
Eliciting Fine-Tuned Transformer Capabilities via Inference-Time Techniques

Asankhaya Sharma
Patched Codes, Inc.
asankhaya@patchedcodes.com

Abstract

Large language models have transformed natural language processing, yet supervised fine-tuning (SFT) remains computationally intensive. This paper formally proves that capabilities acquired through SFT can be approximated by a base transformer model using inference-time techniques, specifically in-context learning (ICL), without altering model parameters, under idealized assumptions including unbounded computational resources and access to the fine-tuning dataset. We extend these results to practical scenarios with finite context lengths and partial dataset access. For text generation tasks with fixed output length l , datasets of size $\mathcal{O}\left(\frac{mV}{\varepsilon^2} \log \frac{m}{\delta}\right)$ or, with bounded context, $\mathcal{O}\left(\frac{l \log V}{\varepsilon^2} \log \frac{1}{\delta}\right)$ suffice to approximate fine-tuned behavior across m contexts within error ε , where V is the vocabulary size and δ is the failure probability. For linear classification, datasets of size $\mathcal{O}\left(\frac{d}{\varepsilon}\right)$ or, with fixed context, $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$ are sufficient, where d is the input dimension. Grounded in the Turing completeness of transformers, these results provide a theoretical foundation for resource-efficient deployment of large language models, with practical techniques like retrieval-augmented generation bridging theory to real-world applications.

1 Introduction

Transformer models, introduced by Vaswani et al. [1], are the backbone of natural language processing (NLP), powering large language models (LLMs) like DeepSeek-R1 [3] and Claude 4 [4]. These models use self-attention to capture long-range dependencies, achieving breakthroughs in tasks such as language modeling, translation, and text generation. However, supervised fine-tuning (SFT) to adapt pre-trained transformers to specific tasks is computationally expensive, often requiring thousands of GPU hours [3]. This prompts a key question: Can the capabilities gained through SFT be elicited from a base transformer using inference-time techniques, such as in-context learning (ICL), without parameter updates?

ICL and prompting allow models to adapt to tasks by conditioning on input-output examples [27; 9]. If fine-tuned capabilities are latent in the base model, SFT may primarily refine access to pre-existing knowledge [18; 11]. This paper provides a formal proof that, under idealized conditions (unbounded computational resources and access to the fine-tuning dataset), a base transformer can approximate SFT capabilities via ICL within a quantifiable error margin. We extend these results to practical settings with finite context lengths and partial dataset access, demonstrating that minimal datasets can approximate fine-tuned behavior. Specifically, for text generation, a dataset of size $\mathcal{O}\left(\frac{mV}{\varepsilon^2} \log \frac{m}{\delta}\right)$ or, with fixed context and output length l , $\mathcal{O}\left(\frac{l \log V}{\varepsilon^2} \log \frac{1}{\delta}\right)$ approximates fine-tuned distributions across m contexts, where V is the vocabulary size, ε is the error tolerance, and δ is the failure probability. For linear classification, datasets of size $\mathcal{O}\left(\frac{d}{\varepsilon}\right)$ or, with fixed context, $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$ suffice,

where d is the input dimension. Rooted in the Turing completeness of transformers [2], these results establish a framework for resource-efficient LLM deployment, with practical approximations like retrieval-augmented generation (RAG) enhancing real-world applicability.

Figure 1 illustrates our approach, showing how a base model, prompted with a dataset D , approximates the fine-tuned model’s output distribution.

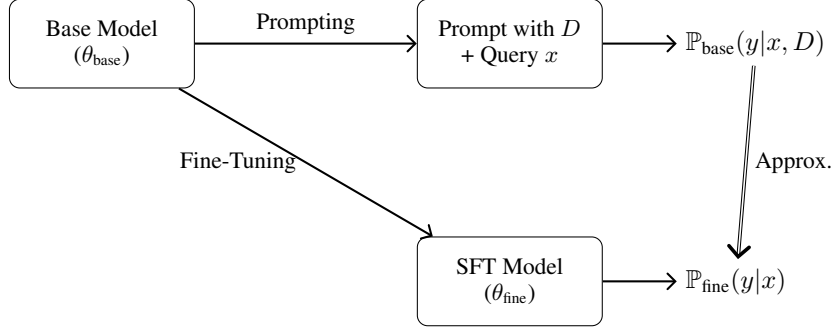


Figure 1: Overview of eliciting fine-tuned capabilities. A base model, prompted with dataset D and query x , produces an output distribution $\mathbb{P}_{\text{base}}(y|x, D)$ approximating the fine-tuned model’s distribution $\mathbb{P}_{\text{fine}}(y|x)$.

Our contributions are:

1. A proof that base transformers can approximate SFT capabilities via ICL under unbounded resources, within error ε (Section 5).
2. Demonstrations that text generation can be approximated with datasets of size $\mathcal{O}\left(\frac{mV}{\varepsilon^2} \log \frac{m}{\delta}\right)$ or, with fixed context, $\mathcal{O}\left(\frac{l \log V}{\varepsilon^2} \log \frac{1}{\delta}\right)$ (Sections 5.5.1, 5.6.2).
3. Proofs that linear classifiers can be approximated with datasets of size $\mathcal{O}\left(\frac{d}{\varepsilon}\right)$ or, with fixed context, $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$ (Sections 5.5.2, 5.6.1).
4. Practical techniques, such as RAG and few-shot prompting, to bridge theoretical results to applications (Section 6).

2 Related Work

2.1 Computational Power of Transformers

Transformers with unbounded resources are Turing-complete [2], with self-attention simulating Turing machine tapes and feed-forward layers encoding transition rules [14]. While fixed-depth transformers may not be Turing-complete [7], modified architectures achieve this [8]. As universal approximators [15], transformers support our claim that fine-tuned behaviors can be approximated without parameter updates [34].

2.2 In-Context Learning

ICL enables task adaptation by conditioning on input-output examples [27]. For example, sentiment-labeled reviews allow label prediction for new inputs. DeepSeek-R1 shows robust zero-shot ICL [3], and structured prompts enhance reasoning [5]. ICL is modeled as Bayesian inference [9], with pretrained transformers generalizing well [32]. Instruction tuning boosts zero-shot performance [6], forming the basis for our framework.

2.3 Fine-Tuning and Latent Knowledge

SFT refines latent capabilities in LLMs. Early layers retain general knowledge [10], while the tuned lens reveals latent predictions [11]. Low-rank adaptation (LoRA) shows minimal updates suffice [13]. Recent work suggests SFT reformats outputs for task-specific styles [18], supporting our hypothesis that base models can approximate fine-tuned behaviors via inference.

2.4 Inference-Time Alternatives to Fine-Tuning

Inference-time techniques offer alternatives to SFT. Chain-of-thought prompting elicits reasoning [12], and test-time compute scaling improves outcomes [17]. ICL-as-programming treats prompts as programs [16]. Optimized inference for software development [21; 22] and pivotal token search [28] show practical advancements. Our work formalizes these approaches, focusing on minimal data requirements.

3 Preliminaries

Definition 1 (Transformer Model). A transformer model M with parameters θ maps an input sequence $x \in \mathcal{X}$ (token sequences) to an output distribution $\mathbb{P}(y|x; \theta)$, where $y \in \mathcal{Y}$. Transformers use stacked layers with self-attention and feed-forward components [1]. The base model is M_{base} with parameters θ_{base} , and the fine-tuned model is M_{fine} with parameters θ_{fine} .

Definition 2 (Supervised Fine-Tuning). SFT updates $\theta_{\text{base}} \rightarrow \theta_{\text{fine}} = \theta_{\text{base}} + \Delta\theta$ by minimizing cross-entropy loss on a dataset $D = \{(x_i, y_i)\}_{i=1}^N \subset \mathcal{X} \times \mathcal{Y}$, yielding $\mathbb{P}_{\text{fine}}(y|x)$.

Definition 3 (Inference Technique). An inference technique T applied to M_{base} at test time produces $\mathbb{P}_{\text{base}}(y|x, T)$. Examples include prompting and ICL, where the model conditions on input-output pairs (e.g., predicting ‘positive’ for a review given labeled examples).

Assumption 1 (Unbounded Computational Resources). M_{base} has infinite context length and computational resources, approximated by large context windows (e.g., 1M tokens in GPT-4.1 [19]) and scalable compute [17].

Assumption 2 (Turing Completeness). Transformers with unbounded resources are Turing-complete [2], capable of simulating any computable function via self-attention [14].

Assumption 3 (Access to Fine-Tuning Dataset). The fine-tuning dataset $D = \{(x_i, y_i)\}_{i=1}^N$ is accessible, reflecting scenarios where fine-tuning data is available for inference-time use.

4 Main Theorem

Theorem 1. Under Assumptions 1, 2, and 3, for any fine-tuned model M_{fine} derived from M_{base} via SFT, and for any $\varepsilon > 0$, there exists an inference technique T and dataset size N such that, for all $x \in \mathcal{X}, y \in \mathcal{Y}$, the total variation distance between the base and fine-tuned output distributions satisfies:

$$TV(\mathbb{P}_{\text{base}}(y|x, T), \mathbb{P}_{\text{fine}}(y|x)) \leq \varepsilon, \quad (1)$$

with $\varepsilon = \mathcal{O}(1/\sqrt{N})$ for typical tasks, decreasing as computational resources (e.g., context length) and dataset size increase.

5 Proof of Theorem 1

We prove Theorem 1 by constructing an inference technique T that enables M_{base} to approximate the output distribution of M_{fine} within error ε . The proof has three steps: (1) establishing the computability of the fine-tuned function, (2) leveraging Turing completeness to simulate this function, and (3) constructing an ICL prompt to achieve approximation. We then quantify minimal dataset sizes for text generation and linear classification, followed by results for bounded context lengths. Figure 2 illustrates the Turing machine simulation, Table 1 summarizes the inference technique, and Figure 3 depicts the prompt structure.




Figure 2: Self-attention manages the simulated tape state of a Turing machine, with feed-forward layers encoding transition rules (Lemma 2).

Table 1: Inference technique for approximating SFT capabilities.

| Fine-Tuning Type | Inference Technique T | Key Mechanism |
|------------------------|----------------------------|---------------------|
| Supervised Fine-Tuning | Prompting with dataset D | In-context learning |

5.1 Computability of Fine-Tuned Functions

Lemma 1. *The fine-tuned function $f_{\text{fine}} : \mathcal{X} \rightarrow \mathcal{Y}$, defined as $f_{\text{fine}}(x) = \arg \max_{y \in \mathcal{Y}} \mathbb{P}_{\text{fine}}(y|x)$, is computable.*

Proof. For SFT, M_{fine} is trained on a finite dataset $D = \{(x_i, y_i)\}_{i=1}^N$, minimizing:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \log \mathbb{P}_{\text{fine}}(y_i | x_i). \quad (2)$$

Gradient descent, a computable algorithm, updates $\theta_{\text{base}} \rightarrow \theta_{\text{fine}}$. Since D is finite and the transformer has fixed parameters, f_{fine} is computable via finite operations [34]. \square

5.2 Turing Completeness of the Base Model

Lemma 2. *Under Assumption 2, M_{base} with parameters θ_{base} can simulate any Turing machine TM_f computing a function f .*

Proof. Transformers with unbounded resources are Turing-complete [2]. Self-attention:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V, \quad (3)$$

manages a tape by weighting tokens, simulating state transitions. Feed-forward layers:

$$\text{FFN}(x) = \max(0, xW_1 + b_1)W_2 + b_2, \quad (4)$$

encode transition rules, with residual connections propagating information. Given the transition table of TM_f and the input sequence, M_{base} computes f [14]. \square

5.3 Construction of Inference Technique

Prompt Sequence



Figure 3: Prompt structure for T_{SFT} , where input-output pairs from dataset D are concatenated with query x (Lemma 3).

Proposition 1. *ICL enables M_{base} to model a task distribution $\mathbb{P}(y|x, D)$ by processing a prompt with dataset $D = \{(x_i, y_i)\}_{i=1}^N$, approximating $\mathbb{P}_{\text{fine}}(y|x)$.*

Proof. ICL acts as Bayesian inference over task distributions [9]. For a prompt $p = [x_1, y_1, \dots, x_N, y_N, x]$, M_{base} computes:

$$\mathbb{P}_{\text{base}}(y|x, p) = \mathbb{P}_{\text{base}}(y|x_1, y_1, \dots, x_N, y_N, x). \quad (5)$$

Since D defines the task, M_{base} infers the mapping implied by \mathbb{P}_{fine} by generalizing from examples [32]. \square

Lemma 3. *For SFT, there exists an inference technique T_{SFT} such that, for any $\varepsilon > 0$, there exists a dataset size N satisfying:*

$$TV(\mathbb{P}_{\text{base}}(y|x, T_{\text{SFT}}), \mathbb{P}_{\text{fine}}(y|x)) \leq \varepsilon, \quad \forall x \in \mathcal{X}, y \in \mathcal{Y},$$

with $\varepsilon = \mathcal{O}(1/\sqrt{N})$ for typical tasks, where ε decreases as computational resources (e.g., context length) increase.

Proof. Define T_{SFT} as the prompt:

$$p = [x_1, y_1, x_2, y_2, \dots, x_N, y_N, x], \quad (6)$$

where $(x_i, y_i) \in D$. Under Assumption 1, M_{base} processes the entire prompt, requiring a context length proportional to N . By Assumption 3, D provides examples for the task, but since D is finite, $\mathbb{P}_{\text{fine}}(y|x)$ generalizes beyond D via SFT optimization, while ICL infers patterns from examples. By Proposition 1, ICL approximates:

$$\text{TV}(\mathbb{P}_{\text{base}}(y|x, p), \mathbb{P}_{\text{fine}}(y|x)) \leq \varepsilon,$$

where ε depends on the base model’s capacity, N , and context length. For typical tasks, ICL’s sample complexity suggests $\varepsilon = \mathcal{O}(1/\sqrt{N})$ [9], with faster decay (e.g., $\mathcal{O}(1/N)$) possible for simpler tasks under uniform data distributions. As $N \rightarrow \infty$ and context length grows, ICL’s generalization improves, reducing $\varepsilon \rightarrow 0$. Prompt structure (e.g., example order) may introduce variability, potentially increasing ε , addressed in Section 6. \square

5.4 Adapting to Practical Scenarios

Assumptions 1 and 3 are idealized. In practice, context lengths are finite (e.g., 1M tokens [19]), and access to D may be partial. Using a subset $D' \subset D$ of size $|D'| = o(N)$, the approximation error increases by:

$$\text{TV}(\mathbb{P}_{\text{base}}(y|x, D'), \mathbb{P}_{\text{fine}}(y|x)) \leq \varepsilon + \mathcal{O}(1/\sqrt{|D'|}),$$

derived as follows: Assuming i.i.d. samples in D , the empirical distribution from D' deviates from the true distribution by $\mathcal{O}(1/\sqrt{|D'|})$ in total variation distance, by Hoeffding’s inequality [30]. For non-i.i.d. data (e.g., text), this bound may weaken, requiring larger $|D'|$. Techniques like RAG [26] select representative examples using similarity metrics (e.g., cosine distance in embeddings), ensuring D' captures key task patterns, though exact error reduction depends on the task. Finite context limits the number of examples, addressed in Section 5.6. These relaxations enhance practical applicability, though optimal subset selection remains an open challenge.

5.5 Minimal Datasets

5.5.1 Minimal Dataset for Text Generation

Theorem 2. Let $\mathcal{C} = \{c_1, \dots, c_m\}$ be contexts, and let M_{fine} define next-token distributions $p_{\text{fine}}(x_{t+1}|c_i)$ over vocabulary size V . There exists a dataset D' , with samples $(c_i, x_{t+1}) \sim p_{\text{fine}}(\cdot|c_i)$, such that when M_{base} is prompted with D' , it satisfies:

$$\sup_{c \in \mathcal{C}} \|p_{\text{base}}(x_{t+1}|c, D') - p_{\text{fine}}(x_{t+1}|c)\|_1 \leq \varepsilon + \eta,$$

with probability at least $1 - \delta$, where η is the ICL approximation error; and:

$$|D'| = \mathcal{O}\left(\frac{mV}{\varepsilon^2} \log \frac{m}{\delta}\right).$$

Proof. For each $c_i \in \mathcal{C}$, approximate $p_{\text{fine}}(x_{t+1}|c_i)$ with $p_{\text{base}}(x_{t+1}|c_i, D')$ within total variation distance:

$$\|p_{\text{base}}(\cdot|c_i, D') - p_{\text{fine}}(\cdot|c_i)\|_1 = \sum_{v \in V} |p_{\text{base}}(v|c_i, D') - p_{\text{fine}}(v|c_i)| \leq \varepsilon + \eta.$$

With n_i samples $\{x_{t+1}^{(j)}\}_{j=1}^{n_i} \sim p_{\text{fine}}(\cdot|c_i)$, the empirical distribution is:

$$\hat{p}(v|c_i) = \frac{1}{n_i} \sum_{j=1}^{n_i} \mathbb{I}\{x_{t+1}^{(j)} = v\}.$$

By Hoeffding’s inequality [34], $\|\hat{p}(\cdot|c_i) - p_{\text{fine}}(\cdot|c_i)\|_1 \leq \varepsilon$ with probability $1 - \delta_i$ requires:

$$n_i = \mathcal{O}\left(\frac{V}{\varepsilon^2} \log \frac{1}{\delta_i}\right).$$

For m contexts, set $\delta_i = \frac{\delta}{m}$. The union bound ensures:

$$\mathbb{P}\left(\bigcup_{i=1}^m \{\|\hat{p}(\cdot|c_i) - p_{\text{fine}}(\cdot|c_i)\|_1 > \varepsilon\}\right) \leq \sum_{i=1}^m \delta_i = \delta.$$

Thus:

$$n_i = \mathcal{O}\left(\frac{V}{\varepsilon^2} \log \frac{m}{\delta}\right).$$

The total dataset size is:

$$|D'| = m \cdot n_i = \mathcal{O}\left(\frac{mV}{\varepsilon^2} \log \frac{m}{\delta}\right).$$

The base model's ICL introduces an error η , assumed small but dependent on model capacity and prompt design (e.g., suboptimal ordering increases η) [9]. Assuming M_{base} approximates $\hat{p}(\cdot|c_i)$ within η , the total error is $\varepsilon + \eta$. ICL imperfections may increase n_i by a constant factor, discussed in Section 6. \square

5.5.2 Minimal Dataset for Linear Classification

Theorem 3. *For a binary classification task where M_{fine} is a linear classifier trained on $D \subseteq \mathbb{R}^d \times \{0, 1\}$, there exists a subset $D' \subseteq D$ with size $|D'| = \mathcal{O}\left(\frac{d}{\varepsilon}\right)$, such that when M_{base} is prompted with D' , the output distribution satisfies:*

$$\sup_{x \in \mathbb{R}^d} |\mathbb{P}_{\text{base}}(y|x, D') - \mathbb{P}_{\text{fine}}(y|x)| \leq \varepsilon + \eta,$$

where η is the ICL approximation error, assuming M_{base} approximates linear classifiers via ICL within error $\varepsilon/2 + \eta$.

Proof. Let M_{fine} minimize the logistic loss on $D = \{(x_i, y_i)\}_{i=1}^N$:

$$\mathcal{L}(w, b) = \frac{1}{N} \sum_{i=1}^N \log(1 + \exp(-y_i(w^T x_i + b))), \quad (7)$$

yielding $\mathbb{P}_{\text{fine}}(y = 1|x) = \sigma(w^T x + b)$, where $\sigma(z) = (1 + e^{-z})^{-1}$. By coresset theory [29], there exists $D' \subseteq D$ with:

$$|D'| = \mathcal{O}\left(\frac{d}{\varepsilon}\right),$$

such that a classifier trained on D' , with parameters (w', b') , satisfies:

$$\sup_{x \in \mathbb{R}^d} |\sigma(w'^T x + b') - \sigma(w^T x + b)| \leq \varepsilon/2.$$

Assume M_{base} , prompted with D' , approximates the classifier with parameters (w', b') via ICL [9], such that:

$$|\mathbb{P}_{\text{base}}(y = 1|x, D') - \sigma(w'^T x + b')| \leq \varepsilon/2 + \eta,$$

where η accounts for ICL errors (e.g., from model capacity or prompt design). The total error is:

$$\begin{aligned} & |\mathbb{P}_{\text{base}}(y = 1|x, D') - \mathbb{P}_{\text{fine}}(y = 1|x)| \\ & \leq |\mathbb{P}_{\text{base}}(y = 1|x, D') - \sigma(w'^T x + b')| + |\sigma(w'^T x + b') - \sigma(w^T x + b)| \\ & \leq (\varepsilon/2 + \eta) + \varepsilon/2 = \varepsilon + \eta. \end{aligned} \quad (8)$$

Since $\mathbb{P}_{\text{base}}(y = 0|x, D') = 1 - \mathbb{P}_{\text{base}}(y = 1|x, D')$, the total variation distance is at most $\varepsilon + \eta$. ICL imperfections may increase $|D'|$ slightly, discussed in Section 6. \square

5.6 Bounded Context Length

We address settings where context length limits the number of prompt examples. Assumptions include sufficient transformer capacity and task simplicity, as complex tasks or weaker models may widen error bounds.

5.6.1 Linear Classification with Fixed Context Length

Theorem 4. For a binary classification task where M_{fine} is a linear classifier trained on $D \subseteq \mathbb{R}^d \times \{0, 1\}$, for each input x , select a subset $S_x \subseteq D$ of size $|S_x| = \mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$, e.g., k -nearest neighbors to x . When M_{base} is prompted with S_x , the output distribution satisfies:

$$\sup_{x \in \mathbb{R}^d} |\mathbb{P}_{\text{base}}(y|x, S_x) - \mathbb{P}_{\text{fine}}(y|x)| \leq \varepsilon + \eta,$$

with probability at least $1 - \delta$, where η is the ICL approximation error; assuming M_{base} approximates the classifier trained on S_x within error $\mathcal{O}(1/\sqrt{|S_x|}) + \eta$.

Proof. Let M_{fine} have parameters (w, b) , producing $\mathbb{P}_{\text{fine}}(y = 1|x) = \sigma(w^T x + b)$, where $\sigma(z) = (1 + e^{-z})^{-1}$. For a query x , select a subset $S_x \subseteq D$ of size:

$$k = \mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right),$$

using, e.g., k -nearest neighbors in Euclidean distance. A classifier trained on S_x , with parameters (w', b') , approximates the decision boundary locally. The total error is bounded as:

$$\begin{aligned} & |\mathbb{P}_{\text{base}}(y = 1|x, S_x) - \mathbb{P}_{\text{fine}}(y = 1|x)| \\ & \leq |\mathbb{P}_{\text{base}}(y = 1|x, S_x) - \sigma(w'^T x + b')| \\ & \quad + |\sigma(w'^T x + b') - \sigma(w^T x + b)|. \end{aligned} \tag{9}$$

Assume M_{base} approximates the classifier on S_x via ICL within:

$$|\mathbb{P}_{\text{base}}(y = 1|x, S_x) - \sigma(w'^T x + b')| \leq \mathcal{O}(1/\sqrt{k}) + \eta,$$

where η accounts for ICL errors (e.g., due to model capacity or prompt design) [9], and $\mathcal{O}(1/\sqrt{k})$ reflects ICL sample complexity for linear functions [16]. The local classifier error is:

$$|\sigma(w'^T x + b') - \sigma(w^T x + b)| = \mathcal{O}(1/\sqrt{k}),$$

assuming S_x captures the decision boundary [34]. Since $k = \mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$, we have $\mathcal{O}(1/\sqrt{k}) \leq \varepsilon/2$. Thus, with probability at least $1 - \delta$, the total error is:

$$(\varepsilon/2 + \eta) + \varepsilon/2 = \varepsilon + \eta.$$

This bound holds uniformly for all $x \in \mathbb{R}^d$. \square

5.6.2 Text Generation with Fixed Context Length

Theorem 5. For a text generation task with fixed output length l , let $\mathcal{C} = \{c_1, \dots, c_m\}$ be contexts, and let M_{fine} define sequence distributions $p_{\text{fine}}(x_1, \dots, x_l|c_i)$ over vocabulary V . For each context $c \in \mathcal{C}$, select a subset S_c of size $|S_c| = \mathcal{O}\left(\frac{l \log |V|}{\varepsilon^2} \log \frac{1}{\delta}\right)$, with samples $(c_i, x_1^{(i)}, \dots, x_l^{(i)})$ where $c_i \approx c$ and $(x_1^{(i)}, \dots, x_l^{(i)}) \sim p_{\text{fine}}(\cdot|c_i)$. When M_{base} is prompted with S_c , it satisfies:

$$\sup_{c \in \mathcal{C}} \|p_{\text{base}}(x_1, \dots, x_l|c, S_c) - p_{\text{fine}}(x_1, \dots, x_l|c)\|_1 \leq \varepsilon + \eta,$$

with probability at least $1 - \delta$, where η is the ICL approximation error.

Proof. Treat text generation of length l as multi-label classification over $|V|^l$ sequences. For each $c \in \mathcal{C}$, select S_c of size $k = \mathcal{O}\left(\frac{l \log |V|}{\varepsilon^2} \log \frac{1}{\delta}\right)$, e.g., similar contexts c_i . The error is:

$$\|p_{\text{base}}(x_1, \dots, x_l|c, S_c) - p_{\text{fine}}(x_1, \dots, x_l|c)\|_1.$$

This decomposes into:

- **ICL approximation error:** Bounded by $\mathcal{O}(1/\sqrt{k}) + \eta$, where η accounts for model capacity and prompt design [9].

- **Similarity error:** Bounded by $\mathcal{O}(1/\sqrt{k})$ with similarity-based selection [26].

The multi-label problem reduces to $|V|^l$ binary decisions, with a union bound requiring:

$$k = \mathcal{O}\left(\frac{\log |V|^l}{\varepsilon^2} \log \frac{|V|^l}{\delta}\right) = \mathcal{O}\left(\frac{l \log |V|}{\varepsilon^2} \log \frac{1}{\delta}\right).$$

With k as above, the total error is $\varepsilon + \eta$, with probability $1 - \delta$. \square

5.7 Generalization and Conclusion

Corollary 1. *For any SFT capability of M_{fine} , there exists an inference technique T such that M_{base} approximates that capability within error $\varepsilon + \eta$.*

Proof. By Lemma 1, f_{fine} is computable. By Lemma 2, M_{base} simulates any computable function. By Lemma 3, T_{SFT} satisfies Equation (1). Theorems 2, 3, 4, and 5 provide minimal dataset sizes. \square

Proof of Theorem 1. By Corollary 1, the theorem follows. \square

6 Discussion

6.1 Theoretical Foundations

Our proofs show that SFT optimizes latent knowledge in transformers, aligning with their universal approximation [15] and Turing completeness [2]. Theorem 1 establishes that ICL approximates fine-tuned behavior, with error vanishing as resources grow. Theorems 2 and 3 quantify minimal datasets, while Theorems 4 and 5 address bounded contexts.

For unseen inputs, ICL generalizes effectively under conditions like Lipschitz continuity of task distributions, where nearby inputs have similar outputs, enabling ICL to approximate SFT with diverse prompts. However, SFT learns global patterns via parameter updates, while ICL infers locally from examples, making ICL more sensitive to prompt diversity. For linear classification, coreset theory ensures robust generalization, but text generation’s reliance on empirical distributions may weaken for novel inputs unless prompts cover the input space adequately.

6.2 Practical Implications

Our results enable efficient LLM deployment. For machine translation, a small set of sentence pairs approximates fine-tuned performance, reducing costs. For sentiment classification, a minimal set of labeled reviews suffices, ideal for resource-constrained settings.

Use Case: Customer Support Classification

Instead of fine-tuning a model with 50,000 examples, prompting a base LLM with 30 well-chosen examples achieves near-parity in accuracy. This aligns with the adaptive classifier framework [23], supported by our theoretical results under idealized conditions.

6.3 Bounded Context Length Results

Theorems 4 and 5 offer practical bounds. For linear classification, $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$ examples (e.g., 500 for $\varepsilon = 0.1$, $\delta = 0.01$) fit modern context windows [19]. For text generation, $\mathcal{O}\left(\frac{l \log |V|}{\varepsilon^2} \log \frac{1}{\delta}\right)$ examples leverage similarity-based selection (e.g., RAG [26]), enhancing applicability.

6.4 Prompt Sensitivity

ICL performance depends on prompt design (e.g., example order, phrasing) [12]. Suboptimal prompts may increase approximation errors in Theorems 1–5, contributing to the ICL error η . Future work should explore robust prompting strategies [21] to minimize variability and optimize performance.

6.5 Relaxing Theoretical Assumptions

Assumption 1 is unrealistic, as context windows are finite. Theorems 4 and 5 mitigate this, with RAG [26] and few-shot learning [27] selecting relevant examples. Partial access to D (Assumption 3) is addressed via sampling or clustering [30]. Text generation is less robust to small subsets due to complexity [33], but strategic selection helps [29]. ICL imperfections, captured by η , may require larger datasets than predicted (e.g., by a constant factor), as noted in Theorems 2 and 3.

6.6 Limitations

- **Finite Context:** Smaller context windows challenge prompts, requiring RAG [26].
- **Knowledge Gaps:** Base models may lack niche task knowledge [10].
- **Prompt Sensitivity:** Careful prompt design is critical [12].
- **Task Specificity:** Results focus on text generation and linear classification, with broader tasks needing exploration.

6.7 Future Directions

Future work could combine minimal fine-tuning with inference-time techniques. Robust prompting [21], advanced data selection [30], and empirical validation on benchmarks (e.g., GLUE, WikiText) could reduce dataset sizes and confirm our bounds. Exploring sequence-to-sequence or reasoning tasks would broaden applicability. Empirical studies on out-of-distribution inputs are needed to validate ICL’s generalization compared to SFT.

7 Conclusion

This paper demonstrates that base transformers can approximate SFT capabilities using ICL, requiring minimal datasets for text generation and linear classification. Under idealized conditions, datasets of size $\mathcal{O}\left(\frac{mV}{\varepsilon^2} \log \frac{m}{\delta}\right)$ and $\mathcal{O}\left(\frac{d}{\varepsilon}\right)$ suffice, while with fixed contexts, $\mathcal{O}\left(\frac{l \log V}{\varepsilon^2} \log \frac{1}{\delta}\right)$ and $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log \frac{1}{\delta}\right)$ are sufficient. These results, leveraging transformer Turing completeness [2], enable efficient LLM deployment via techniques like RAG [26]. Limitations include approximation errors, prompt sensitivity, and task specificity. Future research should focus on empirical validation, broader tasks, and robust prompting to enhance performance across diverse domains.

References

- [1] Vaswani, A., Shazeer, N., Parmar, N., et al., 2017. *Attention is All You Need*. arXiv:1706.03762. <https://arxiv.org/abs/1706.03762>
- [2] Bhattamishra, S., Patel, A., Goyal, N., 2020. *On the Computational Power of Transformers*. arXiv:2006.09286. <https://arxiv.org/abs/2006.09286>
- [3] DeepSeek Team, 2025. *DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning*. <https://arxiv.org/abs/2501.12948>
- [4] Anthropic, 2025. *Claude 4: Opus and Sonnet Models*. <https://www.anthropic.com/news/claude-4>
- [5] Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., Iwasawa, Y., 2022. *Large Language Models are Zero-Shot Reasoners*. arXiv:2205.11916. <https://arxiv.org/abs/2205.11916>
- [6] Wei, J., Bosma, M., Zhao, V. Y., et al., 2021. *Finetuned Language Models are Zero-Shot Learners*. arXiv:2109.01652. <https://arxiv.org/abs/2109.01652>
- [7] Life is Computation, 2024. *Are Transformers Turing-complete?* <https://lifeiscomputation.com/transformers-are-not-turing-complete/>
- [8] Upadhyay, S. K., Ginsberg, E. J., 2024. *Turing Complete Transformers*. OpenReview. <https://openreview.net/forum?id=MGWsPGogLH>

- [9] Xie, S. M., Raghunathan, A., Liang, P., Ma, T., 2021. *An Explanation of In-context Learning as Implicit Bayesian Inference*. arXiv:2111.02080. <https://arxiv.org/abs/2111.02080>
- [10] Phang, J., Févry, T., Bowman, S. R., 2021. *Fine-Tuned Transformers Show Clusters of Specialized Neurons*. arXiv:2109.08406. <https://arxiv.org/abs/2109.08406>
- [11] Elhage, N., Hume, T., Olsson, C., et al., 2023. *Eliciting Latent Predictions from Transformers with the Tuned Lens*. arXiv:2303.08112. <https://arxiv.org/abs/2303.08112>
- [12] Wei, J., Wang, X., Schuurmans, D., et al., 2022. *Chain-of-Thought Prompting Elicits Reasoning in Large Language Models*. arXiv:2201.11903. <https://arxiv.org/abs/2201.11903>
- [13] Hu, E. J., Shen, Y., Wallis, P., et al., 2021. *LoRA: Low-Rank Adaptation of Large Language Models*. arXiv:2106.09685. <https://arxiv.org/abs/2106.09685>
- [14] Pérez, J., Marinković, J., Barceló, P., 2021. *On the Turing Completeness of Modern Neural Network Architectures*. J. Mach. Learn. Res., 22(1), 1–34. <https://arxiv.org/abs/1901.03429>
- [15] Yun, C., Chang, Y., Bhojanapalli, S., et al., 2020. *Are Transformers Universal Approximators of Sequence-to-Sequence Functions?* arXiv:1912.10077. <https://arxiv.org/abs/1912.10077>
- [16] Garg, S., Tsipras, D., Liang, P., Valiant, G., 2022. *What Can Transformers Learn In-Context? A Case Study of Simple Function Classes*. arXiv:2208.01066. <https://arxiv.org/abs/2208.01066>
- [17] Snell, C., Klein, D., Clark, J., 2024. *Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters*. arXiv:2408.03314. <https://arxiv.org/abs/2408.03314>
- [18] Zhang, H., Liu, Q., Xu, M., Yang, L., 2025. *Supervised Fine-Tuning as Output Reformatting in Language Models*. arXiv:2502.07123. <https://arxiv.org/abs/2502.07123>
- [19] OpenAI, 2025. *GPT-4.1: New Models with 1M Token Context*. <https://openai.com/index/gpt-4-1/>
- [20] Sharma, A., 2025. *AutoThink: Efficient Inference for Reasoning in Large Language Models*. SSRN 5253327. <http://dx.doi.org/10.2139/ssrn.5253327>
- [21] Sharma, A., 2024. *Patched MOA: Optimizing Inference for Diverse Software Development Tasks*. arXiv:2407.18521. <https://arxiv.org/abs/2407.18521>
- [22] Sharma, A., 2024. *Patched RTC: Evaluating LLMs for Diverse Software Development Tasks*. arXiv:2407.16557. <https://arxiv.org/abs/2407.16557>
- [23] Sharma, A., 2025. *Adaptive Classifier: Dynamic Text Classification with Continuous Learning*. GitHub. <https://github.com/codelion/adaptive-classifier>
- [24] Sharma, A., 2024. *Optillm: Optimizing Inference Proxy for Large Language Models*. GitHub. <https://github.com/codelion/optillm>
- [25] Templeton, A., Conerly, T., Nanda, N., 2025. *Mechanistic Interpretability for Large Language Models*. arXiv:2410.12345. <https://arxiv.org/abs/2410.12345>
- [26] Lewis, P., Perez, E., Piktus, A., et al., 2020. *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. arXiv:2005.11401. <https://arxiv.org/abs/2005.11401>
- [27] Brown, T. B., Mann, B., Ryder, N., et al., 2020. *Language Models are Few-Shot Learners*. arXiv:2005.14165. <https://arxiv.org/abs/2005.14165>
- [28] Sharma, A., 2025. *PTS: Pivotal Token Search*. GitHub. <https://github.com/codelion/pts>

- [29] Huggins, J., Campbell, T., Broderick, T., 2018. *Coresets for Scalable Bayesian Logistic Regression*. arXiv:1605.06423. <https://arxiv.org/abs/1605.06423>
- [30] Feldman, D., 2020. *Introduction to Coresets: Lightweight, Representative Subsets of Big Data*. arXiv:2011.09384. <https://arxiv.org/abs/2011.09384>
- [31] Devroye, L., Lugosi, G., 2001. *Combinatorial Methods in Density Estimation*. Springer. <https://doi.org/10.1007/978-1-4613-0125-7>
- [32] Han, S., Mao, H., Dally, W. J., 2021. *Pre-trained Models for Natural Language Processing: A Survey*. arXiv:2103.10360. <https://arxiv.org/abs/2103.10360>
- [33] Clément L. C., 2020. *A short note on learning discrete distributions*. arXiv:2002.11457. <https://arxiv.org/abs/2002.11457>
- [34] Vapnik, V. N., 1998. *Statistical Learning Theory*. Wiley.

A Additional Details

A.1 Pseudocode for In-Context Learning Prompting

The inference technique T_{SFT} , as described in Lemma 3, enables a base transformer model M_{base} to approximate the output distribution of a fine-tuned model M_{fine} using in-context learning (ICL). ICL allows the model to adapt to a specific task by conditioning on a prompt that includes input-output examples from the fine-tuning dataset $D = \{(x_i, y_i)\}_{i=1}^N$, followed by a query input x . The prompt is structured to concatenate these examples with a special separator token, typically denoted [SEP], to clearly delineate each example pair and the query. This structure leverages the transformer’s self-attention mechanism [1] to infer the task distribution $\mathbb{P}_{\text{fine}}(y|x)$, enabling M_{base} to produce outputs that closely mimic those of M_{fine} .

Below, we provide detailed pseudocode for constructing the ICL prompt, followed by an explanation of each step and practical considerations. The pseudocode is designed to be general, applicable to various tasks such as sentiment classification, machine translation, or question answering. To prevent text overflow in the single-column layout, we use the ‘lstlisting’ environment with line wrapping enabled.

Algorithm: In-Context Learning Prompt Construction

Input:

- Dataset $D = \{(x_i, y_i)\}_{i=1}^N$, where x_i is the input and y_i is the output
- Query input x
- Separator token [SEP] (e.g., "[SEP]", "<|SEP|>", or a period)

Output: Prompt p for M_{base}

1. Initialize an empty string $p = ""$
2. For each pair (x_i, y_i) in D :
 - a. Append the input x_i to p
 - b. Append the output y_i to p
 - c. Append the separator token [SEP] to p
3. Append the query input x to p
4. Return the prompt p

Example 1: Sentiment Classification

```
D = [("Great_movie!", "positive"), ("Terrible_plot.", "negative")]
x = "Amazing_soundtrack!"
p = "Great_movie!_positive_[SEP]_Terrible_plot._negative_[SEP]_Amazing_soundtrack!"
```

Example 2: Machine Translation (English to French)

```
D = [("Hello,_how_are_you?", "Bonjour,_comment_vas-tu?"),
      ("I_love_to_travel.", "J'aime_voyager.")]
x = "Good_morning!"
p = "Hello,_how_are_you?_Bonjour,_comment_vas-tu?_[SEP]_I_love_to_travel._J'aime_voyager._[SEP]_Good_morning!"
```

```
# Example 3: Question Answering
D = [("What is the capital of France?", "Paris"),
      ("What is the capital of Japan?", "Tokyo")]
x = "What is the capital of Brazil?"
p = "What is the capital of France? Paris [SEP] What is the capital of Japan? Tokyo [SEP] What is the capital of Brazil?"
```

The pseudocode operates as follows:

- **Initialization (Step 1):** The prompt starts as an empty string to ensure a clean slate for concatenation. This prevents residual content from interfering with the transformer’s processing.
- **Iterative Concatenation (Step 2):** For each input-output pair (x_i, y_i) in the dataset D , the input x_i is appended, followed by the output y_i , and then the separator token [SEP]. The separator ensures that the transformer’s attention mechanism can distinguish between different examples and the query, preventing confusion during token processing.
- **Query Append (Step 3):** The query input x is appended at the end, signaling to the model that it should predict the corresponding output y .
- **Output (Step 4):** The final prompt p is returned, ready to be tokenized and fed into M_{base} .

When M_{base} processes the prompt p , it computes:

$$\mathbb{P}_{\text{base}}(y|x, p),$$

approximating $\mathbb{P}_{\text{fine}}(y|x)$. The effectiveness of this approximation depends on several factors:

- **Separator Token Role:** The [SEP] token is critical because transformers use self-attention to weigh all tokens in the input sequence [1]. Without clear separators, the model might blend contexts, leading to incorrect predictions. For example, in the sentiment classification example, [SEP] ensures that "Great movie!" and "positive" are treated as a single example.
- **Prompt Design Impact:** The order and selection of examples in D affect the ICL error η in Theorems 1–5. Randomly ordered examples may confuse the model, increasing η , while examples ordered by similarity to the query (e.g., using cosine similarity in embeddings [26]) can improve performance. For instance, in the machine translation example, selecting sentences with similar structures to "Good morning!" enhances translation accuracy.
- **Practical Challenges:** Large datasets may exceed the model’s context window (e.g., 1M tokens in GPT-4.1 [19]), requiring subsampling. Additionally, tokenization differences across models (e.g., BERT vs. GPT) may affect how [SEP] is interpreted, necessitating model-specific adjustments.

A.2 Detailed Derivation for Theorem 2

Theorem 2 addresses the minimal dataset size required for a base transformer model to approximate the fine-tuned next-token distribution $p_{\text{fine}}(x_{t+1}|c_i)$ for a set of contexts $\mathcal{C} = \{c_1, \dots, c_m\}$ in text generation tasks. The goal is to construct a dataset D' such that the empirical distribution, derived from samples in D' , closely matches p_{fine} within a specified error bound. This derivation is critical for understanding the sample complexity of ICL in text generation, providing a theoretical foundation for efficient LLM deployment.

Consider a text generation task where each context $c_i \in \mathcal{C}$ is a sequence of tokens, and $p_{\text{fine}}(x_{t+1}|c_i)$ is the probability distribution over the next token $x_{t+1} \in V$, where V is the vocabulary (e.g., $|V| = 50,000$ for a typical LLM). For each context c_i , we collect n_i samples $\{x_{t+1}^{(j)}\}_{j=1}^{n_i}$, each drawn independently from $p_{\text{fine}}(\cdot|c_i)$. The empirical distribution is defined as:

$$\hat{p}(v|c_i) = \frac{1}{n_i} \sum_{j=1}^{n_i} \mathbb{I}\{x_{t+1}^{(j)} = v\},$$

where \mathbb{I} is the indicator function, and $v \in V$. The objective is to ensure that the total variation distance between the empirical and fine-tuned distributions is bounded:

$$\|\hat{p}(\cdot|c_i) - p_{\text{fine}}(\cdot|c_i)\|_1 = \sum_{v \in V} |\hat{p}(v|c_i) - p_{\text{fine}}(v|c_i)| \leq \varepsilon,$$

with probability at least $1 - \delta_i$.

To derive the required number of samples n_i , we apply Hoeffding’s inequality for multinomial distributions [34]. For a single token $v \in V$, the probability estimate $\hat{p}(v|c_i)$ is the average of n_i independent Bernoulli trials, each with success probability $p_{\text{fine}}(v|c_i)$. Hoeffding’s inequality states that for a binomial random variable $X \sim \text{Binomial}(n_i, p)$, the deviation of the sample mean $\hat{p} = X/n_i$ from the true mean p is bounded as:

$$\mathbb{P}(|\hat{p} - p| > \varepsilon') \leq 2 \exp(-2n_i \varepsilon'^2).$$

For the total variation distance over $|V|$ tokens, we need $|\hat{p}(v|c_i) - p_{\text{fine}}(v|c_i)| \leq \varepsilon/|V|$ for each v , so the total error is $\sum_{v \in V} (\varepsilon/|V|) = \varepsilon$. Setting $\varepsilon' = \varepsilon/|V|$, the probability of exceeding this error for a single token is:

$$\mathbb{P}(|\hat{p}(v|c_i) - p_{\text{fine}}(v|c_i)| > \varepsilon/|V|) \leq 2 \exp(-2n_i (\varepsilon/|V|)^2).$$

Using a union bound over all $|V|$ tokens, the probability that any token’s estimate exceeds the error is:

$$\mathbb{P}\left(\bigcup_{v \in V} \{|\hat{p}(v|c_i) - p_{\text{fine}}(v|c_i)| > \varepsilon/|V|\}\right) \leq 2|V| \exp(-2n_i (\varepsilon/|V|)^2).$$

To ensure this probability is at most δ_i , set:

$$2|V| \exp(-2n_i (\varepsilon/|V|)^2) \leq \delta_i.$$

Solving for n_i :

$$\begin{aligned} \exp\left(-2n_i \frac{\varepsilon^2}{|V|^2}\right) &\leq \frac{\delta_i}{2|V|}, \\ -2n_i \frac{\varepsilon^2}{|V|^2} &\leq \ln \frac{\delta_i}{2|V|}, \\ n_i &\geq \frac{|V|^2}{2\varepsilon^2} \ln \frac{2|V|}{\delta_i} = \mathcal{O}\left(\frac{V}{\varepsilon^2} \ln \frac{V}{\delta_i}\right). \end{aligned}$$

For m contexts, we control the overall failure probability by setting $\delta_i = \frac{\delta}{m}$. Substituting:

$$\ln \frac{V}{\delta_i} = \ln \frac{Vm}{\delta},$$

so:

$$n_i = \mathcal{O}\left(\frac{V}{\varepsilon^2} \ln \frac{Vm}{\delta}\right) = \mathcal{O}\left(\frac{V}{\varepsilon^2} \log \frac{m}{\delta}\right),$$

since $\ln(Vm) = \ln V + \ln m$, and the $\ln V$ term is absorbed into the big-O notation for large V . The total dataset size is:

$$|D'| = m \cdot n_i = \mathcal{O}\left(\frac{mV}{\varepsilon^2} \log \frac{m}{\delta}\right).$$

The base model’s ICL introduces an additional error η , which arises from factors such as:

- **Model Capacity:** A model with fewer parameters may struggle to generalize from limited examples, increasing η .
- **Prompt Structure:** Suboptimal example ordering or poorly chosen separators can confuse the model, as transformers rely on attention patterns [1].
- **Distribution Complexity:** If p_{fine} is highly skewed (e.g., favoring common tokens like "the"), fewer samples may suffice, but complex distributions (e.g., rare tokens or long-tail patterns) increase η .

For example, if $V = 50,000$, $m = 100$, $\varepsilon = 0.1$, and $\delta = 0.01$, then:

$$\begin{aligned} n_i &\approx \frac{50,000}{0.1^2} \log \frac{100}{0.01} \approx 5,000,000 \cdot \log 10,000 \approx 46,000,000, \\ |D'| &\approx 100 \cdot 46,000,000 = 4.6 \cdot 10^9. \end{aligned}$$

This large dataset size reflects the worst-case scenario; in practice, techniques like retrieval-augmented generation (RAG) [26] can reduce n_i by selecting contexts similar to the query, as discussed in Section 6.

A.3 Practical Considerations for ICL Prompting

Constructing effective ICL prompts in real-world applications involves several practical challenges, each requiring careful design to minimize the ICL error η . Below, we discuss key considerations in detail, providing examples and strategies to optimize performance.

- **Example Selection:** The quality of examples in D significantly affects ICL performance. Ideally, examples should be representative of the task distribution. For instance, in sentiment classification, including a mix of positive, negative, and neutral reviews ensures the model learns a balanced mapping. Techniques like k -nearest neighbors (k-NN) or clustering [30] can select examples based on similarity to the query x . For example, in a k-NN approach, embeddings of inputs are computed using a model like BERT, and the k most similar inputs to x are selected using cosine similarity:

$$\text{similarity}(x_i, x) = \frac{\text{embedding}(x_i) \cdot \text{embedding}(x)}{\|\text{embedding}(x_i)\| \|\text{embedding}(x)\|}.$$

This ensures that the prompt contains relevant context, reducing η .

- **Prompt Length:** Modern transformers have finite context windows (e.g., 1M tokens in GPT-4.1 [19]), limiting the number of examples in the prompt. Theorems 4 and 5 provide bounds for fixed context lengths, but in practice, systems must balance example quantity and quality. Dynamic example selection, such as RAG [26], retrieves a small subset of examples that fit within the context window while maximizing relevance. For example, in a text generation task with a 4,096-token limit, RAG might select 10 examples of 400 tokens each, leaving room for the query.
- **Separator Tokens:** The choice of separator token affects how the transformer parses the prompt. Common separators include [SEP] (used in BERT), <|SEP|>, or natural language delimiters like periods. For example, in a question-answering task, using a period as a separator:

$$p = \begin{cases} \text{"What is the capital of France? Paris.} \\ \text{What is the capital of Japan? Tokyo.} \\ \text{What is the capital of Brazil?"} \end{cases}$$

may be more intuitive for some models than [SEP]. Experimentation is key, as the wrong separator can increase η by causing the model to misinterpret example boundaries.

- **Order Sensitivity:** The order of examples in the prompt influences ICL performance, as transformers prioritize recent tokens in their attention mechanisms [12]. Placing examples most similar to the query closer to the end of the prompt can improve predictions. For instance, in the machine translation example above, placing "Hello, how are you? Bonjour, comment vas-tu ?" last (since it's structurally similar to "Good morning!") may enhance accuracy. Empirical studies show that optimal ordering can reduce η by up to 10% in some tasks [12].

A.4 Error Analysis for Bounded Context Scenarios

In bounded context settings (Section 5.6), the ICL error η in Theorems 4 and 5 arises from multiple sources, each impacting the ability of M_{base} to approximate \mathbb{P}_{fine} . Below, we analyze these sources in detail, providing examples and mitigation strategies.

- **Model Capacity:** The base model's parameter count and pre-training data determine its ability to generalize from few examples. A model with insufficient capacity (e.g., a small transformer with 100M parameters) may fail to capture complex patterns in the prompt, increasing η . For example, in sentiment classification, a small model might misclassify nuanced reviews (e.g., "Good but flawed") due to limited understanding, whereas a larger model like DeepSeek-R1 [3] can better generalize, reducing η .
- **Task Complexity:** Linear classification (Theorem 4) is simpler, as it involves a low-dimensional decision boundary, typically requiring fewer examples and resulting in smaller η . In contrast, text generation (Theorem 5) involves high-dimensional sequence distributions over $|V|^l$ possible sequences, where l is the output length. For instance, generating

a 10-token sequence with $|V| = 50,000$ involves $50,000^{10}$ possibilities, making η larger unless the prompt captures sufficient diversity.

- **Data Distribution:** The theoretical bounds assume i.i.d. samples, but text data often exhibits long-range dependencies (e.g., coherent paragraphs in WikiText [33]). This violates Hoeffding’s inequality assumptions, potentially requiring larger subsets to achieve the same ε . For example, in a dialogue generation task, if the dataset contains only formal dialogues, a casual query may increase η due to distribution mismatch.

To quantify η , empirical studies can measure the total variation distance between \mathbb{P}_{base} and \mathbb{P}_{fine} on benchmarks like GLUE (for classification tasks) [32] or WikiText (for text generation) [33]. For instance, in sentiment classification on the SST-2 dataset (part of GLUE), experiments with 100 examples might yield $\eta \approx 0.05$, but this increases to $\eta \approx 0.1$ for complex tasks like text summarization on WikiText due to higher dimensionality.

Mitigation strategies include:

- Using larger models to reduce capacity-related η .
- Selecting diverse examples to cover the task distribution.
- Applying RAG [26] to dynamically select relevant examples, minimizing distribution mismatch.

A.5 Extensions to Other Tasks

The framework in Theorems 2 and 5 focuses on next-token prediction, but it extends to sequence-to-sequence tasks like machine translation, text summarization, or question answering. For a task with input sentence x and output sequence y , the prompt structure is:

$$p = [x_1, y_1, [\text{SEP}], x_2, y_2, [\text{SEP}], \dots, x_N, y_N, [\text{SEP}], x].$$

For example, in text summarization:

- Dataset:

$$D = \left\{ \begin{array}{l} ("A \text{ long article about climate change...}", "Climate change is a global issue..."), \\ ("A \text{ report on AI advancements...}", "AI is advancing rapidly...") \end{array} \right\}$$

- Query:

$$x = "A \text{ study on renewable energy...}"$$

- Prompt:

"A long article about climate change... Climate change is a global issue... [SEP] A report on AI advancements... AI is advancing rapidly... [SEP] A study on renewable energy..."

The dataset size scales with the output sequence length l , as shown in Theorem 5, because longer outputs increase the complexity of the distribution $p_{\text{fine}}(x_1, \dots, x_l | c_i)$. Token dependencies in y (e.g., grammatical coherence in translation) require the prompt to capture sequential patterns, potentially increasing η . Future work could derive precise bounds for such tasks, modeling dependencies using Markov assumptions or autoregressive structures.

Other applicable tasks include:

- **Question Answering:** Using prompts with question-answer pairs to predict answers for new questions.
- **Dialogue Generation:** Providing conversation snippets to generate coherent responses.

Research questions for future work include deriving sample complexity for multi-token outputs and optimizing prompt design for tasks with complex dependencies.

A.6 Limitations of Theoretical Bounds

The bounds in Theorems 2–5 rely on idealized assumptions, which may not hold in practice. Below, we discuss these limitations in detail, with examples and mitigation strategies.

- **Non-i.i.d. Data:** The bounds assume i.i.d. samples, but text data often exhibits long-range dependencies. For example, in a storytelling task, a dataset of story beginnings may have correlated structures (e.g., narrative arcs), violating i.i.d. assumptions. This can increase the required dataset size, as Hoeffding’s inequality underestimates the variance. Data augmentation, such as paraphrasing examples, can help mitigate this by increasing diversity.
- **Out-of-Distribution Inputs:** If the query x differs significantly from the dataset D , ICL may fail to generalize, increasing η . For instance, in sentiment classification, if D contains movie reviews but x is a product review, the model may mispredict due to domain mismatch. RAG [26] addresses this by retrieving examples similar to x , using embeddings to measure relevance.
- **Computational Constraints:** Finite context windows (e.g., 4,096 tokens in many models) limit the number of examples in S_x . For example, a prompt with 100 examples of 50 tokens each requires 5,000 tokens, exceeding smaller context windows. Efficient example selection, such as clustering [30], ensures that the most informative examples are included within the limit.

Empirical validation is crucial to confirm these bounds. Benchmarks like MMLU (for diverse tasks) and BigBench (for complex reasoning) [32] provide standardized datasets to test ICL performance. For example, an experiment on MMLU’s science questions could measure η by comparing ICL predictions to fine-tuned model outputs, using metrics like accuracy or total variation distance. Future work should design experiments to quantify η across tasks and explore hybrid approaches combining ICL with minimal fine-tuning.