Beyond Greedy Acquisition: A Distribution-Aware Active Learning Framework for Molecular Discovery

Ruixuan Chen [*1] **Zhixuan Song** [*1] **Wentao Li** [1] **Leyi Zhao** [1] **Kai Zhao** [1] **Xiaonan Wang** [1]

[*]Equal contribution

[1]*Beijing Key Laboratory of Artificial Intelligence for Advanced Chemical Engineering Materials, Department of Chemical Engineering, Tsinghua University, Beijing, China*

Correspondence to: Xiaonan Wang wangxiaonan@tsinghua.edu.cn.

1. Introduction

Molecular property prediction accelerates chemical and materials discovery by enabling rapid, low-cost screening of candidate molecules prior to expensive synthesis, characterization, or quantum-chemical computations. While density functional theory (DFT) can provide high-fidelity property labels, its high computational cost makes it difficult to scale to large chemical spaces [1,2]. Active learning (AL) with graph neural networks (GNNs) has therefore become a practical approach to improve sample efficiency under limited labeling budgets in molecular modeling [3,4].

However, recent evaluations indicate that the advantages of common molecular AL pipelines can diminish on out-of-distribution (OOD) samples [5]. In such settings, predictions on OOD samples may become inaccurate due to model's sensitivity to distribution shift between train and test sets, and AL can perform comparably to random sampling [6,7].

Most existing molecular AL methods are model-centric: at each round, they train a predictor on the current labeled set and then they query new molecules by greedily maximizing a one-step acquisition score under the current model state [8]. These approaches do not specify the ordering or distribution of intermediate training subsets, even though learning order can influence deep models and the training trajectory shapes representation learning and calibration [9,10]. As a result, acquisition scores can be biased by trajectory-induced artifacts, leading to redundant queries and limited coverage of chemical space.

We address this gap with DALB (Distribution-Aware Learning Bridge), a data-centric active learning framework that controls the training trajectory at the distribution level before applying standard model-centric acquisition. DALB constructs a distribution-aware curriculum by selecting trajectories that traverse diverse latent-space modes. Importantly, DALB is acquisition-agnostic and can be paired with a wide range of acquisition strategies, enabling improved predictive performance by selecting the best suited strategy.

2. Method

2.1 Molecular representation and prediction model

We consider a SMILES-based molecular dataset [14] where each molecule is associated with target properties (xTB_{S1} and xTB_{T1}). Each SMILES string is parsed into a molecular graph with atom and bond features. We use a message passing neural network (MPNN) [11] to perform atom-level message updates followed by mean pooling to obtain a graph-level embedding. In parallel, a pretrained variational autoencoder (VAE) [12] encoder produces a compact latent descriptor for each molecule. We concatenate the MPNN graph embedding with the VAE latent descriptor and feed the combined representation into a regression feed-forward network for feature prediction.

2.2 Data-level active learning: distribution-aware path filtering

2.2.1 Paths as training curricula

We view the training process as following a curriculum (path) that determines the order in which samples are presented to the model. At training step t, the model observes either: (i) a cumulative prefix subset S_t^π that contains all samples seen up to step t; or (ii) the current mini-batch B_t^π. Our goal is to distinguish paths that remain overly similar to the global distribution throughout training (Aligned) from paths that introduce controlled distributional variation (Exploratory).

2.2.2 Molecule-specific distribution difference calculation

To quantify distribution differences between a training subset and the global set, Optimal Transport (OT) [13] is a natural choice since it compares distributions via global geometry, but repeated OT on high-dimensional molecular representations is costly. We therefore project molecules into a compact latent space using a pretrained VAE encoder, which reduces dimensionality and stabilizes OT/Sinkhorn computations, enabling efficient subset-vs-global deviation evaluation. More generally, this framework is not limited to VAEs. Any pretrained encoder providing a consistent latent space can be utilized.

We embed each sample into the latent space to obtain latent vectors $\{z_i\}_{i=1}^N$. We denote the empirical distribution of any subset (prefix or batch) as μ_S and that of the full training set as μ_G.

2.2.3 Distributional deviation via debiased Sinkhorn distance

To measure how much a subset distribution deviates from the global distribution, we use the debiased

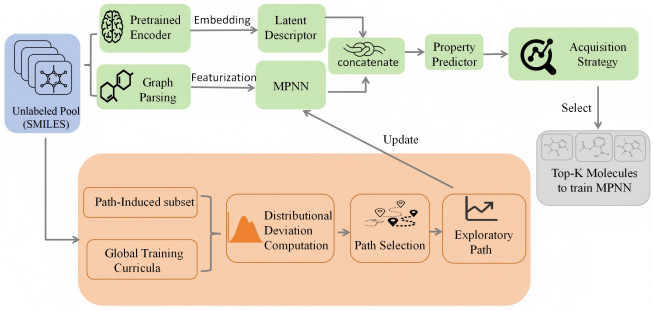

Fig. 1: overall pipeline for DALB framework

Sinkhorn divergence [15,16]:

$$W_{\text{debiased}}(\mu_S, \mu_G) = 2W_\varepsilon(\mu_S, \mu_G) - W_\varepsilon(\mu_S, \mu_S) - W_\varepsilon(\mu_G, \mu_G). \quad (1)$$

We define single-step deviation in cumulative or batchwise mode:

Cumulative mode:

$$D_t^\pi = W_{\text{debiased}}(\mu_{S_t^\pi}, \mu_G). \quad (2)$$

Batchwise mode:

$$D_t^\pi = W_{\text{debiased}}(\mu_{B_t^\pi}, \mu_G). \quad (3)$$

2.2.4 Path deviation score and binary classification

We randomly sample a set of candidate paths and compute D_t^π for each step t. Across sampled paths, we estimate the mean and standard deviation of deviation at each step:

$$\mu_{D,t} = \mathbb{E}_\pi[D_t^\pi], \quad \sigma_{D,t} = \text{Std}_\pi[D_t^\pi]. \quad (4)$$

We score each path by O_π and label Exploratory if $O_\pi > \tau$, otherwise Aligned:

$$O_\pi = \sum_{t=1}^{T} \mathbf{1}\{D_t^\pi \notin [\mu_{D,t} - 2\sigma_{D,t}, \mu_{D,t} + 2\sigma_{D,t}]\}. \quad (5)$$

Aligned paths:

$$O_\pi \le \tau. \quad (6)$$

Exploratory paths:

$$O_\pi > \tau. \quad (7)$$

Exploratory paths are prioritized for curricula.

2.3 Model-level active learning: strategy-agnostic acquisition

Given an unlabeled pool, DALB queries samples using a plug-in acquisition strategy selected from multiple options. In each round, we train the predictor on the labeled set with exploratory curricula prioritized by the data-level module, then score all unlabeled candidates with the chosen strategy and label the top-k.

DALB can pair the strategy that best matches the dataset characteristics and target task, improving sample efficiency and boosting predictive performance under a fixed labeling budget.

3. Motivation and Rationale for DALB

Most AL algorithms are greedy, optimizing a one-step objective under the current model state [17]. Because these scores approximate local, immediate gain, they can become overly sensitive to where the model has recently trained. Early selection errors can be self-reinforcing, producing cold-start behavior where the learner remains stuck in an unrepresentative region [18].

DALB addresses this failure mode by controlling the data exposure trajectory during training. Aligned paths are typically optimization-stable but exploration-poor: the model repeatedly observes subsets that track the global average, limiting exposure to rare clusters, boundary regions, and long-tail modes. In contrast, Exploratory paths inject controlled distributional variation—either at the batch level or at the path level. This increases the diversity of latent modes encountered during optimization, yielding representations that are less tied to the in-distribution "average case." As a result, subsequent acquisition decisions become less trajectory-biased and more reflective of global information, improving robustness and OOD generalization under shift.

4. Expected Outcomes

We expect DALB to improve sample efficiency and robustness under distribution shift. Concretely, (i) our results already show that trajectory control alone outperforms random ordering, confirming the benefit of distribution-aware curricula: models trained with exploratory paths achieve 10.7% and 26.8% lower MAE under the batchwise and cumulative modes, respectively. (ii) We expect that, when paired with common acquisition functions, DALB will yield lower error at the same labeling budget and stronger OOD generalization than greedy baselines. (iii) We further anticipate that DALB will alleviate common limitations of standard active learning, including cold-start behavior and reduced generalization under distribution shift. We will benchmark DALB across multiple datasets to verify the stability of these gains.

References

1. Huang, B., et al. The central role of density functional theory in the AI age. Science, 378(6621):eabn3445, 2023.

2. Jain, A., Shin, Y., and Persson, K. A. Computational predictions of energy materials using density functional theory. Nature Reviews Materials, 1:15004, 2016.

3. Li, B., et al. A diversity maximizing active learning strategy for graph neural network models of chemical properties. Molecular Systems Design & Engineering, 7(4):567–580, 2022.

4. Jiao, Z., et al. Fine-tuning graph neural networks via active learning: Unlocking the potential of graph neural networks trained on nonaqueous systems for aqueous CO_2 reduction. Journal of Chemical Theory and Computation, 21(6):3176–3186, 2025.

5. Shields, M. J., Rowley, C. J., Malone, A., and Meng, Y. S. Evaluating uncertainty-based active learning for accelerating the prediction of aqueous solubility and acute toxicity. Journal of Chemical Information and Modeling, 63(19):6162–6176, 2023.

6. Zhang, Y. and Lee, A. A. Bayesian semi-supervised learning for uncertainty-calibrated prediction of molecular properties and active learning. Chemical Science, 10(44):8154–8163, 2019.

7. Francoeur, P., Penaherrera, D. S., and Koes, D. R. Active learning for small molecule pKa regression: A long way to go. ChemRxiv (preprint), 2022. https://doi.org/10.26434/chemrxiv.2022-8w1q0

8. Subedi, R., Wei, L., Gao, W., Chakraborty, S., and Liu, Y. Empowering active learning for 3D molecular graphs with geometric graph isomorphism. Advances in Neural Information Processing Systems (NeurIPS), 37, 2024.

9. Tousi, E., et al. Learning trajectories of deep neural networks during self-supervised learning. In Proceedings of the 2024 Cognitive Computational Neuroscience Conference, 2024.

10. Hacohen, G. and Weinshall, D. On the power of curriculum learning in training deep networks. In Proceedings of the 36th International Conference on Machine Learning (ICML), volume 97, pages 11059–11069. PMLR, 2019.

11. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In Proceedings of the 34th International Conference on Machine Learning (ICML), volume 70, pages 1263–1272. PMLR, 2017. arXiv:1704.01212

12. Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In Proceedings of the International Conference on Learning Representations (ICLR), 2013.

13. Peyré, G. and Cuturi, M. Computational optimal transport: With applications to data science. Foundations and Trends in Machine Learning, 11(5–6):355–607, 2019.

14. Chen, Y., Verma, S., Greenman, K. P., Yin, H., Wang, Z., Wang, L., Li, J., Gómez-Bombarelli, R., Walsh, A., and Wang, X. A unified active learning framework for photosensitizer design. Chemical Science, 17(2):916–926, 2026. https://doi.org/10.1039/D5SC05749C

15. Feydy, J., Séjourné, T., Vialard, F.-X., Amari, S.-I., Trouvé, A., and Peyré, G. Fast and differentiable Wasserstein barycenter computation. In Proceedings of the 38th International Conference on Machine Learning (ICML), pages 3584–3594. PMLR, 2021.

16. Flamary, R., Courty, N., Gramfort, A., and Trouvé, A. POT: Python optimal transport library. Journal of Machine Learning Research, 2017.

17. Krause, A. and Guestrin, C. Nonmyopic active learning of Gaussian processes: An exploration-exploitation approach. In Proceedings of the 24th International Conference on Machine Learning (ICML), pages 449–456. ACM, 2007.

18. Yehuda, O., Dekel, A., Hacohen, E., and Weinshall, D. Active learning through a covering lens. arXiv:2205.11320, 2022.

19. Musielewicz, J., Wang, X., Tian, T., and Ulissi, Z. FINETUNA: Fine-tuning accelerated molecular simulations. arXiv:2205.01223, 2022.