

Transformers Meet In-Context Learning: A Universal Approximation Theory

Gen Li* Yuchen Jiao* Yu Huang[†] Yuting Wei[†] Yuxin Chen[†]

June 5, 2025

Abstract

Modern large language models are capable of in-context learning, the ability to perform new tasks at inference time using only a handful of input-output examples in the prompt, without any fine-tuning or parameter updates. We develop a universal approximation theory to better understand how transformers enable in-context learning. For any class of functions (each representing a distinct task), we demonstrate how to construct a transformer that, without any further weight updates, can perform reliable prediction given only a few in-context examples. In contrast to much of the recent literature that frames transformers as algorithm approximators — i.e., constructing transformers to emulate the iterations of optimization algorithms as a means to approximate solutions of learning problems — our work adopts a fundamentally different approach rooted in universal function approximation. This alternative approach offers approximation guarantees that are not constrained by the effectiveness of the optimization algorithms being approximated, thereby extending far beyond convex problems and linear function classes. Our construction sheds light on how transformers can simultaneously learn general-purpose representations and adapt dynamically to in-context examples.

Keywords: in-context learning, universal approximation, transformers

Contents

1	Introduction	2
1.1	In-context learning	2
1.2	Approximation theory for in-context learning?	2
1.3	An overview of our main contributions	3
1.4	Related work	3
1.5	Notation	4
2	Problem formulation	5
2.1	Setting: in-context learning	5
2.2	Transformer architecture	5
2.3	Key quantities	7
3	Main results: transformers as universal in-context learners	7
4	Analysis	9
5	Discussion	14

The first two authors contributed equally.

*Department of Statistics, Chinese University of Hong Kong.

[†]Department of Statistics and Data Science, the Wharton School, University of Pennsylvania.

A Proof of key lemmas	14
A.1 Proof of Lemma 1	14
A.2 Proof of Lemma 2	19
A.3 Proof of Lemma 3	23
A.4 Proof of Lemma 4	25

1 Introduction

1.1 In-context learning

The transformer architecture introduced by Vaswani et al. (2017), which leverages a multi-head attention mechanism to capture intricate dependencies between tokens in a sequence, has catalyzed remarkable breakthroughs in large language models and reshaped algorithm designs across diverse AI domains (Khan et al., 2022; Shamshad et al., 2023; Lin et al., 2022; Gillioz et al., 2020). Built upon and powered by the transformer structure, recent pre-trained foundation models (e.g., the Generative Pre-trained Transformer (GPT) series) have unlocked a host of emergent capabilities that were previously unattainable (Bommasani et al., 2021).

One striking example is the emergent capability of “in-context learning” (ICL) — a concept coined by Brown et al. (2020) with the release of GPT-3 — which has since become a cornerstone of modern foundation models (Dong et al., 2022). In a nutshell, in-context learning refers to the ability to perform new tasks at inference time, without any update of the learned model. A contemporary large language model, pretrained in a universal, task-agnostic fashion, can readily handle a new task on the fly when given just a handful of input-output demonstrations. As a concrete example, a new task might be described using some function $f(\cdot)$ (which was unknown *a priori* during pretraining), and be presented in a prompt containing N input-output examples:

$$\text{prompt: } \mathbf{x}_1 \rightarrow f(\mathbf{x}_1), \mathbf{x}_2 \rightarrow f(\mathbf{x}_2), \dots, \mathbf{x}_N \rightarrow f(\mathbf{x}_N), \mathbf{x}_{N+1} \rightarrow ?$$

with the model then asked to predict $f(\mathbf{x}_{N+1})$ for the input instance \mathbf{x}_{N+1} . Remarkably, a pretrained model with ICL capabilities can accomplish so without any fine-tuning or retraining, relying solely on the context provided in the few-shot demonstrations to make high-quality predictions.

1.2 Approximation theory for in-context learning?

The intriguing, phenomenal capability of in-context learning has sparked substantial interest from the theoretical community, motivating a flurry of recent activity to illuminate its fundamental principles and uncover new insights. Such theoretical pursuits have attempted to tackle various facets of ICL, spanning approximation capability, training dynamics, generalization performance, to name just a few (Garg et al., 2022; Von Oswald et al., 2023; von Oswald et al., 2023; Ahn et al., 2023; Bai et al., 2023). In the current paper, we contribute to this growing body of work by investigating the effectiveness of transformers as universal approximators that support in-context learning.

Prior work: transformers as algorithm approximators. Approximation theory has emerged as a powerful lens for demystifying the representation power of transformers for ICL. Towards this end, a predominant approach adopted in recent work is to interpret transformers as algorithm approximators, whereby transformers are constructed to emulate the iterative dynamics of classical optimization algorithms during training, such as gradient descent (Von Oswald et al., 2023), preconditioned gradient descent (Ahn et al., 2023), transfer learning (Hataya et al., 2024), and Newton’s method (Giannou et al., 2023; Fu et al., 2024). The underlying rationale is that: if each iteration of these optimization algorithms can be realized via a few attention and feed-forward layers, then a multi-layer transformer could, in principle, be constructed to emulate the full iterative procedure of an optimization algorithm, as a means to approximate solutions returned by these optimization-based training algorithms. Beyond emulating fixed optimization procedures, transformers can also be constructed to support in-context algorithm selection with the aid of a few more carefully chosen layers, enabling automatic selection of appropriate algorithms based on in-context demonstrations (Bai et al., 2023).

While this algorithm approximator perspective is versatile — owing to the broad applicability of optimization algorithms like gradient descent — its utility is fundamentally constrained by the convergence properties of the algorithms being approximated. Noteworthy, except for [Giannou et al. \(2023\)](#); [Hataya et al. \(2024\)](#), existing analyses from this perspective have been restricted to linear tasks. Indeed, optimization algorithms such as gradient and Newton’s methods enjoy global convergence guarantees primarily in the context of convex loss minimization problems like linear regression, which explains why prior work along this line focused predominantly on simple convex loss minimization settings or on learning linear functions. When this approach is extended to tackle more general problems, the resulting approximation guarantees must account for the optimization error inherent to these algorithms being approximated, thereby limiting the efficacy of this technical approach for addressing broader nonconvex learning problems.

Transformers as universal function approximators? In this work, we follow a fundamentally different route: rather than designing transformers to approximate optimization algorithms as an intermediate step towards approaching desirable solutions, we seek to investigate transformers’ capabilities as direct function approximators within the framework of in-context learning. While function approximation theory has been well established for neural network models (e.g., [Barron \(1993\)](#); [Hornik et al. \(1994\)](#); [Bach \(2017\)](#); [Kurková and Sanguinetti \(2002\)](#)), little work has been done on understanding the universal function approximation capabilities of transformers in the context of ICL. It was largely unclear how transformers can learn universal representation of a general class of functions while being fully adaptive to in-context examples.

1.3 An overview of our main contributions

In this paper, we make progress towards understanding the approximation capability of transformers for in-context learning, with the aim of accommodating general function classes. More concretely, consider a general class \mathcal{F} of functions mapping \mathbb{R}^d to \mathbb{R} , with each function representing a distinct task. Assuming that the gradient of each function from \mathcal{F} has a certain bounded Fourier magnitude norm, we show how to construct a universal transformer, with L layers and input dimension $O(d + n)$ for some large enough parameter n , such that: for every function $f \in \mathcal{F}$, this transformer can make reliable prediction given N in-context examples generated based on f (in a way that works universally across all tasks in \mathcal{F} .) More precisely, the mean squared prediction error after observing N in-context examples is bounded by (up to some logarithmic factor):

$$\sqrt{\frac{1}{N}} + \frac{n}{L} + \left(\frac{\log |\mathcal{N}_\varepsilon|}{n} \right)^{2/3},$$

with \mathcal{N}_ε denoting an ε -cover of the function class and input domain of interest. Clearly, the prediction error can be vanishingly small with judiciously chosen parameters of the transformer. Notably, the term “universal transformer” here refers to a model whose parameters depend only on \mathcal{F} , without prior knowledge about the specific in-context tasks to be performed.

Our universality approximation theory implies the plausibility of designing a transformer that can predict in-context any target function in \mathcal{F} at inference time, without any sort of additional retraining or fine-tuning. From the technical perspective, our analysis consists of (i) identifying a collection of universal general-purpose features to linearly represent any function from the target function class \mathcal{F} , and (ii) constructing transformer layers to perform in-context computation of the optimal linear coefficients for the task on the fly. These design ideas shed light on how transformers can learn general-purpose representations for a complicated function class (far beyond linear functions) while adapting dynamically to in-context examples.

1.4 Related work

In-context learning. The remarkable ICL capability of large language models (LLMs) has inspired an explosion of recent research towards better understanding its emergence and inner workings from multiple different perspectives. Several recent studies attempted to interpret transformer-based ICL through a Bayesian lens ([Xie et al., 2022](#); [Ahuja et al., 2023](#); [Zhang et al., 2023](#); [Hahn and Goyal, 2023](#)), while another strand of work ([Li et al., 2023](#); [Kwon et al., 2025](#); [Cole et al., 2024](#)) analyzed the generalization and stability properties of transformers in the context of ICL. Particularly relevant to the current paper is a seminal line

of recent work exploring the representation power of transformers. For instance, [Akyürek et al. \(2023\)](#); [Bai et al. \(2023\)](#); [Von Oswald et al. \(2023\)](#) demonstrated that transformers can implement gradient descent (GD) to perform linear regression in-context, whereas [Guo et al. \(2024\)](#) extended this capability to more complex scenarios involving linear functions built atop learned representations. Note, however, that empirical analysis conducted by [Shen et al. \(2023\)](#) revealed significant functional differences between ICL and standard GD in practical settings. Further bridging these perspectives, [Vladymyrov et al. \(2024\)](#) showed that linear transformers can implement complex variants of GD for linear regression, while [von Oswald et al. \(2023\)](#); [Dai et al. \(2022\)](#) unveiled connections between ICL and meta-gradient-based optimization. Additionally, [Fu et al. \(2024\)](#); [Giannou et al. \(2023\)](#) constructed transformers capable of executing higher-order algorithms such as the Newton method. Expanding this further, [Giannou et al. \(2024\)](#); [Furuya et al. \(2024\)](#); [Wang et al. \(2024\)](#) showed that transformers can perform general computational operations and learn diverse function classes in-context. More recently, [Cole et al. \(2025\)](#) investigated the representational capabilities of multi-layer linear transformers, uncovering their potential to approximate linear dynamical systems in-context. From a complementary perspective through the lens of loss landscapes, [Ahn et al. \(2023\)](#); [Mahankali et al. \(2024\)](#); [Cheng et al. \(2024\)](#) showed that transformers can implement variants of preconditioned or functional GD in-context. Another important line of research investigated the optimization dynamics underlying transformers trained to perform ICL. For instance, [Zhang et al. \(2024\)](#); [Kim and Suzuki \(2024\)](#) analyzed training dynamics for linear-attention transformers, while [Huang et al. \(2024\)](#); [Li et al. \(2024a\)](#); [Nichani et al. \(2024\)](#); [Yang et al. \(2024\)](#) studied softmax-attention transformers across a variety of ICL tasks, including linear regression ([Huang et al., 2024](#)), binary classification ([Li et al., 2024a](#)), causal structure learning ([Nichani et al., 2024](#)), representation-based learning ([Yang et al., 2024](#)), and chain-of-thought (CoT) reasoning ([Huang et al., 2025](#)). Furthermore, [Chen et al. \(2024b\)](#) explored the optimization dynamics of multi-head attention mechanisms tailored to linear regression settings.

Representation theory of transformers. Substantial theoretical efforts have been recently devoted to characterizing the representational power and capabilities of transformers and self-attention mechanisms across a variety of computational settings and statistical tasks ([Pérez et al., 2019](#); [Elhage et al., 2021](#); [Liu et al., 2022](#); [Likhoshesterov et al., 2021](#); [Wen et al., 2023](#); [Yao et al., 2021](#); [Chen and Li, 2024](#)). A prominent strand of recent research ([Sanford et al., 2023](#); [Wen et al., 2024](#); [Jelassi et al., 2024](#)) revealed notable advantages of transformers over alternative architectures such as RNNs. Despite these advances, several work ([Hahn, 2020](#); [Sanford et al., 2024](#); [Peng et al., 2024](#); [Chen et al., 2024a](#)) also identified inherent limitations of transformers, proving that transformers might fail at certain computational tasks (e.g., parity) and establishing complexity-theoretic lower bounds concerning their representational capabilities. Recently, motivated by the widespread success of the CoT techniques — which explicitly leverage intermediate reasoning steps — several work ([Li et al., 2024b](#); [Merrill and Sabharwal, 2024](#); [Feng et al., 2023](#)) began investigating the theoretical foundations and expressive power of the CoT paradigm.

1.5 Notation

Throughout this paper, bold uppercase letters represent matrices, while bold lowercase letters represent column vectors. For any vector \mathbf{v} , we use $\|\mathbf{v}\|_2$ to denote its ℓ_2 norm, and $\|\mathbf{v}\|_1$ its ℓ_1 norm. For any matrix \mathbf{A} , we denote by $[\mathbf{A}]_{i,j}$ its (i, j) -th entry. The indicator function $\mathbb{1}(\cdot)$ takes the value 1 when the condition in the parentheses is satisfied and zero otherwise. The sign function $\text{sign}(x)$ returns 1 if $x > 0$, -1 if $x < 0$, and 0 if $x = 0$. For any scalar function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, the notation $\sigma(\mathbf{x})$ for $\mathbf{x} \in \mathbb{R}^d$ denotes the elementwise application of σ to each entry of \mathbf{x} . Let $\mathcal{X} = \{N, L, n, \log |\mathcal{N}_\varepsilon|, C_{\mathcal{F}}, \sigma\}$ (and sometimes with the additional inclusion of some precision parameter $\varepsilon_{\text{pred}}$). The notation $f(\mathcal{X}) = O(g(\mathcal{X}))$ or $f(\mathcal{X}) \lesssim g(\mathcal{X})$ (resp. $f(\mathcal{X}) \gtrsim g(\mathcal{X})$) means that there exists a universal constant $C_0 > 0$ such that $f(\mathcal{X}) \leq C_0 g(\mathcal{X})$ (resp. $f(\mathcal{X}) \geq C_0 g(\mathcal{X})$) for any choice of \mathcal{X} . The notation $f(\mathcal{X}) \asymp g(\mathcal{X})$ means $f(\mathcal{X}) \lesssim g(\mathcal{X})$ and $f(\mathcal{X}) \gtrsim g(\mathcal{X})$ hold simultaneously. We define $\tilde{O}(\cdot)$ in the same way as $O(\cdot)$ except that it hides logarithmic factors. We also use $\mathbf{0}$ to denote the all-zero vector. For any positive integer m , we denote $[m] := \{1, \dots, m\}$.

2 Problem formulation

2.1 Setting: in-context learning

To set the stage, we formulate an in-context learning setting that comprises the following components:

- *Function class.* We denote by \mathcal{F} a class of real-valued functions, mapping from \mathbb{R}^d to \mathbb{R} , that we aim to learn. Each function $f \in \mathcal{F}$ represents a distinct prediction task (i.e., given an input \mathbf{x} , predict the output $f(\mathbf{x})$).
- *Input sequence.* The input sequence, typically provided in the prompt, is composed of N input-output pairs — namely, N in-context examples — along with a new input vector for prediction. To be precise, a prompt takes the form of

$$(\mathbf{x}_1, y_1, \mathbf{x}_2, y_2, \dots, \mathbf{x}_N, y_N, \mathbf{x}_{N+1}), \quad (1a)$$

where for every i ,

$$\mathbf{x}_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}_{\mathcal{X}}, \quad z_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}_{\mathcal{Z}}, \quad y_i = f(\mathbf{x}_i) + z_i \text{ for some function } f \in \mathcal{F}. \quad (1b)$$

Here, $\{\mathbf{x}_i\} \subset \mathbb{R}^d$ (resp. $\{z_i\} \subset \mathbb{R}$) are input vectors (resp. noise) sampled randomly from the distribution $\mathcal{D}_{\mathcal{X}}$ (resp. $\mathcal{D}_{\mathcal{Z}}$), and the corresponding output vectors are produced by some function $f \in \mathcal{F}$ not revealed to the learner. Throughout this paper, the noise $\{z_i\}$ is assumed to be independent zero-mean sub-Gaussian with the sub-Gaussian norm upper bounded by σ (Vershynin, 2018), that is,

$$\mathbb{E}[z_i] = 0 \quad \text{and} \quad \mathbb{E}[e^{tz_i}] \leq \exp\left(\frac{\sigma^2 t^2}{2}\right) \text{ for every } t \in \mathbb{R}. \quad (2)$$

For simplicity, we assume throughout that all input vectors lie within a unit Euclidean ball:

$$\mathbf{x} \in \mathcal{B} := \{\mathbf{u} \mid \|\mathbf{u}\|_2 \leq 1\} \quad \text{for any input vector } \mathbf{x}, \quad (3)$$

but this assumption can be easily relaxed and generalized.

It is worth emphasizing that $\{(\mathbf{x}_i, y_i)\}$ should be regarded *not* as training examples, but rather as in-context demonstrations, as they are typically provided within the prompt at inference time instead of being used during the training phase.

Goal. The aim is to design a transformer that, given an input sequence as in (1) produced by any function $f \in \mathcal{F}$, outputs a prediction \hat{y}_{N+1} obeying

$$\hat{y}_{N+1} \approx f(\mathbf{x}_{N+1})$$

in some average sense. Particular emphasis is placed on *universal design*, where the objective is to find a single transformer (to be described next) that performs well simultaneously for all $f \in \mathcal{F}$, without knowing which f to tackle in advance. This universal design requirement aligns closely with the concept of in-context learning, as the goal is for the pre-trained transformer to make reliable predictions based on input-output demonstrations at inference time, without performing any prompt-specific parameter updates.

2.2 Transformer architecture

Next, let us present a precise description of the transformer architecture to be used for in-context learning. Throughout this paper, we would like to use a matrix

$$\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_{N+1}] \in \mathbb{R}^{D \times (N+1)} \quad (4)$$

to encode the input sequence (1a) comprising N input-output examples along with an additional new input. Here, the input dimension D is typically chosen to be larger than d to allow for incorporation of several useful

auxiliary features. For instance, in our construction (to be detailed momentarily), \mathbf{H} takes the following form:

$$\mathbf{H} = \begin{bmatrix} \mathbf{x}_1 & \cdots & \mathbf{x}_N & \mathbf{x}_{N+1} \\ 1 & \cdots & 1 & 1 \\ y_1 & \cdots & y_N & 0 \\ \vdots & \text{auxiliary} & \text{info} & \vdots \\ \hat{y}_1 & \cdots & \hat{y}_N & \hat{y}_{N+1} \end{bmatrix}, \quad (5)$$

where each column entails the original input-output pair (\mathbf{x}_i, y_i) (except the last column where y_{N+1} is replaced with 0), a constant 1, a few dimension containing auxiliary information, and the prediction \hat{y}_i .

Basic building blocks. To begin with, we single out two basic building blocks.

- *Attention layer.* For any input matrix \mathbf{H} , the (self)-attention operator is defined as

$$\text{attn}(\mathbf{H}; \mathbf{Q}, \mathbf{K}, \mathbf{V}) := \frac{1}{N} \mathbf{V} \mathbf{H} \sigma_{\text{attn}}((\mathbf{Q} \mathbf{H})^\top \mathbf{K} \mathbf{H}), \quad (6)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{D \times D}$ represent the parameter matrices, commonly referred to as the query, key, and value matrices, respectively, and the activation function $\sigma_{\text{attn}}(\cdot)$ is applied either columnwise or entrywise to the input. A multi-head (self)-attention layer, which we denote by $\text{Attn}_{\Theta}(\cdot)$ as parameterized by $\Theta = \{\mathbf{Q}_m, \mathbf{K}_m, \mathbf{V}_m\}_{1 \leq m \leq M} \subset \mathbb{R}^{D \times D}$, computes a superposition of the input and the outputs from M attention operators (or attention heads). Namely, given the input matrix \mathbf{H} , the output of the attention layer with M attention heads is defined as

$$\text{Attn}_{\Theta}(\mathbf{H}) := \mathbf{H} + \sum_{m=1}^M \text{attn}(\mathbf{H}; \mathbf{Q}_m, \mathbf{K}_m, \mathbf{V}_m). \quad (7)$$

This attention mechanism plays a pivotal role in the transformer architecture (Vaswani et al., 2017), allowing one to dynamically attend to different parts of the input data.

- *Feed-forward layer (or multilayer perceptron (MLP) layer).* Given an input matrix \mathbf{H} , the feed-forward layer produces an output as follows:

$$\text{FF}_{\Theta}(\mathbf{H}) := \mathbf{H} + \mathbf{U} \sigma_{\text{ff}}(\mathbf{W} \mathbf{H}), \quad (8)$$

where $\Theta = \{\mathbf{U}, \mathbf{W}\} \subset \mathbb{R}^{D \times D}$ bundles the parameter matrices \mathbf{U} and \mathbf{W} together, and the activation function $\sigma_{\text{ff}}(\cdot)$ is applied entrywise to the input.

Throughout this paper, the two activation functions described above are chosen to be the sigmoid function and the ReLU function:

$$\sigma_{\text{attn}}(x) = \frac{e^x}{e^x + 1}, \quad \sigma_{\text{ff}}(x) = x \mathbb{1}(x > 0), \quad (9)$$

each of which is applied entrywise to its respective input.

Multi-layer transformers. With the aforementioned building blocks in place, we can readily introduce the multi-layer transformer architecture. Given an input $\mathbf{H}^{(0)} = \mathbf{H} \in \mathbb{R}^{D \times (N+1)}$, a transformer comprising L attention layers — each coupled with a feed-forward layer — carries out the following computation:

$$\mathbf{H}^{(l)} = \text{FF}_{\Theta_{\text{ff}}^{(l)}} \left(\text{Attn}_{\Theta_{\text{attn}}^{(l)}} (\mathbf{H}^{(l-1)}) \right), \quad l = 1, \dots, L, \quad (10a)$$

with the final output given by

$$\text{TF}_{\Theta}(\mathbf{H}) := \mathbf{H}^{(L)}. \quad (10b)$$

Here, Θ encapsulates all parameter matrices:

$$\Theta = \{\Theta_{\text{attn}}^{(l)}, \Theta_{\text{ff}}^{(l)}\}_{1 \leq l \leq L} \quad \text{with} \quad \Theta_{\text{attn}}^{(l)} = \{Q_m^{(l)}, K_m^{(l)}, V_m^{(l)}\}_{1 \leq m \leq M}, \quad \Theta_{\text{ff}}^{(l)} = \{U^{(l)}, W^{(l)}\} \subset \mathbb{R}^{D \times D}.$$

In particular, the transformer's prediction for the $(N+1)$ -th input can be read out from the very last entry of $\text{TF}_{\Theta}(\mathbf{H})$, i.e.,

$$\hat{y}_{N+1} = \text{ReadOut}(\text{TF}_{\Theta}(\mathbf{H})) := [\mathbf{H}^{(L)}]_{D, N+1}. \quad (11)$$

2.3 Key quantities

Before embarking on our main theory, let us take a moment to isolate a couple of key quantities that play a crucial role in our theoretical development.

For any absolutely integrable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we denote by F_f its Fourier transform, which allows one to express

$$F_f(\boldsymbol{\omega}) = \frac{1}{2\pi} \int_{\mathbf{x}} e^{-j\boldsymbol{\omega}^\top \mathbf{x}} f(\mathbf{x}) d\mathbf{x} \quad (12a)$$

$$f(\mathbf{x}) = \int_{\boldsymbol{\omega}} e^{j\boldsymbol{\omega}^\top \mathbf{x}} F_f(\boldsymbol{\omega}) d\boldsymbol{\omega} \quad (12b)$$

with $j = \sqrt{-1}$ the imaginary unit. It is also helpful to define, for each $\boldsymbol{\omega}$, the maximum magnitude of the Fourier transform over the function class \mathcal{F} as:

$$F^{\text{sup}}(\boldsymbol{\omega}) := \sup_{f \in \mathcal{F}} |F_f(\boldsymbol{\omega})|. \quad (13)$$

Inspired by the seminal work [Barron \(1993\)](#), we introduce the following key quantity:

$$C_{\mathcal{F}} := \sup_{f \in \mathcal{F}} |f(\mathbf{0})| + \int_{\boldsymbol{\omega}} \|\boldsymbol{\omega}\|_2 F^{\text{sup}}(\boldsymbol{\omega}) d\boldsymbol{\omega} < \infty. \quad (14)$$

Informally, this quantity bounds the first moment of the Fourier magnitude distribution over this function class \mathcal{F} . Compared with the quantity $C_f := \int_{\boldsymbol{\omega}} \|\boldsymbol{\omega}\|_2 |F_f(\boldsymbol{\omega})| d\boldsymbol{\omega}$ introduced in [Barron \(1993\)](#) for each function f , the main difference lies in the fact that $C_{\mathcal{F}}$ involves taking the supremum over the entire function class \mathcal{F} . Additionally, recognizing that $j\boldsymbol{\omega} F_f(\boldsymbol{\omega})$ is precisely the Fourier transform of $\nabla f(\mathbf{x})$, one can alternatively express $C_{\mathcal{F}}$ as

$$C_{\mathcal{F}} = \sup_{f \in \mathcal{F}} |f(\mathbf{0})| + \int_{\boldsymbol{\omega}} \sup_{f \in \mathcal{F}} \|F_{\nabla f}(\boldsymbol{\omega})\|_2 d\boldsymbol{\omega}, \quad (15)$$

which tracks the ℓ_1 norm of the maximum Fourier magnitude of the function gradient.

Recall that we have restricted our input space to be within the unit Euclidean ball $\mathcal{B} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \leq 1\}$. A set, denoted by $\mathcal{N}_{\varepsilon}$, is said to be an ε -cover of $\mathcal{F} \times \mathcal{B}$ if, for every $(f, \mathbf{x}) \in \mathcal{F} \times \mathcal{B}$, there exists some $(\hat{f}, \hat{\mathbf{x}}) \in \mathcal{N}_{\varepsilon}$ such that

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq \varepsilon \quad \text{and} \quad |f(\mathbf{x}) - f(\mathbf{0}) - \hat{f}(\hat{\mathbf{x}}) + \hat{f}(\mathbf{0})| \leq C_{\mathcal{F}} \varepsilon. \quad (16)$$

3 Main results: transformers as universal in-context learners

Equipped with the preliminaries and key quantities in Section 2, we are now ready to present our main theoretical findings, as summarized in the theorem below.

Theorem 1. *Let $\mathcal{N}_{\varepsilon}$ be an ε -cover of $\mathcal{F} \times \mathcal{B}$ (see (16)). Then one can construct a transformer such that:*

- i) *it has L layers, $M = O(1)$ attention heads per layer, and input dimension $D = d + 2n + 7$ for some $n \gtrsim \log |\mathcal{N}_{\varepsilon}|$;*

ii) for every $f \in \mathcal{F}$ and any input sequence $\{(\mathbf{x}_i, y_i)\}_{1 \leq i \leq N} \cup \{\mathbf{x}_{N+1}\}$ generated according to (1), with probability at least $1 - O(N^{-10})$, this transformer’s prediction \hat{y}_{N+1} (cf. (11)) satisfies

$$\mathbb{E}[(\hat{y}_{N+1} - f(\mathbf{x}_{N+1}))^2] \lesssim \left(\sqrt{\frac{\log N}{N}} + \frac{n}{L} \right) C_{\mathcal{F}}(C_{\mathcal{F}} + \sigma) + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{\frac{2}{3}}, \quad (17)$$

Here, the precision of the ε -cover is taken to satisfy $\varepsilon \lesssim \sqrt{\frac{\log N}{N}} + \frac{n}{L}$; the expectation in (17) is taken over the randomness of \mathbf{x}_{N+1} ; and the probability that the event (17) occurs is governed by the randomness in the input sequence $\{(\mathbf{x}_i, y_i)\}_{1 \leq i \leq N}$.

It is worth taking a moment to reflect on the interpretation and implications of this theorem.

Universal in-context prediction. Theorem 1 establishes the existence of a pretrained, multi-layer multi-head transformer — configured with judiciously chosen parameters, depth, number of heads, etc. — that can make reliable predictions based on in-context demonstrations. This in-context prediction capability is universal, in the sense that a single transformer can simultaneously handle all functions f in the function class of interest, without requiring any additional training or prompt-specific parameter updates.

Parameter choices. According to Theorem 1, the in-context prediction error depends on the number of layers L , the number of input-output examples N , the transformer’s input dimension D , and the intrinsic data dimension d . Specifically, when both the Fourier quantity $C_{\mathcal{F}}$ and the noise level σ are no larger than $O(1)$, the mean squared prediction error can’t exceed the order $\tilde{O}(1/\sqrt{N} + 1/L + (\log |\mathcal{N}_{\varepsilon}|/n)^{2/3})$, where $2n \approx D - d$. For instance, in the case of the linear function class $\mathcal{F} = \{f : f(\mathbf{x}) = \mathbf{a}^\top \mathbf{x} + b, \|\mathbf{a}\|_2 \leq 1, |b| \leq 1\}$, the Fourier quantity satisfies $C_{\mathcal{F}} = O(1)$ (Barron, 1993). This result implies that properly increasing the model complexity — through the use of deeper architectures (i.e., the ones with larger L) and higher input dimension D — and utilizing more in-context examples (i.e., larger N) — could enhance the in-context learning capability of the transformer, leading to improved prediction accuracy. More concretely, to achieve an $\varepsilon_{\text{pred}}$ -accurate prediction (with $\varepsilon_{\text{pred}}$ the target mean squared prediction error), it suffices to employ a transformer with parameters satisfying (up to log factors)

$$D - d \asymp C_{\mathcal{F}}^3 \varepsilon_{\text{pred}}^{-3/2} \log |\mathcal{N}_{\varepsilon}|, \quad (18a)$$

$$N \gtrsim C_{\mathcal{F}}^2 (C_{\mathcal{F}} + \sigma)^2 \varepsilon_{\text{pred}}^{-2}, \quad (18b)$$

$$L \gtrsim (D - d) C_{\mathcal{F}} (C_{\mathcal{F}} + \sigma) \varepsilon_{\text{pred}}^{-1} \asymp C_{\mathcal{F}}^4 (C_{\mathcal{F}} + \sigma) \varepsilon_{\text{pred}}^{-5/2} \log |\mathcal{N}_{\varepsilon}|. \quad (18c)$$

Logarithmic scaling on the covering number of the function class. Our approximation theory allows the function class of interest to be fairly general. Notably, both the input dimension (including that of auxiliary features) and the depth of the transformer we construct only need to scale logarithmically with the covering number of the target function class (see (18)). In other words, in order to achieve sufficient representation power for ICL, the model complexity needs to grow with the complexity of the target function class — but a logarithmic scaling with the covering number of \mathcal{F} suffices.

Function approximators vs. algorithm approximators. Theorem 1 unveils that transformers can serve as universal function approximators for in-context learning. This perspective contrasts sharply with a substantial body of recent work — e.g., Von Oswald et al. (2023); von Oswald et al. (2023); Bai et al. (2023); Ahn et al. (2023); Giannou et al. (2023); Xie et al. (2022); Cheng et al. (2024) — which has primarily focused on interpreting transformers as algorithm approximators. As alluded to previously, the approximation theory derived from the algorithm approximation perspective is often constrained by the effectiveness of the specific algorithms being approximated. For instance, algorithms like gradient descent and Newton’s method are typically not guaranteed to perform well outside the realm of convex optimization. This limitation partly explains why much of the prior literature has concentrated on relatively simple convex problems, such as linear regression. By contrast, the universal function approximation framework we deliver is not tied to the performance of such (mesa)-optimization algorithms, and as a result, can often deliver direct approximation guarantees for much broader in-context learning problems.

4 Analysis

In this section, we present the key steps for establishing Theorem 1. Informally, our proof comprises the following key ingredients:

- Identify a collection of general-purpose features such that every function (or task) in \mathcal{F} can be (approximately) represented as a combination of these features.
- For each function $f \in \mathcal{F}$, the corresponding linear coefficients can be found by means of a Lasso estimator, which is efficiently solvable via the proximal gradient method.
- A transformer can then be designed to approximate the above proximal gradient iterations.

Throughout the proof, we let $\phi(x)$ denote the following sigmoid function:

$$\phi(z) = \left(z + \frac{1}{2}\right) \mathbf{1} \left\{z + \frac{1}{2} > 0\right\} - \left(z - \frac{1}{2}\right) \mathbf{1} \left\{z - \frac{1}{2} > 0\right\}, \quad (19)$$

which satisfies $\lim_{z \rightarrow -\infty} \phi(z) = 0$, $\lim_{z \rightarrow \infty} \phi(z) = 1$, and $\phi(0) = 1/2$.

Step 1: constructing universal features for the target function class. In this step, we construct a finite collection of features to approximately represent $f(\mathbf{x}) - f(\mathbf{0})$, with the aid of the sigmoid function defined in (19). Our construction is formally presented in the following lemma (and its analysis), whose proof can be found in Appendix A.1.

Lemma 1. *Consider any $\tau > 4$ and any $n \geq c_0 \log |\mathcal{N}_\varepsilon|$ for some large enough constant $c_0 > 0$. There exist a collection of functions $\phi_i^{\text{feature}} : \mathbb{R}^d \rightarrow \mathbb{R}$ ($1 \leq i \leq n$) such that: for every $f \in \mathcal{F}$ and $\mathbf{x} \in \mathcal{B}$, one has*

$$\left| f(\mathbf{x}) - f(\mathbf{0}) - \frac{1}{n} \sum_{i=1}^n \rho_{f,i}^* \phi_i^{\text{feature}}(\mathbf{x}) \right| \lesssim C_{\mathcal{F}} \left(\frac{1}{\tau} + \tau\varepsilon + \left(\frac{\log |\mathcal{N}_\varepsilon|}{n} \right)^{\frac{1}{3}} \right) \quad (20)$$

for some f -dependent coefficients $\{\rho_{f,i}^*\}_{1 \leq i \leq n} \subset \mathbb{R}$ obeying

$$|f(\mathbf{0})| + \frac{1}{n} \sum_{i=1}^n |\rho_{f,i}^*| < 4C_{\mathcal{F}}. \quad (21)$$

Here, the functions $\{\phi_i^{\text{feature}}(\cdot)\}_{1 \leq i \leq n}$ are given by

$$\phi_i^{\text{feature}}(\mathbf{x}) = \phi \left(\tau \left(\frac{1}{\|\boldsymbol{\omega}_i\|_2} \boldsymbol{\omega}_i^\top \mathbf{x} - t_i \right) \right) \quad (22)$$

for some $\{(t_i, \boldsymbol{\omega}_i)\}_{1 \leq i \leq n} \subset \mathbb{R} \times \mathbb{R}^d$ independent of any specific f , where $\phi(\cdot)$ is defined in (19).

In words, for any function $f \in \mathcal{F}$ and any $\mathbf{x} \in \mathcal{B}$, the quantity $f(\mathbf{x}) - f(\mathbf{0})$ can be closely approximated by a linear combination of the features $\{\phi_i^{\text{feature}}(\mathbf{x})\}$, where the ℓ_1 norm of the linear coefficients is well-controlled. This reveals that $\{\phi_i^{\text{feature}}(\cdot)\}$ can serve as general-purpose features capable of linearly representing arbitrary functions in the function class \mathcal{F} . We remark here that the f -dependent coefficients $\{\rho_{f,i}^*\}_{1 \leq i \leq n}$ are not required to be positive. In the rest of the proof, we shall take

$$\tau = 1/\sqrt{\varepsilon} \quad \text{and} \quad \varepsilon_{\text{dis}} := c_{\text{dis}} C_{\mathcal{F}} \left(\sqrt{\varepsilon} + \left(\frac{\log |\mathcal{N}_\varepsilon|}{n} \right)^{\frac{1}{3}} \right) \quad (23)$$

for some large enough constant $c_{\text{dis}} > 0$, which allow one to obtain (see Lemma 1)

$$\left| f(\mathbf{x}) - f(\mathbf{0}) - \frac{1}{n} \sum_{i=1}^n \rho_{f,i}^* \phi_i^{\text{feature}}(\mathbf{x}) \right| \leq \varepsilon_{\text{dis}} \quad \text{for every } f \in \mathcal{F} \text{ and } \mathbf{x} \in \mathcal{B}. \quad (24)$$

Step 2: learning linear coefficients in-context via Lasso. Armed with the general-purpose features $\{\phi_i^{\text{feature}}(\cdot)\}$, we now proceed to show how the linear coefficients $\{\rho_{f,i}^*\}$ in (20) can be approximately located in-context.

Recall from Lemma 1 that $\{\rho_{f,i}^*\}$ satisfy some ℓ_1 -norm constraint (cf. (21)). With this in mind, we attempt estimating $\{\rho_{f,i}^*\}$ from the in-context demonstration $\{(\mathbf{x}_i, y_i)\}_{1 \leq i \leq N}$ by means of the following regularized problem (a.k.a. the Lasso estimator):

$$\underset{\boldsymbol{\rho} \in \mathbb{R}^{n+1}}{\text{minimize}} \quad \ell(\boldsymbol{\rho}) := \frac{1}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho})^2 + \lambda \|\boldsymbol{\rho}\|_1, \quad (25)$$

where λ denotes the regularized parameter to be suitably chosen, and $\boldsymbol{\phi}_i \in \mathbb{R}^{n+1}$ is defined as

$$\boldsymbol{\phi}_i = [\phi_1^{\text{feature}}(\mathbf{x}_i), \phi_2^{\text{feature}}(\mathbf{x}_i), \dots, \phi_n^{\text{feature}}(\mathbf{x}_i), 1]^\top. \quad (26)$$

In general, it is difficult to obtain an exact solution of (25), which motivates us to analyze approximate solutions instead. More specifically, we would like to analyze the prediction error of any $\hat{\boldsymbol{\rho}}$ obeying

$$\ell(\hat{\boldsymbol{\rho}}) - \ell(\boldsymbol{\rho}^*) \leq \varepsilon_{\text{opt}} \quad (27)$$

for some accuracy level ε_{opt} , where $\boldsymbol{\rho}^* \in \mathbb{R}^{n+1}$ collects the f -dependent coefficients $\rho_{f,i}^*$ in Lemma 1 in the following way:

$$\boldsymbol{\rho}^* = \left[\frac{\rho_{f,1}^*}{n}, \dots, \frac{\rho_{f,n}^*}{n}, \rho_{f,0}^* \right]^\top, \quad \text{where} \quad \rho_{f,0}^* = f(\mathbf{0}). \quad (28)$$

Note that in (27) we are comparing $\ell(\hat{\boldsymbol{\rho}})$ with $\ell(\boldsymbol{\rho}^*)$ rather than that of the minimizer of (25), as it facilitates our analysis. The following lemma quantifies the prediction error and the ℓ_1 norm of any $\hat{\boldsymbol{\rho}}$ obeying (27), whose proof is postponed to Appendix A.2.

Lemma 2. Consider any given $\lambda \geq c_\lambda \left(\sqrt{\frac{\log N}{N}} \sigma + C_{\mathcal{F}}^{-1} \varepsilon_{\text{dis}}^2 \right)$ for some sufficiently large constant $c_\lambda > 0$, where ε_{dis} is defined in (23). For any $\hat{\boldsymbol{\rho}}$ that is statistically independent from \mathbf{x}_{N+1} and satisfies (27) with $\varepsilon_{\text{opt}} \geq 0$, we have

$$\mathbb{E} \left[(\boldsymbol{\phi}_{N+1}^\top \hat{\boldsymbol{\rho}} - f(\mathbf{x}_{N+1}))^2 \right] \lesssim \sqrt{\frac{\log N}{N}} (C_{\mathcal{F}}^2 + \lambda^{-2} \varepsilon_{\text{opt}}^2 + \sigma \varepsilon_{\text{dis}}) + \varepsilon_{\text{dis}}^2 + \lambda C_{\mathcal{F}} + \varepsilon_{\text{opt}} \quad (29)$$

$$\|\hat{\boldsymbol{\rho}}\|_1 \lesssim C_{\mathcal{F}} + \lambda^{-1} \varepsilon_{\text{opt}} \quad (30)$$

with probability at least $1 - O(N^{-10})$. Here, the expectation in (29) is taken over the randomness of \mathbf{x}_{N+1} .

Remark 1. In the statement of Lemma 2, $\hat{\boldsymbol{\rho}}$ is allowed to be statistically dependent on $\{(\mathbf{x}_i, y_i)\}_{1 \leq i \leq N}$ but not on \mathbf{x}_{N+1} .

Step 3: solving the Lasso (25) via the inexact proximal gradient method. In light of Lemma 2, it is desirable to make the optimization error ε_{opt} as small as possible. Here, we propose to run the (inexact) proximal gradient method in an attempt to solve (25). More precisely, starting from the initialization $\boldsymbol{\rho}_0^{\text{proximal}} = \mathbf{0}$, the update rule for each iteration $t = 0, \dots, T$ is given by

$$\boldsymbol{\rho}_{t+1}^{\text{proximal}} = \text{prox}_{\eta \lambda \|\cdot\|_1} \left(\boldsymbol{\rho}_t^{\text{proximal}} + \frac{2\eta}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}_t^{\text{proximal}}) \boldsymbol{\phi}_i \right) + \mathbf{e}_{t+1} \quad (31a)$$

$$= \text{ST}_{\eta \lambda} \left(\boldsymbol{\rho}_t^{\text{proximal}} + \frac{2\eta}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}_t^{\text{proximal}}) \boldsymbol{\phi}_i \right) + \mathbf{e}_{t+1}, \quad (31b)$$

where $\eta > 0$ stands for the stepsize, and we have included an additive term \mathbf{e}_{t+1} that allows for inexact updates. Here, the proximal operator $\text{prox}_{\eta\lambda\|\cdot\|_1}(\cdot)$ and the soft thresholding operator $\text{ST}_{\eta\lambda}(\cdot)$ are given respectively by

$$\text{prox}_{\eta\lambda\|\cdot\|_1}^{\text{proximal}}(\mathbf{x}) := \arg \min_{\boldsymbol{\rho}} \left\{ \frac{1}{2} \|\mathbf{x} - \boldsymbol{\rho}\|_2^2 + \eta\lambda \|\boldsymbol{\rho}\|_1 \right\}, \quad (32a)$$

$$\text{ST}_{\eta\lambda}(z) := \text{sign}(z) \max\{|z| - \eta\lambda, 0\}. \quad (32b)$$

Note that $\text{ST}_{\eta\lambda}(\cdot)$ is applied entrywise in (31b).

We now develop convergence guarantees for the above proximal gradient method. It can be shown that: after $T = (L - 1)/2$ iterations, the (inexact) proximal gradient method (31b) produces an iterate $\boldsymbol{\rho}_T^{\text{proximal}}$ enjoying the following performance guarantees; the proof is postponed to Appendix A.3.

Lemma 3. *Take $T = (L - 1)/2$. Assume that*

$$\boldsymbol{\rho}_0^{\text{proximal}} = \mathbf{0}, \quad \|\boldsymbol{\rho}^*\|_1 \lesssim C_{\mathcal{F}}, \quad \lambda \gtrsim \sqrt{\frac{\log N}{N}}(C_{\mathcal{F}} + \sigma), \quad \eta = \frac{1}{2n}, \quad \text{and} \quad \|\mathbf{e}_t\|_1 \leq \varepsilon_{\text{approx}} \lesssim C_{\mathcal{F}} \text{ for all } t \leq T. \quad (33)$$

Then with probability at least $1 - O(N^{-10})$, the output of the algorithm (31b) at the T -th iteration satisfies

$$\ell(\boldsymbol{\rho}_T^{\text{proximal}}) \leq \ell(\boldsymbol{\rho}^*) + \frac{c_1 n C_{\mathcal{F}}^2}{L} + c_1(L + n)\varepsilon_{\text{approx}} \left(C_{\mathcal{F}} + \sigma + \max_{1 \leq k \leq T} \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 + \lambda \right) \quad (34)$$

for some universal constant $c_1 > 0$, and for every $t \leq T$ we have

$$\|\boldsymbol{\rho}_t^{\text{proximal}}\|_1 \lesssim C_{\mathcal{F}} + \frac{n C_{\mathcal{F}}^2}{t\lambda} + \sqrt{N}(t + n) \max_{1 \leq k \leq t} \|\mathbf{e}_k\|_1 + \frac{t + n}{\lambda} \max_{1 \leq k \leq t} \{\|\mathbf{e}_k\|_1 \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1\}. \quad (35)$$

Remark 2. The careful reader might remark that the upper bounds in Lemma 3 appear somewhat intricate, as they depend on the ℓ_1 norm of the previous iterates. Fortunately, once $\{\|\mathbf{e}_k\|_1\}$ are determined, we can apply mathematical induction to derive more concise bounds — an approach to be carried out in Step 4.

Step 4: constructing the transformer to emulate proximal gradient iterations. To build a transformer with favorable in-context learning capabilities, our design seeks to approximate the above proximal gradient iterations, which we elucidate in this step.

Let us begin by describing the input structure for each layer of our constructed transformer. For the l -th hidden layer ($0 \leq l \leq L$), the input matrix $\mathbf{H}^{(l)}$ (see (10)) takes the following form:

$$\mathbf{H}^{(l)} = \begin{pmatrix} \mathbf{x}_1^{(l)} & \mathbf{x}_2^{(l)} & \dots & \mathbf{x}_N^{(l)} & \mathbf{x}_{N+1}^{(l)} \\ y_1^{(l)} & y_2^{(l)} & \dots & y_N^{(l)} & 0 \\ w_1^{(l)} & w_2^{(l)} & \dots & w_N^{(l)} & w_{N+1}^{(l)} \\ \boldsymbol{\phi}_1^{(l)} & \boldsymbol{\phi}_2^{(l)} & \dots & \boldsymbol{\phi}_N^{(l)} & \boldsymbol{\phi}_{N+1}^{(l)} \\ \boldsymbol{\rho}^{(l)} & \boldsymbol{\rho}^{(l)} & \dots & \boldsymbol{\rho}^{(l)} & \boldsymbol{\rho}^{(l)} \\ \lambda^{(l)} & \lambda^{(l)} & \dots & \lambda^{(l)} & \lambda^{(l)} \\ \hat{y}^{(l)} & \hat{y}^{(l)} & \dots & \hat{y}^{(l)} & \hat{y}^{(l)} \end{pmatrix} \in \mathbb{R}^{(d+2n+7) \times (N+1)} \quad (36)$$

where

$$\mathbf{x}_i^{(l)} \in \mathbb{R}^{d+1}, \quad \boldsymbol{\phi}_i^{(l)}, \boldsymbol{\rho}^{(l)} \in \mathbb{R}^{n+1}, \quad \text{and} \quad w_i^{(l)}, \lambda^{(l)}, \hat{y}^{(l)} \in \mathbb{R} \quad \text{for all } 1 \leq i \leq N + 1.$$

Note that the last three row blocks in (36) contain $N + 1$ identical copies of $\boldsymbol{\rho}^{(l)}$, $\lambda^{(l)}$ and $\hat{y}^{(l)}$. In particular, $\mathbf{H}^{(0)}$ admits a simpler form, for which we initialize as follows:

$$\mathbf{x}_i^{(0)} = [\mathbf{x}_i^\top, 1]^\top, \quad \boldsymbol{\phi}_i^{(0)} = \mathbf{0} \quad \text{for all } 1 \leq i \leq N + 1, \quad (37a)$$

$$y_i^{(0)} = y_i, \quad w_i^{(0)} = 1 \quad \text{for all } 1 \leq i \leq N, \quad (37b)$$

$$y_{N+1}^{(0)} = 0, \quad w_{N+1}^{(0)} = 0, \quad \boldsymbol{\rho}^{(0)} = \mathbf{0}, \quad \lambda^{(0)} = \hat{y}^{(0)} = 0. \quad (37c)$$

As a result, the input matrix $\mathbf{H}^{(0)}$ can be simply expressed as

$$\mathbf{H}^{(0)} = \begin{pmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \dots & \mathbf{x}_N & \mathbf{x}_{N+1} \\ 1 & 1 & \dots & 1 & 1 \\ y_1 & y_2 & \dots & y_N & 0 \\ 1 & 1 & \dots & 1 & 0 \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \end{pmatrix} \in \mathbb{R}^{(d+2n+7) \times (N+1)}. \quad (38)$$

Based on the above input structure of $\mathbf{H}^{(l)}$, we are positioned to present our construction, whose proof is postponed to Appendix A.4.

Lemma 4. *One can construct a transformer such that:*

- i) it has L layers, $M = O(1)$ attention heads per layer, and takes the matrix $\mathbf{H}^{(0)}$ (cf. (38)) as input;
- ii) the component $\boldsymbol{\rho}^{(L)}$ in the final output matrix $\mathbf{H}^{(L)}$ coincides with $\boldsymbol{\rho}_{(L-1)/2}^{\text{proximal}}$ (cf. (31b)) for some \mathbf{e}_t , where we choose

$$\lambda \asymp \left(\frac{\log N}{N}\right)^{1/6} C_{\mathcal{F}}^{-1/3} \varepsilon^2_{\text{dis}} + \sqrt{\frac{\log N}{N}} (C_{\mathcal{F}} + \sigma) + C_{\mathcal{F}}^{-1} \varepsilon^2_{\text{dis}}, \quad (39a)$$

$$\max_{1 \leq t \leq T} \|\mathbf{e}_t\|_1 \lesssim \frac{C_{\mathcal{F}}}{(L+n)nN}, \quad (39b)$$

with

$$\hat{\varepsilon} := \sqrt{\frac{\log N}{N}} C_{\mathcal{F}} (\sigma + C_{\mathcal{F}}) + \varepsilon^2_{\text{dis}} + \frac{nC_{\mathcal{F}}^2}{L}; \quad (39c)$$

- iii) for every $f \in \mathcal{F}$, the components $\boldsymbol{\rho}^{(L)}$ and $\hat{y}^{(L)}$ in the final output matrix $\mathbf{H}^{(L)}$ (cf. (36)) satisfy

$$\ell(\boldsymbol{\rho}^{(L)}) - \ell(\boldsymbol{\rho}^*) \lesssim \sqrt{\frac{\log N}{N}} C_{\mathcal{F}} (\sigma + C_{\mathcal{F}}) + \varepsilon^2_{\text{dis}} + \frac{nC_{\mathcal{F}}^2}{L} \quad (40a)$$

$$(\boldsymbol{\phi}_{N+1}^{\top} \boldsymbol{\rho}^{(L)} - \hat{y}^{(L)})^2 \lesssim \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}^2 \quad (40b)$$

with probability at least $1 - O(N^{-10})$.

In words, Lemma 4 demonstrates how a single multi-layer transformer can be constructed to emulate the iterations of the proximal gradient method (cf. (31b)) and achieve high optimization accuracy (see (40a)), while in the meantime controlling the fitting error between $\boldsymbol{\phi}_{N+1}^{\top} \boldsymbol{\rho}^{(L)}$ and $\hat{y}^{(L)}$ (see (40b)). And all this holds simultaneously for all f in the function class of interest.

Step 5: putting everything together. Equipped with the above results, we are ready to put all pieces together towards establishing our theory.

In view of (40a), the output $\boldsymbol{\rho}^{(L)}$ of our constructed transformer in Lemma 4 satisfies

$$\begin{aligned} \ell(\boldsymbol{\rho}^{(L)}) - \ell(\boldsymbol{\rho}^*) &\leq \varepsilon_{\text{opt}}, \\ \text{by taking } \varepsilon_{\text{opt}} &\asymp \sqrt{\frac{\log N}{N}} C_{\mathcal{F}} (\sigma + C_{\mathcal{F}}) + \varepsilon^2_{\text{dis}} + \frac{nC_{\mathcal{F}}^2}{L}. \end{aligned} \quad (41)$$

Invoking Lemma 2 with $\hat{\boldsymbol{\rho}} = \boldsymbol{\rho}^{(L)}$ reveals the following bound concerning the estimate $\boldsymbol{\rho}^{(L)}$:

$$\mathbb{E}[(\boldsymbol{\phi}_{N+1}^{\top} \boldsymbol{\rho}^{(L)} - f(\mathbf{x}_{N+1}))^2] \stackrel{(a)}{\lesssim} \sqrt{\frac{\log N}{N}} (C_{\mathcal{F}}^2 + \lambda^{-2} \varepsilon_{\text{opt}}^2 + \sigma C_{\mathcal{F}}) + \varepsilon^2_{\text{dis}} + \lambda C_{\mathcal{F}} + \varepsilon_{\text{opt}}$$

$$\begin{aligned}
&\stackrel{(b)}{\lesssim} \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(C_{\mathcal{F}} + \sigma) + C_{\mathcal{F}}^2 \varepsilon + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{2/3} \\
&\quad + \varepsilon_{\text{opt}} + \lambda C_{\mathcal{F}} + \sqrt{\frac{\log N}{N}} \lambda^{-2} \varepsilon_{\text{opt}}^2 \\
&\stackrel{(c)}{\lesssim} \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(C_{\mathcal{F}} + \sigma) + C_{\mathcal{F}}^2 \varepsilon + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{2/3} \\
&\quad + \frac{n C_{\mathcal{F}}^2}{L} + \lambda C_{\mathcal{F}} + \sqrt{\frac{\log N}{N}} \lambda^{-2} \varepsilon_{\text{opt}}^2
\end{aligned} \tag{42}$$

holds with probability at least $1 - O(N^{-10})$, where in (a) we have used (29) and a basic bound $\varepsilon_{\text{dis}} \lesssim C_{\mathcal{F}}$ that holds under our assumptions, (b) arises from (23), and (c) results from (41). Next, let us bound the last two terms in (42).

- To begin with, with λ as specified in (39a), we obtain

$$\begin{aligned}
\lambda C_{\mathcal{F}} &\lesssim \left(\frac{\log N}{N} \right)^{1/6} C_{\mathcal{F}}^{2/3} \hat{\varepsilon}^{2/3} + \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(C_{\mathcal{F}} + \sigma) + \varepsilon_{\text{dis}}^2 \\
&\stackrel{(a)}{\lesssim} \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(C_{\mathcal{F}} + \sigma) + \varepsilon_{\text{dis}}^2 + \hat{\varepsilon}
\end{aligned} \tag{43}$$

$$\stackrel{(b)}{\lesssim} \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(C_{\mathcal{F}} + \sigma) + C_{\mathcal{F}}^2 \varepsilon + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{2/3} + \hat{\varepsilon}, \tag{44}$$

where (a) invokes Young's inequality to derive

$$\left(\frac{\log N}{N} \right)^{1/6} C_{\mathcal{F}}^{2/3} \hat{\varepsilon}^{2/3} \leq \frac{1}{3} \left(\left(\frac{\log N}{N} \right)^{1/6} C_{\mathcal{F}}^{2/3} \right)^3 + \frac{2}{3} (\hat{\varepsilon}^{2/3})^{3/2} \leq \frac{1}{3} \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}^2 + \frac{2}{3} \hat{\varepsilon}, \tag{45}$$

and (b) applies (23).

- Moreover, the last term in (42) satisfies

$$\begin{aligned}
\sqrt{\frac{\log N}{N}} \lambda^{-2} \varepsilon_{\text{opt}}^2 &\stackrel{(a)}{\lesssim} \left(\frac{\log N}{N} \right)^{1/2} \left(\frac{\log N}{N} \right)^{-1/3} C_{\mathcal{F}}^{2/3} \hat{\varepsilon}^{-4/3} \varepsilon_{\text{opt}}^2 \stackrel{(b)}{\lesssim} \left(\frac{\log N}{N} \right)^{1/6} C_{\mathcal{F}}^{2/3} \hat{\varepsilon}^{2/3} \\
&\stackrel{(c)}{\lesssim} \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}^2 + \hat{\varepsilon}.
\end{aligned} \tag{46}$$

Here, (a) results from the fact that $\lambda^2 \gtrsim \left(\frac{\log N}{N} \right)^{1/3} C_{\mathcal{F}}^{-2/3} \hat{\varepsilon}^{2/3}$, (b) is valid since $\hat{\varepsilon} \asymp \varepsilon_{\text{opt}}$ (see (39c) and (41)), whereas (c) makes use of (45).

Substituting (44) and (46) into (42), and recalling the definition of $\hat{\varepsilon}$ in (39c), we arrive at

$$\mathbb{E}[(\phi_{N+1}^{\top} \boldsymbol{\rho}^{(L)} - f(\mathbf{x}_{N+1}))^2] \lesssim \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(\sigma + C_{\mathcal{F}}) + C_{\mathcal{F}}^2 \varepsilon + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{2/3} + \frac{n C_{\mathcal{F}}^2}{L} \tag{47}$$

with probability exceeding $1 - O(N^{-10})$, where we have used the fact that (see (39c) and (23))

$$\hat{\varepsilon} \lesssim \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(\sigma + C_{\mathcal{F}}) + C_{\mathcal{F}}^2 \varepsilon + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{2/3} + \frac{n C_{\mathcal{F}}^2}{L}.$$

Combining (47) with (40b) and recalling $\hat{y}_{N+1} = \hat{y}^{(L)}$ (cf. (11)), we reach

$$\begin{aligned}
\mathbb{E}[(\hat{y}_{N+1} - f(\mathbf{x}_{N+1}))^2] &\leq 2\mathbb{E}[(\hat{y}_{N+1} - \phi_{N+1}^{\top} \boldsymbol{\rho}^{(L)})^2] + 2\mathbb{E}[(\phi_{N+1}^{\top} \boldsymbol{\rho}^{(L)} - f(\mathbf{x}_{N+1}))^2] \\
&\lesssim \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}^2 + \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(\sigma + C_{\mathcal{F}}) + C_{\mathcal{F}}^2 \varepsilon + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{2/3} + \frac{n C_{\mathcal{F}}^2}{L} \\
&\asymp \sqrt{\frac{\log N}{N}} C_{\mathcal{F}}(\sigma + C_{\mathcal{F}}) + C_{\mathcal{F}}^2 \varepsilon + C_{\mathcal{F}}^2 \left(\frac{\log |\mathcal{N}_{\varepsilon}|}{n} \right)^{2/3} + \frac{n C_{\mathcal{F}}^2}{L}.
\end{aligned}$$

We can therefore conclude the proof of Theorem 1 by taking $\varepsilon \leq \sqrt{\log N/N} + n/L$.

5 Discussion

In this work, we have investigated the in-context learning capabilities of transformers through the lens of universal function approximation, establishing approximation guarantees that extend far beyond the previously studied convex settings or the problems of learning linear functions. We have demonstrated that: for a fairly general function class \mathcal{F} satisfying mild Fourier-type conditions, one can construct a universal multi-layer transformer achieving the following intriguing property: for every task represented by some function $f \in \mathcal{F}$, the constructed transformer can readily utilize the N input-output examples to achieve the prediction error on the order of $1/\sqrt{N} + n/L + (\log |\mathcal{N}_\varepsilon|/n)^{2/3}$ (up to some logarithmic factor), where \mathcal{N}_ε denotes an ε -cover and we choose the auxiliary input dimension n to exceed the log of a certain covering number. Our analysis imposes only fairly mild assumptions on \mathcal{F} , requiring neither linearity in the function class nor convexity in the learning problem, thereby offering a comprehensive theoretical understanding for the empirical success of transformer-based models in real-world tasks.

Looking forward, we would like to establish tight performance bounds on the prediction error, particularly with respect to its dependence on the number of examples N and the number of layers L . In addition, while our current analysis has focused on the approximation ability of a universal transformer, understanding how pretraining or finetuning affects generalization in in-context learning remains an open question, which we leave for future studies. Finally, the current paper is concerned with constructed approximation, and it would be of great interest to understand end-to-end training dynamics for transformers, which are highly nonconvex and call for innovative technical ideas.

Acknowledgments

G. Li is supported in part by the Chinese University of Hong Kong Direct Grant for Research and the Hong Kong Research Grants Council ECS 2191363. Y. Wei is supported in part by the NSF grants CCF-2106778, CCF-2418156 and CAREER award DMS-2143215. Y. Chen is supported in part by the Alfred P. Sloan Research Fellowship, the ONR grants N00014-22-1-2354 and N00014-25-1-2344, the NSF grants 2221009 and 2218773, and the Amazon Research Award.

A Proof of key lemmas

In this section, we provide complete proofs of the key lemmas introduced in Section 4.

A.1 Proof of Lemma 1

For notational convenience, we shall write the Fourier transform F_f as

$$F_f(\omega) = |F_f(\omega)|e^{j\theta_f(\omega)}, \quad (48)$$

with $\theta_f(\omega)$ representing the angle. By virtue of the Fourier transform of $f(\mathbf{x})$, we have

$$\begin{aligned} f(\mathbf{x}) - f(\mathbf{0}) &= \int_{\omega} e^{j\omega^\top \mathbf{x}} F_f(\omega) d\omega - \int_{\omega} F_f(\omega) d\omega \\ &= \int_{\omega} (e^{j\omega^\top \mathbf{x}} e^{j\theta_f(\omega)} - e^{j\theta_f(\omega)}) |F_f(\omega)| d\omega \\ &\stackrel{(a)}{=} \int_{\omega \neq 0} \left(\cos(\omega^\top \mathbf{x} + \theta_f(\omega)) - \cos(\theta_f(\omega)) \right) |F_f(\omega)| d\omega, \end{aligned} \quad (49)$$

where (a) follows since $f(\mathbf{x})$ is real-valued. In view of (49), we would like to construct a finite collection of features to approximately represent $f(\mathbf{x}) - f(\mathbf{0})$, using the sigmoid function defined in (19).

Towards this end, we first introduce a quantity related to the difference between $\phi(\tau x)$ and the unit step function $\mathbf{1}(x > 0)$ as follows:

$$\delta_\tau := \inf_{0 < \varepsilon \leq 1/2} \left\{ 2\varepsilon + \sup_{|x| \geq \varepsilon} |\phi(\tau x) - \mathbf{1}(x > 0)| \right\}. \quad (50)$$

The proof of Lemma 1 relies heavily upon the following result, whose proof is postponed to Appendix A.1.1.

Lemma 5. *For any $f \in \mathcal{F}$, any \mathbf{x} obeying $\|\mathbf{x}\|_2 \leq 1$, any positive integer m , and any $\tau > 2$, we have*

$$|f(\mathbf{x}) - f(\mathbf{0}) - f^{\text{approx}}(\mathbf{x})| \leq C_{\mathcal{F}} \left(3\delta_{\tau} + \frac{\pi}{m} \right), \quad (51)$$

where for any $f \in \mathcal{F}$ we define

$$f^{\text{approx}}(\mathbf{x}) := \sum_{k=1}^m \int_{\boldsymbol{\omega} \neq \mathbf{0}} \int_t \rho_f(k, t, \boldsymbol{\omega}) \phi(\tau(\|\boldsymbol{\omega}\|_2^{-1} \boldsymbol{\omega}^{\top} \mathbf{x} - t)) \Lambda(k, dt, d\boldsymbol{\omega}) \quad (52)$$

with $\phi(\cdot)$ defined in (19). Here, $\Lambda(k, dt, d\boldsymbol{\omega})$ is a probability measure on $[m] \times \mathbb{R} \times \mathbb{R}^d$ independent from f , while $\rho_f(k, t, \boldsymbol{\omega})$ denotes some weight function depending on $f, k, t, \boldsymbol{\omega}$ such that

$$|\rho_f(k, t, \boldsymbol{\omega})| \leq 3m \left(C_{\mathcal{F}} - \sup_{\tilde{f} \in \mathcal{F}} |\tilde{f}(\mathbf{0})| \right) \quad \text{and} \quad \mathbb{E}_{(k, t, \boldsymbol{\omega}) \sim \Lambda} [|\rho_f(k, t, \boldsymbol{\omega})|] \leq 3 \left(C_{\mathcal{F}} - \sup_{\tilde{f} \in \mathcal{F}} |\tilde{f}(\mathbf{0})| \right). \quad (53)$$

Next, we would like to obtain a more succinct finite-sum approximation of the integration in (52) via random subsampling. Let us draw n independent samples $(k_i, t_i, \boldsymbol{\omega}_i)$ ($1 \leq i \leq n$) from the probability measure $\Lambda(k, dt, d\boldsymbol{\omega})$. Applying the Bernstein inequality and the union bound over $\mathcal{N}_{\varepsilon}$ reveals that: with probability at least $3/4$,

$$\left| \hat{f}^{\text{approx}}(\hat{\mathbf{x}}) - \frac{1}{n} \sum_{i=1}^n \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^{\top} \hat{\mathbf{x}} - t_i)) \right| \lesssim \max \left\{ \sqrt{\frac{m\tilde{C}_{\mathcal{F}}^2 \log |\mathcal{N}_{\varepsilon}|}{n}}, \frac{m\tilde{C}_{\mathcal{F}} \log |\mathcal{N}_{\varepsilon}|}{n} \right\}$$

holds simultaneously for all $(\hat{f}, \hat{\mathbf{x}}) \in \mathcal{N}_{\varepsilon}$, where $\hat{f}^{\text{approx}}(\hat{\mathbf{x}})$ is defined in (52) with f and \mathbf{x} taken respectively to be \hat{f} and $\hat{\mathbf{x}}$, and

$$\tilde{C}_{\mathcal{F}} := C_{\mathcal{F}} - \sup_{\tilde{f} \in \mathcal{F}} |\tilde{f}(\mathbf{0})|. \quad (54)$$

Note that the above application of the Bernstein inequality has used the following bounds for any fixed \hat{f} and $\hat{\mathbf{x}}$:

$$\begin{aligned} \mathbb{E}[\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^{\top} \hat{\mathbf{x}} - t_i))] &= \hat{f}^{\text{approx}}(\hat{\mathbf{x}}), \\ |\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^{\top} \hat{\mathbf{x}} - t_i))| &\leq 3m\tilde{C}_{\mathcal{F}}, \\ \text{Var}(\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^{\top} \hat{\mathbf{x}} - t_i))) &\leq \mathbb{E}[(\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^{\top} \hat{\mathbf{x}} - t_i)))^2] \\ &\leq \sup_{1 \leq k \leq m, t \in \mathbb{R}, \boldsymbol{\omega} \in \mathbb{R}^d} |\rho_{\hat{f}}(k, t, \boldsymbol{\omega})| \mathbb{E}[|\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)|] \leq 9m\tilde{C}_{\mathcal{F}}^2, \end{aligned}$$

where the last relation arises from (53).

Similarly, applying the Bernstein inequality and the union bound once again yields that: with probability at least $3/4$,

$$\left| \mathbb{E}_{(k, t, \boldsymbol{\omega}) \sim \Lambda} [\rho_{\hat{f}}(k, t, \boldsymbol{\omega})] - \frac{1}{n} \sum_{i=1}^n \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \right| \lesssim \max \left\{ \sqrt{\frac{m\tilde{C}_{\mathcal{F}}^2 \log |\mathcal{N}_{\varepsilon}|}{n}}, \frac{m\tilde{C}_{\mathcal{F}} \log |\mathcal{N}_{\varepsilon}|}{n} \right\}$$

holds simultaneously for all $(\hat{f}, \hat{\mathbf{x}}) \in \mathcal{N}_{\varepsilon}$, where we have used the following bounds:

$$\begin{aligned} |\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)| &\leq 3m\tilde{C}_{\mathcal{F}}; \\ \text{Var}(|\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)|) &\leq \mathbb{E}[|\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)|^2] \leq \sup_{1 \leq k \leq m, t \in \mathbb{R}, \boldsymbol{\omega} \in \mathbb{R}^d} |\rho_{\hat{f}}(k, t, \boldsymbol{\omega})| \mathbb{E}[|\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)|] \leq 9m\tilde{C}_{\mathcal{F}}^2. \end{aligned}$$

If $n \gtrsim m \log |\mathcal{N}_\varepsilon|$ and $\tau > 4$, then we can see that $\sqrt{\frac{m\tilde{C}_\mathcal{F}^2 \log |\mathcal{N}_\varepsilon|}{n}} \gtrsim \frac{m\tilde{C}_\mathcal{F} \log |\mathcal{N}_\varepsilon|}{n}$. Thus with probability at least $3/4$, one has

$$\frac{1}{n} \sum_{i=1}^n |\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)| < \mathbb{E}_{(k,t,\boldsymbol{\omega}) \sim \Lambda} [|\rho_{\hat{f}}(k, t, \boldsymbol{\omega})|] + \sqrt{\frac{c_1 m \tilde{C}_\mathcal{F}^2 \log |\mathcal{N}_\varepsilon|}{n}} \stackrel{(a)}{\leq} 3\tilde{C}_\mathcal{F} + \tilde{C}_\mathcal{F} = 4\tilde{C}_\mathcal{F}$$

simultaneously for all $(\hat{f}, \hat{\mathbf{x}}) \in \mathcal{N}_\varepsilon$, where $c_1 > 0$ is some universal constant, and (a) is valid under the condition that $n \geq c_1 m \log |\mathcal{N}_\varepsilon|$. This further shows that, with probability at least $3/4$,

$$|f(\mathbf{0})| + \frac{1}{n} \sum_{i=1}^n |\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)| < |f(\mathbf{0})| + 4\tilde{C}_\mathcal{F} \leq 4C_\mathcal{F}$$

simultaneously for all $(\hat{f}, \hat{\mathbf{x}}) \in \mathcal{N}_\varepsilon$.

Having established several key properties for the epsilon-cover \mathcal{N}_ε , we now extend these properties to all $f \in \mathcal{F}$. For any $f \in \mathcal{F}$, there exists some $(\hat{f}, \hat{\mathbf{x}}) \in \mathcal{N}_\varepsilon$, such that for all $\|\mathbf{x}\|_2 \leq 1$,

$$\begin{aligned} & \left| f(\mathbf{x}) - f(\mathbf{0}) - \frac{1}{n} \sum_{i=1}^n \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \mathbf{x} - t_i)) \right| \\ & \leq |f(\mathbf{x}) - f(\mathbf{0}) - \hat{f}(\hat{\mathbf{x}}) + \hat{f}(\mathbf{0})| + \left| \hat{f}(\hat{\mathbf{x}}) - \hat{f}(\mathbf{0}) - \frac{1}{n} \sum_{i=1}^n \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \hat{\mathbf{x}} - t_i)) \right| \\ & \quad + \left| \frac{1}{n} \sum_{i=1}^n \left(\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \hat{\mathbf{x}} - t_i)) - \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \mathbf{x} - t_i)) \right) \right| \\ & \stackrel{(a)}{\leq} C_\mathcal{F} \left(\varepsilon + 3\delta_\tau + \frac{\pi}{m} + \sqrt{\frac{c_1 m \log |\mathcal{N}_\varepsilon|}{n}} + 4\tau\varepsilon \right), \end{aligned}$$

where c_1 is a universal constant, (a) uses the definition of the epsilon-cover (16), and takes advantage of the fact that

$$\begin{aligned} & \left| \frac{1}{n} \sum_{i=1}^n \left(\rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \hat{\mathbf{x}} - t_i)) - \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \mathbf{x} - t_i)) \right) \right| \\ & = \frac{1}{n} \sum_{i=1}^n \left| \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \right| \left| \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \hat{\mathbf{x}} - t_i)) - \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \mathbf{x} - t_i)) \right| \\ & \leq \frac{1}{n} \sum_{i=1}^n \left| \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \right| \left| \tau \|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top (\hat{\mathbf{x}} - \mathbf{x}) \right| \\ & \leq \frac{1}{n} \sum_{i=1}^n \left| \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \right| \tau \|\hat{\mathbf{x}} - \mathbf{x}\|_2 \leq \frac{1}{n} \sum_{i=1}^n \left| \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \right| \tau \varepsilon \leq 4C_\mathcal{F} \tau \varepsilon. \end{aligned}$$

Recalling that $\tau > 4$, and taking $m = \lceil (\frac{n}{\log |\mathcal{N}_\varepsilon|})^{1/3} \rceil$, we arrive at

$$\left| f(\mathbf{x}) - f(\mathbf{0}) - \frac{1}{n} \sum_{i=1}^n \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i) \phi(\tau(\|\boldsymbol{\omega}_i\|_2^{-1} \boldsymbol{\omega}_i^\top \mathbf{x} - t_i)) \right| \lesssim C_\mathcal{F} \left(\delta_\tau + \left(\frac{\log |\mathcal{N}_\varepsilon|}{n} \right)^{\frac{1}{3}} + \tau\varepsilon \right). \quad (55)$$

To finish up, it suffices to prove that $\delta_\tau \leq 1/\tau$. Recalling the definition of δ_τ in (50) and the function $\phi(\tau x)$, we have

$$|\phi(\tau x) - \mathbf{1}(x > 0)| = (1/2 - \tau|x|) \mathbf{1}\{\tau|x| \leq 1/2\},$$

and as a result,

$$\sup_{|x| \geq \varepsilon'} |\phi(\tau x) - \mathbf{1}(x > 0)| = \sup_{|x| \geq \varepsilon'} (1/2 - \tau|x|) \mathbf{1}\{\tau|x| \leq 1/2\} = \max \left\{ \frac{1}{2} - \tau\varepsilon', 0 \right\}.$$

This reveals that

$$\delta_\tau = \inf_{0 \leq \varepsilon' \leq 1/2} \left\{ 2\varepsilon + \max \left\{ \frac{1}{2} - \tau\varepsilon', 0 \right\} \right\} = \frac{1}{\tau}$$

for any $\tau \geq 2$. Substituting it into (55) and recalling the definition of $\phi_i^{\text{feature}}(\mathbf{x})$ in (22), we establish the claimed result (20) with $\rho_{f,i}^* = \rho_{\hat{f}}(k_i, t_i, \boldsymbol{\omega}_i)$.

A.1.1 Proof of Lemma 5

Consider any positive integer m . According to (49), in order to approximate $f(\mathbf{x}) - f(\mathbf{0})$, it is helpful to approximate $\cos(\boldsymbol{\omega}^\top \mathbf{x} + \theta_f(\boldsymbol{\omega})) - \cos(\theta_f(\boldsymbol{\omega}))$. Towards this end, we first make the following observation

$$\left| \frac{\partial}{\partial \theta} \left(\frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + \theta) - \cos(\theta)}{\|\boldsymbol{\omega}\|_2} \right) \right| = \left| \frac{\sin(\boldsymbol{\omega}^\top \mathbf{x} + \theta) - \sin(\theta)}{\|\boldsymbol{\omega}\|_2} \right| \leq \frac{|\boldsymbol{\omega}^\top \mathbf{x}|}{\|\boldsymbol{\omega}\|_2} \leq \frac{\|\boldsymbol{\omega}\|_2 \|\mathbf{x}\|_2}{\|\boldsymbol{\omega}\|_2} = 1 \quad (56)$$

for any $\theta \in [0, 2\pi)$. Then (56) implies that, for any $\theta \in [0, 2\pi)$, there exists some integer $k(\theta)$ obeying $1 \leq k(\theta) \leq m$ such that

$$\sup_{\mathbf{x} \in \mathcal{B}} \left| \frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + \theta) - \cos(\theta)}{\|\boldsymbol{\omega}\|_2} - \frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + 2\pi k(\theta)/m) - \cos(2\pi k(\theta)/m)}{\|\boldsymbol{\omega}\|_2} \right| \leq \left| \theta - \frac{2\pi k(\theta)}{m} \right| \leq \frac{\pi}{m}. \quad (57)$$

It has been shown by Barron (1993, Lemma 5) that: for any $\boldsymbol{\omega} \in \mathbb{R}^d$, $\theta \in [0, 2\pi)$ and any $\tau \geq 2$, there exists some function $\gamma_{\boldsymbol{\omega}, \theta}(\cdot)$ such that

$$\left| \frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + \theta) - \cos(\theta)}{\|\boldsymbol{\omega}\|_2} - \int_t \gamma_{\boldsymbol{\omega}, \theta}(t) \phi(\tau(\|\boldsymbol{\omega}\|_2^{-1} \boldsymbol{\omega}^\top \mathbf{x} - t)) dt \right| \leq 3\delta_\tau, \quad (58a)$$

$$\int_t |\gamma_{\boldsymbol{\omega}, \theta}(t)| dt \leq 3. \quad (58b)$$

To make it self-contained, we will present the proof of (58) towards the end of this section.

Combine (57) and (58a) and invoke the triangle inequality to reach

$$\begin{aligned} & \left| \frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + \theta) - \cos(\theta)}{\|\boldsymbol{\omega}\|_2} - \int_t \gamma_{\boldsymbol{\omega}, 2\pi k(\theta)/m}(t) \phi(\tau(\|\boldsymbol{\omega}\|_2^{-1} \boldsymbol{\omega}^\top \mathbf{x} - t)) dt \right| \\ & \leq \left| \frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + \theta) - \cos(\theta)}{\|\boldsymbol{\omega}\|_2} - \frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + 2\pi k(\theta)/m) - \cos(2\pi k(\theta)/m)}{\|\boldsymbol{\omega}\|_2} \right| \\ & \quad + \left| \frac{\cos(\boldsymbol{\omega}^\top \mathbf{x} + 2\pi k(\theta)/m) - \cos(2\pi k(\theta)/m)}{\|\boldsymbol{\omega}\|_2} - \int_t \gamma_{\boldsymbol{\omega}, 2\pi k(\theta)/m}(t) \phi(\tau(\|\boldsymbol{\omega}\|_2^{-1} \boldsymbol{\omega}^\top \mathbf{x} - t)) dt \right| \\ & \leq 3\delta_\tau + \frac{\pi}{m}. \end{aligned}$$

Substitution into (49) then gives

$$\left| f(\mathbf{x}) - f(\mathbf{0}) - f^{\text{approx}}(\mathbf{x}) \right| \leq \left(2\delta_\tau + \frac{\pi}{m} \right) \int_{\boldsymbol{\omega}} \|\boldsymbol{\omega}\|_2 |F_f(\boldsymbol{\omega})| d\boldsymbol{\omega} \leq C_{\mathcal{F}} \left(3\delta_\tau + \frac{\pi}{m} \right), \quad (59)$$

where we define

$$f^{\text{approx}}(\mathbf{x}) := \int_{\boldsymbol{\omega} \neq \mathbf{0}} \left(\int_t \phi(\tau(\|\boldsymbol{\omega}\|_2^{-1} \boldsymbol{\omega}^\top \mathbf{x} - t)) \gamma_{\boldsymbol{\omega}, 2\pi k(\theta_f(\boldsymbol{\omega}))/m}(t) dt \right) \|\boldsymbol{\omega}\|_2 |F_f(\boldsymbol{\omega})| d\boldsymbol{\omega}, \quad (60)$$

and recall the definition of $C_{\mathcal{F}}$ in (14).

Next, let us define the following key quantity independent of f :

$$\Gamma_{\mathcal{F}} := \sum_{k=1}^m \int_{\boldsymbol{\omega}} \left(\int_t |\gamma_{\boldsymbol{\omega}, 2\pi k/m}(t)| dt \right) \|\boldsymbol{\omega}\|_2 F^{\text{sup}}(\boldsymbol{\omega}) d\boldsymbol{\omega},$$

which can be bounded by

$$\Gamma_{\mathcal{F}} \leq 3 \sum_{k=1}^m \int_{\omega} \|\omega\|_2 F^{\sup}(\omega) d\omega \leq 3m \left(C_{\mathcal{F}} - \sup_{f' \in \mathcal{F}} |f'(\mathbf{0})| \right) < \infty \quad (61)$$

for any $f \in \mathcal{F}$. As a result,

$$\Lambda(k, dt, d\omega) := \Gamma_{\mathcal{F}}^{-1} |\gamma_{\omega, 2\pi k/m}(t)| \|\omega\|_2 F^{\sup}(\omega) dt d\omega \quad (62)$$

forms a valid probability measure on $[m] \times \mathbb{R} \times \mathbb{R}^d$, given that

$$\Lambda(k, dt, d\omega) \geq 0 \quad \text{and} \quad \sum_{k=1}^m \int_{\omega} \int_t \Lambda(k, dt, d\omega) = 1.$$

Importantly, this probability measure $\Lambda(k, dt, d\omega)$ allows one to express the function $f^{\text{approx}}(\mathbf{x})$ (cf. (60)) as

$$f^{\text{approx}}(\mathbf{x}) = \sum_{k=1}^m \int_{\omega \neq 0} \int_t \rho_f(k, t, \omega) \phi(\tau(\|\omega\|_2^{-1} \omega^{\top} \mathbf{x} - t)) \Lambda(k, dt, d\omega),$$

where

$$\rho_f(k, t, \omega) := \frac{\gamma_{\omega, 2\pi k/m}(t) \mathbb{1}\{k = k(\theta_f(\omega))\} \|\omega\|_2 |F_f(\omega)| dt d\omega}{\Lambda(k, dt, d\omega)}. \quad (63)$$

Substitution into (59) immediately establishes the advertised result (51).

To finish up, it suffices to observe that (i)

$$|\rho_f(k, t, \omega)| \leq \frac{|\gamma_{\omega, 2\pi k/m}(t)| \mathbb{1}\{k = k(\theta_f(\omega))\} \|\omega\|_2 |F_f(\omega)|}{\Gamma_{\mathcal{F}}^{-1} |\gamma_{\omega, 2\pi k/m}(t)| \|\omega\|_2 F^{\sup}(\omega)} \leq \Gamma_{\mathcal{F}} \leq 3m \left(C_{\mathcal{F}} - \sup_{\tilde{f} \in \mathcal{F}} |\tilde{f}(\mathbf{0})| \right)$$

as a consequence of (61), (62) and (63), and (ii)

$$\begin{aligned} \sum_{k=1}^m \int_{\omega \neq 0} \int_t |\rho_f(k, t, \omega)| \Lambda(k, dt, d\omega) &= \int_{\omega \neq 0} \int_t |\gamma_{\omega, 2\pi k(\theta_f(\omega))/m}(t)| dt \|\omega\|_2 |F_f(\omega)| d\omega \\ &\leq 3 \int_{\omega \neq 0} \|\omega\|_2 |F_f(\omega)| d\omega \leq 3 \left(C_{\mathcal{F}} - \sup_{\tilde{f} \in \mathcal{F}} |\tilde{f}(\mathbf{0})| \right) \end{aligned}$$

as a result of (58b).

Proof of (58). This proof is similar to the proof of Lemma 5 in Barron (1993). We begin by claiming the following result (to be proven shortly): for any function $g : [-1, 1] \rightarrow \mathbb{R}$ satisfying $|g'(z)| \leq 1$, $|g(-1)| \leq 1$, and for any $\tau > 0$, there exists a function $\gamma(\cdot)$ such that

$$\left| g(z) - \int \gamma(t) \phi(\tau(z+t)) dt \right| \leq 3\delta_{\tau}, \quad (64a)$$

where

$$\int |\gamma(t)| dt \leq 3. \quad (64b)$$

Here, δ_{τ} and $\phi(\cdot)$ have been defined in (50) and (19), respectively. Suppose for the moment that this claim is valid. To establish (58), we find it convenient to introduce another auxiliary function

$$g_{\omega, \theta}(z) := \frac{\cos(\|\omega\|_2 z + \theta) - \cos(\theta)}{\|\omega\|_2}$$

defined on $[-1, 1]$, which clearly satisfies

$$|g'_{\omega, \theta}(z)| \leq 1 \quad \text{for all } z \in [-1, 1], \quad \text{and} \quad |g_{\omega, \theta}(-1)| = |g_{\omega, \theta}(-1) - g_{\omega, \theta}(0)| \leq 1.$$

Applying inequality (64) to $g_{\omega, \theta}(\cdot)$ and recognizing that $g_{\omega, \theta}\left(\frac{\omega^\top \mathbf{x}}{\|\omega\|_2}\right) = \frac{\cos(\omega^\top \mathbf{x} + \theta) - \cos(\theta)}{\|\omega\|_2}$ establish (58).

To finish up, it remains to establish (64). Define

$$\gamma(t) = \begin{cases} g'(t), & \text{if } t \in [-1, 1], \\ \text{sign}(g(-1)), & \text{if } -1 - |g(-1)| \leq t < -1, \\ 0, & \text{else.} \end{cases} \quad (65)$$

which clearly satisfies

$$|\gamma(t)| \leq 1, \quad \int_t^1 |\gamma(t)| dt = \int_{-1-|g(-1)|}^1 |\gamma(t)| dt \leq 2 + |g(-1)| \leq 3.$$

Observe that for any $z \in [-1, 1]$, it holds that

$$\int_{-\infty}^{\infty} \gamma(t) \mathbf{1}(z - t > 0) dt = \int_{-\infty}^z \gamma(t) dt = \int_{-1-|g(-1)|}^{-1} \text{sign}(g(-1)) dt + \int_{-1}^z g'(t) dt = g(-1) + g(z) - g(-1) = g(z).$$

Recalling the definition of δ_τ in (50), we arrive at

$$\begin{aligned} \left| g(z) - \int_{-\infty}^{\infty} \gamma(t) \phi(\tau(z - t)) dt \right| &= \left| \int_{-\infty}^{\infty} \gamma(t) (\mathbf{1}(z - t > 0) - \phi(\tau(z - t))) dt \right| \\ &\leq \inf_{0 < \varepsilon \leq 1/2} \left\{ \int_{t: |z-t| \leq \varepsilon} |\gamma(t)| dt + \int_{t: |z-t| \geq \varepsilon} |\gamma(t)| |\mathbf{1}(z - t > 0) - \phi(\tau(z - t))| dt \right\} \\ &\leq \inf_{0 < \varepsilon \leq 1/2} \left\{ 2\varepsilon + 3 \sup_{x: |x| \geq \varepsilon} |\phi(\tau x) - \mathbf{1}(x > 0)| \right\} \leq 3\delta_\tau \end{aligned}$$

as claimed.

A.2 Proof of Lemma 2

In view of (21) in Lemma 1, we have

$$\|\rho^*\|_1 \lesssim C_{\mathcal{F}}.$$

A key step in this proof is to establish the following lemma, whose proof is postponed to Appendix A.2.1.

Lemma 6. *With probability at least $1 - O(N^{-11})$, for any $\rho \in \mathbb{R}^{n+1}$ (which can be statistically dependent on $\{(\mathbf{x}_i, y_i)\}_{1 \leq i \leq N}$ but not on \mathbf{x}_{N+1}), we have*

$$\begin{aligned} &\left| \frac{1}{N} \sum_{i=1}^N [(y_i - \phi_i^\top \rho)^2 - z_i^2] - \mathbb{E}[(f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \rho)^2] \right| \\ &\lesssim \sqrt{\frac{\log N}{N}} (\varepsilon_{\text{dis}}^2 + \|\rho - \rho^*\|_1^2 + \sigma(\varepsilon_{\text{dis}} + \|\rho - \rho^*\|_1)) \end{aligned} \quad (66)$$

$$\text{and} \quad \frac{1}{N} \sum_{i=1}^N [(y_i - \phi_i^\top \rho)^2 - (y_i - \phi_i^\top \rho^*)^2] \gtrsim -\varepsilon_{\text{dis}}^2 - \sqrt{\frac{\log N}{N}} \sigma \|\rho - \rho^*\|_1. \quad (67)$$

Here, the expectation in (66) is taken over the randomness of \mathbf{x}_{N+1} .

Armed with this lemma, we can proceed to the proof of Lemma 2.

Proof of (30). Suppose for the moment that

$$\|\hat{\boldsymbol{\rho}}\|_1 > 4\|\boldsymbol{\rho}^*\|_1 + C_{\mathcal{F}} + 4\lambda^{-1}\varepsilon_{\text{opt}}, \quad (68)$$

then it follows from Lemma 6 that

$$\begin{aligned} \ell(\hat{\boldsymbol{\rho}}) - \ell(\boldsymbol{\rho}^*) &= \frac{1}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \hat{\boldsymbol{\rho}})^2 + \lambda \|\hat{\boldsymbol{\rho}}\|_1 - \frac{1}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}^*)^2 - \lambda \|\boldsymbol{\rho}^*\|_1 \\ &\stackrel{(a)}{\geq} \lambda \|\hat{\boldsymbol{\rho}}\|_1 - \lambda \|\boldsymbol{\rho}^*\|_1 - C\varepsilon_{\text{dis}}^2 - C\sqrt{\frac{\log N}{N}}\sigma \|\hat{\boldsymbol{\rho}} - \boldsymbol{\rho}^*\|_1 \\ &\stackrel{(b)}{>} \lambda \|\hat{\boldsymbol{\rho}}\|_1 - \frac{\lambda \|\hat{\boldsymbol{\rho}}\|_1}{4} - C\varepsilon_{\text{dis}}^2 - 2C\sqrt{\frac{\log N}{N}}\sigma \|\hat{\boldsymbol{\rho}}\|_1 \end{aligned} \quad (69)$$

for some universal constant $C > 0$. Here, (a) results from (67), whereas (b) invokes the properties (see (68)) that $\|\hat{\boldsymbol{\rho}}\|_1 > 4\|\boldsymbol{\rho}^*\|_1$ and $\|\hat{\boldsymbol{\rho}} - \boldsymbol{\rho}^*\|_1 \leq \|\hat{\boldsymbol{\rho}}\|_1 + \|\boldsymbol{\rho}^*\|_1 < 2\|\hat{\boldsymbol{\rho}}\|_1$. In addition, under our assumption on λ , we have

$$\lambda \geq 4CC_{\mathcal{F}}^{-1}\varepsilon_{\text{dis}}^2 + 8C\sqrt{\log N/N}\sigma \quad (70)$$

as long as c_λ is large enough, which then gives

$$C\varepsilon_{\text{dis}}^2 \leq \frac{\lambda C_{\mathcal{F}}}{4} < \frac{\lambda \|\hat{\boldsymbol{\rho}}\|_1}{4} \quad \text{and} \quad 2C\sqrt{\frac{\log N}{N}}\sigma \|\hat{\boldsymbol{\rho}}\|_1 < \frac{\lambda \|\hat{\boldsymbol{\rho}}\|_1}{4}.$$

These combined with (69) result in

$$\ell(\hat{\boldsymbol{\rho}}) - \ell(\boldsymbol{\rho}^*) > \lambda \|\hat{\boldsymbol{\rho}}\|_1 - \frac{\lambda \|\hat{\boldsymbol{\rho}}\|_1}{4} - \frac{\lambda \|\hat{\boldsymbol{\rho}}\|_1}{4} - \frac{\lambda \|\hat{\boldsymbol{\rho}}\|_1}{4} = \frac{\lambda}{4} \|\hat{\boldsymbol{\rho}}\|_1 > \varepsilon_{\text{opt}}$$

with the last inequality due to (68). This, however, contradicts the ε_{opt} -optimality of $\hat{\boldsymbol{\rho}}$, which in turn justifies that the assumption (68) cannot possibly hold. As a result, we can conclude that

$$\|\hat{\boldsymbol{\rho}}\|_1 \leq 4\|\boldsymbol{\rho}^*\|_1 + C_{\mathcal{F}} + 4\lambda^{-1}\varepsilon_{\text{opt}} \asymp C_{\mathcal{F}} + \lambda^{-1}\varepsilon_{\text{opt}}, \quad (71)$$

where the last relation is valid since $\|\boldsymbol{\rho}^*\|_1 \lesssim C_{\mathcal{F}}$ (see (21) and (28)).

Proof of (29). Applying (66) in Lemma 6 and making use of the fact that $\|\hat{\boldsymbol{\rho}} - \boldsymbol{\rho}^*\|_1 \leq \|\hat{\boldsymbol{\rho}}\|_1 + \|\boldsymbol{\rho}^*\|_1 \lesssim C_{\mathcal{F}} + \lambda^{-1}\varepsilon_{\text{opt}}$ (see (71)), one can demonstrate that

$$\begin{aligned} &\left| \mathbb{E}[(\boldsymbol{\phi}_{N+1}^\top \hat{\boldsymbol{\rho}} - f(\mathbf{x}_{N+1}))^2] - \frac{1}{N} \sum_{i=1}^N [(y_i - \boldsymbol{\phi}_i^\top \hat{\boldsymbol{\rho}})^2 - z_i^2] \right| \\ &\lesssim \sqrt{\frac{\log N}{N}} (\varepsilon_{\text{dis}}^2 + C_{\mathcal{F}}^2 + \lambda^{-2}\varepsilon_{\text{opt}}^2 + \sigma(\varepsilon_{\text{dis}} + C_{\mathcal{F}} + \lambda^{-1}\varepsilon_{\text{opt}})). \end{aligned} \quad (72)$$

Moreover, we make the observation that

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \hat{\boldsymbol{\rho}})^2 &\leq \ell(\hat{\boldsymbol{\rho}}) \leq \ell(\boldsymbol{\rho}^*) + \varepsilon_{\text{opt}} \leq \frac{1}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}^*)^2 + \lambda \|\boldsymbol{\rho}^*\|_1 + \varepsilon_{\text{opt}} \\ &= \frac{1}{N} \sum_{i=1}^N z_i^2 + \frac{1}{N} \sum_{i=1}^N (f(\mathbf{x}_i) - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}^*)^2 + \frac{2}{N} \sum_{i=1}^N z_i (f(\mathbf{x}_i) - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}^*) + \lambda \|\boldsymbol{\rho}^*\|_1 + \varepsilon_{\text{opt}} \\ &\leq \frac{1}{N} \sum_{i=1}^N z_i^2 + O\left(\varepsilon_{\text{dis}}^2 + \sqrt{\frac{\log N}{N}}\sigma\varepsilon_{\text{dis}} + \lambda C_{\mathcal{F}}\right) + \varepsilon_{\text{opt}}, \end{aligned} \quad (73)$$

where the last inequality follows from the properties $(f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*)^2 \lesssim \varepsilon_{\text{dis}}^2$ and $\|\boldsymbol{\rho}^*\|_1 \lesssim C_{\mathcal{F}}$ (see (24) and (21)), as well as the concentration bound for $\frac{1}{N} \sum_{i=1}^N z_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*)$ to be derived in (79). Combine (72) and (73) to derive

$$\begin{aligned} \mathbb{E}[(\phi_{N+1}^\top \hat{\boldsymbol{\rho}} - f(\mathbf{x}_{N+1}))^2] &\lesssim \sqrt{\frac{\log N}{N}} (\varepsilon_{\text{dis}}^2 + C_{\mathcal{F}}^2 + \lambda^{-2} \varepsilon_{\text{opt}}^2 + \sigma(\varepsilon_{\text{dis}} + C_{\mathcal{F}} + \lambda^{-1} \varepsilon_{\text{opt}})) + \varepsilon_{\text{dis}}^2 + \lambda C_{\mathcal{F}} + \varepsilon_{\text{opt}} \\ &\asymp \sqrt{\frac{\log N}{N}} (C_{\mathcal{F}}^2 + \lambda^{-2} \varepsilon_{\text{opt}}^2 + \sigma \varepsilon_{\text{dis}}) + \varepsilon_{\text{dis}}^2 + \lambda C_{\mathcal{F}} + \varepsilon_{\text{opt}}, \end{aligned}$$

where the last line follows since $\sqrt{\frac{\log N}{N}} \sigma C_{\mathcal{F}} \lesssim \lambda C_{\mathcal{F}}$ and $\sqrt{\frac{\log N}{N}} \sigma \lambda^{-1} \varepsilon_{\text{opt}} \lesssim \varepsilon_{\text{opt}}$ (see (70)).

A.2.1 Proof of Lemma 6

We begin by establishing (66). Given that $y_i = f(\mathbf{x}_i) + z_i$, we make note of the following decomposition:

$$\begin{aligned} &\left| \frac{1}{N} \sum_{i=1}^N [(y_i - \phi_i^\top \boldsymbol{\rho})^2 - z_i^2] - \mathbb{E}[(f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho})^2] \right| \\ &= \left| \frac{1}{N} \sum_{i=1}^N \left\{ (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho})^2 - \mathbb{E}[(f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho})^2] + 2z_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}) \right\} \right| \\ &\leq \left| \frac{1}{N} \sum_{i=1}^N \left\{ (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho})^2 - (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*)^2 \right\} + \mathbb{E}[(f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho}^*)^2] - \mathbb{E}[(f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho})^2] \right| \\ &\quad + \left| \frac{1}{N} \sum_{i=1}^N (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*)^2 - \mathbb{E}[(f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho}^*)^2] \right| + 2 \left| \sum_{i=1}^N z_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}) \right| =: \mathcal{E}_1 + \mathcal{E}_2 + \mathcal{E}_3, \end{aligned} \tag{74}$$

leaving us with three terms to cope with.

- Regarding the term \mathcal{E}_1 , we can apply a little algebra to derive

$$\begin{aligned} \mathcal{E}_1 &\leq \left| \frac{2}{N} \sum_{i=1}^N (\boldsymbol{\rho}^* - \boldsymbol{\rho})^\top \phi_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*) - 2\mathbb{E}[(\boldsymbol{\rho}^* - \boldsymbol{\rho})^\top \phi_{N+1} (f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho}^*)] \right| \\ &\quad + \left| \frac{1}{N} \sum_{i=1}^N ((\boldsymbol{\rho}^* - \boldsymbol{\rho})^\top \phi_i)^2 - \mathbb{E}[(\boldsymbol{\rho}^* - \boldsymbol{\rho})^\top \phi_{N+1}]^2 \right| \\ &\leq 2 \|\boldsymbol{\rho}^* - \boldsymbol{\rho}\|_1 \left\| \frac{1}{N} \sum_{i=1}^N \phi_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*) - \mathbb{E}[\phi_{N+1} (f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho}^*)] \right\|_\infty \\ &\quad + \|\boldsymbol{\rho}^* - \boldsymbol{\rho}\|_1^2 \left\| \frac{1}{N} \sum_{i=1}^N \phi_i \phi_i^\top - \mathbb{E}[\phi_{N+1} \phi_{N+1}^\top] \right\|_\infty, \end{aligned} \tag{75}$$

where the first line arises from the elementary inequality $(a-b)^2 - (a-c)^2 = -(c-b)^2 + 2(c-b)(a-b)$ in conjunction with the triangle inequality. Here, for any matrix \mathbf{A} we denote by $\|\mathbf{A}\|_\infty = \max_{i,j} |A_{i,j}|$ the entrywise ℓ_∞ norm. Recognizing that $\|\phi_i\|_\infty \leq 1$ (see (26) and (22), as well as $\phi(z) \in [-1, 1]$ as in (19)), we obtain

$$\begin{aligned} \|\phi_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*)\|_\infty &\leq \|\phi_i\|_\infty \sup_{\mathbf{x} \in \mathcal{B}} |f(\mathbf{x}) - \tilde{\phi}(\mathbf{x})^\top \boldsymbol{\rho}^*| \stackrel{(a)}{\lesssim} \varepsilon_{\text{dis}}, \\ \|\phi_i \phi_i^\top\|_\infty &\leq \|\phi_i\|_\infty^2 \leq 1 \end{aligned}$$

with $\tilde{\phi}(\mathbf{x}) = [\phi_1^{\text{feature}}(\mathbf{x}), \dots, \phi_n^{\text{feature}}(\mathbf{x}), 1]^\top$ and ε_{dis} defined in (23). Here, (a) follows since

$$\sup_{\mathbf{x} \in \mathcal{B}} |f(\mathbf{x}) - \tilde{\phi}(\mathbf{x})^\top \boldsymbol{\rho}^*| \leq \varepsilon_{\text{dis}} + |f(\mathbf{0}) - \rho_{f,0}^*| = \varepsilon_{\text{dis}}, \quad (76)$$

which holds due to (24) as well as our choice $\rho_{f,0}^* = f(\mathbf{0})$ (see (28)). Then applying Hoeffding's inequality shows that, with probability at least $1 - O(N^{-12})$,

$$\begin{aligned} \left\| \frac{1}{N} \sum_{i=1}^N \phi_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*) - \mathbb{E}[\phi_{N+1} (f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho}^*)] \right\|_\infty &\lesssim \sqrt{\frac{\log N}{N}} \varepsilon_{\text{dis}}, \\ \left\| \frac{1}{N} \sum_{i=1}^N \phi_i \phi_i^\top - \mathbb{E}[\phi_{N+1} \phi_{N+1}^\top] \right\|_\infty &\lesssim \sqrt{\frac{\log N}{N}}. \end{aligned}$$

Substitution into (75) yields

$$\mathcal{E}_1 \lesssim \sqrt{\frac{\log N}{N}} \|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1 (\|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1 + \varepsilon_{\text{dis}}) \quad (77)$$

with probability exceeding $1 - O(N^{-12})$.

- With regards to \mathcal{E}_2 , taking the Hoeffding inequality in conjunction with (76) tell us that, with probability exceeding $1 - O(N^{-12})$,

$$\mathcal{E}_2 \lesssim \sqrt{\frac{\log N}{N}} \varepsilon_{\text{dis}}^2. \quad (78)$$

- We now turn to \mathcal{E}_3 . Recalling that $\|\phi_i\|_\infty \leq 1$ and making use of (76), we apply the triangle inequality to obtain

$$\begin{aligned} \mathcal{E}_3 &\leq 2 \left| \frac{1}{N} \sum_{i=1}^N z_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*) \right| + 2 \left| \frac{1}{N} \sum_{i=1}^N z_i \phi_i^\top (\boldsymbol{\rho} - \boldsymbol{\rho}^*) \right| \\ &\leq 2 \left| \frac{1}{N} \sum_{i=1}^N z_i (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*) \right| + 2 \left\| \frac{1}{N} \sum_{i=1}^N z_i \phi_i \right\|_\infty \|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1 \\ &\lesssim \sqrt{\frac{\log N}{N}} \sigma (\varepsilon_{\text{dis}} + \|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1), \end{aligned} \quad (79)$$

where we have used the sub-Gaussian assumption on $\{z_i\}$.

Substituting (77), (78) and (79) into (74) reveals that, with probability at least $1 - O(N^{-12})$,

$$\begin{aligned} &\left| \frac{1}{N} \sum_{i=1}^N [(y_i - \phi_i^\top \boldsymbol{\rho})^2 - z_i^2] - \mathbb{E}[(f(\mathbf{x}_{N+1}) - \phi_{N+1}^\top \boldsymbol{\rho})^2] \right| \\ &\lesssim \sqrt{\frac{\log N}{N}} (\|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1^2 + \|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1 \varepsilon_{\text{dis}} + \varepsilon_{\text{dis}}^2 + \sigma \|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1 + \sigma \varepsilon_{\text{dis}}) \\ &\asymp \sqrt{\frac{\log N}{N}} (\|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1^2 + \varepsilon_{\text{dis}}^2 + \sigma \|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_1 + \sigma \varepsilon_{\text{dis}}). \end{aligned}$$

Next, we turn to proving (67). In view of the Hoeffding inequality, with probability at least $1 - O(N^{-12})$ we have

$$\frac{1}{N} \sum_{i=1}^N [(y_i - \phi_i^\top \boldsymbol{\rho})^2 - (y_i - \phi_i^\top \boldsymbol{\rho}^*)^2] = \frac{1}{N} \sum_{i=1}^N [(f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho})^2 - (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*)^2] + \frac{2}{N} \sum_{i=1}^N z_i \phi_i^\top (\boldsymbol{\rho}^* - \boldsymbol{\rho})$$

$$\begin{aligned}
&\geq -\frac{1}{N} \sum_{i=1}^N (f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*)^2 - \left\| \frac{2}{N} \sum_{i=1}^N z_i \phi_i \right\|_\infty \|\boldsymbol{\rho}^* - \boldsymbol{\rho}\|_1 \\
&\gtrsim -\varepsilon_{\text{dis}}^2 - \sqrt{\frac{\log N}{N}} \sigma \|\boldsymbol{\rho}^* - \boldsymbol{\rho}\|_1.
\end{aligned}$$

where we have used the following facts (already proven previously):

$$|f(\mathbf{x}_i) - \phi_i^\top \boldsymbol{\rho}^*| \lesssim \varepsilon_{\text{dis}}, \quad \left\| \frac{1}{N} \sum_{i=1}^N z_i \phi_i \right\|_\infty \lesssim \sqrt{\frac{\log N}{N}} \sigma.$$

A.3 Proof of Lemma 3

To begin with, we find it convenient to introduce an auxiliary sequence $\boldsymbol{\rho}_{t+1}^*$ obeying $\boldsymbol{\rho}_0^* = \mathbf{0}$ and

$$\boldsymbol{\rho}_{t+1}^* = \text{ST}_{\eta\lambda} \left(\boldsymbol{\rho}_t^{\text{proximal}} + \frac{2\eta}{N} \sum_{i=1}^N (y_i - \phi_i^\top \boldsymbol{\rho}_t^{\text{proximal}}) \phi_i \right),$$

where $\boldsymbol{\rho}_{t+1}^*$ is obtained by running one *exact* proximal gradient iteration from $\boldsymbol{\rho}_t^{\text{proximal}}$. Standard convergence analysis for the proximal gradient method (e.g., Beck (2017)) reveals that

$$\ell(\boldsymbol{\rho}_{t+1}^*) \leq \ell(\boldsymbol{\rho}_t^{\text{proximal}}), \quad (80a)$$

$$\ell(\boldsymbol{\rho}_{t+1}^*) - \ell(\boldsymbol{\rho}^*) \leq n(\|\boldsymbol{\rho}_t^{\text{proximal}} - \boldsymbol{\rho}^*\|_2^2 - \|\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}^*\|_2^2), \quad (80b)$$

where we recall our choice that $\eta = 1/(2n)$. For completeness, we shall provide the proof of (80) towards the end of this subsection.

Recognizing that $\boldsymbol{\rho}_t^{\text{proximal}} = \boldsymbol{\rho}_t^* + \mathbf{e}_t$ for some additive term \mathbf{e}_t , we can invoke (80b) to show that

$$\begin{aligned}
\ell(\boldsymbol{\rho}_{t+1}^*) - \ell(\boldsymbol{\rho}^*) &\leq n(\|\boldsymbol{\rho}_t^{\text{proximal}} - \boldsymbol{\rho}^*\|_2^2 - \|\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}^*\|_2^2) \\
&= n(\|\boldsymbol{\rho}_t^* - \boldsymbol{\rho}^*\|_2^2 - \|\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}^*\|_2^2 + \|\mathbf{e}_t\|_2^2 + 2\mathbf{e}_t^\top (\boldsymbol{\rho}_t^* - \boldsymbol{\rho}^*)) \\
&\leq n(\|\boldsymbol{\rho}_t^* - \boldsymbol{\rho}^*\|_2^2 - \|\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}^*\|_2^2 + \|\mathbf{e}_t\|_1^2 + 2\|\mathbf{e}_t\|_1(O(C_{\mathcal{F}}) + \|\boldsymbol{\rho}_t^*\|_1)) \\
&\leq n(\|\boldsymbol{\rho}_t^* - \boldsymbol{\rho}^*\|_2^2 - \|\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}^*\|_2^2 + \|\mathbf{e}_t\|_1^2 + 2\|\mathbf{e}_t\|_1(O(C_{\mathcal{F}}) + \|\boldsymbol{\rho}_t^{\text{proximal}}\|_1 + \|\mathbf{e}_t\|_1)) \\
&\leq n(\|\boldsymbol{\rho}_t^* - \boldsymbol{\rho}^*\|_2^2 - \|\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}^*\|_2^2 + c_1\|\mathbf{e}_t\|_1(C_{\mathcal{F}} + \|\boldsymbol{\rho}_t^{\text{proximal}}\|_1))
\end{aligned} \quad (81)$$

for some universal constant $c_1 > 0$, where we have used $\|\boldsymbol{\rho}^*\|_1 \lesssim C_{\mathcal{F}}$ (cf. (21) and (28)) and the assumption $\|\mathbf{e}_t\|_1 \lesssim C_{\mathcal{F}}$. Define

$$k_t = \arg \min_{1 \leq k \leq t} \ell(\boldsymbol{\rho}_k^*).$$

Summing (81) over iterations 0 to t , we obtain a telescoping sum and can then deduce that

$$\begin{aligned}
\ell(\boldsymbol{\rho}_{k_t}^*) - \ell(\boldsymbol{\rho}^*) &= \min_{1 \leq k \leq t} \ell(\boldsymbol{\rho}_k^*) - \ell(\boldsymbol{\rho}^*) \\
&\leq \frac{1}{t} \sum_{k=1}^t (\ell(\boldsymbol{\rho}_k^*) - \ell(\boldsymbol{\rho}^*)) \leq \frac{n\|\boldsymbol{\rho}_0^{\text{proximal}} - \boldsymbol{\rho}^*\|_2^2}{t} + c_1 n \max_{1 \leq k < t} \{\|\mathbf{e}_k\|_1(C_{\mathcal{F}} + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1)\} \\
&= \frac{n\|\boldsymbol{\rho}^*\|_1^2}{t} + c_1 n \max_{1 \leq k < t} \{\|\mathbf{e}_k\|_1(C_{\mathcal{F}} + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1)\},
\end{aligned} \quad (82)$$

where the last line follows since $\boldsymbol{\rho}_0^{\text{proximal}} = \mathbf{0}$.

In addition, it is seen that

$$\ell(\boldsymbol{\rho}_{t+1}^{\text{proximal}}) - \ell(\boldsymbol{\rho}_{t+1}^*) = \frac{1}{N} \sum_{i=1}^N (y_i - \phi_i^\top \boldsymbol{\rho}_{t+1}^{\text{proximal}})^2 - \frac{1}{N} \sum_{i=1}^N (y_i - \phi_i^\top \boldsymbol{\rho}_{t+1}^*)^2 + \lambda \|\boldsymbol{\rho}_{t+1}^{\text{proximal}}\|_1 - \lambda \|\boldsymbol{\rho}_{t+1}^*\|_1$$

$$\begin{aligned}
&\leq \frac{2}{N} \sum_{i=1}^N \phi_i^\top (\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}_{t+1}^{\text{proximal}}) (y_i - \phi_i^\top \boldsymbol{\rho}_{t+1}^{\text{proximal}}) + \lambda \|\boldsymbol{\rho}_{t+1}^{\text{proximal}} - \boldsymbol{\rho}_{t+1}^*\|_1 \\
&\leq \|\boldsymbol{\rho}_{t+1}^* - \boldsymbol{\rho}_{t+1}^{\text{proximal}}\|_1 \left(\left\| \frac{2}{N} \sum_{i=1}^N (y_i - \phi_i^\top \boldsymbol{\rho}_{t+1}^{\text{proximal}}) \phi_i \right\|_\infty + \lambda \right) \\
&\stackrel{(a)}{\leq} c_2 \|\mathbf{e}_{t+1}\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_{t+1}^{\text{proximal}}\|_1 + \lambda)
\end{aligned} \tag{83}$$

for some universal constant $c_2 > 0$, where the first inequality comes from the elementary inequality $(a - b)^2 - (a - c)^2 = -(c - b)^2 + (c - b)(2a - 2b) \leq 2(c - b)(a - b)$ as well as the triangle inequality. Here, (a) follows since

$$\begin{aligned}
\frac{1}{N} \sum_{i=1}^N |y_i| &\leq \frac{1}{N} \sum_{i=1}^N |f(\mathbf{x}_i)| + \frac{1}{N} \sum_{i=1}^N |z_i| \leq \max_{1 \leq i \leq N} |\phi_i^\top \boldsymbol{\rho}^*| + \varepsilon_{\text{dis}} + \left| \frac{1}{N} \sum_{i=1}^N |z_i| - \mathbb{E}[|z|] \right| + \mathbb{E}[|z|] \\
&\lesssim \|\boldsymbol{\rho}^*\|_1 \max_{1 \leq i \leq N} \|\phi_i\|_\infty + \varepsilon_{\text{dis}} + \sigma \lesssim \|\boldsymbol{\rho}^*\|_1 + \varepsilon_{\text{dis}} + C_{\mathcal{F}} + \sigma \asymp C_{\mathcal{F}} + \sigma,
\end{aligned}$$

a consequence of the sub-Gaussian assumption on $\{z_i\}$ and the facts that $\|\phi_i\|_\infty \leq 1$, $\|\boldsymbol{\rho}^*\|_1 \lesssim C_{\mathcal{F}}$, and $\varepsilon_{\text{dis}} = C_{\mathcal{F}} \left(\sqrt{\varepsilon} + \left(\frac{\log |N_\varepsilon|}{n} \right)^{1/3} \right) \lesssim C_{\mathcal{F}}$ for the assumption that $\varepsilon \lesssim \sqrt{\log N/N} + n/L$ and $n \gtrsim \log |N_\varepsilon|$. Recalling that $\ell(\boldsymbol{\rho}_{t+1}^*) \leq \ell(\boldsymbol{\rho}_t^{\text{proximal}})$ (see (80a)), we can invoke the bound (83) recursively to derive

$$\begin{aligned}
\ell(\boldsymbol{\rho}_{t+1}^{\text{proximal}}) &\leq \ell(\boldsymbol{\rho}_{t+1}^*) + c_2 \|\mathbf{e}_{t+1}\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_{t+1}^{\text{proximal}}\|_1 + \lambda) \\
&\leq \ell(\boldsymbol{\rho}_t^{\text{proximal}}) + c_2 \|\mathbf{e}_{t+1}\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_{t+1}^{\text{proximal}}\|_1 + \lambda)
\end{aligned} \tag{84}$$

$$\begin{aligned}
&\leq \ell(\boldsymbol{\rho}_t^*) + 2c_2 \max_{t \leq k \leq t+1} \{ \|\mathbf{e}_k\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 + \lambda) \} \\
&\leq \ell(\boldsymbol{\rho}_{k_t}^*) + c_2(t+1-k_t) \max_{1 \leq k \leq t+1} \{ \|\mathbf{e}_k\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 + \lambda) \}.
\end{aligned} \tag{85}$$

It then follows from (82) that

$$\begin{aligned}
\ell(\boldsymbol{\rho}_{t+1}^{\text{proximal}}) - \ell(\boldsymbol{\rho}^*) &\leq \frac{n\|\boldsymbol{\rho}^*\|_1^2}{t} + c_1 n \max_{1 \leq k \leq t-1} \{ \|\mathbf{e}_k\|_1 (C_{\mathcal{F}} + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1) \} \\
&\quad + c_2(t+1) \max_{1 \leq k \leq t+1} \{ \|\mathbf{e}_k\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 + \lambda) \} \\
&\leq \frac{n\|\boldsymbol{\rho}^*\|_1^2}{t} + c_3(t+n+1) \max_{1 \leq k \leq t+1} \{ \|\mathbf{e}_k\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 + \lambda) \},
\end{aligned}$$

where $c_3 = \max\{c_1, c_2\}$. Recalling that $\|\boldsymbol{\rho}^*\|_1 \lesssim C_{\mathcal{F}}$ (cf. (21) and (28)) as well as our parameter choice $T = \frac{L-1}{2}$, we have thus completed the proof of (34).

In addition, combining (30) in Lemma 2 with the above result, we know that for $\lambda \gtrsim \sqrt{\frac{\log N}{N}}(C_{\mathcal{F}} + \sigma)$,

$$\begin{aligned}
\|\boldsymbol{\rho}_t^{\text{proximal}}\|_1 &\lesssim C_{\mathcal{F}} + \lambda^{-1} \left(\frac{n\|\boldsymbol{\rho}^*\|_1^2}{t-1} + (t+n) \max_{1 \leq k \leq t} \{ \|\mathbf{e}_k\|_1 (C_{\mathcal{F}} + \sigma + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 + \lambda) \} \right) \\
&\lesssim C_{\mathcal{F}} + \frac{n\|\boldsymbol{\rho}^*\|_1^2}{t\lambda} + \sqrt{N}(t+n) \max_{1 \leq k \leq t} \|\mathbf{e}_k\|_1 + \frac{t+n}{\lambda} \max_{1 \leq k \leq t} \{ \|\mathbf{e}_k\|_1 \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 \}, \quad t \geq 2.
\end{aligned}$$

For $t = 0$ and $t = 1$, we have $\|\boldsymbol{\rho}_0^{\text{proximal}}\|_1 = 0$, and with probability at least $1 - O(N^{-20})$,

$$\begin{aligned}
\|\boldsymbol{\rho}_1^{\text{proximal}}\|_1 &\leq \sum_{k=1}^n \max \left\{ \left\| \frac{2\eta}{N} \sum_{i=1}^N y_i \phi_i \right\|_\infty + |e_{1,k}| - \lambda\eta, 0 \right\} \\
&\leq \sum_{k=1}^n \max \left\{ \frac{c}{n} \left(C_{\mathcal{F}} + \sqrt{\frac{\log N}{N}} \sigma \right) + |e_{1,k}| - \frac{\lambda}{2n}, 0 \right\} \lesssim C_{\mathcal{F}} + \|\mathbf{e}_1\|_1,
\end{aligned}$$

provided that $\lambda \geq 2c\sqrt{\frac{\log N}{N}}\sigma$. Here, $e_{1,k}$ denotes the k -th element of \mathbf{e}_1 .

Proof of (80a) and (80b). Define

$$\begin{aligned} g(\boldsymbol{\rho}) &= \frac{1}{N} \sum_{i=1}^N (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho})^2, \\ \psi(\boldsymbol{\rho}) &= g(\boldsymbol{\rho}_t^{\text{proximal}}) + (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}})^\top \nabla g(\boldsymbol{\rho}_t^{\text{proximal}}) + n \|\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}}\|_2^2 + \lambda \|\boldsymbol{\rho}\|_1. \end{aligned}$$

Recall that $\eta = 1/(2n)$. It is self-evident that $\boldsymbol{\rho}_{t+1}^* = \arg \min_{\boldsymbol{\rho}} \psi(\boldsymbol{\rho})$ and $\psi(\cdot)$ is $(2n)$ -strongly convex. Thus, for any $\widehat{\boldsymbol{\rho}}$, we have

$$\psi(\widehat{\boldsymbol{\rho}}) \geq \psi(\boldsymbol{\rho}_{t+1}^*) + n \|\boldsymbol{\rho}_{t+1}^* - \widehat{\boldsymbol{\rho}}\|_2^2. \quad (86)$$

In addition, observe that

$$\begin{aligned} &g(\boldsymbol{\rho}_t^{\text{proximal}}) + (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}})^\top \nabla g(\boldsymbol{\rho}_t^{\text{proximal}}) + n \|\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}}\|_2^2 \\ &= g(\boldsymbol{\rho}_t^{\text{proximal}}) - \frac{2}{N} \sum_{i=1}^N (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}})^\top (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}) \boldsymbol{\phi}_i + n \|\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}}\|_2^2 \\ &\geq g(\boldsymbol{\rho}_t^{\text{proximal}}) - \frac{2}{N} \sum_{i=1}^N (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}})^\top (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}) \boldsymbol{\phi}_i + \frac{1}{N} \sum_{i=1}^N (\boldsymbol{\phi}_i^\top (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}}))^2 = g(\boldsymbol{\rho}), \end{aligned}$$

where the last inequality applies the fact that $\sum_{i=1}^N \|\boldsymbol{\phi}_i\|_2^2 \leq Nn$. Thus, we can conclude that

$$\psi(\boldsymbol{\rho}_{t+1}^*) \geq g(\boldsymbol{\rho}_{t+1}^*) + \lambda \|\boldsymbol{\rho}_{t+1}^*\|_1 = \ell(\boldsymbol{\rho}_{t+1}^*). \quad (87)$$

Similarly, we can also demonstrate that

$$\begin{aligned} g(\boldsymbol{\rho}_t^{\text{proximal}}) + (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}})^\top \nabla g(\boldsymbol{\rho}_t^{\text{proximal}}) &= g(\boldsymbol{\rho}_t^{\text{proximal}}) - \frac{2}{N} \sum_{i=1}^N (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}})^\top (y_i - \boldsymbol{\phi}_i^\top \boldsymbol{\rho}) \boldsymbol{\phi}_i \\ &= g(\boldsymbol{\rho}) - \frac{1}{N} \sum_{i=1}^N (\boldsymbol{\phi}_i^\top (\boldsymbol{\rho} - \boldsymbol{\rho}_t^{\text{proximal}}))^2 \leq g(\boldsymbol{\rho}), \end{aligned}$$

and as a consequence,

$$\psi(\widehat{\boldsymbol{\rho}}) \leq g(\widehat{\boldsymbol{\rho}}) + n \|\widehat{\boldsymbol{\rho}} - \boldsymbol{\rho}_t^{\text{proximal}}\|_2^2 + \lambda \|\widehat{\boldsymbol{\rho}}\|_1 = \ell(\widehat{\boldsymbol{\rho}}) + n \|\widehat{\boldsymbol{\rho}} - \boldsymbol{\rho}_t^{\text{proximal}}\|_2^2. \quad (88)$$

Substituting (87) and (88) into (86), we see that for any $\widehat{\boldsymbol{\rho}}$,

$$\ell(\widehat{\boldsymbol{\rho}}) \geq \ell(\boldsymbol{\rho}_{t+1}^*) + n (\|\widehat{\boldsymbol{\rho}} - \boldsymbol{\rho}_{t+1}^*\|_2^2 - \|\widehat{\boldsymbol{\rho}} - \boldsymbol{\rho}_t^{\text{proximal}}\|_2^2).$$

Taking $\widehat{\boldsymbol{\rho}} = \boldsymbol{\rho}_t^{\text{proximal}}$ yields

$$\ell(\boldsymbol{\rho}_t^{\text{proximal}}) \geq \ell(\boldsymbol{\rho}_{t+1}^*) + n \|\boldsymbol{\rho}_t^{\text{proximal}} - \boldsymbol{\rho}_{t+1}^*\|_2^2 \geq \ell(\boldsymbol{\rho}_{t+1}^*),$$

which completes the proof of (80a). In addition, taking $\widehat{\boldsymbol{\rho}} = \boldsymbol{\rho}^*$ gives

$$\ell(\boldsymbol{\rho}^*) \geq \ell(\boldsymbol{\rho}_{t+1}^*) + n (\|\boldsymbol{\rho}^* - \boldsymbol{\rho}_{t+1}^*\|_2^2 - \|\boldsymbol{\rho}^* - \boldsymbol{\rho}_t^{\text{proximal}}\|_2^2),$$

thus concluding the proof of (80b).

A.4 Proof of Lemma 4

In this proof, we first show how to construct the desirable transformer, followed by an analysis of this construction. In particular, we would like to show that the transformer as illustrated in Figure 1, with parameters specified below, satisfy properties i) and ii) in Lemma 4. For ease of presentation, we denote

$$\mathbf{H}^{(l)} = \text{FF}_{\Theta_{\text{ff}}^{(l)}}(\mathbf{H}^{(l-1/2)}), \quad \mathbf{H}^{(l-1/2)} = \text{Attn}_{\Theta_{\text{attn}}^{(l)}}(\mathbf{H}^{(l-1)}), \quad l = 1, \dots, L. \quad (89)$$

Throughout this subsection, all components in $\mathbf{H}^{(l-1/2)}$ share the same superscript $l-1/2$ (e.g., $\phi_j^{(l-1/2)}$, $\lambda_j^{(l-1/2)}$), which help distinguish between different layers.




Figure 1: Structure of the desirable transformer.

A.4.1 High-level structure

The overall structure of the transformer is depicted in Figure 1.

- The architecture begins with an attention layer **Attn0** and a feedforward layer **FF0**, which serve to initialize certain variables (particularly the features identified in Lemma 1) based on the in-context inputs $\{x_i\}$.
- The remaining $(L - 1)$ layers are divided into $(L - 1)/2$ blocks with identical structure and parameters. More concretely, each block consists of two attention layers and two feed-forward layers, (**Attn1**, **FF1**) and (**Attn2**, **FF2**), which are designed to perform inexact proximal gradient iterations and update the corresponding prediction.

In the sequel, we shall first describe what update rule each layer is designed to implement, followed by detailed explanation about how they can be realized using the transformer architecture.

A.4.2 Intended updates for each layer

Our transformer construction comprises multiple layers (as illustrated in Figure 1) designed to emulate the iterations of the inexact proximal gradient method (31b). Let us begin by describing the desired update to be performed at each layer, abstracting away the specifics of the transformer implementation.

- **FF0**: This feed-forward layer intends to update the components ϕ_j and λ as follows:

$$\phi_j^{(l)} = \phi_j^{(l-1/2)} + [\phi_1^{\text{feature}}(x_j^{(l-1/2)}), \dots, \phi_n^{\text{feature}}(x_j^{(l-1/2)}), 1]^\top, \quad 1 \leq j \leq N + 1; \quad (90a)$$

$$\lambda^{(l)} = c_1 \left(\frac{\log N}{N} \right)^{1/6} C_{\mathcal{F}}^{-1/3} \hat{\varepsilon}^{2/3} + c_1 \sqrt{\frac{\log N}{N}} (C_{\mathcal{F}} + \sigma) + c_1 C_{\mathcal{F}}^{-1} \varepsilon_{\text{dis}}^2 =: \bar{\lambda}, \quad (90b)$$

where $\phi_i^{\text{feature}}(x)$ is defined in (19), $\hat{\varepsilon}$ is defined in Lemma 4 (see (39c)), and $c_1 > 0$ is some large enough universal constant.

- **Attn1**: In this attention layer, we attempt to implement the following updates:

$$\rho^{(l-1/2)} = \rho^{(l-1)} + \frac{2\eta}{N} \sum_{i=1}^N \phi_i^{(l-1)} \left\{ y_i^{(l-1)} - (\phi_i^{(l-1)})^\top \rho^{(l-1)} \right\} + e^{(l-1)} \quad (91)$$

for some residual (or error) term $e^{(l-1)}$, corresponding to an iteration of gradient descent in (31a) before the proximal operator is applied. This residual term $e^{(l-1)}$ shall be bounded shortly.

- **FF1**: This feed-forward layer is designed to implement the following updates:

$$\rho^{(l)} = \text{ST}_{\eta\lambda^{(l-1/2)}}(\rho^{(l-1/2)}), \quad (92a)$$

$$\hat{y}^{(l)} = 0, \quad (92b)$$

which applies the proximal operator (i.e., soft-thresholding) to the output in (91). As a result, this in conjunction with (91) in **Attn1** completes one (inexact) proximal gradient iteration (31a).

- **Attn2**: This attention layer intends to update the prediction \hat{y} based on $\boldsymbol{\rho}^{(l)}$, namely,

$$\hat{y}^{(l+1/2)} = (\boldsymbol{\phi}_{N+1}^{(l)})^\top \boldsymbol{\rho}^{(l)} + \tilde{e}^{(l)}, \quad (93)$$

where $\tilde{e}^{(l)}$ is some residual term that will be bounded momentarily.

- **Attn0, FF2**: These layers do not update the hidden representation $\mathbf{H}^{(l)}$; instead, they are included to ensure consistency with the transformer architecture defined in (10).

On a high level, this transformer implements the inexact proximal gradient method in (31b). In particular, after passing the input through **Attn0** and **FF0**, we obtain

$$\boldsymbol{\phi}_j^{(1)} = \boldsymbol{\phi}_j, \quad \lambda^{(1)} = \bar{\lambda}, \quad 1 \leq j \leq N+1.$$

The remaining layers then proceed as follows:

- All parameters except $\boldsymbol{\rho}^{(l)}$ and $\hat{y}^{(l)}$ will stay fixed throughout the remaining layers, i.e.,

$$\boldsymbol{\phi}_j^{(l)} = \boldsymbol{\phi}_j, \quad \lambda^{(l)} = \lambda^{(l+1/2)} = \bar{\lambda}, \quad y_j^{(l)} = y_j^{(0)} = y_j, \quad w_j^{(l)} = w_j^{(0)}, \quad \forall 1 \leq l \leq L, \quad 1 \leq j \leq N+1. \quad (94)$$

where $w_j^{(0)} = 1$ for $1 \leq j \leq N$ and $w_{N+1}^{(0)} = 0$, as previously defined in (37b) and (37c).

- The components $\boldsymbol{\rho}^{(l)}$ are updated in a way that resembles (31b), namely,

$$\boldsymbol{\rho}^{(2t+1)} = \boldsymbol{\rho}_t^{\text{proximal}} \quad \text{for } t \geq 1 \quad \text{and} \quad \boldsymbol{\rho}^{(0)} = \boldsymbol{\rho}_0^{\text{proximal}} = \mathbf{0}, \quad \lambda = \bar{\lambda}. \quad (95)$$

- The components $\hat{y}^{(l)}$ are computed to approximate the prediction of $f(\mathbf{x}_{N+1})$, namely,

$$\hat{y}^{(2t)} = 0, \quad \hat{y}^{(2t+1)} \approx \boldsymbol{\phi}_{N+1}^\top \boldsymbol{\rho}_t^{\text{proximal}} \quad \text{for } t \geq 1.$$

A.4.3 Parameter design in our transformer construction

Next, we explain how the transformer architecture can be designed to implement the updates described above for each layer.

- **FF0**: Note that the function $\phi(z)$ (cf. (19)) is intimately connected with the ReLU function $\sigma_{\text{ff}}(\cdot)$ as:

$$\phi(z) = \sigma_{\text{ff}}(z + 1/2) - \sigma_{\text{ff}}(z - 1/2). \quad (96)$$

This allows us to decompose (19) as

$$\phi_i^{\text{feature}}(\mathbf{x}) = \sigma_{\text{ff}}\left(\tau_{\text{ff}}\left(\frac{\boldsymbol{\omega}_i^\top \mathbf{x}}{\|\boldsymbol{\omega}_i\|_2} - t_i\right) + \frac{1}{2}\right) - \sigma_{\text{ff}}\left(\tau_{\text{ff}}\left(\frac{\boldsymbol{\omega}_i^\top \mathbf{x}}{\|\boldsymbol{\omega}_i\|_2} - t_i\right) - \frac{1}{2}\right), \quad 1 \leq i \leq n, \quad (97)$$

where we take

$$\tau_{\text{ff}} = 1/\sqrt{\varepsilon}. \quad (98)$$

Moreover, the last entry in $\boldsymbol{\phi}_j^{(l)}$ can be expressed by $1 = \sigma_{\text{ff}}(1) - \sigma_{\text{ff}}(-1)$. In order for this layer to carry out (90), we take the parameter matrices $\mathbf{W}^{(l)} \in \mathbb{R}^{D \times D}$ and $\mathbf{U}^{(l)} \in \mathbb{R}^{D \times D}$ (see (8)) to satisfy

$$\mathbf{W}_{1:2n+2, 1:d+1}^{(l)} = \begin{bmatrix} \|\boldsymbol{\omega}_1\|_2^{-1} \tau_{\text{ff}} \boldsymbol{\omega}_1^\top & -t_1 \tau_{\text{ff}} + \frac{1}{2} \\ \vdots & \vdots \\ \|\boldsymbol{\omega}_n\|_2^{-1} \tau_{\text{ff}} \boldsymbol{\omega}_n^\top & -t_n \tau_{\text{ff}} + \frac{1}{2} \\ \mathbf{0}^\top & 1 \\ \|\boldsymbol{\omega}_1\|_2^{-1} \tau_{\text{ff}} \boldsymbol{\omega}_1^\top & -t_1 \tau_{\text{ff}} - \frac{1}{2} \\ \vdots & \vdots \\ \|\boldsymbol{\omega}_n\|_2^{-1} \tau_{\text{ff}} \boldsymbol{\omega}_n^\top & -t_n \tau_{\text{ff}} - \frac{1}{2} \\ \mathbf{0}^\top & -1 \end{bmatrix}, \quad \mathbf{W}_{2n+3,:}^{(l)} = \bar{\lambda} \mathbf{u}_{d+1}^\top, \quad (99a)$$

$$\mathbf{U}_{d+4:d+n+4, 1:2n+2}^{(l)} = [\mathbf{I}_{n+1}, -\mathbf{I}_{n+1}], \quad \mathbf{U}_{d+2n+6,:}^{(l)} = \mathbf{u}_{2n+3}^\top, \quad (99b)$$

with all remaining entries set to zero. Here, $\bar{\lambda}$ has been defined in (90b), and $\mathbf{W}_{i:j, r:h}$ denotes a submatrix of \mathbf{W} consisting of rows i through j and columns r through h , $\mathbf{W}_{i,:}$ represents the i -th row of \mathbf{W} , whereas $\mathbf{u}_i \in \mathbb{R}^D$ stands for the i -th standard basis vector.

- **Attn1:** First, we discuss how the sigmoid function $\sigma_{\text{attn}}(\cdot)$ defined in (9) can help us implement (91). Note that $\sigma'_{\text{attn}}(0) = 1/4$ and observe the Taylor expansion $\sigma_{\text{attn}}(\tau^{-1}x) = \sigma_{\text{attn}}(0) + \frac{x}{4\tau} + O(\frac{x^2}{\tau^2})$, which allow us to express

$$x \approx 4\tau(\sigma_{\text{attn}}(\tau^{-1}x) - \sigma_{\text{attn}}(0)). \quad (100)$$

In light of this observation, we propose to carry out (91) via the following updates that exploit the sigmoid function:

$$\begin{aligned} \boldsymbol{\rho}^{(l-1/2)} &= \boldsymbol{\rho}^{(l-1)} + \frac{2}{N} \sum_{i=1}^{N+1} 4\tau \left(\sigma_{\text{attn}}(\tau^{-1}\eta y_i^{(l-1)}) - \sigma_{\text{attn}}(0) \right) \boldsymbol{\phi}_i^{(l-1)} \\ &\quad - \frac{2}{N} \sum_{i=1}^{N+1} 4\tau \left(\sigma_{\text{attn}} \left(\tau^{-1}\eta (\boldsymbol{\phi}_i^{(l-1)})^\top \boldsymbol{\rho}^{(l-1)} \right) - \sigma_{\text{attn}}(0) \right) \boldsymbol{\phi}_i^{(l-1)} \\ &\quad + \frac{2}{N} \sum_{i=1}^{N+1} 4\tau \left(\sigma_{\text{attn}}(\tau^{-1}\eta(1 - w_i^{(l-1)})\hat{y}^{(l-1)}) - \sigma_{\text{attn}}(0) \right) \boldsymbol{\phi}_i^{(l-1)}, \end{aligned} \quad (101)$$

where the error vector $\mathbf{e}^{(l-1)}$ can be straightforwardly determined by checking the difference between (91) and (101). Next, to fully realize (101) via a attention layer, we use 4 attention heads in this layer and take the parameter matrices $\mathbf{V}_m^{(l)}$, $\mathbf{Q}_m^{(l)}$, and $\mathbf{K}_m^{(l)} \in \mathbb{R}^{D \times D}$ for $m = 1, 2, 3, 4$ to be:

$$\begin{aligned} \mathbf{V}_{1,(d+n+5:d+2n+5, d+4:d+n+4)}^{(l)} &= 8\tau \mathbf{I}_{n+1}, \quad \mathbf{Q}_{1,(1,:)}^{(l)} = \tau^{-1}\eta \mathbf{u}_{d+2}^\top, \quad \mathbf{K}_{1,(1,:)}^{(l)} = \mathbf{u}_{d+1}^\top, \\ \mathbf{V}_2^{(l)} &= -\mathbf{V}_1^{(l)}, \quad \mathbf{Q}_{2,(1:n+1, d+4:d+n+4)}^{(l)} = \tau^{-1}\eta \mathbf{I}_{n+1}, \quad \mathbf{K}_{2,(1:n+1, d+n+5:d+2n+5)}^{(l)} = \mathbf{I}_{n+1}, \\ \mathbf{V}_3^{(l)} &= \mathbf{V}_1^{(l)}, \quad \mathbf{Q}_{3,(1,:)}^{(l)} = \tau^{-1}\eta(\mathbf{u}_{d+1} - \mathbf{u}_{d+3})^\top, \quad \mathbf{K}_{3,(1,:)}^{(l)} = \mathbf{u}_{d+2n+7}^\top, \\ \mathbf{V}_4^{(l)} &= -\mathbf{V}_1^{(l)}, \quad \mathbf{Q}_4^{(l)} = \mathbf{0}, \quad \mathbf{K}_4^{(l)} = \mathbf{0}, \end{aligned}$$

with all remaining entries set to zero. Here, $\mathbf{V}_{m,(i:j, r:h)}^{(l)}$ denotes a submatrix of $\mathbf{V}_m^{(l)}$ comprising rows i through j and columns r through h , $\mathbf{V}_{m,(i,:)}^{(l)}$ denotes the i -th row of $\mathbf{V}_m^{(l)}$, and $\mathbf{u}_i \in \mathbb{R}^D$ represents the i -th standard basis vector.

Before proceeding, let us bound the residual vector $\mathbf{e}^{(l-1)}$ in (91). Observing that $\sigma'_{\text{attn}}(0) = 1/4$ and $|\sigma''_{\text{attn}}(x)| = |e^x - e^{-x}|/(e^{x/2} + e^{-x/2})^4 < 0.5$ for all x , we can show that

$$\left| x - 4\tau[\sigma_{\text{attn}}(\tau^{-1}x) - \sigma_{\text{attn}}(0)] \right| \leq 0.5 \times 4\tau\tau^{-2}|x|^2 = 2\tau^{-1}|x|^2, \quad \forall x \in \mathbb{R}. \quad (102)$$

Taking this collectively with (101) and the choice that $1 - w_i \neq 0$ only for $i = N + 1$, we arrive at

$$\begin{aligned} \|\mathbf{e}^{(l)}\|_\infty &\leq \frac{4\eta^2}{N\tau} \sum_{i=1}^N (y_i^{(l)})^2 + \frac{4\eta^2}{N\tau} \sum_{i=1}^{N+1} \left((\boldsymbol{\phi}_i^{(l)})^\top \boldsymbol{\rho}^{(l)} \right)^2 \\ &\quad + \frac{4\eta^2 (\hat{y}^{(l)})^2}{N\tau} + \frac{2\eta}{N} \left| \hat{y}^{(l)} - (\boldsymbol{\phi}_{N+1}^{(l)})^\top \boldsymbol{\rho}^{(l)} \right| \\ &\lesssim \frac{\eta^2}{\tau} \left(C_{\mathcal{F}}^2 + \sigma^2 + \|\boldsymbol{\rho}^{(l)}\|_1^2 + (\hat{y}^{(l)})^2 \right) + \frac{2\eta}{N} \left| \hat{y}^{(l)} - (\boldsymbol{\phi}_{N+1}^{(l)})^\top \boldsymbol{\rho}^{(l)} \right|, \end{aligned} \quad (103)$$

where we have used the facts that

$$\begin{aligned} |y_i^{(l)}| &= |y_i| \leq |f(\mathbf{x}_i)| + |z_i| \leq |\boldsymbol{\phi}_i^\top \boldsymbol{\rho}^*| + \varepsilon_{\text{dis}} + |z_1| \lesssim \|\boldsymbol{\rho}^*\|_1 + C_{\mathcal{F}} + \sigma \asymp C_{\mathcal{F}} + \sigma, \\ |(\boldsymbol{\phi}_i^{(l)})^\top \boldsymbol{\rho}^{(l)}| &= |\boldsymbol{\phi}_i^\top \boldsymbol{\rho}^{(l)}| \leq \|\boldsymbol{\phi}_i\|_\infty \|\boldsymbol{\rho}^{(l)}\|_1 \leq \|\boldsymbol{\rho}^{(l)}\|_1. \end{aligned}$$

- **FF1:** In order to implement the soft-thresholding operator using the feed-forward layer, we first need to inspect the connection between the soft-thresholding operator and the ReLU function $\sigma_{\text{ff}}(\cdot)$. Towards this end, observe that

$$\begin{aligned} \text{ST}_{\eta\lambda}(z) &= z + \eta\lambda - (z + \eta\lambda) \mathbf{1}(z + \eta\lambda > 0) + (z - \eta\lambda) \mathbf{1}(z - \eta\lambda > 0) \\ &= z + \sigma_{\text{ff}}(\eta\lambda) - \sigma_{\text{ff}}(z + \eta\lambda) + \sigma_{\text{ff}}(z - \eta\lambda). \end{aligned} \quad (104)$$

Consequently, we propose to design a feed-forward layer capable of implementing the following updates (in an attempt to carry out the proposed update (92)):

$$\begin{aligned} \boldsymbol{\rho}^{(l)} &= \boldsymbol{\rho}^{(l-1/2)} + \sigma_{\text{ff}}(\eta\lambda^{(l-1/2)}) - \sigma_{\text{ff}}(\boldsymbol{\rho}^{(l-1/2)} + \eta\lambda^{(l-1/2)}) + \sigma_{\text{ff}}(\boldsymbol{\rho}^{(l-1/2)} - \eta\lambda^{(l-1/2)}) \\ &= \text{ST}_{\eta\lambda^{(l-1/2)}}(\boldsymbol{\rho}^{(l-1/2)}), \\ \hat{y}^{(l)} &= \hat{y}^{(l-1/2)} - \sigma_{\text{ff}}(\hat{y}^{(l-1/2)}) + \sigma_{\text{ff}}(-\hat{y}^{(l-1/2)}) = 0. \end{aligned}$$

To do so, it suffices to set $\mathbf{W}^{(l)} \in \mathbb{R}^{D \times D}$, $\mathbf{U}^{(l)} \in \mathbb{R}^{D \times D}$ to be

$$\mathbf{W}_{1:2n+5, :}^{(l)} = \begin{bmatrix} \mathbf{0}_{(n+1) \times (d+n+4)} & \mathbf{I}_{n+1} & \eta \mathbf{1} & \mathbf{0} \\ \mathbf{0}_{(n+1) \times (d+n+4)} & \mathbf{I}_{n+1} & -\eta \mathbf{1} & \mathbf{0} \\ & \eta \mathbf{u}_{d+2n+6}^\top & & \\ & \mathbf{u}_{d+2n+7}^\top & & \\ & -\mathbf{u}_{d+2n+7}^\top & & \end{bmatrix}, \quad (105a)$$

$$\mathbf{U}_{d+n+5:d+2n+5, 1:2n+5}^{(l)} = [-\mathbf{I}_{n+1}, \mathbf{I}_{n+1}, \mathbf{1}_{n+1}, \mathbf{0}, \mathbf{0}], \quad \mathbf{U}_{D, 1:2n+5}^{(l)} = [\mathbf{0}_{2n+3}^\top, -1, 1], \quad (105b)$$

with all remaining entries set to zero.

- **Attn2:** Recall that (100) tells us that

$$(\phi_i^{(l)})^\top \boldsymbol{\rho}^{(l)} \approx 4\tau \left\{ \sigma_{\text{attn}} \left(\tau^{-1} (\phi_i^{(l)})^\top \boldsymbol{\rho}^{(l)} \right) - \sigma_{\text{attn}}(0) \right\}. \quad (106)$$

As a result, we would like to design this attention layer to actually implement

$$\hat{y}^{(l+1/2)} = \hat{y}^{(l)} + 4\tau \sum_{i=1}^{N+1} (1 - w_i^{(l)}) \left[\sigma_{\text{attn}} \left(\tau^{-1} (\phi_i^{(l)})^\top \boldsymbol{\rho}^{(l)} \right) - \sigma_{\text{attn}}(0) \right], \quad (107)$$

and hence the residual term $\tilde{e}^{(l)}$ in (93) can be easily determined by comparing (93) with (107).

To realize (107), it suffices to use 2 attention heads, and set $\mathbf{V}_1^{(l+1)}, \mathbf{V}_2^{(l+1)} \in \mathbb{R}^{D \times D}$, $\mathbf{Q}_1^{(l+1)}, \mathbf{Q}_2^{(l+1)} \in \mathbb{R}^{D \times D}$, and $\mathbf{K}_1^{(l+1)}, \mathbf{K}_2^{(l+1)} \in \mathbb{R}^{D \times D}$ to be:

$$\begin{aligned} \mathbf{V}_{1, (D, :)}^{(l+1)} &= 4\tau(\mathbf{u}_{d+1} - \mathbf{u}_{d+3})^\top, \quad \mathbf{Q}_{1, (1:n+1, d+4:d+n+4)}^{(l+1)} = \tau^{-1} \mathbf{I}_{n+1}, \quad \mathbf{K}_{1, (1:n+1, d+n+5:d+2n+5)}^{(l+1)} = \mathbf{I}_{n+1}, \\ \mathbf{V}_2^{(l+1)} &= -\mathbf{V}_1^{(l+1)}, \quad \mathbf{Q}_2^{(l+1)} = \mathbf{0}, \quad \mathbf{K}_2^{(l+1)} = \mathbf{0}, \end{aligned}$$

with all remaining entries set to zero. Here, we recall that \mathbf{u}_i denotes the i -th standard basis vector.

Before moving on, let us single out some bound on $\tilde{e}^{(l)}$. Recalling that $1 - w_i^{(l)} = 0$ for all $1 \leq i \leq N$ and making use of (106) and (102), we can demonstrate that

$$|\tilde{e}^{(l)}| \leq \frac{2((\phi_{N+1}^{(l)})^\top \boldsymbol{\rho}^{(l)})^2}{\tau}. \quad (108)$$

- **Attn0, FF2:** The parameter matrices in these layers are taken to be

$$\mathbf{U}^{(l+1)} = \mathbf{W}^{(l+1)} = \mathbf{Q}^{(1)} = \mathbf{K}^{(1)} = \mathbf{V}^{(1)} = \mathbf{0}, \quad (109)$$

so as to ensure that

$$\mathbf{H}^{(1/2)} = \mathbf{H}^{(0)} \quad \text{and} \quad \mathbf{H}^{(l+1)} = \mathbf{H}^{(l+1/2)}.$$

A.4.4 Verifying (95) and controlling the size of e_t

We now verify (95) by induction, and establish upper bounds on e_t in (31b).

First, it is self-evident that $\boldsymbol{\rho}^{(1)} = \boldsymbol{\rho}_0^{\text{proximal}} = \mathbf{0}$. Now, let us assume that $\boldsymbol{\rho}^{(2t-1)} = \boldsymbol{\rho}_{t-1}^{\text{proximal}}$, and proceed to prove $\boldsymbol{\rho}^{(2t+1)} = \boldsymbol{\rho}_t^{\text{proximal}}$. In the t -th block (which contains **Attn1**, **FF1**, **Attn2**, **FF2**), we have

$$\begin{aligned}\boldsymbol{\rho}^{(2t+1)} &= \boldsymbol{\rho}^{(2t)} = \text{ST}_{\eta\bar{\lambda}}(\boldsymbol{\rho}^{(2t-1/2)}) = \text{ST}_{\eta\bar{\lambda}}\left(\boldsymbol{\rho}^{(2t-1)} + \frac{2\eta}{N} \sum_{i=1}^N \phi_i \{y_i - \phi_i^\top \boldsymbol{\rho}^{(2t-1)}\} + \mathbf{e}^{(2t-1)}\right) \\ &= \text{ST}_{\eta\bar{\lambda}}\left(\boldsymbol{\rho}_{t-1}^{\text{proximal}} + \frac{2\eta}{N} \sum_{i=1}^N \phi_i \{y_i - \phi_i^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}}\}\right) + \mathbf{e}_t = \boldsymbol{\rho}_t^{\text{proximal}},\end{aligned}$$

where

$$\mathbf{e}_t := \text{ST}_{\eta\bar{\lambda}}\left(\boldsymbol{\rho}_{t-1}^{\text{proximal}} + \frac{2\eta}{N} \sum_{i=1}^N \phi_i \{y_i - \phi_i^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}}\} + \mathbf{e}^{(2t-1)}\right) - \text{ST}_{\eta\bar{\lambda}}\left(\boldsymbol{\rho}_{t-1}^{\text{proximal}} + \frac{2\eta}{N} \sum_{i=1}^N \phi_i \{y_i - \phi_i^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}}\}\right).$$

Recalling that $\text{ST}(\cdot)$ is a contraction operator and using (103), we can show that

$$\begin{aligned}\|\mathbf{e}_t\|_\infty &\leq \|\mathbf{e}^{(2t-1)}\|_\infty \lesssim \frac{\eta^2}{\tau} \left(C_{\mathcal{F}}^2 + \sigma^2 + \|\boldsymbol{\rho}^{(2t-1)}\|_1^2 + (\hat{y}^{(2t-1)})^2\right) + \frac{\eta}{N} \left|\hat{y}^{(2t-1)} - \phi_{N+1}^\top \boldsymbol{\rho}^{(2t-1)}\right| \\ &\lesssim \frac{\eta^2}{\tau} \left(C_{\mathcal{F}}^2 + \sigma^2 + \|\boldsymbol{\rho}_{t-1}^{\text{proximal}}\|_1^2 + (\hat{y}^{(2t-1)})^2\right) + \frac{\eta}{N} \left|\hat{y}^{(2t-1)} - \phi_{N+1}^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}}\right|,\end{aligned}\tag{110}$$

which is valid since $\boldsymbol{\rho}^{(2t-1)} = \boldsymbol{\rho}_{t-1}^{\text{proximal}}$. Also, it follows from (93) that

$$\hat{y}^{(2t-1)} = \hat{y}^{(2t-3/2)} = \phi_{N+1}^\top \boldsymbol{\rho}^{(2t-2)} + \tilde{e}^{(2t-2)} = \phi_{N+1}^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}} + \tilde{e}^{(2t-2)},$$

where we have used $\boldsymbol{\rho}^{(2t-2)} = \boldsymbol{\rho}^{(2t-1)} = \boldsymbol{\rho}_{t-1}^{\text{proximal}}$. Moreover, it is seen from (108) that

$$|\tilde{e}^{(2t-2)}| \leq \frac{2(\phi_{N+1}^\top \boldsymbol{\rho}^{(2t-2)})^2}{\tau} = \frac{2(\phi_{N+1}^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}})^2}{\tau}.$$

Substituting these into (110) and using $\|\phi_{N+1}\|_\infty \leq 1$ then yield

$$\begin{aligned}\|\mathbf{e}_t\|_\infty &\lesssim \frac{\eta^2}{\tau} \left(C_{\mathcal{F}}^2 + \sigma^2 + \|\boldsymbol{\rho}_{t-1}^{\text{proximal}}\|_1^2 + (\phi_{N+1}^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}})^2 + (\tilde{e}^{(2t-2)})^2\right) + \frac{\eta}{N} |\tilde{e}^{(2t-2)}| \\ &\lesssim \frac{\eta^2}{\tau} \left(C_{\mathcal{F}}^2 + \sigma^2 + \|\boldsymbol{\rho}_t^{\text{proximal}}\|_1^2 + (\phi_{N+1}^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}})^2 + \frac{(\phi_{N+1}^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}})^4}{\tau^2}\right) + \frac{\eta(\phi_{N+1}^\top \boldsymbol{\rho}_{t-1}^{\text{proximal}})^2}{N\tau} \\ &\lesssim \frac{\eta^2}{\tau} \left(C_{\mathcal{F}}^2 + \sigma^2 + \|\boldsymbol{\rho}_{t-1}^{\text{proximal}}\|_1^2 + \frac{\|\boldsymbol{\rho}_{t-1}^{\text{proximal}}\|_1^4}{\tau^2}\right) + \frac{\eta\|\boldsymbol{\rho}_{t-1}^{\text{proximal}}\|_1^2}{N\tau}.\end{aligned}\tag{111}$$

A.4.5 Proof of property iii) in Lemma 4

Equipped with the above transformer parameters, we are now positioned to establish property iii) in Lemma 4. To this end, we first establish the following lemma, whose proof is postponed to Appendix A.4.6.

Lemma 7. Suppose that $\lambda \geq C\sqrt{\log N/N}(C_{\mathcal{F}} + \sigma) + C_{\mathcal{F}}^{-1}\epsilon_{\text{dis}}^2$, and take τ to be sufficiently large such that

$$\tau \geq CNn^2(L+n)(N+n)C_{\mathcal{F}},$$

with $C > 0$ some large enough constant. Then for all $t \geq 0$, it holds that

$$\|\mathbf{e}_t\|_1 \lesssim \frac{C_{\mathcal{F}}}{(L+n)nN} \quad \text{and} \quad \|\boldsymbol{\rho}_t^{\text{proximal}}\|_1 \lesssim n\sqrt{N}C_{\mathcal{F}}.$$

When $L = 2T + 1$ is an odd number, it holds that $\boldsymbol{\rho}^{(L)} = \boldsymbol{\rho}_T^{\text{proximal}}$ and $\hat{y}^{(L)} = \hat{y}^{(L-1/2)}$. Combining Lemma 7 with (108) then reveals that: by taking

$$\tau \geq C N n^2 (L + n) (N + n) C_{\mathcal{F}} \geq C n^2 N^{5/4} C_{\mathcal{F}}$$

for some large enough constant $C > 0$, we have

$$\begin{aligned} |\boldsymbol{\phi}_{N+1}^\top \boldsymbol{\rho}^{(L)} - \hat{y}^{(L)}| &= |\boldsymbol{\phi}_{N+1}^\top \boldsymbol{\rho}^{(L-1)} - \hat{y}^{(L-1/2)}| = |\tilde{e}^{(L-1)}| \leq \frac{2(\boldsymbol{\phi}_{N+1}^\top \boldsymbol{\rho}_{\frac{L-1}{2}}^{\text{proximal}})^2}{\tau} \\ &\leq \frac{2\|\boldsymbol{\phi}_{N+1}\|_\infty^2 \|\boldsymbol{\rho}_{\frac{L-1}{2}}^{\text{proximal}}\|_1^2}{\tau} \leq \left(\frac{\log N}{N}\right)^{1/4} C_{\mathcal{F}}. \end{aligned}$$

Additionally, taking Lemma 7 together with (34) leads to

$$\begin{aligned} \ell(\boldsymbol{\rho}_T^{\text{proximal}}) - \ell(\boldsymbol{\rho}^*) &\lesssim \frac{nC_{\mathcal{F}}^2}{L} + (L + n) \frac{C_{\mathcal{F}}}{(L + n)nN} (C_{\mathcal{F}} + \sigma + n\sqrt{N}C_{\mathcal{F}} + \lambda) \\ &\lesssim \frac{nC_{\mathcal{F}}^2}{L} + \frac{1}{\sqrt{N}} C_{\mathcal{F}} (C_{\mathcal{F}} + \sigma) + \lambda C_{\mathcal{F}}. \end{aligned} \quad (112)$$

Recalling the bound on $\lambda C_{\mathcal{F}}$ in (43) and the definition of $\hat{\varepsilon}$ in (39c), we have

$$\lambda C_{\mathcal{F}} \lesssim \sqrt{\frac{\log N}{N}} C_{\mathcal{F}} (C_{\mathcal{F}} + \sigma) + \varepsilon_{\text{dis}}^2 + \hat{\varepsilon} \lesssim \sqrt{\frac{\log N}{N}} C_{\mathcal{F}} (C_{\mathcal{F}} + \sigma) + \varepsilon_{\text{dis}}^2 + \frac{nC_{\mathcal{F}}^2}{L},$$

which combined with (112) completes the proof.

A.4.6 Proof of Lemma 7

We intend to prove the following inequalities by induction:

$$\|\mathbf{e}_t\|_1 \leq \frac{c_1 C_{\mathcal{F}}}{(L + n)Nn}, \quad (113)$$

$$\|\boldsymbol{\rho}_t^{\text{proximal}}\|_1 \leq c_2 n \sqrt{N} C_{\mathcal{F}}, \quad (114)$$

for some sufficiently small (resp. large) constant $c_1 > 0$ (resp. $c_2 > 0$).

For the base case, we have

$$\|\mathbf{e}_0\|_1 = \|\boldsymbol{\rho}_0^{\text{proximal}} - \boldsymbol{\rho}_0^*\|_1 = 0,$$

and thus (113) holds for $t = 0$. Next, we intend to prove that (114) holds for $t = k$ under the assumption that (113) holds for $t \leq k$ and (114) holds for $t \leq k - 1$. According to (35), it holds that

$$\begin{aligned} \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 &\leq c_5 C_{\mathcal{F}} + \frac{c_5 n C_{\mathcal{F}}^2}{k\lambda} + c_5 \sqrt{N} (k + n) \max_{1 \leq i \leq k} \|\mathbf{e}_i\|_1 + \frac{c_5 (k + n)}{\lambda} \max_{1 \leq i \leq k} \{\|\mathbf{e}_i\|_1 \|\boldsymbol{\rho}_i^{\text{proximal}}\|_1\} \\ &\stackrel{(a)}{\leq} c_5 C_{\mathcal{F}} + \frac{c_5 n \sqrt{N} C_{\mathcal{F}}}{c_\lambda} + c_5 c_1 C_{\mathcal{F}} + \frac{c_5 c_1 c_2 C_{\mathcal{F}}}{c_\lambda} + \frac{c_5 c_1}{c_\lambda \sqrt{N} n} \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1, \end{aligned}$$

where (a) results from the fact that $\lambda \geq c_\lambda C_{\mathcal{F}} / \sqrt{N}$. For some sufficiently small constant $c_1 > 0$ obeying $c_1 c_5 / c_\lambda \leq 1/4$, we have

$$\|\boldsymbol{\rho}_k^{\text{proximal}}\|_1 \leq \frac{4c_5 C_{\mathcal{F}}}{3} \left(1 + \frac{n\sqrt{N}}{c_\lambda} + c_1\right) + \frac{c_2 C_{\mathcal{F}}}{3} \leq c_2 n \sqrt{N} C_{\mathcal{F}},$$

provided that $c_2 \geq 2c_5 (1 + c_\lambda^{-1} + c_1)$.

Next, we intend to establish that (113) holds for $t = k + 1$. By virtue of (111), one has

$$\begin{aligned}\|\mathbf{e}_{k+1}\|_1 &\leq \frac{c_3 n \eta^2}{\tau} \left(C_{\mathcal{F}}^2 + \sigma^2 + \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1^2 + \frac{\|\boldsymbol{\rho}_k^{\text{proximal}}\|_1^4}{\tau^2} \right) + \frac{c_3 n \eta \|\boldsymbol{\rho}_k^{\text{proximal}}\|_1^2}{N \tau} \\ &\leq \frac{c_3}{4n\tau} (1 + 2c_2^2 n^2 N + 2c_2^2 n^3) C_{\mathcal{F}}^2 + \frac{c_3 \sigma^2}{4n\tau},\end{aligned}$$

for $\tau \geq c_2 n \sqrt{N} C_{\mathcal{F}}$. Without loss of generality, we assume that $\sqrt{\log N / N} \sigma \leq \sqrt{c_4} C_{\mathcal{F}}$, since otherwise the upper bound in Theorem 1 is larger than $C_{\mathcal{F}}^2$ and holds trivially by outputting $\hat{\boldsymbol{\rho}} = \mathbf{0}$. It is then seen that

$$\begin{aligned}\|\mathbf{e}_{k+1}\|_1 &\leq \frac{c_3}{4n\tau} (1 + 2c_2^2 n^2 N + 2c_2^2 n^3) C_{\mathcal{F}}^2 + \frac{c_3 c_4 N C_{\mathcal{F}}^2}{4n\tau} \\ &= \frac{C_{\mathcal{F}}^2}{\tau} \left(\frac{c_3}{4n} (1 + 2c_2^2 n^2 N + 2c_2^2 n^3 + c_4 N) \right) \\ &\leq \frac{c_1 C_{\mathcal{F}}}{(L + n) N n},\end{aligned}$$

with the proviso that

$$\tau \geq \frac{N n (L + n) C_{\mathcal{F}}}{c_1} \left(\frac{c_3}{4n} (1 + 2c_2^2 n^2 N + 2n^3 c_2^2 + c_4 N) \right) + c_2 n \sqrt{N} C_{\mathcal{F}}.$$

This concludes the proof.

References

- Ahn, K., Cheng, X., Daneshmand, H., and Sra, S. (2023). Transformers learn to implement preconditioned gradient descent for in-context learning. *Advances in Neural Information Processing Systems*, 36:45614–45650.
- Ahuja, K., Panwar, M., and Goyal, N. (2023). In-context learning through the bayesian prism. *arXiv preprint arXiv:2306.04891*.
- Akyürek, E., Schuurmans, D., Andreas, J., Ma, T., and Zhou, D. (2023). What learning algorithm is in-context learning? investigations with linear models. In *International Conference on Learning Representations*.
- Bach, F. (2017). Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(19):1–53.
- Bai, Y., Chen, F., Wang, H., Xiong, C., and Mei, S. (2023). Transformers as statisticians: Provable in-context learning with in-context algorithm selection. *Advances in neural information processing systems*, 36:57125–57211.
- Barron, A. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945.
- Beck, A. (2017). *First-order methods in optimization*. SIAM.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., et al. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

- Chen, L., Peng, B., and Wu, H. (2024a). Theoretical limitations of multi-layer transformer. *arXiv preprint arXiv:2412.02975*.
- Chen, S. and Li, Y. (2024). Provably learning a multi-head attention layer. *arXiv preprint arXiv:2402.04084*.
- Chen, S., Sheen, H., Wang, T., and Yang, Z. (2024b). Training dynamics of multi-head softmax attention for in-context learning: Emergence, convergence, and optimality. *arXiv preprint arXiv:2402.19442*.
- Cheng, X., Chen, Y., and Sra, S. (2024). Transformers implement functional gradient descent to learn non-linear functions in context. In *International Conference on Machine Learning*, pages 8002–8037.
- Cole, F., Lu, Y., O’Neill, R., and Zhang, T. (2024). Provable in-context learning of linear systems and linear elliptic pdes with transformers. *arXiv preprint arXiv:2409.12293*.
- Cole, F., Lu, Y., Zhang, T., and Zhao, Y. (2025). In-context learning of linear dynamical systems with transformers: Error bounds and depth-separation. *arXiv preprint arXiv:2502.08136*.
- Dai, D., Sun, Y., Dong, L., Hao, Y., Ma, S., Sui, Z., and Wei, F. (2022). Why can gpt learn in-context? language models implicitly perform gradient descent as meta-optimizers. *arXiv preprint arXiv:2212.10559*.
- Dong, Q., Li, L., Dai, D., Zheng, C., Ma, J., Li, R., Xia, H., Xu, J., Wu, Z., Liu, T., et al. (2022). A survey on in-context learning. *arXiv preprint arXiv:2301.00234*.
- Elhage, N., Nanda, N., Olsson, C., Henighan, T., Joseph, N., Mann, B., Askell, A., Bai, Y., Chen, A., Conerly, T., DasSarma, N., Drain, D., Ganguli, D., Hatfield-Dodds, Z., Hernandez, D., Jones, A., Kernion, J., Lovitt, L., Ndousse, K., Amodei, D., Brown, T., Clark, J., Kaplan, J., McCandlish, S., and Olah, C. (2021). A mathematical framework for transformer circuits. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2021/framework/index.html>.
- Feng, G., Zhang, B., Gu, Y., Ye, H., He, D., and Wang, L. (2023). Towards revealing the mystery behind chain of thought: a theoretical perspective. *Advances in Neural Information Processing Systems*, 36:70757–70798.
- Fu, D., Chen, T.-q., Jia, R., and Sharan, V. (2024). Transformers learn to achieve second-order convergence rates for in-context linear regression. *Advances in Neural Information Processing Systems*, 37:98675–98716.
- Furuya, T., de Hoop, M. V., and Peyré, G. (2024). Transformers are universal in-context learners. *arXiv preprint arXiv:2408.01367*.
- Garg, S., Tsipras, D., Liang, P. S., and Valiant, G. (2022). What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35:30583–30598.
- Giannou, A., Rajput, S., Sohn, J.-y., Lee, K., Lee, J. D., and Papailiopoulos, D. (2023). Looped transformers as programmable computers. In *International Conference on Machine Learning*, pages 11398–11442. PMLR.
- Giannou, A., Yang, L., Wang, T., Papailiopoulos, D., and Lee, J. D. (2024). How well can transformers emulate in-context newton’s method? *arXiv preprint arXiv:2403.03183*.
- Gillioz, A., Casas, J., Mugellini, E., and Abou Khaled, O. (2020). Overview of the transformer-based models for nlp tasks. In *2020 15th Conference on computer science and information systems (FedCSIS)*, pages 179–183. IEEE.
- Guo, T., Hu, W., Mei, S., Wang, H., Xiong, C., Savarese, S., and Bai, Y. (2024). How do transformers learn in-context beyond simple functions? a case study on learning with representations. In *International Conference on Learning Representations*.
- Hahn, M. (2020). Theoretical limitations of self-attention in neural sequence models. *Transactions of the Association for Computational Linguistics*, 8:156–171.

- Hahn, M. and Goyal, N. (2023). A theory of emergent in-context learning as implicit structure induction. *arXiv preprint arXiv:2303.07971*.
- Hataya, R., Matsui, K., and Imaizumi, M. (2024). Automatic domain adaptation by transformers in in-context learning. *arXiv preprint arXiv:2405.16819*.
- Hornik, K., Stinchcombe, M., White, H., and Auer, P. (1994). Degree of approximation results for feedforward networks approximating unknown mappings and their derivatives. *Neural computation*, 6(6):1262–1275.
- Huang, Y., Cheng, Y., and Liang, Y. (2024). In-context convergence of transformers. In *International Conference on Machine Learning*, pages 19660–19722.
- Huang, Y., Wen, Z., Singh, A., Chi, Y., and Chen, Y. (2025). Transformers provably learn chain-of-thought reasoning with length generalization.
- Jelassi, S., Brandfonbrener, D., Kakade, S. M., and Malach, E. (2024). Repeat after me: Transformers are better than state space models at copying. *arXiv preprint arXiv:2402.01032*.
- Khan, S., Naseer, M., Hayat, M., Zamir, S. W., Khan, F. S., and Shah, M. (2022). Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):1–41.
- Kim, J. and Suzuki, T. (2024). Transformers learn nonlinear features in context: Nonconvex mean-field dynamics on the attention landscape. In *Forty-first International Conference on Machine Learning*.
- Kurková, V. and Sanguinetti, M. (2002). Bounds on rates of variable-basis and neural-network approximation. *IEEE Transactions on Information Theory*, 47(6):2659–2665.
- Kwon, S. M., Xu, A. S., Yaras, C., Balzano, L., and Qu, Q. (2025). Out-of-distribution generalization of in-context learning: A low-dimensional subspace perspective. *arXiv preprint arXiv:2505.14808*.
- Li, H., Wang, M., Lu, S., Cui, X., and Chen, P.-Y. (2024a). Training nonlinear transformers for efficient in-context learning: A theoretical learning and generalization analysis. *arXiv preprint arXiv:2402.15607*.
- Li, Y., Ildiz, M. E., Papailiopoulos, D., and Oymak, S. (2023). Transformers as algorithms: Generalization and stability in in-context learning. In *International conference on machine learning*, pages 19565–19594. PMLR.
- Li, Z., Liu, H., Zhou, D., and Ma, T. (2024b). Chain of thought empowers transformers to solve inherently serial problems. *arXiv preprint arXiv:2402.12875*, 1.
- Likhoshervstov, V., Choromanski, K., and Weller, A. (2021). On the expressive power of self-attention matrices. *arXiv preprint arXiv:2106.03764*.
- Lin, T., Wang, Y., Liu, X., and Qiu, X. (2022). A survey of transformers. *AI open*, 3:111–132.
- Liu, B., Ash, J. T., Goel, S., Krishnamurthy, A., and Zhang, C. (2022). Transformers learn shortcuts to automata. *arXiv preprint arXiv:2210.10749*.
- Mahankali, A. V., Hashimoto, T., and Ma, T. (2024). One step of gradient descent is provably the optimal in-context learner with one layer of linear self-attention. In *International Conference on Learning Representations*.
- Merrill, W. and Sabharwal, A. (2024). The expressive power of transformers with chain of thought. In *International Conference on Learning Representations*.
- Nichani, E., Damian, A., and Lee, J. D. (2024). How transformers learn causal structure with gradient descent. In *International Conference on Machine Learning*, pages 38018–38070.
- Peng, B., Narayanan, S., and Papadimitriou, C. (2024). On limitations of the transformer architecture. In *First Conference on Language Modeling*.

- Pérez, J., Marinković, J., and Barceló, P. (2019). On the turing completeness of modern neural network architectures. *arXiv preprint arXiv:1901.03429*.
- Sanford, C., Hsu, D., and Telgarsky, M. (2024). Transformers, parallel computation, and logarithmic depth. *arXiv preprint arXiv:2402.09268*.
- Sanford, C., Hsu, D. J., and Telgarsky, M. (2023). Representational strengths and limitations of transformers. *Advances in Neural Information Processing Systems*, 36:36677–36707.
- Shamshad, F., Khan, S., Zamir, S. W., Khan, M. H., Hayat, M., Khan, F. S., and Fu, H. (2023). Transformers in medical imaging: A survey. *Medical image analysis*, 88:102802.
- Shen, L., Mishra, A., and Khashabi, D. (2023). Do pretrained transformers learn in-context by gradient descent? *arXiv preprint arXiv:2310.08540*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.
- Vladymyrov, M., Von Oswald, J., Sandler, M., and Ge, R. (2024). Linear transformers are versatile in-context learners. *Advances in Neural Information Processing Systems*, 37:48784–48809.
- Von Oswald, J., Niklasson, E., Randazzo, E., Sacramento, J., Mordvintsev, A., Zhmoginov, A., and Vladymyrov, M. (2023). Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pages 35151–35174.
- von Oswald, J., Schlegel, M., Meulemans, A., Kobayashi, S., Niklasson, E., Zucchet, N., Scherrer, N., Miller, N., Sandler, M., Vladymyrov, M., et al. (2023). Uncovering mesa-optimization algorithms in transformers. *arXiv preprint arXiv:2309.05858*.
- Wang, Z., Jiang, B., and Li, S. (2024). In-context learning on function classes unveiled for transformers. In *Forty-first International Conference on Machine Learning*.
- Wen, K., Dang, X., and Lyu, K. (2024). Rnns are not transformers (yet): The key bottleneck on in-context retrieval. *arXiv preprint arXiv:2402.18510*.
- Wen, K., Li, Y., Liu, B., and Risteski, A. (2023). Transformers are uninterpretable with myopic methods: a case study with bounded dyck grammars. *Advances in Neural Information Processing Systems*, 36:38723–38766.
- Xie, S. M., Raghunathan, A., Liang, P., and Ma, T. (2022). An explanation of in-context learning as implicit bayesian inference. In *International Conference on Learning Representations*.
- Yang, T., Huang, Y., Liang, Y., and Chi, Y. (2024). In-context learning with representations: Contextual generalization of trained transformers. *arXiv preprint arXiv:2408.10147*.
- Yao, S., Peng, B., Papadimitriou, C., and Narasimhan, K. (2021). Self-attention networks can process bounded hierarchical languages. *arXiv preprint arXiv:2105.11115*.
- Zhang, R., Frei, S., and Bartlett, P. L. (2024). Trained transformers learn linear models in-context. *Journal of Machine Learning Research*, 25(49):1–55.
- Zhang, Y., Zhang, F., Yang, Z., and Wang, Z. (2023). What and how does in-context learning learn? bayesian model averaging, parameterization, and generalization. *arXiv preprint arXiv:2305.19420*.