
Neural Contextual Bandits without Regret

Parnian Kassraie
ETH Zurich
pkassraie@ethz.ch

Andreas Krause
ETH Zurich
krausea@ethz.ch

Abstract

Contextual bandits are a rich model for sequential decision making given side information, with important applications, e.g., in recommender systems. We propose novel algorithms for contextual bandits harnessing neural networks to approximate the unknown reward function. We resolve the open problem of proving sublinear regret bounds in this setting for general context sequences, considering both fully-connected and convolutional networks. To this end, we first analyze NTK-UCB, a kernelized bandit optimization algorithm employing the Neural Tangent Kernel (NTK), and bound its regret in terms of the NTK maximum information gain γ_T , a complexity parameter capturing the difficulty of learning. Our bounds on γ_T for the NTK may be of independent interest. We then introduce our neural network based algorithm NN-UCB, and show that its regret closely tracks that of NTK-UCB. Under broad non-parametric assumptions about the reward function, our approach converges to the optimal policy at a $\tilde{O}(T^{-1/2d})$ rate, where d is the dimension of the context.

1 Introduction

Contextual bandits are a model for sequential decision making based on noisy observations. At every step, the agent is presented with a context vector and picks an action, based on which it receives a noisy reward. Learning about the reward function with few samples (exploration), while simultaneously maximizing its cumulative payoff (exploitation), are the agent’s

two competing objectives. Our goal is to develop an algorithm whose action selection policy attains sublinear regret, which implies convergence to an optimal policy as the number of observations grows. A celebrated approach for regret minimization is the *optimism* principle: establishing upper confidence bounds (UCB) on the reward, and always selecting a plausibly optimal action. Prior work has developed UCB-based contextual bandit approaches, such as linear or kernelized bandits, for increasingly rich models of reward functions (Abbasi-Yadkori et al., 2011; Auer, 2002; Chowdhury and Gopalan, 2017; Srinivas et al., 2010). There are also several recent attempts to harness deep neural networks for contextual bandit tasks. While these perform well in practice (Gu et al., 2021; Riquelme et al., 2018; Nabati et al., 2021; ZHANG et al., 2021; Zhou et al., 2020), there is a lack of theoretical understanding of neural network-based bandit approaches.

We introduce two optimistic contextual bandit algorithms that employ neural networks to estimate the reward and its uncertainty, NN-UCB and its convolutional variant CNN-UCB. Under the assumption that the unknown reward function f resides in a Reproducing Kernel Hilbert Space (RKHS) with a bounded norm, we prove that both algorithms converge to the optimal policy, if the network is sufficiently wide, or has many channels. To prove this bound, we take a two-step approach. We begin by bounding the regret for NTK-UCB, which simply estimates mean and variance of the reward via Gaussian process (GP) inference. Here, the covariance function of the GP is set to k_{NN} , the Neural Tangent Kernel associated with the given architecture. We then exploit the fact that neural networks trained with gradient descent approximate the posterior mean of this GP (Arora et al., 2019), and generalize our analysis of NTK-UCB to bound NN-UCB’s regret. By drawing a connection between fully-connected and 2-layer convolutional networks, we extend our analysis to include CNTK-UCB and CNN-UCB, the convolutional variants of the algorithms. A key contribution of our work is bounding the *NTK maximum information gain*, a parameter that measures the difficulty of learning about a reward

function when it is a sample from $\text{GP}(0, k_{\text{NN}})$. This result may be of independent interest, as this quantity is integral to sequential decision making approaches.

Related Work This work is inspired by Zhou et al. (2020) who introduce the idea of training a neural network within a UCB style algorithm. They analyze NEURAL-UCB, which bears many similarities to NN-UCB. Relevant treatments of the regret are given by Gu et al. (2021); Yang et al. (2020), ZHANG et al. (2021), and Ban and He (2021) for other neural contextual bandit algorithms. However, as discussed in Section 4, these approaches do not generally guarantee sublinear regret, unless further restrictive assumptions about the context are made. In addition, there is a large literature on kernelized contextual bandits. Closely related to our work are Krause and Ong (2011) and Valko et al. (2013) who provide regret bounds for kernelized UCB methods, with Bayesian and Frequentist perspectives respectively. Srinivas et al. (2010) are the first to tackle the kernelized bandit problem with a UCB based method. Many have then proposed variants of this algorithm, or improved its convergence guarantees under a variety of settings (Berkenkamp et al., 2021; Bogunovic et al., 2020; Candriello et al., 2019; Chowdhury and Gopalan, 2017; Djongla et al., 2013; Kandasamy et al., 2019; Mutn y and Krause, 2019; Scarlett, 2018). The majority of the *maximum information gain*, and Srinivas et al. (2010) establish a priori bounds on this parameter. Their analysis only holds for smooth kernel classes, but has since been extended to cover more complex kernels (Janz et al., 2020; Scarlett et al., 2017; Shekhar et al., 2018; Vakili et al., 2021a). In particular, Vakili et al. (2021a) introduce a technique that applies to smooth Mercer kernels, which we use as a basis for our analysis of the NTK’s maximum information gain. In parallel to UCB methods, online decision making via Thompson Sampling is also extensively studied following Russo and Van Roy (2016).

Our work further builds on the literature linking wide neural networks and Neural Tangent Kernels. Cao and Gu (2019) provide important results on training wide fully-connected networks with gradient descent, which we extend to 2-layer convolutional neural networks (CNNs). Through a non-asymptotic bound, Arora et al. (2019) approximate a trained neural network by the posterior mean of a GP with the NTK as its covariance function. Bietti and Bach (2021) study the Mercer decomposition of the NTK and calculate the decay rate of its eigenvalues, which plays an integral role in our analysis. Little is known about the properties of the Convolutional Neural Tangent Kernel (CNTK), and the extent to which it can be used for

approximating trained CNNs. Bietti (2022) and Mei et al. (2021) are among the first to study this kernel by investigating its invariance towards certain groups of transformations, which we draw inspiration from.

Contributions Our main contributions are:

- To our knowledge, we are the first to give an explicit sublinear regret bound for a neural network based UCB algorithm. We show that NN-UCB’s cumulative regret after a total of T steps is $\tilde{O}(T^{(2d-1)/2d})$, for any arbitrary context sequence on the d -dimensional hyper-sphere. (Theorem 4.2)
- We introduce CNN-UCB, and prove that when the number of channels is large enough, it converges to the optimal policy at the same rate as NN-UCB. (Theorem 5.4)

The \tilde{O} notation omits the terms of order $\log T$ or slower. Along the way, we present intermediate results that may be of independent interest. In Theorem 3.1 we prove that γ_T , the maximum information gain for the NTK after T observations, is $\tilde{O}(T^{(d-1)/d})$. We introduce and analyze NTK-UCB and CNTK-UCB, two kernelized methods with sublinear regret (Theorems 3.2 & 3.3) that can be used in practice or as a theoretical tool. Theorems 3.1 through 3.3 may provide an avenue for extending other kernelized algorithms to neural network based methods.

2 Problem Statement

Contextual bandits are a model of sequential decision making over T rounds, where, at step t , the learner observes a context matrix \mathbf{z}_t , and picks an action \mathbf{a}_t from \mathcal{A} , the finite set of actions. The context matrix consists of a set of vectors, one for each action, i.e., $\mathbf{z}_t = (\mathbf{z}_{t,1}, \dots, \mathbf{z}_{t,|\mathcal{A}|}) \in \mathbb{R}^{d \times |\mathcal{A}|}$. The learner then receives a noisy reward $y_t = f(\mathbf{x}_t) + \epsilon_t$. Here, the input to the reward function is the context vector associated with the chosen action, i.e., $\mathbf{x}_t = \mathbf{z}_t \mathbf{a}_t \in \mathbb{R}^d$, where \mathbf{a}_t is represented as a one-hot vector of length $|\mathcal{A}|$. Then the reward function is defined as $f : \mathcal{X} \rightarrow \mathbb{R}$, where $\mathcal{X} \subseteq \mathbb{R}^d$ denotes the input space. Observation noise is modeled with ϵ_t , an i.i.d. sample from a zero-mean sub-Gaussian distribution with variance proxy $\sigma^2 > 0$. The goal is to choose actions that maximize the cumulative reward over T time steps. This is analogous to minimizing the *cumulative regret*, the difference between the maximum possible (context-dependent) reward and the actual reward received, $R_T = \sum_{t=1}^T f(\mathbf{x}_t^*) - f(\mathbf{x}_t)$, where \mathbf{x}_t is the learner’s pick and

\mathbf{x}_t^* is the maximizer of the reward function at step t

$$\mathbf{x}_t^* = \arg \max_{\mathbf{x}=\mathbf{z}_t, \mathbf{a}, \mathbf{a} \in \mathcal{A}} f(\mathbf{x}).$$

The learner’s goal is to select actions such that $R_T/T \rightarrow 0$ as $T \rightarrow \infty$. This property implies that the learner’s policy will converge to the optimal policy.

2.1 Assumptions

Our regret bounds require some assumptions on the reward function f and the input space \mathcal{X} . Throughout this work, we assume that \mathcal{A} is finite and \mathcal{X} is a subset of \mathbb{S}^{d-1} the d -dimensional unit hyper-sphere. We consider two sets of assumptions on f ,

- *Frequentist Setting:* We assume that f is an arbitrary function residing in $\mathcal{H}_{k_{\text{NN}}}$ the RKHS that is reproducing for the NTK, and has a bounded kernel norm, $\|f\|_{k_{\text{NN}}} \leq B$.
- *Bayesian Setting:* We assume that f is a sample from a zero-mean Gaussian process, that uses the NTK as its covariance function, $\text{GP}(0, k_{\text{NN}})$.

These assumptions are broad, non-parametric and imply that f is continuous on the hyper-sphere. Both the Bayesian and the Frequentist setting impose certain smoothness properties on f via k_{NN} . Technically, the function class addressed by each assumption has an empty intersection with the other. Appendix B.1 provides more insight into the connection between the two assumptions. We require a mild *Sufficient Exploration* assumption on the kernel matrix, exclusive to the results in Sections 4 and 5. This is presented later, under Assumption 4.1.

2.2 The Neural Tangent Kernel

We review important properties of the NTK as relied upon in this work. Training very wide neural networks has similarities to estimation with kernel methods using the NTK. For now, we will focus on fully-connected feed-forward ReLU networks and their corresponding NTK. In Section 5, we extend our result to networks with one convolutional layer. 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,

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

The weights $\mathbf{W}^{(i)}$ are initialized to random matrices with standard normal i.i.d. entries, and

$\boldsymbol{\theta} = (\mathbf{W}^{(i)})_{i \leq L+1}$. Let $\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} f(\mathbf{x}; \boldsymbol{\theta})$ be the gradient of f . Assume that given a fixed dataset, the network is trained with gradient descent using an infinitesimally small learning rate. For networks with large width m , training causes little change in the parameters and, respectively, the gradient vector. For any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, and as m tends to infinity, a limiting behavior emerges: $\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}), \mathbf{g}(\mathbf{x}'; \boldsymbol{\theta}) \rangle / m$, the inner product of the gradients, remains constant during training and converges to $k_{\text{NN}}(\mathbf{x}; \mathbf{x}')$, a deterministic kernel function (Arora et al., 2019; Jacot et al., 2018). This kernel satisfies the conditions of Mercer’s Theorem over \mathbb{S}^{d-1} with the uniform measure (Cao et al., 2021) and has the following Mercer decomposition,

$$k_{\text{NN}}(\mathbf{x}, \mathbf{x}') = \sum_{k=0}^{\infty} \mu_k \sum_{j=1}^{N(d,k)} Y_{j,k}(\mathbf{x}) Y_{j,k}(\mathbf{x}'), \quad (1)$$

where $Y_{j,k}$ is the j -th spherical harmonic polynomial of degree k , and $N(d, k)$ denotes the algebraic multiplicity of μ_k . In other words, each μ_k corresponds to a $N(d, k)$ dimensional eigenspace, where $N(d, k)$ grows with k^{d-2} . Without loss of generality, assume that the distinct eigenvalues μ_k are in descending order. Bietti and Bach (2021) show that there exists an absolute constant $C(d, L)$ such that

$$\mu_k \simeq C(d, L) k^{-d}. \quad (2)$$

Taking the algebraic multiplicity into account, we obtain that the decay rate for the complete spectrum of eigenvalues is of polynomial rate $k^{-1-1/(d-1)}$. This decay is slower than that of the kernels commonly used for kernel methods. The eigen-decay for the squared exponential kernel is $O(\exp(-k^{1/d}))$ (Belkin, 2018), and Matérn kernels with smoothness $\nu > 1/2$ have a $O(k^{-1-2\nu/d})$ decay rate (Santin and Schaback, 2016). The RKHS associated with k_{NN} is then given by

$$\begin{aligned} \mathcal{H}_{k_{\text{NN}}} = \left\{ f : f = \sum_{k \geq 0} \sum_{j=1}^{N(d,k)} \beta_{j,k} Y_{j,k}, \right. \\ \left. \sum_{k \geq 0} \sum_{j=1}^{N(d,k)} \frac{\beta_{j,k}^2}{\mu_k} < \infty \right\}. \end{aligned} \quad (3)$$

Equation 3 explains how the eigen-decay of k controls the complexity of \mathcal{H}_k . Only functions whose coefficients $\beta_{j,k}$ decay at a faster rate than the kernel’s eigenvalues are contained in the RKHS. Therefore, if the eigenvalues of k decay rapidly, \mathcal{H}_k is more limited. The slow decay of the NTK’s eigenvalues implies that the assumptions on the reward function given in Section 2.1 are less restrictive compared to what is often studied in the kernelized contextual bandit literature.

3 Warm-up: NTK-UCB – Kernelized Contextual Bandits with the NTK

Our first step will be to analyze kernelized bandit algorithms that employ the NTK as the kernel. In particular, we focus on the *Upper Confidence Bound* (UCB) exploration policy (Srinivas et al., 2010). Kernelized bandits can be interpreted as modeling the reward function f via a Bayesian prior, namely a Gaussian process $\text{GP}(0, k)$ with covariance function given by k . At each step t , we calculate the posterior mean and variance $\mu_{t-1}(\cdot)$ and $\sigma_{t-1}(\cdot)$, using the samples observed at previous steps. For i.i.d. $\mathcal{N}(0, \sigma^2)$ noise, the posterior GP has a closed form expression,

$$\begin{aligned}\mu_{t-1}(\mathbf{x}) &= \mathbf{k}_{t-1}^T(\mathbf{x})(\mathbf{K}_{t-1} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_{t-1} \\ \sigma_{t-1}^2(\mathbf{x}) &= k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_{t-1}^T(\mathbf{x})(\mathbf{K}_{t-1} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_{t-1}(\mathbf{x})\end{aligned}\quad (4)$$

where $\mathbf{y}_{t-1} = [y_i]_{i < t}$ is the vector of received rewards, $\mathbf{k}_{t-1}(\mathbf{x}) = [k(\mathbf{x}, \mathbf{x}_i)]_{i < t}$, and $\mathbf{K}_{t-1} = [k(\mathbf{x}_i, \mathbf{x}_j)]_{i, j < t}$ is the kernel matrix. We then select the action by maximizing the UCB,

$$\mathbf{x}_t = \arg \max_{\mathbf{x}=\mathbf{z}_t \mathbf{a}, \mathbf{a} \in \mathcal{A}} \mu_{t-1}(\mathbf{x}) + \sqrt{\beta_t} \sigma_{t-1}(\mathbf{x}). \quad (5)$$

The acquisition function balances exploring uncertain actions and exploiting the gained information via parameter β_t which will be detailed later. Our method NTK-UCB, adopts the UCB approach, and uses k_{NN} as the covariance kernel function of the GP for calculating the posteriors in Equation 4.

3.1 Information Gain

The UCB policy seeks to learn about f quickly, while picking actions that also give big rewards. The speed at which we learn about f is quantified by the *maximum information gain*. Assume that for a sequence of inputs $X_T = (\mathbf{x}_1, \dots, \mathbf{x}_T)$, the learner observes noisy rewards $\mathbf{y}_T = (y_1, \dots, y_T)$, and let $\mathbf{f}_T = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_T))$ be the corresponding true rewards. Then the information gain is defined as the mutual information between these random vectors, $I(\mathbf{y}_T; \mathbf{f}_T) := H(\mathbf{y}_T) - H(\mathbf{y}_T | \mathbf{f}_T)$, where H denotes the entropy. Assuming the GP prior $f \sim \text{GP}(0, k_{\text{NN}})$, and in the presence of i.i.d. Gaussian noise,

$$I(\mathbf{y}_T; \mathbf{f}_T) = \frac{1}{2} \log \det(\mathbf{I} + \sigma^{-2} \mathbf{K}_T)$$

with the kernel matrix $\mathbf{K}_T = [k_{\text{NN}}(\mathbf{x}_i, \mathbf{x}_j)]_{i, j \leq T}$. Following Srinivas et al. (2010), we will express our regret bounds in terms of the information gain. The information gain depends on the sequence of points observed. To obtain bounds for arbitrary context sequences, we work with the *maximum information gain* defined as

$\gamma_T := \max_{X_T} I(\mathbf{y}_T; \mathbf{f}_T)$. By bounding $I(\mathbf{y}_T; \mathbf{f}_T)$ with γ_T , we obtain regret bounds that are *independent* of the input sequence.

Many regret bounds in this literature, including ours, are of the form $\tilde{\mathcal{O}}(\sqrt{T\gamma_T})$ or $\tilde{\mathcal{O}}(\sqrt{T}\gamma_T)$. For such a bound not to be vacuous, i.e., for it to guarantee convergence to an optimal policy, γ_T must grow strictly sub-linearly with T . Our first main result is an *a priori* bound on γ_T for Neural Tangent Kernels corresponding to fully-connected networks of depth L .

Theorem 3.1. *Suppose the observation noise is i.i.d., zero-mean and a Gaussian of variance $\sigma^2 > 0$, and the input domain $\mathcal{X} \subset \mathbb{S}^{d-1}$. Then the maximum information gain associated with the NTK of a fully-connected ReLU network is bounded by*

$$\gamma_T = \mathcal{O}\left(\left(\frac{C(d,L)T}{\log(1+\frac{T}{\sigma^2})}\right)^{\frac{d-1}{d}} \log\left(1 + \frac{T}{\sigma^2} \left(\frac{C(d,L)T}{\log(1+\frac{T}{\sigma^2})}\right)^{\frac{d-1}{d}}\right)\right)$$

The parameter γ_T arises not only in the bandit setting, but in a broad range of related sequential decision making tasks (Berkenkamp et al., 2021; Kandasamy et al., 2016; Kirschner et al., 2020; Kirschner and Krause, 2019; Sessa et al., 2019, 2020; Sui et al., 2018). Theorem 3.1 might therefore be of independent interest and facilitate the extension of other kernelized algorithms to neural network based methods. When restricted to \mathbb{S}^{d-1} , the growth rate of γ_T for the NTK matches the rate for a Matérn kernel with smoothness coefficient of $\nu = 1/2$, since both kernels have the same rate of eigen-decay (Chen and Xu, 2021). Srinivas et al. (2010) bound γ_T for smooth Matérn kernels with $\nu \geq 1$, and Vakili et al. (2021a) extend this result to $\nu > 1/2$. From this perspective, Theorem 3.1 pushes the previous literature one step further by bounding the information gain of a kernel with the same eigen-decay as a Matérn kernel with $\nu = 1/2$.¹

Proof Idea Beyond classical analyses of γ_T , additional challenges arise when working with the NTK, since it does not have the smoothness properties required in prior works. As a consequence, we directly use the Mercer decomposition of the NTK (Eq. 1) and break it into two terms, one corresponding to a kernel with a finite-dimensional feature map, and a tail sum. We separately bound the information gain caused by each term. From the Matérn perspective, we are able to extend the previous results, in particular due to our treatment of the Mercer decomposition tail sum. An integral element of our approach is a fine-grained analysis of the NTK’s eigenspectrum over the hyper-

¹Under the assumption that $f \sim \text{GP}(0, k)$ with the covariance function a Matérn $\nu = 1/2$, Shekhar et al. (2018) give a *dimension-type* regret bound for a tree-based bandit algorithm. Their analysis however, is not in terms of the information gain, due to the structure of this algorithm.

sphere, given by Bietti and Bach (2021). The complete proof is given in Appendix C.1.

3.2 Regret Bounds

We now proceed with bounding the regret for NTK-UCB, under both Bayesian and Frequentist assumptions, as explained in Section 2.1. Following Krause and Ong (2011), and making adjustments where needed, we obtain a bound for the Bayesian setting.

Theorem 3.2. *Let $\delta \in (0, 1)$ and suppose f is sampled from $\text{GP}(0, k_{\text{NN}})$. Values of f are observed with a zero-mean Gaussian noise of variance σ^2 , and the exploration parameter is set to $\beta_t = 2 \log(|\mathcal{A}|t^2\pi^2/6\delta)$. Then with probability greater than $1 - \delta$, the regret of NTK-UCB satisfies*

$$R_T \leq C \sqrt{T \beta_T \gamma_T}$$

for any $T \geq 1$, where $C := \sqrt{8\sigma^{-2}/\log(1 + \sigma^{-2})}$.

Crucially, this bound holds for *any* sequence of observed contexts, since γ_T is deterministic and only depends on σ , T , the kernel function k_{NN} , and the input domain \mathcal{X} . A key ingredient in the proofs of regret bounds, including Theorem 3.2, is a concentration inequality of the form

$$|f(\mathbf{x}_t) - \mu_{t-1}(\mathbf{x}_t)| \leq \sqrt{\beta_t \sigma_{t-1}(\mathbf{x}_t)}, \quad (6)$$

which holds with high probability for every \mathbf{x}_t and t , conditioned on context-reward pairs from the previous steps. This inequality holds naturally under the GP assumption, since $f(\mathbf{x})$ is obtained directly from Bayesian inference. Setting β_t to grow with $\log t$ satisfies the inequality and results in a $\tilde{O}(\sqrt{T}\gamma_T)$ regret bound. However, additional challenges arise under the RKHS assumption. For Equation 6 to hold in this setting, we need β_t to grow with $\gamma_t \log t$. The regret would then be $\tilde{O}(\sqrt{T}\gamma_T)$ (Chowdhury and Gopalan, 2017). For NTK-UCB, γ_T is $\tilde{O}(T^{(d-1)/d})$, and the $\gamma_T \sqrt{T}$ rate would no longer imply convergence to the optimal policy for $d \geq 2$. To overcome this open problem (Vakili et al., 2021b), we analyze a variant of our algorithm – called the *Sup* variant – that has been successfully applied in the kernelized bandit literature (Auer, 2002; Chu et al., 2011; Li et al., 2010; Valko et al., 2013). A detailed description of the SUPNTK-UCB, along with its pseudo-code and properties is given in Appendix C.3. Here we give a high-level overview of the Sup variant, and how it resolves the large β_t problem. This variant combines NTK-UCB policy with Random Exploration (RE). At NTK-UCB steps, the UCB is calculated only using the context-reward pairs observed in the previous RE steps. Moreover, the rewards received during the RE steps are *statistically independent* conditioned on the input for those steps.

Together with other properties, this allows a choice of β_t that grows with $\log T$. We obtain the following:

Theorem 3.3. *Let $\delta \in [0, 1]$. Suppose f lies in the RKHS of k_{NN} , with $\|f\|_{k_{\text{NN}}} \leq B$. Samples of f are observed with a zero-mean sub-Gaussian noise with variance proxy σ^2 . Then for a constant $\beta_t = 2 \log(2T|\mathcal{A}|/\delta)$, with probability greater than $1 - \delta$, SUPNTK-UCB algorithm satisfies*

$$R(T) = \mathcal{O}\left(\sqrt{T}\left(\sqrt{\gamma_T \sigma^{-2}(\log T)^3 \log(T \log T |\mathcal{A}|/\delta)} + \sigma B\right)\right).$$

The first term corresponds to regret of the random exploration steps, and the second term results from the steps at which the actions were taken by the NTK-UCB policy. Proofs of Theorems 3.2 and 3.3 are given in Appendices C.2 and C.4 respectively. We finish our analysis of the NTK-UCB with the following conclusion, employing our bound on γ_T from Theorem 3.1 in Theorems 3.2 and 3.3.

Corollary 3.4. *Suppose f satisfies either the GP or RKHS assumption. Then the NTK-UCB (resp. its Sup variant) has sublinear regret $R(T) = \tilde{O}(C_{\text{NN}}(d, L)T^{\frac{2d-1}{2d}})$ with high probability. Hereby, $C_{\text{NN}}(d, L)$ is a coefficient depending on the eigen-decay of the NTK.*

4 Main Result: NN-UCB – Neural Contextual Bandits without Regret

Having analyzed NTK-UCB, we now present our neural net based algorithm NN-UCB, which leverages the connections between NN training and GP regression with the NTK. By design, NN-UCB benefits from the favorable properties of the kernel method which helps us with establishing our regret bound. NN-UCB results from approximating the posterior mean and variance functions appearing in the UCB criterion (Equation 5). First, we approximate the posterior mean μ_{t-1} with $f^{(J)} = f(\mathbf{x}; \boldsymbol{\theta}^{(J)})$, the neural network trained for J steps of gradient descent with some learning rate η with respect to the regularized LSE loss

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^{t-1} (f(\mathbf{x}_i; \boldsymbol{\theta}) - y_i)^2 + m\sigma^2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^0\|_2^2, \quad (7)$$

where m is the width of the network and $\boldsymbol{\theta}^0$ denotes the network parameters at initialization. This choice is motivated by Arora et al. (2019), who show pointwise convergence of the solution of gradient descent on the unregularized LSE loss, to the GP posterior mean when there is no observation noise. We adapt

their result to our setting where we consider ℓ_2 regularized loss and noisy samples. It remains to approximate the posterior variance. Recall from Section 2.2 that the NTK is the limit of $\langle \mathbf{g}(\mathbf{x}), \mathbf{g}(\mathbf{x}') \rangle / m$ as $m \rightarrow \infty$, where $\mathbf{g}(\cdot)$ is the gradient of the network at initialization. This property hints that for a wide network, \mathbf{g} / \sqrt{m} can be viewed as substitute for ϕ , the infinite-dimensional feature map of the NTK, since $k_{\text{NN}}(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$. By re-writing σ_{t-1} in terms of ϕ and substituting ϕ with \mathbf{g} / \sqrt{m} , we get

$$\hat{\sigma}_{t-1}^2(\mathbf{x}) = \frac{\mathbf{g}^T(\mathbf{x})}{\sqrt{m}} \left(\sigma^2 \mathbf{I} + \sum_{i=1}^{t-1} \frac{\mathbf{g}^T(\mathbf{x}_i) \mathbf{g}(\mathbf{x}_i)}{m} \right)^{-1} \frac{\mathbf{g}(\mathbf{x})}{\sqrt{m}}.$$

At the beginning, NN-UCB initializes the network parameters to $\boldsymbol{\theta}^0$. Then at step t , $\hat{\sigma}_{t-1}(\cdot)$ is calculated using $\mathbf{g}(\cdot, \boldsymbol{\theta}^0)$ and the action is chosen via maximizing the *approximate UCB*

$$\mathbf{x}_t = \arg \max_{\mathbf{x} = \mathbf{z}_t \mathbf{a}, \mathbf{a} \in \mathcal{A}} f^{(J)}(\mathbf{x}; \boldsymbol{\theta}_{t-1}) + \sqrt{\beta_t} \hat{\sigma}_{t-1}(\mathbf{x})$$

where $\boldsymbol{\theta}_{t-1}$ is obtained by training $f(\cdot; \boldsymbol{\theta}^0)$, for J steps, with gradient descent on the data observed so far. This algorithm essentially trains a neural network for estimating the reward and combines it with a random feature model for estimating the variance of the reward. These random features arise from the gradient of a neural network with random Gaussian parameters. The pseudo-code to NN-UCB is given in Appendix D. In Appendix A.3, we assess the ability of the approximate UCB criterion to quantify uncertainty in the reward via experiments on the MNIST dataset.

Regret Bound Similar to Theorem 3.3, we make the RKHS assumption on f and establish a regret bound on the Sup variant of NN-UCB. To do so, we need two further technical assumptions. Following Zhou et al. (2020), for convenience we assume that $f(\mathbf{x}; \boldsymbol{\theta}^0) = 0$, for any $\mathbf{x} = \mathbf{z}_t \mathbf{a}$ where $t \leq T$ and $\mathbf{a} \in \mathcal{A}$. As explained in Appendix B.2, this requirement can be fulfilled without loss of generality.

Assumption 4.1 (Sufficient Exploration). The kernel matrix is bounded away from zero, i.e., $\lambda_0 \mathbf{I} \preceq \mathbf{K}_T$.

This assumption is common within the literature (Arora et al., 2019; Cao and Gu, 2019; Du et al., 2019; Zhou et al., 2020) and is satisfied as long as the learner sufficiently explores the input space, such that no two inputs \mathbf{x}_t and $\mathbf{x}_{t'}$ are identical. Further, a weaker version of it is often required to hold for the kernel matrix in the sparse linear bandits literature (Bastani and Bayati, 2020; Hao et al., 2020; Kim and Paik, 2019).

Theorem 4.2. *Let $\delta \in (1, 0)$. Suppose f lies in the RKHS of k_{NN} with $\|f\|_{k_{\text{NN}}} \leq B$. Samples of f are observed with zero-mean sub-Gaussian noise of variance*

proxy σ^2 . Set $J > 1$ and $\beta_t = 2 \log(2T|\mathcal{A}|/\delta)$ constant. Choose the width such that

$$m \geq \text{poly}(T, L, |\mathcal{A}|, \sigma^{-2}, B, \lambda_0^{-1}, \log(1/\delta)),$$

and $\eta = C(LmT + m\sigma^2)^{-1}$ with some universal constant C . Then, with probability greater than $1 - \delta$, the regret of SUPNN-UCB satisfies

$$R(T) = \mathcal{O} \left(\sqrt{T} \left(\sqrt{\gamma_T \sigma^{-2} (\log T)^3 \log(T \log T |\mathcal{A}| / \delta)} + \sigma B \right) \right).$$

The pseudo-code of SUPNN-UCB and the proof are given in Appendix D. The key idea there is to show that given samples with noisy rewards, members of \mathcal{H}_k are well estimated by the solution of gradient descent on the ℓ_2 regularized LSE loss. The following lemma captures this statement.

Lemma 4.3 (Concentration of f and $f^{(J)}$, simplified). *Consider the setting of Theorem 4.2 and further assume that the rewards $\{y_i\}_{i < t}$ are independent conditioned on the contexts $\{\mathbf{x}_i\}_{i < t}$. Let $0 < \delta < 1$ and set m , β_t and η according to Theorem 4.2. Then, with probability greater than $1 - 2e^{-\beta_T/2} - \delta$,*

$$|f^{(J)}(\mathbf{x}_t) - f(\mathbf{x}_t)| \leq \hat{\sigma}_{t-1}(\mathbf{x}_t) \sqrt{\beta_T} \text{Poly}(B, m, t, L, \eta)$$

Theorem 4.2 shows that SUPNN-UCB obeys the same regret guarantee as SUPNTK-UCB. In the theorem, the asymptotic growth of the regret is given for large enough m , and terms that are $o(1)$ with T are neglected. To compare the two algorithms in more detail, we revisit the bound for a fixed m . With a probability greater than $1 - \delta$,

$$\begin{aligned} R_T \leq & \mathcal{O} \left(\sqrt{T} \gamma_T \sqrt{\sigma^{-2} (\log T)^3 \log(T \log T |\mathcal{A}| / \delta)} \right. \\ & + \left(1 + \sigma \sqrt{(m \log(T \log T |\mathcal{A}| / \delta))^{-1}} \right) \sigma B \sqrt{T} \\ & + L^3 \left(\frac{TB}{m\sigma^2} \right)^{5/3} \sqrt{m^3 \log m} \\ & \left. + \frac{\sqrt{B} (1 - m\eta\sigma^2)^{J/2}}{\sqrt{m\eta \log(T \log T |\mathcal{A}| / \delta)}} \right). \end{aligned}$$

The last two terms, which vanish for sufficiently large m , convey the error of approximating GP inference with NN training: The fourth term is the gradient descent optimization error, and the third term is a consequence of working with the linear first order Taylor approximation of $f(\mathbf{x}; \boldsymbol{\theta})$. The first two terms, however, come from selecting explorative actions, as in the regret bound of NTK-UCB (Theorem 3.3). The first term denotes regret from random exploration steps, and the second presents the regret at the steps for which the UCB policy is used to pick actions.

Comparison with Prior Work The NEURAL-UCB algorithm introduced by Zhou et al. (2020) bears resemblance to our method. At step t , NN-UCB approximates the posterior variance via Equation 7 with $\mathbf{g}(\cdot; \boldsymbol{\theta}^0)$, a fixed feature map. NEURAL-UCB, however, updates the feature map at every step t , by using $\mathbf{g}(\cdot; \boldsymbol{\theta}_{t-1})$, where $\boldsymbol{\theta}_{t-1}$ is defined as before. Effectively, Zhou et al. adopt a GP prior that changes with t . Under additional assumptions on f and for $\sigma \geq \max\{1, B^{-1}\}$, they show that for NEURAL-UCB, a guarantee of the following form holds with probability greater than $1 - \delta$.

$$R_T \leq \tilde{O}\left(\sqrt{TI(\mathbf{y}_T; \mathbf{f}_T)}\left[\sigma\sqrt{I(\mathbf{y}_T; \mathbf{f}_T)} + 1 - \log \delta\right] + \sqrt{T}\left(\sigma + \frac{TL}{\sigma}\right)\left(1 - \frac{\sigma^2}{TL}\right)^{J/2} + \sigma B\right)$$

The bound above is *data-dependent* via $I(\mathbf{y}_T; \mathbf{f}_T)$ and in this setting, the only known way of bounding the information gain is through γ_T . The treatment of regret given in Yang et al. (2020) and ZHANG et al. (2021) also results in a bound of the form $\tilde{O}(\sqrt{T}\gamma_T)$. However, the maximum information gain itself grows as $\tilde{O}(T^{(d-1)/d})$ for the NTK covariance function. Therefore, without further assumptions on the sequence of contexts, the above bounds are vacuous. In contrast, our regret bounds for NN-UCB are *sublinear without any further restrictions on the context sequence*. This follows from Theorem 3.1 and Theorem 4.2:

Corollary 4.4. *Under the conditions of Theorem 4.2, for arbitrary sequences of contexts, with probability greater than $1 - \delta$, SUPNN-UCB satisfies,*

$$R(T) = \tilde{O}(C_{\text{NN}}(d, L)T^{\frac{2d-1}{2d}}).$$

The coefficient $C_{\text{NN}}(d, L)$ in Corollary 3.4 and 4.4 denotes the same constant. Figures 1 and 2 in Appendix A plot the information gain and regret obtained for NN-UCB when used on the task of online MNIST classification.

5 Extensions to Convolutional Neural Networks

So far, regret bounds for contextual bandits based on convolutional neural networks have remained elusive. Below, we extend our results to a particular case of 2-layer convolutional networks. Consider a cyclic shift c_l that maps \mathbf{x} to $c_l \cdot \mathbf{x} = (x_{l+1}, x_{l+2}, \dots, x_d, x_1, \dots, x_l)$. We can write a 2-layer CNN, with one convolutional and one fully-connected layer, as a 2-layer NN that is

averaged over all cyclic shifts of the input

$$\begin{aligned} f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta}) &= \sqrt{2} \sum_{i=1}^m v_i \left[\frac{1}{d} \sum_{l=1}^d \sigma_{\text{relu}}(\langle \mathbf{w}_i, c_l \cdot \mathbf{x} \rangle) \right] \\ &= \frac{1}{d} \sum_{l=1}^d f_{\text{NN}}(c_l \cdot \mathbf{x}; \boldsymbol{\theta}). \end{aligned}$$

Let \mathcal{C}_d denote the group of cyclic shifts $\{c_l\}_{l \leq d}$. Then the 2-layer CNN is \mathcal{C}_d -invariant, i.e., $f_{\text{CNN}}(c_l \cdot \mathbf{x}) = f_{\text{CNN}}(\mathbf{x})$, for every c_l . The corresponding CNTK is also \mathcal{C}_d -invariant and can be viewed as an averaged NTK

$$\begin{aligned} k_{\text{CNN}}(\mathbf{x}, \mathbf{x}') &= \frac{1}{d^2} \sum_{l, l'=1}^d k_{\text{NN}}(c_l \cdot \mathbf{x}, c_{l'} \cdot \mathbf{x}') \\ &= \frac{1}{d} \sum_{l=1}^d k_{\text{NN}}(\mathbf{x}, c_l \cdot \mathbf{x}'). \end{aligned} \tag{8}$$

The second equality holds because $k_{\text{NN}}(\mathbf{x}, \mathbf{x}')$ depends on its arguments only through the angle between them. In Appendix B.3, we give more intuition about this equation via the random feature kernel formulation (Chizat et al., 2019; Rahimi and Recht, 2008). Equation 8 implies that the CNTK is a Mercer kernel and in Lemma 5.1 we give its Mercer decomposition. The proof is presented in Appendix B.4.

Lemma 5.1. *The Convolutional Neural Tangent Kernel corresponding to $f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta})$, a 2-layer CNN with standard Gaussian weights, can be decomposed as*

$$k_{\text{CNN}}(\mathbf{x}, \mathbf{x}') = \sum_{k=0}^{\infty} \mu_k \sum_{j=1}^{\bar{N}(d,k)} \bar{Y}_{j,k}(\mathbf{x}) \bar{Y}_{j,k}(\mathbf{x}')$$

where $\mu_k \simeq C(d, L)k^{-d}$. The algebraic multiplicity is $\bar{N}(d, k) \simeq N(d, k)/d$, and the eigenfunctions $\{\bar{Y}_{j,k}\}_{j \leq \bar{N}(d,k)}$ form an orthonormal basis for the space of \mathcal{C}_d -invariant degree- k polynomials on \mathbb{S}^{d-1} .

With this lemma, we show that the 2-layer CNTK has the same distinct eigenvalues as the NTK, while the eigenfunctions and the algebraic multiplicity of each distinct eigenvalue change. The eigenspaces of the NTK are degree- k polynomials, while for the CNTK, they shrink to \mathcal{C}_d -invariant degree- k polynomials. This reduction in the dimensionality of eigenspaces results in a smaller algebraic multiplicity for each distinct eigenvalue.

Information Gain We begin by bounding the maximum information gain $\tilde{\gamma}_T$, when the reward function is assumed to be a sample from $\text{GP}(0, k_{\text{CNN}})$. Proposition 5.2 establishes that the growth rate of $\tilde{\gamma}_T$ matches our result for maximum information gain of

the NTK. The dependence on d however, is improved by a factor of $d^{(d-1)/d}$, indicating that the speed of learning about the reward function is potentially $d^{(d-1)/d}$ times faster for methods that use a CNN. The proof is given in Appendix C.1.

Proposition 5.2. Suppose the input domain satisfies $\mathcal{X} \subset \mathbb{S}^{d-1}$, and samples of f are observed with i.i.d. zero-mean sub-Gaussian noise of variance proxy $\sigma^2 > 0$. Then the maximum information gain of the convolutional neural tangent kernel k_{CNN} satisfies

$$\bar{\gamma}_T = \mathcal{O}\left(\left(\frac{TC(d,L)}{d \log(1+\frac{T}{\sigma^2})}\right)^{\frac{d-1}{d}} \log\left(1 + \frac{T}{\sigma^2} \left(\frac{TC(d,L)}{d \log(1+\frac{T}{\sigma^2})}\right)^{\frac{d-1}{d}}\right)\right).$$

CNTK-UCB This algorithm is set to be the convolutional variant of NTK-UCB. We take the UCB policy (Equation 5) and plug in k_{CNN} as the covariance function for calculating the posterior mean and variance. By Lemma 5.1, the rate of eigen-decay, and therefore, the smoothness properties, are identical between NTK and the 2-Layer CNTK. Through this correspondence, Theorems 3.2 and 3.3 carry over to CNTK-UCB, thus guaranteeing sublinear regret.

Corollary 5.3. *It follows from Theorem 3.2, Theorem 3.3, and Proposition 5.2, that when f satisfies either the GP or the RKHS assumption, CNTK-UCB (resp. its Sup variant) has a sublinear regret*

$$R(T) = \tilde{\mathcal{O}}\left(\frac{C_{\text{NN}}(d, L)}{d^{(d-1)/2d}} T^{\frac{2d-1}{2d}}\right)$$

with high probability. Here $C_{\text{NN}}(d, L)$ is a coefficient depending on the eigen-decay of the NTK.

Corollary 5.3 implies that while regret for both algorithms grows at the same rate with T , CNTK-UCB potentially outperforms NTK-UCB. This claim is investigated on the online MNIST classification task in Figure 3 in Appendix A.

CNN-UCB Here we adopt the structure of NN-UCB and replace the previously used deep fully-connected network by a 2-layer convolutional network. We show that under the same setting as in Theorem 4.2, for each T , there exists a 2-layer CNN with a sufficiently large number of channels, such that the Sup variant of CNN-UCB satisfies the same $\mathcal{O}(T^{(2d-1)/2d})$ regret rate.

Theorem 5.4. *Let $\delta \in (1, 0)$. Suppose f lies in the RKHS of k_{CNN} with $\|f\|_{k_{\text{CNN}}} \leq B$. Set $J > 1$ and $\beta_t = 2 \log(2T|\mathcal{A}|/\delta)$ constant. For any $T \geq 1$, there exists m such that if $\eta = C(LmT + m\sigma^2)^{-1}$ with some universal constant C , then with probability greater than $1 - \delta$, SUPCNN-UCB satisfies,*

$$R(T) = \mathcal{O}\left(\sqrt{T} \left(\sqrt{\gamma_T \sigma^{-2} (\log T)^3 \log(T \log T |\mathcal{A}|/\delta)} + \sigma B\right)\right).$$

The main ingredient in the proof of Theorem 5.4 is a convolutional variant of Lemma 4.3 (Lemma D.14). We prove that a 2-layer CNN trained with gradient descent on the ℓ_2 regularized loss can approximate the posterior mean of a GP with CNTK covariance, calculated from noisy rewards. To this end, we show that training $f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta})$ with gradient descent causes a small change in the network parameters $\boldsymbol{\theta}$ and the gradient vector $\nabla_{\boldsymbol{\theta}} f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta})$. Appendix D.2 presents the complete proof. Comparing Theorem 4.2 and Theorem 5.4, the assumption on m is milder in the former, but the regrets for both algorithms grow at the same rate with T . We do not expect this rate to hold for convolutional networks whose depth is greater than two. In particular, deeper CNNs are no longer \mathcal{C}_d -invariant, and our analysis for CNN-UCB relies on this property for translating the results from the fully-connected setting to the convolutional case. Although the growth rate with T has remained the same, the coefficients in the regret bounds are improved for the convolutional counterparts by a factor of $d^{(d-1)/2d}$. The next corollary presents this observation.

Corollary 5.5. *Under the RKHS assumption and provided that the CNN used in SUPCNN-UCB has enough channels, with high probability, this algorithm satisfies*

$$R(T) = \tilde{\mathcal{O}}\left(\frac{C_{\text{NN}}(d, L)}{d^{(d-1)/2d}} T^{\frac{2d-1}{2d}}\right).$$

Corollary 5.5 establishes the first sublinear regret bound for convolutional contextual bandits. Concurrent to our work, Ban and He (2021) give a $\mathcal{O}(\gamma_T \sqrt{T})$ regret bound for a UCB-inspired algorithm that employs a CNN which has differentiable activation functions with Lipschitz derivatives. This bound suffers from an issue similar to the other works concerning neural contextual bandits, as the growth rate of γ_T for the smooth CNN is not taken into account.

6 Conclusion

We proposed NN-UCB, a UCB based method for contextual bandits when the context is rich or the reward function is complex. Under the RKHS assumption on the reward, and for any arbitrary sequence of contexts, we showed that the regret $R(T)$ grows sub-linearly as $\tilde{\mathcal{O}}(T^{(2d-1)/2d})$, implying convergence to the optimal policy. We extended this result to CNN-UCB, a variant of NN-UCB that uses a 2-layer CNN in place of the deep fully-connected network, yielding the first regret bound for convolutional neural contextual bandits. Our approach analyzed regret for neural network based UCB algorithms through the lens of their respective kernelized methods. A key element in this

approach is bounding the regret in terms of the maximum information gain γ_T . Importantly, we showed that γ_T for both the NTK and the 2-layer CNTK is bounded by $\tilde{O}(T^{(d-1)/d})$, a result that may be of independent interest. We believe our work opens up further avenues towards extending kernelized methods for sequential decision making in a principled way to approaches harnessing neural networks.

Acknowledgments

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. Moreover, we thank Guillaume Wang, Johannes Kirschner, Ya-Ping Hsieh, and Ilija Bogunovic for their valuable comments and feedback.

References

- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Improved algorithms for linear stochastic bandits. In *NIPS*, volume 11, pages 2312–2320, 2011.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pages 8141–8150, 2019.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Yikun Ban and Jingrui He. Convolutional neural bandit: Provable algorithm for visual-aware advertising. *arXiv preprint arXiv:2107.07438*, 2021.
- Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294, 2020.
- Mikhail Belkin. Approximation beats concentration? an approximation view on inference with smooth radial kernels. In *Conference On Learning Theory*, pages 1348–1361. PMLR, 2018.
- Felix Berkenkamp, Andreas Krause, and Angela P Schoellig. Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *Machine Learning*, pages 1–35, 2021.
- Alberto Bietti. Approximation and learning with deep convolutional models: a kernel perspective. In *International Conference on Learning Representations*, 2022.
- Alberto Bietti and Francis Bach. Deep Equals Shallow for ReLU Networks in Kernel Regimes. In *ICLR 2021 - International Conference on Learning Representations*, pages 1–22, Virtual, Austria, 2021. URL <https://hal.inria.fr/hal-02963250>.
- Ilija Bogunovic, Andreas Krause, and Jonathan Scarlett. Corruption-tolerant gaussian process bandit optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 1071–1081. PMLR, 2020.
- Daniele Calandriello, Luigi Carratino, Alessandro Lazaric, Michal Valko, and Lorenzo Rosasco. Gaussian process optimization with adaptive sketching: Scalable and no regret. In *Conference on Learning Theory*, pages 533–557. PMLR, 2019.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, pages 10836–10846, 2019.
- Yuan Cao, Zhiying Fang, Yue Wu, Ding-Xuan Zhou, and Quanquan Gu. Towards understanding the spectral bias of deep learning. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 2205–2211. International Joint Conferences on Artificial Intelligence Organization, 2021.
- Lin Chen and Sheng Xu. Deep neural tangent kernel and laplace kernel have the same {rkhs}. In *International Conference on Learning Representations*, 2021.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alche-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pages 844–853. PMLR, 2017.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.
- Josip Djolonga, Andreas Krause, and Volkan Cevher. High-dimensional gaussian process bandits. In *Neural Information Processing Systems*, 2013.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685. PMLR, 2019.

- Quanquan Gu, Amin Karbasi, Khashayar Khosravi, Vahab Mirrokni, and Dongruo Zhou. Batched neural bandits. *arXiv preprint arXiv:2102.13028*, 2021.
- Botao Hao, Tor Lattimore, and Mengdi Wang. High-dimensional sparse linear bandits. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10753–10763. Curran Associates, Inc., 2020.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- David Janz, David Burt, and Javier González. Bandit optimisation of functions in the matern kernel rkhs. In *International Conference on Artificial Intelligence and Statistics*, pages 2486–2495. PMLR, 2020.
- Kirthevasan Kandasamy, Gautam Dasarathy, Junier Oliva, Jeff Schneider, and Barnabás Póczos. Gaussian process optimisation with multi-fidelity evaluations. In *Proceedings of the 30th/International Conference on Advances in Neural Information Processing Systems (NIPS’30)*, 2016.
- Kirthevasan Kandasamy, Gautam Dasarathy, Junier Oliva, Jeff Schneider, and Barnabas Poczsoz. Multi-fidelity gaussian process bandit optimisation. *Journal of Artificial Intelligence Research*, 66:151–196, 2019.
- Gi-Soo Kim and Myunghee Cho Paik. Doubly-robust lasso bandit. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Johannes Kirschner and Andreas Krause. Stochastic bandits with context distributions. *Advances in Neural Information Processing Systems*, 32:14113–14122, 2019.
- Johannes Kirschner, Ilija Bogunovic, Stefanie Jegelka, and Andreas Krause. Distributionally robust bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 2174–2184. PMLR, 2020.
- Andreas Krause and Cheng S Ong. Contextual gaussian process bandit optimization. In *Advances in neural information processing systems*, pages 2447–2455, 2011.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Learning with invariances in random features and kernel models. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 3351–3418. PMLR, 15–19 Aug 2021.
- Mojmír Mutný and Andreas Krause. Efficient high dimensional bayesian optimization with additivity and quadrature fourier features. *Advances in Neural Information Processing Systems 31*, pages 9005–9016, 2019.
- Ofir Nabati, Tom Zahavy, and Shie Mannor. Online limited memory neural-linear bandits with likelihood matching. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 7905–7915. PMLR, 18–24 Jul 2021.
- 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.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2008.
- Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. In *International Conference on Learning Representations*, 2018.
- Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of thompson sampling. *The Jour-*

- nal of Machine Learning Research*, 17(1):2442–2471, 2016.
- Gabriele Santin and Robert Schaback. Approximation of eigenfunctions in kernel-based spaces. *Advances in Computational Mathematics*, 42(4):973–993, 2016.
- Jonathan Scarlett. Tight regret bounds for bayesian optimization in one dimension. In *International Conference on Machine Learning*, pages 4500–4508. PMLR, 2018.
- Jonathan Scarlett, Ilija Bogunovic, and Volkan Cevher. Lower bounds on regret for noisy gaussian process bandit optimization. In *Conference on Learning Theory*, pages 1723–1742. PMLR, 2017.
- Bernhard Schölkopf, Ralf Herbrich, and Alex J Smola. A generalized representer theorem. In *International conference on computational learning theory*, pages 416–426. Springer, 2001.
- Pier Giuseppe Sessa, Ilija Bogunovic, Maryam Kamgarpour, and Andreas Krause. No-regret learning in unknown games with correlated payoffs. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Pier Giuseppe Sessa, Ilija Bogunovic, Andreas Krause, and Maryam Kamgarpour. Contextual games: Multi-agent learning with side information. *Advances in Neural Information Processing Systems*, 33, 2020.
- Shubhanshu Shekhar, Tara Javidi, et al. Gaussian process bandits with adaptive discretization. *Electronic Journal of Statistics*, 12(2):3829–3874, 2018.
- Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, page 1015–1022, Madison, WI, USA, 2010. Omnipress. ISBN 9781605589077.
- Yanan Sui, Vincent Zhuang, Joel Burdick, and Yisong Yue. Stagewise safe Bayesian optimization with Gaussian processes. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4781–4789. PMLR, 2018.
- 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*, pages 82–90. PMLR, 2021a.
- Sattar Vakili, Jonathan Scarlett, and Tara Javidi. Open problem: Tight online confidence intervals for rkhs elements. In *Conference on Learning Theory*, pages 4647–4652. PMLR, 2021b.
- Michal Valko, Nathan Korda, Rémi Munos, Ilias Flaounas, and Nello Cristianini. Finite-time analysis of kernelised contextual bandits. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, UAI'13, page 654–663, Arlington, Virginia, USA, 2013. AUAI Press.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I. Jordan. Bridging exploration and general function approximation in reinforcement learning: Provably efficient kernel and neural value iterations. *CoRR*, abs/2011.04622, 2020.
- Weitong ZHANG, Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural thompson sampling. In *International Conference on Learning Representations*, 2021.
- Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 2020.

Supplementary Material: Neural Contextual Bandits without Regret

A Experiments

We carry out experiments on the task of online MNIST (LeCun et al., 1998) classification, to assess how well our analysis of information gain and regret matches the practical behavior of our algorithms. We find that the UCB algorithms exhibit a fairly consistent behavior while solving this classification task. In Section A.3 we design two experiments to demonstrate that $\hat{\sigma}_{t-1}$ and the approximate UCB defined for NN-UCB, are a meaningful substitute for the posterior variance and upper confidence bound for NTK-UCB.

A.1 Technical Setup

We formulate the MNIST classification problem such that it fits the bandit setting, following the setup of Li et al. (2010). To create the context matrix \mathbf{z} from a flattened image $\tilde{\mathbf{z}} \in \mathbb{R}^d$, we construct

$$\mathbf{z}^T = \begin{bmatrix} \tilde{\mathbf{z}}^T & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{z}}^T & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \cdots & \mathbf{0} & \tilde{\mathbf{z}}^T \end{bmatrix} \in \mathbb{R}^{K \times Kd}$$

where $K = |\mathcal{A}| = 10$ in the case of MNIST dataset. The action taken \mathbf{a} is a one-hot K -dimensional vector, indicating which class is selected by the algorithm. The context vector corresponding to an action i is

$$\mathbf{x}_i = \mathbf{z}\mathbf{a}_i = [0 \quad \cdots \tilde{\mathbf{z}} \cdots \quad 0] \in \mathbb{R}^{Kd}$$

where $\tilde{\mathbf{z}}$ occupies indices $(i-1)d$ to id . At every step, the context is picked at random and presented to the algorithm. If the correct action is picked, then a noiseless reward of $y_t = 1$ is given, and otherwise $y_t = 0$.

Training Details For NTK-UCB and CNTK-UCB we use the implementation of the NTK from the NEURAL TANGENTS package (Novak et al., 2020). PYTORCH (Paszke et al., 2019) is used for defining and training the networks in NN-UCB and CNN-UCB. To calculate the approximate UCB, we require the computation of the empirical gram matrix $\mathbf{G}^T\mathbf{G}$ where $\mathbf{G}^T = [\mathbf{g}(\mathbf{x}_t; \boldsymbol{\theta}^0)]_{t \leq T}$. To keep the computations light, we always use the diagonalized matrix as a proxy for $\mathbf{G}^T\mathbf{G}$. Algorithm 3 implies that at every step t the network should be trained from initialization. In practice, however, we use a Stochastic Gradient Descent (SGD) optimizer. To train the network at step t of the algorithm, we consider the loss summed over the data points observed so far, and run the optimizer on $f(\mathbf{x}; \boldsymbol{\theta}_{t-1})$ rather than starting from $f(\mathbf{x}; \boldsymbol{\theta}^0)$. At every step t , we stop training when J reaches 1000, or when the average loss gets small, i.e., $\mathcal{L}(\boldsymbol{\theta}^J)/(t-1) \leq 10^{-3}$. While $t \leq 1000$ we train at every step, for $t > 1000$, however, we train the network once every 100 steps.

Hyper-parameter Tuning For the simple task of MNIST classification, we observe that the width of the network does not significantly impact the results, and after searching the exponential space $m \in \{64, 128, 512, 1024\}$ we set $m = 128$ for all experiments. The models are not extensively fine-tuned and hyper-parameters of the algorithms, β, σ , are selected after a light search over $\{10^{-2k}, 0 \leq k \leq 5\}$, and vary between plots. For all experiments, we set the learning rate of the SGD optimizer to $\eta = 0.01$.

A.2 Growth rates: Empirical vs Theoretical

We test our algorithms on the online MNIST classification task. We plot the empirical information gain and regret to verify the tightness of our bounds. For the information gain of the NN-UCB, we let the algorithm




Figure 1: Empirical information gain of NN-UCB with networks of various depths, on the online MNIST classification problem, is of rate $O(T^{0.45})$. Our bound on the maximum information gain (Theorem 3.1) grows as $\tilde{O}(T^{(d-1)/d})$.




Figure 2: The regret of NTK-UCB and NN-UCB with networks of various depths, on the online MNIST classification problem, is of rate $O(T^{0.75})$. In the worst case it grows as $\tilde{O}(T^{(2d-1)/2d})$ (Corollaries 3.4 & 4.4).

run for $T = 10000$ steps. Then we take the sequence $(\mathbf{x}_t)_{t \leq T}$ from this run and plot $\hat{I}(\mathbf{y}_t; \mathbf{f}_t)$ the empirical information gain against time t , where $\hat{I}(\mathbf{y}_t; \mathbf{f}_t) = \frac{1}{2} \log \det(\mathbf{I} + \sigma^{-2} \mathbf{G}_t^T \mathbf{G}_t)$, with $\mathbf{G}_t^T = [\mathbf{g}(\mathbf{x}_\tau; \boldsymbol{\theta}^0)]_{\tau \leq t}$. Figure 1 shows the growth of \hat{I} for NN-UCB with networks of various depth. To calculate the empirical growth rate (given in the figure’s labels), we fit a polynomial to the curve, and obtain a rate of roughly $O(T^{0.45})$. Our theoretical rate is $\tilde{O}(T^{(d-1)/d})$, where $d = 10 \times 784$ is the dimension of \mathbf{x}_t . Note that the theoretical rate bounds the information gain for arbitrary sequences of contexts including adversarial worst cases. Moreover, in the case of MNIST, the contexts reside on a low-dimensional manifold. As for the regret, we run NN-UCB and NTK-UCB for $T = 5000$ steps and plot the regret, which in the case of online MNIST classification, shows the number of misclassified digits. Fitting a polynomial to the data from Figure 2 gives an empirical rate of $O(T^{0.75})$. Our theoretical bound for these algorithms grows as $O(T^{(2d-1)/2d})$.

In Section 5, we conclude that the information gain and the regret grow with the same rate for both NTK-UCB and its convolutional variant. CNTK-UCB however, tends to have a smaller regret for every T . This is due to the fact that $R(T) = \tilde{O}(C_{\text{NN}}(d, L)T^{(2d-1)/2d})$ for the NTK-UCB, while for the convolutional counterpart $R(T) = \tilde{O}(C_{\text{NN}}(d, L)T^{(2d-1)/2d}/d^{(d-1)/2d})$ and $d \geq 1$. The upper bound on the regret being tighter for CNTK-UCB does not imply that the regret will be smaller as well. In Figure 3 we present both algorithms with the same set of contexts, and investigate whether in practice the convolutional variant can outperform NTK-UCB, which seems to be the case for the online MNIST classification task.

A.3 NN-UCB in the face of Uncertainty

The posterior mean and variance have a transparent mathematical interpretation. For designing NN-UCB, however, we approximate μ_{t-1} and σ_{t-1} with $f_{t-1}^{(J)}$ and $\hat{\sigma}_{t-1}$ which are not as easy to interpret. We design two experiments on MNIST to assess how well this approximation reflects the properties of the posterior mean and variance.

Effect of Imbalanced Classes For this experiment, we limit MNIST to only zeros and ones. We create a dataset with underrepresented zeros, such that the ratio of class populations is 1 : 20. This experiment shows that $\hat{\sigma}_{t-1}$ the approximate posterior variance of NN-UCB behaves as expected, similar to σ_{t-1} . The setup is as follows. Using 80% of this dataset, we first run NN-UCB and train the network. On the remaining 20%, we continue to run the algorithm; no longer training the network, but still updating the posterior variance at every step. Presented in Figure 4, we plot the histogram of the posterior variance and upper confidence bound during the *test* phase. The first row presents histograms for $\sigma^{(\text{post})}(\mathbf{x}_t^*)$, the posterior variance calculated for the true digit \mathbf{x}_t^* at step t . The lower plots show histograms of the corresponding UCB(\mathbf{x}_t^*) values. We give two




Figure 3: NN-UCB vs. CNN-UCB for online MNIST classification. Both algorithms exhibit a similar growth rate with T , while CNN-UCB outperforms NN-UCB, as described in Corollary 5.3.




Figure 4: Under-representation Test. Histograms of $\hat{\sigma}_t$ and approximate UCB are plotted for an imbalanced dataset of zeros and ones. In lack of data, NN-UCB chooses the sub-optimal action with a high confidence.

separate plots for the times when the true digit \mathbf{x}^* comes from the under-represented class, and when it comes from the over-represented class, respectively in the left and right columns. In every plot, to distinguish between the steps at which the digit picked by NN-UCB matches the true digit or not, we make a separate histogram for each case. Figure 4 demonstrates that when the ground truth is over-represented and action is picked correctly, the algorithm always has a high confidence (large UCB) and small uncertainty (posterior variance) for the true class. When the true class is over-represented, for both σ^{post} and UCB, the histogram of the steps with correct picks (green) is concentrated and well separated from the histogram of the steps with incorrect picks (red). This implies that the reason behind misclassifying an over-represented digit is having a large variance, and the algorithm is effectively performing exploration because the estimated reward is small for every action. For the under-represented class, however, the red and green histograms are not well separated. Figure 4 shows that sometimes an under-represented digit has been misclassified with a small posterior variance, or a large UCB. In lack of data, the learner is not able to refine its estimation of the posterior mean or variance. It is forced to do explorations more often which results in incorrect classifications.

Effect of Ambiguous Digits We also assess the ability of the UCB to quantify uncertainty in light of ambiguous MNIST samples. To this end, we define an ambiguous digit to be a data point that is classified incorrectly by a well-trained classifier. We first train a 2-layer CNN with 80% of the MNIST dataset as the standard MNIST classifier. The rest of the data we use for testing. We save the misclassified digits from the test set to study NN-UCB. Figure 5 shows a few examples of such digits and the UCB for the top 5 choices of NN-UCB after observing the digit. It can be seen that for these digits there is no action which the algorithm can pick with high confidence. We first run NN-UCB on the training set. On the test set, we continue to run NN-UCB, however we do not train network any further and only update the posterior variance. For any digit, let \mathbf{x}_1^* denote the maximizer of the UCB, and \mathbf{x}_2^* be the class with the second largest UCB value. Figure 6 shows the histogram of $U_{\mathbf{x}_1^*} - U_{\mathbf{x}_2^*}$. The red histogram is for ambiguous samples and the green one for the non-ambiguous ones. Looking at the medians, we see that for clear samples, the algorithm often has a high confidence on its choice, while this is not the case for the ambiguous digits.

B Details of the Main Result

Here we elaborate on a few matters from the main text.




Figure 5: Examples of ambiguous Digits and the top 5 choices of NN-UCB with the largest UCBs.




Figure 6: Ambiguity Test. Histogram of $U_{\mathbf{x}_1^*} - U_{\mathbf{x}_2^*}$ is plotted. For non-ambiguous digits, NN-UCB is significantly more confident about the class it picks.

B.1 On Section 2.1: Connections Between GP and RKHS Assumptions

We explain how the GP and RKHS assumption imposes smoothness on the reward. By assuming that $f \sim \text{GP}(0, k)$, we set $\text{Cov}(f(\mathbf{x}), f(\mathbf{x}'))$ to $k(\mathbf{x}, \mathbf{x}')$. In doing so, we enforce smoothness properties of k onto f . As an example, suppose some normalized kernel k satisfies boundedness or Lipschitz-continuity, then for $\|\mathbf{x} - \mathbf{x}'\| \leq \delta$, $k(\mathbf{x}, \mathbf{x}')$ is close to $k(\mathbf{x}, \mathbf{x}) = 1$. The GP assumption then ensures high correlation for value of f at these points, making a smooth f more likely to be sampled. The NTK is rotationally invariant and can be written as $\kappa(\mathbf{x}^T \mathbf{x}')$, where κ is continuous and C^∞ over $(-1, 1)$, but is not differentiable at ± 1 (Bietti and Bach, 2021). Therefore, our GP assumption only implies that it is more likely for f to be continuous. Regarding the RKHS assumption, the Stone-Weierstrass theorem shows that any continuous function can be uniformly approximated by members of $\mathcal{H}_{k_{\text{NN}}}$. We proceed by laying out the connection between the two assumptions in more detail.

Equipped with the Mercer's theorem, we can investigate properties of $f \sim \text{GP}(0, k)$, in the general case where k is a Mercer kernel and \mathcal{X} is compact. The following proposition shows that sampling f from a GP is equivalent to assuming $f = \sum_i \beta_i \phi_i$, and sampling the coefficients β_i from $\mathcal{N}(0, \lambda_i)$, where ϕ_i and λ_i are the eigenfunctions and eigenvalues of the GP's kernel function.

Proposition B.1. Let k to be a Hilbert-Schmidt continuous positive semi-definite kernel function, with $(\lambda_i)_{i=1}^\infty$, and $(\phi_i)_{i=1}^\infty$ indicating its eigenvalues and eigenfunctions. Assume \mathcal{X} is compact. If $f \sim \text{GP}(0, k)$, then $f = \sum_i \beta_i \phi_i$, where $\beta_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \lambda_i)$.

Proof. It suffices to show that if $f = \sum_i \beta_i \phi_i$, then for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, we have:

$$\mathbb{E}f(\mathbf{x}) = 0, \quad \mathbb{E}f(\mathbf{x})f(\mathbf{x}') = k(\mathbf{x}, \mathbf{x}').$$

Since β_i are Gaussian and independent,

$$\begin{aligned} \mathbb{E}f(\mathbf{x}) &= \sum_i \phi_i(\mathbf{x}) \mathbb{E}\beta_i = 0, \\ \mathbb{E}f(\mathbf{x})f(\mathbf{x}') &= \sum_i \phi_i(\mathbf{x})\phi_i(\mathbf{x}') \mathbb{E}\beta_i^2 = \sum_i \lambda_i \phi_i(\mathbf{x})\phi_i(\mathbf{x}') = k(\mathbf{x}, \mathbf{x}'). \end{aligned}$$

Here we have used orthonormality of ϕ_i s. □

Proposition B.1 suggests that if $f \sim \text{GP}(0, k)$ then $\|f\|_k$ is almost surely unbounded, since by the definition of the inner product on \mathcal{H}_k we have

$$\mathbb{E}\|f\|_k^2 = \sum_i \mathbb{E} \frac{\beta_i^2}{\lambda_i} = \sum_i \frac{\lambda_i}{\lambda_i}.$$

This expectation is unbounded for any kernel with an infinite number of nonzero eigenvalues. Therefore, with probability one, $\|f\|_k$ is unbounded and not a member of \mathcal{H}_k . However, the posterior mean of f after observing T samples lies in \mathcal{H}_k (Proposition B.2). This connection implies that our estimate of f , under both RKHS and GP assumption will be a k -norm bounded function, similarly reflecting the smoothness properties of the kernel.

Proposition B.2. Assume $f \sim \text{GP}(0, k)$, with k Mercer and \mathcal{X} compact. Then μ_T the posterior mean of f given $(\mathbf{x}_i, y_i)_{i=1}^T$ has bounded k -norm, and $\mu_T \in \mathcal{H}_k$.

Proof. We first recall the Representer Theorem (Schölkopf et al., 2001). Consider the loss function $J(f) = Q(\mathbf{f}_T; \mathbf{y}_T) + \sigma^2 \|f\|_k^2$, where $\mathbf{f}_T = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_T)]^T$. Q is a ℓ_2 -loss assessing the fit of \mathbf{f}_T to \mathbf{y}_T . Then $J(f)$ has a unique minimizer, which takes the form:

$$\hat{f} = \sum_{i=1}^T \alpha_i k(\mathbf{x}_i, \cdot)$$

Note that this sum is finite, hence \hat{f} is k -norm bounded and in \mathcal{H}_k . Minimizing J over α , we get

$$\hat{\alpha} = (\mathbf{K}_T + \sigma^2 \mathbf{I})^{-1} \mathbf{y}_T$$

Indicating that $\hat{f}(\mathbf{x}) = \mu_T(\mathbf{x})$. □

B.2 On Assumptions of Theorem 4.2

For technical simplicity, in Theorem 4.2 and its convolutional extension, Theorem 5.4, we require the network at initialization to satisfy $f(\mathbf{x}; \boldsymbol{\theta}^0) = 0$ for all $\mathbf{x} \in \mathcal{X}$. We explain how this assumption can be fulfilled without limiting the problem setting to a specific network or input domain. Without loss of generality, we can initialize the network as follows. For $l \leq L$ let the weights at initialization be,

$$W^{(l)} = \begin{pmatrix} W & 0 \\ 0 & W \end{pmatrix}, \quad W^{(L+1)} = (\mathbf{w}^T, -\mathbf{w}^T)$$

where entries of W and \mathbf{w} are i.i.d and sampled from the normal distribution. Moreover, we assume that $[\mathbf{x}]_j = [\mathbf{x}]_{j+d/2}$ for any $\mathbf{x} = \mathbf{z}_t \mathbf{a}$ where $1 \leq t \leq T$ and any $\mathbf{a} \in \mathcal{A}$. Any input \mathbf{x} can be converted to satisfy this assumption by defining an auxiliary input $\tilde{\mathbf{x}} = [\mathbf{x}, \mathbf{x}]/\sqrt{2}$ in a higher dimension. This mapping together with the initialization method, guarantee that output of the network at initialization is zero for every possible input that the learner may observe during a T -step run of the algorithm. Essentially, this property comes into effect when using result from Arora et al. (2019) or working with the Taylor expansion of the network around initialization. It allows us to write $f(\mathbf{x}_i; \boldsymbol{\theta}) \approx \langle \mathbf{g}^T(\mathbf{x}_i; \boldsymbol{\theta}^0), \boldsymbol{\theta} - \boldsymbol{\theta}^0 \rangle$.

B.3 On Section 5: CNTK as an averaged NTK

This section gives an intuition on, and serves as a sketch for proving Equation 8. Consider a 2-layer ReLU network with width m defined as,

$$f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}^{(0)}) = \frac{2}{\sqrt{m}} \sum_{i=1}^m v_i^{(0)} \sigma(\langle \mathbf{w}_i^{(0)}, \mathbf{x} \rangle)$$

where each weight parameter is an i.i.d sample from $\mathcal{N}(0, 1)$. We denote the complete weight vector by $\boldsymbol{\theta}$ and write the first order Taylor approximation of this function with respect to $\boldsymbol{\theta}$ around the initialization.

$$f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}) \simeq f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta}^{(0)}) + \overbrace{\frac{2}{\sqrt{m}} \sum_{i=1}^m (v_i - v_i^{(0)}) \sigma(\langle \mathbf{w}_i^{(0)}, \mathbf{x} \rangle)}^{f_{\text{NN},1}(\mathbf{x}; \mathbf{v})} + \underbrace{\frac{2}{\sqrt{m}} \sum_{i=1}^m v_i^{(0)} \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, \mathbf{x} \rangle) \langle \mathbf{w}_i - \mathbf{w}_i^{(0)}, \mathbf{x} \rangle}_{f_{\text{NN},2}(\mathbf{x}; \mathbf{W})}$$

We limit the input domain to \mathbb{S}^{d-1} , equipped with the uniform measure. Consider the function class $\mathcal{F}_{\text{NTK}} = \{f(\mathbf{x}) = f_{\text{NN},1}(\mathbf{x}; \mathbf{v}) + f_{\text{NN},2}(\mathbf{x}; \mathbf{W}) \text{ s.t. } \mathbf{v} \in \mathbb{R}^m, \mathbf{W} \in \mathbb{R}^{m \times d}\}$. It is straightforward to show that \mathcal{F}_{NTK} is an RKHS and the following kernel function satisfies the reproducing property for \mathcal{F}_{NTK} (Chizat et al., 2019).

$$h(\mathbf{x}, \mathbf{x}') = \frac{1}{m} \sum_{i=1}^m \sigma(\langle \mathbf{w}_i^{(0)}, \mathbf{x} \rangle) \sigma(\langle \mathbf{w}_i^{(0)}, \mathbf{x}' \rangle) + \frac{1}{m} \sum_{i=1}^m \mathbf{x}^T \mathbf{x}' (v_i^{(0)})^2 \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, \mathbf{x} \rangle) \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, \mathbf{x}' \rangle)$$

Random Feature Derivation of 2-layer CNTK We repeat the first order Taylor approximation for f_{CNN} .

$$f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta}) \simeq f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta}^{(0)}) + \overbrace{\frac{2}{\sqrt{m}} \sum_{i=1}^m (v_i - v_i^{(0)}) \left[\frac{1}{d} \sum_{l=1}^d \sigma(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x} \rangle) \right]}^{f_{\text{CNN},1}(\mathbf{x}; \mathbf{v})} + \underbrace{\frac{2}{\sqrt{m}} \sum_{i=1}^m v_i^{(0)} \left[\frac{1}{d} \sum_{l=1}^d \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x} \rangle) \langle \mathbf{w}_i - \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x} \rangle \right]}_{f_{\text{CNN},2}(\mathbf{x}; \mathbf{W})}$$

And define $\mathcal{F}_{\text{CNTK}}$ to be the convolutional counterpart of \mathcal{F}_{NTK}

$$\mathcal{F}_{\text{CNTK}} = \{f(\mathbf{x}) = f_{\text{CNN},1}(\mathbf{x}; \mathbf{v}) + f_{\text{CNN},2}(\mathbf{x}; \mathbf{W}) \text{ s.t. } \mathbf{v} \in \mathbb{R}^m, \mathbf{W} \in \mathbb{R}^{m \times d}\}.$$

Then $\mathcal{F}_{\text{CNTK}}$ is reproducing for the following kernel function,

$$\begin{aligned} \bar{h}(\mathbf{x}, \mathbf{x}') &= \frac{1}{m} \sum_{i=1}^m \frac{1}{d^2} \sum_{l=1}^d \sum_{l'=1}^d \sigma(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x} \rangle) \sigma(\langle \mathbf{w}_i^{(0)}, c_{l'} \cdot \mathbf{x}' \rangle) \\ &\quad + \frac{1}{m} \sum_{i=1}^m \frac{1}{d^2} \sum_{l=1}^d \sum_{l'=1}^d \langle c_l \cdot \mathbf{x}, c_{l'} \cdot \mathbf{x}' \rangle (v_i^{(0)})^2 \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x} \rangle) \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, c_{l'} \cdot \mathbf{x}' \rangle) \end{aligned}$$

As m the number of channels grows, the average over the parameters converges to the expectation over \mathbf{w} and \mathbf{v} , and the kernel becomes rotation invariant, i.e. only depends on the angle between \mathbf{x} and \mathbf{x}' .

$$\begin{aligned} \lim_{m \rightarrow \infty} \bar{h}(\mathbf{x}, \mathbf{x}') &= \lim_{m \rightarrow \infty} \frac{1}{d} \frac{1}{m} \sum_{i=1}^m \sum_{l=1}^d \sigma(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x} \rangle) \sigma(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x}' \rangle) \\ &\quad + \langle c_l \cdot \mathbf{x}, c_l \cdot \mathbf{x}' \rangle (v_i^{(0)})^2 \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x} \rangle) \dot{\sigma}(\langle \mathbf{w}_i^{(0)}, c_l \cdot \mathbf{x}' \rangle) \\ &= \frac{1}{d} \sum_{l=1}^d h(c_l \cdot \mathbf{x}, c_l \cdot \mathbf{x}') \end{aligned}$$

Equation 8 follows by noting that in the infinite m limit, h and \bar{h} converge to k_{NN} and k_{CNN} respectively. A rigorous proof is given in Proposition 4 from Mei et al. (2021).

B.4 Proof of Lemma 5.1

Let $V_{d,k}$ be the space of degree- k polynomials that are orthogonal to the space of polynomials of degree less than k , defined on \mathbb{S}^{d-1} . Then $V_{d,k}(\mathcal{C}_d)$ denotes the subspace of $V_{d,k}$ that is also \mathcal{C}_d -invariant.

Proof of Lemma 5.1. The NTK kernel satisfies the Mercer condition and has the following Mercer decomposition (Bietti and Bach, 2021),

$$k_{\text{NTK}}(\mathbf{x}, \mathbf{x}') = \sum_{k=0}^{\infty} \mu_k \sum_{j=1}^{N(d,k)} Y_{j,k}(\mathbf{x}) Y_{j,k}(\mathbf{x}')$$

where $\{Y_{j,k}\}_{j \leq N(d,k)}$ form an orthonormal basis for $V_{d,k}$. Bietti and Bach (2021) further show that $Y_{j,k}$ is the j -th spherical harmonic polynomial of degree k , and $N(d,k) = \dim(V_{d,k})$ gives the total count of such polynomials, where

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

For any integer k , it also holds that (Bietti and Bach, 2021)

$$\sum_{j=1}^{N(d,k)} Y_{j,k}(\mathbf{x}) Y_{j,k}(\mathbf{x}') = N(d,k) Q_k^{(d)}(\mathbf{x}^T \mathbf{x}') \quad (\text{B.1})$$

where $Q_k^{(d)}$ is the k -th Gegenbauer polynomial in dimension d . It follows from Equation 8 that k_{CNTK} is also Mercer. Assume that it has a Mercer decomposition of the form

$$k_{\text{CNTK}}(\mathbf{x}, \mathbf{x}') = \sum_{k=0}^{\infty} \bar{\mu}_k \sum_{j=1}^{\bar{N}(d,k)} \bar{Y}_{j,k}(\mathbf{x}) \bar{Y}_{j,k}(\mathbf{x}').$$

We now proceed to identify $\bar{Y}_{j,k}$ and calculate $\bar{N}(d,k)$. From Equations 8 and B.1 we conclude that

$$\bar{k}(\mathbf{x}, \mathbf{x}') = \frac{1}{d} \sum_{l=1}^d k(\mathbf{x}, c_l \cdot \mathbf{x}') = \frac{1}{d} \sum_{l=1}^d \sum_{k=0}^{\infty} \mu_k N(d,k) Q_k^{(d)}(\langle \mathbf{x}, c_l \cdot \mathbf{x}' \rangle).$$

Lemma 1 in Mei et al. (2021) states that for any integer k ,

$$\frac{1}{d} N(d,k) \sum_{l=1}^d Q_k^{(d)}(\langle \mathbf{x}, c_l \cdot \mathbf{x}' \rangle) = \sum_{j=1}^{M(d,k)} Z_{j,k}(\mathbf{x}) Z_{j,k}(\mathbf{x}') \quad (\text{B.2})$$

where $Z_{j,k}$ form an orthonormal basis for $V_{d,k}(\mathcal{C}_d)$. Moreover, for the orthonormal basis $(Z_{j,k})_j$ over \mathbb{S}^{d-1} ,

$$\sum_{j=1}^{M(d,k)} Z_{j,k}(\mathbf{x}) Z_{j,k}(\mathbf{x}) = M(d,k). \quad (\text{B.3})$$

Therefore, $(\mu_k, Z_{j,k})_{j,k}$ is a sequence eigenvalue eigenfunction pairs for \bar{k} and the following assignments satisfy the Mercer decomposition for k_{CNTK} :

$$\begin{aligned} \bar{\mu}_k &= \mu_k, \\ \bar{Y}_{j,k} &= Z_{j,k}, \\ \bar{N}(d,k) &= M(d,k). \end{aligned}$$

It remains to show that $M(d,k) \simeq N(d,k)/d$ when \mathcal{C}_d is the group of cyclic shifts acting on \mathbb{S}^{d-1} . This follows from Equations B.2 and B.3,

$$\begin{aligned} \frac{M(d,k)}{N(d,k)} &= \frac{1}{d} \sum_{l=1}^d Q_k^{(d)}(\langle \mathbf{x}, c_l \cdot \mathbf{x} \rangle) \\ &= \Theta(d^{-1}) \end{aligned}$$

where the last equation holds directly due to Lemma 4 in Mei et al. (2021). \square

C Details of NTK-UCB and CNTK-UCB

Section C.1 gives the proof to our statement on the information gain (Theorem 3.1, Proposition 5.2). The regret bounds of NTK-UCB under the GP (Theorem 3.2) and the RKHS assumptions (Theorem 3.3) are proven in Section C.2 and Section C.4, respectively. In Section C.3, we present SUPNTK-UCB (Algorithm 2), and discuss the exploration policy of this algorithm and its properties.

C.1 Proof of Theorem 3.1 and Proposition 5.2

We begin by giving an overview of the proof. Under the conditions of Theorem 3.1, the NTK is Mercer (Cao et al., 2021) and can be written as $k(\mathbf{x}, \mathbf{x}') = \sum_{i \geq 0} \lambda_i \phi_i(\mathbf{x}) \phi_i(\mathbf{x}')$, with $(\phi_i)_{i \geq 0}$ denoting the orthonormal eigenfunctions. The main idea is to break k into $k_p + k_o$ where $k_p(\mathbf{x}, \mathbf{x}') = \sum_{i \leq \tilde{D}} \lambda_i \phi_i(\mathbf{x}) \phi_i(\mathbf{x}')$ has a finite dimensional feature map $(\phi_i)_{i \leq \tilde{D}}$, corresponding to the sequence of \tilde{D} largest eigenvalues of k . For any arbitrary sequence X_T , we are then able to decompose $I(\mathbf{y}_T; \mathbf{f}_T)$ in two terms, one corresponding to information gain of k_p , and the other, the tail sum of the eigenvalue series $\sum_{i \geq \tilde{D}} \lambda_i$. We proceed by bounding each term separately and picking \tilde{D} such that the second term becomes negligible.

Proof of Theorem 3.1. This proof adapts the finite-dimensional projection idea used in Vakili et al. (2021a). Cao et al. (2021) show that the NTK has a Mercer decomposition on the d -dimensional unit hyper-sphere. Bietti and Bach (2021) establish that

$$k_{\text{NN}}(\mathbf{x}, \mathbf{x}') = \sum_{k \geq 0} \mu_k \sum_{j=1}^{N(d,k)} Y_{j,k}(\mathbf{x}) Y_{j,k}(\mathbf{x}'),$$

where $Y_{j,k}$ is the j -th spherical harmonic polynomial of degree k , and for $k \geq 1$ there are

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

of such k -degree polynomials. Using Stirling's approximation we can show that $N(d, k) = \Theta(k^{d-2})$. The functions $\{Y_{j,k}\}$ are an algebraic basis for $\mathcal{H}_{k_{\text{NN}}}$ the RKHS that is reproducing for k_{NN} . Consider a finite dimensional subspace of $\mathcal{H}_{k_{\text{NN}}}$ that is spanned by the eigenfunctions corresponding to the first D distinct eigenvalues of k_{NN} , $\phi_D = ((Y_{j,0})_{j \leq N(d,0)}, \dots, (Y_{j,D})_{j \leq N(d,D)})$. We decompose the NTK as $k_{\text{NN}} = k_P + k_O$, where k_P is the kernel for the finite dimensional RKHS, and k_O represents the kernel for the Hilbert space orthogonal to it. Let \tilde{D} denote the length of ϕ_D , the feature map corresponding to k_P ,

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

Note that the first \tilde{D} eigenvalues of k_{NN} make up its first D distinct eigenvalues. We write the information gain in terms of eigenvalues of k_p and k_o , and find D such that the finite-dimensional term dominates the infinite-dimensional tail. Assume the arbitrary sequence $X_T = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ is observed. The information gain is $I(\mathbf{y}_T; \mathbf{f}_T) = \frac{1}{2} \log \det(\mathbf{I} + \sigma^{-2} K_{X_T})$, with K_{X_T} being the kernel matrix, $(K_{X_T})_{i,j} = k_{\text{NN}}(\mathbf{x}_i, \mathbf{x}_j)$. Using a similar notation for the kernel matrices of k_P and k_O , we may write

$$\begin{aligned} I(\mathbf{y}_T; \mathbf{f}_T) &= \frac{1}{2} \log \det(\mathbf{I} + \sigma^{-2} (K_{P, X_T} + K_{O, X_T})) \\ &= \frac{1}{2} \log \det(\mathbf{I} + \sigma^{-2} K_{P, X_T}) + \frac{1}{2} \log \det(\mathbf{I} + (\mathbf{I} + \sigma^{-2} K_{P, X_T})^{-1} K_{O, X_T}) \end{aligned} \quad (\text{C.2})$$

We now separately bound the two terms. Let $\Phi_{D,T} = [\phi_D(\mathbf{x}_1), \dots, \phi_D(\mathbf{x}_T)]^T$, then by the mercer decomposition,

$$K_{P, X_T} = \Phi_{D,T} \Lambda_D \Phi_{D,T}^T$$

Where $\Lambda_D \in \mathbb{R}^{\tilde{D} \times \tilde{D}}$ is a diagonal matrix with the first \tilde{D} eigenvalues. Let $\mathbf{H}_T = \Lambda_D^{1/2} \Phi_{D,T}^T \Phi_{D,T} \Lambda_D^{1/2}$, by Weinstein-Aronszajn identity,

$$\begin{aligned} \frac{1}{2} \log \det(\mathbf{I}_T + \sigma^{-2} K_{P,X_T}) &= \frac{1}{2} \log \det(\mathbf{I} + \sigma^{-2} \mathbf{H}_T) \\ &\leq \frac{1}{2} \tilde{D} \log \left(\text{tr}(\mathbf{I} + \sigma^{-2} \mathbf{H}_T) / \tilde{D} \right) \end{aligned}$$

For positive definite matrices $\mathbf{P} \in \mathbb{R}^{n \times n}$, we have $\log \det \mathbf{P} \leq n \log \text{tr}(\mathbf{P}/n)$. The inequality follows from $\mathbf{I}_D + \sigma^{-2} \mathbf{H}_T$ being positive definite. Plugging in the definition of \mathbf{H}_T

$$\begin{aligned} \frac{1}{2} \log \det(\mathbf{I}_T + \sigma^{-2} K_{P,X_T}) &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\sigma^{-2}}{\tilde{D}} \text{tr} \left(\Lambda_D^{1/2} \Phi_{D,T}^T \Phi_{D,T} \Lambda_D^{1/2} \right) \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\sigma^{-2}}{\tilde{D}} \sum_{t=1}^T \Lambda_D^{1/2} \Phi_D^T(\mathbf{x}_t) \phi_D(\mathbf{x}_t) \Lambda_D^{1/2} \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\sigma^{-2}}{\tilde{D}} \sum_{t=1}^T \|\phi_D(\mathbf{x}_t) \Lambda^{1/2}\|_2^2 \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\sigma^{-2}}{D} \sum_{t=1}^T \sum_{k=0}^D \mu_k \sum_{j=1}^{N(d,k)} Y_{j,k}^2(\mathbf{x}_t) \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\sigma^{-2}}{\tilde{D}} \sum_{t=1}^T k_P(\mathbf{x}_t, \mathbf{x}_t) \right) \\ &\leq \frac{1}{2} \tilde{D} \log \left(1 + \frac{\sigma^{-2} T}{\tilde{D}} \right) \end{aligned}$$

The last inequality holds since the NTK is uniformly bounded by 1 on the unit sphere. We now bound the second term in Equation C.2, which corresponds to the infinite-dimensional orthogonal space. Similar to the first term, we bound $\log \det \mathbf{P}$ with $n \log \text{tr}(\mathbf{P}/n)$,

$$\begin{aligned} \frac{1}{2} \log \det(\mathbf{I}_T + (\mathbf{I}_T + \sigma^{-2} K_{P,X_T})^{-1} K_{O,X_T}) &\leq \frac{T}{2} \log \left(1 + \frac{\text{tr} \left((\mathbf{I}_T + \sigma^{-2} K_{P,X_T})^{-1} K_{O,X_T} \right)}{T} \right) \\ &\leq \frac{T}{2} \log (1 + \text{tr}(K_{O,X_T})/T) \end{aligned} \quad (\text{C.3})$$

The second inequality holds due to $(\mathbf{I}_T + \sigma^{-2} K_{P,X_T})^{-1}$ being positive definite, with eigenvalues smaller than 1. We can bound the trace using the Mercer decomposition

$$\begin{aligned} \text{tr}(K_{O,X_T}) &= \sum_{t=1}^T k_O(\mathbf{x}_t, \mathbf{x}_t) \\ &= \sum_{t=1}^T \sum_{k=D+1}^{\infty} \mu_k \sum_{j=1}^{N(d,k)} Y_{j,k}^2(\mathbf{x}_t) \\ &= \sum_{t=1}^T \sum_{k=D+1}^{\infty} \mu_k N(d, k) Q_k^{(d)}(\mathbf{x}_t^T \mathbf{x}_t) \end{aligned}$$

where $Q_k^{(d)}$ is the Gegenbauer polynomial of degree k and \mathbf{x}_t are on \mathbb{S}^{d-1} . We have $Q_K^{(d)}(1) = 1$ which gives,

$$\text{tr}(K_{O,X_T}) = T \sum_{k=D+1}^{\infty} \mu_k N(d, k) \quad (\text{C.4})$$

Plugging in Equation C.4 in Equation C.3, the information gain can be bounded as,

$$I(\mathbf{y}_T, \mathbf{f}_T) \leq \frac{\tilde{D}}{2} \log \left(1 + \frac{T}{\sigma^2 \tilde{D}} \right) + \frac{T}{2} \log \left(1 + \sum_{k=D+1}^{\infty} \mu_k N(d, k) \right)$$

We now bound the second term using Bietti and Bach’s result on decay rate of μ_k . They show that there exists a constant $C_1(d, L)$ such that $\mu_k \leq C_1 k^{-d}$. Using Stirling approximation, there exists C_2 such that $N(d, k) \leq C_2 k^{d-2}$. Then,

$$\sum_{k=D+1}^{\infty} \mu_k N(d, k) \leq C(d, L) \sum_{k=D+1}^{\infty} k^{-2}$$

We can simply bound the series

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

Therefore, there exists there exists $C(d, L)$ such that,

$$\begin{aligned} I(\mathbf{y}_T, \mathbf{f}_T) &\leq \frac{\tilde{D}}{2} \log \left(1 + \frac{T}{\sigma^2 \tilde{D}} \right) + \frac{T}{2} \log \left(1 + \frac{C(d, L)}{D} \right) \\ &\leq \frac{\tilde{D}}{2} \log \left(1 + \frac{T}{\sigma^2 \tilde{D}} \right) + \frac{TC(d, L)}{2D} \end{aligned} \tag{C.5}$$

Note that the first term is increasing with \tilde{D} . We pick \tilde{D} such that the first term in Equation C.5 is dominant. Via Equation C.1 we set,

$$\tilde{D} = \left\lceil \left(\frac{C(d, L)T}{\log(1 + \sigma^{-2}T)} \right)^{\frac{d-1}{d}} \right\rceil$$

The treatment above holds for any arbitrary sequence X_T . Plugging in \tilde{D} with Equation C.5, we then may write,

$$\gamma_T \leq \left(\frac{C(d, L)T}{\log(1 + \sigma^{-2}T)} \right)^{\frac{d-1}{d}} \log \left(1 + \sigma^{-2}T \left(\frac{C(d, L)T}{\log(1 + \sigma^{-2}T)} \right)^{\frac{d}{d-1}} \right)$$

which concludes the proof. \square

Proof of Proposition 5.2. The CNTK is Mercer by Lemma 5.1. Which implies that we may repeat the steps taken in the proof of Theorem 3.1. To avoid confusion, we use the “bar” notation to indicate the convolutional equivalent of the parameters from that proof. The 2-layer CNTK and the NTK share the same eigenvalues. The eigenfunctions of the CNTK also bounded by one over the hyper sphere. The only difference is that $\bar{N}(d, k) = N(d, k)/d$, which comes into effect for calculating $\text{tr}(\bar{K}_{O, X_T})$. For the CNTK we would have,

$$\text{tr}(\bar{K}_{O, X_T}) = T \sum_{k=\tilde{D}+1}^{\infty} \mu_k \bar{N}(d, k) = T \sum_{k=\tilde{D}+1}^{\infty} \mu_k N(d, k)/d$$

Equivalent to Equation C.5 we may write,

$$\bar{I}(\mathbf{y}_T, \mathbf{f}_T) \leq \frac{\tilde{D}}{2} \log \left(1 + \frac{T}{\sigma^2 \tilde{D}} \right) + \frac{TC(d, L)}{2\tilde{D}d}$$

where similar to the proof of Theorem 3.1, we have $\tilde{D} = \Theta(\bar{D}^{d-1})$. For the first term to be larger than the second, \tilde{D} has to be set to

$$\tilde{D} = \left\lceil \left(\frac{C(d, L)T}{d \log(1 + \sigma^{-2}T)} \right)^{\frac{d-1}{d}} \right\rceil$$

which then concludes the proof. \square

C.2 Proof of Theorem 3.2

The proof closely follows the method in Srinivas et al. (2010) and Krause and Ong (2011), with modifications on the assumptions on context domain and actions. The following lemmas will be used.

Lemma C.1. Let $\delta \in (0, 1)$, and set $\beta_t = 2 \log(|\mathcal{A}| \pi_t / \delta)$, where $\sum_{t>1} \pi_t^{-1} = 1$ and $\pi_t > 0$. Then with probability of at least $1 - \delta$

$$|f(\mathbf{z}_t \mathbf{a}) - \mu_{t-1}(\mathbf{z}_t \mathbf{a})| \leq \sqrt{\beta_t} \sigma_{t-1}(\mathbf{z}_t \mathbf{a}), \quad \forall \mathbf{a} \in \mathcal{A}, \forall t > 1.$$

Lemma C.2. Let $\sigma_{t-1}^2(\mathbf{x}_t)$ be the posterior variance, computed at $\mathbf{x}_t = \mathbf{z}_t \mathbf{a}_t$, where \mathbf{a}_t is the action picked by the UCB policy. Then,

$$\frac{1}{2} \sum_{t=1}^T \log(1 + \sigma^{-2} \sigma_{t-1}^2(\mathbf{x}_t)) \leq \gamma T$$

Proof of Theorem 3.2. We use Lemma C.1, to bound the regret at step t . Let $\mathbf{x}_t^* = \mathbf{z}_t \mathbf{a}^*$ denote the optimal point, and $\mathbf{x}_t = \mathbf{z}_t \mathbf{a}_t$ be the maximizer of the UCB. Then by Lemma C.1, with probability of at least $1 - \delta$

$$\begin{aligned} |f(\mathbf{x}_t) - \mu_{t-1}(\mathbf{x}_t)| &\leq \sqrt{\beta_t} \sigma_{t-1}(\mathbf{x}_t) \\ |f(\mathbf{x}_t^*) - \mu_{t-1}(\mathbf{x}_t^*)| &\leq \sqrt{\beta_t} \sigma_{t-1}(\mathbf{x}_t^*) \end{aligned}$$

Therefore, by definition of \mathbf{x}_t (Equation 5) we can write:

$$\begin{aligned} r_t = f(\mathbf{x}_t^*) - f(\mathbf{x}_t) &\leq \mu_{t-1}(\mathbf{x}_t^*) + \sqrt{\beta_t} \sigma_{t-1}(\mathbf{x}_t^*) - f(\mathbf{x}_t) \\ &\leq \mu_{t-1}(\mathbf{x}_t) + \sqrt{\beta_t} \sigma_{t-1}(\mathbf{x}_t) - f(\mathbf{x}_t) \\ &\leq 2\sqrt{\beta_t} \sigma_{t-1}(\mathbf{x}_t). \end{aligned}$$

For the regret over T steps, by Cauchy-Schwartz we have,

$$\begin{aligned} R_T &= \sqrt{\sum_{t=1}^T r_t} \leq \sqrt{T} \sqrt{\sum_{t=1}^T r_t^2} \\ &\leq \sqrt{T} \sqrt{\sum_{t=1}^T 4\beta_t \sigma_{t-1}^2(\mathbf{x}_t)}. \end{aligned} \tag{C.6}$$

Recall that Lemma C.1, holds for any $\beta_t = 2 \log(|\mathcal{A}| \pi_t / \delta)$, where $\sum_{t>1} \pi_t^{-1} = 1$. We pick $\pi_t = \frac{\sigma^{-2} t^2}{6}$, so that β_t is non-increasing and can be upper bounded by β_T . This allows to reduce the problem of bounding regret to the bounding the sum of posterior variances. By Equation 4, and since $k_{\text{NN}}(\mathbf{x}, \mathbf{x}) \leq 1$,

$$\sigma^{-2} \sigma_{t-1}^2(\mathbf{x}_t) \leq \sigma^{-2} k_{\text{NN}}(\mathbf{x}_t, \mathbf{x}_t) \leq \sigma^{-2}$$

For any $r \in [0, a]$, it holds that $r \leq \frac{a \log(1+r)}{\log(1+a)}$. Therefore,

$$\sigma^{-2} \sigma_{t-1}^2 \leq \frac{\sigma^{-2}}{\log(1 + \sigma^{-2})} \log(1 + \sigma^{-2} \sigma_{t-1}^2)$$

Putting together the sum in Equation C.6 we get,

$$\begin{aligned} R_t &\leq \sqrt{4T\beta_T \sigma^2 \sum_{t=1}^T \sigma^{-2} \sigma_{t-1}^2} \\ &\leq \sqrt{\frac{4T\beta_T \sigma^2}{\log(1 + \sigma^{-2})} \sum_{t=1}^T \log(1 + \sigma^{-2} \sigma_{t-1}^2)} \\ &\leq \sqrt{\frac{8\sigma^2}{\log(1 + \sigma^{-2})}} \sqrt{T\beta_T \gamma T} \end{aligned}$$

where the last inequality holds by plugging in Lemma C.2. \square

C.2.1 Proof of Lemma C.1

Proof. Fix $t \geq 1$. Conditioned on $\mathbf{y}_{t-1} = (y_1, \dots, y_{t-1})$, $\mathbf{x}_1, \dots, \mathbf{x}_{t-1}$ are deterministic. The posterior distribution is $f(\mathbf{x}) \sim N(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}^2(\mathbf{x}))$. Applying the sub-Gaussian inequality, and conditioned on the history,

$$\Pr[|f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| > \sqrt{\beta_t} \sigma_{t-1}(\mathbf{x})] \leq e^{-\beta_t/2}$$

\mathcal{A} is finite and $\mathbf{x} = \mathbf{z}_t \mathbf{a}$, then by union bound over S , the following holds with probability of at least $1 - |\mathcal{A}|e^{-\beta_t/2}$,

$$|f(\mathbf{z}_t \mathbf{a}) - \mu_{t-1}(\mathbf{z}_t \mathbf{a})| \leq \sqrt{\beta_t} \sigma_{t-1}(\mathbf{z}_t \mathbf{a}), \quad \forall \mathbf{a} \in \mathcal{A}$$

It is only left to further apply a union bound over all $t \geq 1$. For the statement in lemma to hold, β_t has to be set such that, $\sum_{t \geq 1} |\mathcal{A}|e^{\beta_t/2} \leq \delta$. Setting $\beta_t = 2 \log(|\mathcal{A}|\pi_t/\delta)$ with $\sum_{t \geq 1} \pi_t = 1$ satisfies the condition. \square

C.2.2 Proof of Lemma C.2

Proof. Recall that for a Gaussian random variable entropy is, $H(N(\boldsymbol{\mu}, \boldsymbol{\Sigma})) = \frac{1}{2} \log \det(2\pi e \boldsymbol{\Sigma})$. Let \mathbf{y}_T be the vector of observed rewards and $\mathbf{f}_T = [f(\mathbf{x}_t)]_{t \leq T} \in \mathbb{R}^T$ the true rewards. We have $\mathbf{y}_T = \mathbf{f}_T + \boldsymbol{\epsilon}_T$, therefore, $\mathbf{y}_T | \mathbf{f}_T \sim N(0, \mathbf{I} \sigma^2)$, and $y_T | \mathbf{y}_{T-1} \sim N(\mu_{T-1}, \sigma^2 + \sigma_{T-1}^2(\mathbf{x}_T))$. By the definition of mutual information,

$$\begin{aligned} I(\mathbf{y}_T; \mathbf{f}_T) &= H(\mathbf{y}_T) - H(\mathbf{y}_T | \mathbf{f}_T) \\ &= H(\mathbf{y}_{T-1}) + H(y_T | \mathbf{y}_{T-1}) - \frac{T}{2} \log(2\pi e \sigma^2) \\ &= H(\mathbf{y}_{T-1}) + \frac{1}{2} \log(2\pi e (\sigma^2 + \sigma_{T-1}^2(\mathbf{x}_T))) - \frac{T}{2} \log(2\pi e \sigma^2). \end{aligned}$$

The first equality holds by the chain rule for entropy. By recursion,

$$\gamma_T \geq I(\mathbf{y}_T; \mathbf{f}_T) = \frac{1}{2} \sum_{t=1}^T \log(1 + \sigma^{-2} \sigma_{t-1}^2(\mathbf{x}_t))$$

\square

C.3 The Sup Variant and its Properties

Algorithm 1: GetPosterior

Input: $\Psi_t \subset \{1, \dots, t-1\}$, \mathbf{z}_t

Initialize $\mathbf{K} \leftarrow [k_{\text{NN}}(\mathbf{x}_i, \mathbf{x}_j)]_{i,j \in \Psi_t}$, $\mathbf{Z}^{-1} \leftarrow (\mathbf{K} + \sigma^2 \mathbf{I})^{-1}$, $\mathbf{y} \leftarrow [y_i]_{i \in \Psi_t}^T$

for $\mathbf{a} \in \mathcal{A}$ **do**

 Define $\mathbf{x} := \mathbf{z}_t \mathbf{a}$,

$\mathbf{k} \leftarrow [k_{\text{NN}}(\mathbf{x}_i, \mathbf{x})]_{i \in \Psi_t}^T$

$\mu_{t-1}^{(s)}(\mathbf{x}) \leftarrow \mathbf{k}^T \mathbf{Z}^{-1} \mathbf{y}$

$\sigma_{t-1}^{(s)}(\mathbf{x}) \leftarrow \sqrt{k_{\text{NN}}(\mathbf{x}, \mathbf{x}) - \mathbf{k}^T \mathbf{Z}^{-1} \mathbf{k}}$

end

SUPNTK-UCB combines NTK-UCB policy and Random Exploration, and at every step t , only uses a subset of the previously observed context-reward pairs. These subsets are constructed such that the rewards in each are statistically independent, conditioned on the contexts. Informally put, then the learner chooses an action either if its posterior variance is very high or if the reward is close to the optimal reward. As more steps are played, the criteria for *closeness* to optimal reward and *high variance* is refined. The method is given in Algorithm 2. We give some intuition on the key elements to which the algorithm's desirable properties can be credited.

- The set of indices of the context-reward pairs used for calculating $\mu_{t-1}^{(s)}$ and $\sigma_{t-1}^{(s)}$, is denoted by $\Psi_t^{(s)}$. Once an action is chosen, $\Psi_t^{(s)}$ is updated to $\Psi_{t+1}^{(s)}$ for all s . Each set either grows by one or remains the same.

Algorithm 2: SUPNTK-UCB Algorithm

 T number of total steps, $S = \log T$ number of discretization levels

```

Initialize  $\Psi_1^{(s)} \leftarrow \emptyset \forall s \leq S$ 
for  $t = 1$  to  $T$  do
     $m \leftarrow 1, A_1 \leftarrow \mathcal{A}$ 
    while action  $\mathbf{a}_t$  is not chosen do
         $(\mu_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}), \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a})) \leftarrow \text{GetPosteriors}(\Psi_t^{(s)}, \mathbf{z}_t)$  for all  $\mathbf{a} \in A_s$ 
        if  $\forall \mathbf{a} \in A_s, \sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) \leq \frac{\sigma}{\sqrt{T}}$  then
            Choose  $\mathbf{a}_t = \arg \max_{\mathbf{a} \in A_s} \mu_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) + \sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a})$ .
            Keep the index sets  $\Psi_{t+1}^{(s')} = \Psi_t^{(s')}$  for all  $s' \leq S$ 
        end
        else if  $\forall \mathbf{a} \in A_s, \sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) \leq \sigma 2^{-s}$  then
             $A_{s+1} \leftarrow \left\{ \mathbf{a} \in A_s \mid \mu_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) + \sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) \geq \max_{\mathbf{a} \in A_s} \mu_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) + \sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) - \sigma 2^{1-s} \right\}$ 
             $s \leftarrow s + 1$ 
        end
        else
            Choose  $\mathbf{a}_t \in A_s$  such that  $\sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}_t) > \sigma 2^{-s}$ .
            Update the index sets at all levels  $s' \leq S$ ,
             $\Psi_{t+1}^{(s')} = \begin{cases} \Psi_{t+1}^{(s)} \cup \{t\} & \text{if } s' = s \\ \Psi_t^{(s')} & \text{otherwise} \end{cases}$ 
        end
    end
end
end
    
```

- For every level s , the set A_s includes \mathbf{a}_t^* the true maximizer of the reward with high probability. At every step t we start with A_1 which includes all the actions, and start removing actions which have a small UCB and are unlikely to be \mathbf{a}_t^* .
- The UCB strategy is only used if the learner is certain about the outcome of all actions within A_s , i.e. $\sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) \leq \sigma / \sqrt{T \beta_t}$, for all $\mathbf{a} \in A_s$. The context-reward pairs of these UCB steps are not saved for future estimation of posteriors, i.e. $\Psi_{t+1}^{(s)} = \Psi_t^{(s)}$.
- At step t , if there are actions $\mathbf{a} \in A_s$ for which $\sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) > \sigma 2^{-s}$, then one is chosen at random, and the set $\Psi_t^{(s)}$ is updated with the index t , while all other sets remain the same.
- The last case of the *if* statement in Algorithm 2 considers a middle ground, when the learner is not certain enough to pick an action by maximizing the UCB, but for all $\mathbf{a} \in A_s$ posterior variance is smaller than $\sigma 2^{-s} / \sqrt{\beta_t}$. In this case, the level s is updated as $s \leftarrow s + 1$. In doing so, the learner considers a finer uncertainty level, and updates its criterion for *closeness* to the optimal action.
- The parameter s discretizes the levels of uncertainty. For instance, in the construction of $\Psi_t^{(s)}$, the observed context-reward pairs at steps t are essentially partitioned based on which $[2^{-(s+1)}, 2^{-s}]$ interval $\sigma_{t-1}^{(s)}$ belongs to. If for all $s \leq S$, $\sqrt{\beta_t} \sigma_{t-1}^{(s)} \leq \sigma 2^{-s}$, then that pair is disregarded. Otherwise, it is added to $\Psi_{t+1}^{(s)}$ with the smallest s , for which $\sqrt{\beta_t} \sigma_{t-1}^{(s)} \leq \sigma 2^{-s}$. We set $S = \log T$, ensuring that $\frac{\sigma}{\sqrt{T}} \leq \sigma 2^{-S}$.

Properties of the SupNTK-UCB The construction of this algorithm guarantees properties that will later facilitate the proof of a $(\sqrt{T} \gamma_T)$ regret bound. These properties are given formally in Lemma C.4 and Proposition C.8, here we give an overview. SUPNTK-UCB satisfies that for every $t \leq T$ and $s \leq S$:

1. The true maximizer of reward remains within the set of plausible actions, i.e., $\mathbf{a}_t^* \in A_s$.

2. Given the context \mathbf{z}_t , regret of the actions $\mathbf{a} \in A_s$, is bounded by 2^{3-s} ,

with high probability over the observation noise. Let $\mathbf{X}^{(s)}$ denote sequence of \mathbf{x}_τ with $\tau \in \Psi_t^{(s)}$. Conveniently, the construction of $\Psi_t^{(s)}$ guarantees that,

3. Given $\mathbf{X}^{(s)}$, the corresponding rewards $\mathbf{y}^{(s)}$ are independent random variables and $\mathbb{E}y_\tau = f(\mathbf{x}_\tau)$.

4. Cardinality of each uncertainty set $|\Psi_t^{(s)}|$, is bounded by $\mathcal{O}(\gamma_T \log T)$.

C.4 Proof of Theorem 3.3

Our proof adapts the technique in Valko et al. (2013). Consider the average cumulative regret given the inputs, by property 3 of the algorithm we may write it as,

$$\mathbb{E}[R_T | X_T] = \sum_{t \in \bar{\Psi}} f(\mathbf{x}_t^*) - f(\mathbf{x}_t) + \sum_{t \in [T]/\bar{\Psi}} f(\mathbf{x}_t^*) - f(\mathbf{x}_t)$$

where $\bar{\Psi} := \{t \leq T | \forall s, t \notin \Psi_T^{(s)}\}$ includes the indices of steps with small posterior variance, i.e. $\sigma_t(\mathbf{x}) \leq \sigma/\sqrt{T\beta_t}$. For bounding the first term, we use Azuma-Hoeffding to control $f(\mathbf{x}_t^*) - f(\mathbf{x}_t)$, with $\sigma_t(\mathbf{x}_t)C(B, \sqrt{\gamma_T}, \beta_T)$. Since σ_{t-1} is small for $t \in \bar{\Psi}$, this term grows slower than $\mathcal{O}(\sqrt{\gamma_T T})$. We then use properties 2 and 4, to bound the second term. Having bounded $\mathbb{E}[R_T | X_T]$, we again use Azuma-Hoeffding, to give a bound on the regret R_T .

For this proof, we use both feature map and kernel function notation. Let \mathcal{H}_k be the RKHS corresponding to k and the sequence $\phi = (\sqrt{\lambda_i} \phi_i)_{i=1}^\infty$, be an algebraic orthogonal basis for \mathcal{H}_k , such that $k(\mathbf{x}, \mathbf{x}') = \phi^T(\mathbf{x})\phi(\mathbf{x}')$. For $f \in \mathcal{H}_k$ and there exists a unique sequence θ , such that $f = \phi^T \theta$. If $\|f\|_k \leq B$ then $\|\theta\| \leq B$, since

$$\|f\|_k^2 = \sum_i \frac{(\langle f, \phi_i \rangle)^2}{\lambda_i} = \sum_i \frac{(\langle \sum_j \sqrt{\lambda_j} \theta_j \phi_j, \phi_i \rangle)^2}{\lambda_i} = \sum_i \theta_i^2 = \|\theta\|^2$$

The following lemmas will be used to prove the theorem.

Lemma C.3. Fix $s \leq S$, for any action $\mathbf{a} \in \mathcal{A}$, let $\mathbf{x} = \mathbf{z}_t \mathbf{a}$. Then with probability of at least $1 - 2|\mathcal{A}|e^{-\beta_T/2}$,

$$|\mu_t^{(s)}(\mathbf{x}) - f(\mathbf{x})| \leq \sigma_t^{(s)}(\mathbf{x}) [2\sqrt{\beta_t} + \sigma] B.$$

Lemma C.4. For any $t \leq T$, and $s \leq S$, with probability greater than $1 - 2|\mathcal{A}| \exp^{-\beta_T/2}$,

1. $\mathbf{a}_t^* \in A_s$.

2. For all $\mathbf{a} \in A_s$, given \mathbf{z}_t , $f(\mathbf{x}_t^*) - f(\mathbf{x}) \leq 2^{3-s}$

Lemma C.5. For any $s \leq S$, the cardinality of $\Psi_T^{(s)}$ is grows with T as follows,

$$|\Psi_T^{(s)}| \leq \mathcal{O}\left(4^s \gamma_T \log T\right)$$

Lemma C.6. Consider k_{NN} defined on $\mathcal{X} \subset \mathbb{S}^{d-1}$, and its corresponding RKHS, $\mathcal{H}_{k_{\text{NN}}}$. Any $f \in \mathcal{H}_{k_{\text{NN}}}$ where $\|f\|_{k_{\text{NN}}} \leq B$, is uniformly bounded by B over \mathcal{X} .

Lemma C.7 (Azuma-Hoeffding Inequality). Let X_1, \dots, X_T be random variables with $X_t \leq a_t$ for some $a_t > 0$. Then with probability greater than $1 - 2 \exp\left(\frac{-B^2}{2 \sum_t a_t^2}\right)$,

$$\left| \sum_t X_t - \sum_t \mathbb{E}[X_t | X_1, \dots, X_{t-1}] \right| \leq B$$

Proof of Theorem 3.3. Denote by H_{t-1} the history of the algorithm at time t ,

$$H_{t-1} := \left\{ (\mathbf{z}_i, \mathbf{a}_i, y_i) \right\}_{i < t} \cup \mathbf{z}_t$$

Define $X_t = f(\mathbf{x}_t^*) - f(\mathbf{x}_t)$. By Lemma C.6, f is bounded and we have $|X_t| \leq B$. Then by applying Azuma-Hoeffding (Lemma C.7) on the random variables X_1, \dots, X_T , with probability greater than $1 - 2T|\mathcal{A}|e^{-\beta_T/2}$,

$$|R(T) - \mathbb{E}[R(T)|H_{T-1}]| \leq \sqrt{2TB^2 \log\left(\frac{1}{T|\mathcal{A}|e^{-\beta_T/2}}\right)} \quad (\text{C.7})$$

We now use lemmas C.3, C.5, and C.4 to bound the growth rate of $\mathbb{E}[R(T)|H_{t-1}]$. Recall $\bar{\Psi} := \{t \leq T | \forall s, t \notin \Psi_T^{(s)}\}$.

$$R(T) = \sum_{t \in [T]/\bar{\Psi}} f(\mathbf{x}_t^*) - f(\mathbf{x}_t) + \sum_{t \in \bar{\Psi}} f(\mathbf{x}_t^*) - f(\mathbf{x}_t) \quad (\text{C.8})$$

Bounding the expectation of the first sum gives,

$$\begin{aligned} \sum_{t \in [T]/\bar{\Psi}} \mathbb{E}[f(\mathbf{x}_t^*) - f(\mathbf{x}_t)|H_{t-1}] &= \sum_{s \leq S} \sum_{t \in \Psi_T^{(s)}} f(\mathbf{x}_t^*) - f(\mathbf{x}_t) \\ &\leq \sum_{s \leq S} 2^{3-s} |\Psi_T^{(s)}| \\ &\leq 8S \sqrt{\beta_T(10 + \sigma^{-2}15)} \gamma_T T \log T \end{aligned} \quad (\text{C.9})$$

with a probability of at least $1 - 2ST|\mathcal{A}|e^{-\beta_T/2}$. The First equation holds by Prop. C.8, and the second holds due to Lemma C.4. For the third, we have used the inequality in Lemma C.9 and that $|\Psi_T^{(s)}| \leq T$. We now bound expectation of the second term in Equation C.8. By Lemma C.3, with probability of greater than $1 - 2T|\mathcal{A}|e^{-\beta_T/2}$,

$$\begin{aligned} \sum_{t \in \bar{\Psi}} \mathbb{E}[f(\mathbf{x}_t^*) - f(\mathbf{x}_t)|H_{t-1}] &= \sum_{t \in \bar{\Psi}} f(\mathbf{x}_t^*) - f(\mathbf{x}_t) \\ &\stackrel{(a)}{\leq} \sum_{t \in \bar{\Psi}} \mu_t(\mathbf{x}_t^*) + \sqrt{\beta_t} \sigma_t(\mathbf{x}_t^*) + B(2\sqrt{\beta_t} + \sigma) \sigma_t(\mathbf{x}_t^*) - f(\mathbf{x}_t) \\ &\stackrel{(b)}{\leq} \sum_{t \in \bar{\Psi}} \mu_t(\mathbf{x}_t) + \sqrt{\beta_t} \sigma_t(\mathbf{x}_t) + B(2\sqrt{\beta_t} + \sigma) \sigma_t(\mathbf{x}_t^*) - f(\mathbf{x}_t) \\ &\stackrel{(c)}{\leq} \sum_{t \in \bar{\Psi}} \sqrt{\beta_t} \sigma_t(\mathbf{x}_t) + B(2\sqrt{\beta_t} + \sigma) (\sigma_t(\mathbf{x}_t^*) + \sigma_t(\mathbf{x}_t)) \\ &\stackrel{(d)}{\leq} \sum_{t \in \bar{\Psi}} \frac{\sigma}{\sqrt{T}} + 2B(2 + \frac{\sigma}{\sqrt{\beta_T}}) \frac{\sigma}{\sqrt{T}} \\ &\leq \sigma \sqrt{T} (1 + 2B(2 + \frac{\sigma}{\sqrt{\beta_T}})) \end{aligned} \quad (\text{C.10})$$

For inequalities (a) and (c), we have used Lemma C.3, for \mathbf{x}_t^* and \mathbf{x}_t respectively. By Algorithm 2, if $t \in \bar{\Psi}$, then \mathbf{x}_t is the maximizer of the upper confidence bound, resulting in inequality (b). Lastly for inequality (d), by construction of $\bar{\Psi}$, we have that $\sqrt{\beta_t} \sigma_{t-1}^{(s)}(\mathbf{x}) \leq \frac{\sigma}{\sqrt{T}}$. We plug in $S = \log T$, $\beta_t = \beta_T = 2 \log(2T|\mathcal{A}|/\delta)$ and substitute δ with $\delta/(1 + \log T)$. Putting together equations C.7, C.9, and C.10 gives the result. \square

C.4.1 Proof of Lemma C.3

SUPNTK-UCB is constructed such that the proposition below holds following the results in Valko et al. (2013).

Proposition C.8 (Lemma 6 Valko et al. (2013)). Consider the SUPNTK-UCB algorithm. For all $t \leq T$, $s \leq S$, and for a fixed sequence of \mathbf{x}_t where $t \in \Psi_t^{(s)}$. The corresponding rewards y_t are independent random variables, and we have $\mathbb{E}[y_t|\mathbf{x}_t] = f(\mathbf{x}_t)$.

Proof of Lemma C.3. Let \mathbf{k} , \mathbf{K} , and \mathbf{y} be defined as they are in Algorithm 1, then by definition of ϕ ,

$$\begin{aligned} \mu_t^{(s)}(\mathbf{x}) &= \mathbf{k}^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y} \\ &= \phi^T(\mathbf{x}) (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \Phi^T \mathbf{y} \\ &= \phi^T(\mathbf{x}) (\Phi^T \Phi + \sigma^2 \mathbf{I})^{-1} \Phi^T \mathbf{y} \end{aligned} \quad (\text{C.11})$$

where $\Phi = [\phi^T(\mathbf{x}_i)]_{i \in \Psi_t^{(s)}}^T$. For simplicity let $C = \Phi^T \Phi + \sigma^2 \mathbf{I}$. Similarly for $\sigma_t^{(s)}(\mathbf{x})$ we can write:

$$\begin{aligned} \sigma_t^{(s)}(\mathbf{x}) &= \sqrt{\phi^T(\mathbf{x}) C^{-1} \phi(\mathbf{x})} \\ &= \sqrt{\phi^T(\mathbf{x}) C^{-1} (\Phi^T \Phi + \sigma^2 \mathbf{I}) C^{-1} \phi(\mathbf{x})} \\ &\geq \|\Phi C^{-1} \phi(\mathbf{x})\| \end{aligned} \tag{C.12}$$

Using Equation C.11, we get

$$\begin{aligned} \mu_t^{(s)}(\mathbf{x}) - f(\mathbf{x}) &= \phi^T(\mathbf{x}) (\Phi^T \Phi + \sigma^2 \mathbf{I})^{-1} \Phi^T \mathbf{y} - \phi^T \boldsymbol{\theta} \\ &= \phi^T(\mathbf{x}) C^{-1} \Phi^T \mathbf{y} - \phi^T C^{-1} C \boldsymbol{\theta} \\ &= \phi^T(\mathbf{x}) C^{-1} \Phi^T (\mathbf{y} - \Phi \boldsymbol{\theta}) - \sigma^2 \phi^T(\mathbf{x}) C^{-1} \boldsymbol{\theta} \end{aligned} \tag{C.13}$$

We now bound the first term in Equation C.13. By Proposition C.8, conditioned on \mathbf{x} and $[\mathbf{x}_i]_{i \in \Psi_t}$, $\mathbf{y} - \Phi \boldsymbol{\theta}$ is a vector of zero-mean independent random variables. By Lemma C.6, f is bounded. Similar to Valko et al. (2013), and without loss of generality, we may normalize the vector \mathbf{y} over $\Psi_t^{(s)}$, and assume that $y_i \leq B$. Then, each $|y_i - \phi^T(\mathbf{x}_i) \boldsymbol{\theta}| \leq 2B$. By Equation C.12, and Lemma C.7, with probability greater than $1 - 2 \exp(-\beta_T/2)$,

$$|\phi^T(\mathbf{x}) (\Phi^T \Phi + \sigma^2 \mathbf{I})^{-1} \Phi^T (\mathbf{y} - \Phi \boldsymbol{\theta})| \leq 2B \sqrt{\beta_t} \sigma_t^{(s)}(\mathbf{x}) \tag{C.14}$$

For the second term, by Cauchy-Schwartz we have,

$$\begin{aligned} \sigma^2 \phi^T(\mathbf{x}) (\Phi^T \Phi + \sigma^2 \mathbf{I})^{-1} \boldsymbol{\theta} &= \sigma^2 \phi^T(\mathbf{x}) C^{-1} \boldsymbol{\theta} \\ &\leq B \sigma^2 \sqrt{\phi^T(\mathbf{x}) C^{-1} \sigma^{-2} \sigma^2 \mathbf{I} C^{-1} \phi(\mathbf{x})} \\ &\leq B \sigma \sqrt{\phi^T(\mathbf{x}) C^{-1} C C^{-1} \phi(\mathbf{x})} \\ &\leq B \sigma \sigma_t^{(s)}(\mathbf{x}) \end{aligned} \tag{C.15}$$

The proof is concluded by plugging in Equations C.14 and C.15 into Equation C.13, and taking a union bound over all $|\mathcal{A}|$ actions. \square

C.4.2 Proof of Other Lemmas

Proof of lemma C.4. We prove this lemma by showing the equivalent of NTK-UCB and Kernelized UCB and refer to Lemma 7 in Valko et al. (2013). Consider the GP regression problem with gaussian observation noise, i.e., $y_i = f(\mathbf{x}_i) + \varepsilon_i$, with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$, and $f \sim \text{GP}(0, k)$. Let $\phi(\cdot)$ be the feature map of the GP's covariance kernel function k , i.e. $k(\mathbf{x}, \mathbf{x}') = \phi^T(\mathbf{x}) \phi(\mathbf{x}')$. Denote $\mu_t^{(\text{GP})}(\cdot)$ as the posterior mean function, after observing t samples of (\mathbf{x}_i, y_i) pairs. Then, it is straightforward to show that, for all $\mathbf{x} \in \mathcal{X}$,

$$\mu_t^{(\text{GP})}(\mathbf{x}) = \mu_t^{(\text{Ridge})}(\mathbf{x})$$

Where $\mu_t^{(\text{Ridge})}(\mathbf{x}) = \phi^T(\mathbf{x}) \boldsymbol{\theta}^*$, with $\boldsymbol{\theta}^*$ being the minimizer of the kernel ridge loss,

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^t (y_i - \phi^T(\mathbf{x}_i) \boldsymbol{\theta})^2 + \sigma^2 \|\boldsymbol{\theta}\|^2.$$

Similarly, let $\sigma_t^{(\text{GP})}(\mathbf{x})$ be the posterior variance function of the GP regression. Using matrix identities we have,

$$\sigma_t^{(\text{GP})}(\mathbf{x}) = \sqrt{\phi^T(\mathbf{x}) (K_T + \sigma^2 \mathbf{I})^{-1} \phi(\mathbf{x})} = \sigma_t^{(\text{Ridge})}(\mathbf{x})$$

which is the width of confidence interval used in KERNELIZED UCB (Valko et al., 2013). Their exploration policy is then defined as,

$$x_t^{(\text{Ridge})} = \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_t^{(\text{Ridge})}(\mathbf{x}) + \frac{\beta}{\sigma} \sigma_t^{(\text{GP})}(\mathbf{x}).$$

We conclude that NTK-UCB and KERNELIZED UCB (Valko et al., 2013) are equivalent, up to a constant factor in exploration coefficient β . We modify SUPKERNELUCB to SUPNTK-UCB such that the key lemmas still hold, and Lemma C.4 immediately follows from Valko et al. (2013)'s Lemma 7. \square

Proof of Lemma C.5. Let λ_i denote the eigenvalues of $\mathbf{K}_T + \sigma^2 \mathbf{I}$, in decreasing order. Valko et al. define

$$\tilde{d} := \min\{j : j\sigma^2 \log T \geq \sum_{i \geq j} \lambda_i - \sigma^2\},$$

and show that γ_T gives a data independent upper bound on \tilde{d} ,

$$\gamma_T \geq I(\mathbf{y}_T; \mathbf{f}) \geq \Omega(\tilde{d} \log \log T)$$

Due to the equivalence of SUPKERNELUCB and SUPNTK-UCB, as shown in the proof of lemma C.4, the following lemma holds for SUPNTK-UCB.

Lemma C.9 (Lemma 5, Valko et al. (2013)). *Let $l_T = \max\{\log T, \log(T/(\sigma\tilde{d}))\}$. For all $s \leq S$,*

$$|\Psi_T^{(s)}| \leq 2^s \sqrt{\beta_T(10 + 15\sigma^{-2})\tilde{d}|\Psi_T^{(s)}|l_T}$$

By Lemma C.9, there exists T_0 such that for all $T > T_0$,

$$|\Psi_T^{(s)}| \leq 2^s \sqrt{\beta_T(10 + 15\sigma^{-2})\gamma_T|\Psi_T^{(s)}|\log T}.$$

□

Proof of Lemma C.6. By the Reproducing property of k_{NN} , for any $f \in \mathcal{H}_{k_{\text{NN}}}$ we have,

$$f(\mathbf{x}) = \langle f, k_{\text{NN}}(\mathbf{x}, \cdot) \rangle_{\mathcal{H}_{k_{\text{NN}}}} \leq \|f\|_{k_{\text{NN}}} \|k_{\text{NN}}(\mathbf{x}, \cdot)\|_{k_{\text{NN}}}$$

where the inequality holds due to Cauchy-Schwartz. The NTK is Mercer over \mathcal{X} , with the mercer decomposition $k_{\text{NN}} = \sum_i \lambda_i \phi_i$. Then by definition of inner product in $\mathcal{H}_{k_{\text{NN}}}$,

$$\begin{aligned} \|k_{\text{NN}}(\mathbf{x}, \cdot)\|_k^2 &= \sum_i \frac{\langle k_{\text{NN}}(\mathbf{x}, \cdot), \phi_i \rangle^2}{\lambda_i} \\ &= \sum_i \frac{\langle \sum_j \lambda_j \phi_j(\mathbf{x}) \phi_j, \phi_i \rangle^2}{\lambda_i} \\ &= \sum_i \lambda_i \phi_i(\mathbf{x})^2 \\ &= k_{\text{NN}}(\mathbf{x}, \mathbf{x}) \\ &= 1 \end{aligned}$$

The second to last equality uses the orthonormality of ϕ_i s, and the last equation follows from the definition of the NTK. We have $\|f\|_{k_{\text{NN}}} \leq B$ which concludes the proof. □

D Details of NN-UCB and CNN-UCB

Algorithm 3 and 6 present NN-UCB, and its Sup variant, respectively. The construction of the Sup variant is the same as for NTK-UCB, with minor changes to the conditions of the *if* statements.

D.1 Proof of Theorem 4.2

To bound the regret for NN-UCB, we define an auxiliary algorithm that allows us to use the result from NTK-UCB. Consider a kernelized UCB algorithm, which uses \hat{k} an approximation of the NTK function, where $\hat{k}(\cdot, \cdot) = \mathbf{g}^T(\cdot; \boldsymbol{\theta}^0) \mathbf{g}(\cdot; \boldsymbol{\theta}^0) / m$. We argued in the main text that this kernel can well approximate k_{NN} , and its feature map $\hat{\phi}(\mathbf{x}) = \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}^0) / \sqrt{m}$ can be viewed as a finite approximation of ϕ , the feature map of the NTK. Throughout the proof, we denote the posterior mean and variance calculated via GP($0, \hat{k}$) by $\hat{\mu}_{t-1}$ and $\hat{\sigma}_{t-1}$, respectively. In comparison to NN-UCB, we use $\hat{\mu}_{t-1}$ instead of $f^{(J)}$ to approximate the true posterior mean, however $\hat{\sigma}_{t-1}$ has the same expression as in NN-UCB. Using lemma D.3, we reduce the problem of bounding

Algorithm 3: NN-UCB

σ^2 observation noise, β_t exploration parameter, J number of GD Steps, η GD's learning rate, m width of the network, L depth of the network, T total steps of the bandit

Input: $m, L, J, \eta, \sigma, \beta_t, T$

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

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

 Observe the context \mathbf{z}_t . **for** $\mathbf{a} \in \mathcal{A}$ **do**

 Define $\mathbf{x} := \mathbf{z}_t \mathbf{a}$,

$\hat{\sigma}_{t-1}^2(\mathbf{x}) \leftarrow \mathbf{g}^T(\mathbf{x}; \theta^0) \hat{\mathbf{Z}}_{t-1}^{-1} \mathbf{g}(\mathbf{x}; \theta^0) / m$

$U_{\mathbf{a},t} \leftarrow f(\mathbf{x}; \theta_{t-1}) + \sqrt{\beta_t \hat{\sigma}_{t-1}^2(\mathbf{x})}$

end

$\mathbf{a}_t = \arg \max_{\mathbf{a} \in \mathcal{A}} U_{\mathbf{a},t}$

 Pick \mathbf{a}_t and append the rewards vector \mathbf{y}_t by the observed reward.

$\hat{\mathbf{Z}}_t \leftarrow \sigma^2 \mathbf{I} + \sum_{i \leq t} \mathbf{g}(\mathbf{x}_i; \theta^0) \mathbf{g}^T(\mathbf{x}_i; \theta^0) / m$

$\theta_t \leftarrow \text{TrainNN}(m, L, J, \eta, \sigma^2, \theta^0, [\mathbf{x}_i]_{i \leq t}, \mathbf{y}_t)$

end

Algorithm 4: TrainNN($m, L, J, \eta, \sigma^2, \theta^0, [\mathbf{x}_i]_{i \leq t}, \mathbf{y}_t$)

Input: $[\mathbf{x}_i]_{i \leq t}, \mathbf{y}_t$

Define $\mathcal{L}(\theta) = \sum_{i \in \Psi_t} (f(\mathbf{x}_i; \theta) - y_i)^2 + m \sigma^2 \|\theta - \theta^0\|_2^2$

for $j = 0, \dots, J-1$ **do**

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

end

Output: θ^J

the regret for NN-UCB to the regret of this auxiliary method. We then repeat the technique used for Theorem 3.3 on the auxiliary Sup variant which yields a regret bound depending on $\hat{\gamma}_T$, the information gain of the approximate kernel matrix. Finally, for width m large enough, we bound $\hat{\gamma}_T$ with γ_T , information gain of the exact NTK matrix.

The following lemma provides the grounds for approximating the NTK with the empirical gram matrix of the neural network at initialization. Let $\mathbf{G} = [\mathbf{g}^T(\mathbf{x}_t; \theta^0)]_{t \leq T}^T$.

Lemma D.1 (Arora et al. (2019) Theorem 3.1). *Set $0 < \delta < 1$. If $m = \Omega(L^6 \log(TL/\delta)/\epsilon^4)$, then with probability greater than $1 - \delta$,*

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

The next lemma allows us to write the unknown reward as a linear function of the feature map over the finite set of points that are observed while running the algorithm.

Lemma D.2 (Zhou et al. (2020) Lemma 5.1). *Let f^* be a member of $\mathcal{H}_{k_{\text{NN}}}$ with bounded RKHS norm $\|f\|_{k_{\text{NN}}} \leq B$. If for some constant C ,*

$$m \geq \frac{CT^4 |\mathcal{A}|^4 L^6}{\lambda_0^4} \log(T^2 |\mathcal{A}|^2 L / \delta),$$

then for any $\delta \in (0, 1)$, there exists $\theta^ \in \mathbb{R}^p$ such that*

$$f^*(\mathbf{x}^i) = \langle \mathbf{g}(\mathbf{x}^i; \theta^0), \theta^* \rangle, \quad \sqrt{m} \|\theta^*\|_2 \leq \sqrt{2}B$$

with probability greater than $1 - \delta$, for all $i \leq T|\mathcal{A}|$.

The following lemma acts as the bridge between NN-UCB and the auxiliary UCB algorithm.

Lemma D.3. *Fix $s \leq S$. Consider a given context set, $\{\mathbf{x}_\tau\}_{\tau \in \Psi_t^{(s)}}$. Assume construction of $\Psi_t^{(s)}$ is such that the corresponding rewards, y_τ are statistically independent. Then there exists C_1 , such that for any $\delta > 0$, if the*

Algorithm 5: GetApproxPosterior

 J number of GD Steps, η GD's learning rate, m width of the network, L depth of the network

Input: $\Psi_t \subset \{1, \dots, t-1\}$, \mathbf{z}_t
Initialize $\hat{\mathbf{Z}} \leftarrow \sigma^2 \mathbf{I} + \sum_{i \in \Psi_t} \mathbf{g}(\mathbf{x}_i, \boldsymbol{\theta}^0) \mathbf{g}^T(\mathbf{x}_i, \boldsymbol{\theta}^0) / m$
 $\boldsymbol{\theta}^{(J)} \leftarrow \text{TrainNN}(m, L, J, \eta, \sigma^2, \boldsymbol{\theta}^0, [\mathbf{x}_i]_{i \in \Psi_t}, [y_i]_{i \in \Psi_t})$
for $\mathbf{a} \in \mathcal{A}$ **do**
 Define $\mathbf{x} := \mathbf{z}_t \mathbf{a}$,
 $\hat{\sigma}_{t-1}^2(\mathbf{x}) \leftarrow \frac{\mathbf{g}^T(\mathbf{x}, \boldsymbol{\theta}^0)}{\sqrt{m}} \hat{\mathbf{Z}}^{-1} \frac{\mathbf{g}(\mathbf{x}, \boldsymbol{\theta}^0)}{\sqrt{m}}$
 Output: $f(\mathbf{x}, \boldsymbol{\theta}^{(J)})$ and $\hat{\sigma}_{t-1}^2(\mathbf{x})$
end

learning rate is picked $\eta = C_1(LmT + m\sigma^2)^{-1}$, and

$$m \geq \text{Poly}\left(T, L, |\mathcal{A}|, \sigma^{-2}, \log(1/\delta)\right).$$

Then with probability of at least $1 - \delta$, for all $i \leq T|\mathcal{A}|$,

$$|f(\mathbf{x}^i; \boldsymbol{\theta}^{(J)}) - \hat{\mu}^{(s)}(\mathbf{x}^i)| \leq \hat{\sigma}^{(s)}(\mathbf{x}^i) \sqrt{\frac{TB}{m\eta\sigma^2}} (3 + \sqrt{2}(1 - m\eta\sigma^2)^{J/2}) + \bar{C} \left(\frac{TB}{m\sigma^2}\right)^{2/3} L^3 \sqrt{m \log m}$$

for some constant \bar{C} , where $\hat{\mu}^{(s)}$ and $\hat{\sigma}^{(s)}$ are the posterior mean and variance of the reward with $\text{GP}(0, \hat{k})$ prior, after observing $(\mathbf{x}_\tau, y_\tau)_{\tau \in \Psi_t^{(s)}}$.

The following lemma provides the central concentration inequality and is the *neural* equivalent of Lemma C.3.

Lemma D.4 (Concentration of f and $f^{(J)}$, Formal). *Fix $s \leq S$. Consider a given context set, $\{\mathbf{x}_\tau\}_{\tau \in \Psi_t^{(s)}}$. Assume construction of $\Psi_t^{(s)}$ is such that the corresponding rewards, y_τ are statistically independent. Let $\delta > 0$, $\eta = C_1(LmT + m\sigma^2)^{-1}$, and*

$$m \geq \text{Poly}\left(T, L, |\mathcal{A}|, \lambda_0^{-1}, \sigma^{-2}, \log(1/\delta)\right).$$

Then for any action $\mathbf{a} \in \mathcal{A}$, and for some constant \bar{C} with probability of at least $1 - 2|\mathcal{A}|e^{-\beta T/2} - \delta$,

$$|f(\mathbf{x}; \boldsymbol{\theta}^{(J)}) - f^*(\mathbf{x})| \leq \hat{\sigma}^{(s)}(\mathbf{x}) \left(2B\sqrt{\beta_T} + \sigma\sqrt{\frac{2}{m}}B + \sqrt{\frac{TB}{m\eta\sigma^2}} (3 + \sqrt{2}(1 - m\eta\sigma^2)^{J/2}) \right) + \bar{C} \left(\frac{TB}{m\sigma^2}\right)^{2/3} L^3 \sqrt{m \log m}$$

where $\mathbf{x} = \mathbf{z}_t \mathbf{a}$.

Lemma D.5. *If for any $0 < \delta < 1$*

$$m = \Omega\left(T^6 L^6 \log(TL/\delta)\right),$$

then with probability greater than $1 - \delta$

$$\hat{\gamma}_T \leq \gamma_T + \sigma^{-2}$$

We are now ready to give the proof.

Proof of Theorem 4.2. Construction of $\Psi_t^{(s)}$ is the same in both Algorithm 6 and Algorithm 2. Hence, Proposition C.8 immediately follows. It is straightforward to show that lemmas C.5, C.4, C.9 all apply to SUPNN-UCB. We consider them for the approximate feature map $\hat{\phi}$ and use lemma D.2 to write the unknown reward $f^* = \phi^T(\mathbf{x})\boldsymbol{\theta}^*$ with high probability. By the union bound, all lemmas hold for the SUPNN-UCB setting, with a probability greater than $1 - 2T|\mathcal{A}|e^{-\beta/2} - \delta$ for any $0 < \delta < 1$.

Algorithm 6: Sup variant for NN-UCB

 T number of total steps, $S = 2 \log T$ number of discretization levels

Initialize $\Psi_1^{(s)} \leftarrow \emptyset, \forall s \leq S$
for $t = 1$ **to** T **do**
 $s \leftarrow 1, A_1 \leftarrow \mathcal{A}$
while action \mathbf{a}_t not chosen **do**
 $(f(\mathbf{z}_t \mathbf{a}, \boldsymbol{\theta}^{(j)}), \hat{\sigma}_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a})) \leftarrow \text{GetApproxPosterior}(\Psi_t^{(s)}, \mathbf{z}_t)$ for all $\mathbf{a} \in A_s$
if $\forall \mathbf{a} \in A_s, \sqrt{\beta_t} \hat{\sigma}_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) \leq \frac{\sigma}{T^2}$ **then**
 $\text{Choose } \mathbf{a}_t = \arg \max_{\mathbf{a} \in A_s} f(\mathbf{z}_t \mathbf{a}, \boldsymbol{\theta}^{(j)})(\mathbf{z}_t \mathbf{a}) + \sqrt{\beta_t} \hat{\sigma}_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}),$
 $\text{Keep the index sets } \Psi_{t+1}^{(s')} = \Psi_t^{(s')} \text{ for all } s' \leq S.$
end
else if $\forall \mathbf{a} \in A_s, \sqrt{\beta_t} \hat{\sigma}_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) \leq \sigma 2^{-s}$ **then**
 $A_{s+1} \leftarrow \left\{ \mathbf{a} \in A_s \mid f(\mathbf{z}_t \mathbf{a}, \boldsymbol{\theta}^{(j)}) + \sqrt{\beta_t} \hat{\sigma}_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) \geq \max_{\mathbf{a} \in A_s} f(\mathbf{z}_t \mathbf{a}, \boldsymbol{\theta}^{(j)}) + \sqrt{\beta_t} \hat{\sigma}_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}) - \sigma 2^{1-s} \right\}$
 $s \leftarrow s + 1$
end
else
 $\text{Choose } \mathbf{a}_t \in A_s \text{ such that } \sqrt{\beta_t} \hat{\sigma}_{t-1}^{(s)}(\mathbf{z}_t \mathbf{a}_t) > \sigma 2^{-s}.$
 $\text{Update the index sets at all levels } s' \leq S,$

$$\Psi_{t+1}^{(s')} = \begin{cases} \Psi_{t+1}^{(s)} \cup \{t\} & \text{if } s' = s \\ \Psi_{t+1}^{(s)} & \text{otherwise} \end{cases}$$
end
end
end

 Recall H_{t-1} is the history of the algorithm at time t ,

$$H_{t-1} := \left\{ (\mathbf{z}_i, \mathbf{s}_i, y_i) \right\}_{i < t} \cup \mathbf{z}_t$$

 Similar to proof of Theorem 3.3, we apply the Azuma-Hoeffding bound (Lemma C.7) to the random variables, $X_t =: f^*(\mathbf{x}_t^*) - f^*(\mathbf{x}_t)$. We get, with probability greater than $1 - 2T|\mathcal{A}|e^{-\beta T/2}$,

$$|R(T) - \mathbb{E}[R(T)|H_{T-1}]| \leq \sqrt{4T \log\left(\frac{1}{T|\mathcal{A}|e^{-\beta T/2}}\right)} \quad (\text{D.1})$$

 We now bound $\mathbb{E}[R(T)|H_{t-1}]$. Let $\bar{\Psi} := \{t \leq T \mid \forall s, t \notin \Psi_T^{(s)}\}$. By lemmas C.4 and C.9 applied to \hat{k} the approximate kernel, with probability of at least $1 - ST|\mathcal{A}|e^{-\beta T/2}$,

$$\begin{aligned} \mathbb{E}[R(T)|H_{t-1}] &= \sum_{\substack{\text{all} \\ t \notin \bar{\Psi}}} f^*(\mathbf{x}_t^*) - f^*(\mathbf{x}_t) + \sum_{t \in \bar{\Psi}} f^*(\mathbf{x}_t^*) - f^*(\mathbf{x}_t) \\ &\leq 8S\sqrt{\beta_T(10 + \sigma^{-2}15)}\hat{\gamma}_T T \log T + \underbrace{\sum_{t \in \bar{\Psi}} f^*(\mathbf{x}_t^*) - f^*(\mathbf{x}_t)}_I, \end{aligned}$$

 where $\hat{\gamma}_T$ is the information gain corresponding to the approximate kernel. Applying Lemma D.4, with probability

of greater than $1 - 2T|\mathcal{A}|e^{-\beta_T/2}$,

$$\begin{aligned}
 I &\leq \sum_{t \in \bar{\Psi}} (\hat{\sigma}^{(s)}(\mathbf{x}_t) + \hat{\sigma}^{(s)}(\mathbf{x}_t^*)) \left(2B\sqrt{\beta_T} + \sigma\sqrt{\frac{2}{m}}B + \sqrt{\frac{TB}{m\eta\sigma^2}}(3 + \sqrt{2}(1 - m\eta\sigma^2)^{J/2}) \right) \\
 &\quad + \sqrt{\beta_T}\hat{\sigma}^{(s)}(\mathbf{x}_t) + 2\bar{C}\left(\frac{TB}{m\sigma^2}\right)^{2/3}L^3\sqrt{m\log m} \\
 &\leq \sum_{t \in \bar{\Psi}} 4B\frac{\sigma}{T^2} + \frac{2\sigma^2B}{T^2}\sqrt{\frac{2}{m\beta_T}} + 2\sqrt{\frac{B}{m\eta\beta_T T^2}}(3 + \sqrt{2}(1 - m\eta\sigma^2)^{J/2}) \\
 &\quad + 2\bar{C}\left(\frac{TB}{m\sigma^2}\right)^{2/3}L^3\sqrt{m\log m} \\
 &\leq 4B\frac{\sigma}{T} + 2\sigma^2B\sqrt{\frac{2}{m\beta_T T^2}} + 2\sqrt{\frac{B}{m\eta\beta_T}}(3 + \sqrt{2}(1 - m\eta\sigma^2)^{J/2}) \\
 &\quad + 2\bar{C}T\left(\frac{TB}{m\sigma^2}\right)^{2/3}L^3\sqrt{m\log m}
 \end{aligned}$$

The next to last inequality holds by the construction of $\bar{\Psi}$, and the last one by $|\bar{\Psi}| \leq T$. Set $\delta \leq T|\mathcal{A}|e^{-\beta_T/2}$, $S = \log T$ and choose m large enough according to the conditions of Theorem 4.2. Putting together the two terms,

$$\begin{aligned}
 R(T) &\leq 4B\frac{\sigma}{T} + 2\sqrt{\frac{B(TL + \sigma^2)}{C_1\beta_T}}(3 + \sqrt{2}(1 - \frac{\sigma^2 C_1}{TL + \sigma^2})^{J/2}) \\
 &\quad + 8\sqrt{\beta_T(10 + \sigma^{-2}15)\hat{\gamma}_T T(\log T)^3} + \sqrt{4T\log\left(\frac{1}{T|\mathcal{A}|e^{-\beta_T/2}}\right)}
 \end{aligned}$$

Choosing $\beta_t = \beta_T = 2\log(2T|\mathcal{A}|/\tilde{\delta})$, with probability greater than $1 - (\log T + 1)\tilde{\delta}$,

$$\begin{aligned}
 R(T) &\leq 2(1 + B)\frac{\sigma}{T} + 8\sqrt{2\log(2T|\mathcal{A}|/\tilde{\delta})}\sqrt{\hat{\gamma}_T T(\log T)^3(10 + \sigma^{-2}15)} \\
 &\quad + 2\sqrt{\frac{B(TL + \sigma^2)}{2C_1\log(2T|\mathcal{A}|/\tilde{\delta})}}(3 + \sqrt{2}(1 - \frac{\sigma^2 C_1}{TL + \sigma^2})^{J/2}) + 2\sqrt{T\log(2/\tilde{\delta})}.
 \end{aligned}$$

We conclude the proof by using Lemma D.5, to bound $\hat{\gamma}_T$ with γ_T , and substitute $\tilde{\delta}$ with $\tilde{\delta}/(1 + \log T)$. \square

Before giving the proof of the Lemmas, we present Condition D.6. Lemmas D.2 through D.5, often rely on the width m being large enough, or the learning rate being small enough. For easier readability, we put those inequalities together and present it at m satisfying a certain polynomial rate and η being of order $1/m$.

Condition D.6. The following conditions on m and η , ensure convergence in the lemmas used in this section. For every $t \leq T$,

$$\begin{aligned}
 2\sqrt{t/(m\sigma^2)} &\geq C_1 \left[\frac{mL}{\log(TL^2|\mathcal{A}|/\sigma^2)} \right]^{-3/2}, \\
 2\sqrt{t/(m\sigma^2)} &\leq C_2 \min \{ L^{-6}(\log m)^{-3/2}, (m(\eta\sigma^2)^2 L^{-6} t^{-1} (\log m)^{-1})^{3/8} \}, \\
 \eta &\leq C_3 (m\sigma^2 + tmL)^{-1}, \\
 m^{1/6} &\geq C_4 \sqrt{\log mL}^{7/2} (t/\sigma^2)^{7/6} (1 + \sqrt{t/\sigma^2}).
 \end{aligned}$$

Setting m and η as follows satisfies the inequalities above.

$$\begin{aligned}
 m &\geq \text{Poly}\left(T, L, |\mathcal{A}|, \sigma^{-2}, \log(1/\delta)\right), \\
 \eta &= C(mTL + m\sigma^2)^{-1}
 \end{aligned}$$

D.1.1 Proof of lemma D.3

We first give an overview of the proof. Recall that $\boldsymbol{\theta}^{(J)}$ is the result of running gradient descent for J steps on

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i \leq t} (f(\mathbf{x}_i; \boldsymbol{\theta})^2 - y_i)^2 + m\sigma^2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^0\|_2^2.$$

The network $f(\mathbf{x}; \boldsymbol{\theta})$ is a complex nonlinear function, and it is not feasible to produce the expression for the updates $\boldsymbol{\theta}^j$ at a step j of the gradient descent algorithm. Working with the first order Taylor expansion of the network around initialization, $\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}^0), \boldsymbol{\theta} - \boldsymbol{\theta}^0 \rangle$, we instead consider running the gradient descent on

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}) = \sum_{i \leq t} (\langle \mathbf{g}^T(\mathbf{x}_i; \boldsymbol{\theta}^0), \boldsymbol{\theta} - \boldsymbol{\theta}^0 \rangle - y_i)_2^2 + m\sigma^2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^0\|_2^2.$$

Let $\tilde{\boldsymbol{\theta}}^j$ denote the GD update at step j . It can be proven that gradient descent follows a similar path in both scenarios, i.e. the sequences $(\boldsymbol{\theta}^j)_{j>1}$ and $(\tilde{\boldsymbol{\theta}}^j)_{j>1}$ remain close for all j with high probability. Now gradient descent on $\tilde{\mathcal{L}}$ converges to the global optima, $\tilde{\boldsymbol{\theta}}^*$, for which we have $\langle \mathbf{g}^T(\mathbf{x}_i; \boldsymbol{\theta}^0), \tilde{\boldsymbol{\theta}}^* - \boldsymbol{\theta}^0 \rangle = \hat{\mu}(\mathbf{x})$, concluding the proof.

Proof. Let $\mathbf{f}(\boldsymbol{\theta}) = [f(\mathbf{x}_\tau; \boldsymbol{\theta})]_{\tau \in \Psi_t^{(s)}}$, and $\mathbf{y} = [y_\tau]_{\tau \in \Psi_t^{(s)}}$. Consider the loss function used for training the network

$$\mathcal{L}_1(\boldsymbol{\theta}) = \|\mathbf{f}(\boldsymbol{\theta}) - \mathbf{y}\|_2^2 + m\sigma^2 \|\boldsymbol{\theta}\|_2^2 \quad (\text{D.2})$$

And let the sequence of $(\boldsymbol{\theta}^j)$ denote the gradient descent updates. The following lemma shows that at step j of the GD, $f(\mathbf{x}; \boldsymbol{\theta}^j)$ can be approximated with its first order Taylor approximation at initialization, $\langle \mathbf{g}(\mathbf{x}; \boldsymbol{\theta}^0), \boldsymbol{\theta}^j - \boldsymbol{\theta}^0 \rangle$.

Lemma D.7. *There exists constants $(C_i)_{i \leq 4}$, such that for any $\delta > 0$, if η and m satisfy Condition D.6, then,*

$$|f(\mathbf{x}^i; \boldsymbol{\theta}^J) - f(\mathbf{x}^i; \boldsymbol{\theta}^0) - \langle \mathbf{g}(\mathbf{x}^i; \boldsymbol{\theta}^0), \boldsymbol{\theta}^J - \boldsymbol{\theta}^0 \rangle| \leq C \left(\frac{TB}{m\sigma^2}\right)^{2/3} L^3 \sqrt{m \log m}$$

for some constant C with probability greater than $1 - \delta$, for any $i \leq T|\mathcal{A}|$.

Lemma above holds, as m and η are chosen such that condition D.6 is met. We now show that $\langle \mathbf{g}(\mathbf{x}^i; \boldsymbol{\theta}^0), \boldsymbol{\theta}^J - \boldsymbol{\theta}^0 \rangle$ approximates $\hat{\mu}^{(s)}$ well. Let

$$\begin{aligned} \mathbf{G} &= [\mathbf{g}^T(\mathbf{x}_\tau; \boldsymbol{\theta}^0)]_{\tau \in \Psi_t^{(s)}}^T, \\ \hat{\mathbf{Z}} &= \sigma^2 \mathbf{I} + \sum_{\tau \in \Psi_t^{(s)}} \mathbf{g}(\mathbf{x}_\tau; \boldsymbol{\theta}^0) \mathbf{g}^T(\mathbf{x}_\tau; \boldsymbol{\theta}^0) / m = \sigma^2 \mathbf{I} + \mathbf{G}^T \mathbf{G} / m. \end{aligned}$$

Let $\tilde{\boldsymbol{\theta}}^j$ to be the GD updates of \mathcal{L}_2 , the loss function corresponding to the auxiliary UCB algorithm

$$\mathcal{L}_2(\boldsymbol{\theta}) = \|\mathbf{G}^T(\boldsymbol{\theta} - \boldsymbol{\theta}^0) - \mathbf{y}\|_2^2 + m\sigma^2 \|\boldsymbol{\theta} - \boldsymbol{\theta}^0\|_2^2. \quad (\text{D.3})$$

Lemma D.8 states the nice convergence properties of this optimization problem.

Lemma D.8 (Zhou et al. (2020) Lemma C.4). *If η , m satisfy conditions D.6, then*

$$\begin{aligned} \left\| \tilde{\boldsymbol{\theta}}^j - \boldsymbol{\theta}^0 - \hat{\mathbf{Z}}^{-1} \mathbf{G}^T \mathbf{y} / m \right\|_2 &\leq (1 - \eta m \sigma^2)^{j/2} \sqrt{2TB / (m\sigma^2)} \\ \left\| \tilde{\boldsymbol{\theta}}^j - \boldsymbol{\theta}^0 \right\|_2 &\leq \sqrt{\frac{TB}{m\sigma^2}} \end{aligned}$$

with a probability of at least $1 - \delta$, for any $j \leq J$.

Furthermore, we can show that in space of $\boldsymbol{\theta}$, the path of gradient descent on \mathcal{L}_2 follows GD's path on \mathcal{L}_1 . In other words, as the width m grows the sequences $(\boldsymbol{\theta}^j)_{j \leq J}$ and $(\tilde{\boldsymbol{\theta}}^j)_{j \leq J}$ converge uniformly.

Lemma D.9. *There exists constants $(C_i)_{i \leq 4}$, such that for any $\delta > 0$, if η and m satisfy Condition D.6, then,*

$$\left\| \boldsymbol{\theta}^J - \tilde{\boldsymbol{\theta}}^J \right\|_2 \leq 3 \sqrt{\frac{TB}{m\sigma^2}}$$

with probability greater than $1 - \delta$.

We now have all necessary ingredients to finish the proof. For simplicity, here we denote $\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}^0)$ by \mathbf{g} . Applying Cauchy-Schwartz inequality we have,

$$\begin{aligned} \langle \mathbf{g}, \boldsymbol{\theta}^J - \boldsymbol{\theta}^0 \rangle &= \langle \mathbf{g}, \boldsymbol{\theta}^J - \tilde{\boldsymbol{\theta}}^J \rangle + \langle \mathbf{g}, \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^0 \rangle \\ &\leq \|\mathbf{g}\|_{\hat{\mathbf{Z}}^{-1}} \left\| \boldsymbol{\theta}^J - \tilde{\boldsymbol{\theta}}^J \right\|_2 + \langle \mathbf{g}, \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^0 \rangle \\ &\leq \frac{1}{\sqrt{m\eta}} \|\mathbf{g}\|_{\hat{\mathbf{Z}}^{-1}} \left\| \boldsymbol{\theta}^J - \tilde{\boldsymbol{\theta}}^J \right\|_2 + \langle \mathbf{g}, \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^0 \rangle \\ &\leq 3 \left\| \frac{\mathbf{g}}{\sqrt{m}} \right\|_{\hat{\mathbf{Z}}^{-1}} \sqrt{\frac{TB}{m\eta\sigma^2}} + \langle \mathbf{g}, \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^0 \rangle \end{aligned} \quad (\text{D.4})$$

Recall that $\hat{\mathbf{Z}} = \sigma^2 \mathbf{I} + \mathbf{G}^T \mathbf{G}$. By Lemma D.10, $\|\mathbf{G}\|_F \leq C\sqrt{TLm}$ and we have,

$$\hat{\mathbf{Z}} \preceq (\sigma^2 + CTLm) \mathbf{I} \preceq \frac{1}{m\eta} \mathbf{I}, \quad (\text{D.5})$$

since η is set such that $\eta \leq C(m\sigma^2 + TLM)^{-1}$. Therefore, for any $\mathbf{x} \in \mathbb{R}^p$, $\|\mathbf{x}\|_{\hat{\mathbf{Z}}} \leq \frac{1}{\sqrt{m\eta}} \|\mathbf{x}\|_2$, and follows the second inequality. For the third inequality we have used Lemma D.9. Define $\mathbf{b} = \sum_{i \leq T} y_i \mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}^0) / \sqrt{m}$. Decomposing the second term of the right hand side in equation D.4 we get,

$$\begin{aligned} \langle \mathbf{g}, \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^0 \rangle &= \left\langle \mathbf{g}, \frac{\hat{\mathbf{Z}}\mathbf{b}}{\sqrt{m}} \right\rangle + \left\langle \mathbf{g}, \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^0 - \frac{\hat{\mathbf{Z}}\mathbf{b}}{\sqrt{m}} \right\rangle \\ &\leq \frac{\mathbf{g}^T \hat{\mathbf{Z}}\mathbf{b}}{\sqrt{m}} + \frac{1}{\sqrt{\eta}} \left\| \frac{\mathbf{g}}{\sqrt{m}} \right\|_{\hat{\mathbf{Z}}^{-1}} \left\| \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^0 - \frac{\hat{\mathbf{Z}}\mathbf{b}}{\sqrt{m}} \right\|_2 \\ &\leq \frac{\mathbf{g}^T \hat{\mathbf{Z}}\mathbf{b}}{\sqrt{m}} + \left\| \frac{\mathbf{g}}{\sqrt{m}} \right\|_{\hat{\mathbf{Z}}^{-1}} \sqrt{\frac{2TB}{m\eta\sigma^2}} (1 - \eta m \sigma^2)^{J/2} \end{aligned} \quad (\text{D.6})$$

The second inequality follows from the convergence of GD to \mathcal{L}_2 , given in Lemma D.8. By the definition of posterior mean and variance we have,

$$\begin{aligned} \hat{\mu}^{(s)}(\mathbf{x}) &= \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}^0)^T \hat{\mathbf{Z}}\mathbf{b}}{\sqrt{m}} \\ \hat{\sigma}^{(s)}(\mathbf{x}) &= \left\| \frac{\mathbf{g}(\mathbf{x}; \boldsymbol{\theta}^0)}{\sqrt{m}} \right\|_{\hat{\mathbf{Z}}^{-1}}. \end{aligned}$$

The upper bound on $f(\mathbf{x}; \boldsymbol{\theta}^J) - \hat{\mu}^{(s)}(\mathbf{x})$ follows from plugging in Equation D.6 into Equation D.4, and applying Lemma D.7. Similarly, for the lower bound we have,

$$-f(\mathbf{x}^i; \boldsymbol{\theta}^J) \leq \langle \mathbf{g}, \boldsymbol{\theta}^0 - \boldsymbol{\theta}^J \rangle + C \left(\frac{TB}{m\sigma^2} \right)^{2/3} L^3 \sqrt{m \log m} \quad (\text{D.7})$$

$$\langle \mathbf{g}, \boldsymbol{\theta}^0 - \tilde{\boldsymbol{\theta}}^J \rangle \leq -\hat{\mu}^{(s)}(\mathbf{x}) + \hat{\sigma}^{(s)}(\mathbf{x}) \sqrt{\frac{2TB}{m\eta\sigma^2}} (1 - \eta m \sigma^2)^{J/2} \quad (\text{D.8})$$

$$\langle \mathbf{g}, \tilde{\boldsymbol{\theta}}^J - \boldsymbol{\theta}^J \rangle \leq 3\hat{\sigma}^{(s)}(\mathbf{x}) \sqrt{\frac{TB}{m\eta\sigma^2}} \quad (\text{D.9})$$

Where inequality D.7 holds by Lemma D.7, and the next two inequalities are driven similarly to equations D.4 and D.6. The lower bound results by putting together equations D.7-D.9, and this concludes the proof. \square

D.1.2 Proof of Lemma D.4

Proof. Consider Lemma C.3, and substitute the approximate feature map $\hat{\phi}(x) = \mathbf{g}(x; \boldsymbol{\theta}^0)$ for the NTK feature map $\phi(x)$. For simplicity we denote $\mathbf{g}(x; \boldsymbol{\theta}^0)$ by \mathbf{g} . m is chosen such that lemma D.2 holds. Then, by lemma C.3 applied to $\hat{\phi}$, with probability greater than $1 - \delta - 2|\mathcal{A}|e^{-\beta T/2}$,

$$\begin{aligned} |\hat{\mu}_t^{(s)}(\mathbf{x}) - f^*(\mathbf{x})| &= |\mathbf{g}^T(\mathbf{x})(\mathbf{G}^T \mathbf{G} + \sigma^2 \mathbf{I})^{-1} \mathbf{G}^T \mathbf{y} - \mathbf{g}^T \boldsymbol{\theta}^*| \\ &\leq 2B \sqrt{\beta_t} \hat{\sigma}_t^{(s)}(\mathbf{x}) - \sigma^2 \mathbf{g}^T(\mathbf{x})(\mathbf{G}^T \mathbf{G} + \sigma^2 \mathbf{I})^{-1} \boldsymbol{\theta}^* \end{aligned}$$

for any $\mathbf{x} = \mathbf{z}_t \mathbf{a}$, $\mathbf{s} \in \mathcal{A}$. By Lemma D.2, $\sqrt{m} \|\boldsymbol{\theta}^*\|_2 \leq \sqrt{2}B$, and plugging in Lemma D.3 concludes the proof. \square

D.1.3 Proof of Lemma D.5

Proof. We use some of the inequalities given in the proof of Lemma 5.4 Zhou et al. (2020). Consider an arbitrary sequence $(\mathbf{x}_t)_{t \leq T}$. For the approximate feature map $\mathbf{g}(x; \boldsymbol{\theta}^0)/\sqrt{m}$, recall the definition of information gain after observing T samples,

$$I_g = \frac{1}{2} \log \det (\mathbf{I} + \sigma^{-2} \mathbf{G}_T \mathbf{G}_T^T / m)$$

where $\mathbf{G}_T = [\mathbf{g}(\mathbf{x}_t; \boldsymbol{\theta}^0)]_{t \leq T}^T \in \mathbb{R}^{T \times p}$. Let $[\mathbf{K}_{\text{NN}}]_{i,j \leq T} = k_{\text{NN}}(\mathbf{x}_i, \mathbf{x}_j)$ with k the NTK function of the fully-connected L -layer network.

$$\begin{aligned} I_g &= \frac{1}{2} \log \det (\mathbf{I} + \sigma^{-2} \mathbf{K}_{\text{NN}} + \sigma^{-2} (\mathbf{G}_T \mathbf{G}_T^T / m - \mathbf{K}_{\text{NN}})) \\ &\stackrel{(a)}{\leq} \frac{1}{2} \log \det (\mathbf{I} + \sigma^{-2} \mathbf{K}_{\text{NN}}) + \langle (\mathbf{I} + \sigma^{-2} \mathbf{K}_{\text{NN}})^{-1}, \sigma^{-2} (\mathbf{G}_T \mathbf{G}_T^T / m - \mathbf{K}_{\text{NN}}) \rangle \\ &\leq I_k + \sigma^{-2} \|(\mathbf{I} + \sigma^{-2} \mathbf{K}_{\text{NN}})^{-1}\|_F \|\mathbf{G}_T \mathbf{G}_T^T / m - \mathbf{K}_{\text{NN}}\|_F \\ &\stackrel{(b)}{\leq} I_k + \sigma^{-2} \sqrt{T} \|\mathbf{G}_T \mathbf{G}_T^T / m - \mathbf{K}_{\text{NN}}\|_F \\ &\stackrel{(c)}{\leq} I_k + \sigma^{-2} T \sqrt{T} \epsilon \\ &\stackrel{(d)}{\leq} \gamma_T + \sigma^{-2}. \end{aligned} \tag{D.10}$$

Inequality (a) holds by concavity of $\log \det(\cdot)$. Inequality (b) holds since $\mathbf{I} \preceq \mathbf{I} + \sigma^{-2} \mathbf{K}_{\text{NN}}$. Inequality (c) holds due to Lemma D.1. Finally, inequality (d) arises from the choice of m . Equation D.10 holds for any arbitrary context set, thus it also holds for the sequence which maximizes the information gain. \square

D.1.4 Proof of Other Lemmas in Section D.1

Proof of Lemma D.7. Let m, η satisfy the first two conditions in the lemma. By Cao and Gu (2019) Lemma 4.1, if $\|\boldsymbol{\theta}^J - \boldsymbol{\theta}^0\|_2 \leq \tau$, then

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

Under the given conditions however, Lemma B.2 from Zhou et al. (2020) holds and we have,

$$\|\boldsymbol{\theta}^J - \boldsymbol{\theta}^0\|_2 \leq 2 \sqrt{\frac{TB}{m\sigma^2}}.$$

\square

Proof of lemma D.8. This proof repeats proof of lemma C.4 in Zhou et al. (2020) and is given here for completeness. \mathcal{L}_2 is $m\sigma^2$ -strongly convex, and $C(tmL + m\sigma^2)$ -smooth, since

$$\nabla^2 \mathcal{L}_2 = \mathbf{G}^T \mathbf{G} + m\sigma^2 \mathbf{I} \leq (\|\mathbf{G}\|_2^2 + m\sigma^2) \mathbf{I} \leq C(tmL + m\sigma^2) \mathbf{I},$$

where the inequality follows from lemma D.10. It is widely known that gradient descent on smooth strongly convex functions converges to optima given that the learning rate is smaller than the smoothness coefficient inversed. Moreover, the minima of \mathcal{L}_2 is unique and has the closed form

$$\tilde{\boldsymbol{\theta}}^* = \boldsymbol{\theta}^0 + \hat{\mathbf{Z}}^{-1} \mathbf{G}^T \mathbf{y} / m$$

Having set $\eta \leq C(tmL + m\sigma^2)^{-1}$, we get that $\tilde{\boldsymbol{\theta}}^j$ converges to $\tilde{\boldsymbol{\theta}}^*$ with exponential rate,

$$\begin{aligned} \left\| \tilde{\boldsymbol{\theta}}^j - \boldsymbol{\theta}^0 - \hat{\mathbf{Z}}^{-1} \mathbf{G}^T \mathbf{y} / m \right\|_2^2 &\leq (1 - \eta m \sigma^2)^j \frac{2}{m \sigma^2} (\mathcal{L}_2(\boldsymbol{\theta}^0) - \mathcal{L}_2(\tilde{\boldsymbol{\theta}}^*)) \\ &\leq \frac{2(1 - \eta m \sigma^2)^j}{m \sigma^2} \|\mathbf{y}\|_2^2 \\ &\leq \frac{2TB(1 - \eta m \sigma^2)^j}{m \sigma^2}. \end{aligned}$$

The second inequality holds due to $\mathcal{L}_2 \geq 0$. From the RKHS assumption, the true reward is bounded by B and hence the last inequality follows from $|\Psi_t^{(s)}| \leq T$. Strong Convexity of \mathcal{L}_2 , guarantees a monotonic decrease of the loss and we have,

$$\begin{aligned} m\sigma^2 \left\| \tilde{\boldsymbol{\theta}}^J - \tilde{\boldsymbol{\theta}}^0 \right\|_2^2 &\leq m\sigma^2 \left\| \tilde{\boldsymbol{\theta}}^J - \tilde{\boldsymbol{\theta}}^0 \right\|_2^2 + \left\| \mathbf{G}^T (\tilde{\boldsymbol{\theta}}^J - \tilde{\boldsymbol{\theta}}^0) - \mathbf{y} \right\|_2^2 \\ &\leq m\sigma^2 \left\| \tilde{\boldsymbol{\theta}}^0 - \tilde{\boldsymbol{\theta}}^0 \right\|_2^2 + \left\| \mathbf{G}^T (\tilde{\boldsymbol{\theta}}^0 - \tilde{\boldsymbol{\theta}}^0) - \mathbf{y} \right\|_2^2 \\ &\leq \|\mathbf{y}\|_2^2 \\ &\leq TB \end{aligned}$$

□

Proof of lemma D.9. Choose m, η such that they satisfy condition D.6. By lemma B.2 Zhou et al. (2020), $\left\| \boldsymbol{\theta}^J - \boldsymbol{\theta}^0 \right\|_2 \leq 2\sqrt{TB/(m\sigma^2)}$.

$$\begin{aligned} \left\| \boldsymbol{\theta}^J - \tilde{\boldsymbol{\theta}}^J \right\|_2 &\leq \left\| \boldsymbol{\theta}^J - \boldsymbol{\theta}^0 \right\|_2 + \left\| \tilde{\boldsymbol{\theta}}^J - \tilde{\boldsymbol{\theta}}^0 \right\|_2 \\ &\leq 3\sqrt{\frac{TB}{m\sigma^2}}. \end{aligned}$$

Where the second inequality holds by Lemma D.8.

□

Lemma D.10. Consider the fixed set $\{\mathbf{x}_i\}_{i \leq t}$ of inputs. Let $\mathbf{G} = [\mathbf{g}^T(x_i; \boldsymbol{\theta}^0)]_{i \leq t}^T$, where \mathbf{g} shows the gradients of a L -layer feedforward network of width m at initialization. There exists constants $(C_i)_{i \leq 4}$ such that if for any $\delta > 0$, m satisfies condition D.6, then, with probability greater than $1 - \delta$.

$$\|\mathbf{G}\|_F \leq C\sqrt{tmL}$$

for some constant C .

Proof of Lemma D.10. From Lemma B.3 Cao and Gu (2019), we have $\|\mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}^0)\|_2 \leq \tilde{C}\sqrt{mL}$ with high probability. By the definition of Frobenius norm, it follows,

$$\|\mathbf{G}\|_F \leq \sqrt{t} \max_{i \leq t} C \|\mathbf{g}(\mathbf{x}_i; \boldsymbol{\theta}^0)\|_2 \leq C\sqrt{tmL}$$

□

D.2 Proof of Theorem 5.4

This proof will closely follow the proof of Theorem 4.2, with small adjustments to the condition on m . We begin by giving intuition on why this is the case. In this section, m refers to the number of channels of the convolutional network. Recall that,

$$\bar{f}(\mathbf{x}; \boldsymbol{\theta}) := f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{d} \sum_{l=1}^d f_{\text{NN}}(c_l \cdot \mathbf{x}; \boldsymbol{\theta}).$$

For simplicity in notation, from now on we will refer to $f_{\text{NN}}(\mathbf{x}; \boldsymbol{\theta})$ just as $f(\mathbf{x}; \boldsymbol{\theta})$ and use a “bar” notation to refer to the convolutional counterpart of a variable, to emphasize that the 2-layer convolutional network is the average of the 2-layer fully-connected networks taken over all circular shifts of the input.

$$\bar{k}(\mathbf{x}, \mathbf{x}') = k_{\text{CNN}}(\mathbf{x}, \mathbf{x}') := \frac{1}{d} \sum_{l=1}^d k_{\text{NN}}(c_l \cdot \mathbf{x}, \mathbf{x}'),$$

Then it is straight-forward to show that

$$\bar{g}(\mathbf{x}; \boldsymbol{\theta}^0) := \nabla_{\boldsymbol{\theta}} f_{\text{CNN}}(\mathbf{x}; \boldsymbol{\theta}^0) = \frac{1}{d} \sum_{l=1}^d g(c_l \cdot \mathbf{x}; \boldsymbol{\theta}^0)$$

where the vector $\boldsymbol{\theta}^0$ is referring to the same set of parameter in both case. Starting with an identical initialization, and running gradient descent on the ℓ_2 regularized LSE loss will cause $\bar{\boldsymbol{\theta}}^{(J)} := \boldsymbol{\theta}_{\text{CNN}}^{(J)}$ and $\boldsymbol{\theta}^{(J)}$ not to be equal anymore, but we can still show that $\bar{\boldsymbol{\theta}}^{(J)}$ and $\boldsymbol{\theta}^0$ are close as it was in the fully-connected case. Similarly, we can show that \bar{g} faces a small change during training with gradient descent. We will now present the lemmas needed for proving Theorem 5.4. These are equivalent to the lemmas in Section D.1 repeated for the convolutional network. Some statements, however, are weaker with respect to the condition on m , since convergence of the gram matrix to the CNTK (Lemma D.11) is only proven in the $m \rightarrow \infty$ limit.

Lemma D.11 (Convolutional variant of Lemma D.1). *Let $\bar{\mathbf{G}} = [\bar{g}^T(\mathbf{x}_i; \boldsymbol{\theta}^0)]_{i \leq T}^T$ and $\bar{\mathbf{K}} = [\bar{k}(\mathbf{x}_i, \mathbf{x}_j)]_{i, j \leq T}$. For any $\epsilon > 0$, there exists M such that for every $m \geq M$,*

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

Lemma D.12 (Convolutional variant of Lemma D.2). *Let f^* be a member of $\mathcal{H}_{k_{\text{CNN}}}$ with bounded RKHS norm $\|f\|_{k_{\text{CNN}}} \leq B$. Then there exists M such that for every $m \geq M$, there is a $\boldsymbol{\theta}^* \in \mathbb{R}^p$ that satisfies*

$$f^*(\mathbf{x}^i) = \langle \bar{g}(\mathbf{x}^i; \boldsymbol{\theta}^0), \boldsymbol{\theta}^* \rangle, \quad \sqrt{m} \|\boldsymbol{\theta}^*\|_2 \leq \sqrt{2}B$$

for all $i \leq T|\mathcal{A}|$.

Lemma D.13 (Convolutional variant of Lemma D.3). *Fix $s \leq S$. Consider a given context set, $\{\mathbf{x}_\tau\}_{\tau \in \Psi_t^{(s)}}$. Assume construction of $\Psi_t^{(s)}$ is such that the corresponding rewards, y_τ are statistically independent. Then there exists C_1 , such that for any $\delta > 0$, if the learning rate is picked $\eta = C_1(LmT + m\sigma^2)^{-1}$, and*

$$m \geq \text{Poly}\left(T, L, |\mathcal{A}|, \sigma^{-2}, \log(1/\delta)\right).$$

Then with probability of at least $1 - \delta$, for all $i \leq T|\mathcal{A}|$,

$$|\bar{f}(\mathbf{x}^i; \bar{\boldsymbol{\theta}}^{(J)}) - \hat{\mu}^{(s)}(\mathbf{x}^i)| \leq \hat{\sigma}^{(s)}(\mathbf{x}^i) \sqrt{\frac{TB}{m\eta\sigma^2} (3 + (1 - m\eta\sigma^2)^{J/2})} + \bar{C} \left(\frac{TB}{m\sigma^2}\right)^{2/3} L^3 \sqrt{m \log m}$$

for some constant \bar{C} , where $\hat{\mu}^{(s)}$ and $\hat{\sigma}^{(s)}$ are the posterior mean and variance of the reward with $\text{GP}(0, \hat{K}_{\text{CNN}})$ prior, after observing $(\mathbf{x}_\tau, y_\tau)_{\tau \in \Psi_t^{(s)}}$.

Lemma D.14 (Convolutional variant of D.4). *Fix $s \leq S$. Consider a given context set, $\{\mathbf{x}_\tau\}_{\tau \in \Psi_t^{(s)}}$. Assume construction of $\Psi_t^{(s)}$ is such that the corresponding rewards, y_τ are statistically independent. Let $\delta > 0$ and $\eta = C_1(LmT + m\sigma^2)^{-1}$. Then, there exists M such that for all $m \geq M$, for any action $\mathbf{a} \in \mathcal{A}$, and for some constant \bar{C} with probability of at least $1 - 2|\mathcal{A}|e^{-\beta T/2} - \delta$,*

$$|\bar{f}(\mathbf{x}; \bar{\boldsymbol{\theta}}^{(J)}) - f^*(\mathbf{x})| \leq \hat{\sigma}^{(s)}(\mathbf{x}) \left(2B\sqrt{\beta T} + \sigma\sqrt{\frac{2}{m}}B + \sqrt{\frac{TB}{m\eta\sigma^2}}(3 + (1 - m\eta\sigma^2)^{J/2}) \right) + \bar{C}\left(\frac{TB}{m\sigma^2}\right)^{2/3}L^3\sqrt{m \log m}$$

where $\mathbf{x} = \mathbf{z}_t \mathbf{a}$.

Lemma D.15 (Convolutional variant of Lemma D.5). *There exists M such that for all $m \geq M$,*

$$\hat{\gamma}_T \leq \bar{\gamma}_T + \sigma^{-2}$$

Proof of Theorem 5.4. Repeating the proof of Theorem 4.2, and plugging in Lemmas D.11 through D.15 instead of Lemmas D.1-D.5 concludes the result. \square

D.2.1 Proof of Lemmas in Section D.2

Proof of Lemma D.11. From Arora et al. (2019), for any \mathbf{x}, \mathbf{x}' on the hyper-sphere, with probability one,

$$\lim_{m \rightarrow \infty} \langle \bar{\mathbf{g}}(\mathbf{x}; \boldsymbol{\theta}), \bar{\mathbf{g}}(\mathbf{x}'; \boldsymbol{\theta}) \rangle = k_{\text{CNN}}(\mathbf{x}, \mathbf{x}').$$

In other words, for every $\epsilon > 0$, there exists M such that for all $m \geq M$, with probability one,

$$|\langle \bar{\mathbf{g}}(\mathbf{x}; \boldsymbol{\theta}), \bar{\mathbf{g}}(\mathbf{x}'; \boldsymbol{\theta}) \rangle - k_{\text{CNN}}(\mathbf{x}, \mathbf{x}')| \leq \epsilon$$

Recall that $\bar{\mathbf{G}} = [\bar{\mathbf{g}}^T(\mathbf{x}_t; \boldsymbol{\theta}^0)]_{t \leq T}^T$. Let M_l denote the number of channels that satisfies the equation above for the l -th pair of $(\mathbf{x}_i, \mathbf{x}_j)$, where $i, j \leq T$. Setting

$$M = \max_{l \leq \binom{T}{2}} M_l$$

will ensure that each two elements of $\bar{\mathbf{G}}^T \bar{\mathbf{G}}/m$ and $\bar{\mathbf{K}}$ are closer than ϵ . The proof is concluded via the definition of Frobenius norm. \square

Proof of Lemma D.12. This proof closely tracks the proof of Lemma 5.1 in Zhou et al. (2020). Consider Lemma D.11 and let $\epsilon = \lambda_0/(2TK)$, where $\lambda_0 > 0$ is the smallest eigenvalue of $\bar{\mathbf{K}}$. Then with probability one, we have $\|\bar{\mathbf{G}}^T \bar{\mathbf{G}} - \bar{\mathbf{K}}\|_F \leq \lambda_0/2$. Therefore,

$$\bar{\mathbf{G}}^T \bar{\mathbf{G}} \succcurlyeq \bar{\mathbf{K}} - \|\bar{\mathbf{G}}^T \bar{\mathbf{G}} - \bar{\mathbf{K}}\|_F \mathbf{I} \succcurlyeq \bar{\mathbf{K}} - \lambda_0 \mathbf{I}/2 \succcurlyeq \bar{\mathbf{K}}/2 \succ 0 \quad (\text{D.11})$$

where the first inequality is due to the triangle inequality. This implies that $\bar{\mathbf{G}}$ is also positive definite. Suppose $\bar{\mathbf{G}} = \mathbf{P}\mathbf{A}\mathbf{Q}^T$, with $\mathbf{A} \succ 0$. Setting $\boldsymbol{\theta}^* = \mathbf{P}\mathbf{A}^{-1}\mathbf{Q}^T \mathbf{f}^*$ satisfies the equation in the lemma, where $\mathbf{f}^* = [f^*(\mathbf{x}^i)]_{i \leq T|\mathcal{A}|}$. Moreover, due to Equation D.11,

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

By the assumption on reward, $\|\mathbf{f}^*\|_{k_{\text{CNN}}} \leq B$, completing the proof. \square

Proof of Lemma D.13. It suffices to show that Lemmas D.7, D.8, and D.9 hold for the convolutional variant, and the proof follows by repeating the steps taken in proof of Lemma D.3. We start by showing that lemma D.8 applies to a CNN. Consider the convolutional variant of the auxiliary loss,

$$\bar{\mathcal{L}}_2(\boldsymbol{\theta}) = \|\bar{\mathbf{G}}^T(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^0) - \mathbf{y}\|_2^2 + m\sigma^2\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^0\|_2^2. \quad (\text{D.12})$$

Let $(\tilde{\theta}^j)$ to denote the gradient descent updates. The loss $\tilde{\mathcal{L}}_2$ is strongly convex, which allows us to repeat the proof for the convolutional equivalent of the parameters. We conclude that Lemma D.8 still holds.

By Lemma 4.1 Cao and Gu (2019), if $\|\tilde{\theta}^{(J)} - \theta^0\| \leq \tau$, then,

$$|\bar{f}(\mathbf{x}^i; \bar{\theta}^J) - \bar{f}(\mathbf{x}^i; \tilde{\theta}^J) - \langle \bar{\mathbf{g}}(\mathbf{x}^i; \theta^0), \bar{\theta}^J - \theta^0 \rangle| \leq C\tau^{4/3}L^3\sqrt{m \log m}.$$

To show that Lemma D.7 can be applied to CNN-UCB, it remains to show that $\|\tilde{\theta}^{(J)} - \theta^0\| \leq 2\sqrt{TB/(m\sigma^2)}$. Recall that $\bar{\mathbf{g}}(\mathbf{x}; \theta^0)$ is equal to $\mathbf{g}(c_l \cdot \mathbf{x}; \theta^0)$ averaged over all c_l . Therefore all inequalities that bound a norm of $\mathbf{g}(\mathbf{x}; \theta^0)$ uniformly for all $\mathbf{x} \in \mathcal{X}$, carry over to $\bar{\mathbf{g}}(\mathbf{x}; \theta^0)$. It is straightforward to show that Lemma B.2 from Zhou et al. (2020) also holds in the convolutional case, which completes the proof for convolutional variant of Lemma D.7. From the triangle inequality, and by Lemma D.8 we also get

$$\|\bar{\theta}^J - \tilde{\theta}^J\|_2 \leq 3\sqrt{\frac{TB}{m\sigma^2}}$$

which proves Lemma D.9 for the convolutional setting. \square

Proof of Lemma D.14. We may repeat the proof for Lemma D.4 and use Lemma D.12 and D.13 when needed, in place of Lemmas D.2 and D.3 respectively. \square

Proof of Lemma D.15. We may repeat the recipe of the proof to Lemma D.5. To make it applicable to the CNTK, one step has to be modified, and that is inequality (d) of Equation D.10, which still holds due to Lemma D.11. \square