

---

# Graph Neural Bayesian Optimization for Virtual Screening

---

Miles Wang-Henderson \*  
ETH Zurich

Bartu Soyuer  
ETH Zurich

Parnian Kassraie  
ETH Zurich

Andreas Krause  
ETH Zurich

Ilija Bogunovic  
University College London

## Abstract

Virtual screening is an essential component of early-stage drug and materials discovery. This is challenged by the increasingly intractable size of virtual libraries and the high cost of evaluating properties. We propose GNN-SS, a Graph Neural Network (GNN) powered Bayesian Optimization (BO) algorithm. GNN-SS utilizes random sub-sampling to reduce the computational complexity of the BO problem, and diversifies queries for training the model. We further introduce data-independent projections to efficiently model second-order random feature interactions, and improve uncertainty estimates. GNN-SS is computationally light, sample-efficient, and rapidly narrows the search space by leveraging the generalization ability of GNNs. Our algorithm achieves state-of-the-art performance amongst screening methods on the Practical Molecular Optimization (PMO) benchmark.

## 1 Introduction

Molecular optimization is a key problem in early drug discovery and materials design, wherein the objective is to retrieve candidate compounds with one or more desirable properties. Some properties of interest may include target-ligand binding affinities or quantum-mechanical energies. However, these tasks are often complicated by 1) the combinatorially large size of accessible chemical space, and 2) the prohibitive cost of evaluating objectives *in silico* or via experiment. Thus, a first-line solution is to perform molecular optimization on finite but *large* virtual libraries, where candidates can be readily synthesized and validated.

As libraries grow to cardinalities exceeding  $10^9$ , brute-force enumeration for candidate selection becomes increasingly infeasible. To address this challenge, one promising strategy is to perform Bayesian Optimization, a model-based approach that leverages a surrogate function to evaluate molecular properties. BO techniques are renowned for their ability to rapidly narrow down the search space to the most desirable molecules. In particular, by modeling molecules as graphs, we can use BO methods which exploit the graph structure and utilize Graph Neural Networks (GNNs) to capture the complex molecular properties and improve the sample efficiency of our procedure. However, deployment of GNNs for BO on graphs is a computationally challenging problem.

Specifically, a significant issue arises when estimating uncertainty in such models, as conventional neural approaches necessitate updating a corresponding covariance matrix quadratic in the number of network parameters [Zhou et al., 2020], leading to an inherent trade-off between model expressiveness and computational feasibility. Additionally, while performing BO with a GNN-based acquisition func-

---

\*Correspondence to [mwanghenders@ethz.ch](mailto:mwanghenders@ethz.ch).

tion improves the statistical complexity of the optimization problem, maximizing some acquisition function over vast discrete libraries is still computationally prohibitive.

We address these challenges by introducing GNN-SS, a light GNN-powered BO algorithm that randomly sub-samples the domain before optimizing the acquisition function. This reduces the computational cost and increase the diversity of collected data, which in turn, enhances exploration to improve the generalization ability of the network. To avoid prohibitive matrix inversions required to calculate uncertainty estimates we also employ Johnson-Lindenstrauss (JL) random-feature mappings. In our setting, applying this mapping can be interpreted as sub-sampling the weights of a trained neural network, yielding tractable matrix inversions.

We then evaluate our algorithm with a limited sampling budget on a domain of commercially available small-molecules. We observe that when paired with the Upper Confidence Bound (UCB) acquisition function, our algorithm GNN-SS achieves state-of-the-art performance on the Practical Molecular Optimization (PMO) benchmark among algorithms for ligand-based virtual screening [Gao et al., 2022], even outperforming a number of generative *de novo* algorithms for molecular design.

**Related Works.** Early work on BO for molecular optimization use hand-designed graph kernels [Korovina et al., 2020], or map graph representations to a continuous latent space [Gómez-Bombarelli et al., 2018], to then solve the problem with Gaussian process BO methods, e.g. GP-UCB. Recent work utilize GNNs, as they exploit the structure of objective functions defined on graphs. Such approaches require quantifying the uncertainty of neural network estimates, which may be done by parameterizing the uncertainty itself with a neural network [Graff et al., 2021, Soleimany et al., 2021], probabilistic ensembles [Kim et al., 2021, Hirschfeld et al., 2020], or directly upper bounding the uncertainty in the Neural Tangent Kernel (NTK) regime [Kassraie et al., 2022]. These methods require calculating the acquisition function over the entire domain, and thus, can not be scaled to large molecular libraries. Building upon Kassraie et al. [2022], we propose a scalable algorithm which exhibits competitive performance, while reducing the computational complexity of prior work. To this end, we leverage the simple idea of random sub-sampling, which is grounded in the literature of unstructured many-armed bandits [Mirzasoiman et al., 2015, Bayati et al., 2020].

## 2 Problem Setting

We consider optimization problems on domains of small molecules that emerge in drug or material design. Small-molecules admit a natural representation as undirected graphs with at most  $N$  nodes [Brown et al., 2019, Du et al., 2022]. This allows us to formulate such optimization problems as Bayesian optimization on a domain of graphs, where the aim is to optimize an *unknown* objective function through sequential queries while receiving noisy function evaluations. The BO objective function models the unknown molecular property of interest, and querying a graph corresponds to recommending a molecule to be tested with respect to this property.

At every time step  $t \in \{1, \dots, T\}$ , we select a graph  $G_t$  from a graph domain  $\mathcal{G}$  and observe a noisy evaluation  $y_t = f(G_t) + \epsilon_t$ , where  $f: \mathcal{G} \rightarrow \mathbb{R}$  is the objective function and  $\epsilon_t$  is i.i.d. zero-mean sub-Gaussian noise. We denote the history of observations by  $H_t := \{(G_1, y_1), \dots, (G_t, y_t)\}$ . For a horizon  $T$  (i.e., sampling budget), we seek to obtain a small *simple* regret  $r_T = f(G^*) - \max_{t \leq T} f(G_t)$ , where  $G^* \in \arg \max_{G \in \mathcal{G}} f(G)$ . The aim is to attain a regret that *vanishes* with  $T$ , meaning that  $r_T \rightarrow 0$  as  $T \rightarrow \infty$ , which implies convergence to the optimal graph. Evaluating molecular properties typically requires running costly simulations or experiments with noisy outcomes. Therefore, we seek a sample efficient algorithm that can identify the optimal molecule with the least number of oracle calls.

In our applications of interest the number of accessible molecules ranges between  $10^{10}$ - $10^{20}$  choices [Nicolaou et al., 2016], and popular virtual databases of commercially available compounds include up billions of molecules [e.g. ZINC, Irwin and Shoichet, 2005]. Therefore, we assume that the domain  $\mathcal{G}$  is a very large finite set of graphs, where  $|\mathcal{G}| \gg T$ , meaning that it is infeasible to evaluate every graph in the domain within the sampling budget of  $T$ . Standard BO algorithms for large unstructured domains [e.g., Auer et al., 2002] cannot be applied to this problem setting, as their sample complexity grows with  $\text{poly}(|\mathcal{G}|)$ .

### 3 Method

To obtain a sample efficient algorithm for maximizing the objective function, we need to construct an estimator for  $f$ , using a relatively small number of samples from  $\mathcal{G}$ . To this end, we utilize three key ideas. First, we use Graph Neural Networks which are known to be effective models for learning complex functions defined on graphs. Second, in order to efficiently update our uncertainty over the network’s outputs we use a random-feature approximation to the NTK. This allows us to employ rank-one updates in  $O(tp^2)$  time, where  $p$  is the dimension of our random features and corresponds to the number of network parameters. While previous works [e.g., Zhou et al., 2020, Kassraie et al., 2022] address this by employing a diagonal approximation to the covariance matrix, we aim to model second-order covariance structure by utilizing a low-rank approximation to the empirical NTK. As demonstrated in Figure 2, this helps us draw samples which are informative and yield high reward. Lastly, we devise a sub-sampling (SS) subroutine that allows for sampling diverse data for training the network. These techniques come together in our GNN-powered Sub-sampled BO algorithm (GNN-SS), sketched in Algorithm 1. The sub-routine TRAINGNN is specified in Algorithm 2.

---

#### Algorithm 1 GNN-SS

---

**Input:** acq. function  $\alpha(\cdot)$ , horizon  $T$ , warm-start steps  $T_0$ , random subset size  $K$ , batch size  $b \leq K$ , training epochs  $E$ , learning rate  $\eta$ , regularization parameter  $\lambda$   
**Initialize:**  $\mathbf{K}_0 \leftarrow \lambda \mathbf{I}$ ,  $\boldsymbol{\theta}_0 \sim \mathcal{N}(0, \mathbf{I}_p)$ ,  $H_0 = \emptyset$ .

---

**for** round  $t = 1, \dots, T_0$  **do** ▷ Exploration Stage

**a) Sample random action.**  $G_t \sim \text{Unif}(\mathcal{G})$  and query  $y_t = f(G_t) + \epsilon_t$

**b) Add to history.**  $H_t \leftarrow H_{t-1} \cup \{(G_t, y_t)\}$

**end for**

$\boldsymbol{\theta}_{T_0} \leftarrow \text{TRAINGNN}(H_{T_0}, \boldsymbol{\theta}_0, E, \lambda, \eta)$

$\mathbf{K}_{T_0} \leftarrow \mathbf{K}_0 + \sum_{t=1}^{T_0} \sum_{i=1}^b \mathbf{g}(G_{t,i}) \mathbf{g}^\top(G_{t,i})$

---

**for** round  $t = T_0 + 1, \dots, T$  **do** ▷ Adaptive Stage

**1) Sub-sample.**  $\mathcal{G}_t \sim \text{Unif}(\mathcal{G})$  where  $|\mathcal{G}_t| = K$

**2) Evaluate acquisition.**  $\{G_{t,i}\}_{i=1}^b \leftarrow \arg \max_{G \in \mathcal{G}_t} \alpha(G; H_{t-1})$

**2) Observe noisy reward.**  $y_{t,i} = f(G_{t,i}) + \epsilon_{t,i}$  for  $i \in [b]$

**3) Append to history.**  $H_t \leftarrow H_{t-1} \cup \{(G_{t,1}, y_{t,1}), \dots, (G_{t,b}, y_{t,b})\}$

**4) Update surrogate parameters.**  $\boldsymbol{\theta}_t \leftarrow \text{TRAINGNN}(H_t, \boldsymbol{\theta}_{t-1}, E, \lambda, \eta)$

**4) Update uncertainty.**  $\mathbf{K}_t \leftarrow \mathbf{K}_{t-1} + \sum_{i=1}^b \mathbf{g}(G_i) \mathbf{g}^\top(G_i)$

**end for**

---

**Reward Model.** To model molecular properties, we use a graph neural network  $f_{\text{GNN}}(G; \boldsymbol{\theta}) : \mathcal{G} \rightarrow \mathbb{R}$  of width  $m$ , where  $\boldsymbol{\theta} \in \mathbb{R}^p$  denotes the concatenated vector of network parameters. We initialize the network with  $\boldsymbol{\theta}_0$  which has zero-mean Gaussian i.i.d. entries of unit variance, and update it as we observe more feedback. We maintain a reward surrogate  $f_{\text{GNN}}(G; \boldsymbol{\theta}_t)$ , the GNN trained on the loss

$$\mathcal{L}(\boldsymbol{\theta}; H_t) = \frac{1}{t} \sum_{i=1}^t (f_{\text{GNN}}(G_i, \boldsymbol{\theta}) - y_i)_2^2 + m\lambda \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2$$

where  $\lambda$  is the regularization coefficient. Similar to Kim et al. [2022] and Zhang et al. [2020], we initialize the network at every step of the BO problem with  $\boldsymbol{\theta}_{t-1}$ , and update the parameters via gradient descent for  $E$  epochs. This is in contrast to earlier work on Neural Bandits which re-initialize the network to  $\boldsymbol{\theta}_0$  before running gradient descent [e.g. Zhou et al., 2020, Kassraie et al., 2022].

**Uncertainty Quantification.** For action-selection policies which require uncertainty quantification, e.g. UCB [Srinivas et al., 2010], we may use the gradient of the network at initialization to approximate the uncertainty over  $f$ . Following Kassraie et al. [2022], we set

$$\hat{\sigma}_t^2(G) := \mathbf{g}^\top(G) \left[ \lambda \mathbf{I} + \frac{1}{t} \sum_{i=1}^t \mathbf{g}(G_i) \mathbf{g}^\top(G_i) \right]^{-1} \mathbf{g}(G), \quad (1)$$

where  $\mathbf{g}(G) = \nabla_{\boldsymbol{\theta}} f_{\text{GNN}}(G; \boldsymbol{\theta}_0)$  denotes the gradient at initialization. For very wide networks,  $\hat{\sigma}_t$  presents a provably calibrated uncertainty estimate [Kassraie et al., 2022, Theorem 4.2]. In this

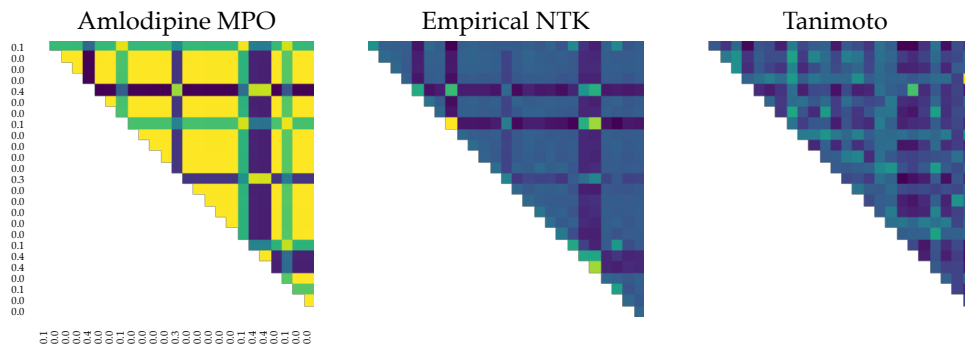


Figure 1: Empirical NTK ( $\hat{K}$ ) effectively captures shared structure between randomly sampled molecules of similar reward, and is more informative than the commonly used Tanimoto kernel ( $K_{\text{Tan}}$ ). To obtain  $\mathbf{g}(G_i)$  we use a Graph Convolutional Network with width 2048, 3 convolutional layers, and mean readout.

formulation, the Gram matrix with elements  $\hat{K}_{ij} = \mathbf{g}^\top(G_i)\mathbf{g}(G_j)$  performs a similar role as the kernel matrix when calculating the posterior variance in given a Gaussian Process. For this reason, among others, the kernel function  $\hat{k}(G, G') = \mathbf{g}^\top(G)\mathbf{g}(G')$  is referred to as the empirical Neural Tangent Kernel. Figure 1 demonstrates that the matrix  $\hat{K}$ , effectively encodes the structural similarity between molecules with similar rewards. In this figure, the left-most matrix has elements  $1 - |y(G_i) - y(G_j)|$ , where  $y(G_i)$  is the reward of molecule  $G_i$ . Here, reward corresponds to the Amlodipine MPO objective, which is a function of the number of aromatic rings as well as the ECP4 fingerprint similarity to the Amlodipine molecule [Brown et al., 2019]. Additionally, we compute a kernel matrix using the Tanimoto kernel [Tanimoto, 1958], a commonly used similarity metric for sets. When applied to molecules, this measures the intersection of shared substructures between  $G$  and  $G'$ , and is computed as  $k_{\text{Tan}}(G, G') = \sum_i \min(\mathbf{h}(G)_i, \mathbf{h}(G')_i) / \sum_i \max(\mathbf{h}(G)_i, \mathbf{h}(G')_i)$ , where  $\mathbf{h}(G) \in \{0, 1\}^{|V|}$  is a vector indicating the presence or absence of substructures in  $G$  from a finite vocabulary  $V$ .

**Numerical Approximation.** Computing the uncertainty estimates introduced in (1) requires  $p \times p$  matrix inversions, where  $p$  denotes the number of trainable parameters in the GNN. This operation is the memory and computation bottleneck of the algorithm, and quickly becomes intractable for even modest-sized networks. A common approach in neural bandit algorithms is to use a diagonal approximation of the matrix  $\mathbf{g}(G_i)\mathbf{g}^\top(G_i)$ , leading to updates linear in the number of parameters  $\mathcal{O}(p)$  [Zhou et al., 2020]. Asymptotically, for large enough width  $m$ , this is a valid mean-field approximation as the parameters become independent [Yang, 2020]. However, for finite neural networks of practical size, this assumption becomes vacuous. Instead, we apply a data-independent random projection  $S: \mathbb{R}^p \rightarrow \mathbb{R}^d$ , with  $d \ll p$ , to reduce the dimensionality of this matrix. We sample  $S$  from an appropriate distribution, such that it is a Johnson-Lindenstrauss (JL) linear map [Johnson and Lindenstrauss, 1984], as JL maps can be designed to satisfy

$$|\langle S\mathbf{g}(G), S\mathbf{g}(G') \rangle - \langle \mathbf{g}(G), \mathbf{g}(G') \rangle| \leq \epsilon_D,$$

if  $d \geq 4\epsilon_D^{-2} \log(1/\delta)$ , where  $\epsilon_D$  is the distortion error and  $\delta \in (0, 1)$  is the confidence level [Johnson and Lindenstrauss, 1984]. Therefore, by specifying a small projected dimension  $d$ , we can control the distortion error that propagates into the uncertainty quantification formula. Setting  $\tilde{\mathbf{g}}(\cdot) = S\mathbf{g}(\cdot)$  we then define the *cheap* uncertainty estimates as

$$\tilde{\sigma}_t^2(G) := \tilde{\mathbf{g}}^\top(G) \left[ \lambda \mathbf{I} + \frac{1}{t} \sum_{i=1}^t \tilde{\mathbf{g}}(G_i)\tilde{\mathbf{g}}^\top(G_i) \right]^{-1} \tilde{\mathbf{g}}(G). \quad (2)$$

Calculating  $\tilde{\sigma}_t^2$  now only requires inverting a  $d \times d$  matrix. Proposition 3 in Appendix A proves that the approximated empirical NTK defined as  $\tilde{k}(G, G') = \tilde{\mathbf{g}}^\top(G)\tilde{\mathbf{g}}(G')$ , is with high probability close to  $k(G, G')$  the empirical NTK for every pair of graphs  $(G, G') \in \mathcal{G} \times \mathcal{G}$ . We also empirically observe that the numerical error introduced due to using  $\tilde{\mathbf{g}}$  instead of  $\mathbf{g}$  is negligible, for practical values of  $d$ . This is expected, due to the closeness of the two corresponding kernels.

**Speeding Up Approximation.** The canonical construction of  $S$  is to use a dense matrix of i.i.d. Gaussian entries, requiring  $\mathcal{O}(pd)$  time to obtain  $\tilde{\mathbf{g}}(G)$  [Johnson and Lindenstrauss, 1984]. Instead, we

employ two computationally-efficient variants. The first is a Fast JL (FJL) variant requiring  $\mathcal{O}(p \log p)$  time to obtain  $\tilde{\mathbf{g}}(G)$  [Ailon and Chazelle, 2009],

$$S_{\text{FJL}} = PHD$$

Here,  $P \in \mathbb{R}^{d \times p}$  is a sparse random matrix with elements  $P_{ij} \in \{0, \sqrt{p/d}\}$  and  $\|P_i\|_\infty = \sqrt{p/d}$ .  $H \in \mathbb{R}^{p \times p}$  is a Walsh-Hadamard matrix with  $H_{ij} \in \{-1, 1\}$ , and  $D \in \mathbb{R}^{p \times p}$  is a diagonal matrix with Rademacher entries  $D_{ii} \in \{-1, 1\}$ . Our second row sub-sampling variant (Sub) is simply

$$S_{\text{Sub}} = P.$$

This variant is inspired by [Weinberger et al., 2009, Yang et al., 2015] and requires only a time of  $\mathcal{O}(d)$ . The distortion using this variant can also be controlled with high probability when the mass along coordinates of  $\mathbf{g}(G)$  is sufficiently uniform, i.e. when  $\|\mathbf{g}(G)\|_\infty / \|\mathbf{g}(G)\|_2$  is sufficiently small for all  $G \in \mathcal{G}$  [Freksen, 2021]. Table 1 summarizes runtime of each of our uncertainty estimators when rank-one updates are employed.

**Acquisition Functions.** Given our reward surrogate  $f_{\text{GNN}}(\cdot; \boldsymbol{\theta}_t)$  and uncertainty estimator  $\tilde{\sigma}_t(\cdot)$ , we construct the following acquisition functions,

$$\begin{aligned} \alpha_{\text{Greedy}}(G; H_t) &= f_{\text{GNN}}(G; \boldsymbol{\theta}_{t-1}), \\ \alpha_{\text{UCB}}(G; H_t) &= f_{\text{GNN}}(G; \boldsymbol{\theta}_{t-1}) + \beta \tilde{\sigma}_{t-1}(G) \end{aligned}$$

Here, the GREEDY acquisition function directly uses the reward surrogate as a proxy for the true unknown reward. The UCB function considers an optimistic estimate of the reward by adding the uncertainty estimate to the reward estimate. The parameter  $\beta$  is the exploration coefficient which controls the level of optimism.

**Random Sub-sampling.** The common approach to action selection is to choose  $G_t$  via  $\max_{G \in \mathcal{G}} \alpha(G; H_t)$ . However, this will require evaluating  $\alpha$  for all members of  $\mathcal{G}$ , which comes with a high computational burden when working on large virtual libraries. We utilize a sub-sampling routine where at every step  $t$  we first draw a random subset  $\mathcal{G}_t \subset \mathcal{G}$ , where  $|\mathcal{G}_t| \ll |\mathcal{G}|$ . We then rank the members of  $\mathcal{G}_t$  according to our acquisition function and query the top- $b$  graphs, i.e.  $\{G_{t,i}\}_{i=1}^b \leftarrow \arg \max_{G \in \mathcal{G}_t} \alpha(G; H_t)$ . Choosing  $b > 1$  yields a simple batched variant of GNN-SS. While this is primarily done to avoid the cost of training, such a procedure could also correspond to a setting in which we sample and rank candidates from a generative model.

The advantage of random sub-sampling is threefold. 1) It saves computations at each round, by allowing us to evaluate  $\alpha(\cdot; H_t)$  only over  $K = |\mathcal{G}_t|$  molecules. 2) Sub-sampling limits the number of graphs available for query, including more promising molecules which drive GNN-SS to exploit. This implicitly encourages exploration, which is believed to be beneficial on large finite domains with complex objectives [Bastani and Bayati, 2020]. 3) This additional exploration helps build a diverse dataset  $H_t$  for training  $f_{\text{GNN}}$  in the future steps, and in turn, yields an accurate reward estimator  $f_{\text{GNN}}(G; \boldsymbol{\theta}_{t-1})$  that generalizes well over the entire domain. The choice of  $K$  is crucial to balancing exploration and exploitation. Moreover, it should be chosen such that the probability of not ever sampling the optimal molecule is small. In our experiments we choose it roughly as  $\mathcal{O}(\sqrt{T})$ , more details are included in Appendix B.

**Exploration Stage.** We start the algorithm with an exploration stage, during which we randomly query a small number ( $\sim 10^2$ ) of molecules. The idea of an exploratory stage stems from the high-dimensional bandit literature [e.g., Hao et al., 2020, Bastani and Bayati, 2020], where additional exploration is required to obtain an accurate estimate of the reward.

Estimator	Runtime
Full	$\mathcal{O}(Np^2)$
Diagonal	$\mathcal{O}(Np)$
Sub	$\mathcal{O}(Nd^2 + Nd)$
FJL	$\mathcal{O}(Nd^2 + Np \log p)$

Table 1: Uncertainty estimates via approximate random-features  $\tilde{\mathbf{g}}$ , yielding covariance matrices with second-order structure.  $N = tb$ , where  $t$  is the number of rounds,  $b$  is batch-size. JL variants project  $\mathbf{g} \in \mathbb{R}^p$  to  $\tilde{\mathbf{g}} \in \mathbb{R}^d$ , where  $d < N$ .

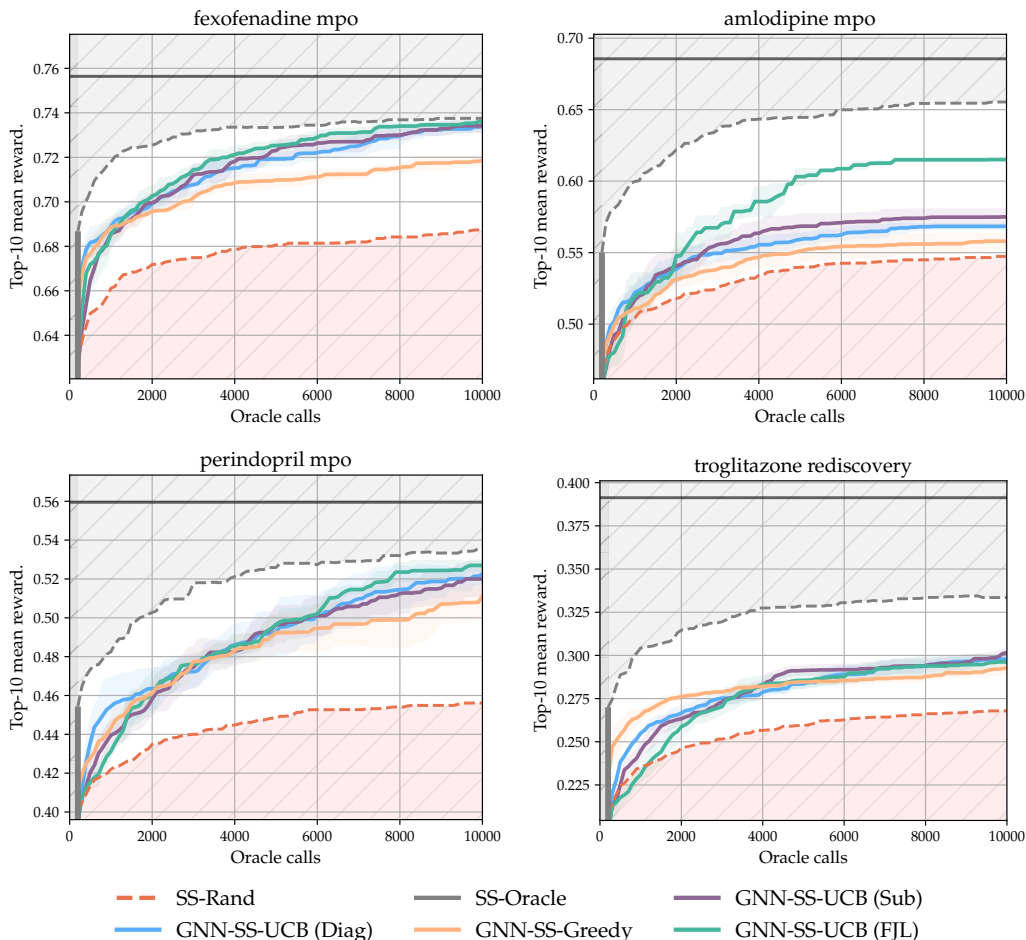


Figure 2: GNN-SS Top-10 Mean Reward. Projection dimension is  $d = 2048$  for both Sub and FJL variants for uncertainty estimator  $\tilde{\sigma}_t(\cdot)$ . Bold horizontal line shows ZINC-250K maximum and the gray region above GNN-SS-Oracle is unreachable by sub-sampled algorithms.

## 4 Results

We perform experiments on ZINC [Irwin and Shoichet, 2005], a dataset of 250k small organic molecules that is commonly used for benchmarking algorithms for molecular optimization [e.g., in Gao et al., 2022, Kim et al., 2022, Graff et al., 2021]. Our experiments demonstrate the application of GNN-SS for virtual screening on these datasets, based on molecular properties which are of interest in drug and materials design. In these experiments, we consider all 23 molecular objectives of the PMO benchmark [Gao et al., 2022], which include commonly used scores such as quantitative estimate of drug-likeness (QED) [Bickerton et al., 2012]. To model these objectives, we employ a vanilla Graph Convolutional Network (GCN) similar to Graff et al. [2021] and for training use different sets of hyper-parameters (details in Appendix B). Reported results are averaged across 5 runs with different random seeds, plotted together with the standard error. Figure 2 shows the performance of GNN-SS paired with UCB, and GREEDY policies on 4 objective functions. Regret plots for additional reward functions are given in Appendix C.

We limit the horizon to a realistic sampling budget, and for every  $t$ , we compute the top-1 reward, i.e.  $\max_{s \leq t} f(G_s)$  the value of the highest scoring molecule queried so far, as well as the top-10 mean reward. As a sanity check, we also run SS-Rand which simply draws a random graph from  $\mathcal{G}_t$ , and SS-Oracle which has oracle knowledge of  $f$  and queries  $\arg \max_{G \in \mathcal{G}_t} f(G)$  after random sub-sampling. A sub-sampled policy cannot achieve a performance higher than SS-Oracle, and should perform better than SS-Rand.



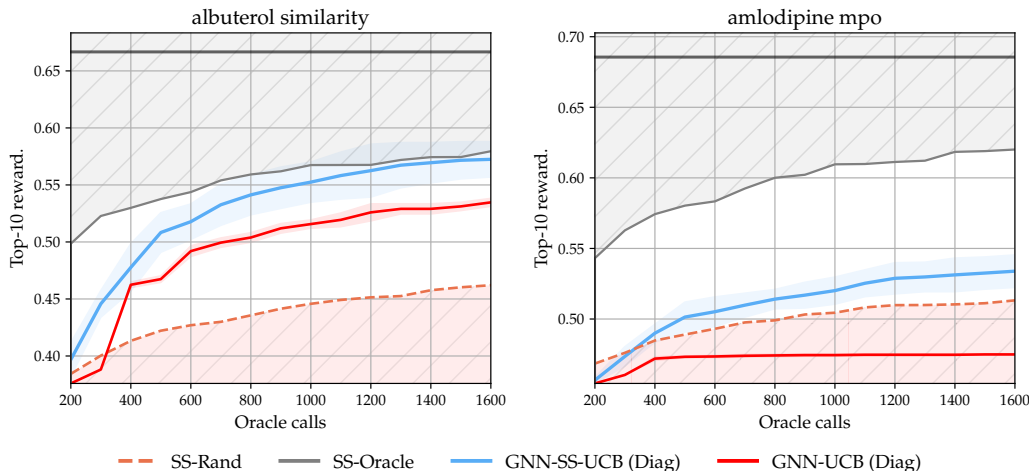


Figure 3: GNN-SS-UCB and GNN-UCB Top-10 Mean Reward. Random subsampling in GNN-SS encourages acquisition of diverse molecules early on during optimization. GNN-UCB evaluates the acquisition function  $\alpha(\cdot; \theta_t)$  on the full ZINC-250K dataset.

To compare our algorithm to prior work we use the PMO benchmark, in which the total number of queried molecules, i.e. the number of oracle calls, is restricted to a budget of  $T = 10,000$ . This benchmark considers 29 algorithms for virtual screening, and evaluates the performance of each by area-under-the-optimization-curve (AUC, as plotted in Figure 2), summed across the 23 objective functions. As reported in Table 3, GNN-SS achieves state-of-the-art performance among non-generative models, including MOLPAL, ranking 4th overall in sample efficiency among the 29 algorithms considered in this benchmark. Importantly, GNN-SS outranks several generative algorithms which go beyond ZINC for maximizing the properties of interest, e.g. SMILES-LSTM Brown et al. [2019] and MARS Xie et al. [2021].

As observed in Figure 2 and Table 2, Sub and FJL variants of our algorithm obtain higher reward molecules overall, ranking 8th and 7th respectively, and first among screening methods. We conjecture that enacting GNN-SS with an uncertainty estimator which models second-order covariance structure allows us to retrieve diverse molecules early during optimization. At the cost of sample efficiency, this improves our surrogate reward estimates, ultimately leading GNN-SS to obtain higher reward molecules.

Our random sub-sampling routine in GNN-SS is partly justified by an experiment in which the acquisition function is evaluated for every member of the domain  $\mathcal{G}$ . As shown in Figure 3, for certain objectives GNN-SS-UCB achieves improved sample efficiency compared to an expensive variant without sub-sampling.

MOLECULAR OBJECTIVE	MOLPAL	GNN-SS GREEDY	GNN-SS UCB (DIAG)	GNN-SS UCB (SUB)	GNN-SS UCB (FJL)
TOP-1 SUM	12.84	13.99	14.19	14.27	<b>14.38</b>
$n/29$ RANK	17	11	10	8	<b>7</b>

Table 2: Summary PMO benchmark Top-1 Reward with different uncertainty estimators. Top-1 Sum aggregates all 23 objectives. GNN-SS-UCB with an FJL approximation of  $\hat{\sigma}_t$  outperforms other variants and screening methods, obtaining molecules with higher reward. Full table in Appendix C.

## 5 Conclusion

We presented GNN-SS, a Graph Neural Network BO algorithm that draws on sub-sampling and random feature projections, to present a scalable solution to the problem of screening large virtual libraries. We demonstrated that GNN-SS achieves competitive performance on the PMO benchmark for sample-efficient molecular optimization. Our results show how GNN-powered

Bayesian Optimization may be deployed with datasets of  $10^6$  molecules. However it remains an open problem to scale BO to virtual databases with billions of molecules, or efficiently go beyond a library and perform *de novo* design using generative models.

## References

- Nir Ailon and Bernard Chazelle. The fast johnson-lindenstrauss transform and approximate nearest neighbors. *SIAM Journal on computing*, 39(1):302–322, 2009.
- Sanjeev Arora, S. Du, Wei Hu, Zhiyuan Li, R. Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *Neural Information Processing Systems*, 2019.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 2002.
- Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 2020.
- Mohsen Bayati, Nima Hamidi, Ramesh Johari, and Khashayar Khosravi. Unreasonable effectiveness of greedy algorithms in multi-armed bandit with many arms. *Advances in Neural Information Processing Systems*, 2020.
- G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 2012.
- Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling*, 59(3): 1096–1108, 2019.
- Yuanqi Du, Tianfan Fu, Jimeng Sun, and Shengchao Liu. Molgensurvey: A systematic survey in machine learning models for molecule design. *arXiv preprint arXiv: Arxiv-2203.14500*, 2022.
- Vijay Prakash Dwivedi, Chaitanya K. Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24 (43):1–48, 2023.
- Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- Casper Benjamin Freksen. An introduction to johnson-lindenstrauss transforms. *arXiv preprint arXiv: 2103.00564*, 2021.
- Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a benchmark for practical molecular optimization. *Advances in Neural Information Processing Systems*, 2022.
- Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 2018.
- David E. Graff, Eugene I. Shakhnovich, and Connor W. Coley. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chem. Sci.*, 2021.
- Botao Hao, Tor Lattimore, and Mengdi Wang. High-dimensional sparse linear bandits. *Advances in Neural Information Processing Systems*, 2020.
- Lior Hirschfeld, Kyle Swanson, Kevin Yang, Regina Barzilay, and Connor W. Coley. Uncertainty quantification using neural networks for molecular property prediction, 2020.
- Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf H. Roohani, J. Leskovec, Connor W. Coley, Cao Xiao, Jimeng Sun, and M. Zitnik. Therapeutics data commons: Machine learning datasets and tasks for drug discovery and development. *Neurips Datasets And Benchmarks*, 2021.



- John J Irwin and Brian K Shoichet. Zinc- a free database of commercially available compounds for virtual screening. *Journal of chemical information and modeling*, 2005.
- William Johnson and Joram Lindenstrauss. Extensions of lipschitz maps into a hilbert space. *Contemporary Mathematics*, 26:189–206, 01 1984.
- Ata Kaban. Improved bounds on the dot product under random projection and random sign projection. *KDD*, 2015.
- Parnian Kassraie, Andreas Krause, and Ilija Bogunovic. Graph neural network bandits. In *Advances in Neural Information Processing Systems*, 2022.
- Samuel Kim, Peter Y Lu, Charlotte Loh, Jamie Smith, Jasper Snoek, and Marin Soljagic. Deep learning for bayesian optimization of scientific problems with high-dimensional structure. *Transactions of Machine Learning Research*, 2021.
- Samuel Kim, Peter Y. Lu, Charlotte Loh, Jamie Smith, Jasper Snoek, and Marin Soljačić. Deep learning for bayesian optimization of scientific problems with high-dimensional structure, 2022.
- Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos, Jeff Schneider, and Eric Xing. Chemo: Bayesian optimization of small organic molecules with synthesizable recommendations. In *International Conference on Artificial Intelligence and Statistics*, 2020.
- Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, Amin Karbasi, Jan Vondrák, and Andreas Krause. Lazier than lazy greedy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015.
- Christos A Nicolaou, Ian A Watson, Hong Hu, and Jibo Wang. The proximal lilly collection: Mapping, exploring and exploiting feasible chemical space. *Journal of chemical information and modeling*, 2016.
- Ava P. Soleimany, Alexander Amini, Samuel Goldman, Daniela Rus, Sangeeta N. Bhatia, and Connor W. Coley. Evidential deep learning for guided molecular property prediction and discovery. *ACS Central Science*, 2021.
- Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *International Conference on Machine Learning*, 2010.
- T.T. Tanimoto. *An Elementary Mathematical Theory of Classification and Prediction*. International Business Machines Corporation, 1958. URL <https://books.google.ch/books?id=yp34HAAACAAJ>.
- Kilian Q. Weinberger, A. Dasgupta, Josh Attenberg, J. Langford, and Alex Smola. Feature hashing for large scale multitask learning. *International Conference on Machine Learning*, 2009.
- Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. {MARS}: Markov molecular sampling for multi-objective drug discovery. In *International Conference on Learning Representations*, 2021.
- Greg Yang. Tensor programs ii: Neural tangent kernel for any architecture. *arXiv preprint*, 2020.
- Yun Yang, Mert Pilanci, and Martin J. Wainwright. Randomized sketches for kernels: Fast and optimal non-parametric regression. *arXiv preprint arXiv: 1501.06195*, 2015.
- Weitong Zhang, Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural thompson sampling. *International Conference On Learning Representations*, 2020.
- Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 2020.

## A Closeness of Random Feature Approximations

As detailed in [Kassraie et al. \[2022\]](#), since a large finite-width neural network behaves approximately as a linear function in feature space, it can be used in place of the posterior mean of a GP, with the NTK associated to the same architecture as its covariance. Training of the former yields an approximate solution to kernel ridge-regression with the NTK, while the latter yields a solution that is exact. Likewise, one may approximate the posterior covariance with explicit features  $\mathbf{g}(G) = \nabla_{\theta} f(G; \theta)$ . We now restate a central lemma that formalizes the error's dependence on hidden-width  $m$  for given depth  $L$ .

**Lemma 1** (Finite-width Approximation to the NTK [[Arora et al., 2019](#)](Thm. 3.1)). *Let node features  $\mathbf{x}_G, \mathbf{x}'_G \in \mathbb{S}^{d-1}$  be restricted to the hyper-sphere, and consider an appropriately initialized neural network of depth  $L$  with ReLU layer-activations of  $\max(0, \mathbf{x})$ . Fix error  $\epsilon_A > 0$ ,  $\delta_A \in (0, 1)$ . Then for minimum hidden width  $m_{\min} \in \Omega(L^6 \epsilon_A^{-4} \log(L/\delta))$ , the following bounds the error of the empirical NTK at initialization, and holds with probability at least  $1 - \delta_A$*

$$|\langle \mathbf{g}(G), \mathbf{g}(G') \rangle - k(G, G')| \leq (L + 1)\epsilon_A.$$

When  $m$  is large, this allows us to employ the following as a provably calibrated uncertainty estimate,

$$\hat{\sigma}_t^2(G) = \mathbf{g}^\top(G) \left( \lambda I + \sum_{i=1}^{t-1} \mathbf{g}(G_i) \mathbf{g}^\top(G_i) \right)^{-1} \mathbf{g}(G) = \mathbf{g}^\top(G) (\lambda I + \mathbf{J}_t^\top \mathbf{J}_t)^{-1} \mathbf{g}(G)$$

Here,  $\mathbf{J}_t = [\nabla_{\theta} f(G_i; \theta)]_{i < t}$  is the Jacobian with  $t - 1$  rows, and each row contains the gradient vector  $\mathbf{g}(G_i)$ . This admits rank-one updates via the matrix inversion lemma. However, the benefit is not immediately clear as the outer-product  $\mathbf{g}(G) \mathbf{g}^\top(G)$  is prohibitive for even shallow, small-width networks, requiring  $\mathcal{O}(m^2)$  memory and  $\mathcal{O}(m^6)$  effort to invert. Practical implementations of neural bandit algorithms that use this construction have therefore addressed expensive inversions by ignoring higher-order covariance structure and taking a diagonal approximation to  $\mathbf{g}(G) \mathbf{g}^\top(G)$ , leading to updates linear in the number of parameters  $\mathcal{O}(p)$ . Instead, we propose an alternative approximation that yields valid confidence sets for optimistic bandit policies.

We now outline Johnson-Lindenstrauss (JL) type projections [[Johnson and Lindenstrauss, 1984](#)]. These projections are *data-independent* and are generated under a random linear map  $S : \mathbb{R}^p \rightarrow \mathbb{R}^d$ , where  $d \ll p$ . Importantly, the distortion of these projections can be controlled for some choice of rank  $d$ , which must be greater than  $d_{\min}$  for some distortion error  $\epsilon_D$  and confidence  $\delta_D$ . The following celebrated lemma bounds the distortion using such a projection.

**Lemma 2** (JLL [[Johnson and Lindenstrauss, 1984](#)]). *Set  $\epsilon_D \in (0, 1/2)$  and let  $d_{\min} \geq 4\epsilon_D^{-2} \log(1/\delta)$ . There exists a matrix  $S \in \mathbb{R}^{d \times p}$  that w.h.p. preserves Euclidean norms, i.e. simultaneously for all  $\mathbf{x} \in \mathbb{R}^p$ ,  $S$  satisfies the following with probability greater than  $1 - \delta_D$*

$$\|S\mathbf{x}\|_2 \in [(1 - \epsilon_D)\|\mathbf{x}\|_2, (1 + \epsilon_D)\|\mathbf{x}\|_2]$$

A corollary of Lemma 2 can also be obtained for squared-norms via the parallelogram law, where the original lemma is applied to norms found in the following,

$$\langle S\mathbf{x}, S\mathbf{x}' \rangle = \frac{1}{4}(\|S(\mathbf{x} - \mathbf{x}')\| + \|S(\mathbf{x} + \mathbf{x}')\|)$$

For any given  $d > d_{\min}$ , this has been shown to preserve the original bound up to the same constant factor [[Kaban, 2015](#)](Thm 2.1). And in the setting where we directly use the network gradients  $\mathbf{g}(G) \in \mathbb{R}^p$  to approximate the NTK, we can make use of this to guarantee that for large enough  $d$ , the following bound on the distortion error  $\epsilon_D$  also holds with high probability,

$$|\langle S\mathbf{g}(G), S\mathbf{g}(G') \rangle - \langle \mathbf{g}(G), \mathbf{g}(G') \rangle| \leq \epsilon_D$$

Now, by the union of events in Lemma 1 and Lemma 2, the following inequality will control the combined distortion  $\epsilon_D$  and approximation error  $\epsilon_A$  of the empirical NTK when obtained by a rank- $d$  projection of the  $p$ -dimensional random feature map  $\mathbf{g}(G)$ .

**Proposition 3** (Approximation Error of Randomly Projected Empirical NTK). *Set  $\delta \in (0, 1)$  and fix a choice of projection dimension  $d \in [\epsilon_A^{-2} \log(1/\delta), p]$ . Denote  $\alpha = 2\delta - \exp(-d\epsilon_D^2/8)$ . Then*

for minimum width  $m_{\min} \in \Omega(L^6 \epsilon_A^{-4} \log(L/\alpha))$ , the following bound holds with probability at least  $1 - \delta$ ,

$$|\langle \tilde{\mathbf{g}}(G), \tilde{\mathbf{g}}(G') \rangle - k(G, G')| \leq \epsilon$$

where  $\epsilon := (L + 1)\epsilon_A + \epsilon_D$ .

We have used  $\tilde{\mathbf{g}}(G) = S\mathbf{g}(G)$  to denote the random projection of  $\mathbf{g}(G)$ . Let  $\exp(-d\epsilon_D^2/8)$  and  $L \exp(-m\epsilon_A^4/L^6)$  be the confidence levels introduced in Lemma 2 and Lemma 1. We require

$$2\delta \leq L \exp(-m\epsilon_A^4/L^6) + 2 \exp(-d\epsilon_D^2/8)$$

which is satisfied by choosing a minimum width  $m_{\min} \geq L^6 \epsilon_A^{-4} \log(L/(2\delta - \exp(-d\epsilon_D^2/8)))$  for some fixed projection dimension  $d$ .

## B Experiment Details

We use the ZINC dataset provided in pyTDC, objectives are evaluated using oracles from the same library Huang et al. [2021]. For a fair comparison to MolPAL, we use similar parameters  $m = 384$ ,  $L = 3$  graph convolutional layers and no positional encodings, corresponding to GNN-SS (Diag, NoLap) in Table 3. The regularization coefficient is chosen via a grid-search  $\lambda \in \{0.1, 0.01, 0.001, 0.0001, 0.00001\}$ . We use  $\lambda = 0.0001$ . We select the best performing exploration coefficient  $\beta$  via a grid-search over  $\{0.001, 0.01, 0.1, 1.0\}$ . The size of randomly sub-sampled set  $|\mathcal{G}_t|$  is set as  $K = 4000$ . Following Gao et al. [2022], we use  $T_0 = 500$ . For remaining experiments that aim to compare different numerical approximations, we use  $k = 4$  absolute graph Laplacian positional encodings [Dwivedi et al., 2023], and choose hidden dimension  $m = 512$ .

Graph Laplacian positional encodings are eigenvectors of the symmetrically normalized graph Laplacian  $\Delta = I - D^{-1/2}AD^{-1/2}$ .  $A$  and  $D$  are the adjacency and degree matrices, respectively. From the matrix of eigenvectors  $U \in \mathbb{R}^{N \times N}$ , we take  $k < N$  eigenvectors, which are concatenated directly to node features as  $X \leftarrow \text{concat}(X, U_k)$ , where  $U_k$  is an  $N \times k$  matrix. To address the sign ambiguity of the eigenvectors we take their absolute value [Dwivedi et al., 2023].

The network parameters are initialized with iid. unit Gaussian entries, and following the procedure of Graff et al. [2021], pre-trained offline for each reward on 500 molecules. We limit training to 250 epochs, with batchsize 250. Early-stopping is used with a patience of 25 iterations without improvement on the validation set. We use MSELoss with  $\text{lr} = 1\text{e-}5$  and a linear decay towards  $1\text{e-}6$ .  $\phi = \text{ReLU}$  non-linearities are applied between each hidden layer. We modify the pytorch\_geometric library Fey and Lenssen [2019] to make use of the neural tangent parameterization for all GCNConv layers, where  $c_\phi = \sqrt{2}$  and  $m_l$  is the width of the  $l$ -layer weights. Using the neural-tangent parameterization ensures a limiting NTK that converges. Each convolutional layer is defined as

$$f_{\text{GCN}}^{(l)}(X, A) = \sqrt{\frac{2}{m}} \text{ReLU} \left( \Delta X^{(l-1)} W^{(l)} \right)$$

The graph representation is obtained as  $\bar{X} = \text{cat}(\bar{X}_{\text{mean}}, \bar{X}_{\text{sum}})$  before passing to a final linear layer, i.e. the concatenation of the mean-readout and sum-readout. We note that for the random seed used in our experiments, the subset  $H_{\text{off}} \subset \mathcal{G}$  for the VALSARTAN SMARTS objective only contains molecules with zero reward, and thus the model is not expected to generalize to unseen molecules with non-zero reward.

---

**Algorithm 2** TrainGNN

---

**Input:** history  $H_t$ , initial parameters  $\theta$ , training epochs  $E$ , batchsize  $b$ , learning rate  $\eta$ , regularization  $\lambda$

---

**Define:** MSE Loss  $\mathcal{L}(\theta_i)$  from Section 3

**Initialize:**  $\theta^{(0)} \leftarrow \theta$ ,  $f_{\text{GNN}}(\cdot; \theta^{(0)})$ , optimizer Adam( $\mathcal{L}$ )

---

**for** epoch  $i = 1, \dots, E$  **do**

    Batch  $\mathcal{B} = \{(y(G_i), G_i)\}^b \sim H_t$

    Update parameters,  $\theta^{(i+1)} \leftarrow \text{Adam}(\mathcal{B}, f_{\text{GNN}}(\cdot; \theta^{(i)}))$

**end for**

---

## C Additional Results

Molecular Objective	MolPAL	GNN-SS (Diag, NoLap)	GNN-SS (Diag)	GNN-SS (SJL)	GNN-SS (FJL)
albuterol sim.	-*	0.640 $\pm$ 0.017	0.639 $\pm$ 0.012	0.630 $\pm$ 0.016	0.642 $\pm$ 0.020
amlodipine mpo	0.621 $\pm$ 0.010	0.543 $\pm$ 0.007	0.581 $\pm$ 0.010	0.597 $\pm$ 0.013	0.660 $\pm$ 0.010
celecoxib redisc.	0.496 $\pm$ 0.002	0.436 $\pm$ 0.008	0.438 $\pm$ 0.008	0.440 $\pm$ 0.008	0.434 $\pm$ 0.008
deco hop	0.804 $\pm$ 0.019	0.874 $\pm$ 0.016	0.871 $\pm$ 0.004	0.871 $\pm$ 0.003	0.872 $\pm$ 0.003
drd2	0.902 $\pm$ 0.007	0.972 $\pm$ 0.011	0.951 $\pm$ 0.019	0.943 $\pm$ 0.024	0.918 $\pm$ 0.010
fexofenadine mpo	0.704 $\pm$ 0.001	0.744 $\pm$ 0.008	0.740 $\pm$ 0.009	0.743 $\pm$ 0.008	0.748 $\pm$ 0.006
gsk3b	0.776 $\pm$ 0.002	0.862 $\pm$ 0.086	0.763 $\pm$ 0.042	0.767 $\pm$ 0.039	0.742 $\pm$ 0.045
isomers c7h8n2o-	0.832 $\pm$ 0.005	0.979 $\pm$ 0.016	0.977 $\pm$ 0.019	0.968 $\pm$ 0.036	0.949 $\pm$ 0.041
isomers c9h10n2-	0.361 $\pm$ 0.009	0.785 $\pm$ 0.023	0.782 $\pm$ 0.027	0.780 $\pm$ 0.022	0.786 $\pm$ 0.021
jnk3	0.457 $\pm$ 0.024	0.552 $\pm$ 0.051	0.479 $\pm$ 0.065	0.500 $\pm$ 0.046	0.437 $\pm$ 0.047
median1	0.301 $\pm$ 0.000	0.303 $\pm$ 0.012	0.311 $\pm$ 0.009	0.313 $\pm$ 0.010	0.312 $\pm$ 0.010
median2	0.266 $\pm$ 0.000	0.275 $\pm$ 0.012	0.275 $\pm$ 0.010	0.273 $\pm$ 0.008	0.264 $\pm$ 0.009
mestranol sim.	0.708 $\pm$ 0.006	0.864 $\pm$ 0.029	0.785 $\pm$ 0.083	0.770 $\pm$ 0.085	0.765 $\pm$ 0.078
osimertinib mpo	0.803 $\pm$ 0.001	0.803 $\pm$ 0.004	0.804 $\pm$ 0.002	0.804 $\pm$ 0.002	0.804 $\pm$ 0.002
perindopril mpo	0.495 $\pm$ 0.003	0.463 $\pm$ 0.004	0.538 $\pm$ 0.012	0.532 $\pm$ 0.011	0.535 $\pm$ 0.010
qed	0.942 $\pm$ 0.000	0.947 $\pm$ 0.000	0.947 $\pm$ 0.000	0.947 $\pm$ 0.000	0.947 $\pm$ 0.000
ranolazine mpo	0.515 $\pm$ 0.007	0.556 $\pm$ 0.011	0.571 $\pm$ 0.006	0.567 $\pm$ 0.011	0.573 $\pm$ 0.005
scaffold hop	0.518 $\pm$ 0.001	0.524 $\pm$ 0.002	0.521 $\pm$ 0.003	0.519 $\pm$ 0.003	0.518 $\pm$ 0.004
sitagliptin mpo	0.100 $\pm$ 0.013	0.478 $\pm$ 0.001	0.393 $\pm$ 0.041	0.397 $\pm$ 0.030	0.367 $\pm$ 0.001
thiothixene red.	0.356 $\pm$ 0.000	0.391 $\pm$ 0.012	0.380 $\pm$ 0.015	0.381 $\pm$ 0.011	0.375 $\pm$ 0.013
troglitazone red.	0.290 $\pm$ 0.000	0.326 $\pm$ 0.000	0.315 $\pm$ 0.022	0.326 $\pm$ 0.023	0.328 $\pm$ 0.021
valsartan smarts	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000
zaleplon mpo	0.262 $\pm$ 0.004	0.514 $\pm$ 0.016	0.521 $\pm$ 0.017	0.534 $\pm$ 0.020	0.513 $\pm$ 0.013
Top-1 AUC sum	12.21	<b>13.83</b>	13.58	13.65	13.49
$n/29$ rank	10	<b>4</b>	7	5	8

Table 3: **Full PMO benchmark results: Top-1 AUC.** The maximum reward for the Albuterol Similarity in the benchmark dataset is 0.666, but reported value in Gao et al. [2022] exceeds it. (NoLap) denotes no Laplacian positional encodings.

Molecular Objective	MolPAL	GNN-SS (Diag, NoLap)	GNN-SS (Diag)	GNN-SS (Sub)	GNN-SS (FJL)
albuterol sim.	-*	0.653 ± 0.017	0.667 ± 0.007	0.660 ± 0.014	0.653 ± 0.017
amlodipine mpo	0.651 ± 0.043	0.556 ± 0.012	0.597 ± 0.012	0.604 ± 0.006	0.678 ± 0.011
celecoxib redisc.	0.511 ± 0.041	0.450 ± 0.008	0.450 ± 0.008	0.451 ± 0.007	0.458 ± 0.000
deco hop	0.860 ± 0.102	0.878 ± 0.001	0.876 ± 0.002	0.877 ± 0.001	0.877 ± 0.002
drd2	0.964 ± 0.709	0.983 ± 0.010	0.976 ± 0.034	0.980 ± 0.005	0.981 ± 0.001
fexofenadine mpo	0.709 ± 0.006	0.756 ± 0.005	0.756 ± 0.002	0.756 ± 0.002	0.756 ± 0.002
gsk3b	0.820 ± 0.128	0.904 ± 0.089	0.847 ± 0.024	0.854 ± 0.022	0.860 ± 0.014
isomers c7h8n2o-	0.882 ± 0.163	1.000 ± 0.006	1.000 ± 0.006	1.000 ± 0.019	1.000 ± 0.017
isomers c9h10n2-	0.391 ± 0.091	0.821 ± 0.013	0.822 ± 0.033	0.822 ± 0.033	0.869 ± 0.000
jnk3	0.608 ± 0.117	0.620 ± 0.030	0.551 ± 0.079	0.566 ± 0.067	0.620 ± 0.005
median1	0.309 ± 0.028	0.315 ± 0.009	0.321 ± 0.006	0.315 ± 0.007	0.321 ± 0.006
median2	0.273 ± 0.021	0.291 ± 0.004	0.291 ± 0.003	0.284 ± 0.009	0.291 ± 0.004
mestranol sim.	0.733 ± 0.081	0.886 ± 0.041	0.845 ± 0.082	0.886 ± 0.012	0.824 ± 0.050
osimertinib mpo	0.816 ± 0.020	0.807 ± 0.004	0.806 ± 0.001	0.806 ± 0.002	0.806 ± 0.002
perindopril mpo	0.504 ± 0.020	0.483 ± 0.005	0.560 ± 0.002	0.555 ± 0.009	0.551 ± 0.006
qed	0.948 ± 0.002	0.947 ± 0.000	0.947 ± 0.000	0.947 ± 0.000	0.948 ± 0.000
ranolazine mpo	0.556 ± 0.064	0.577 ± 0.013	0.586 ± 0.002	0.581 ± 0.009	0.586 ± 0.002
scaffold hop	0.525 ± 0.016	0.526 ± 0.001	0.526 ± 0.001	0.526 ± 0.001	0.526 ± 0.001
sitagliptin mpo	0.117 ± 0.030	0.478 ± 0.000	0.463 ± 0.025	0.474 ± 0.008	0.478 ± 0.006
thiothixene red.	0.361 ± 0.016	0.408 ± 0.005	0.401 ± 0.014	0.404 ± 0.009	0.401 ± 0.013
trogliptazone red.	0.296 ± 0.013	0.370 ± 0.008	0.346 ± 0.030	0.357 ± 0.026	0.370 ± 0.014
valsartan smarts	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
zaleplon mpo	0.286 ± 0.064	0.533 ± 0.014	0.556 ± 0.014	0.561 ± 0.012	0.528 ± 0.007
Top-1 sum	12.84	14.24	14.19	14.27	14.38
$n/29$ rank	17	10	11	8	7

Table 4: **Full PMO benchmark results: Top-1 Reward.** Standard errors are computed with 5 seeded replicates. (NoLap) denotes no Laplacian positional encodings.

Figure 4: GNN-SS Top-10 Mean Reward for all 23 objectives on the PMO benchmark.

