

---

# Why Pool When You Can Flow? Active Learning with GFlowNets

---

Renfei Zhang<sup>◇\*</sup>  
rza104@sfu.ca

Mohit Pandey<sup>♠♣\*</sup>  
mkpandey@student.ubc.ca

Artem Cherkasov<sup>♠♡</sup>  
artc@interchange.ubc.ca

Martin Ester<sup>◇</sup>  
ester@sfu.ca

<sup>◇</sup> School of Computer Science, Simon Fraser University, Burnaby, BC, Canada

<sup>♠</sup> Vancouver Prostate Centre, University of British Columbia, Vancouver, BC, Canada

<sup>♡</sup> Faculty of Medicine, University of British Columbia, Vancouver, BC, Canada

<sup>♣</sup> Diagen AI

## Abstract

The scalability of pool-based active learning is limited by the computational cost of evaluating large unlabeled datasets, a challenge that is particularly acute in virtual screening for drug discovery. While Active learning strategies such as Bayesian Active Learning by Disagreement (BALD) prioritize informative samples, it remains computationally intensive when scaled to libraries containing billions samples. In this work, we introduce BALD-GFlowNet, a generative active learning framework that circumvents this issue. Our method leverages Generative Flow Networks (GFlowNets) to directly sample objects in proportion to the BALD reward. By replacing traditional pool-based acquisition with generative sampling, BALD-GFlowNet achieves scalability that is independent of the size of the unlabeled pool. In our virtual screening experiment, we show that BALD-GFlowNet achieves a performance comparable to that of standard BALD baseline while generating more structurally diverse molecules, offering a promising direction for efficient and scalable molecular discovery.

## 1 Introduction

Active learning strategies are crucial for reducing the oracle labeling cost in large-scale machine learning tasks. However, traditional pool-based active learning faces significant challenges with scalability. Their reliance on scoring every instance in the unlabeled pool introduces a computational overhead that scales linearly with pool size, making them unsuitable for billion-scale datasets.

These limitations are particularly acute in high-throughput virtual screening (HPVS) for drug discovery. Identifying small molecules with high binding affinity to a target protein is a critical step, but exhaustively evaluating massive molecular libraries via docking oracles is computationally prohibitive. For example, performing virtual screening on 1.3 billion ligands can take 28 days to complete, even with 8,000 GPUs [1, 2].

To address the scalability limitation, we propose moving from a selective approach to a generative one. Instead of asking "Which samples from the current pool are most informative?", we ask "What does an informative sample look like?". To this end, we introduce **BALD-GFlowNet**, a generative active learning approach that integrates Generative Flow Networks (GFlowNets) with mutual information (MI)-based rewards [3, 4, 5, 6]. Unlike traditional pool-based methods, BALD-GFlowNet directly generates informative samples guided by a reward function. This approach mitigates the scalability issue by decoupling the acquisition cost from the size of the unlabeled pool. Our main contributions are:

---

\*Equal contribution.

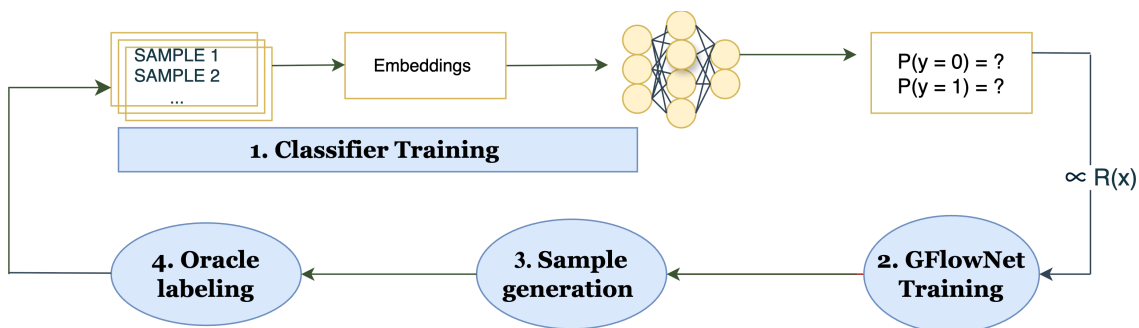


Figure 1: The generative active learning pipeline for BALD-GFlowNet begins by training a binary classifier on an initial training dataset. This classifier is subsequently used in GFlowNet for computing the mutual information to guide the generation of a batch of samples where the classifier’s predictions are most uncertain. These newly generated samples are labeled by an oracle and are subsequently added to the training dataset to retrain the classifier and improve its performance in the next iteration.

- We introduce BALD-GFlowNet, a novel active learning framework that replaces pool-based acquisition with a generative policy trained to sample objects proportional to their BALD reward, making the acquisition cost independent of the unlabeled pool size.
- We show on a synthetic grid task that BALD-GFlowNet identifies high-uncertainty regions more efficiently than exhaustive search, requiring fewer oracle calls to improve model performance. The learned GFlowNet policy demonstrates convergence towards high-uncertainty regions as training progresses, leading to better scalability without sacrificing the quality of acquired samples.
- We apply BALD-GFlowNet to a large-scale virtual screening task for the Janus Kinase 2 (JAK2) protein [7], showing it achieves performance comparable to the BALD baseline with superior efficiency. At a library size of 12 million, BALD-GFlowNet achieves a 12.5% runtime reduction over the pool-based BALD baseline [8]. Here our method not only prioritizes informative samples but also ensures their chemical viability and structural diversity.

## 2 Background and Related Work

### 2.1 Active Learning

Uncertainty-based approaches, such as Bayesian Active Learning by Disagreement (BALD) [3], select the most informative samples with the highest predictive disagreement among the model ensemble. Data points with high mutual information are those where the average prediction is uncertain, but individual models sampled from the posterior are confident. While effective at identifying individual informative samples, such methods do not scale well to large datasets, as they require exhaustively computing the MI score for every sample in the unlabeled pool.

### 2.2 GFlowNets in Drug Discovery

Recent studies have demonstrated the effectiveness of GFlowNets in molecular design and drug discovery, due to their ability to generate diverse, reward-aligned samples. For instance, [9] and [10] have applied GFlowNets to de-novo design of biological sequences and discovery of novel protein binders. [6] demonstrated that atom-level GFlowNets, with rewards combining drug likeliness metrics can generate chemically feasible molecules [11, 12, 13].

## 3 BALD-GFlowNet Framework

**BALD-GFlowNet** is an active learning framework that integrates the BALD objective with GFlowNets to learn a policy for generating informative samples. Instead of exhaustively computing an acquisition score for all objects in the unlabeled pool, our method learns a generative

policy to sample them directly. As shown in Algorithm 1 and Figure 1, the process begins by fitting a surrogate model to a small, initial training dataset. This model then computes a reward for each sample using the BALD acquisition score [3]. We use this principled heuristic because it seeks samples on which models from the posterior distribution are most likely to disagree. A GFlowNet is subsequently trained to sample objects with probability proportional to this BALD reward. The generated candidates are then labeled by an oracle and added to the training set. The surrogate model is updated with this new data, and the cycle repeats. This iterative framework maintains a constant acquisition cost per cycle, making BALD-GFlowNet scalable to billion-sized pools.

---

**Algorithm 1** BALD-GFlowNet Active Learning Framework

---

**Require:** Initial training set size  $N_0$ ; Unlabeled pool of objects  $\mathcal{D}_{\text{pool}}$ ; Surrogate Model  $f$

**Require:** GFlowNet training episodes  $T$ ; AL iterations  $N$ ; Acquisition batch size  $k$

```

1:  $\mathcal{D}_{\text{train}} \leftarrow \text{SampleInitialData}(\mathcal{D}_{\text{pool}}, N_0)$ 
2:  $\mathcal{D}_{\text{pool}} \leftarrow \mathcal{D}_{\text{pool}} \setminus \mathcal{D}_{\text{train}}$ 
3: Initialize a GFlowNet Policy  $\pi$ 
4: Train an initial surrogate model  $f$  with  $\mathcal{D}_{\text{train}}$ 

5: for  $t = 1, \dots, N$  do
6:   for episode = 1, ...,  $T$  do
7:      $\pi \leftarrow \text{TrainGFNPoly}(\pi, \text{BALD})$ 
8:   end for

9:    $b_t \leftarrow \text{SampleFromGFNPoly}(\pi, k)$ 
10:   $y_t \leftarrow \text{QueryOracle}(b) \forall b \in b_t$ 
11:   $\mathcal{D}_{\text{train}} \leftarrow \mathcal{D}_{\text{train}} \cup \{(b_t, y_t)\}$ 
12:   $\mathcal{D}_{\text{pool}} \leftarrow \mathcal{D}_{\text{pool}} \setminus \mathcal{D}_{\text{train}}$ 
13:  Update the surrogate model  $f$  on  $\mathcal{D}_{\text{train}}$ 
14: end for
15: return Trained surrogate model  $f$ 
```

---

## 4 Empirical Results

We first evaluate our proposed framework, BALD-GFlowNet, on a synthetic two-dimensional (2D) grid that is small enough to allow for exact analysis, and compare its sampling efficiency against traditional active learning baselines. We find that BALD-GFlowNet (i) effectively navigates the grid to sample points with high uncertainty, (ii) identifies informative data points more efficiently than an exhaustive BALD search. We apply BALD-GFlowNet to a large-scale virtual screening domain to improve the performance of a classifier that predicts the binding affinity of small molecules to a target protein. Here, the GFlowNet agent directly generates molecules for acquisition, guided by a composite reward function that combines the BALD score with key drug-likeness properties. We find that BALD-GFlowNet samples high-reward molecules faster than baselines, when scaling to large chemical spaces with billions of compounds.

### 4.1 2D Grid Experiment


To rigorously evaluate our proposed method in a controlled setting, we first verify its performance on a synthetic task with 2D grid. This grid environment allows for a precise analysis of the GFlowNet-based policy’s sampling and computational scalability and comparison against established active learning baselines.

The BALD-GFlowNet framework here consists of a surrogate model and an active learning agent learning the GFlowNet policy. The GFlowNet agent starts at the origin (0,0). It then navigates the grid using an action space comprising four directional moves (up, down, left, right) and a terminal "stop" action. The policy is trained to sample terminal states  $x$  with probability proportional to a reward  $R(x)$  defined by the BALD acquisition score at that state.

To analyze the GFlowNet’s learning behavior and sampling efficiency, we focus on a single active learning step where the reward landscape is fixed. To understand how the GFlowNet’s exploration

evolves, we analyzed the distribution of terminal states from trajectories sampled at different training episodes. In early training (Figure 2, left and middle panels), as the policy learns basic navigation, terminal states are spread across the state space with no clustering patterns. In contrast, later stage (Figure 2, right) shows trajectories converging toward regions with high MI reward, indicating more focused and reward-driven exploration.

Figure 5 shows the efficiency gains of our method. While the BALD baseline requires  $O(n^2)$  oracle calls with respect to grid size  $n$ , our method uses substantially fewer calls. Although the number of oracle calls for our method increases with training, it quickly plateaus at a much lower level than the baseline, highlighting its computational efficiency without exhaustive evaluation.



**Figure 2: Terminal State Distribution Over Time** The background shows the mutual information (MI) reward landscape (colorbar, top right), with warmer colors indicating higher rewards. Density heatmaps (colorbar, bottom right) represent the distribution of episode terminal states over the most recent 300 trajectories at three training checkpoints: 300, 1000, and 2000 episodes. White contour lines delineate regions of equal terminal density. The green star indicates the initial state position. Early training (left) shows broad exploration, whereas later training stages show convergence toward high-reward regions of the state space.

## 4.2 Virtual Screening (VS) Case Study

The objective of our virtual screening case study is to generate a diverse and informative set of molecules to improve docking score classification for Janus Kinase 2 (JAK2) [7]. To achieve this, we finetune a pretrained GFlowNet [6], which uses a graph-based Transformer policy to construct molecules atom-by-atom, from a chemical action space of atom and bond additions. We then guide the generation process by integrating the BALD objective into the reward function.

To replicate the low hit rate characteristic of VS campaigns, we create a highly imbalanced classification task. Then we finetune the GFlowNet using the following composite reward function designed to balance informativeness with drug-likeness properties:

$$\text{Reward}(x) = MI(x) \cdot TPSA(x) \cdot QED(x) \cdot SAS(x) \cdot Rings(x). \quad (1)$$

Here  $MI(x)$  is the mutual information ( $I(y; \omega \mid x, D)$ ) responsible for prioritizing molecules that are informative for the surrogate classifier, while the remaining terms ensure the generated molecules possess desirable chemical properties.

We compare BALD-GFlowNet to BALD baseline on the JAK2 virtual screening task. Our results show that BALD-GFlowNet provides a scalable solution for active learning in massive molecular libraries without sacrificing performance. The convergence point in Figure 3a highlights the efficiency gain: BALD-GFlowNet reaches over 90% of BALD’s maximum F1 score using the same number of oracle calls that BALD needs for a single iteration (also see Table 1). Although BALD baseline achieves marginally higher peak performance, this gain comes at a cost that scales linearly with library size  $O(n)$ . In contrast, BALD-GFlowNet’s runtime remains constant, and an extrapolation to a molecule pool of size 50 million leads to an estimated 2.5 times wall-clock speedup (Figure 3b).

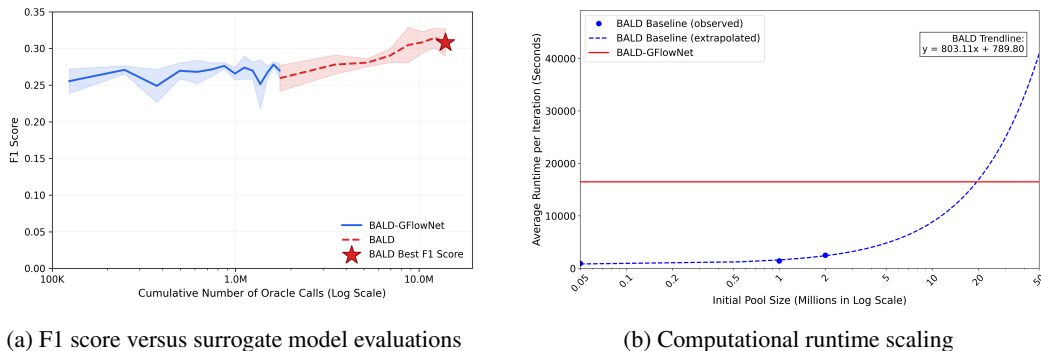


Figure 3: **Efficiency analysis of BALD-GFlowNet.** (a) F1 score versus the cumulative number of calls to the surrogate model (log scale). Random baseline does not query the surrogate model. BALD-GFlowNet (blue) achieves peak F1 score of 0.30 after around 1.74M evaluations (124K per iteration  $\times$  14 iterations). In contrast, BALD (red dashed) reaches its peak F1=0.32 at iteration 7, requiring 12.2M evaluations (1.74M per iteration  $\times$  7). The red star marks BALD’s peak performance. Gaussian smoothing ( $\sigma = 1.5$ ) is applied. (b) Runtime scaling analysis across different pool sizes (log scale). Standard BALD shows  $\mathcal{O}(n)$  complexity, with blue circles showing empirically measured runtimes and the dashed line representing extrapolated values ( $y = 803.11x + 789.80$ ). BALD-GFlowNet maintains constant  $\mathcal{O}(1)$  complexity (red), showing orders of magnitude improvement for large-scale molecular discovery tasks.

## 5 Discussion & Limitations

Our empirical results on both the synthetic grid and large-scale VS tasks show that BALD-GFlowNet is a scalable alternative to traditional pool-based active learning. The framework achieves classification performance comparable to the strong BALD baseline while providing distinct advantages in sampling efficiency. By learning a generative policy, BALD-GFlowNet’s acquisition cost remains constant regardless of the unlabeled pool’s size. The framework’s scalability was confirmed by the number of oracle calls on the synthetic dataset and 2.5 times runtime reduction on a pool of 50 million molecules. The success of this generative approach for active acquisition suggests a promising direction for other scientific discovery problems characterized by vast search spaces where exhaustive evaluation is impractical.

Despite its promising results, our framework has several limitations. First, the effectiveness of BALD-GFlowNet depends on the surrogate model quality. A low-quality surrogate can provide a misleading reward signal, which prevents the GFlowNet from identifying informative regions. Second, the performance of BALD-GFlowNet is closely linked to the design of the reward function. In VS, an imbalance between MI-based rewards and drug-likeness metrics can bias the model to generate either uninformative or chemically infeasible molecules. Finally, our current implementation is tailored to drug discovery. Extending BALD-GFlowNet to other domains would require the development of domain-specific GFlowNet policy networks.

## 6 Acknowledgments and Funding Disclosure

This research was enabled by funds provided by a Canadian Institutes of Health Research (CIHR) doctoral award (FRN: FBD-187593) to MP. AC and ME would like to acknowledge the support by their respective NSERC Discovery grants.

## References

- [1] Christoph Gorgulla, Andras Boeszoermenyi, Zi-Fu Wang, Patrick D. Fischer, Paul W. Coote, Krishna M. Padmanabha Das, Yehor S. Malets, Dmytro S. Radchenko, Yuri S. Moroz, David A. Scott, Konstantin Fackeldey, Moritz Hoffmann, Iryna Iavniuk, Gerhard Wagner, and Haribabu Arthanari. An open-source drug discovery platform enables ultra-large virtual screens. *Nature*, 580:663–668, 2020.

- [2] Jeonghyeon Kim, Juno Nam, and Seongok Ryu. Understanding active learning of molecular docking and its applications, 2024.
- [3] Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- [4] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394, 2021.
- [5] Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.
- [6] Mohit Pandey, Gopeshh Subbaraj, and Emmanuel Bengio. Gflownet pretraining with inexpensive rewards. *arXiv preprint arXiv:2409.09702*, 2024.
- [7] Karoline Gäbler, Iris Behrmann, and Claude Haan. Jak2 mutants (eg, jak2v617f) and their importance as drug targets in myeloproliferative neoplasms. *Jak-Stat*, 2(3):e25025, 2013.
- [8] John J Irwin and Brian K Shoichet. Zinc- a free database of commercially available compounds for virtual screening. *Journal of chemical information and modeling*, 45(1):177–182, 2005.
- [9] Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure FP Dossou, Chanakya Ajit Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, et al. Biological sequence design with gflownets. In *International Conference on Machine Learning*, pages 9786–9801. PMLR, 2022.
- [10] Maksym Korablyov, Cheng-Hao Liu, Moksh Jain, Almer M van der Sloot, Eric Jolicoeur, Edward Ruediger, Andrei Cristian Nica, Emmanuel Bengio, Kostiantyn Lapchevskyi, Daniel St-Cyr, et al. Generative active learning for the search of small-molecule protein binders. *arXiv preprint arXiv:2405.01616*, 2024.
- [11] Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1:1–11, 2009.
- [12] Sivaprakasam Prasanna and Robert J Doerksen. Topological polar surface area: a useful descriptor in 2d-qsar. *Current medicinal chemistry*, 16(1):21–41, 2009.
- [13] Sheng Tian, Junmei Wang, Youyong Li, Dan Li, Lei Xu, and Tingjun Hou. The application of in silico drug-likeness predictions in pharmaceutical research. *Advanced drug delivery reviews*, 86:2–10, 2015.
- [14] Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory balance: Improved credit assignment in gflownets. *Advances in Neural Information Processing Systems*, 35:5955–5967, 2022.
- [15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [16] Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018.
- [17] Julia Boone, Tolunay Seyfi, and Fatemeh Afghah. A joint reconstruction-triplet loss autoencoder approach towards unseen attack detection in iov networks. *IEEE Internet of Things Journal*, 2025.
- [18] Y Gal and Z Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. arxiv [stat. ml]. Retrieved fro <http://arxiv.org/abs/1506.02142>, 2015.
- [19] T Charnock, L Perreault-Levasseur, and F Lanusse. Bayesian neural networks. arxiv e-prints, art. *arXiv preprint arXiv:2006.01490*, 10, 2020.

- [20] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
  - [21] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.
  - [22] Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup help with robustness and generalization? *arXiv preprint arXiv:2010.04819*, 2020.
  - [23] Jerret Ross, Brian Belgodere, Vijil Chenthamarakshan, Inkit Padhi, Youssef Mroueh, and Payel Das. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(12):1256–1264, 2022.
  - [24] Alexander N Shvanyuk, Sergey V Ryabukhin, A Tolmachev, AV Bogolyubsky, DM Mykytenko, AA Chupryna, W Heilman, and AN Kostyuk. Enamine real database: Making chemical diversity real. *Chemistry today*, 25(6):58–59, 2007.
  - [25] Francesco Gentile, Vibudh Agrawal, Michael Hsing, Anh-Tien Ton, Fuqiang Ban, Ulf Norinder, Martin E Gleave, and Artem Cherkasov. Deep docking: a deep learning platform for augmentation of structure based drug discovery. *ACS central science*, 6(6):939–949, 2020.
  - [26] Amr Alhossary, Stephanus Daniel Handoko, Yuguang Mu, and Chee-Keong Kwoh. Fast, accurate, and reliable molecular docking with quickvina 2. *Bioinformatics*, 31(13):2214–2216, 2015.
  - [27] Seul Lee, Seanie Lee, Kenji Kawaguchi, and Sung Ju Hwang. Drug discovery with dynamic goal-aware fragments. *arXiv preprint arXiv:2310.00841*, 2023.
  - [28] Siddarth Venkatraman, Moksh Jain, Luca Scimeca, Minsu Kim, Marcin Sendera, Mohsin Hasan, Luke Rowe, Sarthak Mittal, Pablo Lemos, Emmanuel Bengio, et al. Amortizing intractable inference in diffusion models for vision, language, and control. *arXiv preprint arXiv:2405.20971*, 2024.
  - [29] Gao Huang, Zhuang Liu, and Kilian Q Weinberger. Densely connected convolutional networks. corr. *arXiv preprint arXiv:1608.06993*, 2016.
  - [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
  - [31] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31, 2018.
- [utf8]inputenc [T1]fontenc hyperref url booktabs amsfonts nicefrac microtype xcolor graphicx

## A Technical Appendices and Supplementary Material

### A.1 Synthetic Dataset

The 2D 30x30 grid experiment is conducted on dataset constructed by function  $f(i, j) = \sin(\frac{2\pi i}{50}) \cos(\frac{2\pi j}{50}) + \frac{ij}{10000} + \exp(-\frac{(i-70)^2 + (j-30)^2}{2 \cdot 20^2}) + \varepsilon$ , where  $\varepsilon \sim \mathcal{N}(0, 0.1^2)$  is added as Gaussian noise (see Figure 4 for the full landscape).




Figure 4: The landscape for ground truth function defined for the 2D grid domain.

### A.2 Synthetic Grid Experimental Details

The surrogate BNN uses Monte Carlo Dropout for uncertainty estimation. The positive class for the classification task was defined as any point with a ground-truth function value above the 1st percentile of all outputs. The active learning loop was initialized by training the BNN on a set of 100 points sampled uniformly at random. Then it proceeds for 20 iterations, where each iteration involves acquiring 10 new data points.

The GFlowNet policy network is optimized using the Adam optimizer with a Trajectory Balance loss [14] (see Table 2 for detail). To provide meaningful and structured spatial embeddings for both the BNN and the GFlowNet policy, we train an autoencoder which takes positional encodings of 2D coordinates as input, using reconstruction and triplet margin loss functions (see A.3).

### A.3 Autoencoder Architecture and Training

**Positional Encoding and Autoencoder** Instead of directly inputting the two-dimensional coordinates into the autoencoder, they are first transformed a higher-dimensional embedding. We create separate positional encodings for the x and y coordinates using the method from Vaswani et al. [15]. The two resulting vectors are then concatenated to form the final input for the autoencoder. As shown Table 2, both the encoder and decoder are MultiLayer Perceptrons consisting of multiple linear layers with ReLU activation [16].

**Loss Function** The autoencoder is trained with a composite loss function to produce a spatially aware latent space:

$$\mathcal{L} = \sum_{i=1}^N \|x_i - \hat{x}_i\|^2 + \lambda_{TML} \cdot \sum_{i=1}^N \max(0, \|E(x_i) - E(x_i^p)\|^2 - \|E(x_i) - E(x_i^n)\|^2 + m) \quad (2)$$

The first term is the Mean Squared Error (MSE) between the original coordinate vectors  $x_i$  and their reconstructions  $\hat{x}_i$ . The second term is a Triplet Margin Loss, which encourages the encoder  $E(\cdot)$  to produce a spatially aware latent space [17]. This loss operates on triplets consisting of an anchor  $x_i$ , a positive sample  $x_i^p$ , and a negative sample  $x_i^n$ . See Table 2 for detail.



#### A.4 2D Grid Surrogate Classifier

The surrogate model is a Bayesian neural network (BNN) modeled as a Multi-Layer Perceptron with Monte Carlo dropout for uncertainty estimation [3, 18, 19]. The network consists of three hidden layers with ReLU activations and Dropout for Bayesian inference [16, 20]. The final output layer produces logits over the two classes (see Table 2). The classifier uses F1 accuracy as its performance metric due to the skewed data distribution.

#### A.5 Policy Network Architecture and Training

The policy network is designed to map a latent state representation to a probability distribution over a discrete action made up of eight directions and a stop action. A linear projection first transforms the latent coordinates into a fixed hidden dimension, followed by a LeakyReLU activation [20, 21]. The output is then processed by a stack of transformer encoder layers composed of multi-head self-attention and feedforward blocks, regularized with dropout and LeakyReLU for stable training. A final linear layer maps the hidden features to logits over the action space (refer to Table 2).

#### A.6 2D Grid Training Objective and Exploration

The GFlowNet policy is trained using the Trajectory Balance objective [14]. To encourage exploration and prevent policy collapse, we introduce three mechanisms. First, an  $\epsilon$ -greedy strategy is used during trajectory sampling, allowing the agent to select a random valid action with probability  $\epsilon$ . Second, we enforce a minimum trajectory length during training by masking the stop action with probability  $\epsilon_{stop}$ . On top of this, we implement depth-aware masking on the stop action to encourage longer trajectories. Stop action is allowed with a probability that increases with the current trajectory’s length. In addition, we implement a one-step lookahead that skips any move whose next state would lead to repeated actions or boundaries.

#### A.7 Mixup Strategy

To address class imbalance, we apply mixup directly to the learned latent representations [22]. In the 2D grid scenario, mixup is performed on encoded spatial embeddings from raw coordinates, whereas in the VS case study, it is applied to molecular embeddings generated by MoLFormer [23]. Synthetic examples are generated by creating a weighted average of two embeddings within the same class:

$$\mathbf{v}_{mixup} = \lambda \cdot \mathbf{v}_1 + (1 - \lambda) \cdot \mathbf{v}_2 \quad (3)$$

where the mixing coefficient  $\lambda \sim Uniform[0, 1]$ . In the case study, the attention mask is formed by the union of the parent masks, preserving the contextual scope of both source molecules.

#### A.8 Virtual Screening Implementation Details

**Dataset Acquisition and Processing** Our dataset is derived from the Enamine REAL database, a collection of synthetically feasible molecules well-suited for virtual screening campaigns [24]. To construct our target-specific dataset, we first selected 2 million of these molecules while using the Deep Docking protocol for Janus Kinase 2 (JAK2) to ensure most potential hits are included in the dataset [25, 7].

From the dataset, we randomly sampled 10,000 molecules as training dataset and 100,000 molecules as held-out test set, with the remaining compounds forming the unlabeled pool. At each active learning iteration, we actively acquire 100 molecules which are subsequently labeled by QuickVina 2 [26, 27]. To adapt the dataset for classification, we convert docking scores into binary labels using a cutoff at the 0.01 quantile of the most negative docking scores in the initial training dataset. The positive class denotes strong binding affinity and the negative class indicates weak or no interaction.

**Framework Adaptation for the Virtual Screening Task** We first train a surrogate classifier that adopts a MoLFormer-based neural network, on a dataset of 10,000 randomly sampled molecules from the Enamine REAL database [24, 23]. Docking scores are converted to binary labels using a stringent 1% cutoff, and we use mixup technique to handle the resulting imbalance so the model can generalize better to rare small molecules with high affinity [22]. The classifier is then re-trained in each iteration of the generation process.

We then compare three acquisition strategies: (i) BALD, which exhaustively scores all unlabeled points and selects the top-k by mutual information [3]; (ii) Random, which selects points uniformly from the unlabeled pool; and (iii) GFlowNet, which samples points from a trained generative policy.

**GFlowNet Fine-tuning and Regularization** To further maintain the feasibility of generated compounds, we regularize the finetuning process with a Relative Trajectory Balance (RTB) objective and an offline dataset of feasible molecules [6, 28].

**Surrogate Classifier** Our surrogate model utilizes the MoLFormer encoder architecture to generate embeddings from molecular SMILES strings [23]. These tokenized embeddings are then processed by a classification head, consisting of Multi-Layer Perceptron and Multi-Head Attention layers, to yield a final prediction. First, the attention pooling layer assigns a different weight to each token, creating a unifying contextualized representation of the molecule [15]. Then this pooled representation goes through a DenseNet-like architecture that allows the model to build progressively more complex and hierarchical features [29]. The refined output passes through a self-attention mechanism to refine its representation by re-weighting the importance of different internal features. See Table 3 for a summary of the model hyperparameters.

## B Additional Plots and Results

### B.1 Oracle Call Comparison on the Synthetic Grid Task

Figure 5 compares the number of oracle calls between BALD-GFlowNet and the BALD baseline. Here the number of oracle calls is defined as the cumulative count of unique terminal states selected for labeling, with each state corresponding to a unique query to the surrogate model. Compared to the baseline’s constant  $O(n)$  calls, BALD-GFlowNet requires much fewer oracle calls.




Figure 5: Oracle calls for BALD-GFlowNet and the BALD baseline in the 2D grid environment for a single AL step. The Random baseline does not query the surrogate to compute acquisition rewards for selected samples, so its call count is zero.

### B.2 Peak F1 Score and Query Number Comparison on VS Task

This table summarizes the peak F1 Score against the total computational cost, measured in cumulative number of queries to the oracle required to reach that peak. BALD-GFlowNet achieves a competitive F1 score while being approximately 7 times more computationally efficient than the standard BALD baseline.

Table 1: Peak performance and acquisition queries for each strategy.

Method	Peak F1	Peak Iteration	Queries at Peak
<b>BALD-GFlowNet</b>	0.30	14	1.74M
BALD	0.32	7	12.2M
Random	0.28	5	0

Table 2: Hyperparameter configuration for models in the grid experiment.

Hyperparameter	Value
<i>Active Learning Pipeline</i>	
Total AL Steps	20
Initial Dataset Size	100
Acquisition Size per Step	10
Test Set Size	100
Grid Size	$30 \times 30$
<i>Coordinate Autoencoder</i>	
Positional Encoding Dimension	128
Number of Layers (Encoder/Decoder)	6
Hidden Dimension	512
Latent Dimension	50
Activation Function	ReLU [16]
Loss Function	MSE + Triplet Margin Loss [17]
Triplet Loss Weight ( $\lambda_{\text{TML}}$ )	0.1
Triplet Loss Margin ( $m$ )	1.0
Optimizer	Adam [30]
Learning Rate	1e-3
Training Epochs	50
Batch Size	512
<i>Surrogate BNN Model</i>	
Input Dimension	50
Hidden Dimension	256
Number of Hidden Layers	1
Dropout Rate	0.1
Number of Dropout Models	3
Optimizer	Adam [30]
Learning Rate	1e-3
Weight Decay	1e-4
Training Epochs	7
Loss Function	Cross Entropy Loss [31]
<i>GFlowNet Policy (Transformer)</i>	
Input Dimension	50
Hidden Dimension	256
Number of Encoder Layers	6
Number of Attention Heads	8
Feed-Forward Dimension	1024
Dropout Rate	0.1
Number of Actions	5
Optimizer	Adam [30]
Learning Rate	1e-4
Loss Function	Trajectory Balance (TB) [14]
Training Episodes per AL Step	50000
Min Trajectory Length	50
Max Trajectory Length	100
$\epsilon$ -Greedy	0.1
$\epsilon$ -Stop	0.5
Initial Partition	10

Table 3: Hyperparameter configuration for models involved in the BALD-GFlowNet pipeline.

Hyperparameter	Value
<i>BALD-GFlowNet Pipeline</i>	
Active Learning Iterations	30
Initial Dataset Size	10000
Acquisition Size	100
Initial Pool Size	1741777
<i>GFlowNet Model</i>	
Loss Function	RTB [28]
Loss Coefficient	0.04
MLE Coefficient	20
MI Reward Range	[0, 1]
MI Task Slope	1 [6]
Beta	96
Number of GNN Layers	15
Number of Heads per GNN Layer	4
Number of Hidden Dimensions	128
Number of MLP Layers	10
Sampling Temperature $T$	1
Optimizer	Adam [30]
Learning Rate	1e-6
Weight Decay	1e-8
Momentum	0.9
Second Momentum	0.999
Epsilon	1e-8
Batch Size	64
Training Steps	3881
<i>Classifier Model</i>	
Loss Function	Cross Entropy Loss [31]
Training Epochs	100
Batch Size	1024
Optimizer	Adam [30]
Patience	10
Learning Rate	0.000001
Learning Rate Scheduler	gamma 0.1, step size 40
Weight Decay	0.01
Base Model	MoLFormer [23]
Number of MC Dropout Models	3
Dropout Rate	0.1
Number of Dense Layers	3
Hidden Dimensionality	768
Intermediate Dimensionality	3072
Multi-head Attention Heads	8
Residual Blocks	2

## C Additional Information

### C.1 Compute Resources

All experiments were conducted on a private institutional GPU cluster managed by the Slurm workload manager using NVIDIA A100-PCIE-40GB GPUs (40,960 MiB memory, CUDA 12.4). The three methods from the virtual screening case study used 3 GPUs for 13 days, requiring approximately 25,800 MiB of GPU memory. The hypergrid experiment ran on a single GPU for 1 day with memory

usage of 5,000 MiB. In total, the reported experiments required approximately 40 GPU-days of computation. No significant additional compute was used for unreported preliminary experiments.

## C.2 License

The table below provides a list of softwares, datasets, packages and similar with their license and terms of use:

Table 4: Licenses and terms of use for assets used in this work.

Asset	License	Terms of Use / License URL
IBM MoLFormer	Apache-2.0	<a href="https://www.apache.org/licenses/LICENSE-2.0">https://www.apache.org/licenses/LICENSE-2.0</a>
Enamine REAL DB	Not Applicable	<a href="https://enamine.net/terms-of-use">https://enamine.net/terms-of-use</a>
QuickVina2	Apache-2.0	Not Applicable
PyTorch	BSD Style License	<a href="https://github.com/pytorch/pytorch/blob/main/LICENSE">https://github.com/pytorch/pytorch/blob/main/LICENSE</a>
RDKit	BSD 3-Clause	<a href="https://github.com/rdkit/rdkit/blob/master/license.txt">https://github.com/rdkit/rdkit/blob/master/license.txt</a>

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer:

Justification: We claim that our method BALD-GFlowNet achieves comparable F1 score with significantly fewer oracle calls than the BALD baseline. We confirm this claim with experiments on the two-dimensional grid and virtual screening case study (see section 4).

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer:

Justification: Our paper discusses that our method depends on the quality of the surrogate model. Additionally, under the virtual screening case study, our method requires a balance between the BALD reward and drug likeliness metrics (see section 5).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer:

Justification: Not applicable, our paper does not involve any theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer:

Justification: The synthetic dataset from the grid experiment and the dataset used for virtual screening are easy to generate and we describe how to generate it in the Appendix A.1 and A.8. For the surrogate model architecture and policy networks in both cases, we describe them in detail (Appendix A.4, A.5, and A.8) with hyperparameters from Table 2 and Table 3. We are planning to release the datasets and Github repository shortly.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).



- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer:

Justification: We are planning to release the datasets and Github repository shortly.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer:

Justification: We have provided details about hyperparameter, optimizer type, dataset construction and splits in the Appendix A.1 and A.8, as well as Table 2 and 3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer:

Justification: In addition the mean F1 score, we have also reported the standard deviation across the three different seeded runs for the virtual screening case study along with the method for computing the statistics in Figure 1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

#### 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer:

Justification: We mentioned the GPU, type of cluster, memory, storage, estimated total compute time in GPU days as well as the amount of total compute for each individual experiment in subsection C.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

#### 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer:

Justification: We respect the code of ethics listed under NeurIPS.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer:

Justification: Paper presents a method for improving the scalability of active learning. As a generic method, it could be used in applications with positive or negative societal impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer:

Justification: We do not release pretrained language models, image generators, scraped datasets, or similar.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer:

Justification: We explicitly mention the licenses of assets with their terms of use in Table 4.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

### 13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer:

Justification: We do not introduce any new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer:

Justification: We do not perform crowdsourcing experiments and research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer:

Justification: We do not perform crowdsourcing or experiments related to human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

#### 16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer:

Justification: The core method in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.