
Technical Debt in In-Context Learning: Diminishing Efficiency in Long Context

Taejong Joo¹ Diego Klabjan¹

Abstract

Transformers have demonstrated remarkable in-context learning (ICL) capabilities, adapting to new tasks by simply conditioning on demonstrations without parameter updates. Compelling empirical and theoretical evidence suggests that ICL, as a general-purpose learner, could outperform task-specific models. However, it remains unclear to what extent the transformers optimally learn in-context compared to principled learning algorithms. To bridge this gap, we introduce a new framework for quantifying optimality of ICL as a learning algorithm in stylized settings. Our findings reveal a striking dichotomy: while ICL initially matches the efficiency of a Bayes optimal estimator, its efficiency significantly deteriorates in long context. Through an information-theoretic analysis, we show that the diminishing efficiency is inherent to ICL. These results clarify the trade-offs in adopting ICL as a universal problem solver, motivating a new generation of on-the-fly adaptive methods without the diminishing efficiency.

1. Introduction

Transformers, particularly large language models (LLMs), are able to perform *in-context learning* (ICL) (Brown et al., 2020); they can adapt to new tasks simply by conditioning on demonstrations in their input prompt (Xie et al., 2022). Not only conveniently operated without any explicit parameter updates, but ICL even with just a few demonstrations (a.k.a. *few-shot* ICL) surprisingly outperforms task-specific state-of-the-art models in diverse tasks, from question answering to common sense reasoning (Chowdhery et al., 2023; Touvron et al., 2023; Brown et al., 2020).

This raises a fundamental question about how we shape artificial intelligence systems: Could ICL serve as a universal learner, obviating the need for task-specific models? To answer this, we must first address a more precise question:

How optimal is ICL as a learning algorithm, compared to principled learning algorithms?

At first glance, the impressive performance of few-shot ICL and more recently many-shot ICL (Bertsch et al., 2024;

Agarwal et al., 2024) might seem to be an affirmative answer. However, this conclusion would be premature. Even when ICL outperforms state-of-the-art task-specific models or matches (super) human-level performances, it may still not be an optimal learning algorithm. This is evidenced by the results that carefully fine-tuned LLMs often outperform ICL when provided with the same amount of demonstrations (Min et al., 2021; Zhao et al., 2024).

The question in principle can be accurately answered by comparing ICL with principled learning algorithms across LLMs with different data and model scales (Wei et al., 2023; Raventós et al., 2023) on diverse types of tasks (Srivastava et al., 2022; Wei et al., 2022). However, the computational demands for training modern LLMs pose significant challenges for evaluating optimality of ICL as a learning algorithm. The goal of this work is to answer the question without such prohibitive computational demands.

Previous attempts. To answer the question, theoretical studies have analyzed *asymptotic behavior of ICL* using rich tools from statistics and learning theory, such as regret and generalization bounds (Jeon et al., 2024; Zhang et al., 2023; Bai et al., 2023; Li et al., 2023b). However, these asymptotic results fall short of fully characterizing real-world LLM behavior. For instance, the regret upper bound for LLMs become nearly vacuous in few-shot regimes (Langford & Caruana, 2001; Dziugaite & Roy, 2017), which cannot explain the striking few-shot ICL performances. Moreover, because other principled learning algorithms have the similar asymptotic behavior, it remains unclear whether ICL is a *better* learning algorithm than such learning algorithms.

Physics-style or synthetic benchmarking approaches have provided valuable insights that *transformers might optimally learn in-context*, isolating core aspects of LLM training in controlled environments (Allen-Zhu & Li, 2023; Garg et al., 2022; Ahn et al., 2023). These approaches by nature can enable an efficient, comprehensive comparison between ICL and principled learning algorithms with arbitrarily high levels of statistical significances, while providing insights that often generalize to real-world LLMs despite inherent simplifications (cf. §A.1). Concretely, by examining ICL performances across different demonstration sizes in a stylized benchmark, Garg et al. (2022) and follow-up works (Akyürek et al., 2022; Von Oswald et al., 2023) show that

ICL seemingly learns new tasks with an efficiency comparable to provably optimal algorithms. However, these works have not yet provided an explicit relationship between relevant quantities (e.g., sample complexity and the optimality gap). Thus, the question of *to what extent* transformers can learn optimally in-context remains unanswered.

New benchmarking framework. We revisit the performance profiles (Dolan & Moré, 2002)—classic benchmarking framework for optimization software—for *benchmarking* optimality of ICL as a learning algorithm in the stylized ICL setting (Garg et al., 2022). Our framework can quantify *how many more demonstrations* are required for ICL to achieve a certain performance compared to principled learning algorithms. Thus, our analysis can accurately describe optimality of ICL with a more intuitive measure, making fundamental progress in physics-style approaches for ICL.

Unveiling diminishing efficiency in long context. As a result, we uncover a new insight on optimality of ICL in §3:

While ICL initially matches the efficiency of the Bayes optimal estimator, its efficiency significantly deteriorates in long context.

More precisely, for low performance requirements, ICL achieves near optimal sample complexity comparable to the Bayes optimal estimator, aligning with its strong few-shot performance observed in practice. However, ICL’s sample complexity sharply deteriorates beyond a certain threshold, often requiring 1.5 times more demonstrations to achieve high performance requirements than the Bayes optimal estimator. Further, we provide evidence that ICL may lack fundamental statistical properties (e.g., consistency and asymptotic efficiency) unlike the principled learning algorithms, which allow learning algorithms to benefit from large sample sizes. Crucially, this novel insight would be difficult to uncover through many-shot ICL experiments on real-world LLMs due to intractability (Agarwal et al., 2024) or analysis tools in the stylized setting (Garg et al., 2022), as ICL errors generally decrease with more demonstrations.

Intrinsic suboptimality of ICL. We prove that ICL without diminishing efficiency has stringent necessary conditions (e.g., negligible excess risk) using information-theoretic tools. Crucially, the result is independent to particular instantiation of models and environments, suggesting the diminishing efficiency is intrinsic to the ICL mechanism itself.

This discovery unveils a hidden *technical debt* in the ICL mechanism: the price we pay for its training-free adaptability is a fundamental inefficiency in sample complexity that compounds as we push toward higher performance targets with the current ICL mechanism *as is*.

Impact and outlook. Taken together, our findings suggest a more nuanced view of ICL than the prevailing excitement

for replacing task-specific fine-tuned models with ICL as a universal problem solver. While ICL’s ability to adapt without training remains attractive, our work reveals the fundamental technical debt that must be considered in AI system designs. Crucially, this debt appears intrinsic to the ICL mechanism and thus unlikely to be serviced by simply scaling data and model sizes. We hope these insights clarify the trade-offs in adopting ICL as a universal problem solver and motivate a new generation of “on-the-fly” adaptive methods without the diminishing efficiency.

2. Setup

In §2.1, we describe the meta ICL environment for evaluating ICL as a learning algorithm, followed by designs of a transformer for solving the meta ICL task (§2.2). We then devise principled predictors (§2.3) and compare them with transformers using performance measures defined in §2.4.

2.1. Meta ICL Environment

In the meta ICL (Garg et al., 2022), each prompt characterizes an instance of a learning problem. Specifically, a prompt H_T consists of demonstrations with a test input, i.e., $H_T \triangleq (X_1, Y_1, \dots, X_T, Y_T, X_{T+1})$, and each output is generated by some function f^* , i.e., $Y_t = f^*(X_t)$ for $t \in [T + 1] \triangleq \{1, 2, \dots, T + 1\}$. Here, the goal of a transformer is formalized as accurately predicting Y_{T+1} with H_t , which requires to (implicitly) infer the underlying function f^* from the demonstrations. We denote the set of demonstrations as $D_T \triangleq (X_1, Y_1, \dots, X_T, Y_T)$.

For the data generating distribution of a prompt H_T , we follow the approach of sampling target functions f^* from a *hierarchical distribution* (Panwar et al., 2024) to capture a more interesting aspect of a learning algorithm—model selection. Under the hierarchical f^* , the prompt H_T is realized by the following sampling process, which is denoted as $H_T \sim \mathbb{P}(\cdot; \mathcal{E})$ with parameters $\mathcal{E} \triangleq ([M], \sigma_w^2, \sigma_\epsilon^2)$.

1) Sample the implicit dimension $m \sim \mathcal{U}([M])$ from a uniform distribution over set $[M]$ and construct the (unobservable) feature space $\Phi_m(\mathcal{X})$:

$$\Phi_m(x) \triangleq [1, \cos(\frac{\pi x}{\mathcal{T}}), \sin(\frac{\pi x}{\mathcal{T}}), \dots, \cos(\frac{m\pi x}{\mathcal{T}}), \sin(\frac{m\pi x}{\mathcal{T}})]$$

where $\mathcal{T} > 0$ controls the frequency of the trigonometric functions.

2) Sample weight $w_m \sim \mathcal{N}(0, \sigma_w^2 \mathbf{I}_{2m+1})$, where \mathbf{I}_{2m+1} is the identity matrix with rank $2m + 1$. The weight w_m defines the target function f^* :

$$f^*(x) \triangleq w_m^\top \Phi_m(x) / \sqrt{2m + 1},$$

where the normalization factor $\sqrt{2m + 1}$ ensures that the variance of f^* remains constant across different m . We

let $\mathcal{F}_m \triangleq \{w^T \Phi_m(\cdot) | w \in \mathbb{R}^{2m+1}, \Phi_m : \mathcal{X} \rightarrow \mathbb{R}^{2m+1}\}$ denote the set of all possible target functions with implicit dimension m .

3) Construct a prompt H_T with a test output y_{T+1} by $x_t \sim \mathcal{U}([x_{\min}, x_{\max}])$, $y_t = f^*(x_t) + \epsilon_t$ for $t \in [T + 1]$, where $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is a random observation noise.

This hierarchical sampling ensures a diverse range of target functions with varying complexities, where a different realization of (m, w_m) instantiates a new learning problem.

In this work, we benchmark ICL with respect to different configurations of \mathcal{E} , called *scenario*, to enable comprehensive evaluations that could be encountered in practical scenarios (e.g., low signal-to-noise ratio (SNR), defined as $\text{Var}(f^*)/\sigma_\epsilon^2$, for emulating a highly noisy environment). Following Panwar et al. (2024), we set $\mathcal{T} = x_{\max} - x_{\min} = 5$ (our findings are indifferent to these values). We denote \mathcal{S} as a set of scenarios and \mathcal{E}_s as parameters of a scenario $s \in \mathcal{S}$. We also have $H_T^s \triangleq (X_1^s, Y_1^s, \dots, X_{T+1}^s)$ generated from $\mathbb{P}(\cdot; \mathcal{E}_s)$ for each scenario s , where we omit superscripts when there is no ambiguity.

2.2. Transformers

For a transformer TF_θ , we adopt the setup from Garg et al. (2022) and follow-up works (Panwar et al., 2024; Von Oswald et al., 2023; Akyürek et al., 2022; Raventós et al., 2023) that use the GPT-2 (Radford et al., 2019) architecture (cf. details in Appendix A.2). For optimizing θ in the pretraining stage, we use the following minimization objective

$$\mathcal{L}(\theta) \triangleq \mathbb{E}_{H_{T_{\text{train}}}} \left[\frac{1}{T_{\text{train}}} \sum_{t=0}^{T_{\text{train}}-1} l(\text{TF}_\theta(H_t), Y_{t+1}) \right] \quad (1)$$

where $H_{T_{\text{train}}}$ is generated by the prompt distribution described in §2.1. We use the squared loss function for l , following previous works in the regression setting. Also, we set $T_{\text{train}} = 50$ for all scenarios, which is roughly $2 \cdot (2M+1)$ as in the previous works (Garg et al., 2022; Panwar et al., 2024). We train TF_θ separately for each scenario.

2.3. Principled Baselines

To benchmark ICL, we derive principled baselines that learn from demonstrations D_t and produce a prediction function $f_b(\cdot; D_t)$, where b is the identifier of a particular baseline. We denote $f_b^t(x) \triangleq f_b(x; D_t)$ and $f_{\text{ICL}}^t(X_{t+1}) \triangleq \text{TF}_\theta(H_t)$ when ever there is no ambiguity.

The optimal baseline is Bayesian model averaging (BMA), which makes prediction by aggregating models from different hypothesis classes:

$$f_{\text{BMA}}^t(x) = \sum_{m \in [M]} p(\mathcal{F}_m | D_t) \hat{w}_m^\top(D_t) \Phi_m(x), \quad (2)$$

where $p(\mathcal{F}_m | D_t)$ is the posterior probability of model class \mathcal{F}_m and \hat{w}_m is the ridge regression estimator for \mathcal{F}_m ,

defined as $\hat{w}_m(D_t) = (\Phi_{m,t}^\top \Phi_{m,t} + \frac{\sigma_w^2}{\sigma_\epsilon^2} \mathbf{I}_{2m+1})^{-1} \Phi_{m,t}^\top \mathbf{Y}_t$ with $\Phi_{m,t} \in \mathbb{R}^{t \times (2m+1)}$ whose k -th row is $\Phi_m^\top(X_k)$ and $\mathbf{Y}_t = (Y_1, \dots, Y_t) \in \mathbb{R}^t$. It is a standard result that

$$f_{\text{BMA}}^t \in \arg \min_{f \in \mathcal{F}} \mathbb{E}_{Y_{t+1}} [l(f(X_{t+1}; D_t), Y_{t+1}) | H_t] \quad (3)$$

holds almost everywhere for all $t \in \mathbb{N}$, where \mathcal{F} is the set of all functions from H_t to \mathbb{R} (Ahuja & Lopez-Paz, 2023; Bishop, 2007).

In addition, we consider a family of principled baselines that embodies different model selection strategies while having the same model fitting capacity as the optimal predictor. Such baselines make predictions by

$$f_b^t(x; D_t) = \hat{w}_{m_b^\dagger}^\top(D_t) \Phi_{m_b^\dagger}(x), \quad (4)$$

where $m_b^\dagger \in \arg \max_{m \in [M]} \{\text{Score}_b(m)\}$ with $\text{Score}_b(\cdot)$ being some model selection criterion of b .

2.4. Measures for Benchmarking Optimality of ICL

Inspired by seminal work (Dolan & Moré, 2002) that benchmarks (deterministic) optimization software, we first define the base metric measuring the optimality of a learning algorithm in $s \in \mathcal{S}$. Then, we present the performance measures summarizing the base metric across \mathcal{S} . In the following, we let \mathcal{B} contain all baseline learning algorithms and ICL. We set the test prompt length as $T = 2T_{\text{train}} = 100$, which is within the length generalization regime (Zhou et al., 2024).

Base metric. Our base metric is the *performance ratio*, which normalizes the sample complexity of a learning algorithm by that of the best algorithm among all baselines.

Definition 2.1. For $b \in \mathcal{B}$ at $s \in \mathcal{S}$, the *performance ratio* of a requirement r against $\tilde{\mathcal{B}} \subseteq \mathcal{B}$ is defined as $R_b^s(r; \tilde{\mathcal{B}}) = N_b^s(r) / \min_{\tilde{b} \in \tilde{\mathcal{B}}} \{N_{\tilde{b}}^s(r)\}$, where $N_b^s(r) \triangleq \min \{t | \mathbb{E}[l(f_b^t(X_{t+1}^s), Y_{t+1}^s)] \leq r\}$ is the sample complexity of achieving the performance r .

The performance ratio quantifies the relative efficiency of a learning algorithm, addressing data-dependent nature of the sample complexity. Specifically, $R_b^s(r; \tilde{\mathcal{B}})$ indicates that the learning algorithm b requires $R_b^s(r; \tilde{\mathcal{B}})$ times more demonstrations to achieve performance r at scenario s compared to the best learner among $\tilde{\mathcal{B}}$. When $\text{BMA} \in \tilde{\mathcal{B}}$, algorithms with $R_b^s(r; \tilde{\mathcal{B}}) = 1$ have optimal efficiency at s due to (3).

Performance measures. Based on the performance ratio across different scenarios, our goal is to report a “single” score that summarizes how optimal ICL is across \mathcal{S} . However, naively summarizing the performance ratio for a requirement r is inappropriate because the difficulty of achieving r varies across learning problems, making comparisons inconsistent. Therefore, we define the *reference*

performance quantile $\psi_{\mathcal{B}^{\text{ref}}}^Q(s)$ as the Q -th quantile of reference performances at s for $Q \in (0, 1)$. Here, we measure the performance quantile in a reverse order, for making higher performance quantile analogous to higher performance. The reference performances at s is defined as a set of performances achieved by reference models $\mathcal{B}^{\text{ref}} \subseteq \mathcal{B}$; that is, $\{\mathbb{E}[l(f_b^t(X_{t+1}^s), Y_{t+1}^s)] | b \in \mathcal{B}^{\text{ref}}, t \in [T]\}$.

With this idea, the performance ratios across \mathcal{S} is summarized by the *mean performance ratio* and the *performance profile*, which are defined as follows.

Definition 2.2. For the performance quantile $\psi_{\mathcal{B}^{\text{ref}}}^Q$, the *mean performance ratio* of $b \in \mathcal{B}$ against $\tilde{\mathcal{B}} \subseteq \mathcal{B}$ is defined as $\text{MPR}(b; \psi_{\mathcal{B}^{\text{ref}}}^Q, \tilde{\mathcal{B}}) \triangleq \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} R_b^s(\psi_{\mathcal{B}^{\text{ref}}}^Q(s); \tilde{\mathcal{B}})$.

Definition 2.3. For the performance quantile $\psi_{\mathcal{B}^{\text{ref}}}^Q$, the *performance profile* of $b \in \mathcal{B}$ against $\tilde{\mathcal{B}} \subseteq \mathcal{B}$ at a ratio $\tau \geq 1$ is defined as

$$\rho_b(\tau; \psi_{\mathcal{B}^{\text{ref}}}^Q, \tilde{\mathcal{B}}) = \frac{1}{|\mathcal{S}|} |\{s \in \mathcal{S} : R_b^s(\psi_{\mathcal{B}^{\text{ref}}}^Q(s); \tilde{\mathcal{B}}) \leq \tau\}|.$$

The two measures capture complementary aspects of optimality of ICL. Specifically, the mean performance quantile quantifies the average inefficiency of a model b in attaining a certain performance, which is assumed to be achievable by b . In contrast, the performance profile measures the frequency with which a model b can achieve the performance quantile given a tolerance for inefficiency. These intuitive measures provide novel insights into optimality of ICL that are not apparent in previous error rates-based comparisons and asymptotic analyses.

2.5. On Usage of Stylized Setting

The stylized setting discussed in this section provides a rigorous benchmark for studying optimality of ICL with, in theory, arbitrarily high levels of statistical significance, including performance comparisons against BMA as a fundamental limit. While it simplifies certain real-world elements (e.g., autoregressive loss), empirical evidence suggests that insights from this setting generalize remarkably well to real-world tasks (Ahn et al., 2024; Li et al., 2023c). Though the setting necessarily simplifies real-world LLMs, its ability to deepen our understanding of ICL demonstrates its value. Refer to Appendix A.1 for more details.

3. Benchmarking ICL Efficiency

We measure to what extents transformers efficiently learn a new task through ICL compared to the optimal learning algorithm (§3.1) and principled baselines (§3.2).

3.1. Can Transformer Optimally Learn In Context?

We first examine the efficiency of ICL compared to the Bayes optimal predictor, which learns new concepts with op-




Figure 1. Mean performance ratio of ICL against BMA across different performance requirements. The shaded areas represent the standard deviation of the corresponding performance ratio.

timal efficiency (cf. (3)). For comprehensive evaluation, we design the test scenarios with various levels of SNRs: $\mathcal{S} = \{([M], \sigma_y^2, \sigma_w^2) \mid M = 10, \sigma_y^2 \in \{0.003, 0.03, 0.3\}, \sigma_w^2 \in \{0.1, 1, 10\}\}$ (cf. §2.1). Also, to minimize the impacts of stochasticity of the sampling process of H_t , we evaluate performances for each scenario 512 times. Then, we analyze the mean performance ratio of ICL against BMA for all quantiles of performances achieved by ICL; that is, we measure $\text{MPR}(\text{ICL}; \psi_{\mathcal{B}_1^{\text{ref}}}^Q, \tilde{\mathcal{B}}_1)$ with $\mathcal{B}_1^{\text{ref}} \triangleq \{\text{ICL}\}$, $Q \in \{0.01, 0.1, \dots, 0.9, 0.99\}$, and $\tilde{\mathcal{B}}_1 \triangleq \{\text{ICL}, \text{BMA}\}$. In this way, we measure the efficiency of ICL in achieving each performance level under various difficulties in extracting information from prompts. In the following, we regard prompts with more than 40 demonstrations as the many-shot regime where the average performance quantile is approximately 0.5 (cf. Figure A3 in Appendix).

Figure 1 reveals a striking dichotomy in optimality of ICL.

Near optimal few-shot efficiency. For low performance quantiles ($Q \leq 0.3$), ICL demonstrates its remarkable near optimal efficiency. Specifically, the mean performance ratio is at most 1.1, which means that it requires only 10% more demonstrations on average than the optimal learning algorithm to achieve the performance lower than $\psi_{\mathcal{B}^{\text{ref}}}^{0.3}(s)$ for $s \in \mathcal{S}$. Considering the average sample complexity for the performance quantile of 0.3 is 19, this explains ICL’s impressive few-shot performance observed in practice (e.g., demonstration sizes of 5 and 15 in Brown et al. (2020)).

Increasingly suboptimal many-shot efficiency. Starting from $Q = 0.3$ or more apparently from $Q = 0.7$ onward, the performance ratio grows almost monotonically with Q , increasing from around 1.1 at $Q = 0.3$ to around 1.2 at $Q = 0.7$ and to around 1.45 at $Q = 0.99$. That is, ICL becomes increasingly suboptimal compared to the optimal learning algorithm when pursuing high performance requirements. Considering higher performance quantiles requires larger demonstration sizes, this indicates suboptimal scaling in long context that suboptimality increases with the demonstration size.

Importantly, these findings do *not* contradict established benefits of many-shot ICL (Bertsch et al., 2024; Agarwal et al.,




Figure 2. Performance profiles ρ_b across different performance ratios τ under different target performance quantiles Q . Each curve represents the probability that a method achieves the desired performance within a factor τ of the best method’s sample complexity (x -axes). Figure A4 in Appendix illustrates results for all Q .

2024); as we analyze later in Figure 3, ICL still achieves monotonic improvement in MSE with more demonstrations. Rather, our novel evaluation framework reveals that this improvement comes at an increasingly inefficient sample complexity, indicating significant diminishing returns in extracting information from demonstrations.

3.2. Benchmarking ICL Against Principled Baselines

We have shown that ICL is significantly inefficient compared to BMA in high performance regimes. While BMA is learnable by minimizing (1), it might seem unrealistic for ICL to compete with BMA that performs the expensive model averaging operation. Therefore, we compare ICL with more practical baselines with a computational constraint that select a single model using principled criteria (cf. (4)): Akaike Information Criterion (AIC) (Akaike, 1974) as a minimax-rate optimal model selection mechanism, Bayesian Information Criterion (BIC) (Schwarz, 1978) as a consistent model selection mechanism, and Bayesian Model Comparison (BMC) as an efficient BMA alternative selecting maximum a posteriori model class. These baselines represent the spectrum of principled model selection methods, which often asymptotically converge to either AIC or BIC (Ding et al., 2018).

To quantitatively assess relative efficiency, we measure performance profiles $\rho_b(\tau; \psi_{\mathcal{B}_2^{\text{ref}}}^Q, \tilde{\mathcal{B}}_2)$ with $\mathcal{B}_2^{\text{ref}} = \{\text{ICL}, \text{AIC}\}$ and $\tilde{\mathcal{B}}_2 = \{\text{ICL}, \text{AIC}, \text{BIC}, \text{BMC}\}$. This allows us to measure the probability that each method achieves a reference performance level within given sample complexity budgets, which evaluates both efficiency and effectiveness (i.e., maximum achievable performances) of learning algorithms.

Superiority of ICL in few-shot regimes. Perhaps not surprisingly (given the results from comparison with BMA), ICL dominates the baselines with restricted capacity under low performance requirements. Specifically, it achieves the perfect performance profile at $\tau = 1$ for $Q \leq 0.3$. This means that it optimally attains the performance requirement in *all* scenarios when $Q \leq 0.3$. Given that each baseline has its own strength in certain scenarios, this guarantee is




Figure 3. Mean squared errors for different demonstration sizes. Figure A5 in Appendix illustrates results for all $s \in \mathcal{S}$ and $b \in \mathcal{B}$.

quite strong and not observed in other baselines. Further, for $Q = 0.4$, ICL reaches a perfect performance profile within $\tau \leq 1.2$. This means that ICL attains the required performance of $Q = 0.4$ in all scenarios by using at most 20% more demonstrations on average compared to the best method in each scenario. Conversely, all baselines selecting a single model struggle in the low-performance regime due to high uncertainty under a small number of demonstrations preventing them from selecting the proper model class (Hoeting et al., 1999; Wasserman, 2000).

Inferiority of ICL in many-shot regimes. Figure 2 illustrates diminishing efficiency of ICL in long context regimes. Specifically, as the performance requirement increases, the initial performance profile at $\tau = 1$ is reduced, indicating the reduced probability that ICL learns the most efficiently among $\tilde{\mathcal{B}}_2$. Beside, the computational budget τ required to reach perfect performance profile increases as the performance requirement increases. Eventually for $Q \geq 0.8$, even at $\tau = 3$, ICL achieves the performance profile around 0.8, which means that ICL cannot reach the performance requirements for 20% of cases by using even 3 times more demonstrations than other models.

Crucially, this increasingly suboptimal behavior is opposite to the behaviors of principled baselines. In Figure 2, as opposed to ICL, the principled learning algorithms significantly reduce the time to reach the (near) perfect performance profiles as Q increases. Eventually, despite their significant deficiencies in few-shot regimes, all such baselines become more effective (achieving higher performance profiles at $\tau = 3$) and more efficient (sharply improving the performance profiles with respect to τ) than ICL in many-shot regimes. Therefore, some characteristics enabling learning algorithms to leverage large number of demonstrations might be missing in the ICL mechanism.

To gain further insights, we qualitatively analyze MSEs across different numbers of demonstrations for each scenario. As a trivial baseline, we also consider an ensemble that aggregates the ridge estimators $\{\hat{w}_m\}_{m \in [M]}$ using equal weights. Figure 3 shows that while all methods show decreasing MSEs with more demonstrations, ICL exhibits persistent discrepancies from the principled learning algorithms in many-shot regimes. Further, in Figure 4, we analyze the squared prediction difference between each model




Figure 4. Squared prediction differences between BMA and other methods for different demonstration sizes. Figure A6 in Appendix illustrates results for all $s \in \mathcal{S}$.

and the Bayes optimal predictor for each scenario. Critically, it reveals that while consistent estimators (BMC, BIC) seem to converge in L^2 (albeit at different rates), ICL’s L^2 distance to f_{BMA}^t plateaus after receiving few demonstrations. This behavior mirrors the trivial ensemble, which does not update its hypothesis about the model class with demonstrations. This suggests another fundamental limitation: ICL may lack asymptotic efficiency and consistency (cf. Ding et al. (2018) for formal definitions). These findings challenge the prevailing optimism about ICL’s scalability.

3.3. On Sources of the Diminishing Efficiency

We observe a significant suboptimality of ICL under high performance requirements, which typically requires longer context sizes than the pretraining prompt (cf. Figure A3 in Appendix). Given universally observed deficiencies of machine learning models in the out-of-distribution regimes (Hendrycks & Dietterich, 2019; Koh et al., 2021), it is tempting to attribute the diminishing efficiency to the deficiencies in out-of-distribution regimes.

We take a closer look at this in Figure 3, which corresponds to the achievable error due to the bias-variance decomposition. Recalling that $T_{\text{train}} = 50$ was used for pretraining, Figure 3 and Figure A6 in Appendix show no apparent differences in the achievable error between in-distribution and out-of-distribution regimes, except in low SNR scenarios $(\sigma_w^2, \sigma_\epsilon^2) = (0.1, 0.03)$ and $(\sigma_w^2, \sigma_\epsilon^2) = (1, 0.3)$. This finding aligns with the length generalization literature, which suggests that transformers often generalize to contexts up to 2.5 times longer than those seen during pretraining (Zhou et al., 2024). Further, given that the average performance quantile at T_{train} is 0.6, Figure 1 reveals that fundamental inefficiency already emerges in the in-distribution regime.

Therefore, the diminishing efficiency observed in §3.1 and §3.2 cannot be fully attributed to the transformers’ out-of-distribution generalization capability. Rather, as we analyze next in §4, it is intrinsic to the ICL mechanism itself.

4. Analyzing Suboptimality of ICL

Using information-theoretic tools, we explain why ICL’s efficiency as a learning algorithm diminishes in long context.

4.1. ICL Error Decomposition

Adopting a Bayesian viewpoint (Jeon et al., 2024), we denote the oracle distribution with e drawn from an environment \mathcal{E} by $\bar{P}_e^t(\cdot) \triangleq \mathbb{P}(Y_{t+1} \in \cdot | H_t, e) = \mathbb{P}(Y_{t+1} \in \cdot | X_{t+1}, e)$ (e.g., \mathcal{E} characterizes the sampling process in §2.1 with $e = (m, w_m)$). Similarly, we let TF_θ models the conditional distribution of outputs, i.e., $\text{TF}_\theta(H_t) \triangleq P_\theta(Y_{t+1} \in \cdot | H_t) \triangleq P_\theta^t(\cdot)$. All subsequent discussions in this section assumes no distribution shift; that is, \mathcal{E} is the environment under which TF_θ was pretrained. We assume that Y_{t+1} is either discrete or continuous.

With this notation, the ICL performance with t demonstrations from \mathcal{E} is defined as $\mathbb{E}[-\log P_\theta^t(Y_{t+1})] = \mathbb{E}[-\log \bar{P}_e^t(Y_{t+1})] + \mathbb{E}[D_{\text{KL}}(\bar{P}_e^t \| P_\theta^t)]$ (Jeon et al., 2024). Here, the first term is the (irreducible) aleatoric uncertainty and constant with respect to t in our setting. The second term can be further decomposed as

$$\begin{aligned} \mathbb{E}[D_{\text{KL}}(\bar{P}_e^t \| P_\theta^t)] &= \mathbb{E}\left[\int \log \frac{d\bar{P}_e^t}{dP_\theta^t}(y) \bar{P}_e^t(dy)\right] \\ &= \underbrace{\mathbb{E}\left[D_{\text{KL}}(\bar{P}_e^t \| \hat{P}_\mathcal{E}^t)\right]}_{\triangleq \epsilon_{\text{Bayes}}^t \text{ (Bayes risk)}} + \underbrace{\mathbb{E}\left[\log \frac{\hat{P}_\mathcal{E}^t(Y_{t+1})}{P_\theta^t(Y_{t+1})}\right]}_{\triangleq \epsilon_{\text{XS}}^t \text{ (Excess risk)}}, \end{aligned} \quad (5)$$

where the second equality comes from the law of total expectation and $\hat{P}_\mathcal{E}^t(Y_{t+1}) \triangleq \mathbb{P}(Y_{t+1} \in \cdot | H_t, \mathcal{E})$ is the posterior over Y_{t+1} given H_t .

In (5), the Bayes risk $\epsilon_{\text{Bayes}}^t$ measures how well the Bayes-optimal predictor performs under uncertainty on e . It is non-negative and *decreases monotonically* with more demonstrations; that is, $\epsilon_{\text{Bayes}}^{t+1} \leq \epsilon_{\text{Bayes}}^t$ for all $t \in \mathbb{N}$ (Jeon et al., 2022). Demonstration size t required to bring this risk below a threshold q is captured by $N_{\text{BMA}}(q) \triangleq \min_{t \in \mathbb{N}} \{\epsilon_{\text{Bayes}}^t \leq q\}$. Here, q represents the absolute value of the performance requirement (e.g., MSE), whereas \mathcal{Q} in §3 denotes the performance quantile.

The excess risk ϵ_{XS}^t measures the performance of the transformer relative to the Bayes optimal predictor. Due to the non-negativity of excess risk and independence between TF_θ and $\epsilon_{\text{Bayes}}^t$, this term determines when ICL emerges and how well it can perform. For instance, if TF_θ achieves an excess risk curve such that $\epsilon_{\text{XS}}^t - \epsilon_{\text{XS}}^0 \leq \epsilon_{\text{Bayes}}^0 - \epsilon_{\text{Bayes}}^t$, non-trivial ICL performance emerges, improving upon the zero-shot performance with demonstrations. Further, if $\epsilon_{\text{XS}}^t \rightarrow 0$ as $t \rightarrow \infty$, then ICL is Bayes-risk consistent and asymptotically matches BMA. In §4.2, we dissect ϵ_{XS}^t based on our empirical results (§3).

4.2. On Excess Risk

Interpreting the transformer’s prediction in the meta-ICL setup as the Gaussian distribution (e.g., by adding a small

random Gaussian noise to the prediction), the squared prediction difference in Figure 4 is directly proportional to the excess risk, up to a constant scale and shift. The same applies to each baseline’s squared prediction difference, interpreted as its own excess risks.

In this regard, Figure 4 illustrates that the transformer’s excess risk remains roughly bounded within a modest interval in a certain length generalization regime (e.g., $t \leq 2T_{\text{train}}$), suggesting that it would perform ICL non-trivially due to the monotonicity of $\epsilon_{\text{Bayes}}^t$. However, once the context length becomes much longer than the one seen during pretraining (e.g., $t > 2T_{\text{train}}$ in Figure A1), the excess risk deteriorates sharply. This explains why ICL is not a consistent learner, being dominated by the principled learning algorithms in large sample regimes, as we observed in §3.2.

We formally encode the above empirical observations about the non-vanishing excess risk curve into Assumption 4.1.

Assumption 4.1. There exist constants $(\bar{t}, \Delta_{\text{XS}}) \in (\mathbb{N}, \mathbb{R}_+)$ such that $0 \leq \Delta_{\text{XS}} \leq \epsilon_{\text{XS}}^{t'}$ for all $t' \geq \bar{t}$.

The assumption states that, after some reference point \bar{t} , the excess risks of TF_θ can be lower bounded. In other words, it assumes that TF_θ does not magically reduce its excess risk in the out-of-distribution context length regimes.

Crucially, as we show in §4.3, Δ_{XS} controls a lower bound of ICL’s suboptimal efficiency in learning from demonstrations. For a transformer with a strong length generalization ability, $\epsilon_{\text{XS}}^{t'}$ in the assumption can also be upper bounded, making the subsequent suboptimality analysis nearly tight. In this regard, our analysis encompasses plausible (near) future advances in length generalization capability. Therefore, our analysis under Assumption 4.1 is a general result highlighting the ICL mechanism’s intrinsic flaws, isolating them from the transformer’s length generalization capability.

4.3. Analyzing Suboptimality of ICL

Next, we explain the critical suboptimality of ICL observed in §3, where ICL initially matches the efficiency of the optimal learning algorithm but starts to significantly deteriorate in many-shot regimes. To this end, we define suboptimality of ICL at performance requirement q as the additional number of demonstrations required for ICL to achieve requirement q compared to the Bayes optimal estimator, denoted as $\text{SubOpt}(q) \triangleq \min_t \{t - N_{\text{BMA}}(q) \mid \epsilon_{\text{Bayes}}^t + \epsilon_{\text{XS}}^t \leq q\}$. Here, we define suboptimality at q with respect to the reducible part of the ICL performance (i.e., $\mathbb{E}[D_{\text{KL}}(\bar{P}_e^t \parallel P_\theta^t)]$), which is equivalent to defining it with respect to the ICL performance up to constant scaling in q .

The following theorem constructs a lower bound of $\text{SubOpt}(q)$ under Assumption 4.1 where \mathbb{I} denotes the mutual information.

Theorem 4.2. Let us assume a tuple $(\bar{t}, \Delta_{\text{XS}})$ satisfies Assumption 4.1. For a sufficiently small q such that $N_{\text{BMA}}(q) \geq \bar{t}$, it holds that

$$\text{SubOpt}(q) \geq \text{LB}(q) \triangleq \min_{t \in \mathbb{N}} \left\{ t \mid \mathbb{I}(Y_{N_{\text{BMA}}(q)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q)-1}) > \Delta_{\text{XS}} \right\} \quad (6)$$

where \tilde{D}_{t+1} is a sample from the same distribution as D_{t+1} .

Theorem 4.2 intuitively characterizes suboptimality (cf. Figure A2 in Appendix for an illustration of the concept). Specifically, suppose the Bayes optimal learner requires $N_{\text{BMA}}(q)$ demonstrations to achieve the performance q . Then, $\text{SubOpt}(q)$ represents the additional demonstrations required for ICL to compensate for the excess risk ϵ_{XS}^t . Here, the compensation represents how much the new demonstrations \tilde{D}_{t+1} reduce the uncertainty about $Y_{N_{\text{BMA}}(q)}$ given a prompt $H_{N_{\text{BMA}}(q)-1}$, which corresponds to the conditional mutual information in (6). The theorem is proven in §B.1.

Characterizing suboptimality with $\mathbb{I}(Y_{N_{\text{BMA}}(q)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q)-1})$ provides clear insights into ICL’s suboptimality. Specifically, transformers with small excess risks in the non-vanishing regime are less subject to suboptimality. Besides, since a higher performance requirement (i.e., a smaller q) increases $N_{\text{BMA}}(q)$, suboptimality naturally increases due to reduced conditional mutual information. The following theorem, which is proven in §B.2, makes this intuition precise by establishing necessary conditions for $\text{SubOpt}(q)$ being constant with respect to q .

Theorem 4.3. Let us assume a tuple $(\bar{t}, \Delta_{\text{XS}})$ satisfies Assumption 4.1 and let q be such that $N_{\text{BMA}}(q) \geq \bar{t}$. If $\text{LB}(q') = \text{LB}(q)$ for all $\Delta_{\text{XS}} < q' < q$, then either of the following condition holds:

1. **Negligible excess risk:** $\Delta_{\text{XS}} \leq \mathbb{I}(Y_t; \tilde{D}_1 \mid H_{t-1})$ for all $t \geq N_{\text{BMA}}(q)$, and $\text{LB}(q) = 0$,
2. **Negligible diminishing returns:** $\mathbb{I}(Y_t; \tilde{D}_1 \mid H_{t-1}) < \left(1 + \frac{1}{\text{LB}(q)}\right) \mathbb{I}(Y_t; \tilde{D}_1 \mid H_{t-1})$ for all $t \geq N_{\text{BMA}}(q)$, where $\tilde{t} \triangleq N_{\text{BMA}}(q) + \text{LB}(q)$ and $\text{LB}(q) > 0$.

Non-deteriorating suboptimality has stringent necessary conditions that rarely hold in practice. Specifically, the *negligible excess risk* condition requires that the information gain from a single demonstration, regardless of demonstration size, dominates the excess risk. While this may hold for few-shot regimes (explaining the significant efficiency of few-shot ICL), ensuring this assumption across all prompt lengths is quite strong given the diminishing nature of $\mathbb{I}(Y_t; \tilde{D}_1 \mid H_{t-1})$ with t in most learning scenarios (Rissanen, 1984; Clarke & Barron, 1990). For a similar reason, the *negligible diminishing returns* condition,

which requires a constant lower bound of $\mathbb{I}(Y_t; \tilde{D}_1 | H_{t-1})$ for all demonstration sizes t , is quite strong. Therefore, $\text{SubOpt}(q)$ inevitably grows as q decreases, leading to increasing suboptimality of ICL under a high performance requirement as observed in §3.

As a concrete intuition on suboptimality, we consider the following crude approximations: (A1) $\epsilon_{\text{Bayes}}^t \approx C_1/\sqrt{t}$ for some constant factor C_1 and (A2) $\epsilon_{\text{XS}} \lesssim \epsilon_{\text{XS}}^t$ for all $t \in \mathbb{N}_+$. Here, (A1) corresponds to sublinear convergence of the Bayes posterior estimator, which holds in many cases (Rissanen, 1984; Clarke & Barron, 1990), and (A2) corresponds to Assumption 4.1 with $(\bar{t}, \Delta_{\text{XS}}) = (0, \epsilon_{\text{XS}})$. Replacing (A1) with other common bounds, such as $\epsilon_{\text{Bayes}}^t \approx C_1/t$ or $\epsilon_{\text{Bayes}}^t \approx C_1 \exp(-t)$, yields similar results.

Under (A1) and (A2), for performance achievable by the transformer (i.e., $q > \epsilon_{\text{XS}}$), a simple calculation gives $\text{SubOpt}(q) \gtrsim \frac{C_1^2}{(q-\epsilon_{\text{XS}})^2} - \frac{C_1^2}{q^2} \geq \frac{C_1^2 \epsilon_{\text{XS}}}{q^2(q-\epsilon_{\text{XS}})}$. Here, the rapid growth of $\text{SubOpt}(q)$ as q decreases highlights the significant inefficiency of ICL in achieving high performance requirement. Moreover, another way of improving suboptimality by reducing ϵ_{XS} , from the perspective of the rough power law estimations from the scaling laws (Kaplan et al., 2020), would require an exponential increase in pretraining data size or computational resources. Thus, in either way, a transformer exhibits significant suboptimality in achieving high performance through ICL compared to principled learning algorithms.

4.4. Impacts of Scaling Computations

We find that simply scaling model size or pretraining prompt length does not fundamentally eliminate inefficiency in long context even though these modifications can reduce overall excess risk. As shown in Figure A1 in Appendix, larger capacities and longer pretraining contexts lower the magnitude of the excess risk but do not change its “shape” in many-shot regimes, leading to persistent suboptimality. Thus, enhancing transformers’ ability to handle longer contexts alone does not resolve this inefficiency. See Appendix A.3 for experimental settings and results.

5. Related Work

Asymptotic Behavior Analysis. Xie et al. (2022) show that ICL predictions converge to posterior probabilities in asymptotic demonstration size regimes. Subsequent works expand these results to encompass finite-sample guarantees (Li et al., 2023b; Zhang et al., 2023; Bai et al., 2023), broader prompt distribution structures (Li et al., 2023b;c; Zhang et al., 2023), and structural characteristics of transformers (Zhang et al., 2023). Recent studies analyze the average cumulative regret across demonstrations (Zhang et al., 2023; Jeon et al., 2024), treating ICL as an online learning algorithm. However, prac-

tical applications prioritize test sample performance over demonstration set performance. In this work, we directly analyze suboptimality of ICL in achieving a specific performance requirement through the excess sample complexity compared to the Bayes optimal learning algorithm.

Stylized ICL Benchmarks. With the meta-ICL framework (cf. §2.1), Garg et al. (2022) demonstrate that transformers are capable of learning simple function classes (e.g., linear models and random neural networks) from demonstrations, achieving error curves qualitatively similar to those of optimal learning algorithms under asymptotic pretraining sample conditions. Subsequent works extend the results to finite pretraining sample scenarios (Raventós et al., 2023) and mixture function classes (Pathak et al., 2023; Panwar et al., 2024). Further, new analytical frameworks that directly analyze ICL predictions reveal that ICL exhibits behavior similar to gradient descent (Von Oswald et al., 2023; Akyürek et al., 2022). In this work, we measure how many demonstrations are required for ICL to achieve a certain performance level, rather than analyzing ICL performance as a function of demonstration size. This new perspective unveils the fundamental inefficiency of ICL in the many-shot regimes, which is subtle to discover with previous analyses.

Scaling ICL. Recent advances in handling long-context prompts (Chen et al., 2023; Su et al., 2024; Press et al., 2021) have enabled studies demonstrating (near) monotonic improvements in ICL performance with increased demonstrations (Li et al., 2023a; Agarwal et al., 2024; Anil et al., 2024). Notably, Bertsch et al. (2024) show that many-shot ICL can surpass parameter-efficient fine-tuning methods (Hu et al., 2022) given the same number of demonstrations, highlighting ICL’s sample efficiency. Our work extends these findings by examining optimality of performance gains from a learning-theoretic perspective, revealing that ICL’s sample complexity diminishes as sample sizes increase.

6. Conclusion

The surprisingly strong ICL performance of LLMs suggest its potential to eliminate the need for task-specific models. To rigorously examine this potential, we developed a novel framework for benchmarking optimality of ICL as a learning algorithm against principled learning algorithms. We found that while few-shot ICL’s efficiency is comparable to the Bayes optimal learning algorithm, its efficiency quickly diminishes with more demonstrations. Our information-theoretic analysis showed that the non-vanishing excess curve in long context causes fundamental inefficiency in many-shot regimes. This highlights the need for a new adaptation method that can reduce excess risk with more demonstrations, enabling sample-efficient learning of novel tasks while preserving the update-free nature of ICL.

Impact Statement

In this work, we study optimality of in-context learning as a learning algorithm against principled learning algorithms. Because our study focuses on theoretical aspects rather than practical applications, we do not foresee direct ethical concerns or societal impacts arising from our findings.

References

- Agarwal, R., Singh, A., Zhang, L. M., Bohnet, B., Rosias, L., Chan, S., Zhang, B., Anand, A., Abbas, Z., Nova, A., et al. Many-shot in-context learning. In *Neural Information Processing Systems*, 2024.
- Ahn, K., Bubeck, S., Chewi, S., Lee, Y. T., Suarez, F., and Zhang, Y. Learning threshold neurons via edge of stability. *Advances in Neural Information Processing Systems*, 2023.
- Ahn, K., Cheng, X., Song, M., Yun, C., Jadbabaie, A., and Sra, S. Linear attention is (maybe) all you need (to understand transformer optimization). In *International Conference on Learning Representations*, 2024.
- Ahuja, K. and Lopez-Paz, D. A closer look at in-context learning under distribution shifts. *arXiv preprint arXiv:2305.16704*, 2023.
- Akaike, H. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.
- Akyürek, E., Schuurmans, D., Andreas, J., Ma, T., and Zhou, D. What learning algorithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*, 2022.
- Allen-Zhu, Z. and Li, Y. Physics of language models: Part 1, learning hierarchical language structures. *arXiv preprints, abs/2305.13673*, 2023.
- Anil, C., Wu, Y., Andreassen, A., Lewkowycz, A., Misra, V., Ramasesh, V., Slone, A., Gur-Ari, G., Dyer, E., and Neyshabur, B. Exploring length generalization in large language models. In *Advances in Neural Information Processing Systems*, 2022.
- Anil, C., Durmus, E., Rinsky, N., Sharma, M., Benton, J., Kundu, S., Batson, J., Tong, M., Mu, J., Ford, D. J., et al. Many-shot jailbreaking. In *Neural Information Processing Systems*, 2024.
- Bai, Y., Chen, F., Wang, H., Xiong, C., and Mei, S. Transformers as statisticians: Provable in-context learning with in-context algorithm selection. In *Advances in Neural Information Processing Systems*, 2023.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *International Conference on Machine Learning*, 2009.
- Bertsch, A., Ivgi, M., Alon, U., Berant, J., Gormley, M. R., and Neubig, G. In-context learning with long-context models: An in-depth exploration. *arXiv preprint arXiv:2405.00200*, 2024.
- Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer, 2007.
- Bishop, C. M. and Bishop, H. Transformers. In *Deep Learning: Foundations and Concepts*, pp. 357–406. Springer, 2023.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 2020.
- Chen, S., Wong, S., Chen, L., and Tian, Y. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*, 2023.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- Clarke, B. S. and Barron, A. R. Information-theoretic asymptotics of Bayes methods. *IEEE Transactions on Information Theory*, 36(3):453–471, 1990.
- Ding, J., Tarokh, V., and Yang, Y. Model selection techniques: An overview. *IEEE Signal Processing Magazine*, 35(6):16–34, 2018.
- Dolan, E. D. and Moré, J. J. Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91:201–213, 2002.
- Dziugaite, G. K. and Roy, D. M. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017.
- Garg, S., Tsipras, D., Liang, P. S., and Valiant, G. What can transformers learn in-context? a case study of simple function classes. In *Advances in Neural Information Processing Systems*, 2022.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.

- Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. Bayesian model averaging: A tutorial. *Statistical Science*, 14(4):382–401, 1999.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.
- Jeon, H. J., Zhu, Y., and Van Roy, B. An information-theoretic framework for supervised learning. *arXiv preprint arXiv:2203.00246*, 2022.
- Jeon, H. J., Lee, J. D., Lei, Q., and Van Roy, B. An information-theoretic analysis of in-context learning. In *International Conference on Machine Learning*, 2024.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Kazemnejad, A., Padhi, I., Natesan Ramamurthy, K., Das, P., and Reddy, S. The impact of positional encoding on length generalization in transformers. In *Advances in Neural Information Processing Systems*, 2023.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Gao, I., et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, 2021.
- Langford, J. and Caruana, R. (not) bounding the true error. In *Advances in Neural Information Processing Systems*, 2001.
- Li, M., Gong, S., Feng, J., Xu, Y., Zhang, J., Wu, Z., and Kong, L. In-context learning with many demonstration examples. *arXiv preprint arXiv:2302.04931*, 2023a.
- Li, Y., Ildiz, M. E., Papailiopoulos, D., and Oymak, S. Transformers as algorithms: Generalization and stability in in-context learning. In *International Conference on Machine Learning*, 2023b.
- Li, Y., Li, Y., and Risteski, A. How do transformers learn topic structure: Towards a mechanistic understanding. In *International Conference on Machine Learning*, 2023c.
- Min, S., Lewis, M., Zettlemoyer, L., and Hajishirzi, H. Metaicl: Learning to learn in context. *arXiv preprint arXiv:2110.15943*, 2021.
- Panwar, M., Ahuja, K., and Goyal, N. In-context learning through the Bayesian prism. In *International Conference on Learning Representations*, 2024.
- Pathak, R., Sen, R., Kong, W., and Das, A. Transformers can optimally learn regression mixture models. In *International Conference on Learning Representations*, 2023.
- Press, O., Smith, N. A., and Lewis, M. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2021.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf, 2019. [Online; accessed 24-November-2024].
- Raventós, A., Paul, M., Chen, F., and Ganguli, S. Pretraining task diversity and the emergence of non-Bayesian in-context learning for regression. In *Advances in Neural Information Processing Systems*, 2023.
- Rissanen, J. Universal coding, information, prediction, and estimation. *IEEE Transactions on Information Theory*, 30(4):629–636, 1984.
- Schwarz, G. Estimating the dimension of a model. *The Annals of Statistics*, pp. 461–464, 1978.
- Srivastava, A., Rastogi, A., Rao, A., Shoeb, A. A. M., Abid, A., Fisch, A., Brown, A. R., Santoro, A., Gupta, A., Garriga-Alonso, A., et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Turner, R. E. An introduction to transformers. *arXiv preprint arXiv:2304.10557*, 2023.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- Von Oswald, J., Niklasson, E., Randazzo, E., Sacramento, J., Mordvintsev, A., Zhmoginov, A., and Vladymyrov,

-
- M. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, 2023.
- Wasserman, L. Bayesian model selection and model averaging. *Journal of Mathematical Psychology*, 44(1):92–107, 2000.
- Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- Wei, J., Wei, J., Tay, Y., Tran, D., Webson, A., Lu, Y., Chen, X., Liu, H., Huang, D., Zhou, D., et al. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846*, 2023.
- Xie, S. M., Raghunathan, A., Liang, P., and Ma, T. An explanation of in-context learning as implicit Bayesian inference. In *International Conference on Learning Representations*, 2022.
- Zhang, Y., Zhang, F., Yang, Z., and Wang, Z. What and how does in-context learning learn? Bayesian model averaging, parameterization, and generalization. *arXiv preprint arXiv:2305.19420*, 2023.
- Zhao, H., Andriushchenko, M., Croce, F., and Flammarion, N. Is in-context learning sufficient for instruction following in llms? *arXiv preprint arXiv:2405.19874*, 2024.
- Zhou, Y., Alon, U., Chen, X., Wang, X., Agarwal, R., and Zhou, D. Transformers can achieve length generalization but not robustly. *arXiv preprint arXiv:2402.09371*, 2024.

A. Additional Details

A.1. On Usage of Stylized Setting

Comprehensive analyses with statistical significance. The benchmark in the stylized settings in principle enables comprehensive comparisons across different environments (e.g., \mathcal{S}) and architectures (e.g., different TF_θ), achieving arbitrarily high levels of statistical significance. In empirical studies, these factors are constrained to the configurations of the datasets or the computational budgets.

Comparison with the optimal method. The stylized setting enables comparison with principled learning algorithms. Specifically, BMA considered in (2) provides the minimum achievable performance of any learning algorithms *at all prompt lengths*. This strong guarantee is typically not possible in empirical studies, as even human performances could not be an oracle or simply not possible to attain with only the data provided to the transformer. Also, the theoretical studies themselves do not allow for precise performance comparison, except analyzing the general asymptotic behavior that is shared among reasonable learning algorithms.

From stylized settings to practical LLMs. Although we study stylized settings in a rigorous manner, it does not capture all aspects of LLMs. For example, the ICL objective in (1) is not an autoregressive loss used for pretraining LLMs, omitting the losses of predictions at each Y_t . Therefore, one potential concern is the generalization of results obtained in stylized settings. While it cannot be shown precisely, the findings from such stylized settings have been surprisingly well generalized to the real-world tasks (Ahn et al., 2024; Li et al., 2023c). For instance, Ahn et al. (2024) perform synthetic experiments even with simplified transformers to study optimization methods for LLMs that surprisingly well reproduce the results from the real-world natural language data.

Given the significance of actionable insights from the stylized settings such as foretelling impacts of scaling ICL to the asymptotic region of the demonstration size, which is extremely challenging with real-world LLMs, we hold positive views on the role of stylized settings in LLM research whose significant advantages outweigh the potential concerns on its generalization to the LLMs in practice.

A.2. Detailed Configurations

Model. For the model, we use the GPT-2 (Radford et al., 2019) architecture for TF_θ , which is a standard architecture in the meta ICL and other stylized experimental settings; that is, we define TF_θ as a decoder-only transformers (Vaswani et al., 2017) with 12 layers, 8 attention heads, and 256-dimensional embedding space. For readers unfamiliar with transformers, we refer to the excellent tutorials (Turner, 2023; Bishop & Bishop, 2023). We remark that viewing TF_θ as a function from a sequence of vectors with an arbitrary length to a vector with the same dimension does not significantly impact the understanding of core findings in this paper.

Optimization. For minimizing the ICL objective $l(\theta)$, we compute the stochastic gradient with 64 prompts and update θ by using the Adam optimizer (Kingma & Ba, 2015) with fixed learning rate of 10^{-4} for one million training iterations. Also, in order to boost the convergence speed, we use curriculum learning (Bengio et al., 2009) as recommended in (Garg et al., 2022; Panwar et al., 2024) by increasing the length of the prompt by 2 every 2,000 training iterations until it reaches $(2M + 1)$ (and the order of Fourier series by 1 until it reaches M).

A.3. Impacts of Scaling Computations

We show that the non-vanishing excess risk curve of the transformer in long context causes the efficiency of learning to diminish with more demonstrations. Therefore, a natural question is whether enhancing transformers’ capacity of handling longer context can make the excess risk decrease with more demonstrations and thus resolve the fundamental inefficiency. We analyze the impacts of scaling the pretraining context lengths (by setting T_{train} to 100 and 200) and the model sizes (by scaling the number of layers, the number of heads, the embedding dimension by factors of 0.5, 2, and 3) on the excess risk. Note that we did not explore different positional encoding methods since we already use no positional encoding scheme that is effective at length generalization (Anil et al., 2022; Kazemnejad et al., 2023), which is from an inductive bias for the sample order-stable learning algorithms.

Figure A1 (left) shows that increasing T_{train} significantly reduces the excess risk values, especially for long-context regimes as desired. However, overall shape of the excess risk curve remains non-vanishing in the long-context regime. We observe from Figure A1 (right) similar effects of increasing the model sizes. Interestingly, larger models do not increase the length




Figure A1. Impacts of the pretraining prompt length (left) and the model size (right) on the excess risk curve in $\mathcal{E} = ([M], \sigma_y^2, \sigma_w^2) = ([10], 0.03, 10)$. For (left), the pretraining losses are 1.06, 0.58, and 0.34 for models trained with $T_{\text{train}} = 50$, $T_{\text{train}} = 100$, and $T_{\text{train}} = 200$, respectively. For (right), the pretraining losses are 1.36, 1.19, 0.99, and 0.90 for half-capacity, standard, double-capacity, and triple-capacity models respectively.

generalization regime, which is consistent with previous results (Zhou et al., 2024).

The results suggest that simply increasing computations with a larger model and a longer pretraining prompt length does not fundamentally change the *shape* of the excess risk, even though their overall scales improve. Therefore, while the degree of suboptimality can be relaxed with reduced excess risk, the inefficiency in many-shot regimes persist.

B. Proof of Claims

B.1. Proof of Theorem 4.2

Proof. We first characterize suboptimality by the Bayes risk as follows:

$$\text{SubOpt}(q) = \min_{t \in \mathbb{Z}_+} \{t - N_{\text{BMA}}(q) \mid \epsilon_{\text{Bayes}}^t + \epsilon_{\text{XS}}^t \leq q\} \quad (7)$$

$$= \min_{t \in \mathbb{Z}_+} \{t \mid \epsilon_{\text{Bayes}}^t \leq q - \epsilon_{\text{XS}}^t\} - N_{\text{BMA}}(q). \quad (8)$$

Since $q < \epsilon_{\text{Bayes}}^{N_{\text{BMA}}(q)-1}$, the monotonicity of $\epsilon_{\text{Bayes}}^t$ and the non-negativity of ϵ_{XS}^t give

$$\min_{t \in \mathbb{Z}_+} \{t \mid \epsilon_{\text{Bayes}}^t \leq q - \epsilon_{\text{XS}}^t\} = \min_{t \geq N_{\text{BMA}}(q)} \{t \mid \epsilon_{\text{Bayes}}^t \leq q - \epsilon_{\text{XS}}^t\} \geq \min_{t \geq N_{\text{BMA}}(q)} \{t \mid \epsilon_{\text{Bayes}}^t < \epsilon_{\text{Bayes}}^{N_{\text{BMA}}(q)-1} - \epsilon_{\text{XS}}^t\}. \quad (9)$$

To prove the theorem, we note the following.

(N1). Bayes error reduction as the conditional mutual information: The Bayes error can be expressed as the reduction of (differential) entropy as follows.

$$\epsilon_{\text{Bayes}}^t = \mathbb{E} \left[D_{\text{KL}}(\bar{P}_e^t \parallel \hat{P}_{\mathcal{E}}^t) \right] = h(Y_{t+1}|H_t) - h(Y_{t+1}|H_t, e), \quad \text{for continuous } Y_{t+1} \quad (10)$$

$$\epsilon_{\text{Bayes}}^t = \mathbb{E} \left[D_{\text{KL}}(\bar{P}_e^t \parallel \hat{P}_{\mathcal{E}}^t) \right] = \mathbb{H}(Y_{t+1}|H_t) - \mathbb{H}(Y_{t+1}|H_t, e), \quad \text{for discrete } Y_{t+1} \quad (11)$$

where h is the differential entropy and \mathbb{H} is the Shannon entropy.

Therefore, for any $u \leq v$ and continuous Y_{t+1} , we have

$$\epsilon_{\text{Bayes}}^u - \epsilon_{\text{Bayes}}^v = h(Y_{u+1}|H_u) - h(Y_{u+1}|H_u, e) - (h(Y_{v+1}|H_v) - h(Y_{v+1}|H_v, e)) \quad (12)$$

$$= h(Y_{u+1}|X_{u+1}, D_u) - h(Y_{v+1}|X_{v+1}, D_v) \quad (13)$$

$$= \mathbb{I}(Y_{u+1}; \tilde{D}_{v-u}|X_{u+1}, D_u) \quad (14)$$

where $\tilde{D}_{v-u} \triangleq (\tilde{X}_1, \tilde{Y}_1, \dots, \tilde{X}_{v-u}, \tilde{Y}_{v-u})$ is independently sampled from the same distribution as D_{v-u} , the second equality comes from the conditional independence $Y_{n+1} \perp D_n | X_{n+1}, e$ for any $n \in \mathbb{N}_+$, and the last equality comes from the chain rule. For the discrete Y 's, the same process can be applied by replacing h with \mathbb{H} .

(N2). Lower bound of the excess risk: Let q be such that $N_{\text{BMA}}(q) \geq \bar{t}$. Therefore, by Assumption 4.1, we have

$$\left\{ t \in \mathbb{N} \mid t \geq N_{\text{BMA}}(q), \epsilon_{\text{Bayes}}^t < \epsilon_{\text{Bayes}}^{N_{\text{BMA}}(q)-1} - \epsilon_{\text{XS}}^t \right\} \subseteq \left\{ t \in \mathbb{N} \mid t \geq N_{\text{BMA}}(q), \epsilon_{\text{Bayes}}^{N_{\text{BMA}}(q)-1} - \epsilon_{\text{Bayes}}^t > \Delta_{\text{XS}} \right\}. \quad (15)$$

By applying **(N1)** and **(N2)** to (8), we get the desired result as

$$\begin{aligned} \text{SubOpt}(q) &= \min_{t \geq N_{\text{BMA}}(q)} \left\{ t \mid \epsilon_{\text{Bayes}}^t \leq q - \epsilon_{\text{XS}}^t \right\} - N_{\text{BMA}}(q) \\ &\geq \min_{t \geq N_{\text{BMA}}(q)} \left\{ t \mid \epsilon_{\text{Bayes}}^{N_{\text{BMA}}(q)-1} - \epsilon_{\text{Bayes}}^t > \Delta_{\text{XS}} \right\} - N_{\text{BMA}}(q) = \min_{t \in \mathbb{N}} \left\{ t \mid \epsilon_{\text{Bayes}}^{N_{\text{BMA}}(q)-1} - \epsilon_{\text{Bayes}}^{t+N_{\text{BMA}}(q)} > \Delta_{\text{XS}} \right\} \\ &= \min_{t \in \mathbb{N}} \left\{ t \mid \mathbb{I}(Y_{N_{\text{BMA}}(q)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q)-1}) > \Delta_{\text{XS}} \right\}. \end{aligned} \quad (16)$$

□

B.2. Proof of Theorem 4.3

Proof. Consider $q_1, q_2 \in (\Delta_{\text{XS}}, q)$ such that $q_1 < q_2 < q$ and $N_{\text{BMA}}(q_1) > N_{\text{BMA}}(q_2)$. The goal is to show necessary conditions for $LB(q_1) \leq LB(q_2)$.

Note that $LB(q_1) < LB(q_2)$ is impossible because $\mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_1)-1}) \leq \mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_2)-1})$ for any $t \in \mathbb{N}$. Specifically, we have

$$\mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_1)-1}) \leq \mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_2)-1}), \quad \forall t \in \mathbb{N} \quad (17)$$

, which implies

$$\left\{ t \in \mathbb{N} \mid \mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_1)-1}) > \Delta_{\text{XS}} \right\} \subseteq \left\{ t \in \mathbb{N} \mid \mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_2)-1}) > \Delta_{\text{XS}} \right\} \quad (18)$$

, and in turn $LB(q_1) \geq LB(q_2)$.

Therefore, we next show the necessary condition for $LB(q_1) = LB(q_2)$.

(NC 1). Negligible excess risk: Let us suppose $\Delta_{\text{XS}} \leq \mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q_1)-1}) \leq \mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q_2)-1})$. In this case, $LB(q_1) = LB(q_2) = 0$ as desired. Since q_1 and q_2 are chosen arbitrary, the first necessary condition is given by

$$\Delta_{\text{XS}} \leq \mathbb{I}(Y_t; \tilde{D}_1 \mid H_{t-1}), \quad t \geq \bar{t}. \quad (19)$$

(NC 2). No diminishing returns: If **(NC 1)** does not hold, we have $\Delta_{\text{XS}} > \mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q_1)-1})$. In this case, we rule out the possibility $\mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q_1)-1}) < \Delta_{\text{XS}} \leq \mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q_2)-1})$ because this gives $LB(q_2) = 0$ and $LB(q_1) > 0$, which contradicts $LB(q_1) = LB(q_2)$.

Thus, we consider the case $\mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q_1)-1}) \leq \mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q_2)-1}) < \Delta_{\text{XS}}$. In this case, $LB(q_1) = LB(q_2)$ requires the following condition

$$\mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_{LB(q_2)} \mid H_{N_{\text{BMA}}(q_2)-1}) < \mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_{LB(q_2)+1} \mid H_{N_{\text{BMA}}(q_1)-1}), \quad (20)$$

where the condition comes from $\mathbb{I}(Y_{N_{\text{BMA}}(q_1)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_1)-1}) \leq \mathbb{I}(Y_{N_{\text{BMA}}(q_2)}; \tilde{D}_{t+1} \mid H_{N_{\text{BMA}}(q_2)-1})$ for any $t \in \mathbb{N}$.

By the construction of q_1 and q_2 , we get

$$\mathbb{I}(Y_{N_{\text{BMA}}(q)}; \tilde{D}_{LB(q)} \mid H_{N_{\text{BMA}}(q)-1}) < \mathbb{I}(Y_{N_{\text{BMA}}(q)+k}; \tilde{D}_{LB(q)+1} \mid H_{N_{\text{BMA}}(q)-1+k}), \quad \forall k \in \mathbb{N}_+. \quad (21)$$

Due to the chain rule of the mutual information, for any $\tilde{k} \in \mathbb{N}_+$, it holds that

$$\mathbb{I}(Y_{N_{\text{BMA}}(q)}; \tilde{D}_{\tilde{k}} \mid H_{N_{\text{BMA}}(q)-1}) = \sum_{i=0}^{\tilde{k}-1} \mathbb{I}(Y_{N_{\text{BMA}}(q)+i}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q)-1+i}) \geq \tilde{k} \mathbb{I}(Y_{N_{\text{BMA}}(q)+\tilde{k}-1}; \tilde{D}_1 \mid H_{N_{\text{BMA}}(q)+\tilde{k}-2}). \quad (22)$$

Similarly,

$$\mathbb{I}(Y_{N_{\text{BMA}}(q)+k}; \tilde{D}_{\tilde{k}+1} | H_{N_{\text{BMA}}(q)-1+k}) = \sum_{i=0}^{\tilde{k}} \mathbb{I}(Y_{N_{\text{BMA}}(q)+k+i}; \tilde{D}_1 | H_{N_{\text{BMA}}(q)-1+k+i}) \leq (1 + \tilde{k}) \mathbb{I}(Y_{N_{\text{BMA}}(q)+k}; \tilde{D}_1 | H_{N_{\text{BMA}}(q)-1+k}). \quad (23)$$

Therefore, we get the second necessary condition as

$$\mathbb{I}(Y_t; \tilde{D}_1 | H_{t-1}) \leq \mathbb{I}(Y_{\tilde{t}+\tilde{k}-1}; \tilde{D}_1 | H_{\tilde{t}+\tilde{k}-2}) < \left(1 + \frac{1}{\tilde{k}}\right) \mathbb{I}(Y_t; \tilde{D}_1 | H_{t-1}), \quad \forall t \geq \tilde{t}, \quad (24)$$

where $\tilde{k} = LB(q) > 1$ for q such that $N_{\text{BMA}}(q) \geq \tilde{t}$. □

C. Additional Figures

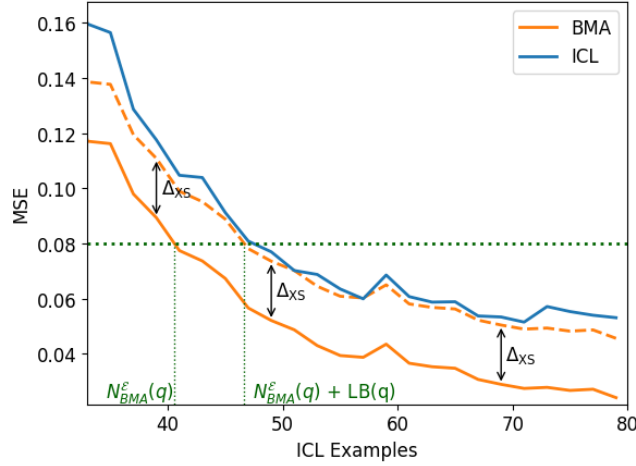


Figure A2. Graphical illustration of Theorem 4.2 when $q = 0.08 - \sigma^2$, where $\sigma^2 = \mathbb{E}[-\log \tilde{P}_e^t(Y_{t+1})]$ is the irreducible aleatoric uncertainty. The solid orange and blue lines represent MSEs of BMA and ICL, respectively. Here, the dashed orange line corresponds to the $\sigma^2 + \epsilon_{\text{Bayes}}^t + \Delta_{\text{XS}}$, which serves as a lower bound on MSEs of ICL. The shift by Δ_{XS} induces suboptimality that requires at least $LB(q)$ additional number of demonstrations for ICL to achieve the requirement q , compared to BMA.




Figure A3. The number of demonstrations (y-axis) required to achieve each performance quantile (x-axis). The shaded area represents the standard error. We note that performance quantile $Q = 0.6$ is achieved by T_{Train} number of demonstrations on average.

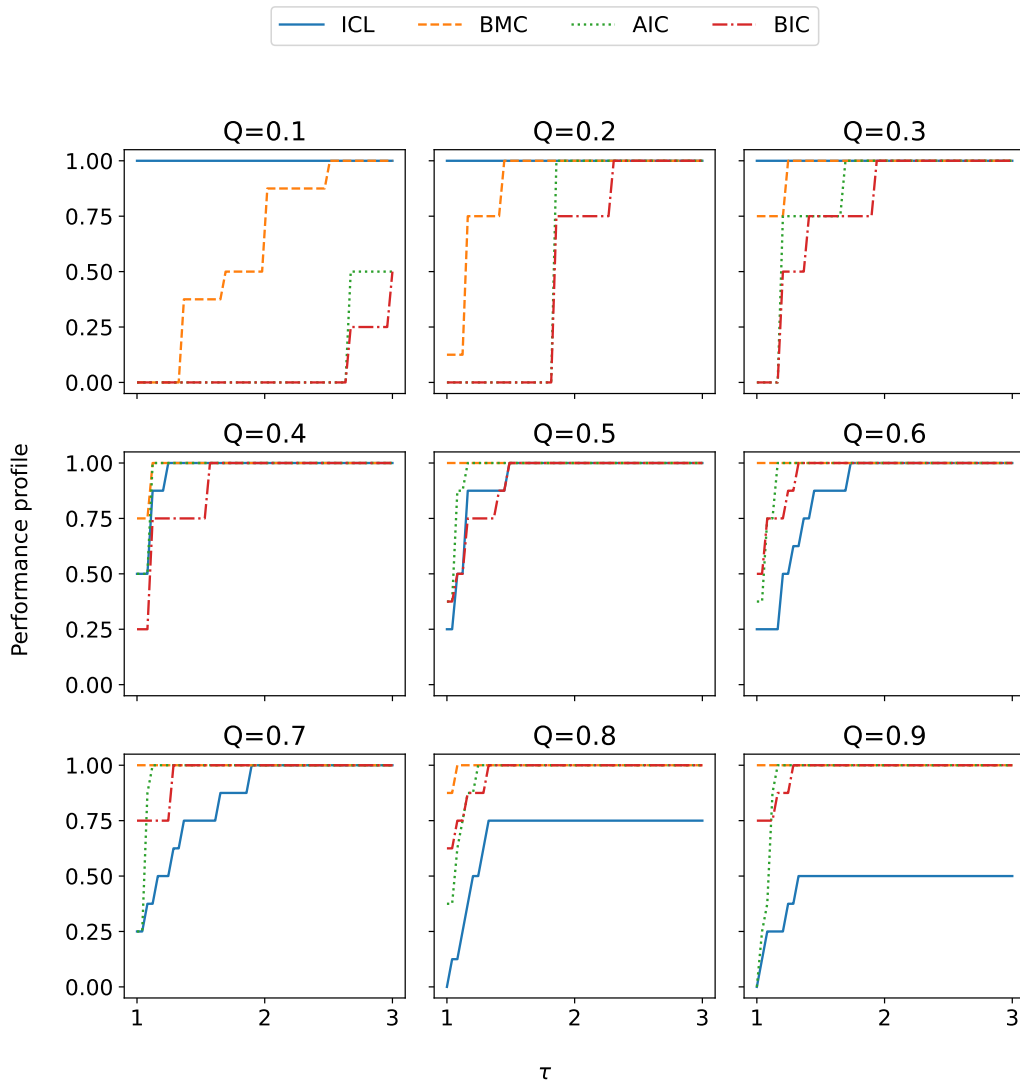


Figure A4. Performance profiles ρ_b across different performance ratios τ under different target performance quantiles Q . Each curve represents the probability that a method achieves the desired performance within a factor τ of the best method's sample complexity (x-axes).

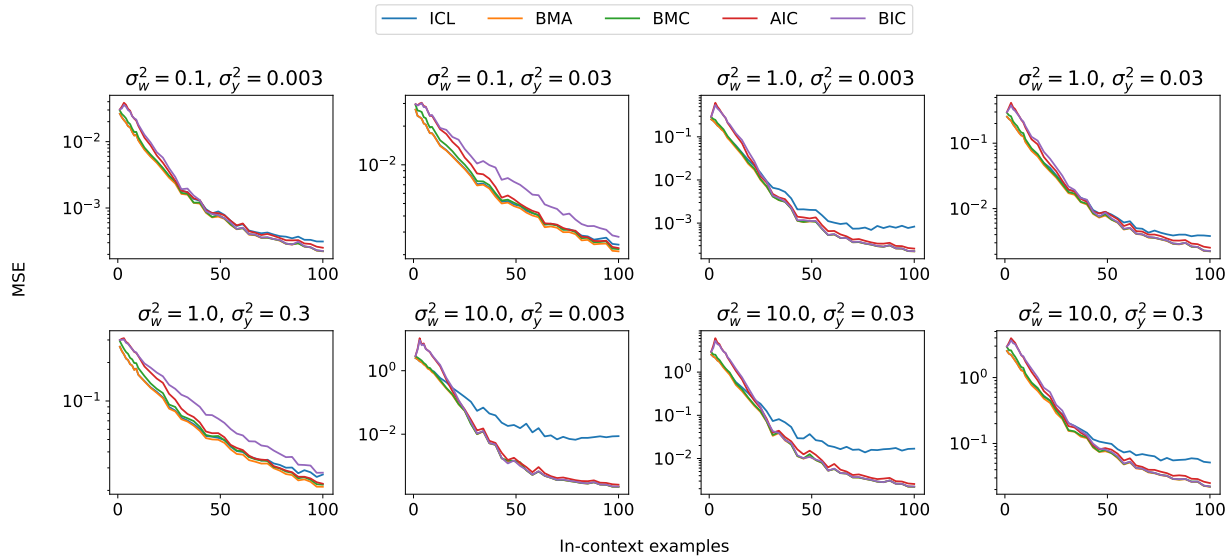


Figure A5. Mean squared errors for different demonstration sizes.




Figure A6. Squared prediction differences between BMA and other methods for different demonstration sizes.