
Graph Neural Network Bandits

Parnian Kassarai
ETH Zurich
pkassarai@ethz.ch

Andreas Krause
ETH Zurich
krausea@ethz.ch

Ilija Bogunovic
University College London
i.bogunovic@ucl.ac.uk

Abstract

We consider the bandit optimization problem with the reward function defined over graph-structured data. This problem has important applications in molecule design and drug discovery, where the reward is naturally invariant to graph permutations. The key challenges in this setting are scaling to large domains, and to graphs with many nodes. We resolve these challenges by embedding the permutation invariance into our model. In particular, we show that graph neural networks (GNNs) can be used to estimate the reward function, assuming it resides in the Reproducing Kernel Hilbert Space of a permutation-invariant additive kernel. By establishing a novel connection between such kernels and the graph neural tangent kernel (GNTK), we introduce the first GNN confidence bound and use it to design a phased-elimination algorithm with sublinear regret. Our regret bound depends on the GNTK’s maximum information gain, which we also provide a bound for. While the reward function depends on all N node features, our guarantees are *independent* of the number of graph nodes N . Empirically, our approach exhibits competitive performance and scales well on graph-structured domains.

1 Introduction

Contemporary bandit optimization approaches consider problems on large or continuous domains and have successfully been applied to a significant number of machine learning and real-world applications, e.g., in mobile health, environmental monitoring, economics, and hyperparameter tuning, to name a few. The main idea behind them is to exploit the correlations between the rewards of “similar” actions. This in turn, has resulted in increasingly rich models of reward functions (e.g., in linear and kernelized bandits [37, 12]), including several recent attempts to harness deep neural networks for bandit tasks (see, e.g, [45, 29]). A vast majority of previous works only focus on standard input domains and obtaining theoretical regret bounds.

Learning on graph-structured data, such as molecules or biological graph representations, requires designing sequential methods that can effectively exploit the structure of graphs. Consequently, graph neural networks (GNNs) have received attention as a rapidly expanding class of machine learning models. They deem remarkably well-suited for prediction tasks in applications such as designing novel materials [20], drug discovery [22], structure-based protein function prediction [16], etc. This rises the question of how to bridge the gap, and design bandit optimization algorithms on graph-structured data that can exploit the power of graph neural networks to approximate a graph reward.

In this paper, we consider bandit optimization over graphs and propose to employ graph neural networks with one convolutional layer to estimate the unknown reward. To scale to *large* graph domains (both in the number of graphs and number of nodes), we propose practical structural assumptions to model the reward function. In particular, we propose to use *permutation invariant additive* kernels. We show a novel connection between such kernels and the *graph neural tangent kernel* (GNTK) that we define in Section 3. Our main result are *GNN confidence bounds* that can be readily used in sequential decision-making algorithms to achieve sublinear regret bounds (see Section 4.3).

Related Work. Our work extends the rich toolbox of methods for kernelized bandits and Bayesian optimization (BO) that work under the norm bounded Reproducing Kernel Hilbert Space assumption [37, 14, 42, 12]). The majority of these methods are designed for general Euclidean domains and rely on kernelized confidence sets to select which action to query next. The exception is [43], that consider the spectral setting in which the reward function is a linear combination of the eigenvectors of the graph Laplacian and the bandit problem is defined over nodes of a single graph. In contrast, our focus is on optimizing over graph domains (i.e., set of graphs), and constructing confidence sets that can quantify the uncertainty of graph neural networks estimates.

This work contributes to the literature on *neural bandits*, in which a fully-connected [45, 44, 19], or a single hidden layer convolutional network [25] is used to estimate the reward function. These works provide sublinear cumulative regret bounds in their respective settings, however, when applied directly to graph features (as we demonstrate in Section 4.3), these approaches do not scale well with the number of graph nodes.

Due to its important applications in molecule design, sequential optimization on graphs has recently received considerable attention. For example, in [28], the authors propose a kernel to capture similarities between graphs, and at every step, select the next graph through a kernelized random walk. Other works (e.g., [17, 18, 23, 38]) encode graph representations to the (continuous) latent space of a variational autoencoder and perform BO in the latent space. While practically relevant for discovering novel molecules with optimized properties, these approaches lack theoretical guarantees and deem computationally demanding.

A primary focus in our work is on embedding the natural structure of the data, i.e., permutation invariance, into the reward model. This is inspired by the works of [6, 32] that consider invariances in kernel-based supervised learning. Consequently, the graph neural tangent kernel plays an integral role in our theoretical analysis. Du et al. [15] provide a recursive expression for the tangent kernel of a GNN, without showing that the obtained expression is the limiting tangent kernel as defined in Jacot et al. [21] (i.e., as in Eq. (3)). In contrast, we analyze the learning dynamics of the GNN and properties of the GNTK by exploiting the connection between the structure of a graph neural network and that of a neural network (in Section 3). We recover that the graph neural tangent kernel also encodes additivity. Additive models for bandit optimization have been previously studied in [24] and [35], however, these works only focus on Euclidean domains and standard base kernels.

Finally, we build upon the recent literature on elimination-based algorithms that make use of maximum variance reduction sampling [13, 8, 7, 9, 40, 30]. One of our proposed algorithms, GNN-PE, employs a phased elimination strategy together with our GNN confidence sets.

Main Contributions. We introduce a bandit problem over graphs and propose to capture prior knowledge by modeling the unknown reward function using a permutation invariant additive kernel. We establish a key connection between such kernel assumptions and the graph neural tangent kernel (Proposition 3.2). By exploiting this connection, we provide novel statistical confidence bounds for the graph neural network estimator (Theorem 4.2). We further prove that a phased elimination algorithm that uses our GNN-confidence bounds (GNN-PE) achieves sublinear regret (Theorem 4.3). Importantly, our regret bound scales favorably with the number of graphs and is *independent* of the number of graph nodes (see Table 1). Finally, we empirically demonstrate that our algorithm consistently outperforms baselines across a range of problem instances.

2 Problem Statement

We consider a bandit problem where the learner aims to optimize an *unknown* reward function via sequential interactions with a stochastic environment. At every time step $t \in \{1, \dots, T\}$, the learner selects a graph G_t from a graph domain \mathcal{G} and observes a noisy reward $y_t = f^*(G_t) + \epsilon_t$, where $f^* : \mathcal{G} \rightarrow \mathbb{R}$ is the reward function and ϵ_t is i.i.d. zero-mean sub-Gaussian noise with known variance proxy σ^2 . Over a time horizon T , the learner seeks a small *cumulative* regret $R_T = \sum_{t=1}^T f^*(G^*) - f^*(G_t)$, where $G^* \in \arg \max_{G \in \mathcal{G}} f^*(G)$. The aim is to attain regret that is *sublinear* in T , meaning that $R_T/T \rightarrow 0$ as $T \rightarrow \infty$, which implies convergence to the optimal graph. As an example application, consider drug or material design, where molecules may be represented with graph structures (e.g., from SMILES representations [1]) and the reward $f^*(G)$ can correspond to an unknown molecular property of interest, e.g., atomization energy. Evaluating such properties typically requires running

costly simulations or experiments with noisy outcomes. To identify the most promising candidate, e.g., the molecule with the highest atomization energy, molecules are sequentially recommended for testing and the goal is to find the optimal molecule with the least number of evaluations.

Graph Domain. We assume that the domain \mathcal{G} is a finite set of undirected graphs with N nodes.¹ Without exploiting structure, standard bandit algorithms (e.g., [3]) cannot generalize across graphs, and their regret linearly depends on $|\mathcal{G}|$. To capture the structure, we consider reward functions depending on features associated with the graph nodes. Similar to Du et al. [15], we associate each node $j \in [N]$ with a feature vector $\mathbf{h}_{G,j} \in \mathbb{R}^d$, for every graph $G \in \mathcal{G}$. We use $\mathbf{h}_G = (\mathbf{h}_{G,j})_{j=1}^N \in \mathbb{R}^{Nd}$ to denote the concatenated vector of all node features, and $\mathcal{N}(j)$ as the neighborhood of node j , including itself. We define the aggregated node feature $\bar{\mathbf{h}}_{G,j} = \sum_{i \in \mathcal{N}(j)} \mathbf{h}_{G,i} / \|\sum_{i \in \mathcal{N}(j)} \mathbf{h}_{G,i}\|_2$ as the normalized sum of the neighboring nodes’ features. Similarly, $\bar{\mathbf{h}}_G \in \mathbb{R}^{Nd}$ denotes the aggregated features, stacked across all nodes. Lastly, we let P_N be the group of all permutations of length N , and use $c \cdot G$ to denote a permuted graph, where a permutation $c \in P_N$ is a bijective mapping from $\{1, \dots, N\}$ onto itself. Permuting the nodes of a graph $c \cdot G$ produces a permuted feature vector $\mathbf{h}_{c \cdot G} := (\mathbf{h}_{G,c(j)})_{j=1}^N$, and the same holds for the aggregated features $\bar{\mathbf{h}}_{c \cdot G}$.

Reward Model. Practical graph optimization problems, such as drug discovery and materials optimization often do not depend on how the graphs’ nodes in the dataset are ordered. We incorporate this structural prior into modeling the reward function, and consider functions that are *invariant to node permutations*. We assume that f^* depends on the graph only through the aggregated node features and gives the same reward for all permutations of a graph, i.e., $f^*(c \cdot G) = f^*(G)$, for any $G \in \mathcal{G}$ and $c \in P_N$. To guarantee such an invariance, we assume that the reward belongs to the reproducing kernel Hilbert space (RKHS) corresponding to a permutation invariant kernel

$$\bar{k}(G, G') = \frac{1}{|P_N|^2} \sum_{c, c' \in P_N} k(\bar{\mathbf{h}}_{c \cdot G}, \bar{\mathbf{h}}_{c' \cdot G'}),$$

where k can be any kernel defined on graph representations $\bar{\mathbf{h}}_G$. This assumption further restricts the hypothesis space to permutation invariant functions defined on Nd -dimensional vector representations of graphs. This is due to the reproducing property of the RKHS which allows us to write $f(G) = \langle f, \bar{k}(G, \cdot) \rangle = \langle f, \bar{k}(c \cdot G, \cdot) \rangle = f(c \cdot G)$. To make progress when optimizing over graphs with a *large* number of nodes N , we assume that k decomposes additively over node features, i.e.,

$$k(\bar{\mathbf{h}}_G, \bar{\mathbf{h}}_{G'}) = \frac{1}{N} \sum_{j=1}^N k^{(j)}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j}).$$

Thereby, we obtain an *additive* graph kernel that is *invariant* to node permutations:

$$\bar{k}(G, G') = \frac{1}{|P_N|^2} \sum_{c, c' \in P_N} \frac{1}{N} \sum_{j=1}^N k^{(j)}(\bar{\mathbf{h}}_{G,c(j)}, \bar{\mathbf{h}}_{G',c'(j)}). \quad (1)$$

For an arbitrary choice of $k^{(j)}$, calculating \bar{k} requires a costly sum over $(N!)^2$ operands, since $|P_N| = N!$. In Section 3, we select a base kernel for which the sum can be reduced to N^2 terms. We are now in a position to state our main assumption. We assume that f^* belongs to the RKHS of \bar{k} and has a B -bounded RKHS norm. The norm-bounded RKHS regularity assumption is typical in the kernelized and neural bandits literature [37, 12, 45, 25]. Note that Eq. (1) only puts a structural prior on the kernel function, i.e., it describes the generic form of an additive permutation invariant graph kernel. Specifying the base kernels $k^{(j)}$ determines the representation power of \bar{k} . The smoother the base kernels are, the less complex the RKHS of \bar{k} will be. In Section 3, we set the base kernels $k^{(j)}$ such that \bar{k} becomes the expressive graph neural tangent kernel.

3 Graph Neural Networks

Graph neural networks are effective models for learning complex functions defined on graphs. As in Du et al. [15], we consider graph networks that have a *single* graph convolutional layer and L

¹This assumption is for ease of exposition. Graphs with fewer than N nodes can be treated by adding auxiliary nodes with no features that are disconnected from the rest of the graph.

fully-connected ReLU layers of equal width m . Such a network $f_{\text{GNN}}(G; \boldsymbol{\theta}) : \mathcal{G} \rightarrow \mathbb{R}$ may be recursively defined as follows:

$$\begin{aligned} f^{(1)}(\bar{\mathbf{h}}_{G,j}) &= \mathbf{W}^{(1)} \bar{\mathbf{h}}_{G,j}, \\ f^{(l)}(\bar{\mathbf{h}}_{G,j}) &= \sqrt{\frac{2}{m}} \mathbf{W}^{(l)} \sigma_{\text{relu}}(f^{(l-1)}(\bar{\mathbf{h}}_{G,j})) \in \mathbb{R}^m, 1 < l \leq L \\ f_{\text{GNN}}(G; \boldsymbol{\theta}) &= \frac{1}{N} \sum_{j=1}^N \sqrt{2} \mathbf{W}^{(L+1)} \sigma_{\text{relu}}(f^{(L)}(\bar{\mathbf{h}}_{G,j})), \end{aligned} \quad (2)$$

where $\boldsymbol{\theta} := (\mathbf{W}^{(i)})_{i \leq L+1}$ is initialized randomly with standard normal i.i.d. entries, and $\sigma_{\text{relu}}(\mathbf{x}) := \max(\mathbf{0}, \mathbf{x})$. The network operates on aggregated node features $\bar{\mathbf{h}}_{G,j}$ as typical in Graph Convolutional Networks [27]. For convenience, we assume that at initialization $f_{\text{GNN}}(G; \boldsymbol{\theta}^0) = 0$, for all $G \in \mathcal{G}$, similar to [25, 45]. This assumption can be fulfilled without loss of generality, with a similar treatment as in [25, Appendix B.2].

Embedded Invariances. In this work, we use graph neural networks to estimate the unknown reward function f^* . This choice is motivated by the expressiveness of the GNN, the fact that it scales well with graph size, and particularly due to the invariances embedded in its structure. We observe that the graph neural network f_{GNN} is invariant to node permutations, i.e., for all $G \in \mathcal{G}$ and $c \in P_N$,

$$f_{\text{GNN}}(G; \boldsymbol{\theta}) = f_{\text{GNN}}(c \cdot G; \boldsymbol{\theta}).$$

The key step to show this property is proving that f_{GNN} can be expressed as an additive model of L -layer fully-connected ReLU networks,

$$f_{\text{GNN}}(G; \boldsymbol{\theta}) = \frac{1}{N} \sum_{j=1}^N f_{\text{NN}}(\bar{\mathbf{h}}_{G,j}; \boldsymbol{\theta}),$$

where f_{NN} has a similar recursive definition as f_{GNN} (see Equation A.1). The above properties are formalized in Lemma A.1 and Lemma A.2.

Lazy (NTK) Regime. We initialize and train f_{GNN} in the well-known lazy regime [11]. In this initialization regime, when the width m is large, training with gradient descent using a small learning rate causes little change in the network’s parameters. Let $\mathbf{g}_{\text{GNN}}(G, \boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} f_{\text{GNN}}(G, \boldsymbol{\theta})$ denote the gradient of the network. It can be shown that during training, for all $G \in \mathcal{G}$, the network remains close to $f_{\text{GNN}}(G, \boldsymbol{\theta}^0) + \mathbf{g}_{\text{GNN}}^T(G, \boldsymbol{\theta}^0)(\boldsymbol{\theta} - \boldsymbol{\theta}^0)$, that is, its first order approximation around initialization parameters $\boldsymbol{\theta}^0$. Training this linearized model with a squared error loss is equivalent to kernel regression with a *tangent kernel* $\tilde{k}_{\text{GNN}}(G, G') := \mathbf{g}_{\text{GNN}}^T(G; \boldsymbol{\theta}^0) \mathbf{g}_{\text{GNN}}(G'; \boldsymbol{\theta}^0)$. For networks of finite width, this kernel function is random since it depends on the random network parameters at initialization. We show in Proposition 3.1, that in the infinite width limit, the tangent kernel converges to a deterministic kernel. This proposition introduces the Graph Neural Tangent Kernel as the limiting kernel, and links it to the Neural Tangent Kernel ([21], also defined in Appendix A).

Proposition 3.1. *Consider any two graphs G and G' with N nodes and d -dimensional node features. In the infinite width limit, the tangent kernel $\tilde{k}_{\text{GNN}}(G, G')/m$ converges to a deterministic kernel,*

$$k_{\text{GNN}}(G, G') := \lim_{m \rightarrow \infty} \tilde{k}_{\text{GNN}}(G, G')/m. \quad (3)$$

which we refer to as the Graph Neural Tangent Kernel (GNTK). Moreover,

$$k_{\text{GNN}}(G, G') = \frac{1}{N^2} \sum_{j,j'=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j'}) \quad (4)$$

where $k_{\text{NN}} : \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ is the Neural Tangent Kernel.

The proof is given in Appendix A.1. We note that $\bar{\mathbf{h}}_{G,j}$ lies on the d -dimensional sphere, since the aggregated node features are normalized. The NTK is bounded by 1 for any two points on the sphere [5]. Therefore, Proposition 3.1 implies that the GNTK is also bounded, i.e., $k_{\text{GNN}}(G, G') \leq 1$ for any $G, G' \in \mathcal{G}$. This proposition yields a kernel which captures the behaviour of the lazy GNN. While defined on graphs with Nd dimensional representations, the effective input domain of this

kernel is d -dimensional. This advantage directly stems from the additive construction of the GNTK. The next proposition uncovers the embedded structure of the GNTK by showing a novel connection between the GNTK and \bar{k} , the permutation invariant additive kernel from Eq. (1). The proof is presented in Appendix A.1.

Proposition 3.2. *Consider \bar{k} from Eq. (1), where for every $1 \leq j \leq N$ the base kernel $k^{(j)}$ is set to be equal to k_{NN} ,*

$$\bar{k}(G, G') = \frac{1}{|P_N|^2} \sum_{c, c' \in P_N} \frac{1}{N} \sum_{j=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G, c(j)}, \bar{\mathbf{h}}_{G', c'(j)}).$$

Then the permutation invariant additive kernel and the GNTK are identical, i.e., for all $G, G' \in \mathcal{G}$,

$$\bar{k}(G, G') = k_{\text{GNN}}(G, G').$$

This result implies that k_{GNN} inherits the favorable properties of the permutation invariant additive kernel class. Hence, functions residing in \mathcal{H}_{GNN} , the RKHS of k_{GNN} , are additive, invariant to node permutations, and act on G through its aggregated node features. While we use the GNTK as an analytical tool, this kernel can be of independent interest in kernel methods over graph domains. In particular, calculating k_{GNN} requires significantly fewer operations compared to a kernel \bar{k} with an arbitrary choice of $k^{(j)}$, for which calculating $\bar{k}(G, G')$ requires super-exponentially many operations in N (See Eq. (1)). In contrast, due to the decomposition in Eq. (4), calculating k_{GNN} only costs a quadratic number of summations.

4 GNN Bandits

The bandit literature is rich with algorithms that effectively balance exploration and exploitation to achieve sublinear regret. Two components are common in kernelized bandit optimization algorithms. The *maximum information gain*, for characterizing the worst-case complexity of the learning problem [37, 24, 12, 40]; and *confidence sets*, for quantifying the learner’s uncertainty [4, 39, 36, 12, 31]. Our first main result is an upper bound for the maximum information gain when the hypothesis space is \mathcal{H}_{GNN} (Theorem 4.1). We then propose valid confidence sets that utilize GNNs in Theorem 4.2. These theorems may be of independent interest, as they can be used towards bounding the regret for a variety of bandit algorithms on graphs. Lastly, we introduce the GNN-PE algorithm, together with its regret guarantee.

4.1 Information Gain

In bandit tasks, the learner seeks actions that give a large reward while, at the same time, provide information about the unknown reward function. The speed of learning about f^* is commonly quantified via the maximum information gain. Assume that the learner chooses a sequence of actions (G_1, \dots, G_T) and observes noisy rewards, where the noise is i.i.d. and drawn from a zero-mean sub-Gaussian distribution with a variance proxy λ . The information gain of this sequence calculated via the GNTK is

$$I(G_1, \dots, G_T; k_{\text{GNN}}) = \frac{1}{2} \log \det(\mathbf{I} + \lambda^{-1} \mathbf{K}_{\text{GNN}, T})$$

with the kernel matrix $\mathbf{K}_{\text{GNN}, T} = [k_{\text{GNN}}(G_i, G_j)]_{i, j \leq T}$. The maximum information gain (MIG) [37] is then defined as:

$$\gamma_{\text{GNN}, T} = \max_{\substack{(G_1, \dots, G_T) \\ \forall t: G_t \in \mathcal{G}}} I(G_1, \dots, G_T; k_{\text{GNN}}). \quad (5)$$

In Section 4.3, we express regret bounds in terms of this quantity, as common in kernelized and neural bandits. In Theorem 4.1 we obtain a data-independent bound on the MIG. The proof is given in Appendix B.

Theorem 4.1 (GNTK Information Gain Bound). *Suppose the observation noise is i.i.d., and drawn from a zero-mean sub-Gaussian distribution, and the input domain is \mathcal{G} . Then the maximum information gain associated with k_{GNN} is bounded by*

$$\gamma_{\text{GNN}, T} = \mathcal{O}\left(T^{\frac{d-1}{d}} \log^{\frac{1}{d}} T\right).$$

We observe that the obtained MIG bound *does not* depend on N the number of nodes in the graphs. To highlight this advantage, we compare Theorem 4.1 to the equivalent bound for the vanilla neural tangent kernel which ignores the graph structure. We consider the neural tangent kernel that operates on graphs through the Nd -dimensional vector of aggregated node features $\bar{\mathbf{h}}_G$,

$$\kappa_{\text{NN}}(G, G') = \kappa_{\text{NN}}\left(\frac{\bar{\mathbf{h}}_G}{N}, \frac{\bar{\mathbf{h}}_{G'}}{N}\right). \quad (6)$$

For κ_{NN} the maximum information gain scales as $\gamma_{\text{NN}, T} = \mathcal{O}(T^{(Nd-1)/Nd} \log^{1/Nd} T)$, where N appears in the exponent [25]. This results in poor scalability with graph size in the bandit optimization task, as we further demonstrate in Section 4.3. Table 1 summarizes this comparison.

4.2 Confidence Sets

Quantifying the uncertainty over the reward helps the learner to guide exploration and balance it against exploitation. Confidence sets are an integral tool for uncertainty quantification. Conditioned on the history $H_{t-1} = (G_i, y_i)_{i < t}$, for any $G \in \mathcal{G}$, the set $\mathcal{C}_{t-1}(G, \delta)$ defines an interval to which $f^*(G)$ belongs with a high probability such that,

$$\mathbb{P}(\forall G \in \mathcal{G} : f^*(G) \in \mathcal{C}_{t-1}(G, \delta)) \geq 1 - \delta. \quad (7)$$

An approach common to the kernelized bandit literature is to construct sets of the form $\mathcal{C}_{t-1}(G, \delta) = [\mu_{t-1}(G) \pm \beta_t \sigma_{t-1}(G)]$ where β_t depends on the confidence level δ . The center of the interval, characterized by $\mu_{t-1}(\cdot)$, is the estimate of the reward, and the width $\beta_t \sigma_{t-1}(\cdot)$, reflects the uncertainty. In this work, we utilize GNNs for construction of such sets. To this end, we train a graph neural network to estimate the reward. We use the gradient of this network at initialization to approximate the uncertainty over the reward, as in [45]. Let $f_{\text{GNN}}(G; \boldsymbol{\theta}_{t-1}^{(J)})$ be the GNN trained with gradient descent for J steps and by using learning rate η on the loss

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{t} \sum_{i < t} (f_{\text{GNN}}(G_i, \boldsymbol{\theta}) - y_i)_2^2 + m\lambda \|\boldsymbol{\theta} - \boldsymbol{\theta}^0\|_2^2,$$

where λ is the regularization coefficient, and $\boldsymbol{\theta}^0$ the network parameters at initialization. We propose confidence sets of the form

$$\mathcal{C}_{t-1}(G, \delta) = [\hat{\mu}_{t-1}(G) \pm \beta_t \hat{\sigma}_{t-1}(G)],$$

where the center and width of the set are calculated via,

$$\begin{aligned} \hat{\mu}_{t-1}(G) &:= f_{\text{GNN}}(G; \boldsymbol{\theta}_{t-1}^{(J)}), \\ \hat{\sigma}_{t-1}^2(G) &:= \frac{\mathbf{g}_{\text{GNN}}^T(G; \boldsymbol{\theta}^0)}{\sqrt{m}} \left(\lambda \mathbf{I} + \frac{1}{t} \sum_{i=1}^{t-1} \frac{\mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^0) \mathbf{g}_{\text{GNN}}^T(G_i; \boldsymbol{\theta}^0)}{m} \right)^{-1} \frac{\mathbf{g}_{\text{GNN}}(G; \boldsymbol{\theta}^0)}{\sqrt{m}}. \end{aligned} \quad (8)$$

Here $\mathbf{g}_{\text{GNN}}(G; \boldsymbol{\theta}^0) = \nabla_{\boldsymbol{\theta}} f_{\text{GNN}}(G; \boldsymbol{\theta}^0)$ denotes the gradient at initialization. Moreover, we use $\lambda_0 := \lambda_{\min}(\mathbf{K}_{\text{GNN}}) > 0$ to denote the minimum eigenvalue of the kernel matrix calculated for the entire domain, i.e., $\mathbf{K}_{\text{GNN}} = [k_{\text{GNN}}(G, G')]_{G, G' \in \mathcal{G}}$. Theorem 4.2 shows that this construction gives valid confidence intervals, i.e., it satisfies Eq. (7), when the reward function lies in \mathcal{H}_{GNN} and has a bounded RKHS norm.

Theorem 4.2 (GNN Confidence Bound). *Set $\delta \in (0, 1)$. Suppose $f^* \in \mathcal{H}_{\text{GNN}}$ with a bounded norm $\|f^*\|_{\text{GNN}} \leq B$. Assume that the random sequences $(G_i)_{i < t}$ and $(\epsilon_i)_{i < t}$ are statistically independent. Let the width $m = \text{poly}(t, L, B, |\mathcal{G}|, \lambda, \lambda_0^{-1}, \log(N/\delta))$, learning rate $\eta = C(Lm + m\lambda)^{-1}$ with some universal constant C , and $J \geq 1$. Then for all graphs $G \in \mathcal{G}$, with probability of at least $1 - \delta$,*

$$|f^*(G) - \hat{\mu}_{t-1}(G)| \lesssim \beta_t \hat{\sigma}_{t-1}(G),$$

where $\hat{\mu}_{t-1}$ and $\hat{\sigma}_{t-1}$ are defined in Eq. (8) and

$$\beta_t \approx \sqrt{2}B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{2 \log 2 |\mathcal{G}| / \delta}.$$

The " \approx " notation in Theorem 4.2 omits the terms that vanish with t , i.e., are $o(1)$. An exact version of the theorem without the aforementioned approximations is given in Appendix C.1.

Setting	MIG Bound, γ_T	Cumulative Regret (Phased Elimination)
Neural	$\mathcal{O}\left(T^{\frac{Nd-1}{Nd}} \log^{\frac{1}{Nd}} T\right)$	$\tilde{\mathcal{O}}\left(T^{\frac{2Nd-1}{2Nd}} \log^{\frac{1}{2Nd}} T\right)$
Graph Neural	$\mathcal{O}\left(T^{\frac{d-1}{d}} \log^{\frac{1}{d}} T\right)$	$\tilde{\mathcal{O}}\left(T^{\frac{2d-1}{2d}} \log^{\frac{1}{2d}} T\right)$

Table 1: Summary of main bounds for the NN and GNN approaches. Here T denotes the bandit horizon, N the number of nodes in each graph, and d the dimension of node features. The GNN guarantees are *independent* of N .

4.3 Bandit Optimization with Graph Neural Networks

The developed confidence sets can be used to assist the learner with controlling the growth of regret. In this section, we give a concrete example on how our GNN confidence sets (Equation 8) can be used by an algorithm to solve bandit optimization tasks on graphs.

We introduce GNN-Phased Elimination (GNN-PE; see Algorithm 1) that consists of episodes of pure exploration over a set of plausible maximizer graphs, similar to [7, 30]. Each episode is followed by an elimination step, that makes use of GNN confidence bounds to shrink the set of plausible maximizers. More formally, at step t during an episode e , the learner selects actions via $G_{e,t} = \arg \max_{G \in \mathcal{G}_e} \hat{\sigma}_{e,t-1}(G)$, where \mathcal{G}_e is the set of graphs that might maximize f^* according to the learner’s current knowledge. Once the episode is over after T_e steps, the set \mathcal{G}_e is updated to contain graphs that still have a chance of being a maximizer according to the confidence bounds $[\hat{\mu}_{e,T_e}(G) \pm \beta_{T_e} \hat{\sigma}_{e,T_e}(G)]$ where $\hat{\mu}_{e,T_e}$ and $\hat{\sigma}_{e,T_e}$ are only computed based on the points within episode e .

Theorem 4.3 shows that GNN-PE incurs a sublinear control over the cumulative regret. We provide the proof in Appendix C. We use $\tilde{\mathcal{O}}(\cdot)$ notation to hide $\text{polylog}(T)$ factors.

Theorem 4.3. *Set $\delta \in (0, 1)$. Suppose $f^* \in \mathcal{H}_{k_{\text{GNN}}}$ with a bounded norm $\|f^*\|_{k_{\text{GNN}}} \leq B$. Let the width $m = \text{poly}(t, L, B, |\mathcal{G}|, \lambda, \lambda_0^{-1}, \log(N/\delta))$, learning rate $\eta = C(Lm + m\lambda)^{-1}$ with some universal constant C , and $J \geq 1$. Then with probability at least $1 - \delta$, GNN-PE satisfies*

$$R_T = \tilde{\mathcal{O}}\left(\sqrt{T\gamma_{T,\text{GNN}}}\left(B + \frac{\sigma}{\sqrt{\lambda}}\sqrt{\log|\mathcal{G}|/\delta}\right)\right).$$

We can observe the benefit of working with a graph neural network by comparing the bound in Theorem 4.3 with the regret for a structure-agnostic algorithm. Recall κ_{NN} the vanilla NTK, defined over the concatenated feature vectors (Equation 6). For the sake of this comparison, we ignore the geometric structure and assume that $f^* \in \mathcal{H}_{\text{NN}}$. Swapping out k_{GNN} for κ_{NN} , and respectively the GNN with an NN as defined in Eq. (A.1), we obtain NN-PE, the neural network counterpart of GNN-PE. This algorithm accepts Nd -dimensional input vectors as actions. Similar to Theorem 4.3, we can show that NN-PE can satisfy a guarantee of $\mathcal{O}(T^{(2Nd-1)/2Nd} \log^{1/2Nd} T)$ for the regret. This bound suggests that as N grows, finding the optimal graph can become more challenging for the learner. Working with k_{GNN} to encode the structure of the bandit problem, and consequently using the GNN to solve it, removes the dependency on N in the exponent. This result is summarized in Table 1.

We provide some intuition on why working with a permutation invariant model is beneficial for bandit optimization on graphs. Confidence sets which are constructed for member of \mathcal{H}_{NN} are larger, and result in sub-optimal action selection. Further, training the neural network is a more challenging task, since permutation invariance is not hard coded in the network architecture and has to be learned from the data. This results in less accurate reward estimates. We refer the reader to Appendix A.3 for a more rigorous discussion. There we compare \mathcal{H}_{GNN} and \mathcal{H}_{NN} , the hypothesis spaces corresponding to the two models, through the Mercer decomposition of their kernels.

5 Experiments

We create synthetic datasets which may be of independent interest and can be used for evaluating and benchmarking machine learning algorithms on graph domains. Each dataset is constructed from a finite graph domain together with a reward function. The domains are generated randomly and

differ in properties of the member graphs that influence the problem complexity, e.g., number of nodes and edge density. Each domain $\mathcal{G}_{p,N}$ consists of Erdős-Rényi random graphs, where each graph has N nodes, and between each two nodes there exists an edge with probability p . The node features are i.i.d. $d = 10$ dimensional standard Gaussian vectors. We choose $N \in \{5, 20, 100\}$, $p \in \{0.05, 0.2, 0.95\}$, and thereby sample a total of 9 different domains each containing 10000 graphs. For instance, $\mathcal{G}_{0.05,5}$ denotes the domain with sparse and small graphs, while $\mathcal{G}_{0.95,100}$ is the domain of dense graphs with many nodes. For every domain, we sample a random reward function $f : \mathcal{G}_{p,N} \rightarrow \mathbb{R}$ that is invariant to node permutations. We use $\text{GP}(0, k_{\text{GNN}})$ as a prior, and sample f from its posterior-mean GP. The posterior is calculated using a small random dataset $(G_i, y_i)_{i \leq 5}$, where y_i are drawn independently from $\mathcal{N}(0, 1)$ and G_i are randomly chosen from $\mathcal{G}_{p,N}$. The corresponding dataset is then $D_{p,N} = \{(G_i, f(G_i)) \mid G_i \in \mathcal{G}_{p,N}\}$.

Experiment Setup. Every performance curve in the paper shows an average over 20 runs of the corresponding bandit problem, each with a different action set sampled from $\mathcal{G}_{p,N}$. The shaded areas in all figures show the standard error across runs. In all experiments, the reward is observed with a zero-mean Gaussian noise of variance $\sigma = 10^{-2}$. We always set width $m = 2048$ and layers $L = 2$, for every type of network architecture. Four algorithms appear in our experiments. In addition to our main algorithm GNN-PE, we introduce GNN-UCB, which selects actions via $G_t = \arg \max \hat{\mu}_{t-1}(G) + \beta_t \hat{\sigma}_{t-1}(G)$, the classic UCB policy based on the GNN confidence sets. The pseudo-code is given in Appendix D.3. NN-UCB, introduced by [45], is the neural counterpart of GNN-UCB, and NN-PE as discussed in Section 4.3. To configure these algorithms, we only tune λ and $\beta = \beta_t$, and we do so by using the simplest dataset $D_{0.05,5}$. We find that the algorithms are not sensitive to domain configurations and work for all $D_{p,N}$ out of the box. Therefore, the same values for λ and β are used across all experiments. We include the complete result of our hyperparameter search in Figure 5.

Lazy training. We initialize the graph neural networks (and the NNs) in the lazy regime as described in Eq. (2) (and Eq. A.1). Training a network in this regime with gradient descent causes little change in the weights. Consequently, it is challenging to effectively train a lazy network in practice. Therefore, the stopping criterion for gradient descent plays a crucial role in achieving sublinear regret. Inaccurate estimation of the reward function disturbs the balance of exploration and exploitation, and leads the learner to poor optima. To prevent this issue, we devise a stopping criterion that depends on the history H_{t-1} , such that, as t grows, the network is often trained for more gradient descent steps J . This criterion can be employed by any neural bandit algorithm and may be of independent practical interest. The details of training with gradient descent, stopping and batching are given in Appendix D.2.

Regret Experiments. We assess the performance of the algorithms on bandit optimization tasks over different domains. In Figure 1, we show the inference cumulative regret $\hat{R}_T = \sum_{t \leq T} f^*(G^*) - \max_{G \in \mathcal{G}} \hat{\mu}_{t-1}(G)$, for which we select graph domains with $N = 20$ nodes and edge probability $p = 0.2$. Figure 6 shows the regret for all dataset configurations. To verify scalability with $|\mathcal{G}|$, we run the algorithms on action sets of increasing size $|\mathcal{G}| \in \{200, 500, 1000\}$. Figure 1 presents the results: GNN-PE consistently outperforms the other methods. It is evident that the algorithms built with GNN confidence sets find the optimal graph, regardless of the size of the domain. The GNN algorithms exhibit competitive performance, and attain sublinear regret for all dataset configurations. The neural methods however, may fail to scale and find the optima in limited time.

Scalability with Graph Size. In Section 4.3, we argue that using a neural network which takes $\bar{h}_G \in \mathbb{R}^{Nd}$ as the input, causes the regret to grow with $\mathcal{O}(T^{(2Nd-1)/2Nd})$. The additive structure of the GNN, however, allows the learner to work on a d -dimensional domain, independent of graph size. Figure 2 reflects this behaviour. Fixing $p = 0.2$, and $|\mathcal{G}| = 200$, we run the algorithm over domains with two graph sizes $N \in \{20, 100\}$. GNN-PE achieves sublinear regret in both cases, and manages to find a global maxima within roughly the same number of steps. This is in contrast to NN-PE, which is more affected by increasing graph size. A similar comparison for all configurations and algorithms is plotted in Figure 7, and the same behaviour is observed across all settings: the performance of GNN methods scales well with N , while this is not the case for NN methods.

Effect of Graph Density. As a final observation, we discuss the effect of edge density. Consider a complete graph G with $\binom{N}{2}$ edges. The neighborhoods are symmetric and the aggregated node features $\bar{h}_{G,i}$ are identical for all $i \leq N$. Permutations on this graph will not change the output of either f_{GNN} or f_{NN} . Therefore, we expect that for dense graphs, i.e., large values of p , using a permutation invariant model comes with fewer benefits for the learner. This is opposed to when the graph is sparse

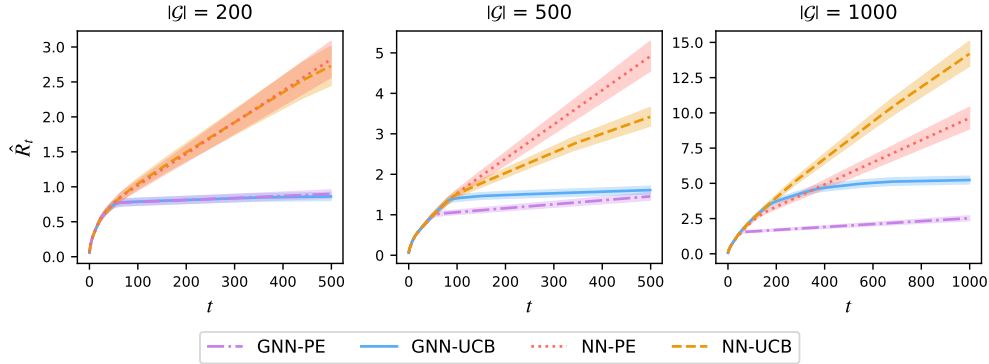


Figure 1: Regret \hat{R}_T over a time horizon of 500 and 1000 steps with $N = 20$ and $p = 0.2$. GNN-PE consistently outperforms other algorithms and scales well with size of the action set $|\mathcal{G}|$.




Figure 2: Increasing N has little effect on GNN-PE.




Figure 3: Increasing the edge density of the graphs reduces the performance gap between GNN-UCB and NN-UCB.

and the neighborhoods are asymmetric. To verify this conjecture, we fix $N = 20$, $|\mathcal{G}| = 200$, and run the algorithm over domains with graphs of different edge probability $p \in \{0.05, 0.2, 0.95\}$. Figure 3 shows that while GNN-UCB always achieves sublinear regret, it takes longer to find the optima when the graphs are more dense. NN-UCB however, improves as the edge probability p grows, since, roughly put, the graphs in the domain are becoming invariant to permutations. Therefore Figure 3 confirms that the performance gap between the two method is reduced for graphs that are more dense. In Figure 8, we plot the effect of graph density for other dataset configurations and bandit algorithms. This behaviour is observed predominantly for the UCB algorithms.

6 Conclusion

We analyze the use of graph neural networks in bandit optimization tasks over large graph domains. The main takeaway is that encoding the natural structure of the environment into the model, reduces the complexity of the task for the learner. By selecting a kernel that embeds invariances, we introduce structure into the algorithm in a principled manner. Importantly, we propose key structural assumptions on the graph reward function and establish a novel connection between additive permutation invariant kernels and the GNTK. We construct valid graph neural network confidence sets, and use it to build a GNN bandit algorithm that achieves sublinear regret. While all node features contribute to the graph's reward, our bounds are independent of the number of nodes. This result holds for GNNs with a single convolutional layer and graphs with node feature representation. An immediate next step is to generalize this approach to other more complex graph neural network architectures and representations (e.g., by including information about graph edges) and investigating their effectiveness for bandit optimization. Our analysis opens up two avenues of future research. The proposed kernel and the graph confidence sets may be used in other algorithms for sequential decision-making tasks on graphs.

Additionally, our approach of embedding the environment’s permutation invariant structure into the algorithm may inspire further work on structured bandit optimization in presence of invariances.

Acknowledgments and Disclosure of Funding

We thank Jonas Rothfuss for his valuable suggestions regarding the experiments. We acknowledge Deepak Narayanan’s effort on an earlier version of the code. We thank Nicolas Emmenegger and Scott Sussex for their thorough feedback, and lastly, we thank Alex Hägele for fruitful discussions regarding the writing. This research was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and Innovation Program Grant agreement no. 815943.

References

- [1] Eric Anderson, Gilman D Veith, and David Weininger. *SMILES, a line notation and computerized interpreter for chemical structures*. US Environmental Protection Agency, Environmental Research Laboratory, 1987.
- [2] Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *Advances in Neural Information Processing Systems*, 2019.
- [3] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 2002.
- [4] Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. *Advances in Neural Information Processing Systems*, 2008.
- [5] Alberto Bietti and Francis Bach. Deep equals shallow for ReLU networks in kernel regimes. In *International Conference on Learning Representations*, 2021.
- [6] Alberto Bietti, Luca Venturi, and Joan Bruna. On the sample complexity of learning under geometric stability. *Advances in Neural Information Processing Systems*, 2021.
- [7] Ilija Bogunovic and Andreas Krause. Misspecified Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, 2021.
- [8] Ilija Bogunovic, Jonathan Scarlett, Andreas Krause, and Volkan Cevher. Truncated variance reduction: A unified approach to Bayesian optimization and level-set estimation. In *Advances in Neural Information Processing Systems*, 2016.
- [9] Ilija Bogunovic, Zihan Li, Andreas Krause, and Jonathan Scarlett. A robust phased elimination algorithm for corruption-tolerant Gaussian process bandits. In *Advances in Neural Information Processing Systems*, 2022.
- [10] Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. *Advances in Neural Information Processing Systems*, 2019.
- [11] Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 2019.
- [12] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, 2017.
- [13] Emile Contal, David Buffoni, Alexandre Robicquet, and Nicolas Vayatis. Parallel gaussian process optimization with upper confidence bound and pure exploration. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2013.
- [14] Nando De, Alex Smola, and Masrour Zoghi. Exponential regret bounds for Gaussian process bandits with deterministic observations. In *International Conference on Machine Learning*, 2012.

- [15] Simon S Du, Kangcheng Hou, Russ R Salakhutdinov, Barnabas Poczos, Ruosong Wang, and Keyulu Xu. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. *Advances in Neural Information Processing Systems*, 2019.
- [16] Vladimir Gligorijević, P Douglas Renfrew, Tomasz Kosciolatek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C Taylor, Ian M Fisk, Hera Vlamakis, et al. Structure-based protein function prediction using graph convolutional networks. *Nature communications*, 2021.
- [17] 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.
- [18] Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science*, 2020.
- [19] Quanquan Gu, Amin Karbasi, Khashayar Khosravi, Vahab Mirrokni, and Dongruo Zhou. Batched neural bandits. *arXiv preprint arXiv:2102.13028*, 2021.
- [20] Kai Guo and Markus J Buehler. A semi-supervised approach to architected materials design using graph neural networks. *Extreme Mechanics Letters*, 2020.
- [21] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, 2018.
- [22] Dejun Jiang, Zhenxing Wu, Chang-Yu Hsieh, Guangyong Chen, Ben Liao, Zhe Wang, Chao Shen, Dongsheng Cao, Jian Wu, and Tingjun Hou. Could graph neural networks learn better molecular representation for drug discovery? a comparison study of descriptor-based and graph-based models. *Journal of cheminformatics*, 2021.
- [23] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*, 2018.
- [24] Kirthevasan Kandasamy, Jeff Schneider, and Barnabás Póczos. High dimensional Bayesian optimisation and bandits via additive models. In *International Conference on Machine Learning*. PMLR, 2015.
- [25] Parnian Kassraie and Andreas Krause. Neural contextual bandits without regret. In *International Conference on Artificial Intelligence and Statistics*, 2022.
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [27] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [28] Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos, Jeff Schneider, and Eric Xing. Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations. In *International Conference on Artificial Intelligence and Statistics*, 2020.
- [29] Andreas Krause and Cheng Ong. Contextual Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, 2011.
- [30] Zihan Li and Jonathan Scarlett. Gaussian process bandit optimization with few batches. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022.
- [31] Xiuyuan Lu and Benjamin Van Roy. Information-theoretic confidence bounds for reinforcement learning. *Advances in Neural Information Processing Systems*, 2019.
- [32] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Learning with invariances in random features and kernel models. *CoRR*, 2021.

- [33] Roman Novak, Lechao Xiao, Jiri Hron, Jaehoon Lee, Alexander A. Alemi, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. Neural tangents: Fast and easy infinite neural networks in Python. In *International Conference on Learning Representations*, 2020.
- [34] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch, 2017.
- [35] Paul Rolland, Jonathan Scarlett, Ilija Bogunovic, and Volkan Cevher. High-dimensional Bayesian optimization via additive models with overlapping groups. In *International Conference on Artificial Intelligence and Statistics*, 2018.
- [36] Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 2014.
- [37] 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.
- [38] Samuel Stanton, Wesley Maddox, Nate Gruver, Phillip Maffettone, Emily Delaney, Peyton Greenside, and Andrew Gordon Wilson. Accelerating Bayesian optimization for biological sequence design with denoising autoencoders. In *International Conference on Machine Learning*, 2022.
- [39] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 1933.
- [40] Sattar Vakili, Nacime Bouziani, Sepehr Jalali, Alberto Bernacchia, and Da shan Shiu. Optimal order simple regret for gaussian process bandits. In *Advances in Neural Information Processing Systems*, 2021.
- [41] Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in Gaussian process bandits. In *International Conference on Artificial Intelligence and Statistics*, 2021.
- [42] Michal Valko, Nathan Korda, Rémi Munos, Ilias Flaounas, and Nello Cristianini. Finite-time analysis of kernelised contextual bandits. In *Conference on Uncertainty in Artificial Intelligence*, 2013.
- [43] Michal Valko, Rémi Munos, Branislav Kveton, and Tomáš Kocák. Spectral bandits for smooth graph functions. In *International Conference on Machine Learning*, 2014.
- [44] Weitong ZHANG, Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural Thompson Sampling. In *International Conference on Learning Representations*, 2021.
- [45] Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with UCB-based exploration. In *International Conference on Machine Learning*, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] We often refer to the Theorem or section number when making a claim about the contribution of the paper.
 - (b) Did you describe the limitations of your work? [Yes] Section 2 explains the theoretical setting for which our result is valid. Any remaining assumption is stated in the theorem statements.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] The impact is primarily theoretical.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] Section 2 explains the theoretical setting for which our result is valid. Section 3 and Section 4 spell out the scope of our algorithm. Any remaining assumption is stated in the theorem statements.
 - (b) Did you include complete proofs of all theoretical results? [Yes] Appendix A contains the proof to Proposition 3.1 and Proposition 3.2. Appendix B present the proof of Theorem 4.1, and lastly Appendix C the proof to Theorem 4.2 and Theorem 4.3.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Appendix D.2 includes the practical details and the instructions required to produce the results. We use random synthetic data which we generate ourselves, instruction to create similar random datasets is also given.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] This is covered in Section 5 and Appendix D.2.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] All experiments are run for 20 random seeds. All plots report the standard error across the runs.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] About 5 hour per CPU core on 100 internal cluster nodes.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] We use the Neural Tangents and PyTorch libraries. Both cited.
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Supplementary Material: Graph Neural Network Bandits

A The Neural Tangent Kernel and its Connection to the GNTK

Let $f(\mathbf{x}; \boldsymbol{\theta}) : \mathbb{R}^d \rightarrow \mathbb{R}$ be a fully-connected network, with L hidden layers of equal width m , and ReLU activations, recursively defined as follows:

$$\begin{aligned} f^{(1)}(\mathbf{x}) &= \mathbf{W}^{(1)}\mathbf{x}, \\ f^{(l)}(\mathbf{x}) &= \sqrt{\frac{2}{m}} \mathbf{W}^{(l)} \sigma_{\text{relu}}(f^{(l-1)}(\mathbf{x})) \in \mathbb{R}^m, \quad 1 < l \leq L \\ f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}) &= \sqrt{2} \mathbf{W}^{(L+1)} \sigma_{\text{relu}}(f^{(L)}(\mathbf{x})) \in \mathbb{R}. \end{aligned} \tag{A.1}$$

The weights $\mathbf{W}^{(i)}$ are initialized to random matrices with standard normal i.i.d. entries, and $\boldsymbol{\theta}^0 = (\mathbf{W}^{(i)})_{i \leq L+1}$. Consider the first order approximation of $f_{\text{NN}}(\mathbf{x}, \boldsymbol{\theta})$ around the initial parameters $\boldsymbol{\theta}^0$, i.e.,

$$\tilde{f}_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}) = \mathbf{g}_{\text{NN}}^T(\mathbf{x}, \boldsymbol{\theta}^0)(\boldsymbol{\theta} - \boldsymbol{\theta}^0),$$

since the network is defined to be zero at initialization. By considering a fixed dataset and a square loss, training with the linear model $\tilde{f}_{\text{NN}}(\mathbf{x}, \boldsymbol{\theta})$, is equivalent to regression with the *tangent kernel* [21], defined as

$$\tilde{k}_{\text{NN}}(\mathbf{x}, \mathbf{x}') = \mathbf{g}_{\text{NN}}^T(\mathbf{x}; \boldsymbol{\theta}^0) \mathbf{g}_{\text{NN}}(\mathbf{x}'; \boldsymbol{\theta}^0). \tag{A.2}$$

The tangent kernel is random since it depends on $\boldsymbol{\theta}^0$. Jacot et al. [21] show that in the infinite width limit, $\tilde{k}_{\text{NN}}(G, G')$ converges to a deterministic kernel, which they call the Neural Tangent Kernel (NTK),

$$\lim_{m \rightarrow \infty} \tilde{k}_{\text{NN}}(\mathbf{x}, \mathbf{x}')/m = k_{\text{NN}}(\mathbf{x}, \mathbf{x}').$$

The NTK satisfies the Mercer condition and has the following Mercer decomposition [5],

$$k_{\text{NN}}(\mathbf{x}, \mathbf{x}') = \sum_{r=0}^{\infty} \lambda_r \sum_{s=1}^{M(d,r)} Y_{s,r}(\mathbf{x}) Y_{s,r}(\mathbf{x}'), \tag{A.3}$$

where $\{Y_{s,r}\}_{s \leq M(d,r)}$ form an orthonormal basis for $V(d, r)$ the space of degree- r polynomials on \mathbb{S}^{d-1} . They eigenvalues λ_r decay at a r^{-d} rate [5]. The eigenfunction $Y_{s,r}$ is the s -th spherical harmonic polynomial of degree r , and $M(d, r) = \dim(V_{d,r})$ gives the total count of such polynomials, where

$$M(d, r) = \frac{2k + d - 2}{r} \binom{r + d - 3}{d - 2}.$$

The NTK adopts a recursive definition (see Appendix A.2). Its properties and connections to infinite-width fully-connected networks are studied in detail [2, 5, 10].

A.1 Properties of GNN and GNTK

We first note the connection between f_{GNN} and f_{NN} .

Lemma A.1 (GNN as sums of NNs). *Consider f_{GNN} the graph neural network defined in Eq. (2), and the feedforward network f_{NN} as given in Eq. (A.1). Then,*

$$f_{\text{GNN}}(G, \boldsymbol{\theta}) = \frac{1}{N} \sum_{j=1}^N f_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \boldsymbol{\theta}).$$

Proof of Lemma A.1. According to Eq. (A.1), the two layer NN with width m decomposes as:

$$f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}) = \sqrt{2} \sum_{j=1}^m w_j^{(2)} \sigma_{\text{relu}}(\langle \mathbf{w}_j^{(1)}, \mathbf{x} \rangle), \quad (\text{A.4})$$

where $\mathbf{w}_j^{(1)} \in \mathbb{R}^d$ are weights in the first layer and $w_j^{(2)} \in \mathbb{R}$ are weights in the second layer. Similarly, the two layer GNN (see Eq. (2)) is given by:

$$f_{\text{GNN}}(G; \boldsymbol{\theta}) = \frac{\sqrt{2}}{N} \sum_{i=1}^N \sum_{j=1}^m w_j^{(2)} \sigma_{\text{relu}}(\langle \mathbf{w}_j^{(1)}, \bar{\mathbf{h}}_{G,i} \rangle) \quad (\text{A.5})$$

$$= \frac{1}{N} \sum_{i=1}^N f_{\text{NN}}(\bar{\mathbf{h}}_{G,i}; \boldsymbol{\theta}), \quad (\text{A.6})$$

where Eq. (A.6) follows from Eq. (A.4). The relation in Eq. (A.6) holds trivially for arbitrary L . \square

We are now ready to show the permutation invariance property.

Lemma A.2 (Geometric Invariance of f_{GNN}). *The graph neural network f_{GNN} is invariant to node permutations, i.e., for all $G \in \mathcal{G}$ and $c \in P_N$,*

$$f_{\text{GNN}}(G; \boldsymbol{\theta}) = f_{\text{GNN}}(c \cdot G; \boldsymbol{\theta})$$

Proof of Lemma A.2. Consider any permutation $c \in P_N$. By Lemma A.1,

$$f_{\text{GNN}}(c \cdot G; \boldsymbol{\theta}) = \frac{1}{N} \sum_{j=1}^N f_{\text{NN}}(\bar{\mathbf{h}}_{G,c(j)}; \boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N f_{\text{NN}}(\bar{\mathbf{h}}_{G,i}; \boldsymbol{\theta}) = f_{\text{GNN}}(G; \boldsymbol{\theta}). \quad (\text{A.7})$$

Since the summation over $f_{\text{NN}}(\bar{\mathbf{h}}_{G,c(j)}; \boldsymbol{\theta})$ for all j , contains the same terms as a sum over $f_{\text{NN}}(\bar{\mathbf{h}}_{G,i}; \boldsymbol{\theta})$ for all i . \square

We now prove that k_{GNN} as defined in Section 2, is deterministic and can be written as a double sum of k_{NN} 's evaluated on $\bar{\mathbf{h}}_{G,j}$ aggregated features of different nodes of the graph.

Proof of Proposition 3.1. For any $\boldsymbol{\theta}^0$ we first show that,

$$\tilde{k}_{\text{GNN}}(G, G') = \frac{1}{N^2} \sum_{j,j'=1}^N \tilde{k}_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j'}), \quad (\text{A.8})$$

where $\tilde{k}_{\text{NN}}(\cdot, \cdot)$ is from Eq. (A.2). Then, we take the $m \rightarrow \infty$ limit. Starting from the definition of \tilde{k}_{GNN} (see Section 3) and by omitting $\boldsymbol{\theta}^0$ for simplicity of notation, we have:

$$\begin{aligned} \tilde{k}_{\text{GNN}}(G, G') &= \mathbf{g}_{\text{GNN}}^T(G) \mathbf{g}_{\text{GNN}}(G') \\ &\stackrel{\text{Lemma A.1}}{=} \left[\sum_{j=1}^N \frac{1}{N} \mathbf{g}_{\text{NN}}^T(\bar{\mathbf{h}}_{G,j}) \right] \left[\sum_{j'=1}^N \frac{1}{N} \mathbf{g}_{\text{NN}}(\bar{\mathbf{h}}_{G',j'}) \right] \\ &= \frac{1}{N^2} \sum_{j,j'=1}^N \mathbf{g}_{\text{NN}}^T(\bar{\mathbf{h}}_{G,j}) \mathbf{g}_{\text{NN}}^{(j)}(\bar{\mathbf{h}}_{G',j'}) \\ &\stackrel{\text{A.2}}{=} \frac{1}{N^2} \sum_{j,j'=1}^N \tilde{k}_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j'}). \end{aligned}$$

The chain of equations above prove Eq. (A.8). Plugging in the definition of the GNTK, we obtain:

$$\begin{aligned}
k_{\text{GNN}}(G, G') &= \lim_{m \rightarrow \infty} \tilde{k}_{\text{GNN}}(G, G')/m \\
&= \lim_{m \rightarrow \infty} \frac{1}{N^2} \sum_{j, j'=1}^N \tilde{k}_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j'})/m \\
&= \frac{1}{N^2} \sum_{j, j'=1}^N \lim_{m \rightarrow \infty} \tilde{k}_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j'})/m \\
&= \frac{1}{N^2} \sum_{j, j'=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j'}),
\end{aligned}$$

where the second equality holds since \tilde{k}_{NN} is continuous, and for continuous functions, limit of finite sums is equal to sum of the limits. This concludes the proof. \square

Proof of Proposition 3.2. From Proposition 3.1, we have

$$k_{\text{GNN}}(G, G') = \frac{1}{N^2} \sum_{j, j'=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,j}, \bar{\mathbf{h}}_{G',j'}).$$

It then suffices to show that $\bar{k}(G, G')$ (as defined in Eq. (1)) is equal to the right hand side of the above equation. Consider P_N the set of permutations of N elements. Every permutation $c \in P_N$ gives a mapping from $(1, \dots, j, \dots, N)$ to $(c(1), \dots, c(j), \dots, c(N))$, where $c(j) \in [N]$ denotes the element that is placed at the j -th position. We define a restricted set of permutations $P_{N|j \rightarrow i} = \{c \in P_N : c(j) = i\}$, such that

$$P_N = \bigcup_{i=1}^N P_{N|j \rightarrow i}. \quad (\text{A.9})$$

Moreover, for any $1 \leq j \leq N$, $\{P_{N|j \rightarrow i}\}_{i=1}^N$ are disjoint sets and the cardinality of each restricted permutation set is

$$|P_{N|j \rightarrow i}| = (N-1)!, \quad (\text{A.10})$$

which implies that the mapping $j \rightarrow i$ is repeated $(N-1)!$ times across the elements of P_N . Back to definition of $\bar{k}(G, G')$ we may decompose P_N and write

$$\begin{aligned}
\bar{k}(G, G') &= \frac{1}{N!} \sum_{c' \in P_N} \left[\frac{1}{N!} \sum_{c \in P_N} \frac{1}{N} \sum_{j=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,c(j)}, \bar{\mathbf{h}}_{G',c'(j)}) \right] \\
&= \frac{1}{N!} \sum_{c' \in P_N} \left[\frac{1}{N!} \frac{1}{N} \sum_{j=1}^N \sum_{c \in P_N} k_{\text{NN}}(\bar{\mathbf{h}}_{G,c(j)}, \bar{\mathbf{h}}_{G',c'(j)}) \right] \\
&\stackrel{\text{Eq. (A.9)}}{=} \frac{1}{N!} \sum_{c' \in P_N} \left[\frac{1}{N!} \frac{1}{N} \sum_{j=1}^N \sum_{i=1}^N \sum_{c \in P_{N|j \rightarrow i}} k_{\text{NN}}(\bar{\mathbf{h}}_{G,c(j)}, \bar{\mathbf{h}}_{G',c'(j)}) \right].
\end{aligned}$$

Now, by definition of $P_{N|i}$, we have $c(j) = i$ for all c in this set. Therefore,

$$\begin{aligned}
\bar{k}(G, G') &= \frac{1}{N!} \sum_{c' \in P_N} \left[\frac{1}{N!} \frac{1}{N} \sum_{j=1}^N \sum_{i=1}^N \sum_{c \in P_{N|j \rightarrow i}} k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',c'(j)}) \right] \\
&\stackrel{\text{Eq. (A.10)}}{=} \frac{1}{N!} \sum_{c' \in P_N} \left[\frac{1}{N!} \frac{1}{N} \sum_{j=1}^N (N-1)! \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',c'(j)}) \right] \\
&= \frac{1}{N!} \sum_{c' \in P_N} \left[\frac{1}{N^2} \sum_{j=1}^N \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',c'(j)}) \right].
\end{aligned}$$

Now, we consider the restricted permutations $P_N|j \rightarrow i'$ and repeat a similar treatment for $c' \in P_N$,

$$\begin{aligned}
\bar{k}(G, G') &= \frac{1}{N^2} \sum_{j=1}^N \frac{1}{N!} \sum_{c' \in P_N} \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',c'(j)}) \\
&= \frac{1}{N^2} \sum_{j=1}^N \frac{1}{N!} \sum_{i'=1}^N \sum_{c' \in P_N|j \rightarrow i'} \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',c'(j)}) \\
&= \frac{1}{N^2} \sum_{j=1}^N \frac{1}{N!} \sum_{i'=1}^N \sum_{c' \in P_N|j \rightarrow i'} \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',i'}) \\
&= \frac{1}{N^2} \sum_{j=1}^N \frac{1}{N!} \sum_{i'=1}^N (N-1)! \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',i'}) \\
&= \frac{1}{N^2} \frac{1}{N!} \sum_{i'=1}^N N! \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',i'}) \\
&= \frac{1}{N^2} \sum_{i'=1}^N \sum_{i=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_{G,i}, \bar{\mathbf{h}}_{G',i'}).
\end{aligned}$$

□

Lemma A.3 (Mercer Decomposition of the GNTK). *The GNTK is Mercer and can be decomposed as*

$$k_{\text{GNN}}(G, G') = \sum_{r=0}^{\infty} \lambda_r \sum_{s=1}^{M(d,r)} Z_{s,r}(\bar{\mathbf{h}}_G) Z_{s,r}(\bar{\mathbf{h}}_{G'})$$

where λ_k are identical to eigenvalues of k_{NN} . The algebraic multiplicity of each λ_r is $M(d, r)$. The eigenfunctions $\{Z_{s,r}\}_{s \leq M(d,r)}$ are degree- r polynomials with the permutation invariant additive structure

$$Z_{s,r}(\bar{\mathbf{h}}_G) := \frac{1}{N} \sum_{j=1}^N Y_{s,r}(\bar{\mathbf{h}}_{G,j}).$$

where $Y_{s,r}$ are degree- r spherical harmonics.

Proof of Lemma A.3. Plugging in the Mercer decomposition of k_{NN} as given in Eq. (A.3) into Proposition 3.1 we get,

$$\begin{aligned}
k_{\text{GNN}}(G, G') &= \frac{1}{N^2} \sum_{j,j'=1}^N \sum_{r=0}^{\infty} \lambda_r \sum_{s=1}^{M(d,r)} Y_{s,r}(\bar{\mathbf{h}}_{G,j}) Y_{s,r}(\bar{\mathbf{h}}_{G',j'}) \\
&= \sum_{r=0}^{\infty} \lambda_r \sum_{s=1}^{M(d,r)} \left(\frac{1}{N} \sum_{j=1}^N Y_{s,r}(\bar{\mathbf{h}}_{G,j}) \right) \left(\frac{1}{N} \sum_{j'=1}^N Y_{s,r}(\bar{\mathbf{h}}_{G',j'}) \right) \quad (\text{A.11}) \\
&= \sum_{r=0}^{\infty} \lambda_r \sum_{s=1}^{M(d,r)} Z_{s,r}(\bar{\mathbf{h}}_G) Z_{s,r}(\bar{\mathbf{h}}_{G'}).
\end{aligned}$$

□

A.2 Recursive Expression for the NTK

For the sake of completeness, we provide a closed-form expression for the NTK function used in Eq. (4) (for more details, see Section 2.1 in [5]). We limit the input space to \mathbb{S}^{d-1} since, by the definition, our feature vectors are always normalized, i.e., $\|\bar{\mathbf{h}}_u\|_2 = 1$ for every $u \in V(G)$. For a ReLU network with L layers considered in Eq. (A.1) with inputs on the sphere (by taking appropriate

limits on the widths), the corresponding $k_{\text{NN}}(\mathbf{x}, \mathbf{x}')$ ([21]) depends on $\angle(\mathbf{x}, \mathbf{x}')$ and is given by $k_{\text{NN}}(\mathbf{x}, \mathbf{x}') = \kappa_{\text{NN}}^{(L)}(\mathbf{x}^T \mathbf{x}')$ where $\kappa_{\text{NN}}^{(L)}(\cdot)$ is defined recursively as follows:

$$\begin{aligned}\kappa_{\text{NN}}^{(1)}(u) &= \kappa^{(1)}(u) = u, \\ \kappa^{(l)}(u) &= \kappa_1(\kappa^{(l-1)}(u)), \\ \kappa_{\text{NN}}^{(l)}(u) &= \kappa_{\text{NN}}^{(l-1)}(u) \cdot \kappa_0(\kappa^{(l-1)}(u)) + \kappa^{(l)}(u) \quad \text{for } 2 \leq l \leq L,\end{aligned}\tag{A.12}$$

where

$$\begin{aligned}\kappa_0(u) &= \frac{1}{\pi}(\pi - \arccos(u)), \\ \kappa_1(u) &= \frac{1}{\pi}\left(u(\pi - \arccos(u)) + \sqrt{1 - u^2}\right).\end{aligned}$$

Finally, we note that $\kappa_{\text{NN}}^{(L)}(1) = 1$ (Bietti and Bach [5]), and hence $k_{\text{NN}}(\mathbf{x}, \mathbf{x}') \leq 1$ for all $\mathbf{x}, \mathbf{x}' \in \mathbb{S}^{d-1}$.

A.3 Effect of Structure on the Hypothesis Space

In Section 4, we demonstrated that the additive permutation invariant structure of k_{GNN} help produce tighter bandit regret and information gain bounds, when the reward function is also permutation invariant. We now characterize how this invariance alters the hypothesis space, independent of the bandit setup. Bietti and Bach [5] give the Mercer decomposition of an NTK defined on a Nd -dimensional sphere. Applying their result we may decompose κ_{NN} as

$$\kappa_{\text{NN}}(G, G') = \sum_{r \geq 0}^{\infty} \lambda_{\text{NN},r} \sum_{s=1}^{N(Nd,r)} Y_{s,r,Nd}(\bar{\mathbf{h}}_G) Y_{s,r,Nd}(\bar{\mathbf{h}}_{G'})$$

where $Y_{s,r} : \mathbb{S}^{Nd-1} \rightarrow \mathbb{R}$ is the s -th degree- r spherical harmonic polynomial. Each eigenvalue $\lambda_{\text{NN},r}$ corresponds to the eigenspace $V_{Nd,r}$, the space of degree- r spherical harmonics, defined on an Nd -dimensional domain. The algebraic multiplicity of each eigenvalue is $M(Nd, r)$ and equal to the dimension of its eigenspace,

$$\dim(V_{Nd,r}) = M(Nd, r) = \frac{2k + d - 2}{r} \binom{r + d - 3}{d - 2} = cr^{Nd-2}.$$

Lastly, the eigenvalues decay at a $\lambda_{\text{NN},r} \simeq c(Nd, L)r^{-Nd}$ rate. The dependence of $c(Nd, L)$ on Nd is exponential, but linear in L , as shown by [5].

We compare this kernel to the GNTK, which has the invariances encoded in its construction. In Lemma A.3 we show that it may be written as

$$k_{\text{GNN}}(G, G') = \sum_{r \geq 0}^{\infty} \lambda_{\text{GNN},r} \sum_{s=1}^{N(d,r)} Z_{s,r,Nd}(\bar{\mathbf{h}}_G) Z_{s,r,Nd}(\bar{\mathbf{h}}_{G'}),$$

where the eigenvalues $\lambda_{\text{GNN},r}$ decay at a $c(d, L)r^{-d}$ rate, and have an algebraic multiplicity of $M(d, r) \simeq cr^{d-2}$. The eigenvectors $Z_{s,r}$ are degree- r polynomials with the permutation invariant additive structure

$$Z_{s,r,Nd}(\bar{\mathbf{h}}_G) := \frac{1}{N} \sum_{j=1}^N Y_{s,r,d}(\bar{\mathbf{h}}_{G,j}).$$

where $Y_{s,r,d}$ are degree- r spherical harmonics, defined on \mathbb{S}^{d-1} . Let $\bar{V}_{Nd,r}$ be the eigenspace corresponding to $\lambda_{\text{GNN},r}$ the r -th eigenvalue. Due to the specific structure of the $Z_{s,r,Nd}$ polynomials, there exists a bijection between $\bar{V}_{Nd,r}$ and $V_{d,r}$ the space of degree- r spherical harmonics defined on \mathbb{S}^{d-1} . The two vector spaces are isomorphic and thus have the same (finite) dimensionality, $\dim(\bar{V}_{Nd,r}) = \dim(V_{d,r})$. This implies that the r -th eigenspace of the GNTK is smaller than the r -th eigenspace of the NTK

$$\frac{\dim(\bar{V}_{Nd,r})}{\dim(V_{Nd,r})} = \frac{\dim(V_{d,r})}{\dim(V_{Nd,r})} \leq \frac{r^d}{r^{Nd}} = \frac{1}{r^{d(N-1)}}.$$

Further, note that $V_{d,r} \subsetneq V_{Nd,r}$. This connection is shown in Figure 4.

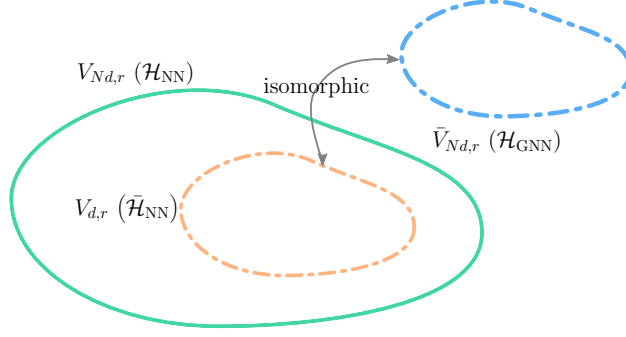


Figure 4: Relation between the GNN and NN hypothesis space

B Information Gain Bounds

In this section, we present a bound for the maximum information gain (defined in Eq. (5)) of the graph neural tangent kernel using the fact that a GNTK can be decomposed as an average of lower-dimensional NTKs (see Eq. (4)).

Proof of Theorem 4.1. We follow a similar technique as in Vakili et al. [41]. Consider an arbitrary sequence of graphs $(G_i)_{i \leq T}$, where each $G_i \in \mathcal{G}$. Let $\mathbf{K}_{\text{GNN}} \in \mathbb{R}^{T \times T}$ denote the corresponding GNTK matrix where

$$[\mathbf{K}_{\text{GNN}}]_{i,j} = k_{\text{GNN}}(G_i, G_j) = \frac{1}{N^2} \sum_{u,u'=1}^N \sum_{r=0}^{\infty} \lambda_r \sum_{s=1}^{M(d,r)} Y_{s,r}(\bar{\mathbf{h}}_{G_i,u}) Y_{s,r}(\bar{\mathbf{h}}_{G_j,u'}),$$

holds by Lemma A.3. We decompose $k_{\text{GNN}}(G_i, G_j)$ into $k_D(G_i, G_j) + k_O(G_i, G_j)$ where

$$k_D(G_i, G_j) = \frac{1}{N^2} \sum_{u,u'=1}^N \sum_{r \leq D} \lambda_r \sum_{s=1}^{M(d,r)} Y_{s,r}(\bar{\mathbf{h}}_{G_i,u}) Y_{s,r}(\bar{\mathbf{h}}_{G_j,u'}),$$

$$k_O(G_i, G_j) = \frac{1}{N^2} \sum_{u,u'=1}^N \sum_{r \geq D+1} \lambda_r \sum_{s=1}^{M(d,r)} Y_{s,r}(\bar{\mathbf{h}}_{G_i,u}) Y_{s,r}(\bar{\mathbf{h}}_{G_j,u'}).$$

The kernel k_D is reproducing for \mathcal{H}_{k_D} , a finite-dimensional subspace of $\mathcal{H}_{k_{\text{GNN}}}$ that is spanned by the eigenfunctions corresponding to the first D distinct eigenvalues. The kernel k_O is reproducing for \mathcal{H}_{k_O} which is orthogonal to \mathcal{H}_{k_D} . Moreover $k_D(G_i, G_j) = \phi_D^T(G_i) \phi_D(G_j)$ where the concatenated feature vector $\phi_D(G)$ is defined as

$$\phi_D(G) = \left(\left(\frac{\sqrt{\lambda_0}}{N} \sum_{u=1}^N Y_{s,0}(\bar{\mathbf{h}}_{G,u}) \right)_{s \leq N(d,0)}, \dots, \left(\frac{\sqrt{\lambda_D}}{N} \sum_{u=1}^N Y_{s,D}(\bar{\mathbf{h}}_{G,u}) \right)_{s \leq N(d,D)} \right).$$

Here, $\phi_D(G) \in \mathbb{R}^{\tilde{D}}$ where

$$\tilde{D} = \sum_{r=0}^D M(d,r) \simeq C \sum_{r=0}^D r^{d-2} \leq C \frac{(D+1)^{d-1}}{D-1}, \quad (\text{B.1})$$

since by Stirling's approximation $M(d,r)$ grows with r^{d-2} [5].

Recall that the information gain is $I(\mathbf{y}_T; \mathbf{f}_T) = \frac{1}{2} \log \det(\mathbf{I} + \lambda^{-1} \mathbf{K}_{\text{GNN}})$. Defining \mathbf{K}_D and \mathbf{K}_O such that $[\mathbf{K}_D]_{i,j} = k_D(G_i, G_j)$ and $[\mathbf{K}_O]_{i,j} = k_O(G_i, G_j)$, we get $\mathbf{K}_{\text{GNN}} = \mathbf{K}_D + \mathbf{K}_O$ and therefore,

$$\begin{aligned} I(\mathbf{y}_T; \mathbf{f}_T) &= \frac{1}{2} \log \det(\mathbf{I} + \lambda^{-1} (\mathbf{K}_D + \mathbf{K}_O)) \\ &= \frac{1}{2} \log \det(\mathbf{I} + \lambda^{-1} \mathbf{K}_D) + \frac{1}{2} \log \det(\mathbf{I} + (\mathbf{I} + \lambda^{-1} \mathbf{K}_D)^{-1} \mathbf{K}_O). \end{aligned} \quad (\text{B.2})$$

We bound each term separately, starting with the first term.

Consider the $T \times \tilde{D}$ feature matrix $\Phi_D = [\phi_D(G_1), \dots, \phi_D(G_T)]$. Then $\mathbf{K}_D = \Phi_D \Phi_D^T$ and by the Weinstein-Aronszajn identity,

$$\frac{1}{2} \log \det(\mathbf{I} + \lambda^{-1} \mathbf{K}_D) = \frac{1}{2} \log \det(\mathbf{I} + \lambda^{-1} \Phi_D^T \Phi_D).$$

For positive definite matrices $\mathbf{P} \in \mathbb{R}^{n \times n}$, we have $\log \det \mathbf{P} \leq n \log \text{tr}(\mathbf{P}/n)$. Applying this identity we get,

$$\begin{aligned} \frac{1}{2} \log \det(\mathbf{I} + \lambda^{-1} \mathbf{K}_D) &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\lambda^{-1}}{\tilde{D}} \text{tr}(\Phi_D^T \Phi_D) \right) \\ &= \frac{1}{2} \tilde{D} \log \left(1 + \frac{\lambda^{-1}}{\tilde{D}} \sum_{t=1}^T \phi_D^T(G_t) \phi_D(G_t) \right) \\ &= \frac{1}{2} \tilde{D} \log \left(1 + \frac{\lambda^{-1}}{\tilde{D}} \sum_{t=1}^T \|\phi_D(G_t)\|_2^2 \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\lambda^{-1}}{\tilde{D}} \sum_{t=1}^T \sum_{r=0}^D \frac{\lambda_r}{N^2} \sum_{s=1}^{M(d,r)} \left(\sum_{u=1}^N Y_{s,r}(\bar{\mathbf{h}}_{G_t,u}) \right)^2 \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\lambda^{-1}}{\tilde{D}} \sum_{t=1}^T \sum_{u,u'=1}^N \sum_{r=0}^D \frac{\lambda_r}{N^2} \sum_{s=1}^{M(d,r)} Y_{s,r}(\bar{\mathbf{h}}_{G_t,u}) Y_{s,r}(\bar{\mathbf{h}}_{G_t,u'}) \right) \\ &= \frac{1}{2} \tilde{D} \log \left(1 + \frac{\lambda^{-1}}{\tilde{D}} \sum_{t=1}^T k_D(G_t, G_t) \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{T/\lambda}{\tilde{D}} \right). \end{aligned} \tag{B.3}$$

The last inequality holds since by definition, k_{GNN} is uniformly bounded by 1 on the unit sphere (this holds since k_{NN} is also uniformly bounded by 1 on the same domain; see Appendix A.2).

For bounding the second term, we again use the $\log \det \mathbf{P} \leq n \log \text{tr}(\mathbf{P}/n)$ inequality and write,

$$\begin{aligned} \frac{1}{2} \log \det(\mathbf{I} + (\mathbf{I} + \lambda^{-1} \mathbf{K}_D)^{-1} \mathbf{K}_O) &\leq \frac{T}{2} \log \left(1 + \frac{\text{tr}((\mathbf{I} + \lambda^{-1} \mathbf{K}_D)^{-1} \mathbf{K}_O)}{T} \right) \\ &\leq \frac{T}{2} \log(1 + \text{tr}(\mathbf{K}_O)/T). \end{aligned}$$

The second inequality holds due to $(\mathbf{I} + \lambda^{-1} \mathbf{K}_D)^{-1}$ being positive definite, with eigenvalues smaller than 1. To bound $\text{tr}(\mathbf{K}_O)$, note that

$$\begin{aligned} [\mathbf{K}_O]_{i,j} &= \frac{1}{N^2} \sum_{u,u'=1}^N \sum_{r \geq D+1} \lambda_r \sum_{s=1}^{M(d,r)} Y_{s,r}(\bar{\mathbf{h}}_{G_i,u}) Y_{s,r}(\bar{\mathbf{h}}_{G_j,u'}) \\ &= \frac{1}{N^2} \sum_{u,u'=1}^N \sum_{r \geq D+1} \lambda_r M(d,r) P_r(\langle \bar{\mathbf{h}}_{G_i,u}, \bar{\mathbf{h}}_{G_j,u'} \rangle) \\ &\leq \frac{1}{N^2} \sum_{u,u'=1}^N \sum_{r \geq D+1} \lambda_r M(d,r) \\ &\leq \sum_{r \geq D+1} \lambda_r M(d,r) \end{aligned}$$

The second equality follows from $\sum_{s=1}^{M(d,r)} Y_{s,r}(\mathbf{x}) Y_{s,r}(\mathbf{x}') = M(d,r) P_r(\mathbf{x}^T \mathbf{x}')$, where P_r is the degree- r Legendre polynomial [5]. The Legendre basis is bounded in $[0, 1]$, resulting in the

first inequality. By Bietti and Bach [5, Corollary 3], there exists a constant $c_1(d, L)$ such that $\lambda_r \leq c_1(d, L)r^{-d}$. Stirling's approximation states that $n! \sim \sqrt{2\pi n}(n/e)^n$. Therefore, there exists a constant c_2 such that $M(d, r) \leq c_2r^{d-2}$. Therefore, there exists a constant $c(d, L)$ such that,

$$[\mathbf{K}_O]_{i,j} \leq \sum_{r \geq D+1} \lambda_r M(d, r) \leq c(d, L) \sum_{r \geq D+1} r^{-2} \leq \frac{c(d, L)}{D}$$

where the second inequality comes from

$$\sum_{r \geq D+1} r^{-2} \leq \int_D^\infty z^{-2} dz = \frac{1}{D}.$$

Therefore, we may bound the second term of the information gain as follows,

$$\frac{1}{2} \log \det(\mathbf{I} + (\mathbf{I} + \lambda^{-1} \mathbf{K}_D)^{-1} \mathbf{K}_O) \leq \frac{T}{2} \log \left(1 + \frac{c(d, L)}{D} \right) \quad (\text{B.4})$$

From Eq. (B.2), Eq. (B.3) and Eq. (B.4),

$$I(\mathbf{y}_T; \mathbf{f}_T) \leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{T\lambda^{-1}}{\tilde{D}} \right) + \frac{T}{2} \log \left(1 + \frac{c(d, L)}{D} \right). \quad (\text{B.5})$$

For the first term to dominate the second, \tilde{D} has to be set to

$$\tilde{D} = \left\lceil \left(\frac{c(d, L)T}{\log(1 + T/\lambda)} \right)^{\frac{d-1}{d}} \right\rceil.$$

This results in

$$\gamma_T = \mathcal{O} \left(\left(\frac{T}{\log(1 + \frac{T}{\lambda})} \right)^{\frac{d-1}{d}} \log \left(1 + \frac{T}{\lambda} \left(\frac{\log(1 + \frac{T}{\lambda})}{T} \right)^{\frac{d-1}{d}} \right) \right),$$

that is itself $\mathcal{O}(T^{\frac{d-1}{d}} \log^{\frac{1}{d}} T)$, therefore concluding the proof. \square

C Proof of Theorem 4.2 and Theorem 4.3

Proof of Theorem 4.3. The proposed algorithm GNN-PE (see Algorithm 1) is a variant of the Phased GP Uncertainty algorithm proposed in [7]. The following analysis closely follows the one of [7] (but ignores misspecification), with important differences pertained to the introduction of GNN estimator and GNTK analysis. The algorithm runs in episodes of exponentially increasing length T_e , and maintains a set of potentially optimal graphs \mathcal{G}_e . To compute the set of potentially optimal graphs after every episode, it uses the confidence bounds from Theorem 4.2. The total number of episodes is denoted with E , and it holds that $E \leq \lceil \log_2 T \rceil$, since the length of the episode is growing exponentially.

To bound the regret of GNN-PE, we make use of the finite-dimensional tangent kernel. With a slightly different notation from Section 3, we set $\hat{k}_{\text{GNN}}(\cdot, \cdot) = \mathbf{g}_{\text{GNN}}^T(\cdot; \boldsymbol{\theta}^0) \mathbf{g}_{\text{GNN}}(\cdot; \boldsymbol{\theta}^0) / m$, where $\mathbf{g}_{\text{GNN}}^T(\cdot; \boldsymbol{\theta}^0)$ denotes the gradient of the GNN at initialization. We argued in the main text that this kernel can well approximate k_{GNN} . The feature map corresponding to this kernel, $\hat{\phi}(G) = \mathbf{g}_{\text{GNN}}(G) / m$ can be viewed as a finite-dimensional approximation of ϕ_{GNN} , the (infinite length) feature map of the GNTK.

Throughout the proof, we denote the posterior mean and variance calculated via $\text{GP}(0, \hat{k}_{\text{GNN}})$ by $\hat{\mu}_{t-1}$ and $\hat{\sigma}_{t-1}$, respectively. Recall that the posterior mean and variance function after observing the data $(G_i, y_i)_{i < t}$ is calculated via

$$\begin{aligned} \hat{\mu}_{t-1}(G; \hat{k}_{\text{GNN}}) &= \hat{\mathbf{k}}_{t-1}^T(G) (\hat{\mathbf{K}}_{t-1} + \lambda \mathbf{I})^{-1} \mathbf{y}_{t-1}, \\ \hat{\sigma}_{t-1}^2(G; \hat{k}_{\text{GNN}}) &= \hat{k}_{\text{GNN}}(G, G) - \hat{\mathbf{k}}_{t-1}^T(G) (\hat{\mathbf{K}}_{t-1} + \lambda \mathbf{I})^{-1} \hat{\mathbf{k}}_{t-1}(G), \end{aligned} \quad (\text{C.4})$$

where the constant λ is the variance proxy for observation noise. Here $\mathbf{y}_{t-1} = [y_i]_{i < t}$ is the vector of observed values, $\hat{\mathbf{k}}_{t-1}(\mathbf{x}) = [\hat{k}_{\text{GNN}}(\mathbf{x}, \mathbf{x}_\tau)]_{i < t}$, and $\hat{\mathbf{K}}_{t-1} = [\hat{k}_{\text{GNN}}(\mathbf{x}_i, \mathbf{x}_j)]_{i, j < t}$ is the kernel matrix.

Algorithm 1: GNN PHASED ELIMINATION (GNN-PE)

Input: m, J, η, λ, T Set episode index $e = 1$, episode length $T_e = 1$, and set of potentially optimal graphs $\mathcal{G}_e = \mathcal{G}$ **Initialize** network parameters to a random θ^0 .**for** $t = 1, \dots, T_e$ **do** For all $G \in \mathcal{G}$, calculate $\hat{\sigma}_{t-1}(G)$ as defined in Eq. (8).

Select

$$G_t = \arg \max_{G \in \mathcal{G}_e} \hat{\sigma}_{t-1}^2(G) \quad (\text{C.1})$$

endReceive $\{y_1, \dots, y_{T_e}\}$, such that

$$y_t = f^*(G_t) + \epsilon_t \quad \text{for } t \in \{1, \dots, T_e\}$$

Calculate $\theta_e^{(J)} = \text{TrainGNN} \left(m, J, \eta, \lambda, \theta^0, (G_t, y_t)_{t=1}^{T_e} \right)$ Use β and ε as defined in Equation (C.7) and Equation (C.8) and update

$$\mathcal{G}_{e+1} \leftarrow \left\{ G \in \mathcal{G}_e : f_{\text{GNN}}(G; \theta_e^{(J)}) + \beta \hat{\sigma}_{T_e}(G) + 2\varepsilon \geq \max_{G \in \mathcal{G}_e} (f_{\text{GNN}}(G; \theta_e^{(J)}) - \beta \hat{\sigma}_{T_e}(G)) \right\}, \quad (\text{C.2})$$

Set $T_{e+1} \leftarrow 2T_e$, $e \leftarrow e + 1$ and return to the **Initialize** step**Output:** Terminate after T total evaluations (E total episodes) and return

$$\hat{G}_T = \arg \max_{G \in \mathcal{G}} f_{\text{GNN}}(G; \theta_{E-1}^{(J)}) \quad (\text{C.3})$$

Algorithm 2: TrainGNN

Input: $m, J, \eta, \lambda, \theta^0, (G_i, y_i)_{i < t}$ Define $\mathcal{L}(\theta) = \frac{1}{t} \sum_{i < t} (f_{\text{GNN}}(G_i; \theta) - y_i)^2 + m\lambda \|\theta - \theta^0\|_2^2$ Initialize $\theta^{(0)} = \theta^0$ **for** $j = 1, \dots, J$ **do** $\theta^{(j)} = \theta^{(j-1)} - \eta \nabla \mathcal{L}(\theta^{(j-1)})$ **end****Output:** $\theta^{(J)}$

We note that GNN-PE uses $\hat{\sigma}_{t-1}$ as the variance estimate, while instead of $\hat{\mu}_{t-1}$, the algorithm makes use of the GNN predictions, i.e., $f_{\text{GNN}}(G; \theta_t^{(J)})$ for the center of the confidence set.

As will become clear soon, since GNN-PE uses $\hat{\sigma}_{t-1}$, this yields a regret bound depending on $\hat{\gamma}_T$, the information gain corresponding to the approximate kernel \hat{k}_{GNN} . Lastly, in Lemma C.1, for appropriately set width m , we bound $\hat{\gamma}_T$ with $\gamma_{\text{GNN}, T}$ (maximum information gain corresponding to the exact GNTK). We use this result in our steps below.

Step 1 (Max variance bound) Consider any fixed episode e , and recall that T_e denotes the episode length. By the exploration policy of GNN-PE (Equation (C.1)), at any step t (within an episode) and for any graph $G \in \mathcal{G}_e$, we have $\hat{\sigma}_{t-1}(G) \leq \hat{\sigma}_{t-1}(G_t)$. From Eq. (C.4), since the covariance matrix is positive definite, conditioning on a larger set of points reduces the posterior variance and thus $\hat{\sigma}_{T_e}(G) \leq \hat{\sigma}_{t-1}(G)$, for all $G \in \mathcal{G}_e$ and $t \leq T_e$. Putting the two inequalities together, $\hat{\sigma}_{T_e}(G) \leq \min_{t \leq T_e} \hat{\sigma}_{t-1}(G_t)$, which gives

$$\hat{\sigma}_{T_e}^2(G) \leq \frac{1}{T_e} \sum_{t=1}^{T_e} \hat{\sigma}_{t-1}^2(G_t).$$

For any $s \in [0, 1/\lambda]$, it holds that

$$s^2 \leq \frac{1}{\lambda \log(1 + 1/\lambda)} \log(1 + s^2).$$

For any G_t , we have $\hat{\sigma}_{t-1}^2(G_t)/\lambda \leq \hat{k}_{\text{GNN}}(G_t, G_t)/\lambda \leq 1/\lambda$. Therefore,

$$\hat{\sigma}_{t-1}^2(G_t) \leq \frac{1}{\log(1 + 1/\lambda)} \log(1 + \lambda^{-2} \hat{\sigma}_{t-1}^2(G_t)).$$

From Srinivas et al. [37, Lemma 5.3], we can conclude that,

$$\sum_{t=1}^{T_e} \log(1 + \lambda^{-1} \hat{\sigma}_{t-1}^2(G_t)) = 2I(G_1, \dots, G_{T_e}; \hat{k}_{\text{GNN}}) \leq 2\hat{\gamma}_{T_e},$$

where $\hat{\gamma}_{T_e}$ is the maximum information gain corresponding to this episode and the kernel \hat{k}_{GNN} . This inequality allows us to bound the posterior variance at the end of episode e as follows. For any $G \in \mathcal{G}_e$,

$$\hat{\sigma}_{T_e}^2(G) \leq \frac{2\hat{\gamma}_{T_e}}{T_e \log(1 + \lambda^{-1})}. \quad (\text{C.5})$$

Step 2 (Confidence bounds) Consider an episode e , and let $\tilde{\delta} = \delta/(3E)$, where $E \leq \log T$ is the number of episode. From Theorem C.2, for all $G \in \mathcal{G}$, with probability at least $1 - \tilde{\delta}$,

$$|f_{\text{GNN}}(G; \boldsymbol{\theta}_e^{(J)}) - f^*(G)| \leq \beta \hat{\sigma}_{T_e}(G) + \varepsilon, \quad (\text{C.6})$$

where for simplicity we use

$$\beta := \sqrt{2}B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{2 \log 2|\mathcal{G}|/\tilde{\delta}} + \sqrt{\frac{2B}{m\eta\lambda}} \left(\sqrt{2} + (1 - m\eta\lambda)^{J/2} \right), \quad (\text{C.7})$$

$$\varepsilon := \tilde{C}L^3 \left(\frac{B}{m\lambda} \right)^{2/3} \sqrt{m \log m}. \quad (\text{C.8})$$

Applying the union bound Eq. (C.6) holds for every $G \in \mathcal{G}$ and $e \in [E]$ with probability at least $1 - \delta/3$. Finally, by using $\tilde{\delta} = \delta/3$ in Lemma C.1 and by applying the union bound, we have that both events in Lemma C.1 and in Eq. (C.6) hold jointly with probability at least $1 - 2\delta/3$. In the rest of the proof, we condition on the joint event holding true. This implies that $G^* \in \mathcal{G}_e$ for every $e \in [E]$, i.e., according to the rule in Eq. (C.2), the algorithm will not eliminate G^* .

Step 3 (Cumulative Regret) We use R_e to denote the episodic regret, and write

$$R_T = \sum_{e=1}^E R_e \leq m_1 B + \sum_{e=2}^E \sum_{t=1}^{T_e} (f^*(G^*) - f^*(G_t^{(e)})). \quad (\text{C.9})$$

The first inequality follows since k_{GNN} is uniformly bounded by 1 and the RKHS norm of f^* is bounded by B . We also add an additional superscript e in $G_t^{(e)}$, to denote a graph selected in episode e at time step t .

Consider any episode e , the following holds due to Eq. (C.6):

$$\begin{aligned} f^*(G^*) - f^*(G_t^{(e)}) &\leq f_{\text{GNN}}(G^*; \boldsymbol{\theta}_{e-1}^{(J)}) + \beta \hat{\sigma}_{T_{e-1}}(G^*) + 2\varepsilon \\ &\quad - (f_{\text{GNN}}(G_t^{(e)}; \boldsymbol{\theta}_{e-1}^{(J)}) - \beta \hat{\sigma}_{T_{e-1}}(G_t^{(e)})). \end{aligned}$$

Moreover, we use the elimination rule in Eq. (C.2) to obtain:

$$\begin{aligned} f^*(G^*) - f^*(G_t^{(e)}) &\leq (f_{\text{GNN}}(G^*; \boldsymbol{\theta}_{e-1}^{(J)}) - \beta \hat{\sigma}_{T_{e-1}}(G^*) - \varepsilon) \\ &\quad + 2\beta \hat{\sigma}_{T_{e-1}}(G^*) + 4\varepsilon \\ &\quad - (f_{\text{GNN}}(G_t^{(e)}; \boldsymbol{\theta}_{e-1}^{(J)}) + \beta \hat{\sigma}_{T_{e-1}}(G_t^{(e)}) + \varepsilon) \\ &\quad + 2\beta \hat{\sigma}_{T_{e-1}}(G_t^{(e)}) \\ &\leq 4\beta_{T_{e-1}+1} \max_{G \in \mathcal{G}_{e-1}} \hat{\sigma}_{T_{e-1}}(G) + 4\varepsilon. \end{aligned} \quad (\text{C.10})$$

Next, we combine Eq. (C.10) and Eq. (C.9) and write:

$$R_T \leq m_1 B + \sum_{e=2}^E \sum_{t=1}^{T_e} 4\beta_{T_{e-1}+1} \max_{G \in \mathcal{G}_{e-1}} \hat{\sigma}_{T_{e-1}}(G) + 4\epsilon \quad (\text{C.11})$$

$$\leq m_1 B + \sum_{e=2}^E 4T_e \beta \sqrt{\frac{2\hat{\gamma}_{T_{e-1}}}{T_{e-1} \log(1+\lambda^{-1})}} + 4T_e \epsilon \quad (\text{C.12})$$

$$= m_1 B + \sum_{e=2}^E 8\beta \sqrt{\frac{2T_{e-1}\hat{\gamma}_{T_{e-1}}}{\log(1+\lambda^{-1})}} + 4T_e \epsilon \quad (\text{C.13})$$

$$\leq m_1 B + \sum_{e=2}^E 8\beta \sqrt{\frac{2T\hat{\gamma}_T}{\log(1+\lambda^{-1})}} + 4T_e \epsilon \quad (\text{C.14})$$

$$\leq m_1 B + 8(\log(T) + 1) \left(\beta \sqrt{\frac{2T\hat{\gamma}_T}{\log(1+\lambda^{-1})}} + 4T\epsilon \right) \quad (\text{C.15})$$

$$\leq m_1 B + 8(\log(T) + 1) \left(\beta \sqrt{\frac{2T(\gamma_T + \epsilon(m))}{\log(1+\lambda^{-1})}} + 4T\epsilon \right), \quad (\text{C.16})$$

where Eq. (C.12) follows since $T > m_e$ for every e and Eq. (C.5) and Eq. (C.13) since $T_e = 2T_{e-1}$. To obtain Eq. (C.14), we use that $T > T_{e-1}$ and $\hat{\gamma}_T \geq \hat{\gamma}_{T_{e-1}}$. Finally, Eq. (C.15) follows since the number of episodes $E \leq \lceil \log T \rceil$, and Eq. (C.16) follows from Lemma C.1.

Step 4 (Putting everything together) Plugging in the expressions for β and ϵ we obtain

$$\begin{aligned} R_T \leq & 8\sqrt{2}(\log T + 1) \sqrt{\frac{2T(\gamma_T + \epsilon(m))}{\log(1+\lambda^{-1})}} \left(B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{\log \frac{|\mathcal{G}|(\log T + 1)}{\delta}} \sqrt{\frac{B}{m\eta\lambda}} \left(\sqrt{2} + (1 - m\eta\lambda)^{J/2} \right) \right) \\ & + B + 32(\log T + 1) \tilde{C} T L^3 \left(\frac{B}{m\lambda} \right)^{2/3} \sqrt{m \log m} \end{aligned}$$

Since $m = \text{poly}(t)$ and $\eta \sim 1/m$, the last term and $\epsilon(m)$ vanishes with T at a $o(1)$ rate, and the gradient descent error term becomes a constant factor. Then we obtain with probability at least $1 - \delta$

$$R_T = \mathcal{O} \left(\sqrt{T\gamma_T} \left(B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{\log \frac{|\mathcal{G}| \log T}{\delta}} \right) \right). \quad (\text{C.17})$$

□

Lemma C.1 (Bounding MIG with its approximation). Set $\delta \in (0, 1)$. If $m = \text{poly}(t, L, |\mathcal{G}|, \lambda, B, \lambda_0^{-1}, \log(N/\delta))$, then with probability at least $1 - \delta$,

$$\hat{\gamma}_T \leq \gamma_T + \epsilon(m),$$

where γ_T is the maximum information gain of the GNTK over \mathcal{G} as defined in Eq. (5), and $\epsilon(m) = o(m^{-1/4})$.

Proof of Lemma C.1. The proof follows from Lemma C.6, by repeating the technique given in Lemma D.5, [25]. Here we repeat it for the sake of completeness.

Consider an arbitrary sequence of graphs $(G_t)_{t \leq T}$. Consider the feature map $\hat{\phi}(G) = \mathbf{g}_{\text{GNN}}(G; \theta^0) / \sqrt{m}$. For the kernel \hat{k}_{GNN} , which corresponds to this feature map, the information gain after observing T samples is

$$\hat{I}_T = \frac{1}{2} \log \det (\mathbf{I} + \lambda^{-1} \bar{\mathbf{G}}_T \bar{\mathbf{G}}_T^T / m),$$

where $\bar{\mathbf{G}}_T = [\bar{\mathbf{g}}(G_t)]_{t \leq T}^T$. Let $[\mathbf{K}_{\text{GNN}}]_{i,j \leq T} = k_{\text{GNN}}(G_i, G_j)$ with k the NTK function of the fully-connected L -layer network.

$$\begin{aligned}
\hat{I}_T &= \frac{1}{2} \log \det (\mathbf{I} + \lambda^{-1} \mathbf{K}_{\text{GNN}} + \lambda^{-1} (\bar{\mathbf{G}}_T \bar{\mathbf{G}}_T^T / m - \mathbf{K}_{\text{GNN}})) \\
&\stackrel{(a)}{\leq} \frac{1}{2} \log \det (\mathbf{I} + \lambda^{-1} \mathbf{K}_{\text{GNN}}) + \langle (\mathbf{I} + \lambda^{-1} \mathbf{K}_{\text{GNN}})^{-1}, \lambda^{-1} (\bar{\mathbf{G}}_T \bar{\mathbf{G}}_T^T / m - \mathbf{K}_{\text{GNN}}) \rangle \\
&\leq I_T + \lambda^{-1} \|(\mathbf{I} + \lambda^{-1} \mathbf{K}_{\text{GNN}})^{-1}\|_F \|\bar{\mathbf{G}}_T \bar{\mathbf{G}}_T^T / m - \mathbf{K}_{\text{GNN}}\|_F \\
&\stackrel{(b)}{\leq} I_T + \lambda^{-1} \sqrt{T} \|\bar{\mathbf{G}}_T \bar{\mathbf{G}}_T^T / m - \mathbf{K}_{\text{GNN}}\|_F \\
&\stackrel{(c)}{\leq} I_T + \lambda^{-1} T \sqrt{T} \epsilon \\
&\stackrel{(d)}{\leq} \gamma_T + \epsilon(m).
\end{aligned} \tag{C.18}$$

Inequality (a) holds by concavity of $\log \det(\cdot)$. Inequality (b) holds since $\mathbf{I} \preceq \mathbf{I} + \lambda^{-1} \mathbf{K}_{\text{GNN}}$. Inequality (c) holds due to Lemma C.6. Finally, inequality (d) uses the polynomial choice of m , and requires that m grows with at least $O(T^6)$. Equation C.18 holds for any arbitrary context set, thus it also holds for the sequence which maximizes the information gain. \square

C.1 Proof of Theorem 4.2

We first present the formal version of Theorem 4.2.

Theorem C.2 (GNN Confidence Bound, Formal). *Set $\delta \in (0, 1)$. Suppose $f^* \in \mathcal{H}_{k_{\text{GNN}}}$ with a bounded norm $\|f^*\|_{k_{\text{GNN}}} \leq B$. Samples of f are observed with zero-mean σ^2 -sub-Gaussian noise. Assume that the random sequences $(G_i)_{i < t}$ and $(\epsilon_i)_{i < t}$ are statistically independent. Set $J > 1$, choose the width $m = \text{poly}(t, L, |\mathcal{G}|, \lambda, \lambda_0^{-1}, \log(N/\delta))$, and learning rate $\eta = C(Lm + m\lambda)^{-1}$ with some universal constant C . Then for all graphs $G \in \mathcal{G}$, with probability of at least $1 - \delta$,*

$$|f^*(G) - \hat{\mu}_{t-1}(G)| \leq \beta \hat{\sigma}_t(G) + \tilde{C} L^3 \left(\frac{B}{m\lambda}\right)^{2/3} \sqrt{m \log m}$$

where

$$\beta = \sqrt{2}B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{2 \log(2|\mathcal{G}|/\delta)} + \sqrt{\frac{2B}{m\eta\lambda}} (\sqrt{2} + (1 - m\eta\lambda)^{J/2})$$

for some constant \tilde{C} .

To prove the theorem, first we state the necessary lemmas.

Lemma C.3 (Confidence interval for f_{GNN} around $\hat{\mu}_{t-1}$). *Assume history $H_t = \{(G_i, y_i)\}_{i \leq t}$ with $(G_i)_{i \leq t}$ and $(\epsilon_i)_{i \leq t}$ statistically independent. Let $m = \text{poly}(t, L, \lambda, \log(|\mathcal{G}|N/\delta))$. There exists C_1 , such that for any $\delta > 0$, if the learning rate is picked $\eta = C_1(Lm + m\lambda)^{-1}$, then for a graph $G \in \mathcal{G}$, with probability of at least $1 - \delta$,*

$$|f_{\text{GNN}}(G; \boldsymbol{\theta}^{(j)}) - \hat{\mu}_t(G)| \leq \hat{\sigma}_t(G) \sqrt{\frac{2B}{m\eta\lambda}} (\sqrt{2} + (1 - m\eta\lambda)^{J/2}) + \tilde{C} L^3 \left(\frac{B}{m\lambda}\right)^{2/3} \sqrt{m \log m}$$

for some constant \tilde{C} , where $\hat{\mu}_t$ and $\hat{\sigma}_t$ are as defined in Eq. (C.4).

Proof of Lemma C.3. We define $\mathbf{Z} := \lambda \mathbf{I} + \bar{\mathbf{G}}^T \bar{\mathbf{G}} / (Tm)$ and $\mathbf{b} = \sum_{i \leq t} y_i \bar{\mathbf{g}}(G_i) / (T\sqrt{m})$. Recall that $\bar{\mathbf{G}}_T = [\bar{\mathbf{g}}(G_t)]_{t \leq T}^T$. Let the sequence $(\tilde{\boldsymbol{\theta}}^{(j)})_{j=1}^J$ denote the gradient descent updates on the following loss function

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}) = \frac{1}{t} \sum_{i \leq t} (\langle \bar{\mathbf{g}}(G_i), \boldsymbol{\theta} - \boldsymbol{\theta}^0 \rangle - y_i)_2^2 + m\lambda \|\boldsymbol{\theta} - \boldsymbol{\theta}^0\|_2^2.$$

Note that $\tilde{\boldsymbol{\theta}}^0 = \boldsymbol{\theta}^0$ and $\tilde{\boldsymbol{\theta}}^{(j)}$ also depends on t the number of data points. We omit the t index, for simplicity of the notation during the proof of this lemma.

By Lemma C.7, $\|\bar{\mathbf{G}}\|_F \leq C\sqrt{TLm}$ and we have,

$$\mathbf{Z} \stackrel{\text{w.h.p}}{\preceq} (\lambda + CL)\mathbf{I} \preceq \frac{1}{m\eta}\mathbf{I}, \quad (\text{C.19})$$

since η is set such that $\eta \leq C(m\lambda + Lm)^{-1}$. Therefore, for any $\mathbf{x} \in \mathbb{R}^p$, $\|\mathbf{x}\|_{\mathbf{Z}} \leq \frac{1}{\sqrt{m\eta}}\|\mathbf{x}\|_2$. Now consider $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^p$, using Eq. (C.19) together with Cauchy-Schwarz implies,

$$\langle \mathbf{x}, \mathbf{x}' \rangle \leq \|\mathbf{x}\|_{\mathbf{Z}}\|\mathbf{x}'\|_{\mathbf{Z}^{-1}} \stackrel{\text{w.h.p}}{\leq} \frac{1}{\sqrt{m\eta}}\|\mathbf{x}\|_2\|\mathbf{x}'\|_{\mathbf{Z}^{-1}}. \quad (\text{C.20})$$

Applying the inequality above, we may write

$$\begin{aligned} \langle \bar{\mathbf{g}}(G), \boldsymbol{\theta}^{(J)} - \boldsymbol{\theta}^0 \rangle &= \langle \bar{\mathbf{g}}(G), \boldsymbol{\theta}^{(J)} - \tilde{\boldsymbol{\theta}}^{(J)} \rangle + \langle \bar{\mathbf{g}}(G), \tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0 \rangle \\ &\stackrel{\text{w.h.p}}{\leq} \frac{1}{\sqrt{m\eta}}\|\bar{\mathbf{g}}(G)\|_{\mathbf{Z}^{-1}}\|\boldsymbol{\theta}^{(J)} - \tilde{\boldsymbol{\theta}}^{(J)}\|_2 + \langle \bar{\mathbf{g}}(G), \tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0 \rangle \\ &\stackrel{\text{w.h.p}}{\leq} 2\left\|\frac{\bar{\mathbf{g}}(G)}{\sqrt{m}}\right\|_{\mathbf{Z}^{-1}}\sqrt{\frac{B}{m\eta\lambda}} + \langle \bar{\mathbf{g}}(G), \tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0 \rangle \end{aligned} \quad (\text{C.21})$$

For the last inequality of Eq. (C.21) we have used Lemma C.8. Decomposing the second term of the right hand side in Eq. (C.21) gives,

$$\begin{aligned} \langle \bar{\mathbf{g}}(G), \tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0 \rangle &= \langle \bar{\mathbf{g}}(G), \frac{\mathbf{Z}\mathbf{b}}{\sqrt{m}} \rangle + \langle \bar{\mathbf{g}}(G), \tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0 - \frac{\mathbf{Z}\mathbf{b}}{\sqrt{m}} \rangle \\ &\stackrel{\text{w.h.p}}{\leq} \frac{\bar{\mathbf{g}}^T(G)\mathbf{Z}\mathbf{b}}{\sqrt{m}} + \frac{1}{\sqrt{\eta}}\left\|\frac{\bar{\mathbf{g}}(G)}{\sqrt{m}}\right\|_{\mathbf{Z}^{-1}}\left\|\tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0 - \frac{\mathbf{Z}\mathbf{b}}{\sqrt{m}}\right\|_2 \\ &\stackrel{\text{w.h.p}}{\leq} \frac{\bar{\mathbf{g}}^T(G)\mathbf{Z}\mathbf{b}}{\sqrt{m}} + \left\|\frac{\bar{\mathbf{g}}(G)}{\sqrt{m}}\right\|_{\mathbf{Z}^{-1}}\sqrt{\frac{2B}{m\eta\lambda}}(1 - \eta m\lambda)^{J/2} \end{aligned} \quad (\text{C.22})$$

where the first inequality is a consequence of Eq. (C.20). The second inequality follows from the convergence of GD on the proxy loss $\tilde{\mathcal{L}}$, given in Lemma C.8. By the definition of posterior mean and variance (Eq. C.4) when the regularization parameter is set to $\lambda \leftarrow t\lambda$ we have,

$$\begin{aligned} \hat{\mu}_t(G) &= \frac{\bar{\mathbf{g}}^T(G)\mathbf{Z}\mathbf{b}}{\sqrt{m}}, \\ \hat{\sigma}_t(G) &= \left\|\frac{\bar{\mathbf{g}}(G)}{\sqrt{m}}\right\|_{\mathbf{Z}^{-1}}. \end{aligned}$$

The final upper bound on $f_{\text{GNN}}(G; \boldsymbol{\theta}^{(J)}) - \hat{\mu}_t(G)$ follows from plugging in Equation C.22 into Equation C.21, and applying Lemma C.10. Similarly, for the lower bound we have,

$$-f_{\text{GNN}}(G; \boldsymbol{\theta}^{(J)}) \stackrel{\text{w.h.p}}{\leq} \langle \bar{\mathbf{g}}(G), \boldsymbol{\theta}^0 - \boldsymbol{\theta}^{(J)} \rangle + \tilde{C}L^3 \left(\frac{B}{m\lambda}\right)^{2/3} \sqrt{m \log m} \quad (\text{C.23})$$

$$\langle \bar{\mathbf{g}}(G), \boldsymbol{\theta}^0 - \tilde{\boldsymbol{\theta}}^{(J)} \rangle \stackrel{\text{w.h.p}}{\leq} -\hat{\mu}_t(G) + \hat{\sigma}_t(G)\sqrt{\frac{2B}{m\eta\lambda}}(1 - \eta m\lambda)^{J/2} \quad (\text{C.24})$$

$$\langle \bar{\mathbf{g}}(G), \boldsymbol{\theta}^0 - \boldsymbol{\theta}^{(J)} \rangle \stackrel{\text{w.h.p}}{\leq} 2\hat{\sigma}_t(G)\sqrt{\frac{B}{m\eta\lambda}} + \langle \bar{\mathbf{g}}(G), \tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0 \rangle \quad (\text{C.25})$$

Where inequality C.23 holds by Lemma C.7, and the next two inequalities are driven similarly to equations C.21 and C.22. The lower bound results by putting together equations C.23-C.25, and this concludes the proof. Note that we are implicitly taking a union bound over the 6 inequalities that all hold with high probability. The conditions of the used lemmas require that m is picked at a $\text{poly}(\log(N/\delta))$ rate, and a constant number of union bounds do not affect this rate. \square

The next lemma gives a confidence interval over members of $\mathcal{H}_{k_{\text{GNN}}}$.

Lemma C.4 (RKHS Confidence Interval from Vakili et al. [40]). *Let $f^* \in \mathcal{H}_k$ with $\|f^*\|_k \leq B$, and the observation noise to be sub-gaussian with parameter σ^2 . Assume $H_t = \{(\mathbf{x}_\tau, y_\tau)\}$ with $(\mathbf{x}_\tau)_{\tau \leq t}$ and $(\varepsilon_\tau)_{\tau \leq t}$ statistically independent. Then for a fixed input \mathbf{x} , with probability greater than $1 - \delta$,*

$$|f^*(\mathbf{x}) - \mu_t(\mathbf{x})| \leq \left(B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{2 \log(2/\delta)}\right) \sigma_t(\mathbf{x}).$$

The following lemma shows that members of $\mathcal{H}_{k_{\text{GNN}}}$ are well described by the first-order Taylor approximation of a GNN around initialization.

Lemma C.5 (Approximation by a linearized GNN). *Let f be a member of $\mathcal{H}_{k_{\text{GNN}}}$ with bounded RKHS norm $\|f\|_{k_{\text{GNN}}} \leq B$. Set $\delta \in (0, 1)$ and let N denote an upper bound on the possible number of nodes for a graph. If $m = \mathcal{O}(L^6 |\mathcal{G}|^4 / \lambda_0^4 \log(|\mathcal{G}|^2 LN / \delta))$, then with probability greater than $1 - \delta$, there exists $\boldsymbol{\theta}^* \in \mathbb{R}^p$ such that for all $G \in \mathcal{G}$*

$$f(G) = \langle \mathbf{g}_{\text{GNN}}(G; \boldsymbol{\theta}^0), \boldsymbol{\theta}^* \rangle, \quad \sqrt{m} \|\boldsymbol{\theta}^*\|_2 \leq \sqrt{2} B.$$

Proof of Lemma C.5. The proof follows the technique for Lemma 5.1 Zhou et al. [45] with some modifications.

From Eq. (C.30) proof of Lemma C.6, for $m = \mathcal{O}(L^6 / \epsilon^4 \log(LN / \delta))$, and for G_i and G_j in the domain,

$$|k_{\text{GNN}}(G_i, G_j) - \mathbf{g}_{\text{GNN}}^T(G_i; \boldsymbol{\theta}^0) \mathbf{g}_{\text{GNN}}(G_j; \boldsymbol{\theta}^0)| \leq (L + 1) \epsilon$$

with probability greater than $1 - \delta$. Let \mathbf{K}_{full} be the GNTK matrix calculated for all $G \in \mathcal{G}$ and $\bar{\mathbf{G}}_{\text{full}} = [\mathbf{g}_{\text{GNN}}^T(G; \boldsymbol{\theta}^0)]_{G \in \mathcal{G}}$. Then applying a union bound over all $G \in \mathcal{G}$, and setting $\delta \leftarrow \delta / |\mathcal{G}|^2$, if $m = \mathcal{O}(L^6 / \epsilon^4 \log(|\mathcal{G}|^2 LN / \delta))$, then

$$\|\mathbf{K}_{\text{full}} - \bar{\mathbf{G}}_{\text{full}}^T \bar{\mathbf{G}}_{\text{full}} / m\|_F \leq |\mathcal{G}| \epsilon$$

with probability greater than $1 - \delta$. Now applying this inequality when $\epsilon = \lambda_0 / 2 |\mathcal{G}|$, we get that if

$$m = \mathcal{O}(L^6 |\mathcal{G}|^4 / \lambda_0^4 \log(|\mathcal{G}|^2 LN / \delta))$$

then $\|\mathbf{K}_{\text{full}} - \bar{\mathbf{G}}_{\text{full}}^T \bar{\mathbf{G}}_{\text{full}} / m\|_F \leq \lambda_0 / 2$, with probability greater than $1 - \delta$. Via the Triangle inequality we get that

$$\bar{\mathbf{G}}_{\text{full}}^T \bar{\mathbf{G}}_{\text{full}} / m \succcurlyeq \mathbf{K}_{\text{full}} - \|\mathbf{K}_{\text{full}} - \bar{\mathbf{G}}_{\text{full}}^T \bar{\mathbf{G}}_{\text{full}} / m\|_F \succcurlyeq \mathbf{K}_{\text{full}} - \frac{\lambda_0}{2} \succcurlyeq \mathbf{K}_{\text{full}} / 2 \succ 0. \quad (\text{C.26})$$

Since $\lambda_0 > 0$, $\bar{\mathbf{G}}_{\text{full}}$ is positive definite and may be decomposed as $\bar{\mathbf{G}}_{\text{full}} = \mathbf{P} \mathbf{A} \mathbf{Q}^T$, where $\mathbf{P} \in \mathbb{R}^{p \times |\mathcal{G}|}$, $\mathbf{P} \in \mathbb{R}^{|\mathcal{G}| \times |\mathcal{G}|}$ are unitary and $\mathbf{A} \succ 0$. Let $\mathbf{f} = [f(G)]_{G \in \mathcal{G}}$ be the vector of function values. We show that $\boldsymbol{\theta}^* = \mathbf{P} \mathbf{A}^{-1} \mathbf{Q}^T \mathbf{f}$ satisfies the statement of the lemma. By definition of $\boldsymbol{\theta}^*$,

$$\bar{\mathbf{G}}_{\text{full}}^T \boldsymbol{\theta}^* = \mathbf{Q} \mathbf{A} \mathbf{P}^T \mathbf{P} \mathbf{A}^{-1} \mathbf{Q}^T \mathbf{f} = \mathbf{f}$$

which implies for all $G \in \mathcal{G}$, $\langle \mathbf{g}_{\text{GNN}}(G; \boldsymbol{\theta}^0), \boldsymbol{\theta}^* \rangle = f_{\text{GNN}}(G)$. As for the norm of $\boldsymbol{\theta}^*$ we may write,

$$\|\boldsymbol{\theta}^*\|_2^2 \leq \mathbf{f}^T \mathbf{Q} \mathbf{A}^{-2} \mathbf{Q}^T \mathbf{f} = \mathbf{f}^T (\bar{\mathbf{G}}_{\text{full}}^T \bar{\mathbf{G}}_{\text{full}})^{-1} \mathbf{f} \leq \frac{2}{m} \mathbf{f}^T \mathbf{K}_{\text{full}}^{-1} \mathbf{f} \leq \frac{2B^2}{m}$$

where the next to last inequality holds due to Eq. (C.26), and the last inequality follows from $\|f\|_{k_{\text{GNN}}}^2 \leq B^2$. \square

We are now ready to present the proof of our main confidence interval bound.

Proof of Theorem 4.2. Consider Lemma C.4, when the kernel function is $\tilde{k}(G, G') = \bar{\mathbf{g}}^T(G) \bar{\mathbf{g}}(G') / m$ and choose the regularization parameter $t\lambda$. Subsequently, the posterior mean and variance after observing t samples will be $\hat{\mu}_t(G)$ and $\hat{\sigma}_t(G)$. Then this lemma states that for $f \in \mathcal{H}_{\tilde{k}}$ with a norm bounded by B , with probability greater than $1 - \delta/2$,

$$|f(G) - \hat{\mu}_t(G)| \leq \hat{\sigma}_t \left(B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{2 \log 4 / \delta} \right).$$

By Lemma C.5, for m large enough the reward function can be written as

$$f^*(G) = \langle \mathbf{g}_{\text{GNN}}(G; \boldsymbol{\theta}^0), \boldsymbol{\theta}^* \rangle, \quad \sqrt{m} \|\boldsymbol{\theta}^*\|_2 \leq \sqrt{2}B.$$

for all $G \in \mathcal{G}$, indicating that $f^* \in \mathcal{H}_{\hat{k}_{\text{GNN}}}$ with $\|f\|_{\hat{k}_{\text{GNN}}} \leq \sqrt{2}B$. Therefore, following Lemma C.4 with probability greater than $1 - \delta/2$,

$$|f^*(G) - \hat{\mu}_t(G)| \leq \hat{\sigma}_t \left(\sqrt{2}B + \frac{\sigma}{\sqrt{\lambda}} \sqrt{2 \log 4/\delta} \right)$$

for some fixed G . Further, Lemma C.3 bounds the difference between $\hat{\mu}_{t-1}(\cdot)$ and $\hat{\mu}_{t-1}(\cdot)$ with probability higher than $1 - \delta/2$. Plugging in $\hat{\mu}_{t-1}$ and $\hat{\sigma}_{t-1}$ gives,

$$\begin{aligned} |\hat{\mu}_{t-1}(G) - f^*(G)| &\leq \hat{\sigma}_t(G) \left[\sqrt{2}B + \frac{\sigma}{\lambda} \sqrt{2 \log 4/\delta} + \sqrt{\frac{2B}{m\eta\lambda}} (\sqrt{2} + (1 - m\eta\lambda)^{J/2}) \right] \\ &\quad + \tilde{C}L^3 \left(\frac{B}{m\lambda} \right)^{2/3} \sqrt{m \log m} \end{aligned}$$

with probability greater than $1 - \delta$. Setting $\delta \leftarrow \delta/|\mathcal{G}|$ and taking a union bound over all $G \in \mathcal{G}$ concludes the proof of Theorem C.2. The informal version of the theorem is achieved by considering that $m = \text{poly}(t)$ and omitting all the terms that are $o(t^{-1})$ with t . \square

C.2 GNN Helper Lemmas

Lemma C.6 (Norm concentration of Gram matrix and GNTK matrix at initialization). *Set $\epsilon > 0$, $\delta \in (0, 1)$, and let $N = |V(G)|$ denote the number of nodes for every graph $G \in \mathcal{G}$. For width $m = \Omega(L^6/\epsilon^4 \log(LNT^2/\delta))$ in Eq. (2), the following holds with probability at least $1 - \delta$:*

$$\|\mathbf{K}_{\text{GNN}} - \bar{\mathbf{G}}^T \bar{\mathbf{G}}/m\|_F \leq T\epsilon.$$

Proof of Lemma C.6. We make use of the connection between the GNN as defined in Eq. (2) and a fully-connected neural network. In particular, let $f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}) : \mathbb{R}^d \rightarrow \mathbb{R}$ be a fully-connected network, with L hidden layers of equal width m , and ReLU activations as defined in Eq. (A.1). Then, we have

$$f_{\text{GNN}}(G; \boldsymbol{\theta}) = \frac{1}{\sqrt{N}} \sum_{v \in V(G)} f_{\text{NN}}(\bar{\mathbf{h}}_v; \boldsymbol{\theta}).$$

Let $\bar{\mathbf{G}}$ be a matrix with T columns where each column $i \in [T]$ contains gradient vector $\mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^0)$. Then, the matrix $\bar{\mathbf{G}}^T \bar{\mathbf{G}}$ is a $T \times T$ matrix and represents the Gram matrix of the network for the parameters $\boldsymbol{\theta}^0$. For all $i, j \leq T$, we have

$$[\bar{\mathbf{G}}^T \bar{\mathbf{G}}]_{i,j} = \frac{1}{N} \sum_{v=1}^N \mathbf{g}^T(\bar{\mathbf{h}}_v) \mathbf{g}(\bar{\mathbf{h}}_{v'}), \quad (\text{C.27})$$

where $\mathbf{g}(\mathbf{x}) = \nabla_{\boldsymbol{\theta}} f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}^0)$ denotes the gradient of f_{NN} at initialization. It follows by the definition of the neural tangent kernel that

$$k_{\text{GNN}}(G_i, G_j) = \frac{1}{N} \sum_{v=1}^N k_{\text{NN}}(\bar{\mathbf{h}}_v, \bar{\mathbf{h}}_{v'}). \quad (\text{C.28})$$

For some fixed $\epsilon > 0$ and $\delta \in (0, 1)$, the result of Arora et al. [2, Theorem 3.1] states that when $m = \Omega(L^6/\epsilon^4 \log(L/\delta))$, the following holds for any \mathbf{x}, \mathbf{x}' with unit norms and probability at least $1 - \delta$:

$$|k_{\text{NN}}(\mathbf{x}, \mathbf{x}') - \mathbf{g}^T(\mathbf{x}) \mathbf{g}(\mathbf{x}')/m| \leq (L+1)\epsilon. \quad (\text{C.29})$$

Next, by using equations C.28 and C.27, and the triangle inequality for any two input graphs G_i, G_j we get

$$|k_{\text{GNN}}(G_i, G_j) - \mathbf{g}_{\text{GNN}}^T(G_i; \boldsymbol{\theta}^0) \mathbf{g}_{\text{GNN}}(G_j; \boldsymbol{\theta}^0)| \leq \frac{1}{N} \sum_{v=1}^N |k_{\text{NN}}(\bar{\mathbf{h}}_v, \bar{\mathbf{h}}_{v'}) - \mathbf{g}^T(\bar{\mathbf{h}}_v) \mathbf{g}(\bar{\mathbf{h}}_{v'})|. \quad (\text{C.30})$$

Then, since $\bar{\mathbf{h}}_v \in \mathbb{S}^{d-1}$ for every node v irrespective of the corresponding graph, we can use Eq. (C.29) together with the union bound over all $(\bar{\mathbf{h}}_{G,u}, \bar{\mathbf{h}}_{G',u})$ pairs and $m = \Omega(L^6/\epsilon^4 \log(LN/\delta))$, to obtain that for any G_i and G_j with probability at least $1 - \delta$:

$$|k_{\text{GNN}}(G_i, G_j) - \mathbf{g}_{\text{GNN}}^T(G_i; \boldsymbol{\theta}^0) \mathbf{g}_{\text{GNN}}(G_j; \boldsymbol{\theta}^0)/m| \leq (L+1)\epsilon.$$

To arrive at the main result we consider the difference in the Frobenius norm:

$$\|\mathbf{K}_{\text{GNN}} - \bar{\mathbf{G}}^T \bar{\mathbf{G}}/m\|_F = \sqrt{\sum_{i,j \leq T} \left(k_{\text{GNN}}(G_i, G_j) - \mathbf{g}_{\text{GNN}}^T(G_i; \boldsymbol{\theta}^0) \mathbf{g}_{\text{GNN}}(G_j; \boldsymbol{\theta}^0)/m \right)^2}.$$

By setting $\epsilon \leftarrow \epsilon/(L+1)$ and again applying the union bound over each (G_i, G_j) pair, for $m = \Omega(L^6/\epsilon^4 \log(LNT^2/\delta))$, the following holds with probability at least $1 - \delta$:

$$\|\mathbf{K}_{\text{GNN}} - \bar{\mathbf{G}}^T \bar{\mathbf{G}}/m\|_F \leq T\epsilon. \quad \square$$

Lemma C.7 (Gradient descent norm bounds). *Consider the fixed set $\{G_i\}_{i \leq t}$ of inputs. Let $\bar{\mathbf{G}} = [\mathbf{g}_{\text{GNN}}^T(G_i; \boldsymbol{\theta}^0)]_{i \leq t}^T$ be the matrix of gradients at initialization and $\bar{\mathbf{G}}^{(j)} = [\mathbf{g}_{\text{GNN}}^T(G_i; \boldsymbol{\theta}^{(j)})]_{i \leq t}^T$. The vector of network outputs after the j -th update is denoted by $\mathbf{f}_{\text{GNN}}^{(j)} = [f_{\text{GNN}}(G_i; \boldsymbol{\theta}^{(j)})]_{i \leq t}$. Assume τ is set such that $\|\boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0\|_2 \leq \tau$ for all $j \leq J$. If $m = \text{poly}(t, L, \lambda^{-1}, \log(N/\delta))$, then with probability greater than $1 - \delta$,*

$$\|\bar{\mathbf{G}}\|_F \leq C_1 \sqrt{tmL} \quad (\text{C.31})$$

$$\|\bar{\mathbf{G}} - \bar{\mathbf{G}}^{(j)}\|_F \leq C_2 \tau^{1/3} L^{7/2} \sqrt{tm \log m} \quad (\text{C.32})$$

$$\|\mathbf{f}_{\text{GNN}}^{(j)} - \bar{\mathbf{G}}^{(j)}(\boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0)\|_2 \leq C_3 \tau^{4/3} L^3 \sqrt{tm \log m} \quad (\text{C.33})$$

for some constants C_1, C_2, C_3 .

Proof of Lemma C.7. We follow the recipe introduced in Zhou et al. [45], and reproduce gradient norm bounds for when the network is a GNN.

From Lemma B.3 Cao and Gu [10], we get $\|\mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^0)\|_2 \leq \bar{C} \sqrt{mL}$ with probability of at least $1 - \delta/3$. By the definition of Frobenius norm, it follows,

$$\|\bar{\mathbf{G}}\|_F \leq \sqrt{t} \max_{i \leq t} C \|\mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^0)\|_2 \stackrel{\text{w.h.p.}}{\leq} C \sqrt{tmL}.$$

For Eq. (C.32) we may write,

$$\begin{aligned} \|\bar{\mathbf{G}}^{(j)} - \bar{\mathbf{G}}\|_F &\leq \sqrt{t} \max_{i \leq t} \|\mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^{(j)}) - \mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^0)\|_2 \\ &\leq \frac{\sqrt{t}}{N} \max_{i \leq t} \sum_{u \in V(G_i)} \|\mathbf{g}(\bar{\mathbf{h}}_u; \boldsymbol{\theta}^{(j)}) - \mathbf{g}(\bar{\mathbf{h}}_u; \boldsymbol{\theta}^0)\|_2 \\ &\stackrel{\text{w.h.p.}}{\leq} \frac{\sqrt{t}}{N} \max_{i \leq t} \sum_{u \in V(G_i)} \tilde{C}_2 \tau^{1/3} L^3 \sqrt{\log m} \|\mathbf{g}(\bar{\mathbf{h}}_u; \boldsymbol{\theta}^0)\|_2 \\ &\stackrel{\text{w.h.p.}}{\leq} C_2 \tau^{1/3} L^{7/2} \sqrt{tm \log m} \end{aligned}$$

with probability greater than $1 - \delta/3$, where the next to last inequality holds by Lemma B.5 Zhou et al. [45] and the last inequality follows directly from Lemma B.6 Zhou et al. [45].

As for Eq. (C.33), by definition of the single-BLOCK GNN and Lemma B.4 Zhou et al. [45],

$$\begin{aligned} \|\mathbf{f}_{\text{GNN}}^{(j)} - \bar{\mathbf{G}}^{(j)}(\boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0)\|_2 &\leq \sqrt{t} \max_{i \leq t} \left| f_{\text{GNN}}(G_i; \boldsymbol{\theta}^{(j)}) - \langle \mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^{(j)}), \boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0 \rangle \right| \\ &\leq \frac{\sqrt{t}}{N} \max_{i \leq t} \sum_{u \in V(G_i)} \left| f(\bar{\mathbf{h}}_u; \boldsymbol{\theta}^{(j)}) - \langle \mathbf{g}(\bar{\mathbf{h}}_u; \boldsymbol{\theta}^{(j)}), \boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0 \rangle \right| \\ &\stackrel{\text{w.h.p.}}{\leq} C_3 \tau^{4/3} L^3 \sqrt{tm \log m} \end{aligned}$$

with probability greater than $1 - \delta/3$. □

Lemma C.8 (Convergence properties of the proxy optimization problem). *Let the sequence $(\tilde{\theta}^{(j)})$ denote the gradient descent updates on the following loss function,*

$$\tilde{\mathcal{L}}(\theta) = \frac{1}{t} \sum_{i \leq t} (\langle \mathbf{g}_{\text{GNN}}(G_i; \theta^0), \theta - \theta^0 \rangle - y_i)_2^2 + m\lambda \|\theta - \theta^0\|_2^2.$$

then if $m = \text{poly}(T, L, B, |\mathcal{G}|, \lambda_0^{-1}, \lambda^{-1}, \log(N/\delta))$ and the learning rate $\eta \leq C(mL + m\lambda)^{-1}$

$$\begin{aligned} \|\tilde{\theta}^{(j)} - \theta^0\|_2 &\leq \sqrt{B/m\lambda}, \\ \|\tilde{\theta}^{(j)} - \theta^0 - (\lambda\mathbf{I} + \bar{\mathbf{G}}^T \bar{\mathbf{G}}/(tm))^{-1} \bar{\mathbf{G}}^T \mathbf{y}/(tm)\|_2 &\leq (1 - \eta m\lambda)^{j/2} \sqrt{2B/m\lambda}. \end{aligned}$$

Proof of Lemma C.8. This lemma adapts Lemma D.8 Kassraie and Krause [25] to our setting. We repeat the proof for the sake of completeness. Note that $\tilde{\mathcal{L}}$ is $m\lambda$ -strongly convex, and $C(mL + m\lambda)$ -smooth, since

$$\nabla^2 \tilde{\mathcal{L}} = \frac{2\bar{\mathbf{G}}^T \bar{\mathbf{G}}}{t} + 2m\lambda\mathbf{I} \leq 2\left(\frac{\|\bar{\mathbf{G}}\|_F^2}{t} + m\lambda\right)\mathbf{I} \leq C(mL + m\lambda)\mathbf{I},$$

where the second inequality follows from Lemma C.7. Strong Convexity of $\tilde{\mathcal{L}}$ guarantees a monotonic decrease of the loss if the learning rate is smaller than the smoothness coefficient inversed. Therefore,

$$\begin{aligned} m\lambda \|\tilde{\theta}^{(j)} - \tilde{\theta}^0\|_2^2 &\leq m\lambda \|\tilde{\theta}^{(j)} - \tilde{\theta}^0\|_2^2 + \frac{1}{t} \|\bar{\mathbf{G}}(\tilde{\theta}^{(j)} - \tilde{\theta}^0) - \mathbf{y}\|_2^2 \\ &\leq m\lambda \|\tilde{\theta}^0 - \tilde{\theta}^0\|_2^2 + \frac{1}{t} \|\bar{\mathbf{G}}(\tilde{\theta}^0 - \tilde{\theta}^0) - \mathbf{y}\|_2^2 \\ &\leq \frac{\|\mathbf{y}\|_2^2}{t} \\ &\leq B \end{aligned}$$

From the RKHS assumption, the true reward is bounded by B and hence the last inequality follows since the size of the training set is t .

Gradient descent on smooth and strongly convex functions converges to optima if the learning rate is smaller than the smoothness coefficient inversed. Under this condition the minima of $\tilde{\mathcal{L}}$ is unique and has the closed form

$$\tilde{\theta}^* = \theta^0 + (\lambda\mathbf{I} + \bar{\mathbf{G}}^T \bar{\mathbf{G}}/(mt))^{-1} \bar{\mathbf{G}}^T \mathbf{y}/(mt)$$

Having set $\eta \leq C(mL + m\lambda)^{-1}$, we get that $\tilde{\theta}^{(j)}$ converges to $\tilde{\theta}^*$ with the following exponential rate,

$$\begin{aligned} \|\tilde{\theta}^{(j)} - \theta^0 - (\lambda\mathbf{I} + \bar{\mathbf{G}}^T \bar{\mathbf{G}}/(mt))^{-1} \bar{\mathbf{G}}^T \mathbf{y}/(mt)\|_2^2 &\leq (1 - \eta m\lambda)^j \frac{2}{m\lambda} (\tilde{\mathcal{L}}(\theta^0) - \tilde{\mathcal{L}}(\tilde{\theta}^*)) \\ &\leq \frac{2(1 - \eta m\lambda)^j}{m\lambda} \frac{\|\mathbf{y}\|_2^2}{t} \\ &\leq \frac{2B(1 - \eta m\lambda)^j}{m\lambda}. \end{aligned}$$

□

Lemma C.9 (Gradient descent parameters bound). *Let the sequence $\theta^{(J)}$ denote the J -th gradient descent update on the GNN loss,*

$$\tilde{\mathcal{L}}(\theta) = \frac{1}{t} \sum_{i \leq t} (f_{\text{GNN}}(G_i; \theta^0) - y_i)_2^2 + m\lambda \|\theta - \theta^0\|_2^2.$$

If $m = \text{poly}(T, L, B, |\mathcal{G}|, \lambda_0^{-1}, \lambda^{-1}, \log(N/\delta))$ and $\eta \leq C(mL + m\lambda)^{-1}$ for some C , then with probability greater than $1 - \delta$,

$$\|\theta^{(J)} - \theta^0\|_2 \leq 2\sqrt{B/m\lambda}.$$

Proof of Lemma C.9. Following Zhou et al. [45] we introduce the sequence $(\tilde{\boldsymbol{\theta}}^{(j)})$ which denotes the gradient descent updates on the following proxy loss,

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}) = \frac{1}{t} \sum_{i \leq t} (\langle \mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^0), \boldsymbol{\theta} - \boldsymbol{\theta}^0 \rangle - y_i)_2^2 + m\lambda \|\boldsymbol{\theta} - \boldsymbol{\theta}^0\|_2^2.$$

By Lemma C.8,

$$\|\tilde{\boldsymbol{\theta}}^{(J)} - \boldsymbol{\theta}^0\|_2 \leq \sqrt{\frac{B}{m\lambda}}$$

It remains to show that $\|\boldsymbol{\theta}^{(j)} - \tilde{\boldsymbol{\theta}}^{(j)}\|_2 \leq \sqrt{B/m\lambda}$, which concludes the proof due to triangle inequality. By writing out the gradient descent updates of the two sequences we get,

$$\begin{aligned} \|\boldsymbol{\theta}^{(j+1)} - \tilde{\boldsymbol{\theta}}^{(j+1)}\|_2 &= \left\| (1 - \eta m\lambda)(\boldsymbol{\theta}^{(j)} - \tilde{\boldsymbol{\theta}}^{(j)}) - \frac{\eta}{t} (\bar{\mathbf{G}}^{(j)} - \bar{\mathbf{G}})^T (\mathbf{f}_{\text{GNN}}^{(j)} - \mathbf{y}) \right. \\ &\quad \left. - \frac{\eta}{t} \bar{\mathbf{G}}^T (\mathbf{f}_{\text{GNN}}^{(j)} - \bar{\mathbf{G}}(\boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0) + \bar{\mathbf{G}}(\boldsymbol{\theta}^{(j)} - \tilde{\boldsymbol{\theta}}^{(j)})) \right\|_2 \\ &\leq \frac{\eta}{t} \|(\bar{\mathbf{G}}^{(j)} - \bar{\mathbf{G}})\|_2 \|\mathbf{f}_{\text{GNN}}^{(j)} - \mathbf{y}\|_2 + \frac{\eta}{t} \|\bar{\mathbf{G}}\|_2 \|\mathbf{f}_{\text{GNN}}^{(j)} - \bar{\mathbf{G}}(\boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0)\|_2 \\ &\quad + \|\mathbf{I} - \eta(m\lambda\mathbf{I} + \bar{\mathbf{G}}\bar{\mathbf{G}}^T/t)\|_2 \|\boldsymbol{\theta}^{(j)} - \tilde{\boldsymbol{\theta}}^{(j)}\|_2 \end{aligned}$$

We bound each term separately. In the rest of the proof, Lemma C.7 is always used with $\tau = \sqrt{B/m\lambda}$. Lemma C.3 Zhou et al. [45] directly holds for f_{GNN} , and states that $\|\mathbf{f}_{\text{GNN}}^{(j)} - \mathbf{y}\|_2 \leq (B+1)\sqrt{t}$. Then Eq. (C.32) Lemma C.7 gives,

$$\frac{\eta}{t} \|(\bar{\mathbf{G}}^{(j)} - \bar{\mathbf{G}})\|_2 \|\mathbf{f}_{\text{GNN}}^{(j)} - \mathbf{y}\|_2 \stackrel{\text{w.h.p.}}{\leq} \eta C_1 \left(\frac{B}{m\lambda}\right)^{1/6} L^{7/2} (B+1) \sqrt{m \log m}.$$

Recall that $f_{\text{GNN}}(\bar{G}; \boldsymbol{\theta}^0) = 0$ by design. Then for the second term, using Equations C.31 and C.33 from Lemma C.7,

$$\frac{\eta}{t} \|\bar{\mathbf{G}}\|_2 \|\mathbf{f}_{\text{GNN}}^{(j)} - \bar{\mathbf{G}}(\boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0)\|_2 \stackrel{\text{w.h.p.}}{\leq} \eta C_2 \left(\frac{B}{m\lambda}\right)^{2/3} L^{7/2} m \sqrt{\log m}.$$

As for the last term, first note that by Eq. (C.31) Lemma C.7, with high probability

$$\eta(m\lambda\mathbf{I} + \bar{\mathbf{G}}\bar{\mathbf{G}}^T/t) \stackrel{\text{w.h.p.}}{\preceq} \eta(m\lambda\mathbf{I} + C_1 m L \mathbf{I}) \preceq \mathbf{I}$$

where the last inequality holds since η is chosen to be small enough. Therefore,

$$\|\mathbf{I} - \eta(m\lambda\mathbf{I} + \bar{\mathbf{G}}\bar{\mathbf{G}}^T/t)\|_2 \|\boldsymbol{\theta}^{(j)} - \tilde{\boldsymbol{\theta}}^{(j)}\|_2 \stackrel{\text{w.h.p.}}{\leq} (1 - \eta m\lambda) \|\boldsymbol{\theta}^{(j)} - \tilde{\boldsymbol{\theta}}^{(j)}\|_2.$$

We put the three terms back together and unroll the recursive inequality. Then if m is picked to be large enough at the above stated polynomial rate,

$$\|\boldsymbol{\theta}^{(j)} - \tilde{\boldsymbol{\theta}}^{(j)}\|_2 \stackrel{\text{w.h.p.}}{\leq} \sqrt{\frac{B}{m\lambda}}.$$

□

The next lemma shows that the first order approximation of a GNN at initialization can still describe the network after it has been trained with gradient descent for J steps.

Lemma C.10 (Taylor approximation of a trained GNN). *If $m = \text{poly}(T, L, B, |\mathcal{G}|, \lambda_0^{-1}, \lambda^{-1}, \log(N/\delta))$ and for some constant C , $\eta = C(mL + m\lambda)^{-1}$, then*

$$\left| f_{\text{GNN}}(G_t; \boldsymbol{\theta}^{(j)}) - f_{\text{GNN}}(G_t; \boldsymbol{\theta}^0) - \langle \mathbf{g}_{\text{GNN}}(G_t; \boldsymbol{\theta}^0), \boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}^0 \rangle \right| \leq \tilde{C} L^3 \left(\frac{B}{m\lambda}\right)^{2/3} \sqrt{m \log m}$$

with some constant \tilde{C} and any $t \leq T$, with probability greater than $1 - \delta$.

Proof of lemma C.10. By Lemma 4.1 Cao and Gu [10], if $m = \text{poly}(T, L, B, |\mathcal{G}|, \lambda_0^{-1}, \lambda^{-1}, \log(1/\delta))$ and η is set according to the statement of the lemma, then for a fixed \mathbf{x} with probability greater than $1 - \delta$,

$$\left| f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}^{(J)}) - f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}^0) - \langle g(\mathbf{x}; \boldsymbol{\theta}^0), \boldsymbol{\theta}^{(J)} - \boldsymbol{\theta}^0 \rangle \right| \leq C\tau^{4/3}L^3\sqrt{m \log m}$$

where $\|\boldsymbol{\theta}^{(J)} - \boldsymbol{\theta}^0\|_2 \leq \tau$. We use this inequality with $\mathbf{x} = (\bar{\mathbf{h}}_v^{(i)})$ for all $u \in V(G_i)$. Setting $\delta \leftarrow \delta/2N$ and applying the union bound gives

$$\left| \sum_{u \in V(G_i)} f_{\text{NN}}(\bar{\mathbf{h}}_v^{(i)}; \boldsymbol{\theta}^{(J)}) - \sum_{u \in V(G_i)} f_{\text{NN}}(\bar{\mathbf{h}}_v^{(i)}; \boldsymbol{\theta}^0) - \left\langle \sum_{u \in V(G_i)} \mathbf{g}(\bar{\mathbf{h}}_v^{(i)}; \boldsymbol{\theta}^0), \boldsymbol{\theta}^{(J)} - \boldsymbol{\theta}^0 \right\rangle \right| \leq CN\tau^{4/3}L^3\sqrt{m \log m}$$

Therefore, if $m = \text{poly}(T, L, B, |\mathcal{G}|, \lambda_0^{-1}, \lambda^{-1}, \log(N/\delta))$ with probability greater than $1 - \delta$,

$$\left| f(G_i; \boldsymbol{\theta}^{(J)}) - f(G_i; \boldsymbol{\theta}^0) - \langle \mathbf{g}_{\text{GNN}}(G_i; \boldsymbol{\theta}^0), \boldsymbol{\theta}^{(J)} - \boldsymbol{\theta}^0 \rangle \right| \leq C\tau^{4/3}L^3\sqrt{m \log m} \quad (\text{C.34})$$

It remains to bound $\|\boldsymbol{\theta}^{(J)} - \boldsymbol{\theta}^0\|_2$ in order to calculate η in Equation C.34. From Lemma C.9, we have $\|\boldsymbol{\theta}^{(J)} - \boldsymbol{\theta}^0\|_2 \leq 2\sqrt{B/m\lambda}$. Setting $\delta \leftarrow \delta/T$ and taking a union bound over all $t \leq T$ concludes the proof. Note that the added $\log T$ term from union bound does not change rate for m , since it is already growing polynomially with T . \square

D Experiments

We include the details of the experiments in Section 5, together with the supplementary plots.

D.1 Synthetic Permutation Invariant Datasets

To test our permutation invariant additive model, we pick the GNTK as the kernel function and create 18 datasets that inherit this structure. As explained in Section 5, each dataset consists of a finite domain of size $\mathcal{G} = 10000$ together with a reward function, both of which are generated randomly. The domains are sets of Erdős-Rényi random graphs, where each graph has N nodes, and between each two nodes there exists an edge with probability p . The node features are i.i.d. d -dimensional standard Gaussian vectors.

For every domain, we sample a random reward function. We use $\text{GP}(0, k_{\text{GNN}})$ as a prior, and sample f from its posterior GP. The posterior is calculated using a small random dataset $(G_i, y_i)_{i \leq 5}$, where y_i are drawn independently from $\mathcal{N}(0, 1)$ and G_i are randomly chosen from $\mathcal{G}_{p, N}$. We choose the posterior GP over the prior as it produces somewhat smoother samples. We note that functions drawn from this GP *do not* reside in \mathcal{H}_{GNN} . Table 2 shows the characteristics of the datasets, which will be released together with the code to generate them from scratch.

	$N = 5$	$N = 20$	$N = 100$
$p = 0.05$	$d \in \{10, 100\}$	$d \in \{10, 100\}$	$d \in \{10\}$
$p = 0.2$	$d \in \{10, 100\}$	$d \in \{10, 100\}$	$d \in \{10, 100\}$
$p = 0.95$	$d \in \{10, 100\}$	$d \in \{10, 100\}$	$d \in \{10, 100\}$

Table 2: Parameters of the synthetic datasets

D.2 Practical Details

The python code to our algorithms, bandit environment, and experiments will be released.

Algorithm. There are some differences between how we utilize the algorithm in practice and the pseudo-code in Algorithm 1. We list these modifications for transparency.

- When calculating $\hat{\sigma}_{t-1}$ we approximate $\hat{\mathbf{K}}_{t-1}$ with its diagonal so that the matrix inversion takes $o(t)$ operations.

- GNN-PE suggests to discard data from previous episodes, so that the decisions are non-adaptive. In practice we keep the history for training the network.
- We set all $T_e = 1$.
- Only from $t \geq T_2 = 80$ we follow Eq. (C.2) and intersect the sets of plausible maximizers. For the first T_2 steps construct them via

$$\mathcal{G}_{e+1} \leftarrow \{G \in \mathcal{G} : f_{\text{GNN}}(G; \boldsymbol{\theta}_e^{(J)}) + \beta_{T_e} \hat{\sigma}_{T_e}(G) \geq \max_{G \in \mathcal{G}} (f_{\text{GNN}}(G; \boldsymbol{\theta}_e^{(J)}) - \beta_{T_e} \hat{\sigma}_{T_e}(G))\}.$$

Network Architecture. We set the width of all architectures to $m = 2048$ and depth to $L = 2$. This combination is picked primarily to keep computations light, while somewhat adhering to the theoretical setup. To calculate $\hat{\sigma}_{t-1}$ we approximate the gram matrix $\mathbf{G}^T \mathbf{G}$ with its diagonal, which gets worse as the number of network parameters grow. The picked values for m and L producing a descriptive network, and allow us to use this diagonal approximation with a negligible cumulative error.

Graph Neural Tangent Kernel To implement this kernel function, we use the NTK class from the Neural Tangents library [33], and sum the base NTK via Eq. (4). This library offers the tangent kernels of every network architecture, however it is unclear how the kernel is derived for a GNN, therefore we use our own expression.

Initialization. We initialize the networks by directly following the definition of Eq. (2). The scaling with $1/\sqrt{m}$ is crucial in activating networks in the lazy regime. If this condition is not met, the confidence sets $[f_{\text{GNN}}(\cdot; \boldsymbol{\theta}) \pm \hat{\sigma}_{t-1}(\cdot)]$ may not be valid, since $\hat{\sigma}_{t-1}$ no longer accurately describes the posterior variance of f_{GNN} .

Training. When analyzing the training dynamics of f_{GNN} , we consider SGD on the ℓ_2 -regularized loss. In practice, however, we train the network with the Adam optimizer [26] from PYTORCH [34], and without weight decay. The learning rate is set to $\eta = 0.001$. We allow $T_0 = 40$ steps of random exploration, to mimic some form of pre-training. The random exploration steps are included in our regret plots. For the first $T_1 = 100$ steps, we train the network from scratch (using the same initialization $\boldsymbol{\theta}_0$) at every step t , as described by the algorithm, and then in batches of $T_B = 20$ just to keep computations light. At every step t we run the Adam optimizer for J_t gradient descent steps, where J_t is calculated via the following stopping criteria

$$\begin{aligned} J_t = \min \quad & J \\ \text{s.t.} \quad & \mathcal{L}(\boldsymbol{\theta}_{t-1}^{(J)}) \leq \mathcal{L}_0 \quad \text{or,} \quad \frac{\Delta \mathcal{L}_{t-1}^{(J-1)} - \Delta \mathcal{L}_{t-1}^{(J)}}{\Delta \mathcal{L}_{t-1}^{(J-1)}} \leq \delta_0 \end{aligned}$$

where we set $\mathcal{L}_0 = 10^{-4}$, $\delta_0 = 10^{-3}$, and

$$\Delta \mathcal{L}_{t-1}^{(J)} := \mathcal{L}(\boldsymbol{\theta}_{t-1}^{(J)}) - \mathcal{L}(\boldsymbol{\theta}_{t-1}^{(J-1)}).$$

The above criterion targets both value of the loss function and the rate at which it is decaying. Effectively, this rule stops training if either the loss is lower than a threshold \mathcal{L}_0 , or if the loss has plateaued, i.e. the relative change in the the loss is lower than a threshold δ_0 . Roughly put, the two conditions on value and decay of the loss, cause the training algorithm to run longer for larger t and prevent over-fitting when t is small. The hyperparameters of the optimizer, i.e., η , δ_0 , \mathcal{L}_0 , T_B , T_0 , and T_1 are selected by hand and not automatically tuned.

D.3 GNN-UCB & NN-UCB

In Section 5, we compare GNN-PE with NN-PE, GNN-UCB and NN-UCB as baselines. The pseudo-code is laid out in Algorithm 3 and Algorithm 4.

Algorithm 3: GNN-UCB

Input: $m, J, \eta, \lambda, \beta_t, T$

Initialize network parameters to a random θ^0 , and $\hat{\mathbf{K}}_0 = \sigma^2 \mathbf{I}$.

for $t = 1 \dots T$ **do**

for $G \in \mathcal{G}$ **do**

$$\begin{aligned} & \hat{\sigma}_{t-1}^2(G) \leftarrow \mathbf{g}_{\text{GNN}}^T(G; \theta^0) \hat{\mathbf{K}}_{t-1}^{-1} \mathbf{g}_{\text{GNN}}(G; \theta^0) / m \\ & U_{G,t} \leftarrow f_{\text{GNN}}(G; \theta_{t-1}^{(J)}) + \beta_t \hat{\sigma}_{t-1}(G) \end{aligned}$$

end

$$G_t = \arg \max_{G \in \mathcal{G}} U_{G,t}$$

Select G_t and append the rewards vector \mathbf{y}_t by the observed reward.

$$\text{Set } \hat{\mathbf{K}}_t \leftarrow \lambda \mathbf{I} + \sum_{i \leq t} \mathbf{g}_{\text{GNN}}(G_i; \theta^0) \mathbf{g}_{\text{GNN}}^T(G_i; \theta^0) / mt$$

$$\text{Calculate } \theta_t^{(J)} = \text{TrainGNN}(m, J, \eta, \lambda, \theta^0, (G_i, y_i)_{i \leq t})$$

end

Algorithm 4: NN-UCB

Input: $m, J, \eta, \lambda, \beta_t, T$

Initialize network parameters to a random θ^0 , and $\hat{\mathbf{K}}_0 = \sigma^2 \mathbf{I}$.

for $t = 1 \dots T$ **do**

for $G \in \mathcal{G}$ **do**

$$\begin{aligned} & \hat{\sigma}_{t-1}^2(G) \leftarrow \mathbf{g}_{\text{NN}}^T(\bar{\mathbf{h}}_G; \theta^0) \hat{\mathbf{K}}_{t-1}^{-1} \mathbf{g}_{\text{NN}}(\bar{\mathbf{h}}_G; \theta^0) / m \\ & U_{G,t} \leftarrow f_{\text{NN}}(\bar{\mathbf{h}}_G; \theta_{t-1}^{(J)}) + \beta_t \hat{\sigma}_{t-1}(G) \end{aligned}$$

end

$$G_t = \arg \max_{G \in \mathcal{G}} U_{G,t}$$

Select G_t and append the rewards vector \mathbf{y}_t by the observed reward.

$$\text{Set } \hat{\mathbf{K}}_t \leftarrow \lambda \mathbf{I} + \sum_{i \leq t} \mathbf{g}_{\text{NN}}(\bar{\mathbf{h}}_{G_i}; \theta^0) \mathbf{g}_{\text{NN}}^T(\bar{\mathbf{h}}_{G_i}; \theta^0) / mt$$

$$\text{Calculate } \theta_t^{(J)} = \text{TrainNN}(m, J, \eta, \lambda, \theta^0, (\bar{\mathbf{h}}_{G_i}, y_i)_{i \leq t})$$

end

Algorithm 5: TrainNN

Input: $m, J, \eta, \lambda, \theta^0, (\bar{\mathbf{h}}_{G_i}, y_i)_{i < t}$

Define $\mathcal{L}(\theta) = \frac{1}{t} \sum_{i < t} (f_{\text{NN}}(\bar{\mathbf{h}}_{G_i}; \theta) - y_i)^2 + m\lambda \|\theta - \theta^0\|_2^2$

Initialize $\theta^{(0)} = \theta^0$

for $j = 1, \dots, J$ **do**

$$\theta^{(j)} = \theta^{(j-1)} - \eta \nabla \mathcal{L}(\theta^{(j-1)})$$

end

Output: $\theta^{(J)}$

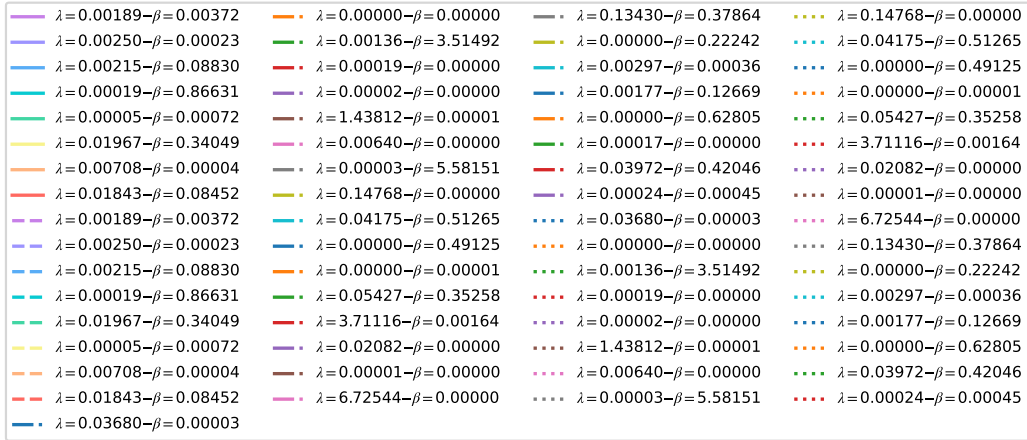



Figure 5: Results of hyper-parameter search for all algorithms. The GNN methods then to perform well for many configurations of λ and β .




Figure 6: Comparing performance of GNN-PE, GNN-UCB, NN-PE, and NN-UCB for all dataset configurations. The GNN methods consistently outperform NN methods. Hyper-parameter tuning is done only for $N = 5, p = 0.05$ and the same is used across all setting.

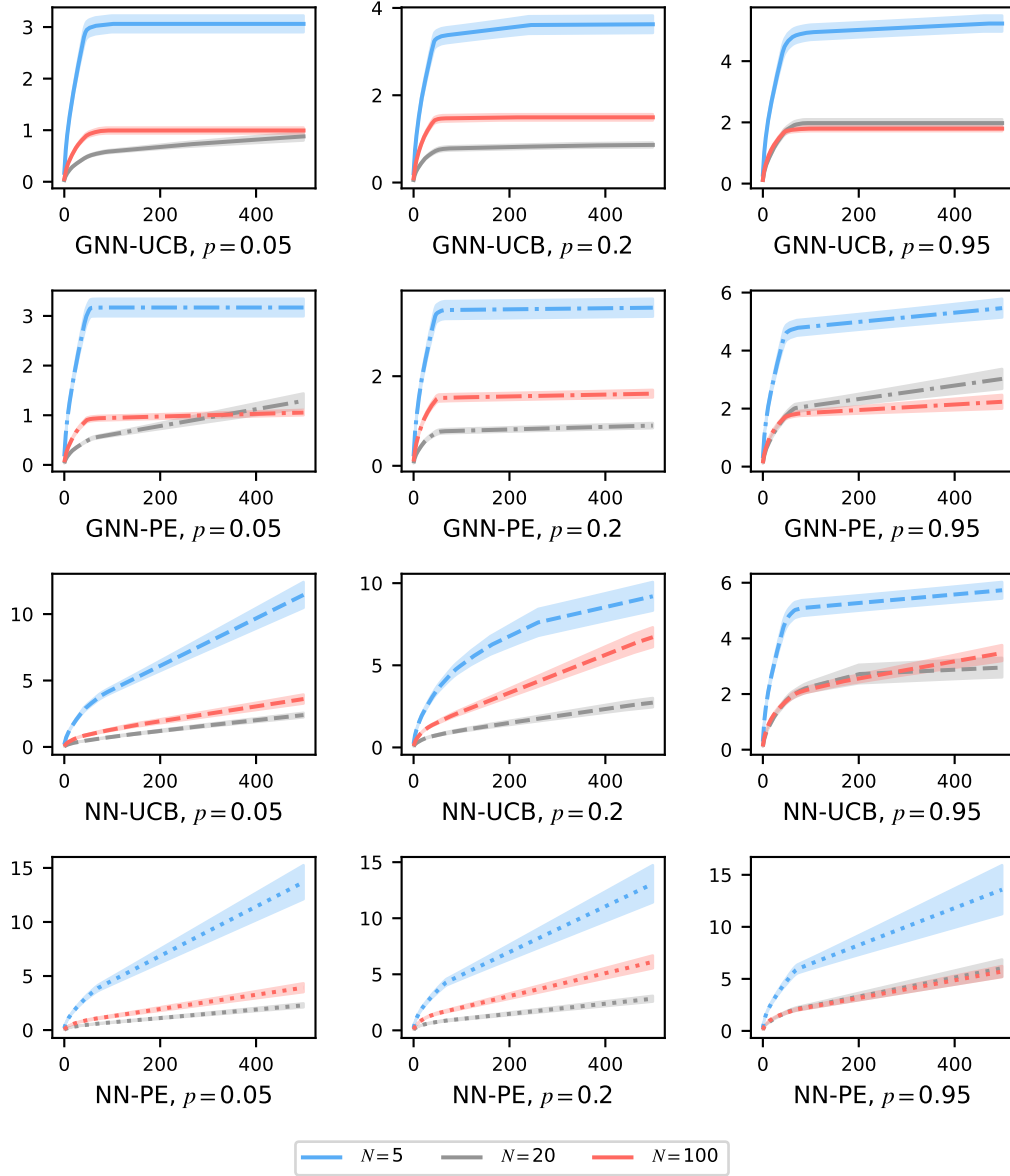


Figure 7: Effect of graph size on performance of GNN-PE, GNN-UCB, NN-PE, and NN-UCB. GNN methods perform well regardless of value of N . Inference on sparse small graphs is challenging since the random graphs tend to have very few edges.




Figure 8: Effect of edge density on performance of GNN-PE, GNN-UCB, NN-PE, and NN-UCB. The NN algorithms tend to improve as p grows.