscape, are anchored to experimental binding affinities, maintaining their alignment with empirical observations.

*DSM Loss.* Training neural networks with denoising is a well-established technique aimed at enhancing the robustness and generalization capabilities of models (Bishop, 1995; Vincent et al., 2008; Kong et al., 2020; Godwin et al., 2021). More specifically, the denoising score matching technique has been employed to pretrain models for 3D molecular tasks, demonstrating that such objective not only simulates learning a molecular force field but also significantly boosts performance across various downstream tasks (Zaidi et al., 2022). Furthermore, within the framework of energy-

based models (EBMs), the DSM objective can be seen as a way of modeling the likelihood of the data (Song & Ermon, 2019; Ho et al., 2020; Song et al., 2021; Jin et al., 2023a). By training an EBM to denoise perturbed data, we implicitly maximize the likelihood of observing the original, unperturbed data under the modeled distribution. Thus, in the context of protein-ligand binding energy modeling, the DSM objective ensures that a crystal structure is guided toward the local minima of the learned energy landscape.

Given such generalization capability and ability to accurately shape local minima for crystal structures, we incorporate the DSM objective into our DualBind model. Following previous DSM-based studies (Zaidi et al., 2022; Jin et al., 2023b), given a crystal structure  $(\mathbf{A}, \mathbf{X})$  from the training dataset, we first perturb it by adding Gaussian noise specifically to the ligand atoms’ coordinates. Formally, for each ligand atom coordinate  $x_i$ , the perturbed coordinate  $\tilde{x}_i$  can be obtained by

$$\tilde{x}_i = x_i + \sigma \epsilon_i, \quad \text{where } \epsilon_i \sim \mathcal{N}(0, \mathbf{I}_3). \quad (2)$$

$\sigma$  is a hyperparameter for tuning the noise scale. Note that Gaussian noise is added only to the coordinates of the ligand atoms in  $\mathbf{X}$ . However, for simplicity in notation, we denote the resulting matrix, which includes both perturbed ligand atom coordinates and unperturbed protein atom coordinates, as  $\tilde{\mathbf{X}}$ .

As illustrated in the lower branch in Figure 2, we feed the perturbed complex structure into the model and obtain its predicted energy  $E_{\theta}(\mathbf{A}, \tilde{\mathbf{X}})$ . Then the DSM loss matches the score of the model distribution  $p_{\theta}$  at the perturbed data  $\tilde{\mathbf{X}}$  with the score of the distribution  $q$  where  $\tilde{\mathbf{X}}$  is sampled. Score is defined as the gradient of the log-probability *w.r.t.*

$\tilde{\mathbf{X}}$ . Formally,

$$\mathcal{L}_{\text{DSM}} = \mathbb{E}_{q(\tilde{\mathbf{X}})} \left[ \left\| \nabla_{\tilde{\mathbf{X}}} \log p_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}) - \nabla_{\tilde{\mathbf{X}}} \log q(\tilde{\mathbf{X}}) \right\|^2 \right]. \quad (3)$$

As shown by Vincent (2011), this is equivalent to

$$\mathcal{L}_{\text{DSM}} = \mathbb{E}_{q(\tilde{\mathbf{X}}|\mathbf{X})p_{\text{data}}(\mathbf{X})} \left[ \left\| \nabla_{\tilde{\mathbf{X}}} \log p_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}) - \nabla_{\tilde{\mathbf{X}}} \log q(\tilde{\mathbf{X}}|\mathbf{X}) \right\|^2 \right]. \quad (4)$$

The distribution given by  $E_{\theta}(\mathbf{A}, \tilde{\mathbf{X}})$  is defined as  $p_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}) = \frac{\exp(-E_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}))}{Z_{\theta}}$ , where  $Z_{\theta}$  is the normalizing constant, which is typically intractable but a constant *w.r.t.*  $\tilde{\mathbf{X}}$  (LeCun et al., 2006; Song & Kingma, 2021). Therefore, the first term in Eq. (4) can be computed by

$$\begin{aligned} \nabla_{\tilde{\mathbf{X}}} \log p_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}) &= -\nabla_{\tilde{\mathbf{X}}} E_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}) - \nabla_{\tilde{\mathbf{X}}} \log Z_{\theta} \\ &= -\nabla_{\tilde{\mathbf{X}}} E_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}). \end{aligned} \quad (5)$$

The second term in Eq. (4) can be easily computed as  $\nabla_{\tilde{\mathbf{X}}} \log q(\tilde{\mathbf{X}}|\mathbf{X}) = -\frac{(\tilde{\mathbf{X}} - \mathbf{X})}{\sigma^2}$  since  $q(\tilde{\mathbf{X}}|\mathbf{X})$  is a Gaussian distribution. Hence,

$$\mathcal{L}_{\text{DSM}} = \mathbb{E}_{q(\tilde{\mathbf{X}}|\mathbf{X})p_{\text{data}}(\mathbf{X})} \left[ \left\| \nabla_{\tilde{\mathbf{X}}} E_{\theta}(\mathbf{A}, \tilde{\mathbf{X}}) - \frac{(\tilde{\mathbf{X}} - \mathbf{X})}{\sigma^2} \right\|^2 \right]. \quad (6)$$

Intuitively, by applying this DSM loss, the model learns to shape its energy landscape such that the energy valleys (*a.k.a.*, local minima) align with the original unperturbed crystal structures.

The total training loss is defined as the weighted summation of the MSE loss and the DSM loss, *i.e.*,

$$\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda \mathcal{L}_{\text{DSM}}, \quad (7)$$

where  $\lambda$  is a hyperparameter that balances the influence of two losses.

**The Architecture of  $E_{\theta}(\mathbf{A}, \mathbf{X})$ .** To facilitate a meaningful comparison with DSMBind, we adopt the architecture for  $E_{\theta}(\mathbf{A}, \mathbf{X})$  as used in DSMBind. It incorporates an SE(3)-invariant model, built on a frame averaging neural network (Puny et al., 2021), to obtain the atom representations. Energy predictions are then made by considering pairwise atom interactions within a predefined distance threshold. For an in-depth understanding of this architecture, we direct readers to Jin et al. (2023b). It is important to note that our dual-loss framework is model-agnostic. This flexibility ensures that our framework can be easily adapted across different modeling approaches for  $E_{\theta}(\mathbf{A}, \mathbf{X})$ .

**Advantages of DualBind.** DualBind offer substantial benefits over models that uses only DSM or MSE.

- Compared to DSM-only models such as DSMBind, DualBind’s integration with actual training affinity data significantly boosts the prediction performance.
- More importantly, this alignment enables the prediction of absolute affinity values, rather than merely comparative values provided by DSM-only models. The ability to quantify absolute binding affinities is essential for drug efficacy study.
- By synthesizing the strengths of both DSM and MSE, DualBind also exhibits better generalization capability compared to MSE-only models, leading to improved prediction performance.
- Another significant advantage of DualBind is its ability to effectively utilize both labeled and unlabeled data simultaneously. Unlike MSE-only methods which are heavily dependent on the presence of labeled data, DualBind’s dual-loss framework significantly reduces this dependency. This flexibility is particularly valuable in scenarios where labeled data are scarce. To be specific, for data samples where affinity labels are available, the model leverages a combination of MSE and DSM losses, as formulated in Eq. (7). For unlabeled crystal structures, we can ignore the upper branch in Figure 2 and only use the DSM loss to train the model. The effectiveness of this approach is validated through experiments detailed in Section 3. We demonstrate that DualBind, by using additional unlabeled data, significantly outperforms MSE-only models, when both are trained with the same quantity of labeled data. This shows DualBind’s unique capability to harness the full potential of both labeled and unlabeled data.

### 3. Experiments

In this section, we evaluate DualBind empirically to demonstrate the previously outlined advantages.

**Dataset.** Following Jin et al. (2023b), we use the refined subset of PDBbind v2020 (Su et al., 2018) as our training set, which has 4643 protein-ligand complexes. The validation set curated by Stärk et al. (2022) has 357 complexes. For testing, we employ the widely used CASF-2016 benchmark (Su et al., 2018), which consists of 285 high-quality protein-ligand complexes known for their reliable binding affinity measurements. We have excluded any training samples that overlap in ligands with those in the validation and test sets, resulting in a slightly smaller training set size for DualBind (4643 complexes) compared to DSMBind (4806 complexes).

**Setup.** The noise scale  $\sigma$  in Eq. (2) is uniformly sampled from the interval  $[0.1, 1]$  during training. We use  $\lambda = 2$  as the weight for DSM loss. We measure the Pearson cor-

Table 1. Performance comparison on the CASF-2016 benchmark.  $\uparrow$  ( $\downarrow$ ) represents that a higher (lower) value denotes better performance. “N/A” in the RMSE column represents that the metric is not applicable, as certain methods yield only comparative values instead of absolute affinity values. The results for scoring functions are sourced from Su et al. (2018), results for  $K_{\text{DEEP}}$  from Kwon et al. (2020), and results for other methods from their respective publications.

| Method             | Affinity labels | $R_p^\uparrow$           | RMSE $^\downarrow$       | $\rho^\uparrow$          |
|--------------------|-----------------|--------------------------|--------------------------|--------------------------|
| Glide-XP           | $\times$        | 0.467                    | 1.95                     | -                        |
| Glide-SP           | $\times$        | 0.513                    | 1.89                     | -                        |
| Autodock Vina      | $\times$        | 0.604                    | 1.73                     | -                        |
| DSMBind (Gaussian) | $\times$        | 0.638                    | N/A                      | -                        |
| DSMBind (SE(3))    | $\times$        | 0.656                    | N/A                      | -                        |
| $K_{\text{DEEP}}$  | $\checkmark$    | 0.738                    | 1.462                    | -                        |
| PIGNet             | $\checkmark$    | 0.749                    | -                        | -                        |
| DualBind           | $\checkmark$    | <b>0.757</b> $\pm 0.006$ | <b>1.461</b> $\pm 0.013$ | <b>0.742</b> $\pm 0.008$ |
| MSE-only           | $\checkmark$    | 0.749 $\pm 0.008$        | 1.491 $\pm 0.017$        | 0.736 $\pm 0.006$        |
| DSM-only           | $\times$        | 0.646 $\pm 0.005$        | N/A                      | 0.652 $\pm 0.007$        |

relation coefficient ( $R_p$ ) to assess the linear relationship between the predicted and actual binding affinity. The root mean square error (RMSE) is used to quantify the average magnitude of the prediction errors, offering a direct measure of prediction error across all test cases. Additionally, the Spearman correlation coefficient ( $\rho$ ) is reported to assess the rank-order accuracy of our predictions, reflecting how well the model can predict the relative order of binding affinities in the dataset. We further assess the impact of each component in DualBind by evaluating two ablated models: one retaining only the MSE loss and the other only the DSM loss, referred to as MSE-only and DSM-only, respectively. For DualBind and its ablated models, we report the average performance across three random runs. The model checkpoints are selected based on achieving the highest Pearson correlation coefficient on the validation set.

**Baselines.** We consider several empirical scoring functions as baselines, including Glide-XP and Glide-SP (Friesner et al., 2004), along with AutoDock Vina (Trott & Olson, 2010; Eberhardt et al., 2021). Additionally, we compare against two versions of DSMBind that use Gaussian noise and SE(3) noise respectively (Jin et al., 2023b). In the realm of supervised machine learning methods, we include  $K_{\text{DEEP}}$  (Jiménez et al., 2018) and PIGNet (Moon et al., 2022), which use training sets similar in size to ours. It is important to note that direct comparisons across all models are challenging due to differences in dataset sizes and splits used for training. We emphasize that the primary aim of our study is not to outperform all supervised methods in the literature but to demonstrate the effectiveness of our dual-loss approach, which is designed to be model-agnostic.

**Results.** As summarized in Table 1, DualBind achieves a Pearson correlation of 0.757, significantly outperforming scoring functions and DSM-based methods. Unlike DSM-only models, which do not provide absolute affinity predic-

tions (as indicated by “N/A” for RMSE), DualBind achieves an RMSE of 1.461, affirming its capability to deliver accurate absolute affinity values. Furthermore, as demonstrated in Figure 3, the rank fit of our DualBind model on training complexes significantly surpasses that of DSM-only models (Figure 1 (b)), achieving a Spearman correlation of 0.820 compared to 0.522. This highlights DualBind’s enhanced capability in accurately ranking protein-ligand interactions, overcoming the limitations posed by the unrealistic Boltzmann distribution assumption in DSM-only models. Hence, on the test set, DualBind continues to exhibit superior performance with a Spearman correlation of 0.742, compared to 0.652 for DSM-only models. Additionally, when compared to existing supervised models and our MSE-only variant, DualBind exhibits enhanced linear correlation, reduced RMSE, and improved ranking power. These results validate the key advantages outlined in Section 2.

### Experiment on DualBind’s Flexible Data Use Strategy.

As outlined in Section 2, because of the dual-loss framework, DualBind can flexibly employ the DSM loss for unlabeled data while using a combination of DSM and MSE losses for labeled data. This allows for the simultaneous training on both labeled

and unlabeled datasets. Thus, in this experiment, we employ labels from only 50% of the training samples (*i.e.*, 2321 complexes) to train both the MSE-only model and DualBind. For DualBind, we additionally leverage the remaining 50% samples without using their affinity labels.




Figure 3. Rank fit of DualBind on training complexes.

Table 2. Experimental results on DualBind’s flexible data use strategy. In this comparison, DualBind and the MSE-only model use the same amount of labeled data. DualBind enhances the performance by additionally using unlabeled data, leveraging its unique data use capability.  $\uparrow$  ( $\downarrow$ ) represents that a higher (lower) value denotes better performance.

| Method   | #Labeled | #Unlabeled | $R_p^\uparrow$           | RMSE $^\downarrow$       | $\rho^\uparrow$          |
|----------|----------|------------|--------------------------|--------------------------|--------------------------|
| MSE-only | 2321     | $\times$   | 0.664 $\pm$ 0.037        | 1.694 $\pm$ 0.086        | 0.666 $\pm$ 0.028        |
| DualBind | 2321     | 2321       | <b>0.731</b> $\pm$ 0.007 | <b>1.684</b> $\pm$ 0.087 | <b>0.732</b> $\pm$ 0.006 |
| MSE-only | 4643     | $\times$   | 0.749 $\pm$ 0.008        | 1.491 $\pm$ 0.017        | 0.736 $\pm$ 0.006        |

Specifically, for DualBind, we randomly sampled data from both the labeled and unlabeled subsets to compute their respective loss values, followed by back-propagation to train the model. The goal of this experimental setup is to effectively demonstrate DualBind’s unique capacity to utilize both labeled data and unlabeled data, a functionality not available in MSE-only models.

The average results over three random runs are reported in Table 2. Notably, DualBind significantly outperforms the MSE-only model when using an equivalent amount of labeled data, achieving a Pearson correlation of 0.731 compared to 0.664. This improvement shows DualBind’s effective use of additional unlabeled data. For comparison, we also include results from an MSE-only model trained with 100% labeled data. Impressively, DualBind, utilizing only half labeled data and half unlabeled data, achieves results nearly comparable to the MSE-only model trained on the full labeled dataset, in terms of both linear and ranking correlation coefficients. It clearly demonstrates DualBind’s superior ability to enhance model performance through its unique data use strategy. This strategic use of data resources also broadens the practical applicability of DualBind, especially in scenarios where acquiring extensive labeled data is challenging.

## 4. Conclusion

In this study, we propose DualBind, a novel dual-loss framework which can perform accurate absolute binding affinity prediction. Our evaluations on standard benchmark demonstrate that DualBind consistently outperforms both baseline models and its ablated variants. A key feature of DualBind is its innovative strategy of integrating unlabeled data with labeled examples, which has shown considerable benefits in our initial experiments. Looking ahead, we plan to expand our research into hybrid training approaches and explore pretraining techniques around the DualBind framework.

## Acknowledgements

We thank the entire NVIDIA BioNeMo team for the discussions and support throughout this research.

## References

- Bishop, C. M. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116, 1995.
- Chen, L., Tan, X., Wang, D., Zhong, F., Liu, X., Yang, T., Luo, X., Chen, K., Jiang, H., and Zheng, M. Transformer-CPI: improving compound–protein interaction prediction by sequence-based deep learning with self-attention mechanism and label reversal experiments. *Bioinformatics*, 36(16):4406–4414, 2020.
- Eberhardt, J., Santos-Martins, D., Tillack, A. F., and Forli, S. AutoDock Vina 1.2.0: New docking methods, expanded force field, and python bindings. *Journal of chemical information and modeling*, 61(8):3891–3898, 2021.
- Friesner, R. A., Banks, J. L., Murphy, R. B., Halgren, T. A., Klicic, J. J., Mainz, D. T., Repasky, M. P., Knoll, E. H., Shelley, M., Perry, J. K., et al. Glide: a new approach for rapid, accurate docking and scoring. 1. method and assessment of docking accuracy. *Journal of medicinal chemistry*, 47(7):1739–1749, 2004.
- Godwin, J., Schaarschmidt, M., Gaunt, A., Sanchez-Gonzalez, A., Rubanova, Y., Veličković, P., Kirkpatrick, J., and Battaglia, P. Simple gnn regularisation for 3d molecular property prediction & beyond. *arXiv preprint arXiv:2106.07971*, 2021.
- Hernández-Garrido, C. A. and Sánchez-Cruz, N. Experimental uncertainty in training data for protein-ligand binding affinity prediction models. *Artificial Intelligence in the Life Sciences*, 4:100087, 2023.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- 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):18209–18232, 2021.

- Jiménez, J., Skalic, M., Martinez-Rosell, G., and De Fabritiis, G. K deep: protein–ligand absolute binding affinity prediction via 3d-convolutional neural networks. *Journal of chemical information and modeling*, 58(2):287–296, 2018.
- Jin, W., Chen, X., Veticaden, A., Sarzikova, S., Raychowdhury, R., Uhler, C., and Hacoheh, N. DSMBind: Se (3) denoising score matching for unsupervised binding energy prediction and nanobody design. *bioRxiv*, pp. 2023–12, 2023a.
- Jin, W., Sarkizova, S., Chen, X., Hacoheh, N., and Uhler, C. Unsupervised protein–ligand binding energy prediction via neural euler’s rotation equation. *Advances in Neural Information Processing Systems*, 36, 2023b.
- Kong, K., Li, G., Ding, M., Wu, Z., Zhu, C., Ghanem, B., Taylor, G., and Goldstein, T. Flag: Adversarial data augmentation for graph neural networks. 2020.
- Kwon, Y., Shin, W.-H., Ko, J., and Lee, J. Ak-score: accurate protein–ligand binding affinity prediction using an ensemble of 3d-convolutional neural networks. *International journal of molecular sciences*, 21(22):8424, 2020.
- LeCun, Y., Chopra, S., and Hadsell, R. A tutorial on energy-based learning. 2006.
- Li, J., Guan, X., Zhang, O., Sun, K., Wang, Y., Bagni, D., and Head-Gordon, T. Leak proof pdbname: A reorganized dataset of protein–ligand complexes for more generalizable binding affinity prediction. *ArXiv*.
- Li, S., Wan, F., Shu, H., Jiang, T., Zhao, D., and Zeng, J. MONN: a multi-objective neural network for predicting compound–protein interactions and affinities. *Cell Systems*, 10(4):308–322, 2020.
- Lin, H., Wang, S., Zhu, J., Li, Y., Pei, J., and Lai, L. DeepRLI: A multi-objective framework for universal protein–ligand interaction prediction. *arXiv preprint arXiv:2401.10806*, 2024.
- Lu, W., Wu, Q., Zhang, J., Rao, J., Li, C., and Zheng, S. TankBind: Trigonometry-aware neural networks for drug–protein binding structure prediction. *Advances in neural information processing systems*, 35:7236–7249, 2022.
- 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):3661–3673, 2022.
- Nguyen, T., Le, H., Quinn, T. P., Nguyen, T., Le, T. D., and Venkatesh, S. GraphDTA: predicting drug–target binding affinity with graph neural networks. *Bioinformatics*, 37(8):1140–1147, 2021.
- Öztürk, H., Özgür, A., and Ozkirimli, E. DeepDTA: deep drug–target binding affinity prediction. *Bioinformatics*, 34(17):i821–i829, 2018.
- Puny, O., Atzmon, M., Smith, E. J., Misra, I., Grover, A., Ben-Hamu, H., and Lipman, Y. Frame averaging for invariant and equivariant network design. In *International Conference on Learning Representations*, 2021.
- Song, Y. and Ermon, S. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Song, Y. and Kingma, D. P. How to train your energy-based models. *arXiv preprint arXiv:2101.03288*, 2021.
- Song, Y., Durkan, C., Murray, I., and Ermon, S. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34: 1415–1428, 2021.
- Stärk, H., Ganea, O., Pattanaik, L., Barzilay, R., and Jaakkola, T. Equibind: Geometric deep learning for drug binding structure prediction. In *International conference on machine learning*, pp. 20503–20521. PMLR, 2022.
- 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):895–913, 2018.
- Trott, O. and Olson, A. J. AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of computational chemistry*, 31(2):455–461, 2010.
- Vincent, P. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103, 2008.
- Zaidi, S., Schaarschmidt, M., Martens, J., Kim, H., Teh, Y. W., Sanchez-Gonzalez, A., Battaglia, P., Pascanu, R., and Godwin, J. Pre-training via denoising for molecular property prediction. In *The Eleventh International Conference on Learning Representations*, 2022.
- Zhang, X., Gao, H., Wang, H., Chen, Z., Zhang, Z., Chen, X., Li, Y., Qi, Y., and Wang, R. Planet: a multi-objective graph neural network model for protein–ligand binding affinity prediction. *Journal of Chemical Information and Modeling*, 2023a.

Zhang, X., Wang, L., Helwig, J., Luo, Y., Fu, C., Xie, Y., Liu, M., Lin, Y., Xu, Z., Yan, K., et al. Artificial intelligence for science in quantum, atomistic, and continuum systems. *arXiv preprint arXiv:2307.08423*, 2023b.