
On Expressive Power of Looped Transformers: Theoretical Analysis and Enhancement via Timestep Encoding

Kevin Xu¹ Issei Sato¹

Abstract

Looped Transformers provide advantages in parameter efficiency, computational capabilities, and generalization for reasoning tasks. However, their expressive power regarding function approximation remains underexplored. In this paper, we establish the approximation rate of Looped Transformers by defining the modulus of continuity for sequence-to-sequence functions. This reveals a limitation specific to the looped architecture. That is, the analysis prompts the incorporation of scaling parameters for each loop, conditioned on timestep encodings. Experiments validate the theoretical results, showing that increasing the number of loops enhances performance, with further gains achieved through the timestep encoding. Code is available at <https://github.com/kevin671/tmlt>.

1. Introduction

Transformers (Vaswani et al., 2017) have become the standard architecture for a wide range of machine learning tasks, including natural language processing and computer vision. However, they exhibit certain limitations, particularly when applied to complex tasks. The expressive power of Transformers is theoretically constrained (Merrill & Sabharwal, 2023; Feng et al., 2023), and they empirically struggle with reasoning and planning problems (Kambhampati et al., 2024). Although chain-of-thought reasoning (Wei et al., 2022) can mitigate these challenges in some cases, it typically relies on manually crafted prompts or costly intermediate supervision. Moreover, Transformers encounter difficulties with length generalization (Deletang et al., 2023) and require substantial computational resources as the number of model parameters increases (Pope et al., 2022).

¹The University of Tokyo. Correspondence to: Kevin Xu <kevinxu@g.ecc.u-tokyo.ac.jp>, Issei Sato <sato@g.ecc.u-tokyo.ac.jp>.

To address these limitations, Looped Transformers presents a promising approach. The architecture consists of fixed-size Transformer layers, in which the output is recursively fed back into the input. Looped Transformers exhibit advantages in parameter efficiency thanks to their weight-tying structure (Lan et al., 2020; Takase & Kiyono, 2021; Csordás et al., 2024; Bae et al., 2025), achieving performance comparable to standard Transformers while using fewer parameters. Additionally, they are well suited for size generalization by adjusting the loop count based on task complexity (Dehghani et al., 2019; Fan et al., 2024b). Their recursive structure endows them with the expressive power to emulate iterative algorithms and universal computational capabilities, akin to programmable computers (Giannou et al., 2023). Furthermore, their inductive bias enhances performance on reasoning tasks (Saunshi et al., 2025).

In contrast, the expressive power of Looped Transformers and the properties unique to the looped architecture in function approximation remain unexplored. The expressive power of standard Transformers, on the other hand, has been examined extensively in prior studies. These studies show that Transformers can be universal approximators for continuous permutation-equivariant functions on compact domains (Yun et al., 2020; Kajitsuka & Sato, 2024). Furthermore, their approximation rate has been analyzed by identifying specific properties of the target functions (Takakura & Suzuki, 2023; Jiang & Li, 2024; Wang & E, 2024), providing insights into the underlying characteristics of Transformer architectures. However, these findings cannot be directly extended due to the weight-tying constraints. Although the approximation rate of looped ReLU networks has been established only recently (Zhang et al., 2023), that of Looped Transformers remains unknown.

Our contributions are summarized as follows:

- We establish the approximation rate of Looped Transformers for fixed-length continuous sequence-to-sequence functions, with respect to the number of loops and three newly defined types of moduli of continuity.
- We identify an inherent limitation of Looped Transformers and address it by proposing Timestep-Modulated Looped Transformers (TMLT), which incorporate scaling parameters that are conditioned on timestep encodings.

2. Background

This section defines the Transformer and Looped Transformer architectures, reviews related work, and examines prior theoretical studies on the function approximation capabilities of Transformers and weight-tied networks, thereby clarifying the research question addressed in this paper.

Notations: Vectors are represented by lowercase boldface letters *e.g.*, \mathbf{v} , and matrices are denoted by uppercase boldface letters *e.g.*, \mathbf{X} . The i -th element of a vector \mathbf{v} is denoted by v_i , and the (i, j) -th element of a matrix \mathbf{X} is denoted by $X_{i,j}$. The n -th column of a matrix \mathbf{X} is denoted by $\mathbf{X}_{:,n}$.

Given an input sequence $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N] \in \mathbb{R}^{m \times N}$, where $\mathbf{x}_i \in \mathbb{R}^m$, and a function $f : \mathbb{R}^m \rightarrow \mathbb{R}^m$, the token-wise application of f is denoted by the bold symbol \mathbf{f} *i.e.*

$$\mathbf{f}(\mathbf{X}) = [f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_N)] \in \mathbb{R}^{m \times N}.$$

For $p \in [1, \infty)$, the p -norm, denoted by $\|\cdot\|_p$, represents the entry-wise p -norm. This norm applies to both vectors and matrices *e.g.*, $\|\mathbf{X}\|_p$. The L^p -norm of a function is defined for $p \in [1, \infty)$ as:

$$\|f\|_{L^p} := \left(\int_{\Omega} \|f(\mathbf{X})\|_p^p d\mathbf{X} \right)^{1/p},$$

where Ω represents the domain of the function f .

2.1. Transformer Architecture

Given an input sequence $\mathbf{X} \in \mathbb{R}^{m \times N}$, composed of N token embedding of dimension size m , the self-attention layers with h heads and head size s , and the feed-forward layer with width size q , are defined as follows:

$$\text{Attn}(\mathbf{X}) = \sum_{i=1}^h \mathbf{W}_O^i \mathbf{W}_V^i \mathbf{X} \sigma_s \left[(\mathbf{W}_K^i \mathbf{X})^\top \mathbf{W}_Q^i \mathbf{X} \right],$$

$$\text{FF}(\mathbf{X}_{:,n}) = \mathbf{W}_2 \sigma_R(\mathbf{W}_1 \mathbf{X}_{:,n} + \mathbf{b}_1) + \mathbf{b}_2,$$

where $\mathbf{W}_O^i \in \mathbb{R}^{m \times s}$, \mathbf{W}_V^i , \mathbf{W}_K^i , $\mathbf{W}_Q^i \in \mathbb{R}^{s \times m}$, $\mathbf{W}_1 \in \mathbb{R}^{q \times m}$, $\mathbf{W}_2 \in \mathbb{R}^{m \times q}$, $\mathbf{b}_1 \in \mathbb{R}^q$, $\mathbf{b}_2 \in \mathbb{R}^m$ are parameters, σ_R denotes ReLU function, and σ_s denotes a softmax operator applied to the columns of the matrix.

Transformer block $\text{TF} : \mathbb{R}^{m \times N} \rightarrow \mathbb{R}^{m \times N}$ is defined by

$$\begin{aligned} \mathbf{X}' &= \mathbf{X} + \text{Attn}(\mathbf{X}), \\ \text{TF}(\mathbf{X}) &= \mathbf{X}' + \mathbf{FF}(\mathbf{X}'), \end{aligned}$$

where \mathbf{FF} represents token-wise FF. In other words,

$$\text{TF} = (\text{id} + \mathbf{FF}) \circ (\text{id} + \text{Attn}),$$

where id denotes the identity mapping, where we omit the domain of definition for simplicity. For the analysis of expressive power in Section 3, we exclude layer normalization and our constructive proof relies on the softmax function to approximate the hardmax function as in previous studies (Yun et al., 2020; Kim et al., 2023).

2.2. Looped Transformer

Looped Transformer with a single layer is represented as:

$$\mathcal{L}_2 \circ \text{TF}^{\circ r} \circ \mathcal{L}_1,$$

where \mathcal{L}_2 and \mathcal{L}_1 represent token-wise affine linear layers, and $\text{TF}^{\circ r}$ denotes the composition of TF applied r times. While we focus on single-layer as (Dehghani et al., 2019; Lan et al., 2020; Yang et al., 2024; Fan et al., 2024b), they can also be implemented with multiple layers as (Csordás et al., 2024; Bae et al., 2025; Saunshi et al., 2025).

Overview of Previous Work The recursive structure was introduced into Transformers (Dehghani et al., 2019), where the number of loops can be adaptively adjusted, allowing for size generalization (Fan et al., 2024a). Looped Transformers are closely related to weight-tying Transformers (Lan et al., 2020; Takase & Kiyono, 2021), achieving performance comparable to standard Transformers using fewer parameters. Deep equilibrium models (Bai et al., 2019), which compute fixed points of iterative layers, are also related. In addition, the recursive structure enables the model to emulate iterative algorithms, including basic computational primitives (Giannou et al., 2023) and learning algorithms (Giannou et al., 2024; Yang et al., 2024). Furthermore, recent studies have demonstrated that Looped Transformers exhibit an inductive bias towards reasoning tasks (Saunshi et al., 2025). To improve performance, more sophisticated architectures, such as mixture-of-experts (Csordás et al., 2024) and relaxed weight-tying (Bae et al., 2025), have been introduced.

2.3. Theoretical Analysis on Expressive Power

We review related work and summarize the comparisons between our problem setting and previous studies in Table 1.

Universality of Transformers The universal approximation theorem for fully connected neural networks (Cybenko, 1989; Hornik et al., 1989) shows that networks of sufficient size can approximate certain classes of functions with arbitrarily low error. For Transformers, the target function class extends to sequence-to-sequence functions. Transformers compute a *contextual mapping* of the input, which requires capturing the entire sequence and computing the token embedding within context (Yun et al., 2020), formulated as:

Definition 2.1 (Yun et al., 2020). Consider a finite set $\mathbb{L} \subset \mathbb{R}^{d \times N}$. A map $\text{CM} : \mathbb{L} \rightarrow \mathbb{R}^{1 \times N}$ defines a *contextual mapping* if the map satisfies the following:

1. For any $\mathbf{L} \in \mathbb{L}$, the N entries in $\text{CM}(\mathbf{L})$ are all distinct.
2. For any $\mathbf{L}, \mathbf{L}' \in \mathbb{L}$, with $\mathbf{L} \neq \mathbf{L}'$, all entries of $\text{CM}(\mathbf{L})$ and $\text{CM}(\mathbf{L}')$ are distinct.

Prior studies have shown that Transformers can compute contextual mappings, enabling memorization (Kim et al.,

Table 1. Comparisons of our problem setting with related work on the theoretical analysis of function approximation.

Paper	Model Type	Function Class	Approximation Rate	Looped (Weight-Tying)
Yarotsky (2018)	FFN	Continuous functions	✓	×
Yun et al. (2020)	Transformer	Continuous seq-to-seq functions	×	×
Takakura & Suzuki (2023)	Transformer	γ -smooth infinite-length	✓	×
Kajitsuka & Sato (2024)	Transformer	Continuous seq-to-seq functions	×	×
Jiang & Li (2024)	Transformer	Temporal coupled functions	✓	×
Wang & E (2024)	Transformer	Long but sparse memories	✓	×
Zhang et al. (2023)	FFN	Continuous functions	✓	✓
Ours	Transformer	Continuous seq-to-seq functions	✓	✓

2023) and universal approximation (Yun et al., 2020; Kajitsuka & Sato, 2024).

For Looped Transformers, as the fixed parameters of a single Transformer layer are used, the results of previous studies cannot be directly applied. This leads to the question: *Can Looped Transformers compute contextual mappings?* and *Are they universal approximators?*

Approximation Rate of Transformers Beyond the universality, the approximation rate provides deeper insights into the characteristics of models (Barron, 1993; Yarotsky, 2018). This rate is derived as an upper bound of error in terms of the properties of the target functions and the complexity of the networks. For Transformers, recent studies have investigated these rates and the nature of the target functions (Takakura & Suzuki, 2023; Jiang & Li, 2024; Wang & E, 2024). Specifically, they have shown conditions under which Transformers can overcome the curse of dimensionality (Takakura & Suzuki, 2023) and revealed structures in target functions that Transformers can effectively approximate (Jiang & Li, 2024; Wang & E, 2024).

Our study focuses on understanding the architectural properties of Looped Transformers, particularly in comparison to standard Transformers. To this end, we explore the approximation rate and investigate the properties of target functions that determine their approximation errors.

Expressive Power of Weight-Tied Neural Networks Recently, it has been shown that single *fixed-size* networks can serve as universal approximators in a parameter-efficient manner; that is, the parameter count depends solely on the input dimension, not the approximation error (Zhang et al., 2023). Furthermore, the approximation rate of weight-tied ReLU networks has been established with respect to the number of loops and the modulus of continuity of continuous functions (Zhang et al., 2023). The modulus of continuity for $g : \mathbb{R}^d \rightarrow \mathbb{R}$ and $\delta \geq 0$ is defined as:

$$\omega_g(\delta) := \sup \{ |g(\mathbf{x}) - g(\mathbf{x}')| : \|\mathbf{x} - \mathbf{x}'\|_2 \leq \delta \}.$$

Our question is whether the results can be extended to sequence-to-sequence functions and Transformers, which

require contextual mappings. For a sequence-to-sequence function $f : \mathbb{R}^{d \times N} \rightarrow \mathbb{R}^{d \times N}$, the modulus of continuity can be generalized as:

$$\omega_f(\delta) := \sup \{ \|f(\mathbf{X}) - f(\mathbf{X}')\|_p : \|\mathbf{X} - \mathbf{X}'\|_2 \leq \delta \}$$

We investigate whether this modulus of continuity alone can determine the approximation rate.

For Looped Transformers, it has been shown that they can represent standard Transformers, although their parameter count grows with both the desired approximation accuracy and the sequence length (Saunshi et al., 2025). Moreover, no existing work has established their approximation rate.

3. Approximation Rate Analysis

In this section, we establish the approximation rate of Looped Transformers. We define three types for the modulus of continuity in Section 3.2 that determine the approximation rate. The main results are presented in Section 3.3, followed by a proof sketch in Section 3.4.

3.1. Preliminaries

The target function class of our analysis is continuous functions that Transformers can represent. Specifically, these are *permutation-equivariant* functions, defined as follows:

Definition 3.1 (Yun et al., 2020). A function $f : \mathbb{R}^{d \times N} \rightarrow \mathbb{R}^{d \times N}$ is said to be *permutation equivariant* if $f(\mathbf{X}\mathbf{P}) = f(\mathbf{X})\mathbf{P}$ holds for any permutation matrix \mathbf{P} . Let $\mathcal{F}_{\text{PE}}(\Omega)$ denote the set of continuous functions, defined on Ω , that are permutation equivariant.

We evaluate both the number of parameters and the *bit complexity*, the maximum number of bits required to represent the network’s weights (Vardi et al., 2022; Kim et al., 2023).

In our proofs, we introduce *IDs* for tokens, sequences, and tokens within sequences as theoretical constructs to formalize contextual mappings.

Definition 3.2. A *token ID* is a unique integer assigned to each token. A *sequence ID* uniquely identifies each sentence.

A *contextual token ID* uniquely identifies a specific token within a specific sentence. We denote the set of contextual token IDs as $\mathcal{K} = 0, 1, \dots, K-1$, with corresponding embeddings $\mathbf{y}_k \in \mathbb{R}^d$ for each $k \in \mathcal{K}$.

This notion is defined in Kim et al. (2023), to which we refer for further details, for constructive proofs of contextual mappings. The actual construction of contextual token IDs may vary depending on the specific proof. In our case, we adopt a different construction from that of Kim et al. (2023).

3.2. Definition of modulus of Continuity

As briefly mentioned in the preliminary discussion, we define the modulus of continuity in Eq. 2.3 as:

Definition 3.3 (Modulus of Sequence Continuity). Given a sequence-to-sequence continuous function $f : \mathbb{R}^{d \times N} \rightarrow \mathbb{R}^{d \times N}$, the modulus of *sequence continuity* is defined by:

$$\omega_f(\delta) := \sup \{ \|f(\mathbf{X}) - f(\mathbf{X}')\|_p : \|\mathbf{X} - \mathbf{X}'\|_2 \leq \delta \}.$$

We omit the subscript p for simplicity. This quantifies how the output sequence shifts relative to differences in input, hence referred to as *sequence continuity*.

We found that this alone is insufficient to determine the approximation rate of Looped Transformers, in contrast to the case of ReLU networks (Zhang et al., 2023). Informally, this issue arises because Transformers compute contextual mappings. We notably identified two additional types of modulus of continuity, defined as follows.

Definition 3.4 (Modulus of Contextual Continuity). Given a sequence-to-sequence continuous function $f : \mathbb{R}^{d \times N} \rightarrow \mathbb{R}^{d \times N}$, the modulus of *contextual continuity* is defined by:

$$\omega_f^{\text{cont}}(\delta) := \sup_{n, \mathbf{X}, \mathbf{X}'} \{ \|f(\mathbf{X})_{:,n} - f(\mathbf{X}')_{:,n}\|_p : \|\mathbf{X} - \mathbf{X}'\|_2 \leq \delta, \mathbf{X}_{:,n} = \mathbf{X}'_{:,n} \}.$$

Definition 3.5 (Modulus of Token Continuity). Given a sequence-to-sequence continuous function $f : \mathbb{R}^{d \times N} \rightarrow \mathbb{R}^{d \times N}$, the modulus of *token continuity* is defined by:

$$\omega_f^{\text{tok}}(\delta) := \sup_{n, \mathbf{X}, \mathbf{X}'} \{ \|f(\mathbf{X})_{:,n} - f(\mathbf{X}')_{:,n}\|_p : \|\mathbf{X}_{:,n} - \mathbf{X}'_{:,n}\|_2 \leq \delta, \mathbf{X}_{:,q} = \mathbf{X}'_{:,q} (\forall q \neq n) \}.$$

The *modulus of contextual continuity* quantifies the variation in the contextual token embeddings induced by perturbations of context. For example, consider the sentences: (1) “I write papers” and (2) “You write books”. It measures the difference in the contextual token embedding of the same word ‘write’ within different contexts.

On the other hand, the *modulus of token continuity* quantifies the variation in the output embedding caused by perturbations to the token itself within the same context such as (1) “I write papers” and (2) “I draft papers”.

3.3. Main Result

The result establishes the approximation rate of Looped Transformers in terms of the number of loops and the three types of moduli of continuity of the target function.

Theorem 3.6. Given a function $f \in \mathcal{F}_{\text{PE}}([0, 1]^{d \times N})$, $r > N$, there exists a Looped Transformer, composed of $\text{TF} : \mathbb{R}^{(17d+9) \times N} \rightarrow \mathbb{R}^{(17d+9) \times N}$ with two heads, head size 1, and width size of $q = 49d + 25$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R}^d \rightarrow \mathbb{R}^{17d+9}$ and $\mathcal{L}_2 : \mathbb{R}^{17d+9} \rightarrow \mathbb{R}^d$ s.t.

$$\begin{aligned} & \|\mathcal{L}_2 \circ \text{TF}^{or} \circ \mathcal{L}_1 - f\|_{L^p} \\ & \leq (Nd)^{\frac{1}{p}} \left(\omega_f^{\text{tok}}(\delta\sqrt{d}) + \omega_f^{\text{cont}}(\delta\sqrt{Nd}) \right) + \omega_f(\delta\sqrt{Nd}) \\ & \quad + \mathcal{O}(N^{2/p}\delta^{d/p}) + \mathcal{O}\left((M\delta)^{-1}dN\right)^{1/p}, \\ & \text{where } \delta = ((r - N)/2)^{-1/((N+1)d+1)}, \end{aligned}$$

where M is the maximum absolute value of the model parameters, and the bit complexity is $\mathcal{O}(\delta^{-(N+1)d})$.

Theorem 3.6 shows that increasing the number of loops r reduces the approximation error. Under infinite-precision weights, this leads to a universal approximation theorem.

Corollary 3.7 (Universality). The hypothesis space of Looped Transformers, defined by

$$\begin{aligned} \mathcal{H} := & \{ \mathcal{L}_2 \circ \text{TF}^{or} \circ \mathcal{L}_1 : [0, 1]^{d \times N} \rightarrow [0, 1]^{d \times N} \mid \\ & m, q \leq Cd, h = 2, s = 1, r \in \mathbb{N}, \mathbf{W} \in \mathbb{R}^{n_w} \}, \end{aligned}$$

where C is a positive constant, \mathbf{W} denotes the flattened set of all weights in the network, and n_w represents the total number of these weights, is dense in $\mathcal{F}_{\text{PE}}([0, 1]^{d \times N})$ w.r.t. the L^p norm.

This approximation analysis highlights the characteristics of Looped Transformers, including both their capabilities and limitations, as summarized below:

- While the number of parameters remains fixed at $\mathcal{O}(d)$, independent of the desired approximation accuracy and the sequence length, the error can be reduced by increasing the number of loops.
- Looped Transformers, even with weight-tied self-attention using a hard-max function, can compute contextual mappings and become universal approximators.
- The approximation rate depends on three types of continuity, with contextual and token dependencies unique to Looped Transformers; these dependencies are not present in standard Transformers or looped ReLU networks.

Our contribution lies in establishing the approximation rate with respect to the number of loops, based on novel moduli of continuity that are unique to Looped Transformers.

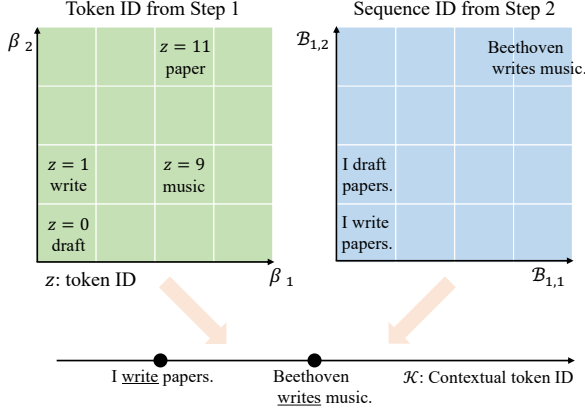


Figure 1. The networks construct contextual token IDs by combining token IDs with sequence IDs.

Furthermore, the additional dependency can amplify the approximation error, revealing a limitation inherent to Looped Transformers. A detailed discussion of this issue, along with improvement methods, is provided in Section 4.

3.4. Proof Sketch

This section presents a proof sketch, emphasizing distinctions from prior studies and challenges unique to the looped architecture. The formal proof is provided in Appendix A.

The basic strategy involves approximating the continuous target function f with a piecewise constant function \tilde{f} , which is approximated by the network, denoted by \hat{f} . For $\delta^{-1} \in \mathbb{N}$, $\delta^{-1} \geq 2$, the input space $[0, 1]^{d \times N}$ is divided into discretized cubes with width δ , denoted by $\{Q_{\mathcal{B}}\}_{\mathcal{B} \in \{0, 1, \dots, \delta^{-1}-1\}^{d \times N}}$. Each cube is assigned a representative point $\hat{X}_{\mathcal{B}} \in Q_{\mathcal{B}}$, and the piecewise constant function \tilde{f} is then defined as:

$$\tilde{f}(X) = f(\hat{X}_{\mathcal{B}}) \quad \text{where } \mathcal{B} \text{ satisfies } X \in Q_{\mathcal{B}}. \quad (1)$$

The approximation with networks consists of three steps. First, the network assigns a *token ID* to each token. Second, it assigns a *sequence ID*. The combination of the *token ID* and *sequence ID* constitutes the *contextual token IDs* as in Fig. 1. Finally, these are mapped to embeddings that represent the output of the target function at each token.

Step 1. Token-wise Quantization. The network uses the feed-forward network to assign each input token, denoted by $X_{:,n}$, to a token ID, denoted by z , in a token-wise manner.

$$X_{:,n} \in [0, 1]^d \rightarrow z \in \{0, 1, \dots, \delta^{-d} - 1\}. \quad (2)$$

Step 2. Contextual Mapping. The network, given N token IDs computes their sequence ID. We notice that the result of previous studies on Transformers (Yun et al., 2020;

Kim et al., 2023) cannot be directly applied to Looped Transformers due to the following distinctions:

- Yun et al. (2020) employed both sparse and uniform attention mechanisms, whereas Looped Transformers are limited to a single fixed attention layer.
- Kim et al. (2023) used N layers to store N parameters required for representing the target function, whereas Looped Transformers have a fixed parameter size.

Notably, we found that Looped Transformers with N -loops can compute contextual mapping. Let $z \in \{0, 1, \dots, \delta^{-d} - 1\}^N$ represent a sequence consist of N ordered and distinct token IDs, satisfying $z_1 > z_2 > \dots > z_N$. The network then maps z to a sequence ID through an inner product with $u = (\delta^{-d(N-1)}, \dots, \delta^{-d}, 1)$, which satisfies

$$|u^\top z - u^\top z'| > 1, \quad \text{if } z \neq z'. \quad (3)$$

This guarantees that the network assigns distinct sequence IDs for different z . Combined with token IDs, the network computes contextual mapping. The key idea is that the network requires only δ^{-d} to represent u , allowing it to be implemented with Looped Transformers.

Step 3. Function Value Mapping. The network maps the contextual token IDs into the target embeddings in a token-wise manner, using $K - 1$ loops to sequentially map $k = 0, 1, \dots, K - 1$ to $\tilde{y}_k \in \mathbb{R}^d$, which approximates y_k , in each iteration. In our constructive proofs, we design both the set of contextual token IDs and their ordering.

Weight-tied feed-forward networks cannot map accurately, and the error can only be bounded by the maximum difference between adjacent contextual token embeddings, *i.e.*

$$|(\tilde{y}_k - y_k)_i| \leq \max_{k' \in \mathcal{K}} |(y_{k'} - y_{k'-1})_i| \quad (4)$$

holds for $k \in \mathcal{K}$ and $i = 1, \dots, d$.

Generally, the following inequality holds, for $x \in \mathbb{R}^m$,

$$\max_i |x_i| \leq \|x\|_p \leq m^{\frac{1}{p}} \max_i |x_i|. \quad (5)$$

That is, by controlling the p -norm, $\|y_k - y_{k-1}\|_p$, the error in Eq. 4 can be bounded. We require \mathcal{K} to be designed such that the differences between neighboring contextual token embeddings are bounded *w.r.t.* the p -norm.

To illustrate our idea, consider the following sentences:

- (1) I write papers. ; I write papers. (different token ID with same sequence ID)
- (2) I write papers. ; You write books. (same token ID with different sequence ID)




Figure 2. Approximation error and modulus of continuity. The linear interpolation technique reduces the error by a factor of $1/\delta^{-1}$.

We found that none of the moduli of continuity, defined in Section 3.2, alone can bound the difference between ‘write’ and ‘papers’ in (1). In contrast, the error of ‘write’ in (2) can be bounded by the contextual continuity, ω_f^{cont} . Thus, we designed contextual token IDs such that, basically, identical or similar tokens with different sequence IDs are positioned adjacent to each other, as shown in Fig. 2. To reduce errors in corner cases, linear interpolation is applied; further details are provided in Appendix A. This allows us to obtain the following error bound.

$$\max_{k' \in \mathcal{K}} \|\mathbf{y}_{k'} - \mathbf{y}_{k-1}\|_p \leq \omega_f^{\text{tok}}(\delta\sqrt{d}) + \omega_f^{\text{cont}}(\delta\sqrt{Nd}). \quad (6)$$

Substituting $\mathbf{x} = \mathbf{y}_k - \mathbf{y}_{k-1}$ into Eq. 5, with Eq. 4 and Eq. 6, the following result holds:

$$|(\tilde{\mathbf{y}}_k - \mathbf{y}_k)_i| \leq \omega_f^{\text{tok}}(\delta\sqrt{d}) + \omega_f^{\text{cont}}(\delta\sqrt{Nd}), \quad (7)$$

for $i = 1, \dots, d$ and $k \in \mathcal{K}$.

Concatenated into a Single Transformer Layer In the final construction, we show that the composition of the three sub-networks from Steps 1, 2, and 3 can be implemented within a single Transformer block. While our proof strategy follows Zhang et al. (2023), their approach necessitates an additional layer. In contrast, we show that a single Transformer block suffices, as detailed in Appendix A.

Deriving Approximation Rate Lastly, we analyze the approximation error of our construction and establish the approximation rate in terms of the number of loops.

With the triangle inequality, we obtain the following:

$$\|\tilde{f} - f\|_{L^p} \leq \int \|\tilde{f}(\mathbf{X}) - f(\mathbf{X})\|_p d\mathbf{X} \quad (8)$$

$$\begin{aligned} &\leq \int \|\tilde{f}(\mathbf{X}) - \bar{f}(\mathbf{X})\|_p d\mathbf{X} + \int \|\bar{f}(\mathbf{X}) - f(\mathbf{X})\|_p d\mathbf{X} \\ &\quad + \mathcal{O}(N^{2/p}\delta^{d/p}) + \mathcal{O}\left(\left((M\delta)^{-1}dN\right)^{1/p}\right), \end{aligned} \quad (9)$$

where $\mathcal{O}(N^{2/p}\delta^{d/p})$ arises from the case where identical tokens appear in sequences, and $\mathcal{O}\left(\left((M\delta)^{-1}dN\right)^{1/p}\right)$ results from the restriction on the norm of weights.

Considering the error within cubes in Eq. 1, we obtain

$$\int \|\bar{f}(\mathbf{X}) - f(\mathbf{X})\|_p d\mathbf{X} \leq \omega_f(\delta\sqrt{Nd}). \quad (10)$$

Since, generally, the norm of sequences can be bounded by the maximum norm of the token-wise vectors as

$$\|f(\mathbf{X})\|_p \leq (Nd)^{\frac{1}{p}} \max_{i,n} |f(\mathbf{X})_{i,n}|, \quad (11)$$

the error between \tilde{f} and \bar{f} can be bounded by

$$\int \|\tilde{f}(\mathbf{X}) - \bar{f}(\mathbf{X})\|_p d\mathbf{X} \leq (Nd)^{\frac{1}{p}} \max_{k' \in \mathcal{K}} |\tilde{\mathbf{y}}_{k'} - \mathbf{y}_{k'}|. \quad (12)$$

Substituting $\mathbf{x} = \tilde{\mathbf{y}}_k - \mathbf{y}_k$ into Eq. 5, and using Eq. 7 and Eq. 12, we obtain:

$$\begin{aligned} &\int \|\tilde{f}(\mathbf{X}) - \bar{f}(\mathbf{X})\|_p d\mathbf{X} \\ &\leq (Nd)^{\frac{1}{p}} \left(\omega_f^{\text{tok}}(\delta\sqrt{d}) + \omega_f^{\text{cont}}(\delta\sqrt{Nd}) \right). \end{aligned} \quad (13)$$

We then express δ in terms of the number of loops r to determine the approximation rate. We use $\delta^{-1} - 1$ loops for Step 1, N loops for Step 2, and $2\delta^{-(N+1)d} - 1$ loops for Step 3, with 1 loop required to connect each step *i.e.*

$$r = \delta^{-1} + 2\delta^{-(N+1)d} + N. \quad (14)$$

Now, δ can be bounded in terms of the number of loops as:

$$\delta^{-1} + 2\delta^{-(N+1)d} = r - N \quad (15)$$

$$\Rightarrow \delta^{-1} \cdot 2\delta^{-(N+1)d} \geq r - N \quad (\delta^{-1} \geq 2) \quad (16)$$

$$\Leftrightarrow \delta \leq \left(\frac{r - N}{2} \right)^{-1/((N+1)d+1)}. \quad (17)$$

By combining Eq. 10 and Eq. 13 with Eq. 9, and substituting Eq. 17, we obtain Theorem 3.6.

4. From Theory to Practice: Introducing Timestep Encoding

The theoretical result in Section 3 highlights a limitation of the looped architecture. We show that a variant of architecture can overcome this limitation.

4.1. Motivation

Limitation Specific to Looped Transformers The approximation rate in Theorem 3.6 includes two additional moduli of continuity, which can lead to increased errors, reflecting a limitation inherent to Looped Transformers.

We can identify the cause of additional dependency in the error in Eq. 4, caused by weight-tied feed-forward networks. This can be formalized as follows:

Lemma 4.1. *Given $\mathbf{y}_k \in \mathbb{R}^d$ for $k = 0, 1, \dots, K - 1$ with*

$$|(\mathbf{y}_k - \mathbf{y}_{k-1})_i| \leq \varepsilon_i \quad \text{for } i = 1, \dots, d,$$

there exists a feed-forward layer $\text{FF} : \mathbb{R}^{12d} \rightarrow \mathbb{R}^{12d}$ with a width size of $18d$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^{12d}$ and $\mathcal{L}_2 : \mathbb{R}^{12d} \rightarrow \mathbb{R}^d$ s.t.

$$|(\mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(K-1)} \circ \mathcal{L}_1(k) - \mathbf{y}_k)_i| \leq \varepsilon_i, \quad (18)$$

for $i = 1, \dots, d$ and $k = 0, 1, \dots, K - 1$.

This shows that large variations in the target function may lead to approximation errors, raising the question of whether inequality in Eq. 18 can be replaced with equality.

Improving Approximation Rate of Looped Transformers

To eliminate this dependency, we introduce *time-dependent* parameters. Specifically, we modify the feed-forward layers by adding a scaling vector for each loop step as follows:

$$\text{FF}(\mathbf{X}) \rightarrow \boldsymbol{\eta}(t) \odot \text{FF}(\mathbf{X}) \quad \text{for the } t\text{-th loops,}$$

where \odot is an element-wise product, $t \in \mathbb{N}$ is the loop index, and $\boldsymbol{\eta}(t) \in \mathbb{R}^d$ is the scaling parameter for each loop. This method is analogous to Hypernetworks (Ha et al., 2016). With the definition

$$(\text{id} + \boldsymbol{\eta} \odot \text{FF})^r := (\text{id} + \boldsymbol{\eta}(r) \odot \text{FF}) \circ \dots \circ (\text{id} + \boldsymbol{\eta}(1) \odot \text{FF}),$$

we show that this model can memorize labels exactly.

Theorem 4.2. *Given $\mathbf{y}_k \in \mathbb{R}^d$ for $k = 0, 1, \dots, K - 1$, there exists a feed-forward layer $\text{FF} : \mathbb{R}^{4d} \rightarrow \mathbb{R}^{4d}$ with a width size of $6d$, $\boldsymbol{\eta}(t) \in \mathbb{R}^{4d}$ for $t = 1, \dots, K - 1$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^{4d}$ and $\mathcal{L}_2 : \mathbb{R}^{4d} \rightarrow \mathbb{R}^d$ s.t.*

$$|(\mathcal{L}_2 \circ (\text{id} + \boldsymbol{\eta} \odot \text{FF})^{\circ(K-1)} \circ \mathcal{L}_1(k) - \mathbf{y}_k)_i| = 0,$$

for $i = 1, \dots, d$ and $k = 0, 1, \dots, K - 1$.

The proof is provided in Appendix B. For implementation, adding parameters per loop increases the total parameter count proportionally. Thus, we introduce timestep encoding.

4.2. Timestep-Modulated Looped Transformer

We employ timestep encodings to condition scaling parameters on the loop index (timestep). This method is inspired by adaptive instance normalization (Peebles & Xie, 2023).

To condition on timesteps, frequency embeddings are processed through a two-layer MLP with hidden size matching the Transformer block and SiLU activation. Let $\text{TE}(t) \in \mathbb{R}^d$ denote timestep embeddings, defined as:

$$\text{TE}(t) = \mathbf{W}_3 \cdot \text{SiLU}(\mathbf{W}_4 \cdot \text{PE}(t) + \mathbf{b}_4) + \mathbf{b}_3,$$

where $\mathbf{W}_3, \mathbf{W}_4 \in \mathbb{R}^{d \times d}$, $\mathbf{b}_3, \mathbf{b}_4 \in \mathbb{R}^d$, and $\text{PE}(t) \in \mathbb{R}^d$ denotes the timestep encoding function that maps the timestep into a d -dimensional embedding, s.t.

$$\text{PE}(t)_{2i} = \sin(t/10000^{2i/d}),$$

$$\text{PE}(t)_{2i+1} = \cos(t/10000^{2i/d}).$$

We use the root mean square layer normalization (RMSNorm) (Zhang & Sennrich, 2019), which is widely used in several recent LLMs (et al., 2023; Team, 2024), defined as:

$$\bar{\mathbf{x}} = \boldsymbol{\alpha} \odot \frac{\mathbf{x}}{\text{RMS}(\mathbf{x})}, \quad \text{where } \text{RMS}(\mathbf{x}) = \sqrt{\frac{1}{d} \sum_{i=1}^d x_i^2},$$

where $\boldsymbol{\alpha} \in \mathbb{R}^d$ is a gain parameter for rescaling. We define time-dependent RMSNorm, denoted by RMSN, as:

$$\text{RMSN}(\mathbf{x}, t) = \boldsymbol{\alpha}(t) \odot \frac{\mathbf{x}}{\text{RMS}(\mathbf{x})}$$

where $\boldsymbol{\alpha}(t) \in \mathbb{R}^d$ is a time-dependent parameter generated by a network. With scaling parameters, the time-dependent Transformer block is defined as follows:

$$\mathbf{X}' = \mathbf{X} + \gamma_1(t) \odot \text{Attn}(\text{RMSN}_1(\mathbf{X}, t)),$$

$$\text{TF}(\mathbf{X}, t) = \mathbf{X}' + \gamma_2(t) \odot \text{FF}(\text{RMSN}_2(\mathbf{X}', t)),$$

where $\gamma_1(t), \gamma_2(t) \in \mathbb{R}^d$ are time-dependent parameters applied token-wise, as well as RMSNorm.

The time-dependent vector parameters are generated as:

$$\boldsymbol{\alpha}_1(t), \boldsymbol{\alpha}_2(t), \gamma_1(t), \gamma_2(t) = \mathbf{W}_5 \cdot \text{SiLU}(\text{TE}(t)) + \mathbf{b}_5,$$

where $\mathbf{W}_5 \in \mathbb{R}^{4d \times d}$ and $\mathbf{b}_5 \in \mathbb{R}^d$.

5. Experiments

This section presents experimental results supporting our theoretical findings. We used Looped Transformers with varying numbers of loops, both with and without timestep encoding, and compared to standard Transformers. We assess approximation capabilities based on test evaluation, as we observe a strong correlation between train and test performance. The details are provided in Appendix C.

Table 2. Test accuracy for reasoning tasks. Performance improves as the number of loops increases..

Task	TF	Looped TF				w/ Timestep Encoding			
	L=6	r=4	r=8	r=16	r=32	r=4	r=8	r=16	r=32
Sudoku	0.0	0.0	0.0	65.6	87.9	0.0	0.0	62.0	90.2
Countdown	53.8	28.3	52.7	81.0	88.1	33.2	54.4	80.2	90.5
	L=12	r=5	r=10	r=50	r=100	r=5	r=10	r=50	r=100
LCS (60)	70.0	66.0	81.8	98.6	96.9	68.5	80.5	99.3	97.1
LCS (100)	39.8	39.6	45.1	93.5	98.2	36.7	45.6	98.1	98.6
ED (40)	54.2	41.4	57.9	85.4	90.4	44.8	63.5	94.5	96.1
ED (60)	41.4	23.8	32.6	47.3	47.7	26.6	38.9	57.3	88.3

5.1. Problem Setting

We evaluate the model on two types of tasks. The first consists of reasoning problems known to be challenging for standard Transformers. These are used to examine whether increasing the number of loops and incorporating timestep encodings can enhance performance. The second includes core Transformer benchmarks, such as in-context learning and language modeling.

5.1.1. REASONING TASKS

Dynamic Programming is a method for solving complex problems by breaking them down into simpler sub-problems. We use edit distance (ED) and longest common subsequence (LCS) tasks with varying input lengths. Each task has 10^6 train samples and 10^3 test samples.

Sudoku is a constrained satisfaction problem that involves filling a 9×9 grid with digits from 1 to 9, such that each digit appears exactly once in every row, column, and predefined 3×3 sub-grid. The grid is flattened into a sequence representation. Unlike (Yang et al., 2023), we use the dataset from (Radcliffe, 2020), sampling 3M instances for training and 100K for testing.

Countdown is a game in which a given set of input numbers must be combined using basic arithmetic operations to reach a target number (Yao et al., 2023; Gandhi et al., 2024). We consider cases with 4 input numbers and target numbers ranging from 10 to 100, where 10% of the target numbers are reserved for evaluation. We generate 5M samples for training and 1K samples for testing.

5.1.2. IN-CONTEXT AND LANGUAGE MODELING

The in-context learning problem is to learn the function class from a given sequence, which was investigated with Looped Transformers (Yang et al., 2024) without timestep encodings. We use decision tree functions. For the language modeling task, we use the WikiText-103 (Merity et al., 2017) dataset, containing over 100 million tokens from Wikipedia articles. Details are in Appendix C.2 and Appendix C.3.

5.2. Results

The results in Table 2 demonstrate that increasing the number of loops improves performance on reasoning tasks, with higher loop counts significantly outperforming standard Transformers. Furthermore, incorporating timestep encodings leads to additional gains; in particular, for the edit distance task with input size $n = 60$, the model with loop counts $r = 100$ achieves significantly better performance when timestep encodings are incorporated.

 Table 3. MSE (\downarrow) on the in-context learning task.

	TF L=12	Looped r=12	w/ Timestep r=12
Test	8.6e-03	1.4e-02	1.7e-03

 Table 4. Perplexity (\downarrow) on the WikiText-103 dataset.

	TF L=12	Looped r=24	w/ Timestep r=24
Train	15.9	17.1	15.9
Test	20.5	20.6	19.6

As evidenced by the results in Table 3 and Table 4, the use of timestep encodings leads to performance gains in both in-context learning and language modeling.

6. Conclusion

We establish the approximation rate of Looped Transformers with respect to the number of loops and the moduli of continuity of the target function. Our analysis reveals a limitation of Looped Transformers, which is addressed by timestep encodings. To the best of our knowledge, this study is the first to investigate the function approximation capabilities of Looped Transformers. Extending the analysis to multiple layers, varying input lengths, and characterizing optimal memorization capacity presents promising avenues for future research. Beyond expressivity, investigating estimation performance and enhancing training stability constitute important challenges moving forward.

Acknowledgements

This work was supported by JSPS KAKENHI Grant Number JP24H00709.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Bae, S., Fisch, A., Harutyunyan, H., Ji, Z., Kim, S., and Schuster, T. Relaxed recursive transformers: Effective parameter sharing with layer-wise loRA. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=WwpYSOkkCt>.
- Bai, S., Kolter, J. Z., and Koltun, V. Deep equilibrium models. *Advances in neural information processing systems*, 32, 2019.
- Barron, A. R. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.
- Bartlett, P., Maiorov, V., and Meir, R. Almost linear vc dimension bounds for piecewise polynomial networks. *Advances in neural information processing systems*, 11, 1998.
- Csordás, R., Irie, K., Schmidhuber, J., Potts, C., and Manning, C. D. MoEUT: Mixture-of-experts universal transformers. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=ZxVrkm7Bjl>.
- Cybenko, G. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, L. Universal transformers. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HyzdRiR9Y7>.
- Deletang, G., Ruoss, A., Grau-Moya, J., Genewein, T., Wenliang, L. K., Catt, E., Cundy, C., Hutter, M., Legg, S., Veness, J., and Ortega, P. A. Neural networks and the chomsky hierarchy. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=WbxHAzkeQcn>.
- et al., H. T. Llama: Open and efficient foundation language models, 2023. URL <https://arxiv.org/abs/2302.13971>.
- Fan, Y., Du, Y., Ramchandran, K., and Lee, K. Looped transformers for length generalization, 2024a. URL <https://arxiv.org/abs/2409.15647>.
- Fan, Y., Du, Y., Ramchandran, K., and Lee, K. Looped transformers for length generalization, 2024b. URL <https://arxiv.org/abs/2409.15647>.
- Feng, G., Zhang, B., Gu, Y., Ye, H., He, D., and Wang, L. Towards revealing the mystery behind chain of thought: A theoretical perspective. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=qHrADgAdYu>.
- Gandhi, K., Lee, D. H. J., Grand, G., Liu, M., Cheng, W., Sharma, A., and Goodman, N. Stream of search (sos): Learning to search in language. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=2cop2jmQVL>.
- Garg, S., Tsipras, D., Liang, P., and Valiant, G. What can transformers learn in-context? a case study of simple function classes. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho, K. (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=f1NZJ2eOet>.
- Giannou, A., Rajput, S., Sohn, J.-Y., Lee, K., Lee, J. D., and Papailiopoulos, D. Looped transformers as programmable computers. In *Proceedings of the 40th International Conference on Machine Learning*, 2023. URL <https://proceedings.mlr.press/v202/giannou23a.html>.
- Giannou, A., Yang, L., Wang, T., Papailiopoulos, D., and Lee, J. D. How well can transformers emulate in-context newton’s method? *arXiv preprint arXiv:2403.03183*, 2024. URL <https://arxiv.org/abs/2403.03183>.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- Hornik, K., Stinchcombe, M., and White, H. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- Jiang, H. and Li, Q. Approximation rate of the transformer architecture for sequence modeling. *arXiv preprint arXiv:2305.18475*, 2024. URL <https://arxiv.org/abs/2305.18475>.

- Kajitsuka, T. and Sato, I. Are transformers with one layer self-attention using low-rank weight matrices universal approximators? In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=nJnky5K944>.
- Kambhampati, S., Valmeekam, K., Guan, L., Verma, M., Stechly, K., Bhambri, S., Saldyt, L. P., and Murthy, A. B. Position: LLMs can’t plan, but can help planning in LLM-modulo frameworks. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=Th8JPEmH4z>.
- Kim, J., Kim, M., and Mozafari, B. Provable memorization capacity of transformers. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=8JCg5xJCTPR>.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1eA7AEtvS>.
- Loshchilov, I. and Hutter, F. Fixing weight decay regularization in adam, 2018. URL <https://openreview.net/forum?id=rk6qdGgCZ>.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Byj72udxe>.
- Merrill, W. and Sabharwal, A. The parallelism tradeoff: Limitations of log-precision transformers. *Transactions of the Association for Computational Linguistics*, 11:531–545, 2023.
- Peebles, W. and Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- Pope, R., Douglas, S., Chowdhery, A., Devlin, J., Bradbury, J., Levskaya, A., Heek, J., Xiao, K., Agrawal, S., and Dean, J. Efficiently scaling transformer inference, 2022. URL <https://arxiv.org/abs/2211.05102>.
- Radcliffe, D. G. 3 million sudoku puzzles with ratings, 2020.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
- Saunshi, N., Dikkala, N., Li, Z., Kumar, S., and Reddi, S. J. Reasoning with latent thoughts: On the power of looped transformers. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=din0lGfZFd>.
- Takakura, S. and Suzuki, T. Approximation and estimation ability of transformers for sequence-to-sequence functions with infinite dimensional input. In *Proceedings of the 40th International Conference on Machine Learning*, 2023. URL <https://proceedings.mlr.press/v202/takakura23a.html>.
- Takase, S. and Kiyono, S. Lessons on parameter sharing across layers in transformers. *arXiv preprint arXiv:2104.06022*, 2021.
- Team, G. Gemma 2: Improving open language models at a practical size, 2024. URL <https://arxiv.org/abs/2408.00118>.
- Vardi, G., Yehudai, G., and Shamir, O. On the optimal memorization power of reLU neural networks. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=MkTPtnjeYTV>.
- Vaswani, A., Shazeer, N. M., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In *Neural Information Processing Systems*, 2017.
- Wang, M. and E, W. Understanding the expressive power and mechanisms of transformer for sequence modeling. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=0o7Rd5jngV>.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., brian ichter, Xia, F., Chi, E. H., Le, Q. V., and Zhou, D. Chain of thought prompting elicits reasoning in large language models. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho, K. (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=_VjQlMeSB_J.
- Yang, L., Lee, K., Nowak, R. D., and Papailiopoulos, D. Looped transformers are better at learning learning algorithms. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=HHbRxoDTxE>.
- Yang, Z., Ishay, A., and Lee, J. Learning to solve constraint satisfaction problems with recurrent transformer. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=udNhDCr2KQe>.

- Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., and Narasimhan, K. R. Tree of thoughts: Deliberate problem solving with large language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=5Xc1ecx0lh>.
- Yarotsky, D. Optimal approximation of continuous functions by very deep relu networks. In *Conference on learning theory*, pp. 639–649. PMLR, 2018.
- Yun, C., Bhojanapalli, S., Rawat, A. S., Reddi, S., and Kumar, S. Are transformers universal approximators of sequence-to-sequence functions? In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ByxRM0Ntvr>.
- Zhang, B. and Sennrich, R. Root mean square layer normalization. *Advances in Neural Information Processing Systems*, 32, 2019.
- Zhang, S., Lu, J., and Zhao, H. On enhancing expressive power via compositions of single fixed-size relu network. In *International Conference on Machine Learning*, pp. 41452–41487. PMLR, 2023.

A. Proofs for Theorem 3.6

The main theorem incorporates a restriction on the norm of weights, leading to errors when approximating discontinuous functions, such as step functions with ReLU or hardmax functions with softmax. We first establish the approximation rate assuming that weights can take arbitrary precision real values, as outlined below. Then, we account for the bit complexity of bounded weights to complete the proof of the main Theorem 3.6.

Theorem A.1. *Given a function $f \in \mathcal{F}_{\text{PE}}([0, 1]^{d \times N})$, $r > N$, there exists a Looped Transformer, composed of $\text{TF} : \mathbb{R}^{(17d+9) \times N} \rightarrow \mathbb{R}^{(17d+9) \times N}$ with two heads, head size 1, a width size of $q = 49d + 25$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R}^d \rightarrow \mathbb{R}^{17d+9}$ and $\mathcal{L}_2 : \mathbb{R}^{17d+9} \rightarrow \mathbb{R}^d$ s.t.*

$$\|\mathcal{L}_2 \circ \text{TF}^{or} \circ \mathcal{L}_1 - f\|_{L^p} \leq (Nd)^{\frac{1}{p}} \left(\omega_f^{\text{tok}}(\delta\sqrt{d}) + \omega_f^{\text{cont}}(\delta\sqrt{Nd}) \right) + \omega_f(\delta\sqrt{Nd}) + \mathcal{O}(N^{2/p}\delta^{d/p}), \quad (19)$$

where $\delta = ((r - N)/2)^{-1/((N+1)d+1)}$.

A.1. Proof of Theorem A.1

Proof. Since any continuous function can be approximated by a piecewise constant function with arbitrarily small errors, we approximate $f \in \mathcal{F}_{\text{PE}}([0, 1]^{d \times N})$ with piece-wise constant function $\bar{f} : [0, 1]^{d \times N} \rightarrow \mathbb{R}^{d \times N}$. We choose $\delta^{-1} \in \mathbb{N}$, $\delta^{-1} \geq 2$, determining how finely the input is divided: we divide the input space $[0, 1]^{d \times N}$ into δ -discretized cubes, denoted by $\{Q_{\mathbf{B}}\}$ for $\mathbf{B} \in \{0, 1, \dots, \delta^{-1} - 1\}^{d \times N}$ defined by

$$Q_{\mathbf{B}} := \left\{ \mathbf{X} \in [0, 1]^{d \times N} : \mathbf{X}_{i,n} \in [\mathbf{B}_{i,n}\delta, (\mathbf{B}_{i,n} + 1)\delta), \quad i = 1, 2, \dots, d, n = 1, 2, \dots, N \right\}. \quad (20)$$

Each cube $Q_{\mathbf{B}}$ is associated with a representative point $\hat{\mathbf{X}}_{\mathbf{B}}$, defined as the vertex of $Q_{\mathbf{B}}$ with the minimal ℓ_1 norm. Then, we define the piecewise constant function \bar{f} for $\mathbf{X} \in [0, 1]^{d \times N}$ as

$$\bar{f}(\mathbf{X}) := f(\hat{\mathbf{X}}_{\mathbf{B}}). \quad (21)$$

Since we can bound the error within each cube, we have:

$$\max_{\mathbf{X} \in [0, 1]^{d \times N}} \{\|\bar{f}(\mathbf{X}) - f(\mathbf{X})\|_p\} \leq \omega_f(\delta\sqrt{Nd}). \quad (22)$$

Our construction consists of three steps to approximate \bar{f} , as outlined below.

1. The network, with $\delta^{-1} - 1$ loops, maps the input space $[0, 1]^d$ token-wise to the coordinates $\beta \in \{0, 1, \dots, \delta^{-1} - 1\}^d$ of discretized cubes, and then bijectively maps these coordinates to integers, representing *token IDs* in the set $\{0, 1, \dots, \delta^{-1} - 1\}$, using a δ^{-1} -base system; for example, if $d = 2$ and $\delta^{-1} = 3$, then coordinates $(\beta_1, \beta_2) = (2, 1)$ are mapped to the integer $2 \times 3^1 + 1 \times 3^0 = 7$.
2. The network, with N loops, computes a contextual mapping from the set of N distinct token IDs into the set of *contextual token ID*. *Contextual token IDs* refer to token IDs assigned to each token within the context of a *sequence ID*.
3. The network, with $2\delta^{-(N+1)d} - 1$ loops, approximately maps *contextual token IDs* into the output embeddings of each token in a token-wise manner. To achieve a small approximation error, the network has to be designed so that neighboring IDs correspond to similar output token embeddings. Furthermore, *dummy indices* are used to reduce the error.

The details for each step are provided below.

Step 1. Token-wise Quantization. The input space for each token $\mathbf{x} \in [0, 1]^d$ are divided into δ -discretized cubes denoted by $\{Q_{\beta}\}$ for $\beta \in \{0, 1, \dots, \delta^{-1} - 1\}^d$, defined as

$$Q_{\beta} := \left\{ \mathbf{x} \in [0, 1]^d : \mathbf{x}_i \in [\beta_i\delta, (\beta_i + 1)\delta), \quad \text{for all } i = 1, 2, \dots, d \right\}. \quad (23)$$

By Lemma A.5, there exists a feed-forward layer $\text{FF}^{(1)} : \mathbb{R}^{5d} \rightarrow \mathbb{R}^{5d}$ of width size $q = 7d$, and two affine linear maps $\mathcal{L}_1^{(1)} : \mathbb{R}^d \rightarrow \mathbb{R}^{5d}$ and $\mathcal{L}_2^{(1)} : \mathbb{R}^{5d} \rightarrow \mathbb{R}^d$ such that

$$\mathcal{L}_2'^{(1)} \circ (\text{id} + \text{FF}^{(1)})^{\circ(\delta^{-1}-1)} \circ \mathcal{L}_1^{(1)}(\mathbf{x}) = \boldsymbol{\beta} \quad \text{s.t.} \quad \mathbf{x} \in Q_{\boldsymbol{\beta}}. \quad (24)$$

In addition, we need to bijectively map the d -dimensional vector $\boldsymbol{\beta}$ to an integer *token ID*, denoted by z . We use a δ^{-1} -base system: we define the vector $\mathbf{u}_{(\delta^{-1})} \in \mathbb{R}^d$ as

$$\mathbf{u}_{(\delta^{-1})} := (\delta^{-(d-1)}, \delta^{-(d-2)}, \dots, \delta^{-1}, 1)^\top, \quad (25)$$

and define z as

$$z := \mathbf{u}_{(\delta^{-1})}^\top \boldsymbol{\beta} \in \{0, 1, \dots, \delta^{-d} - 1\}. \quad (26)$$

To implement this, we define an affine linear map $\mathcal{L}_2^{(1)} : \mathbb{R}^{5d} \rightarrow \mathbb{R}$ via

$$\mathcal{L}_2^{(1)}(\mathbf{x}) = \mathbf{u}_{(\delta^{-1})}^\top \mathcal{L}_2'^{(1)}(\mathbf{x}). \quad (27)$$

Thus, we have

$$\left(\mathcal{L}_2^{(1)} \circ (\text{FF}_1^{(1)})^{\circ(\delta^{-1}-1)} \circ \mathcal{L}_1^{(1)}(\mathbf{x}) \right)_n = \mathbf{u}_{(\delta^{-1})}^\top \boldsymbol{\beta} = z \quad \text{s.t.} \quad \mathbf{x} \in Q_{\boldsymbol{\beta}}. \quad (28)$$

We establish an upper bound on the maximum distance in input space between adjacent token IDs to derive the approximation error for the following steps. Define the input cubes corresponding to each token ID z as follows:

$$Q_z := \left\{ \mathbf{x} \in [0, 1]^d : \mathbf{x}_i \in [\beta_i \delta, (\beta_i + 1) \delta] \quad \text{for } i = 1, 2, \dots, d \quad \text{s.t.} \quad z = \mathbf{u}_{(\delta^{-1})}^\top \boldsymbol{\beta} \right\}. \quad (29)$$

Then we have

$$\max_{z, \mathbf{x} \in Q_z, \mathbf{x}' \in Q_{z-1}} \|\mathbf{x} - \mathbf{x}'\|_2 \leq \begin{cases} \delta \sqrt{d}, & \text{if } \beta_d \in \{1, 2, \dots, \delta^{-1} - 1\}, \\ \sqrt{d}, & \text{if } \beta_d = 0. \end{cases} \quad (30)$$

Informally, in this δ^{-1} -based representation, the least significant digit corresponds to the index of the d -th dimension, β_d . As the token ID increments sequentially, the index in the d -th dimension increases as $0, 1, 2, \dots, \delta^{-1} - 1$, while the higher-order digits remain unchanged. Consequently, consecutive token IDs correspond to tokens that are “similar” in the d -dimensional space, with a maximum distance of $\delta \sqrt{d}$. However, when a carry occurs, the higher-order digits may change significantly, leading to cases where tokens that are not adjacent in input space become adjacent in their indices. In such cases, the distance is only bounded by \sqrt{d} .

Step 2. Contextual Mapping. The networks, with N -loops, map the list of N token IDs, denoted by $\mathbf{z} \in \{0, 1, \dots, \delta^{-d} - 1\}^N$, into sequence IDs bijectively. Combined with token IDs, the network computes contextual mapping.

We consider only the case where all N input tokens are distinct, disregarding other cases, as they can be treated as negligible when δ is small. The number of subsets in which at least one of the N tokens is duplicated is given by

$$(\delta^{-d})^N - (\delta^{-d} \cdot (\delta^{-d} - 1) \dots (\delta^{-d} - N + 1)) < \frac{N(N-1)}{2} \delta^{-(N-1)d}, \quad (31)$$

when δ is sufficiently small. The volume of these subsets is bounded by

$$\frac{C \delta^{-(N-1)d}}{\delta^{-Nd}} = \mathcal{O}(N^2 \delta^d). \quad (32)$$

Thus, the error with respect to the L^p norm is bounded by $\mathcal{O}(N^{2/p} \delta^{d/p})$.

Let \mathbb{L}_δ denote the set of N distinct token IDs, *i.e.*

$$\mathbb{L}_\delta := \{ \mathbf{z} \in \{0, 1, \dots, \delta^{-d} - 1\}^N \mid z_i \neq z_j \text{ for all } i \neq j \}. \quad (33)$$

Due to permutation equivariance, we can assume without loss of generality that elements of $\mathbf{z} \in \mathbb{L}_\delta$ are ordered, *i.e.*, $\mathbf{z}_1 > \mathbf{z}_2 > \dots > \mathbf{z}_N$. Define $\mathbf{u}_{(\delta^{-d})} := (\delta^{-(N-1)d}, \dots, \delta^{-d}, 1)^\top$, which satisfy

$$|\mathbf{u}_{(\delta^{-d})}^\top \mathbf{z} - \mathbf{u}_{(\delta^{-d})}^\top \mathbf{z}'| > 1, \quad \text{for any } \mathbf{z}, \mathbf{z}' \in \mathbb{L}_\delta \text{ with } \mathbf{z} \neq \mathbf{z}'. \quad (34)$$

This mapping, $\mathbf{u}_{(\delta^{-d})}^\top \mathbf{z}$, represents \mathbf{z} in a δ^{-d} -base system. Then, we define sequence ID for $\mathbf{z} \in \mathbb{L}_\delta$ as:

$$s(\mathbf{z}) := \mathbf{u}_{(\delta^{-d})}^\top \mathbf{z} = \sum_{n=1}^N \mathbf{z}_n \delta^{-(N-n)d}. \quad (35)$$

By Lemma A.6, there exists a Transformer block $\text{TF}'^{(2)} : \mathbb{R}^{5 \times N} \rightarrow \mathbb{R}^{5 \times N}$ with single head, head size $s = 1$, and width size $q = 3$, and two affine linear maps $\mathcal{L}_1'^{(2)} : \mathbb{R} \rightarrow \mathbb{R}^5$ and $\mathcal{L}_2'^{(2)} : \mathbb{R}^5 \rightarrow \mathbb{R}$ such that

$$\mathcal{L}_2'^{(2)} \circ (\text{TF}'^{(2)})^{\circ N} \circ \mathcal{L}_1'^{(2)}(\mathbf{z}^\top) = s(\mathbf{z}) \cdot \mathbf{1}_N^\top. \quad (36)$$

Furthermore, we have to add *dummy indices* to alleviate the approximation error caused by the looped architecture in Step 3. Recall that $\mathbf{B} \in \{0, 1, \dots, \delta^{-1} - 1\}^{d \times N}$ represents the coordinates of the inputs. Let $\mathbf{Z}_\mathbf{B} \in \{0, 1, \dots, \delta^{-1} - 1\}^{d \times N}$ denote the ordered coordinates of \mathbf{B} where the tokens are ordered by their token IDs, *i.e.*, $(\mathbf{u}_{(\delta^{-1})}^\top \mathbf{Z}_\mathbf{B})_1 > (\mathbf{u}_{(\delta^{-1})}^\top \mathbf{Z}_\mathbf{B})_2 > \dots > (\mathbf{u}_{(\delta^{-1})}^\top \mathbf{Z}_\mathbf{B})_N$, in other words, $\mathbf{z} = \mathbf{u}_{(\delta^{-1})}^\top \mathbf{Z}_\mathbf{B}$. Recall that we consider only the case where all N input tokens are distinct. By redefining the sequence ID of Eq. 35 for \mathbf{B} instead of \mathbf{z} , sequence IDs in δ^{-d} -base can be rewritten in the δ^{-1} -base system as follows:

$$s(\mathbf{B}) := \mathbf{u}_{(\delta^{-d})}^\top (\mathbf{u}_{(\delta^{-1})}^\top \mathbf{Z}_\mathbf{B}) \quad (37)$$

$$= \sum_{i=1}^d \sum_{n=1}^N (\mathbf{Z}_\mathbf{B})_{i,n} \delta^{-((N-n)d + (d-i))}. \quad (38)$$

Then, we define *extended sequence IDs* as:

$$s_{\text{valid}}(\mathbf{B}) := 2s(\mathbf{B}) - (\mathbf{Z}_\mathbf{B})_{d,N} \quad (39)$$

and the *dummy sequence IDs* as:

$$s_{\text{dummy}}^b(\mathbf{B}) := s_{\text{valid}}(\mathbf{B}) + b. \quad (40)$$

for $b \in \{\delta^{-1}, \delta^{-1} + 1, \dots, 2\delta^{-1} - 1\}$. Then, define each set as follows:

$$\begin{aligned} \mathcal{S}_{\text{valid}} &:= \{s_{\text{valid}}(\mathbf{B}) : \mathbf{B} \in \{0, 1, \dots, \delta^{-1} - 1\}^{d \times N}\}, \\ \mathcal{S}_{\text{dummy}} &:= \{s_{\text{dummy}}^b(\mathbf{B}) : \mathbf{B} \in \{0, 1, \dots, \delta^{-1} - 1\}^{d \times N}, b \in \{\delta^{-1}, \delta^{-1} + 1, \dots, 2\delta^{-1} - 1\}\}. \end{aligned}$$

Recalling that $(\mathbf{Z}_\mathbf{B})_{d,N} \in \{0, \dots, \delta^{-1} - 1\}$, we observe that

$$\mathcal{S}_{\text{valid}} \cap \mathcal{S}_{\text{dummy}} = \emptyset, \quad (41)$$

$$\mathcal{S}_{\text{valid}} \cup \mathcal{S}_{\text{dummy}} = \{0, 1, \dots, 2\delta^{-Nd} - 1\}. \quad (42)$$

We define the input cubes for each valid sequence ID $s \in \mathcal{S}_{\text{valid}}$ as follows:

$$\mathcal{Q}_s := \left\{ \mathbf{X} \in [0, 1]^{d \times N} : \mathbf{X} \in \mathcal{Q}_\mathbf{B}, \quad \text{s.t.} \quad s = s'(\mathbf{B}) \right\}. \quad (43)$$

Analogous to Eq. 30, we have

$$\max_{s \in \mathcal{S}_{\text{valid}}, \mathbf{X} \in \mathcal{Q}_s, \mathbf{X}' \in \mathcal{Q}_{s-1}} \|\mathbf{X} - \mathbf{X}'\|_2 \leq \begin{cases} \delta\sqrt{Nd} & \text{if } \mathbf{B}_{d,N} \in \{1, 2, \dots, \delta^{-1} - 1\}, \\ \sqrt{Nd} & \text{if } \mathbf{B}_{d,N} = 0. \end{cases} \quad (44)$$

To implement this, we slightly modified $\text{TF}'^{(2)}$. By Corollary A.7, there exists a Transformer block $\text{TF}^{(2)} : \mathbb{R}^{8 \times N} \rightarrow \mathbb{R}^{8 \times N}$ with two heads, head size $s = 1$, and width size $q = 5$, and two affine linear maps $\mathcal{L}_1^{(2)} : \mathbb{R}^2 \rightarrow \mathbb{R}^8$ and $\mathcal{L}_2^{(2)} : \mathbb{R}^8 \rightarrow \mathbb{R}$ s.t.

$$\mathcal{L}_2^{(2)} \circ (\text{TF}^{(2)})^{\circ N} \circ \mathcal{L}_1^{(2)} \left(\begin{bmatrix} \mathbf{z}^\top \\ \mathbf{Z}_{d,:} \end{bmatrix} \right) = (2s(\mathbf{z}) - \mathbf{Z}_{d,N}) \cdot \mathbf{1}_N^\top. \quad (45)$$

Step 3. Token-wise Function Value Mapping. In Steps 1 and 2, the network receives a token ID and extended sequence ID as input for each token, collectively forming a *contextual token ID*. With $2\delta^{-(N+1)d} - 1$ loops, the network approximately maps *contextual token IDs* to the output token embeddings of the target function.

To construct contextual token IDs, we define a bijective affine linear mapping $\mathcal{L}_0^{(3)} : \mathbb{N}^2 \rightarrow \mathbb{N}$ as follows:

$$\mathcal{L}_0^{(3)}(z, s) := 2z\delta^{-Nd} + s, \quad (46)$$

where z represents a token ID, defined in Eq. 26, and s represents a sequence ID. Recall that $z \in \{0, 1, \dots, \delta^{-d} - 1\}$ and sequence IDs are less than $2\delta^{-Nd}$, so informally, it's as if we are adding another digit, z , as the most significant digit in a δ^{-d} -based system. Define the set of contextual token IDs as:

$$\mathcal{K}_{\text{valid}} := \left\{ \mathcal{L}_0^{(3)}(z, s) : z \in \{0, 1, 2, \dots, \delta^{-d} - 1\}, s \in \mathcal{S}_{\text{valid}} \right\}. \quad (47)$$

and dummy contextual token ID as

$$\mathcal{K}_{\text{dummy}} := \left\{ \mathcal{L}_0^{(3)}(z, s) : z \in \{0, 1, 2, \dots, \delta^{-d} - 1\}, s \in \mathcal{S}_{\text{dummy}} \right\}. \quad (48)$$

From Eq. 41 and Eq. 42, the following holds:

$$\mathcal{K}_{\text{valid}} \cap \mathcal{K}_{\text{dummy}} = \emptyset, \quad (49)$$

$$\mathcal{K} := \mathcal{K}_{\text{valid}} \cup \mathcal{K}_{\text{dummy}} = \{0, 1, \dots, 2\delta^{-(N+1)d} - 1\}. \quad (50)$$

We now define the target output embedding for each ID. Let $\mathbf{y}_k \in \mathbb{R}^d$ denote the contextual token embedding corresponding to each contextual token ID, defined as follows:

$$\mathbf{y}_k := \begin{cases} \bar{f}(\hat{\mathbf{X}}_{\mathcal{B}})_{:,n} & \text{s.t. } \mathcal{L}_0^{(3)}(z_n, s'(\mathcal{B})) = k \\ \text{lin_interp}(\mathbf{y}_{\text{near}^-(k, \mathcal{K}_{\text{valid}})}, \mathbf{y}_{\text{near}^+(k, \mathcal{K}_{\text{valid}})}, k - \text{near}^-(k, \mathcal{K}_{\text{valid}}), \delta^{-1}) & \text{for } k \in \mathcal{K}_{\text{dummy}}, \end{cases} \quad (51)$$

where the nearest functions are defined as

$$\text{near}^+(a, \mathcal{S}) := \arg \min_{b \in \mathcal{S}, b > a} |a - b|, \quad \text{near}^-(a, \mathcal{S}) := \arg \min_{b \in \mathcal{S}, b < a} |a - b|, \quad (52)$$

and a function `lin_interp` is defined by

$$\text{lin_interp}(a, b, t, n) := a + \frac{t}{n}(b - a). \quad (53)$$

The illustration of \mathbf{y} is shown in Fig. 2.

Thanks to our design of \mathcal{K} , the error between neighboring contextual token embeddings can be bounded as follows. There are two types of error: the variation induced by contextual perturbation and the variation induced by token perturbation, or both. The examples of each pattern are shown in Fig. 2, such that

- (1) I draft papers. ; I write papers. (perturbation of token)
- (2) He writes papers. ; Mozart writes music. (perturbation of context)
- (3) He drinks coffee. ; He drinks coffee. (perturbation of both token and context)

Recall that there are two types of adjacency that generally have small errors, with a few instances causing large errors when ‘carryover’ occurs, as stated in Eq. 30 and Eq. 44. The key point of our design of \mathcal{K} is that when a large variation occurs, linear interpolation inevitably takes place to smooth out the steep changes between adjacent indices.

Thus for token ID in Eq. 30, if a small variation of token ID, with same context, in input space, $\delta\sqrt{d}$, occurs, the error $e_k = \|\mathbf{y}_k - \mathbf{y}_{k-1}\|_p$ in the output contextual token embedding can be bounded by the modulus of token continuity as

$$e_k \leq \omega_f^{\text{tok}}(\delta\sqrt{d}). \quad (54)$$

In contrast, if a large variation in token input space, \sqrt{d} , occurs in the token input space, the error can be bounded using linear interpolation with δ^{-1} intermediate points as:

$$e_k \leq \delta\omega_f^{\text{tok}}(\sqrt{d}). \quad (55)$$

The same holds for sequence IDs in Eq. 44. That is, since the variation in context is bounded by sequence variation, the difference in adjacent contextual token IDs caused by perturbations in context is bounded as

$$e_k \leq \omega_f^{\text{cont}}(\delta\sqrt{Nd}), \quad \text{or} \quad e_k \leq \delta\omega_f^{\text{cont}}(\sqrt{Nd}). \quad (56)$$

for $k = 0, 1, \dots, K-1$.

Since that

$$\omega_f^{\text{cont, tok}}(n \cdot t) \leq n \cdot \omega_f^{\text{cont, tok}}(t) \quad (57)$$

for any $n \in \mathbb{N}$ and $t \in [0, \infty)$, with $\delta < 1$, it follows that (note that it holds with opposite inequality due to $\delta < 1$)

$$\delta\omega_f^{\text{cont}}(\sqrt{Nd}) \leq \omega_f^{\text{cont}}(\delta\sqrt{Nd}) \quad \text{and} \quad \delta\omega_f^{\text{cont}}(\sqrt{Nd}) \leq \omega_f^{\text{cont}}(\delta\sqrt{Nd}). \quad (58)$$

Considering the maximum difference when both token and context perturbations occur, we have, with the triangle inequality,

$$\max_{k' \in \mathcal{K}} \|\mathbf{y}_{k'} - \mathbf{y}_{k'-1}\|_p \leq \delta \left(\omega_f^{\text{tok}}(\sqrt{d}) + \omega_f^{\text{cont}}(\sqrt{Nd}) \right) \leq \omega_f^{\text{tok}}(\delta\sqrt{d}) + \omega_f^{\text{cont}}(\delta\sqrt{Nd}). \quad (59)$$

Generally, the following inequality holds, for any vector $\mathbf{x} \in \mathbb{R}^d$,

$$\max_i |\mathbf{x}_i| \leq \|\mathbf{x}\|_p \leq d^{\frac{1}{p}} \max_i |\mathbf{x}_i|. \quad (60)$$

Substituting $\mathbf{x} = \mathbf{y}_k - \mathbf{y}_{k-1}$ into the above inequality results in

$$\max_i |(\mathbf{y}_k - \mathbf{y}_{k-1})_i| \leq \|\mathbf{y}_k - \mathbf{y}_{k-1}\|_p. \quad (61)$$

By Lemma 4.1, there exists a feed-forward layer $\text{FF}^{(3)} : \mathbb{R}^{12d} \rightarrow \mathbb{R}^{12d}$ of width size $18d$ and two affine linear maps $\mathcal{L}_1^{(3)} : \mathbb{R} \rightarrow \mathbb{R}^{12d}$ and $\mathcal{L}_2^{(3)} : \mathbb{R}^{12d} \rightarrow \mathbb{R}^d$ such that

$$\left| \left(\mathcal{L}_2^{(3)} \circ (\text{id} + \text{FF}^{(3)})^{(K-1)} \circ \mathcal{L}_1^{(3)}(k) - \mathbf{y}_k \right)_i \right| \leq \max_{k' \in \mathcal{K}} |(\mathbf{y}_{k'} - \mathbf{y}_{k'-1})_i|, \quad (62)$$

for $i = 1, 2, \dots, d$ and $k = 0, 1, \dots, K-1$, where $K = 2\delta^{-(N+1)d}$.

Let $\tilde{\mathbf{y}}_k \in \mathbb{R}^d$ be defined as $\tilde{\mathbf{y}}_k := \mathcal{L}_2^{(3)} \circ (\text{id} + \text{FF}^{(3)})^{(K-1)} \circ \mathcal{L}_1^{(3)}(k)$. Then we have

$$|(\tilde{\mathbf{y}}_k - \mathbf{y}_k)_i| \leq \max_{k' \in \mathcal{K}} |(\mathbf{y}_{k'} - \mathbf{y}_{k'-1})_i| \quad (63)$$

$$\leq \max_{k' \in \mathcal{K}} \|\mathbf{y}_{k'} - \mathbf{y}_{k'-1}\|_p \quad \text{because of Eq. 61} \quad (64)$$

$$\leq \omega_f^{\text{tok}}(\delta\sqrt{d}) + \omega_f^{\text{cont}}(\delta\sqrt{Nd}) \quad \text{because of Eq. 59,} \quad (65)$$

for $i = 1, 2, \dots, d$ and $k = 0, 1, \dots, K-1$.

Concatenated into a Single Transformer Layer Define the input space for each step as:

$$\mathbf{X}^{(0)} \in \mathbb{R}^{1 \times N}, \quad \mathbf{X}^{(1)} \in \mathbb{R}^{5d \times N}, \quad \mathbf{X}^{(2)} \in \mathbb{R}^{8 \times N}, \quad \mathbf{X}^{(3)} \in \mathbb{R}^{12d \times N}, \quad (66)$$

where $\mathbf{X}^{(0)}$ act as counter. Define $\text{Attn} : \mathbb{R}^{(17d+9) \times N} \rightarrow \mathbb{R}^{(17d+9) \times N}$ with two heads, head size $s = 1$, and width size $q = 5$ via

$$\text{Attn} \left(\begin{bmatrix} \mathbf{X}^{(0)} \\ \mathbf{X}^{(1)} \\ \mathbf{X}^{(2)} \\ \mathbf{X}^{(3)} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{0}_{1 \times N} \\ \mathbf{0}_{5d \times N} \\ \text{Attn}^{(2)}(\mathbf{X}^{(2)}) \\ \mathbf{0}_{12d \times N} \end{bmatrix}, \quad (67)$$

where $\text{Attn}^{(2)}$ denote the self-attention layer of $\text{TF}^{(2)}$ and $\mathbf{0}_{m \times N}$ denote $m \times N$ zero matrix. Let

$$x_0 \in \mathbb{R}, \quad \mathbf{x}_1 \in \mathbb{R}^{5d}, \quad \mathbf{x}_2 \in \mathbb{R}^8, \quad \mathbf{x}_3 \in \mathbb{R}^{12d}$$

denote the token-wise input space. Define $\text{FF} : \mathbb{R}^{17d+9} \rightarrow \mathbb{R}^{17d+9}$, with the impulse function in Proposition A.4, as:

$$\text{FF} \left(\begin{bmatrix} x_0 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix} \right) = \begin{bmatrix} 1 \\ \text{FF}^{(1)}(\mathbf{x}_1) \\ \text{FF}^{(2)}(\mathbf{x}_2) + \text{impulse}_{(\delta^{-1}-1)}(\mathcal{L}'(\mathbf{x}_1), x_0) \\ \text{FF}^{(3)}(\mathbf{x}_3) + \text{impulse}_{(\delta^{-1}+N)}(\mathcal{L}''(\mathbf{x}_2), x_0) \end{bmatrix} \quad (68)$$

where $\text{FF}^{(2)}$ denotes the feed-forward layer of $\text{TF}^{(2)}$, two linear maps $\mathcal{L}' : \mathbb{R}^{5d} \rightarrow \mathbb{R}^8$ and $\mathcal{L}'' : \mathbb{R}^8 \rightarrow \mathbb{R}^{12d}$ are defined respectively as follows:

$$\mathcal{L}'(\mathbf{x}) := \mathcal{L}_1^{(2)} \left(\begin{bmatrix} \mathcal{L}_2^{(1)}(\mathbf{x}) \\ \mathbf{x}_{5d} \end{bmatrix} \right) \quad (69)$$

$$\mathcal{L}''(\mathbf{x}) := \mathcal{L}_1^{(3)} \left(\mathcal{L}_0^{(3)}((\mathbf{x}_2)_1, \mathcal{L}_2^{(2)}(\mathbf{x}_2)) \right) \quad (70)$$

and **impulse** refers to the dimension-wise application of impulse function. Note that x_0 serves the role of a counter.

Each step should always be zero or set to an appropriate initial value at the beginning of the corresponding loop iteration. If the bias term causes it to deviate from this value before the iteration starts, the offset can be subtracted in advance to compensate.

As shown in Proposition A.4, the impulse function requires 2 ReLU functions per dimension and 2 ReLU functions for the corresponding loop iteration. Since the total dimension of \mathbf{x}_2 and \mathbf{x}_3 is $12d + 8$, this results in an additional width size of $24d + 16 + 4 = 24d + 20$. Since the width size is $7d$ for $\text{FF}^{(1)}$, 5 for $\text{TF}^{(2)}$, and $18d$ for $\text{FF}^{(3)}$, resulting in a total width size of $49d + 25$.

Define two affine linear maps $\mathcal{L}_1 : \mathbb{R}^d \rightarrow \mathbb{R}^{17d+9}$ and $\mathcal{L}_2 : \mathbb{R}^{17d+9} \rightarrow \mathbb{R}^d$ via

$$\mathcal{L}_1(\mathbf{x}) = (0, \mathcal{L}_1^{(1)}(\mathbf{x}), \mathbf{0}_{12d+8}^\top)^\top, \quad \mathcal{L}_2((x_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)^\top) = \mathcal{L}_2^{(3)}(\mathbf{x}_3). \quad (71)$$

Thus, the network, consisting of three steps, is defined as:

$$\tilde{f}(\mathbf{X}) := \mathcal{L}_2 \circ \text{TF}^{or} \circ \mathcal{L}_1(\mathbf{X}) \quad (72)$$

where $r = \delta^{-1} + N + 2\delta^{-(N+1)d}$ and $\text{TF} : \mathbb{R}^{(17d+9) \times N} \rightarrow \mathbb{R}^{(17d+9) \times N}$ consists of Attn and FF .

Deriving Approximation Error Generally, the following inequality holds.

$$\sum_{i=1}^n x_i^p \leq \left(\sum_{i=1}^n x_i \right)^p \quad \text{for } x_i \geq 0 \text{ and } p \geq 1. \quad (73)$$

Substituting $x_i = \|\tilde{f}(\mathbf{X}) - f(\mathbf{X})\|_p$ into the above inequality results in

$$\|\tilde{f} - f\|_{L^p} = \left(\int \|\tilde{f}(\mathbf{X}) - f(\mathbf{X})\|_p^p d\mathbf{X} \right)^{1/p} \leq \int \|\tilde{f}(\mathbf{X}) - f(\mathbf{X})\|_p d\mathbf{X}. \quad (74)$$

Also, generally, the following inequality holds, for $\mathbf{x} \in \mathbb{R}^m$,

$$\max_i |\mathbf{x}_i| \leq \|\mathbf{x}\|_p \leq m^{\frac{1}{p}} \max_i |\mathbf{x}_i|. \quad (75)$$

With the triangle inequality, we can bound the approximation error as

$$\|\tilde{f} - f\|_{L^p} \leq \int \|\tilde{f}(\mathbf{X}) - f(\mathbf{X})\|_p d\mathbf{X} \quad \text{because of Eq. 74} \quad (76)$$

$$\leq \int \|\tilde{f}(\mathbf{X}) - \bar{f}(\mathbf{X})\|_p d\mathbf{X} + \int \|\bar{f}(\mathbf{X}) - f(\mathbf{X})\|_p d\mathbf{X} \quad (77)$$

$$\leq (Nd)^{\frac{1}{p}} \max_{k' \in \mathcal{K}} |\tilde{\mathbf{y}}_{k'} - \mathbf{y}_{k'}| + \int \|\bar{f}(\mathbf{X}) - f(\mathbf{X})\|_p d\mathbf{X} + \mathcal{O}(N^{2/p} \delta^{d/p}) \quad \text{because of Eq. 75} \quad (78)$$

$$\leq (Nd)^{\frac{1}{p}} \max_{k' \in \mathcal{K}} |\tilde{\mathbf{y}}_{k'} - \mathbf{y}_{k'}| + \omega_f(\delta \sqrt{Nd}) + \mathcal{O}(N^{2/p} \delta^{d/p}) \quad \text{because of Eq. 22} \quad (79)$$

$$\leq (Nd)^{\frac{1}{p}} (\omega_f^{\text{tok}}(\delta \sqrt{d}) + \omega_f^{\text{cont}}(\delta \sqrt{Nd})) + \omega_f(\delta \sqrt{Nd}) + \mathcal{O}(N^{2/p} \delta^{d/p}) \quad \text{because of Eq. 65.} \quad (80)$$

Then, δ is expressed in terms of the number of loops as:

$$r = \delta^{-1} + N + 2\delta^{-(N+1)d} \Leftrightarrow \delta^{-1} + 2\delta^{-(N+1)d} = r - N \quad (81)$$

$$\Rightarrow \delta^{-1} \cdot 2\delta^{-(N+1)d} \geq r - N \quad (82)$$

$$\Leftrightarrow \delta \leq \left(\frac{r - N}{2} \right)^{-1/((N+1)d+1)}. \quad (83)$$

Thus, we have completed the proof of Theorem A.1. \square

A.2. Proof of Theorem 3.6

Extending the construction in Theorem A.1, we then account for the boundedness of the weights and bit complexity.

Proof. Due to the use of bounded weights to approximate discontinuous functions, there inevitably exist regions where quantization errors arise in Step 1 of our construction. We define these regions, for $0 < \epsilon < \delta$, as

$$\Omega([0, 1]^{d \times N}, \delta^{-1}, \epsilon) := \left\{ \mathbf{X} \in [0, 1]^{d \times N} : \exists i, n \quad \text{such that} \quad \mathbf{X}_{i,n} \in \bigcup_{k=1}^{\delta^{-1}} (k\delta - \epsilon, k\delta) \right\}. \quad (84)$$

That is, Ω consists of all inputs for which at least one coordinate $\mathbf{X}_{i,n}$ lies within an ϵ -neighborhood of a quantization discontinuity point $k\delta$ for some $k \in \{1, \dots, \delta^{-1}\}$.

Outside the region Ω , the quantization function is piecewise constant and can be precisely approximated using bounded weights. By the construction in Lemma A.5, the maximum magnitude of weights required is proportional to $1/\epsilon$. We now estimate the Lebesgue measure of Ω . For each coordinate (i, n) , the set

$$\bigcup_{k=1}^{\delta^{-1}} (k\delta - \epsilon, k\delta)$$

is a union of δ^{-1} intervals, each of length ϵ , so the total measure of this set is $\delta^{-1}\epsilon$. Since the input \mathbf{X} has $d \times N$ coordinates, and since we consider the event that at least one coordinate lies in a bad region, we apply the union bound for measures to obtain $\text{meas}(\Omega) \leq dN \times \delta^{-1}\epsilon$. Substituting the bound on $\text{meas}(\Omega)$, and using the fact that the maximum magnitude of the weights M satisfies $M = 1/\epsilon$, we have $\text{meas}(\Omega) \leq dN \times \delta^{-1}M^{-1}$. Consequently, the L^p -norm of the quantization error is bounded by $\mathcal{O}\left((M\delta)^{-1}dN\right)^{1/p}$.

When replacing hardmax with softmax, it is required that the error in step 2 remains sufficiently small so that it does not affect step 3. Specifically, a step function is used, in Lemma 4.1 for step 3, to map the index, defined for $\epsilon > 0$ as

$$\text{step}_\epsilon(x) = \begin{cases} 1 & \text{if } x \geq 0, \\ 0 & \text{if } x < 0 - \epsilon, \end{cases} \quad (85)$$

The construction of Lemma 4.1 use this function for indices $k \in \mathcal{K}$ obtained from step 2 as $\text{step}_\epsilon(k-i)$ for $i = 0, 1 \dots K-1$. Thus, the error in step 2 does not affect step 3 if the perturbed indices, denoted by \tilde{k} , satisfies

$$\text{step}_\epsilon(\tilde{k} + 1 - i) = \text{step}_\epsilon(k - i) \quad \text{for all } i = 0, 1 \dots K-1. \quad (86)$$

To estimate the error introduced by replacing hardmax with softmax in step 2, we revisit the construction of Lemma A.6. Specifically, we extract and reformulate the key components necessary for this estimation. In particular, we consider a simplified form of the attention computation in Eq. 135, denoted by $g_H : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$, which is defined as

$$g_H(\mathbf{v}, \mathbf{a}) := \mathbf{v}_{\arg \max_i \mathbf{a}_i} \quad (87)$$

When hardmax in Eq. 135 is replaced with softmax, the function can be expressed as

$$g_S(\mathbf{v}, \mathbf{a}, \lambda) := \sum_{i=1}^N \frac{\exp(\lambda \mathbf{a}_i)}{\sum_{j=1}^N \exp(\lambda \mathbf{a}_j)} \mathbf{v}_i, \quad (88)$$

where $\lambda > 0$. Note that $\lim_{\lambda \rightarrow \infty} g_S(\mathbf{v}, \mathbf{a}, \lambda) = g_H(\mathbf{v}, \mathbf{a})$. According to the construction of Lemma A.6, such as Eq. 147, the domains, denoted by $\mathbf{v} \in \mathbb{L}_\delta$ and $\mathbf{a} \in \mathbb{A}_\delta$, are restricted on

$$\mathbb{L}_\delta := \{\mathbf{z} \in \{0, 1, \dots, \delta^{-d} - 1\}^N \mid \mathbf{z}_i \neq \mathbf{z}_j \text{ for all } i \neq j\}, \quad (89)$$

$$\mathbb{A}_\delta := \{\mathbf{a} \in \mathbb{R}^N \mid \mathbf{a}_i \in \{0, 1, \dots, \delta^{-d} - 1\} \text{ or } \mathbf{a}_i < 0, \quad \text{if } \mathbf{a}_i, \mathbf{a}_j \geq 0, \text{ then } \mathbf{a}_i \neq \mathbf{a}_j \text{ for all } i \neq j, \quad \exists i \text{ s.t. } \mathbf{a}_i > 0\}. \quad (90)$$

We impose the following two additional assumptions on $\mathbf{v} \in \mathbb{L}_\delta$ and $\mathbf{a} \in \mathbb{A}_\delta$

$$n = \arg \max_{1 \leq i \leq N} \mathbf{a}_i, \quad v_n = \max_{1 \leq i \leq N} v_i. \quad (91)$$

Under these assumptions, for any finite $\lambda > 0$ we have

$$0 < g_S(\mathbf{v}, \mathbf{a}, \lambda) < g_H(\mathbf{v}, \mathbf{a}) \quad (92)$$

and

$$g_S(\mathbf{v}, \mathbf{a}, \lambda) \geq \frac{\exp(\lambda \mathbf{a}_n)}{\sum_{j=1}^N \exp(\lambda \mathbf{a}_j)} \mathbf{v}_n, \quad (93)$$

Thus we have

$$g_H(\mathbf{v}, \mathbf{a}) - g_S(\mathbf{v}, \mathbf{a}, \lambda) = \mathbf{v}_n - g_S(\mathbf{v}, \mathbf{a}, \lambda) \quad (94)$$

$$\leq \frac{\sum_{j=1}^N \exp(\lambda \mathbf{a}_j) - \exp(\lambda \mathbf{a}_n)}{\sum_{j=1}^N \exp(\lambda \mathbf{a}_j)} \mathbf{v}_n \quad (95)$$

$$= \frac{\sum_{j \neq n} \exp(\lambda \mathbf{a}_j)}{\sum_{j \neq n} \exp(\lambda \mathbf{a}_j) + \exp(\lambda \mathbf{a}_n)} \mathbf{v}_n \quad (96)$$

$$\leq \frac{\sum_{j \neq n} \exp(\lambda \mathbf{a}_j)}{\exp(\lambda \mathbf{a}_n)} \mathbf{v}_n \quad (97)$$

$$= \mathbf{v}_n \sum_{j \neq n} \exp(\lambda(\mathbf{a}_j - \mathbf{a}_n)) \quad (98)$$

$$\leq \mathbf{v}_n N \exp(-\lambda) \quad \text{because } \mathbf{a}_j - \mathbf{a}_n \leq -1 \quad (99)$$

$$\leq \delta^{-d} N \exp(-\lambda). \quad (100)$$

Thus, if we aim to bound the error within $\epsilon > 0$, λ must satisfies

$$\delta^{-d} N \exp(-\lambda) < \epsilon \Leftrightarrow \lambda > \log\left(\frac{\delta^{-d} N}{\epsilon}\right). \quad (101)$$

From Eq. 46, the error of contextual ID $k - \tilde{k}$ can be bounded in terms of the error of $\epsilon = g_H(\mathbf{v}, \mathbf{a}) - g_S(\mathbf{v}, \mathbf{a}, \lambda)$ as:

$$k - \tilde{k} \leq (\mathbf{u}_{(\delta^{-d})})^\top (\epsilon \mathbf{1}_N) \leq \epsilon N \delta^{-(N-1)d} \quad (102)$$

where $\mathbf{u}_{(\delta^{-d})} := (\delta^{-(N-1)d}, \dots, \delta^{-d}, 1)^\top$ and $\mathbf{1}_N \in \mathbb{R}^N$ denotes the all-ones vector.

Since Eq. 86 holds if $0 < k - \tilde{k} < 1$, it follows that

$$\lambda > \log \left((\delta^{-d} N) \cdot (N \delta^{-(N-1)d}) \right) \Rightarrow \epsilon N \delta^{-(N-1)d} < 1, \quad (103)$$

and Eq. 86 holds. Thus, the bit complexity of λ can be bounded by

$$\mathcal{O}(\log \log(\delta^{-Nd} N^2)), \quad (104)$$

while ensuring that no error occurs in step 3 when using the softmax function instead of the hardmax function.

The bit complexity at each step of the computation can be analyzed as follows. In Step 1 and Step 2, the bit complexity is bounded by $\mathcal{O}(\log \delta^{-1})$, reflecting the cost of maintaining precision within error δ . In contrast, Step 3 incurs a significantly higher cost, with a bit complexity bounded by $\mathcal{O}(2\delta^{-(N+1)d})$, due to the need to evaluate higher-order terms accurately.

As a result, the overall bit complexity of the Looped Transformer is dominated by Step 3 and can be bounded by

$$\mathcal{O}(\max\{\log \log(\delta^{-Nd} N^2), 2\delta^{-(N+1)d}\}) = \mathcal{O}(\delta^{-(N+1)d}) \quad (\delta \rightarrow 0). \quad (105)$$

With Theorem A.1, the proof of Theorem 3.6 is completed. \square

A.3. Approximating Discontinuous Piecewise Functions

We define three utility functions using ReLU activations. Since the target function is discontinuous, there are negligible ‘trifling regions’ introduced by the bounded weights of the networks.

Proposition A.2 (Rectangular function). *Given $t_1, t_2 \in \mathbb{R}$, there exist four ReLU functions that can approximate the following function, denote by $\text{rect}_t : \mathbb{R} \rightarrow \mathbb{R}$, for $0 < \epsilon < t_2 - t_1$, such that:*

$$\text{rect}_{(t_1, t_2, \epsilon)}(x) = \begin{cases} 1 & \text{if } x \in [t_1, t_2 - \epsilon], \\ 0 & \text{if } x \in (-\infty, t_1 - \epsilon] \cup [t_2, \infty), \end{cases} \quad (106)$$

which is represented by

$$\text{rect}_{(t_1, t_2, \epsilon)}(x) = \sigma_R\left(\frac{x - t_1 + \epsilon}{\epsilon}\right) - \sigma_R\left(\frac{x - t_1}{\epsilon}\right) + \sigma_R\left(\frac{-x + t_2}{\epsilon}\right) - \sigma_R\left(\frac{-x + t_2 - \epsilon}{\epsilon}\right) - 1. \quad (107)$$

Note that maximum magnitude of the weights is $1/\epsilon$ and ‘trifling regions’ are $(t - \epsilon, t) \cup (t + 1 - \epsilon, t + 1)$.

Proposition A.3 (Step function). *There exist four ReLU functions that can approximate the following function, denote by $\text{step} : \mathbb{R} \rightarrow \mathbb{R}$, for $\epsilon > 0$, such that:*

$$\text{step}_\epsilon(x) = \begin{cases} 1 & \text{if } x \geq 0, \\ 0 & \text{if } x < 0 - \epsilon, \end{cases} \quad (108)$$

which is represented via


$$\text{step}_\epsilon(x) = \sigma_R\left(\frac{x}{\epsilon} + 1\right) - \sigma_R\left(\frac{x}{\epsilon}\right). \quad (109)$$

Proposition A.4 (Impulse function). *Given $\theta \in \mathbb{N}$, there exist four ReLU functions that can approximate the following function, denote by $\text{impulse}_\theta : \mathbb{R} \times \mathbb{N} \rightarrow \mathbb{R}$ for $x \in [-M, M]$ and $t \in \mathbb{N}$, such that:*

$$\text{impulse}_\theta(x, t) = \begin{cases} x & \text{if } t = \theta, \\ 0 & \text{otherwise.} \end{cases} \quad (110)$$

which is represented via

$$\begin{aligned} \text{impulse}_\theta(x, t) := & \sigma_R(x + 2M(t - \theta + 1/2)) - 2M\sigma_R(t - \theta + 1/2) \\ & - \sigma_R(x + 2M(t - \theta - 1/2)) + 2M\sigma_R(t - \theta - 1/2). \end{aligned} \quad (111)$$


 Figure 3. An illustration of $h_k(x)$.

A.4. Step 1. Token-wise Quantization

We aim to construct quantization function $\mathbf{g} : [0, 1]^d \rightarrow \{0, 1, \dots, \delta^{-1}\}^d$, for $\epsilon < \delta$, for each dimension as

$$\mathbf{g}(\mathbf{x}) = (g(x_1), g(x_2), \dots, g(x_d))^T, \quad \text{where } g(x) = k \quad \text{if } x \in [k\delta, (k+1)\delta - \epsilon] \text{ for } k = 0, \dots, \delta^{-1} - 1. \quad (112)$$

This function $g : \mathbb{R} \rightarrow \mathbb{R}$ can be expressed as

$$g(x) = \sum_{i=0}^{n-1} i \cdot \text{rect}_{(i\delta, (i+1)\delta - \epsilon)}(x) \quad (113)$$

for any $n \in \mathbb{N}$ and $x \in \mathbb{R}$. The illustration of $h_k(x) := k \cdot \text{rect}_{(k\delta, (k+1)\delta - \epsilon)}(x)$ is shown in Fig 3. The key idea is that $h_k(x)$ can be represented with a single function h in the form of $h(kx, k^2, k)$. Lemma A.5 implement $h(kx, k^2, k)$ with a feed-forward layer and perform the summation through a skip connection.

Lemma A.5. *Given any $\delta^{-1} \in \mathbb{N}$ and $x \in \mathbb{R}^d$, there exists a feed-forward layer $\text{FF} : \mathbb{R}^{5d} \rightarrow \mathbb{R}^{5d}$ of width size $q = 7d$ with the maximum magnitude of the weights $1/\epsilon$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R}^d \rightarrow \mathbb{R}^{5d}$ and $\mathcal{L}_2 : \mathbb{R}^{5d} \rightarrow \mathbb{R}^d$ s.t.*

$$\left(\mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(\delta^{-1}-1)} \circ \mathcal{L}_1(\mathbf{x}) \right)_i = k \quad \text{if } x \in [k\delta, (k+1)\delta - \epsilon], \quad \text{for } k = 0, \dots, \delta^{-1} - 1 \quad (114)$$

for any $i = 1, 2, \dots, d$.

Proof. On the basis of Proposition A.2, define function $h_k(x) = k \cdot \text{rect}_k(x)$ via

$$h_k(x) := \sigma_R\left(\frac{k}{\epsilon}(x - k\delta + \epsilon)\right) - \sigma_R\left(\frac{k}{\epsilon}(x - k\delta)\right) + \sigma_R\left(\frac{k}{\epsilon}(-x + k\delta + 1)\right) - \sigma_R\left(\frac{k}{\epsilon}(-x + k\delta + 1 - \epsilon)\right) - k, \quad (115)$$

which satisfies

$$h_k(x) = \begin{cases} k & \text{if } x \in [k\delta, (k+1)\delta - \epsilon], \\ 0 & \text{if } x \in (-\infty, k\delta - \epsilon] \cup [(k+1)\delta, \infty), \end{cases} \quad (116)$$

For any $x \in [k\delta, (k+1)\delta - \epsilon]$ for $k = 0, 1, \dots, \delta^{-1} - 1$, we have

$$\sum_{i=0}^{\delta^{-1}-1} h_i(x) = h_k(x) = k. \quad (117)$$

Define function $h : \mathbb{R}^3 \rightarrow \mathbb{R}$ to represent h_k via

$$h(kx, k^2, k) := \sigma_R\left(\frac{kx}{\epsilon} - \frac{k^2\delta}{\epsilon} + k\right) - \sigma_R\left(\frac{kx}{\epsilon} - \frac{k^2\delta}{\epsilon}\right) + \sigma_R\left(-\frac{kx}{\epsilon} + \frac{k^2\delta}{\epsilon} + \frac{k}{\epsilon}\right) - \sigma_R\left(-\frac{kx}{\epsilon} + \frac{k^2\delta}{\epsilon} + k\frac{1-\epsilon}{\epsilon}\right) - \sigma_R(k) = h_k(x). \quad (118)$$

Define ξ_k as

$$\xi_k = \left(kx, k^2, k, x, \sum_{i=0}^{k-1} h_i(x) \right)^\top. \quad (119)$$

Then, construct a feed-forward layer $\text{FF} : \mathbb{R}^5 \rightarrow \mathbb{R}^5$ with a skip connection such that

$$(\text{id} + \text{FF})(\xi_k) = (\text{id} + \text{FF}) \left(\left(kx, k^2, k, x, \sum_{i=0}^{k-1} h_i(x) \right)^\top \right) \quad (120)$$

$$= \left((k+1)x, (k+1)^2, k+1, x, \sum_{i=0}^k h_i(x) \right)^\top \quad (121)$$

$$= \xi_{k+1}. \quad (122)$$

via

$$\begin{aligned} (\text{id} + \text{FF}) \left(\begin{bmatrix} kx \\ k^2 \\ k \\ x \\ \sum_{i=0}^{k-1} h_i(x) \end{bmatrix} \right) &= \begin{bmatrix} kx \\ k^2 \\ k \\ x \\ \sum_{i=0}^{k-1} h_i(x) \end{bmatrix} + \\ &\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 1 & -1 & -1 \end{bmatrix} \sigma_R \left(\begin{bmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ \frac{1}{\epsilon} & -\frac{\delta}{\epsilon} & 1 & 0 & 0 \\ \frac{1}{\epsilon} & -\frac{\delta}{\epsilon} & 0 & 0 & 0 \\ -\frac{1}{\epsilon} & \frac{\delta}{\epsilon} & \frac{1}{\epsilon} & 0 & 0 \\ -\frac{1}{\epsilon} & \frac{\delta}{\epsilon} & \frac{1-\epsilon}{\epsilon} & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} kx \\ k^2 \\ k \\ x \\ \sum_{i=0}^{k-1} h_i(x) \end{bmatrix} \right) + \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} \end{aligned} \quad (123)$$

$$= \begin{bmatrix} kx \\ k^2 \\ k \\ x \\ \sum_{i=0}^{k-1} h_i(x) \end{bmatrix} + \begin{bmatrix} x \\ 2k+1 \\ 1 \\ 0 \\ h_k(x) \end{bmatrix} \quad (124)$$

$$= \begin{bmatrix} kx + x \\ k^2 + 2k + 1 \\ k + 1 \\ x \\ \sum_{i=0}^{k-1} h_i(x) + h_k(x) \end{bmatrix} \quad (125)$$

$$= \begin{bmatrix} (k+1)x \\ (k+1)^2 \\ k+1 \\ x \\ \sum_{i=0}^k h_i(x) \end{bmatrix}. \quad (126)$$

Then, define two affine linear maps $\mathcal{L}_1 : \mathbb{R}^1 \rightarrow \mathbb{R}^5$ and $\mathcal{L}_2 : \mathbb{R}^5 \rightarrow \mathbb{R}^1$ by

$$\mathcal{L}_1(x) := (0, 0, 0, x, 0), \quad \mathcal{L}_2(x_1, x_2, x_3, x_4, x_5) := x_5. \quad (127)$$

Thus, we have

$$\mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(\delta^{-1}-1)} \circ \mathcal{L}_1(x) = \mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(\delta^{-1}-1)}(\xi_0) \quad (128)$$

$$= \mathcal{L}_2(\xi_{\delta^{-1}}) \quad (129)$$

$$= \sum_{i=0}^{\delta^{-1}-1} h_i(x). \quad (130)$$

For d -dimensional inputs, we need d -times more parameters. \square

A.5. Step 2. Contextual Mapping

The network takes token IDs as inputs, denoted by $\mathbf{z} \in \{0, 1, \dots, \delta^{-d} - 1\}^N$. We only consider cases where all token IDs are distinct. The network maps token IDs into a sequence ID using the inner product with the vector $\mathbf{u} \in \mathbb{R}^N$ defined as $\mathbf{u} := (\delta^{-(N-1)d}, \delta^{-(N-2)d}, \dots, \delta^{-d}, 1)^\top$ i.e.

$$s(\mathbf{z}) := \mathbf{u}^\top \mathbf{z}. \quad (131)$$

Due to permutation equivariance, we can assume without loss of generality that elements of $\mathbf{z} \in \mathbb{L}_\delta$ are ordered, i.e., $z_1 > z_2 > \dots > z_N$. Then the map s satisfies

$$|\mathbf{u}^\top \mathbf{z} - \mathbf{u}^\top \mathbf{z}'| > 1, \quad \text{if } \mathbf{z} \neq \mathbf{z}'. \quad (132)$$

In other words, s represent \mathbf{z} in δ^{-d} -base system. The network computes $\mathbf{u}^\top \mathbf{z}$ in the form of $\sum_{i=1}^N \delta^{-(N-i)d} z_i$. The network computes $\mathbf{s}^{(k)} := \sum_{i=1}^k \delta^{-(k-i)d} z_i$ in each loop, and after N -loops, it outputs $\mathbf{s}^{(N)} = \mathbf{u}^\top \mathbf{z}$. To implement this, the self-attention layer selects z_k in the k -th loop iteration. We design the key, query, and value weights to select the maximum token ID. The feed-forward layer post-processes the token ID in such a way that if it is selected, then it is replaced with a negative value to prevent selection in subsequent iterations, i.e., the post-processed token IDs for the k -th loop are

$$\mathbf{z}_i^{(k)} = z \quad \text{s.t.} \quad \begin{cases} z < 0 & \text{if } i \leq k, \\ z = z_i & \text{otherwise.} \end{cases}$$

We focus on self-attention layers that employ the hardmax function.

Lemma A.6. Consider the set of distinct indices corresponding to the d -dimensional δ -discretized regions of N tokens, i.e.

$$\mathbb{L}_\delta := \{\mathbf{z} \in \{0, 1, \dots, \delta^{-d} - 1\}^N \mid z_i \neq z_j \text{ for all } i \neq j\}. \quad (133)$$

There exists a function $s : \mathbb{R}^N \rightarrow \mathbb{R}$ composed of Transformer block $\text{TF} : \mathbb{R}^{5 \times N} \rightarrow \mathbb{R}^{5 \times N}$ with the hardmax function, single head, head size $s = 1$, and width size $q = 3$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^5$ and $\mathcal{L}_2 : \mathbb{R}^5 \rightarrow \mathbb{R}$, such that

$$\mathcal{L}_2 \circ \text{TF}^{\circ N} \circ \mathcal{L}_1(\mathbf{z}^\top) = (\mathbf{u}^\top \mathbf{z}) \cdot \mathbf{1}_N^\top,$$

for any $\mathbf{z} \in \mathbb{L}_\delta$, where $\mathbf{u} := (\delta^{-(N-1)d}, \delta^{-(N-2)d}, \dots, \delta^{-d}, 1)^\top$.

Proof. Due to permutation equivariance, we can assume without loss of generality that elements of $\mathbf{z} \in \mathbb{L}_\delta$ are ordered, i.e., $z_1 > z_2 > \dots > z_N$. Define $\mathbf{u} \in \mathbb{R}^N$ as $\mathbf{u} := (\delta^{-(N-1)d}, \dots, \delta^{-d}, 1)^\top$, which satisfy

$$|\mathbf{u}^\top \mathbf{z} - \mathbf{u}^\top \mathbf{z}'| > 1, \quad \text{if } \mathbf{z} \neq \mathbf{z}' \text{ for any } \mathbf{z}, \mathbf{z}' \in \mathbb{L}_\delta. \quad (134)$$

We construct Transformer block $\text{TF} : \mathbb{R}^{5 \times N} \rightarrow \mathbb{R}^{5 \times N}$ with single head and head size $s = 1$ such that, for any $\mathbf{z} \in \mathbb{L}_\delta$,

$$\text{TF}^{\circ N} \left(\begin{bmatrix} \mathbf{z}^\top \\ \mathbf{z}^\top \\ \mathbf{1}_N^\top \\ \mathbf{0}_N^\top \\ \mathbf{0}_N^\top \end{bmatrix} \right) = \begin{bmatrix} \mathbf{0}_N^\top \\ \mathbf{0}_N^\top \\ \mathbf{1}_N^\top \\ \mathbf{0}_N^\top \\ (\mathbf{u}^\top \mathbf{z}) \cdot \mathbf{1}_N^\top \end{bmatrix}. \quad (135)$$

where $\mathbf{0}_N, \mathbf{1}_N \in \mathbb{R}^N$ denote the all-zero and all-one vectors, respectively. For $\mathbf{z} \in \mathbb{L}_\delta$, define two series for $k = 0, \dots, N$ as:

$$\mathbf{z}_i^{(k)} := z \quad \text{s.t.} \quad \begin{cases} z < 0 & \text{if } i \leq k, \\ z = z_i & \text{otherwise,} \end{cases} \quad \text{for } i = 1, \dots, N, \quad \in \mathbb{R}^N, \quad (136)$$

$$\mathbf{s}^{(k)} := \sum_{i=0}^k \delta^{-(k-i)d} \mathbf{z}_i \quad \in \mathbb{R}. \quad (137)$$

While $\mathbf{z}^{(k)}$ is not uniquely determined, any vector that satisfies the conditions is accepted as $\mathbf{z}^{(k)}$. The series $\mathbf{s}^{(k)}$ satisfies

$$\mathbf{s}^{(k+1)} = \sum_{i=1}^{k+1} \delta^{-(k+1-i)d} \mathbf{z}_i \quad (138)$$

$$= \left(\sum_{i=1}^{k+1-1} \delta^{-(k+1-i)d} \mathbf{z}_i \right) + \mathbf{z}_k \quad (139)$$

$$= \left(\sum_{i=0}^k \delta^{-d} \cdot \delta^{-(k-i)d} \mathbf{z}_i \right) + \mathbf{z}_{k+1} \quad (140)$$

$$= \delta^{-d} \cdot \mathbf{s}^k + \mathbf{z}_{k+1}, \quad (141)$$

for $k = 0, \dots, N-1$. Note that $\mathbf{s}^{(N)} = \mathbf{u}^\top \mathbf{z}$. Define a single-head self-attention $\text{Attn} : \mathbb{R}^{5 \times N} \rightarrow \mathbb{R}^{5 \times N}$ such that

$$\text{Attn} \left(\begin{bmatrix} \mathbf{v}^\top \\ \mathbf{a}^\top \\ \mathbf{1}_N^\top \\ *^\top \\ *^\top \end{bmatrix} \right) = \begin{bmatrix} \mathbf{0}_N^\top \\ \mathbf{0}_N^\top \\ \mathbf{0}_N^\top \\ (\mathbf{v}_{\arg \max_j \mathbf{a}_j}) \cdot \mathbf{1}_N^\top \\ \mathbf{0}_N^\top \end{bmatrix}, \quad (142)$$

where $\mathbf{v}, \mathbf{a} \in \mathbb{R}^N$, and $*$ $\in \mathbb{R}^N$ denotes arbitrary vectors, via the weight parameters

$$\mathbf{W}_O = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \quad \mathbf{W}_V = [1 \ 0 \ 0 \ 0 \ 0], \quad \mathbf{W}_K = [0 \ 1 \ 0 \ 0 \ 0], \quad \mathbf{W}_Q = [0 \ 0 \ 1 \ 0 \ 0]. \quad (143)$$

Define $\text{FF} : \mathbb{R}^5 \rightarrow \mathbb{R}^5$ of width size $q = 3$ via:

$$\text{FF} \left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} \right) = \begin{bmatrix} 0 & 0 & 0 \\ -M & 0 & 0 \\ 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 1 & \delta^{-d} - 1 \end{bmatrix} \sigma_R \left(\begin{bmatrix} 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} + \begin{bmatrix} 0 \\ \epsilon \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \quad (144)$$

$$= \begin{bmatrix} 0 \\ -M\sigma_R(x_2 - x_4 + \epsilon) \\ 0 \\ -\sigma_R(x_4) \\ (\delta^{-d} - 1)\sigma_R(x_5) + \sigma_R(x_4) \end{bmatrix}, \quad (145)$$

where $0 < \epsilon < 1$ and $M > \frac{\delta^{-d}-1}{\epsilon}$.

For $x_2 \in \{0, 1, \dots, \delta^{-d} - 1\} \cup \{x \mid x \leq 0\}$ and $x_4 \in \{0, 1, \dots, \delta^{-d} - 1\}$ with $x_4 \geq x_2$, we have

$$x_2 - M\sigma_R(x_2 - x_4 + \epsilon) = z, \quad s.t. \quad \begin{cases} z = x_1 & \text{if } x_4 > x_2, \\ z < 0 & \text{if } x_4 = x_2. \end{cases} \quad (146)$$

This post-processes the token ID in such a way that if it is selected, then it is replaced with a negative value *i.e.*

$$\begin{aligned} \mathbf{z}_i^{(k)} - M\sigma_R(\mathbf{z}_i^{(k)} - \mathbf{z}_k + \epsilon) &= z \quad s.t. \quad \begin{cases} z < 0 & \text{if } i \leq k+1, \\ z = \mathbf{z}_i & \text{otherwise,} \end{cases} \\ &= \mathbf{z}_i^{(k+1)} \quad \text{for } i = 1, \dots, N. \end{aligned} \quad (147)$$

We confirm that the Transformer block $\text{TF} : \mathbb{R}^{5 \times N} \rightarrow \mathbb{R}^{5 \times N}$, composed of Attn and FF, satisfies, for $k = 0, \dots, N-1$,

$$\text{TF} \left(\begin{bmatrix} \mathbf{z}^\top \\ (\mathbf{z}^{(k)})^\top \\ \mathbf{1}_n^\top \\ \mathbf{0}_n^\top \\ (\mathbf{s}^{(k)})^\top \end{bmatrix} \right) = (\text{id} + \mathbf{FF}) \circ (\text{id} + \text{Attn}) \left(\begin{bmatrix} \mathbf{z}^\top \\ (\mathbf{z}^{(k)})^\top \\ \mathbf{1}_n^\top \\ \mathbf{0}_n^\top \\ (\mathbf{s}^{(k)})^\top \end{bmatrix} \right) \quad (148)$$

$$= (\text{id} + \mathbf{FF}) \left(\begin{bmatrix} \mathbf{z}^\top \\ (\mathbf{z}^{(k)})^\top \\ \mathbf{1}_n^\top \\ \mathbf{z}_{k+1} \cdot \mathbf{1}_N^\top \\ (\mathbf{s}^{(k)})^\top \end{bmatrix} \right) \quad (149)$$

$$= \begin{bmatrix} \mathbf{z}^\top \\ (\mathbf{z}^{(k)})^\top \\ \mathbf{1}_n^\top \\ \mathbf{z}_{k+1} \cdot \mathbf{1}_N^\top \\ (\mathbf{s}^{(k)})^\top \end{bmatrix} + \begin{bmatrix} \mathbf{0}_n^\top \\ -M \sigma_R((\mathbf{z}^{(k)})^\top - \mathbf{z}_k \cdot \mathbf{1}_N^\top + \epsilon \mathbf{1}_N^\top) \\ \mathbf{0}_n^\top \\ -\mathbf{z}_{k+1} \cdot \mathbf{1}_N^\top \\ (\delta^{-d} - 1)(\mathbf{s}^{(k)})^\top + \sigma_R(\mathbf{z}_{k+1} \mathbf{1}_N^\top) \end{bmatrix} \quad (150)$$

$$= \begin{bmatrix} \mathbf{z}^\top \\ (\mathbf{z}^{(k+1)})^\top \\ \mathbf{1}_n^\top \\ \mathbf{0}_n^\top \\ \delta^{-d}(\mathbf{s}^{(k)})^\top + \mathbf{z}_{k+1} \mathbf{1}_N^\top \end{bmatrix} \quad \text{because of Eq. 147} \quad (151)$$

$$= \begin{bmatrix} \mathbf{z}^\top \\ (\mathbf{z}^{(k+1)})^\top \\ \mathbf{1}_n^\top \\ \mathbf{0}_n^\top \\ (\mathbf{s}^{(k+1)})^\top \end{bmatrix} \quad \text{because of Eq. 141.} \quad (152)$$

Define two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^5$ and $\mathcal{L}_2 : \mathbb{R}^5 \rightarrow \mathbb{R}$ via $\mathcal{L}_1(x) := (x, x, 1, 0, 0)$ and $\mathcal{L}_2(x_1, x_2, x_3, x_4, x_5) := x_5$ respectively. Thus, we have

$$\mathcal{L}_2 \circ \text{TF}^{\circ N} \circ \mathcal{L}_1(\mathbf{z}^\top) = (\mathbf{u}^\top \mathbf{z}) \cdot \mathbf{1}_N^\top.$$

□

Combining the sequence ID with token ID, the network computes contextual mapping.

Corollary A.7. *There exists a Transformer block $\text{TF}_2 : \mathbb{R}^{8 \times N} \rightarrow \mathbb{R}^{8 \times N}$ with the hardmax function, two heads, head size $s = 1$, and width size $q = 5$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R}^2 \rightarrow \mathbb{R}^8$ and $\mathcal{L}_2 : \mathbb{R}^8 \rightarrow \mathbb{R}$ s.t.*

$$\mathcal{L}_2 \circ \text{TF}_2^{\circ N} \circ \mathcal{L}_1 \left(\begin{bmatrix} \mathbf{z}^\top \\ \mathbf{Z}_{d,:} \end{bmatrix} \right) = (2\mathbf{u}^\top \mathbf{z} - \mathbf{Z}_{d,N}) \cdot \mathbf{1}_N^\top \quad \text{for any } \mathbf{z} \in \mathbb{L}_\delta.$$

Proof. Define a self-attention $\text{Attn} : \mathbb{R}^{8 \times N} \rightarrow \mathbb{R}^{8 \times N}$ such that

$$\text{Attn} \left(\begin{bmatrix} \mathbf{v}^\top \\ \mathbf{a}^\top \\ \mathbf{1}_N^\top \\ *^\top \\ *^\top \\ \mathbf{Z}_{d,:}^\top \\ *^\top \\ *^\top \end{bmatrix} \right) = \begin{bmatrix} \mathbf{0}_N^\top \\ \mathbf{0}_N^\top \\ \mathbf{0}_N^\top \\ (\mathbf{v}_{\arg \max_j \mathbf{a}_j}) \cdot \mathbf{1}_N^\top \\ \mathbf{0}_N^\top \\ \mathbf{0}_N^\top \\ \mathbf{Z}_{d,N} \cdot \mathbf{1}_N^\top \\ \mathbf{0}_N^\top \end{bmatrix}, \quad (153)$$

where $\mathbf{v}, \mathbf{a} \in \mathbb{R}^N$, and $\ast \in \mathbb{R}^N$ denotes arbitrary vectors, via the weight parameters

$$\mathbf{W}_O^1 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{W}_V^1 = [1 \ 0 \ \dots \ 0], \quad \mathbf{W}_K^1 = [0 \ 1 \ 0 \ \dots \ 0], \quad \mathbf{W}_Q^1 = [0 \ 0 \ 1 \ 0 \ \dots \ 0] \quad (154)$$

and

$$\mathbf{W}_O^2 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \quad \mathbf{W}_V^2, \mathbf{W}_K^2 = [0 \ \dots \ 0 \ 1 \ 0 \ 0 \ 0], \quad \mathbf{W}_Q^2 = [0 \ 0 \ 1 \ 0 \ \dots \ 0], \quad (155)$$

where \dots denotes a sequence of zeros.

Define $\text{FF} : \mathbb{R}^8 \rightarrow \mathbb{R}^8$ of width size $q = 5$ via:

$$\text{FF} \left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \end{bmatrix} \right) = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ -M & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 1 & \delta^{-d} - 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 & -1 \end{bmatrix} \sigma_R \left(\begin{bmatrix} 0 & 1 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \end{bmatrix} + \begin{bmatrix} 0 \\ \epsilon \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right) \quad (156)$$

$$= \begin{bmatrix} 0 \\ -M\sigma_R(x_2 - x_4 + \epsilon) \\ 0 \\ -\sigma_R(x_4) \\ (\delta^{-d} - 1)\sigma_R(x_5) + \sigma_R(x_4) \\ 0 \\ 0 \\ -\sigma_R(x_7) \\ \sigma_R(x_7) - \sigma_R(x_8) \end{bmatrix}, \quad (157)$$

where $0 < \epsilon < 1$ and $M > \frac{\delta^{-d}-1}{\epsilon}$. Then, we define two affine linear maps $\mathcal{L}_1 : \mathbb{R}^2 \rightarrow \mathbb{R}^8$ and $\mathcal{L}_2 : \mathbb{R}^8 \rightarrow \mathbb{R}$ respectively:

$$\mathcal{L}_1(x_1, x_2) := (x_1, x_1, 1, 0, 0, x_2, 0, 0), \quad \mathcal{L}_2(x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8) := 2x_5 - x_6 \quad (158)$$

From Lemma A.6, the corollary holds for this construction. \square

A.6. Step 3. Function Value Mapping with Bit Extraction

We employ a bit extraction technique (Bartlett et al., 1998), as used (Zhang et al., 2023) for weight-tied ReLU networks, to approximately memorize the label set. Given $K \in \mathbb{N}$ input indices $k = 1, 2, \dots, K$ with associated values $y_1, y_2, \dots, y_K \in \mathbb{R}$, the network approximately memorizes the differences $y_i - y_{i-1}$ using their base-2 representation. Since the binary

representation is limited to $\{0, 1\}$, the differences $y_i - y_{i-1}$ must be rescaled by a factor $\epsilon := \max_i |y_i - y_{i-1}|$ as

$$a_i = \lfloor \frac{y_i}{\epsilon} \rfloor, \quad (159)$$

where $\lfloor x \rfloor := \max\{n : n \leq x, n \in \mathbb{Z}\}$. Then, the difference $b_i = a_i - a_{i-1}$ satisfies $b_i \in \{-1, 0, 1\}$ and can be represented using two binary values $c_i, d_i \in \{0, 1\}$ as follows:

$$b_i = c_i - d_i, \quad (160)$$

and we have

$$a_k = a_0 + \sum_{i=0}^k b_i = a_0 + \sum_{i=0}^k c_i - \sum_{i=0}^k d_i \quad \text{for } k = 0, 1, \dots, n-1. \quad (161)$$

Lemma A.9 and Lemma 4.1 show that $\sum_{i=0}^k c_i$ and $\sum_{i=0}^k d_i$ can be realized by composition of single feed-forward layer. Thus, the network can approximate y_i using ϵa_i , denoted as \tilde{y}_i , with the following accuracy:

$$|\tilde{y}_i - y_i| = |\epsilon \lfloor \frac{y_i}{\epsilon} \rfloor - \frac{y_i}{\epsilon}| = \epsilon |\lfloor \frac{y_i}{\epsilon} \rfloor - \frac{y_i}{\epsilon}| \leq \epsilon. \quad (162)$$

For a d -dimensional input-output pair, we construct the networks for each dimension *i.e.*

$$\tilde{\mathbf{y}} = (\tilde{y}^1, \tilde{y}^2, \dots, \tilde{y}^d) \quad (163)$$

The basic strategy of our lemma and proof follows Lemma D.1 from Zhang et al. (2023), as shown below and Proposition 3.2. However, their result cannot be directly applied here, as it requires depth-2 networks.

Proposition A.8 (Lemma D.1 in Zhang et al. (2023)). *Given any $r \in \mathbb{N}^+$, there exists a function $\text{FF} : \mathbb{R}^{3d} \rightarrow \mathbb{R}^{3d}$ with width 8 and depth 2, utilizing two affine linear maps $\mathcal{L}_1 : \mathbb{R}^2 \rightarrow \mathbb{R}^5$ and $\mathcal{L}_2 : \mathbb{R}^5 \rightarrow \mathbb{R}$, such that for any $\theta_1, \theta_2, \dots, \theta_r \in \{0, 1\}$, the following holds:*

$$\mathcal{L}_2 \circ \text{FF}^{\circ r} \circ \mathcal{L}_1(k, \text{bin } 0.\theta_1\theta_2 \dots \theta_r) = \sum_{\ell=1}^k \theta_\ell \quad \text{for } k = 1, 2, \dots, r, \quad (164)$$

where $\text{bin } 0.\theta_1\theta_2 \dots \theta_r$ denote the binary representation of $\theta = \sum_{l=1}^r \theta_l 2^{-l}$.

We found that the *loop unrolling* technique allows us to reduce the number of layers from 2 to 1 by replacing $x^{k+1} = \text{ReLU}(\text{ReLU}(x'^k))$ with $(x^{k+1}, x'^k) = \text{ReLU}(x'^k, x^k)$. Although our method makes the weights dependent on $\theta_1, \theta_2, \dots, \theta_r$, this does not present an issue for our construction in function approximation. Specifically, $\theta_1, \theta_2, \dots, \theta_r$ is fixed for each target function, and the role of the network is to learn the weights tailored to that single function.

Lemma A.9. *Given $\theta_1, \theta_2, \dots, \theta_r \in \{0, 1\}$ for some $r \in \mathbb{N}^+$, there exists a feed-forward layer $\text{FF} : \mathbb{R}^6 \rightarrow \mathbb{R}^6$ of width size 9 and two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^6$ and $\mathcal{L}_2 : \mathbb{R}^6 \rightarrow \mathbb{R}$ s.t.*

$$\mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ r} \circ \mathcal{L}_1(k) = \sum_{l=1}^k \theta_l \quad \text{for } k = 1, 2, \dots, r, \quad (165)$$

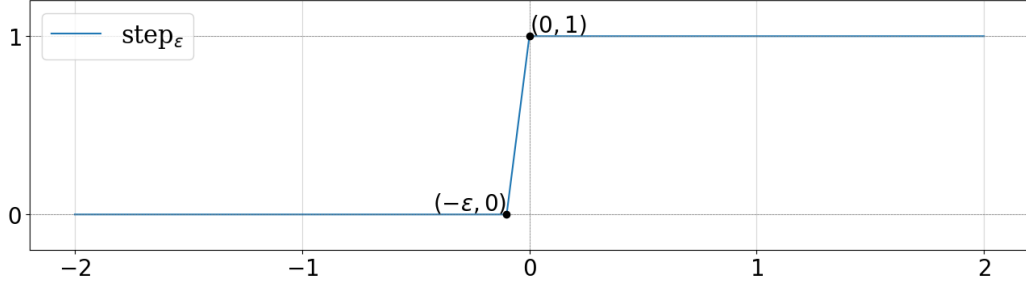
where the bit complexity is bounded by $\mathcal{O}(r)$.

Proof. From Proposition A.3, we have a function $\text{step}_\epsilon(x)$, for $\epsilon > 0$, defined by

$$\text{step}_\epsilon(x) = \sigma_R\left(\frac{x}{\epsilon} + 1\right) - \sigma_R\left(\frac{x}{\epsilon}\right), \quad (166)$$

as shown in Figure 4, and it satisfies

$$\text{step}_\epsilon(x) = \begin{cases} 1 & \text{if } x \geq 0, \\ 0 & \text{if } x \leq 0 - \epsilon. \end{cases} \quad (167)$$


 Figure 4. An illustration of $\text{step}_\epsilon(x)$.

Define β_l for $l = 0, 1, \dots, r$ as

$$\beta_l = \text{bin } 0.\theta_l \dots \theta_r, \quad (168)$$

where $\text{bin } 0.\theta_l \dots \theta_r$ denote the binary representation of $\theta = \sum_{i=l}^r \theta_i 2^{-i}$ and $\beta_0 := 0$. If we set $\epsilon < 2^{-r}$, it follows that

$$\theta_l = \text{step}_\epsilon(\text{bin } 0.\theta_l \dots \theta_r - \frac{1}{2}) \quad (169)$$

$$= \text{step}_\epsilon(\beta_l - \frac{1}{2}), \quad (170)$$

implying, for $l = 0, 1, \dots, r-1$,

$$\beta_{l+1} = 2\beta_l - \theta_l \quad (171)$$

$$= 2\beta_l - \text{step}_\epsilon(\beta_l - \frac{1}{2}). \quad (172)$$

For all $l = 1, \dots, r$, since the product xy satisfies

$$xy = \max\{0, x + y - 1\} \quad (173)$$

$$= \sigma_R(x + y - 1), \quad (174)$$

for $x, y \in \{0, 1\}$, it follows that

$$\sum_{l=1}^k \theta_l = \sum_{l=1}^k \theta_l + \sum_{l=k+1}^r 0 \quad (175)$$

$$= \sum_{l=1}^r \theta_l \cdot \text{step}_\epsilon(k-l) \quad (176)$$

$$= \sum_{l=1}^r \sigma_R(\theta_l + \text{step}_\epsilon(k-l) - 1) \quad (177)$$

$$= \sum_{l=1}^r \sigma_R(\text{step}_\epsilon(\beta_l - \frac{1}{2}) + \text{step}_\epsilon(k-l) - 1). \quad (178)$$

To compute the right-hand side, we require two nested ReLU functions. By employing loop unrolling, we precompute $\mathcal{T}(\beta_l - \frac{1}{2})$ and $\mathcal{T}(k-l)$ in the previous iterations, reducing the requirement to a single layer.

Define ξ_l for $l = 0, 1, \dots, r-1$ as

$$\xi_l = \left(k-l, \beta_l, \beta_{l+1}, \text{step}_\epsilon(\beta_l - \frac{1}{2}), \text{step}_\epsilon(k-l), \text{sum}(l) \right)^\top, \quad (179)$$

where $\text{sum}(l) := \sum_{i=1}^l \sigma_R(\text{step}_\epsilon(\beta_i - \frac{1}{2}) + \text{step}_\epsilon(k-i) - 1).$

Note that we have β_{l+1} in the l -th loop to precompute $\text{step}_\epsilon(\beta_{l+1} - \frac{1}{2})$ and $\text{step}_\epsilon((k-(l+1)))$ for the $(l+1)$ -th loop.

Define $\text{FF} : \mathbb{R}^6 \rightarrow \mathbb{R}^6$ with a width size of 9 such that

$$(\text{id} + \text{FF})(\xi_l) = (\text{id} + \text{FF}) \left(\begin{bmatrix} k-l \\ \beta_l \\ \beta_{l+1} \\ \text{step}_\epsilon(\beta_l - \frac{1}{2}) \\ \text{step}_\epsilon(k-l) \\ \text{sum}(l) \end{bmatrix} \right) \quad (180)$$

$$= \begin{bmatrix} k-l \\ \beta_l \\ \beta_{l+1} \\ \text{step}_\epsilon(\beta_l - \frac{1}{2}) \\ \text{step}_\epsilon(k-l) \\ \text{sum}(l) \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 & -1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} k-l \\ \beta_l \\ \beta_{l+1} \\ \text{step}_\epsilon(\beta_l - \frac{1}{2}) \\ \text{step}_\epsilon(k-l) \\ \text{sum}(l) \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ -1/(2\epsilon) + 1 \\ -1/(2\epsilon) \\ 0 \\ -1/\epsilon + 1 \\ -1/\epsilon \\ -1 \end{bmatrix} + \begin{bmatrix} -1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (181)$$

$$= \begin{bmatrix} k-l \\ \beta_l \\ \beta_{l+1} \\ \text{step}_\epsilon(\beta_l - \frac{1}{2}) \\ \text{step}_\epsilon(k-l) \\ \text{sum}(l) \end{bmatrix} + \begin{bmatrix} -1 \\ \sigma_R(\beta_l) - \sigma_R(\text{step}_\epsilon(\beta_l - \frac{1}{2})) \\ \sigma_R(\beta_{l+1}) - (\sigma_R(\frac{\beta_{l+1}-1/2}{\epsilon} + 1) - \sigma_R(\frac{\beta_{l+1}-1/2}{\epsilon})) \\ -\sigma_R(\text{step}_\epsilon(\beta_l - \frac{1}{2})) + \sigma_R(\frac{\beta_{l+1}-1/2}{\epsilon} + 1) - \sigma_R(\frac{\beta_{l+1}-1/2}{\epsilon}) \\ -\sigma_R(\text{step}_\epsilon(k-l)) + \sigma_R(\frac{k-(l+1)}{\epsilon} + 1) - \sigma_R(\frac{k-(l+1)}{\epsilon}) \\ \sigma_R(\text{step}_\epsilon(k-l) + \text{step}_\epsilon(\beta_l - \frac{1}{2}) - 1) \end{bmatrix} \quad (182)$$

$$= \begin{bmatrix} k-(l+1) \\ 2\beta_l - \text{step}_\epsilon(\beta_l - \frac{1}{2}) \\ 2\beta_{l+1} - \text{step}_\epsilon(\beta_{l+1} - \frac{1}{2}) \\ \text{step}_\epsilon(\beta_{l+1} - \frac{1}{2}) \\ \text{step}_\epsilon((k-(l+1))) \\ \text{sum}(l+1) \end{bmatrix} = \begin{bmatrix} k-(l+1) \\ \beta_{l+1} \\ \beta_{l+2} \\ \text{step}_\epsilon(\beta_{l+1} - \frac{1}{2}) \\ \text{step}_\epsilon((k-(l+1))) \\ \text{sum}(l+1) \end{bmatrix} = \xi_{l+1}, \quad (183)$$

Define $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^6$ and $\mathcal{L}_2 : \mathbb{R}^6 \rightarrow \mathbb{R}$ via

$$\mathcal{L}_1(k) := (k, \beta_0, \beta_1, 0, 0, 0)^\top = \xi_0, \quad \mathcal{L}_2(x_1, x_2, x_3, x_4, x_5, x_6) := x_6, \quad (184)$$

respectively. The lemma holds for this construction. \square

Then, we prove Lemma 4.1 with Lemma A.9.

Lemma 4.1. *Given $\mathbf{y}_k \in \mathbb{R}^d$ for $k = 0, 1, \dots, K-1$ with*

$$|(\mathbf{y}_k - \mathbf{y}_{k-1})_i| \leq \varepsilon_i \quad \text{for } i = 1, \dots, d,$$

there exists a feed-forward layer $\text{FF} : \mathbb{R}^{12d} \rightarrow \mathbb{R}^{12d}$ with a width size of $18d$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^{12d}$ and $\mathcal{L}_2 : \mathbb{R}^{12d} \rightarrow \mathbb{R}^d$ s.t.

$$|(\mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(K-1)} \circ \mathcal{L}_1(k) - \mathbf{y}_k)_i| \leq \varepsilon_i, \quad (18)$$

for $i = 1, \dots, d$ and $k = 0, 1, \dots, K-1$.

Proof. We prove this for the case where $d = 1$, considering $y_k \in \mathbb{R}$ for $k = 0, 1, \dots, K - 1$. Define

$$a_k = \lfloor \frac{y_k}{\varepsilon} \rfloor \quad \text{for } k = 0, 1, \dots, K - 1, \quad (185)$$

where $\lfloor x \rfloor = \max\{n : n \leq x, n \in \mathbb{Z}\}$ and set

$$b_k = a_k - a_{k-1} \quad \text{for } k = 1, 2, \dots, K - 1. \quad (186)$$

Since $b_k \in \{-1, 0, 1\}$, there exist $c_k \in \{0, 1\}$ and $d_k \in \{0, 1\}$ such that

$$b_k = c_k - d_k \quad \text{for } k = 1, 2, \dots, K - 1. \quad (187)$$

Thus, we have

$$a_k = a_0 + \sum_{i=1}^k c_i - \sum_{i=1}^k d_i \quad \text{for any } k \in \{1, 2, \dots, K - 1\} \quad (188)$$

From Lemma A.9, there exist $\text{FF}^{(c)}, \text{FF}^{(d)} : \mathbb{R}^6 \rightarrow \mathbb{R}^6$ of width size 9 and affine linear maps $\mathcal{L}'_2 : \mathbb{R}^6 \rightarrow \mathbb{R}$ and $\mathcal{L}_1^{(c)}, \mathcal{L}_1^{(d)} : \mathbb{R} \rightarrow \mathbb{R}^6$ s.t.

$$\mathcal{L}'_2 \circ (\text{id} + \text{FF}^{(c)})^{\circ(m-1)} \circ \mathcal{L}_1^{(c)}(k) = \sum_{i=1}^k c_i, \quad \mathcal{L}'_2 \circ (\text{id} + \text{FF}^{(d)})^{\circ(m-1)} \circ \mathcal{L}_1^{(d)}(k) = \sum_{i=1}^k d_i, \quad (189)$$

for $k = 0, 1, \dots, K - 1$. Then, define $\text{FF} : \mathbb{R}^{12} \rightarrow \mathbb{R}^{12}$ of width size 18, for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^6$,

$$\text{FF} \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \right) := \begin{bmatrix} \text{FF}^{(c)}(\mathbf{x}) \\ \text{FF}^{(d)}(\mathbf{y}) \end{bmatrix}. \quad (190)$$

Define $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^{12}$ and $\mathcal{L}_2 : \mathbb{R}^{12} \rightarrow \mathbb{R}$ as

$$\mathcal{L}_1(x) := \begin{bmatrix} \mathcal{L}_1^{(c)}(x) \\ \mathcal{L}_1^{(d)}(x) \end{bmatrix}, \quad \mathcal{L}_2 \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \right) := \epsilon(a_0 + \mathcal{L}'_2(\mathbf{x}) - \mathcal{L}'_2(\mathbf{y})). \quad (191)$$

We can confirm that

$$\mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(K-1)} \circ \mathcal{L}_1(k) \quad (192)$$

$$= \mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(K-1)} \left(\begin{bmatrix} \mathcal{L}_1^{(c)}(k) \\ \mathcal{L}_1^{(d)}(k) \end{bmatrix} \right) \quad (193)$$

$$= \mathcal{L}_2 \left(\begin{bmatrix} (\text{id} + \text{FF}^{(c)})^{\circ(K-1)} \circ \mathcal{L}_1^{(c)}(k) \\ (\text{id} + \text{FF}^{(d)})^{\circ(K-1)} \circ \mathcal{L}_1^{(d)}(k) \end{bmatrix} \right) \quad (194)$$

$$= \epsilon(a_0 + \sum_{i=1}^k c_i - \sum_{i=1}^k d_i) = \epsilon a_k. \quad (195)$$

Thus we have

$$|\mathcal{L}_2 \circ (\text{id} + \text{FF})^{\circ(K-1)} \circ \mathcal{L}_1(k) - y_k| = |\epsilon a_k - y_k| \leq \varepsilon. \quad (196)$$

For d -dimensional inputs, we need d -times more parameters. \square

B. Role of Time-dependent Scaling Parameters

We show that time-dependent scaling parameters overcome the limitations inherent to the looped architecture and eliminate the dependence of the modulus of continuity. We use the architecture defined in Section 4 as:

$$\text{FF}(\mathbf{x}) \rightarrow \boldsymbol{\eta}(t) \odot \text{FF}(\mathbf{x}) \quad \text{for the } t\text{-th loops,} \quad (197)$$

The following lemma demonstrates that time-dependent scaling parameters can exactly map indices to output vectors.

Theorem 4.2. *Given $\mathbf{y}_k \in \mathbb{R}^d$ for $k = 0, 1, \dots, K-1$, there exists a feed-forward layer $\text{FF} : \mathbb{R}^{4d} \rightarrow \mathbb{R}^{4d}$ with a width size of $6d$, $\boldsymbol{\eta}(t) \in \mathbb{R}^{4d}$ for $t = 1, \dots, K-1$, and two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^{4d}$ and $\mathcal{L}_2 : \mathbb{R}^{4d} \rightarrow \mathbb{R}^d$ s.t.*

$$|(\mathcal{L}_2 \circ (\text{id} + \boldsymbol{\eta} \odot \text{FF})^{\circ(K-1)} \circ \mathcal{L}_1(k) - \mathbf{y}_k)_i| = 0,$$

for $i = 1, \dots, d$ and $k = 0, 1, \dots, K-1$.

Proof. We consider the case when $d = 1$, where $y_k \in \mathbb{R}$ for $k = 0, 1, \dots, K-1$. We update y_k as follows:

$$y_k \rightarrow y_k + \epsilon, \quad (198)$$

where ϵ is chosen such that none of the y_l values are zero.

Next, we define $\eta(l) \in \mathbb{R}^4$ as:

$$\eta(l) := (0, 1, \frac{y_l}{y_{l-1}} - 1, \frac{y_l}{y_{l-1}})^\top \quad \text{for } l = 1, 2, \dots, K-1. \quad (199)$$

By Proposition A.4, we have, $x \in [-M, M]$ and $t \in \mathbb{N}$,

$$\begin{aligned} \text{impulse}_0(x, t) &= \sigma_R(x + 2M(t + 1/2)) - 2M\sigma_R(t + 1/2) \\ &\quad - \sigma_R(x + 2M(t - 1/2)) + 2M\sigma_R(t - 1/2) \end{aligned} \quad (200)$$

$$= \begin{cases} x & \text{if } t = 0, \\ 0 & \text{otherwise,} \end{cases} \quad (201)$$

where $M > \max_{k \in \{0, 1, \dots, K-1\}} y_k$.

Let $k \in \{0, 1, \dots, K-1\}$ be the input index that specifies which y_k to extract. Define

$$s(l) := \sum_{i=0}^l \text{impulse}_0(y_i, k - i), \quad (202)$$

for $l = 0, 1, \dots, K-1$, which satisfies

$$s(K-1) = y_k. \quad (203)$$

Define $\boldsymbol{\xi}_l \in \mathbb{R}^4$ via

$$\boldsymbol{\xi}_l := (k, k - l - 1, y_l, s(l))^\top. \quad (204)$$

for $l = 0, 1, \dots, K-1$.

Then, define $\text{FF} : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ of width size $q = 6$ via:

$$(\text{id} + \eta(l) \odot \text{FF})(\xi_{l-1}) = \xi_{l-1} + \eta(l) \odot \left(\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & -2M & 2M \end{bmatrix} \sigma_R \left(\begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 2M & 1 & 0 \\ 0 & 2M & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \xi_{l-1} + \begin{bmatrix} 0 \\ -1 \\ M \\ -M \\ 1/2 \\ -1/2 \end{bmatrix} \right) + \begin{bmatrix} 0 \\ -1 \\ 0 \\ 0 \end{bmatrix} \right) \quad (205)$$

$$= \xi_{l-1} + \begin{bmatrix} 0 \\ 1 \\ \frac{y_l}{y_{l-1}} - 1 \\ \frac{y_l}{y_{l-1}} \end{bmatrix} \odot \begin{bmatrix} 0 \\ -1 \\ \sigma_R(y_{l-1}) - \sigma_R(-y_{l-1}) \\ \left(\sigma_R(y_{l-1} + 2M((k-l) + 1/2)) \right. \\ \left. - 2M\sigma_R((k-l) + 1/2) \right. \\ \left. - \sigma_R(y_{l-1} + 2M(k-l-1/2)) + 2M\sigma_R(k-l-1/2) \right) \end{bmatrix} \quad (206)$$

$$= \begin{bmatrix} k \\ k-l \\ y_{l-1} \\ s(l-1) \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ \frac{y_l}{y_{l-1}} - 1 \\ \frac{y_l}{y_{l-1}} \end{bmatrix} \odot \begin{bmatrix} 0 \\ -1 \\ y_{l-1} \\ \text{impulse}_0(y_{l-1}, k-l) \end{bmatrix} \quad (207)$$

$$= \begin{bmatrix} k \\ k-l-1 \\ y_l \\ s(l) \end{bmatrix} = \xi_l. \quad (208)$$

for $l = 1, 2, \dots, K-1$. Thus we have

$$(\text{id} + \eta(K-1) \odot \text{FF}) \circ \dots \circ (\text{id} + \eta(1) \odot \text{FF})(\xi_0) = \xi_{K-1} \quad (209)$$

Then, define two affine linear maps $\mathcal{L}_1 : \mathbb{R} \rightarrow \mathbb{R}^4$ and $\mathcal{L}_2 : \mathbb{R}^4 \rightarrow \mathbb{R}$ by

$$\mathcal{L}_1(x) := (k, k, y_0, 0), \quad \mathcal{L}_2(x_1, x_2, x_3, x_4) := x_4 - \epsilon. \quad (210)$$

We can extend this to d -dimensional input by using d times more parameters, by applying the above to each dimension \square

C. Details of Experiments

This appendix section provides additional details on the experiments for each task.

C.1. Reasoning Tasks

C.1.1. PROBLEM SETTINGS

Longest Common Subsequence (LCS) is the longest common to a given set of sequences. We use problems with input lengths of 60 and 100. Two sequences are sampled uniformly from the alphabet.

Edit Distance (ED) problem, also known as Levenshtein distance, is to find the minimum cost required to change one sequence into the other. We adopted the problem setting and data generation approach from [Feng et al. \(2023\)](#), but applied larger input lengths. The costs for insertion, deletion, and replacement were set to 2, 2, and 3, respectively. They generate instances of the edit distance problem as shown in Algorithm 1. The first string is randomly selected, while the second is generated in two ways: (1) a random string yielding a large edit distance, and (2) a corrupted copy of the first string, resulting in a small edit distance.

Algorithm 1 ED Data Generation from Feng et al. (2023)

```

1: Input: Length of the First String  $n$ 
2: Input: Alphabet  $V = \{a, b, \dots, z\}$ 
3: Output: Sequence  $s_1, s_2$  Sample  $t$  uniformly from  $\{3, 4, \dots, 10\}$   $T \leftarrow$  Sample  $t$  letters from  $V$   $s_1 \leftarrow$  Sample  $n$  letters uniformly from  $T$  Sample  $p$  uniformly from  $[0, 1]$ 
4: if  $p < 0.4$  then
5:   Sample  $l$  uniformly from  $\{n - 3, n - 2, \dots, n + 2\}$ 
6:    $s_2 \leftarrow$  Sample  $l$  letters uniformly from  $T$ 
7: else
8:   while  $\text{len}(s_2)$  not in  $[n - 3, n + 2]$  do
9:      $s_2 \leftarrow s_1$ 
10:    for  $i \leftarrow 1$  to  $n$  do
11:      Sample  $p$  uniformly from  $\{0, 1, \dots, \text{len}(s_2) - 1\}$ 
12:      Sample  $l$  uniformly from  $T$ 
13:      Randomly conduct one of the followings: pop  $s_2[p]$ , substitute  $s_2[p]$  with  $l$ , insert  $l$  into  $s_2[p]$ 
14:    end for
15:  end while
16: end if
    
```

Sudoku We use the dataset from Radcliffe (2020), which contains over 3 million Sudoku puzzles. The puzzles are flattened, with 0 representing blank grids. The input sequence is formatted as:

100503700603008090000009800010000000876100000000006000000000007080907604700060312.

Countdown To generate the dataset, we randomly sampled the target and input numbers for each instance. Pairs that have no solution were excluded. For tokenization, we assigned unique tokens to each number and symbol. The target sequence is represented as:

58 84 48 62 96 62 - 58 = 4 48 / 4 = 12 84 + 12 = 96 .

The model learns to predict the target sequence from the input:

58 84 48 62 96 0 0 0 0 0 0 0 0 0 0 0 0 0 0 .

C.1.2. TRAINING AND TEST ACCURACY CORRELATION FOR EDIT DISTANCE

Given that our study focuses on function approximation capabilities, one might question whether it is appropriate to rely on test evaluations, which are influenced by generalization. Here, we confirm that there is a strong correlation between training and test results, validating this approach. Figure 5 demonstrates a strong positive correlation between training and test accuracy, enabling the evaluation of approximation power through test accuracy.




Figure 5. Training and test accuracy for the edit distance task with a sequence length of 60.

C.1.3. MODEL AND TRAINING CONFIGURATION

We used Looped Transformers with 4 attention heads and a 256-dimensional embedding. The AdamW optimizer (Loshchilov & Hutter, 2018) was used with $\beta_1 = 0.9$, $\beta_2 = 0.999$, a weight decay of 0.01, and a linear learning rate decay scheduler starting at $\text{lr} = 10^{-4}$ and ending at 0, with 5 warm-up steps. Training consisted of 50 epochs for reasoning, 200K steps for in-context learning, and 100K iterations for language modeling, using a batch size of 64. For time-dependent models, $\gamma(t)$ and $\alpha(t)$ were initialized as zero and one vectors, respectively, following Peebles & Xie (2023). The input embeddings are added at each loop iteration. Furthermore, for Sudoku and in-context learning tasks, the output of each intermediate loop is incorporated into the loss as (Yang et al., 2023).

C.2. In-Context Learning

We followed the setting of Yang et al. (2024). The problem is to learn the function class from a given sequence composed of the pairs of input \mathbf{x}_i and output values $f(\mathbf{x}_i)$. The input for model is $(\mathbf{x}_1, f(\mathbf{x}_1), \dots, \mathbf{x}_k, f(\mathbf{x}_k), \mathbf{x}_{\text{test}})$, and model learns to predict $f(\mathbf{x}_{\text{test}})$. The model is trained on $f(\mathbf{x}_k)$ and its performance is evaluated on $f(\mathbf{x}_{\text{test}})$ using the squared error.

We use depth-4 decision trees with 20-dimensional inputs. Each function in this class is represented by a full binary tree with 16 leaf nodes. Non-leaf nodes are associated with specific input coordinates, while leaf nodes are assigned target values. To evaluate $f(\mathbf{x})$, the tree is traversed from the root, moving to the right if the coordinate value is positive and to the left otherwise. Inputs and leaf node values are sampled from $N(0, I)$, and the coordinates for non-leaf nodes are drawn uniformly at random. Our training setup follows the approach of Yang et al. (2024). Following the curriculum training approach of Garg et al. (2022); Yang et al. (2024), we progressively increase the task dimensionality from 5 to 20 in steps of 1 every 5000 steps, while the sequence length increases from 26 to 101 in increments of 5 over the same interval.

C.3. Language Modeling

Tokenization is performed using byte-pair encoding, following GPT-2 (Radford et al., 2019). The Transformer model is based on the GPT-2 decoder architecture (Radford et al., 2019). The baseline standard Transformer model consists of 6 layers, 8 attention heads, and an embedding size of 512. The Looped Transformer has a 1 layer, 12 attention heads, and a hidden dimension of 768, which were chosen to match the parameter size of the baseline. We initialize $\gamma(t)$ as zero vectors and $\alpha(t)$ as one vector for time-dependent models. The AdamW optimizer (Loshchilov & Hutter, 2018) is used with $\beta_1 = 0.9$, $\beta_2 = 0.95$, a weight decay of 1×10^{-4} , and a learning rate schedule with 2000 warmup steps. The maximum learning rate is set to 2×10^{-4} and decays to 6×10^{-5} using a cosine schedule. Training is conducted for 100k iterations with a batch size of 48 and a block size of 1024.

D. Discussion

Multiple Layers A natural question is whether our analysis, which focuses on single-layer Looped Transformers, can be extended to multi-layer architectures. We restricted our analysis to a single layer in order to highlight a key strength of Looped Transformers—namely, their universality as function approximators even with just one layer. A more specific question is whether multi-layer Looped Transformers can overcome potential limitations inherent to the single-layer design. While it is conceivable that deeper architectures but with fixed-depth feedforward layers may achieve better approximation accuracy, this remains an open question. The difficulty lies in the fact that such improvements are not captured in terms of asymptotic order but rather in constants, which are harder to analyze theoretically. For instance, if we allow a logarithmic number of layers depending on the desired approximation precision, then even our current construction may overcome the limitations of the looped architecture. However, this deviates from our main objective, which is to characterize the approximation rate solely in terms of the number of loops.

Additional Experiments To assess the model’s sensitivity to input continuity, we designed a perturbed version of the WikiText-103 dataset, where 10% of the tokens were randomly replaced. We trained models with and without timestep encoding and evaluated both their memorization performance and continuity behavior. Continuity was measured by applying small perturbations to the input and quantifying the change in output embeddings. The model with timestep encoding showed improved memorization (cross-entropy loss reduced from 4.32 to 4.18) and a significant reduction in continuity coefficients (from 130.6 to 21.5). These results suggest that timestep encoding not only enhances stability under perturbations but also enables more faithful input-output mappings, thereby improving both robustness and learning efficiency.