

Transformers as Algorithms: Generalization and Implicit Model Selection in In-context Learning

Yingcong Li* M. Emrullah Ildiz* Dimitris Papailiopoulos† Samet Oymak*

Abstract

In-context learning (ICL) is a type of prompting where a transformer model operates on a sequence of (input, output) examples and performs inference on-the-fly. This *implicit training* is in contrast to explicitly tuning the model weights based on examples. In this work, we formalize in-context learning as an *algorithm learning* problem, treating the transformer model as a learning algorithm that can be specialized via training to implement—at inference-time—another target algorithm. We first explore the statistical aspects of this abstraction through the lens of multitask learning: We obtain generalization bounds for ICL when the input prompt is (1) a sequence of i.i.d. (input, label) pairs or (2) a trajectory arising from a dynamical system. The crux of our analysis is relating the excess risk to the stability of the algorithm implemented by the transformer, which holds under mild assumptions. Secondly, we use our abstraction to show that transformers can act as an adaptive learning algorithm and perform model selection across different hypothesis classes. We provide numerical evaluations that (1) demonstrate transformers can indeed implement near-optimal algorithms on classical regression problems with i.i.d. and dynamic data, (2) identify an inductive bias phenomenon where the transfer risk on unseen tasks is independent of the transformer complexity, and (3) empirically verify our theoretical predictions.

Keywords. in-context learning, generalization, model-selection, system identification, transformers

1 Introduction

Transformer (TF) models were originally developed for NLP problems to address long-range dependencies through the attention mechanism. In recent years, language models have become increasingly large, with some boasting billions of parameters (e.g., GPT-3 has 175B, and PaLM has 540B parameters [7, 10]). It is perhaps not surprising that these large language models (LLMs) have achieved state-of-the-art performance on a wide range of natural language processing tasks. What is surprising is the ability of some of these LLMs to perform *in-context learning* (ICL), i.e., to adapt and perform a specific task given a short prompt, in the form of instructions, and a small number of examples [7]. The ability of ICL to arise in these models without explicit training, allows them to efficiently perform new tasks without the need for model weight updates.

Specifically, ICL shows the ability of a TF to provide a prediction on an example given a few (input, output) examples within a prompt. As depicted in Figure 1, for NLP, the examples may correspond to pairs of

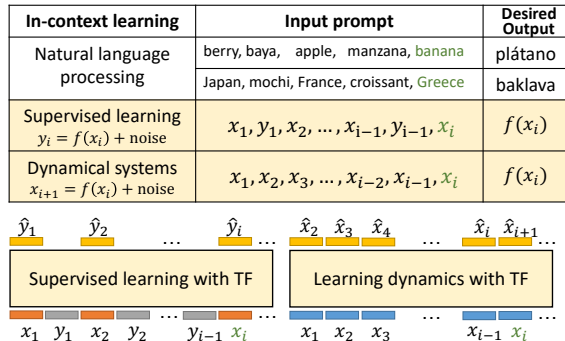


Figure 1: Examples of in-context learning. We focus on the lower two settings in the table where a transformer admits a supervised dataset or dynamical system trajectory as a prompt. Then, it auto-regressively predicts the output following an input example x_i based on the prompt (x_1, \dots, x_i) .

*University of California, Riverside, {yli692@mildi001@,oymak@ece.}ucr.edu.

†University of Wisconsin-Madison, dimitris@papail.io




Figure 2: Examples of implicit model-selection in three ICL settings: **(a) Noisy linear regression:** $y_i \sim \mathcal{N}(\mathbf{x}_i^\top \boldsymbol{\beta}, \sigma^2)$ with $\mathbf{x}_i, \boldsymbol{\beta} \sim \mathcal{N}(0, \mathbf{I})$. **(b) Linear data with covariance prior:** $y_i = \mathbf{x}_i^\top \boldsymbol{\beta}$ with $\boldsymbol{\beta} \sim \mathcal{N}(0, \Sigma)$ with non-isotropic Σ . **(c) Partially observed linear dynamics:** $\mathbf{x}_t = \mathbf{C} \mathbf{s}_t$ and $\mathbf{s}_{t+1} \sim \mathcal{N}(\mathbf{A} \mathbf{s}_t, \sigma^2 \mathbf{I})$ with randomly sampled \mathbf{C}, \mathbf{A} . Each setting trains a transformer with large number of random regression tasks and evaluates on a new task from the same distribution. In (a) and (b), ICL performances match Bayes-optimal decision (weighted linear ridge regression) that adapt to noise level σ and covariance prior Σ on the tasks. (c) shows that ICL outperforms auto-regressive least-squares estimators with varying memory H . Implicit model selection refers to transformer’s ability to implement a competitive ML algorithm by leveraging the task prior learned during its training. See Sec 7 for experimental details.

(question, answer)’s or translations. Recent works [20, 26] demonstrate that ICL can also be used to infer general functional relationships. For instance, [22] aim to solve supervised learning problems where we feed an entire training dataset $(\mathbf{x}_i, f(\mathbf{x}_i))_{i=1}^{n-1}$ as the input prompt, expecting that conditioning the TF model on this prompt would allow it to infer an accurate prediction for a new input point \mathbf{x}_n . As discussed in [20], this provides a supervised-learning flavor to ICL, where the model *implicitly trains* on the data provided within the prompt, and performs inference on test points.

Our work formalizes in-context learning from a statistical lens, abstracting the transformer as a learning algorithm where the goal is inferring the correct (input, output) functional relationship from prompts. We focus on a meta-learning setting where the model is trained on many tasks, allowing ICL to generalize to both new and previously-seen tasks. Our main contributions are:

- **Generalization bounds (Sec. 3 & 6):** Suppose the model is trained on T tasks each with a data-sequence containing n examples. During training, each sequence is fed to the model auto-regressively as depicted in Figure 1. By abstracting ICL as an *algorithm learning* problem, we establish a multitask (MTL) generalization rate of $1/\sqrt{nT}$. In order to achieve the proper dependence on the sequence length ($1/\sqrt{n}$ factor), we overcome temporal dependencies by relating generalization to algorithmic stability and we provide numerical and theoretical evidence for the stability of transformer-induced algorithms. Experiments demonstrate that ICL indeed benefits from learning across the full task sequence in line with theory.
- **From multitask to meta-learning (Sec. 4):** We provide insights into how our MTL bounds can inform generalization ability of ICL on previously unseen tasks (i.e. transfer learning). Our experiments reveal an intriguing *inductive bias phenomenon*: The gap between the transfer risk and MTL risk is *independent of the model complexity* and only depends on the *task complexity* and the number of tasks T .
- **Insights on model selection (Sec. 5):** The *algorithm learning* viewpoint motivates model-selection as a natural feature of in-context learning (e.g. whether to implement a linear model or neural net for prediction). We discuss how this enables the selection of the correct hypothesis set for the tasks when algorithms correspond to empirical risk minimization procedures. We verify this model-selection ability through extensive experiments as illustrated in Figure 2. Among others, the experiments reveal that transformer-based algorithm can learn task covariance as a prior to compete with optimally-weighted ridge and outperform least-squares of different orders to learn partially-observed dynamical systems.

The remainder of the paper is organized as follows. The next section discusses connections to prior art and Section 2 introduces the problem setup. Section 3 provides our main theoretical guarantees for ICL and

stability of transformers. Section 4 extends our arguments and experiments to the transfer learning setting and Section 5 elucidates the implicit model/algorithm selection ability of transformers via ICL. Section 6 extends our results to learning stable dynamical systems where prompts correspond to temporally-dependent system trajectories. Section 7 provides numerical evaluations, and we discuss future work in Section 8.

1.1 Related work

In-context learning and transformers: With the success of large language models, prompting methods has witnessed immense interest. ICL [7, 37] is a prompting strategy where a transformer serves as an on-the-fly predictive model through conditioning on a sequence of input/output examples $(\mathbf{x}_1, f(\mathbf{x}_1), \dots, \mathbf{x}_{n-1}, f(\mathbf{x}_{n-1}), \mathbf{x}_n)$. Our work is inspired by [20] which studies ICL in synthetic settings and demonstrates transformers can serve as complex classifiers through ICL. In parallel, [22] uses ICL as an AutoML (i.e. model-selection, hyperparameter tuning) framework where they plug in a dataset to transformer and use it as a classifier for new test points. Our formalism on *algorithm learning* provides a justification on how transformers can accomplish this with proper meta-training. [57] interprets ICL as implicit Bayesian inference and develops guarantees when the training distribution is a mixture of HMMs. Very recent works [54, 2, 12] aim to relate ICL to running gradient descent algorithm over the input prompt. [2] also provides related observations regarding the optimal decision making ability of ICL for linear models. We remark that, for linear regression problems, Bayes risk and optimal regularization schemes have been investigated by earlier works [28, 41, 56]. [14] shows that language models like GPT3 can be fine-tuned to perform competitively across a variety of non-NLP tasks (e.g., classification and regression on tabular/vision/synthetic data) where the training data is provided in the form of a language prompts, in a similar manner to that of ICL. Unlike prior work, we provide finite sample generalization guarantees for ICL and our theory extends to temporally-dependent prompts (e.g. when prompts are trajectories of dynamical systems). Dynamical systems in turn relate to the recent work by [26] who use ICL for distilling reinforcement learning algorithms.

Multitask/Meta learning: In order for ICL to work well, the transformer model needs to train with large amounts of related prompt instances. This makes it inherently connected to meta learning [18]. However, a key distinction is that, in ICL, adaptation to a new task happens implicitly through input prompt. Our analysis has some parallels with recent literature on multitask representation learning [33, 15, 50, 9, 27, 23, 40, 49, 11, 34, 16] since we develop excess MTL risk bounds by training the model with T tasks and quantify these bounds in terms of complexity of the hypothesis space (i.e. transformer architecture), the number of tasks T , and the number of samples per task. In relation to [48, 8], our experiments on linear regression with covariance-prior (Figure 2(b)) demonstrate ICL’s ability to implicitly implement optimally-weighted linear representations.

Learning with temporal dependencies: Finally, our results require statistical analysis of temporally dependent data as we create n prompts from a dataset of n examples in an autoregressive fashion. This makes it related to recent line of work on statistical aspects of time-series prediction and (non)linear system identification ([46, 30, 13, 43, 25, 31, 52, 51, 44, 38, 19, 42] among others). Our work has two distinguishing aspects over these. First, in the ICL setting, the transformer is trained with multiple tasks simultaneously in contrast to the traditional single-task supervised learning. Secondly, our *algorithm learning* formulation leads to new challenges and insights when verifying the conditions to apply Azuma-type inequalities (e.g. Theorem 2). Specifically, our results are facilitated through connections to *algorithmic stability* conditions [5, 6]. We also provide experiments and theory that justify our stability conditions for transformer architectures both empirically and theoretically (see Theorem 1).

2 Problem Setup

Notation. Let \mathcal{X} be the input feature space, and \mathcal{Y} be the output/label space. We use boldface for vector variables. $[n]$ denotes the set $\{1, 2, \dots, n\}$. $c, C > 0$ denote absolute constants and $\|\cdot\|_{\ell_p}$ denotes the ℓ_p -norm.

In-context learning setting: We denote a length- m prompt containing $m - 1$ in-context examples and the m ’th input by $\mathbf{x}_{\text{prompt}}^{(m)} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{m-1}, \mathbf{x}_m)$. Here $\mathbf{x}_m \in \mathcal{X}$ is the input to predict and $\mathbf{z}_i \in \mathcal{Z}$ is the i ’th

in-context example provided within prompt. Let $\text{TF}(\cdot)$ denote a transformer (more generally an auto-regressive model) that admits $\mathbf{x}_{\text{prompt}}^{(m)}$ as its input and outputs a label $\hat{\mathbf{y}}_m = \text{TF}(\mathbf{x}_{\text{prompt}}^{(m)})$ in \mathcal{Y} . Then the model aims to predict the next output denoted by $\hat{\mathbf{y}}_m$. In this work, we investigate two scenarios where prompts are obtained from a sequence of i.i.d. (input, label) pairs or a trajectory of a dynamical system.

- *Independent (\mathbf{x}, \mathbf{y}) pairs.* Similar to [20], we draw i.i.d. samples $(\mathbf{x}_i, \mathbf{y}_i)_{i=1}^n \in \mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ from a data distribution. Then a length- m prompt is written as $\mathbf{x}_{\text{prompt}}^{(m)} = (\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_{m-1}, \mathbf{y}_{m-1}, \mathbf{x}_m)$, and the model predicts $\hat{\mathbf{y}}_m = \text{TF}(\mathbf{x}_{\text{prompt}}^{(m)}) \in \mathcal{Y}$ for $1 \leq m \leq n$.

- *Dynamical systems.* In this setting, the prompt is simply the trajectory generated by a dynamical system, namely, $\mathbf{x}_{\text{prompt}}^{(m)} = (\mathbf{x}_1, \dots, \mathbf{x}_{m-1}, \mathbf{x}_m)$ and therefore, $\mathcal{Z} = \mathcal{X} = \mathcal{Y}$. Specifically, we investigate the state observed setting that is governed by dynamics $f(\cdot)$ via $\mathbf{x}_{m+1} = f(\mathbf{x}_m) + \text{noise}$. Here, $\mathbf{y}_m := \mathbf{x}_{m+1}$ is the label associated to \mathbf{x}_m , and the model admits $\mathbf{x}_{\text{prompt}}^{(m)}$ as input and predicts the next state $\hat{\mathbf{y}}_m := \hat{\mathbf{x}}_{m+1} = \text{TF}(\mathbf{x}_{\text{prompt}}^{(m)}) \in \mathcal{X}$.

We first consider the training phase of ICL where we wish to learn a good $\text{TF}(\cdot)$ model through multitask learning (MTL). Suppose we have T tasks associated with data distributions $\{\mathcal{D}_t\}_{t=1}^T$. Each task independently samples a training sequence $\mathcal{S}_t = (\mathbf{z}_{ti})_{i=1}^n$ according to its distribution. $\mathcal{S}_{\text{all}} = \{\mathcal{S}_t\}_{t=1}^T$ denote the set of all training sequences. We use $\mathcal{S}_t^{(m)} = (\mathbf{z}_{t1}, \dots, \mathbf{z}_{tm})$ to denote a subsequence of $\mathcal{S}_t := \mathcal{S}_t^{(n)}$ for $m \leq n$ and $\mathcal{S}^{(0)}$ denotes an empty subsequence.

ICL can be interpreted as an implicit optimization on the subsequence $\mathcal{S}^{(m)} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_m)$ to make prediction on \mathbf{x}_{m+1} . To model this, we abstract the transformer model as a learning algorithm that maps a sequence of data to a prediction function (e.g. gradient descent, empirical risk minimization). Concretely, let \mathcal{A} be a set of algorithm hypotheses such that algorithm $\text{Alg} \in \mathcal{A}$ maps a sequence of form $\mathcal{S}^{(m)}$ into a prediction function $f_{\mathcal{S}^{(m)}}^{\text{Alg}} : \mathcal{X} \rightarrow \mathcal{Y}$. With this, we represent TF via

$$\text{TF}(\mathbf{x}_{\text{prompt}}^{(m+1)}) = f_{\mathcal{S}^{(m)}}^{\text{Alg}}(\mathbf{x}_{m+1}). \quad (1)$$

This abstraction is without losing generality as we have the explicit relation $f_{\mathcal{S}^{(m)}}^{\text{Alg}}(\mathbf{x}) := \text{TF}((\mathcal{S}^{(m)}, \mathbf{x}))$. Given training sequences \mathcal{S}_{all} and a loss function $\ell(\mathbf{y}, \hat{\mathbf{y}})$, the ICL training can be interpreted as searching for the optimal algorithm $\text{Alg} \in \mathcal{A}$, and the training objective is formulated as follows:

$$\begin{aligned} \text{Alg} &= \arg \min_{\text{Alg} \in \mathcal{A}} \mathfrak{L}_{\mathcal{S}_{\text{all}}}(\text{Alg}) := \frac{1}{T} \sum_{t=1}^T \mathfrak{L}_t(\text{Alg}) && \text{(ICL-ERM)} \\ \text{where } \mathfrak{L}_t(\text{Alg}) &= \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{y}_{ti}, f_{\mathcal{S}_t^{(i-1)}}^{\text{Alg}}(\mathbf{x}_{ti})). \end{aligned}$$

Here, $\mathfrak{L}_t(\text{Alg})$ is the training loss of task t and $\mathfrak{L}_{\mathcal{S}_{\text{all}}}(\text{Alg})$ is the task-averaged MTL loss. Let $\mathcal{L}_t(\text{Alg}) = \mathbb{E}_{\mathcal{S}_t}[\mathfrak{L}_t(\text{Alg})]$ and $\mathcal{L}_{\text{MTL}}(\text{Alg}) = \mathbb{E}[\mathfrak{L}_{\mathcal{S}_{\text{all}}}(\text{Alg})] = \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\text{Alg})$ be the corresponding population risks. Observe that, task-specific loss $\mathfrak{L}_t(\text{Alg})$ is an empirical average of n terms, one for each prompt $\mathbf{x}_{\text{prompt}}^{(i)}$ and the corresponding label \mathbf{y}_i .

To develop generalization bounds, our primary interest is controlling the gap between empirical and population risks. For the problem (ICL-ERM), we wish to bound the excess MTL risk

$$R_{\text{MTL}}(\text{Alg}) = \mathcal{L}_{\text{MTL}}(\text{Alg}) - \min_{\text{Alg} \in \mathcal{A}} \mathcal{L}_{\text{MTL}}(\text{Alg}). \quad \text{(MTL-risk)}$$

Following the MTL training (ICL-ERM), we also evaluate the model on previously-unseen tasks; this can be thought of as the transfer learning problem. Concretely, let $\mathcal{D}_{\text{task}}$ be a distribution over tasks and draw a target task $\mathcal{T} \sim \mathcal{D}_{\text{task}}$ with data distribution $\mathcal{D}_{\mathcal{T}}$ and a sequence $\mathcal{S}_{\mathcal{T}} = \{\mathbf{z}_i\}_{i=1}^n \sim \mathcal{D}_{\mathcal{T}}$. Define the empirical and population risks on \mathcal{T} as $\mathfrak{L}_{\mathcal{T}}(\text{Alg}) = \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{y}_i, f_{\mathcal{S}_{\mathcal{T}}^{(i-1)}}^{\text{Alg}}(\mathbf{x}_i))$ and $\mathcal{L}_{\mathcal{T}}(\text{Alg}) = \mathbb{E}_{\mathcal{S}_{\mathcal{T}}}[\mathfrak{L}_{\mathcal{T}}(\text{Alg})]$. Then the transfer risk of an algorithm Alg is defined as $\mathcal{L}_{\text{TFR}}(\text{Alg}) = \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Alg})]$. With this setup, we are ready to state our main contributions.

3 Generalization Guarantees for In-context Learning

In this section, we study the i.i.d. data setting with training sequences $\mathcal{S}_t = (\mathbf{x}_{ti}, \mathbf{y}_{ti})_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}_t$ for $t \in [T]$ and establish generalization bounds for in-context learning. Section 6 extends our results to dynamical systems.

Algorithmic Stability

In ICL a training example $(\mathbf{x}_i, \mathbf{y}_i)$ in the prompt impacts all future decisions of the algorithm from predictions $i + 1$ to n . This necessitates us to control the stability to input perturbation of the learning algorithm emulated by the transformer. Our stability condition is borrowed from the algorithmic stability literature. As stated in [5, 6], the stability level of an algorithm is typically in the order of $1/m$ (for realistic generalization guarantees) where m is the training sample size (in our setting prompt length). This is formalized in the following assumption that captures the variability of the transformer output.

Assumption 1 (Error Stability [6]) *Let $\mathcal{S} = (\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_m, \mathbf{y}_m)$ be a sequence with $m \geq 1$ and \mathcal{S}^j be the sequence where the j 'th sample of \mathcal{S} is replaced by $(\mathbf{x}'_j, \mathbf{y}'_j)$. Error stability holds for a distribution $(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}$ if there exists a constant $K > 0$ such that for any $\mathcal{S}, (\mathbf{x}'_j, \mathbf{y}'_j) \in (\mathcal{X} \times \mathcal{Y})$, $j \in [m]$, and $\text{Alg} \in \mathcal{A}$ we have*

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y})} \left| \ell(\mathbf{y}, f_{\mathcal{S}}^{\text{Alg}}(\mathbf{x})) - \mathbb{E}_{(\mathbf{x}, \mathbf{y})} \left| \ell(\mathbf{y}, f_{\mathcal{S}^j}^{\text{Alg}}(\mathbf{x})) \right| \leq \frac{K}{m}. \quad (2)$$

Let ρ be a distance metric on \mathcal{A} . Pairwise error stability holds if for all $\text{Alg}, \text{Alg}' \in \mathcal{A}$ we have

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y})} \left| \ell(\mathbf{y}, f_{\mathcal{S}}^{\text{Alg}}(\mathbf{x})) - \ell(\mathbf{y}, f_{\mathcal{S}}^{\text{Alg}'}(\mathbf{x})) \right| - \mathbb{E}_{(\mathbf{x}, \mathbf{y})} \left| \ell(\mathbf{y}, f_{\mathcal{S}^j}^{\text{Alg}}(\mathbf{x})) - \ell(\mathbf{y}, f_{\mathcal{S}^j}^{\text{Alg}'}(\mathbf{x})) \right| \leq \frac{K\rho(\text{Alg}, \text{Alg}')}{m}. \quad (3)$$

Here (2) is our primary stability condition borrowed from [6] and ensures that all algorithms $\text{Alg} \in \mathcal{A}$ are K -stable. We will also use the stronger pairwise stability condition to develop tighter generalization bounds. The following theorem shows that, under mild assumptions, a multilayer transformer obeys the stability condition (2). The proof is deferred to Appendix A.1 and Theorem 4.

Theorem 1 *Let $\mathbf{x}_{\text{prompt}}^{(m)}, \mathbf{x}'_{\text{prompt}}{}^{(m)}$ be two prompts that differ only at the inputs $\mathbf{z}_j = (\mathbf{x}_j, \mathbf{y}_j)$ with $\mathbf{z}'_j = (\mathbf{x}'_j, \mathbf{y}'_j)$. Assume inputs and labels lie within the unit Euclidean ball in \mathbb{R}^d ¹. Shape these prompts into matrices $\mathbf{X}_{\text{prompt}}, \mathbf{X}'_{\text{prompt}} \in \mathbb{R}^{(2m-1) \times d}$ respectively. Let $\text{TF}(\cdot)$ be a transformer composed of D layers followed by a linear head: $\text{TF}(\mathbf{x}_{\text{prompt}}^{(m)}) = \mathbf{H}, \mathbf{X}_{(D)}$. Here $\mathbf{X}_{(i)}$ denotes the i 'th layer output with $\mathbf{X}_{(0)} := \mathbf{X}_{\text{prompt}}$ and*

$$\mathbf{X}_{(i)} = \text{Parallel_MLPs}(\text{ATTN}(\mathbf{X}_{(i-1)})) \quad \text{where} \quad \text{ATTN}(\mathbf{X}) := \text{softmax}(\mathbf{X}^\top \mathbf{W}_i \mathbf{X}) \mathbf{X} \mathbf{V}_i,$$

and softmax is row-wise. Assume TF is normalized to output $[-1, 1]$ as follows: $\|\mathbf{V}\| \leq 1, \|\mathbf{W}\| \leq 1/2$, MLPs obey $\text{MLP}(\mathbf{x}) = \text{ReLU}(\mathbf{M}\mathbf{x})$ with $\|\mathbf{M}\| \leq 1$, and each row of \mathbf{H} has Euclidean norm at most $1/m$. Then,

$$|\text{TF}(\mathbf{x}_{\text{prompt}}^{(m)}) - \text{TF}(\mathbf{x}'_{\text{prompt}}{}^{(m)})| \leq \frac{2}{2m-1} ((1 + e)^D).$$

Thus, assuming loss $\ell(\mathbf{y}, \cdot)$ is L -Lipschitz, the algorithm induced by $\text{TF}(\cdot)$ obeys (2) with $K = 2L((1 + e)^D)$.

A few remarks are in place. First, the dependence on depth is exponential. However this is not as prohibitive for typical transformer architectures which tend to not be very deep. For example the different variants of GPT-2 and BERT have between 12-48 [1] layers). Secondly, if the altered token has an excessive influence on the prediction, this would disrupt the stability condition. Our theorem avoids this via row-wise norm control of the output layer \mathbf{H} which ensures that input tokens contribute equally. We anticipate the analysis can be extended to more realistic transformer architectures (e.g. GPT-2 architecture used in our experiments) by incorporating multiple attention-heads and causal masks. Finally, Figure 7 provides numerical evidence for multiple ICL settings that demonstrate that stability of GPT-2 architecture's predictions with respect to inputs indeed improves with longer prompt length m in line with our theory.

¹Here, we assume $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^d$, otherwise, inputs and labels are both embedded into d -dimensional vectors of proper size.

Generalization Bounds

We are ready to establish generalization bounds by leveraging our stability conditions. We use covering numbers (i.e. metric entropy) to control model complexity.

Definition 1 (Covering number) Let \mathcal{Q} be any hypothesis set and $d(q, q') \geq 0$ be a distance metric over $q, q' \in \mathcal{Q}$. Then, $\bar{\mathcal{Q}} = \{q_1, \dots, q_N\}$ is an ε -cover of \mathcal{Q} with respect to $d(\cdot, \cdot)$ if for any $q \in \mathcal{Q}$, there exists $q_i \in \bar{\mathcal{Q}}$ such that $d(q, q_i) \leq \varepsilon$. The ε -covering number $\mathcal{N}(\mathcal{Q}, d, \varepsilon)$ is the cardinality of the minimal ε -cover.

To cover the algorithm space \mathcal{A} , we need to introduce a distance metric. We formalize this in terms of the prediction difference between two algorithms on the worst-case data-sequence.

Definition 2 (Algorithm distance) Let \mathcal{A} be an algorithm hypothesis set and $\mathcal{S} = (\mathbf{x}_i, \mathbf{y}_i)_{i=1}^n$ be a sequence with n examples that is admissible for some task $t \in [T]$. For any pair $\text{Alg}, \text{Alg}' \in \mathcal{A}$, define the distance metric $\rho(\text{Alg}, \text{Alg}') := \sup_{\mathcal{S}} \frac{1}{n} \sum_{i=1}^n \|f_{\mathcal{S}^{(i-1)}}^{\text{Alg}}(\mathbf{x}_i) - f_{\mathcal{S}^{(i-1)}}^{\text{Alg}'}(\mathbf{x}_i)\|_{\ell_2}$.

We note that the distance ρ is controlled by the Lipschitz constant of the transformer architecture (i.e. largest gradient norm with respect to the model weights). Following Definitions 1&2, the ε -covering number of the hypothesis set \mathcal{A} is $\mathcal{N}(\mathcal{A}, \rho, \varepsilon)$. This brings us to our main result on MTL risk of (ICL-ERM).

Theorem 2 Suppose \mathcal{A} is K -stable per Assumption 1 and the loss function $\ell(\mathbf{y}, \cdot)$ is L -Lipschitz for all $\mathbf{y} \in \mathcal{Y}$ and takes values over $[0, 1]$. Let $\mathcal{A}g$ be the empirical solution of (ICL-ERM). Then, with probability at least $1 - 2\delta$, the excess MTL test risk obeys

$$R_{\text{MTL}}(\mathcal{A}g) \leq \inf_{\varepsilon > 0} 4L\varepsilon + 2(1 + K \log n) \frac{\log \mathcal{N}(\mathcal{A}, \rho, \varepsilon) + \log(1/\delta)}{cnT}. \quad (4)$$

Additionally suppose \mathcal{A} is K -pairwise-stable and set diameter $D := \sup_{\text{Alg}, \text{Alg}' \in \mathcal{A}} \rho(\text{Alg}, \text{Alg}')$. Using the convention $x_+ = \max(x, 1)$, with probability at least $1 - 4\delta$,

$$R_{\text{MTL}}(\mathcal{A}g) \leq \inf_{\varepsilon > 0} 8L\varepsilon + \frac{L_+ + K \log n}{\sqrt{cnT}} \int_{\varepsilon}^{D/2} \rho \frac{D/2}{\log \mathcal{N}(\mathcal{A}, \rho, u)} du + D_+ \frac{\rho}{\log(\log(D/\varepsilon)/\delta)}. \quad (5)$$

The first bound (4) achieves $1/\sqrt{nT}$ rate by covering the algorithm space with resolution ε . For Lipschitz architectures with $\dim(\mathcal{A})$ trainable weights we have $\log \mathcal{N}(\mathcal{A}, \rho, \varepsilon) \sim \dim(\mathcal{A}) \log(1/\varepsilon)$. Thus, up to logarithmic factors, the excess risk is bounded by $\frac{\dim(\mathcal{A})}{nT}$ and will vanish as $n, T \rightarrow \infty$. Recalling (1), this can also be seen as a generalization bound for general auto-regressive models that learns from sequential data.

Note that ρ in Def. 2 is task-dependent. For instance, suppose tasks are realizable with labels $\mathbf{y} = f(\mathbf{x})$ and admissible task sequences have the form $\mathcal{S} = (\mathbf{x}_i, f(\mathbf{x}_i))_{i=1}^n$. Then, ρ will depend on the function class of f (e.g. whether tasks f are linear models, neural nets, etc), specifically, as the function class becomes richer, both ρ and the covering number becomes larger. Under the stronger pairwise-stability, we can obtain a bound in terms of the Dudley entropy integral which arises from a chaining argument. This bound is typically in the same order as the Rademacher complexity of the function class with $T \times n$ samples [55]. Note that achieving $1/\sqrt{T}$ dependence is rather straightforward as tasks are sampled independently. Thus, the main feature of Theorem 2 is obtaining the multiplicative $1/\sqrt{n}$ term by overcoming temporal dependencies. Figure 3 shows that training with full sequence is indeed critical for ICL accuracy.

Proof sketch. Let Alg^* be the population MTL minima. We first decompose $R_{\text{MTL}}(\mathcal{A}g)$ as follows:

$$R_{\text{MTL}}(\mathcal{A}g) = \underbrace{\mathcal{L}_{\text{MTL}}(\mathcal{A}g) - \mathcal{P}_{\mathcal{S}_{\text{all}}}(\mathcal{A}g)}_a + \underbrace{\mathcal{P}_{\mathcal{S}_{\text{all}}}(\mathcal{A}g) - \mathcal{P}_{\mathcal{S}_{\text{all}}}(\text{Alg}^*)}_b + \underbrace{\mathcal{P}_{\mathcal{S}_{\text{all}}}(\text{Alg}^*) - \mathcal{L}_{\text{MTL}}(\text{Alg}^*)}_c$$

where $b \leq 0$ by definition. Regarding (a) and (c), we need a concentration bound on $|\mathcal{L}_{\text{MTL}}(\text{Alg}) - \mathcal{P}_{\mathcal{S}_{\text{all}}}(\text{Alg})|$ for a fixed algorithm $\text{Alg} \in \mathcal{A}$. To achieve this bound, we introduce the sequence of variables $X_{t,i} = \mathbb{E} \frac{1}{n} \sum_{j=1}^n \ell(\mathbf{y}_{tj}, f_{\mathcal{S}_t^{(j-1)}}^{\text{Alg}}(\mathbf{x}_{tj})) - \mathcal{S}_t^{(i)}$ for $0 \leq i \leq n$, which form a martingale by construction.

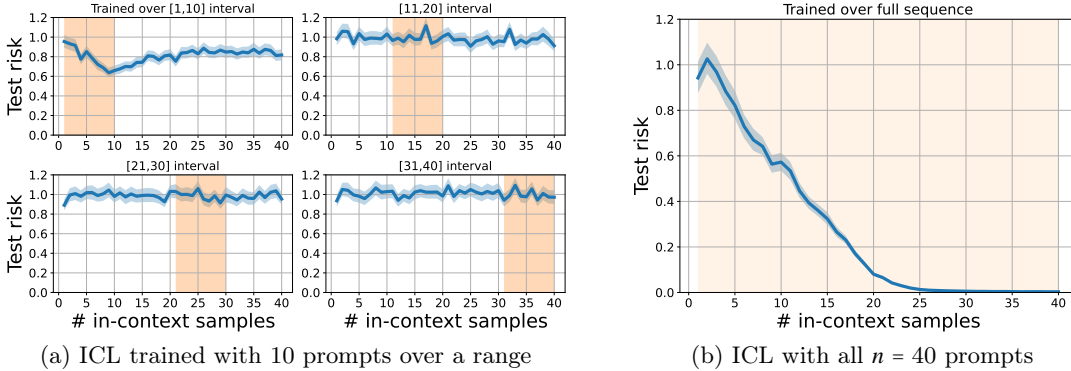


Figure 3: The value of learning across the full sequence: **Right side:** Standard **ICL-ERM** where each task trains with all $n = 40$ prompts (full task sequence). **Left side:** **ICL-ERM** focuses on different parts of the trajectory by fitting $n/4 = 10$ prompts per task over $i \in [1, 10]$ to $[31, 40]$ (highlighted as the orange ranges). We train with $T = 500000$ random linear regression tasks and display the performance on new tasks from the same prior (i.e. transfer risk). Right side learns to solve linear regression via ICL whereas left side fails to do so even when restricted to their target ranges.

The critical stage is upper bounding the martingale differences $|X_{t,i} - X_{t,i-1}|$ through our stability assumption, namely, increments are at most $1 + \prod_{j=1}^n \frac{K}{j} \cdot 1 + K \log n$. Then, we utilize Azuma-Hoeffding’s inequality to achieve a concentration bound on $\frac{1}{T} \prod_{t=1}^T X_{t,0} - X_{t,n}$, which is equivalent to $|\mathcal{L}_{\text{MTL}}(\text{Alg}) - \mathcal{L}_{\text{Sall}}(\text{Alg})|$. To conclude, we make use of covering/chaining arguments to establish a uniform concentration bound on $|\mathcal{L}_{\text{MTL}}(\text{Alg}) - \mathcal{L}_{\text{Sall}}(\text{Alg})|$ for all $\text{Alg} \in \mathcal{A}$. ■

Multiple sequences per task. Finally consider a setting where each task is associated with M independent sequences with size n . This typically arises in reinforcement learning problems (e.g. dynamical systems in Sec. 6) where we collect data through multiple rollouts each leading to independent sequences. In this setting, statistical error rate improves to $1/\sqrt{nMT}$ as discussed in Appendix B.1. In the next section, we will contrast MTL vs transfer learning by letting $M \rightarrow \infty$. This way, even if n and T are fixed, the model will fully learn the T source tasks during the MTL phase as the excess risk vanishes with $M \rightarrow \infty$.

4 Generalization and Inductive Bias on Unseen Tasks

In this section, we explore transfer learning to assess the performance of in-context learning across new tasks: the MTL phase generates a pretrained model Alg trained on T source tasks and we use Alg to predict a target task \mathcal{T} . To start with, we consider a standard meta-learning setup where T sources are drawn from the task distribution $\mathcal{D}_{\text{task}}$ and we evaluate the transfer risk on new tasks $\mathcal{T} \sim \mathcal{D}_{\text{task}}$. We are interested in controlling the transfer risk $\mathcal{L}_{\text{TFR}}(\text{Alg}) = \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Alg})]$ in terms of the MTL risk $\mathcal{L}_{\text{MTL}}(\text{Alg})$. Since the source tasks are i.i.d, we have $\mathbb{E}_{\text{source}}[\mathcal{L}_{\text{MTL}}(\text{Alg})] = \mathcal{L}_{\text{TFR}}(\text{Alg})$. More specifically, one can use a standard generalization analysis to control the transfer risk in terms of MTL risk as follows

$$\mathcal{L}_{\text{TFR}}(\text{Alg}) - \mathcal{L}_{\text{MTL}}(\text{Alg}) \leq \frac{\sqrt{\dim(\mathcal{A})}}{T}.$$

Here, $\dim(\mathcal{A})$ captures the complexity of the algorithm space (e.g. as formalized in Theorem 7 in Appendix B.3, one can use covering numbers similar to Theorem 2). Here, an important distinction with MTL is that transfer risk decays as $1/\text{poly}(T)$ because the unseen tasks induce a distribution shift which, typically, cannot be mitigated with more samples n or more sequences M .

• **Inductive Bias in Transfer Risk.** Before investigating distribution shift, let us consider the following question: While $1/\text{poly}(T)$ behavior may be unavoidable, is it possible that dependence on architectural

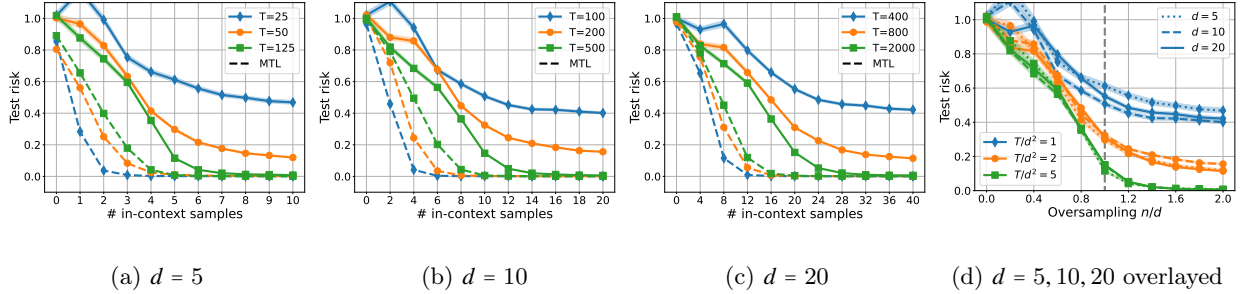


Figure 4: In Figures (a,b,c), we plot the $d \in \{5, 10, 20\}$ -dimensional results for transfer and MTL risk curves. Figure (d) overlays (a,b,c) to reveal that transfer risks are aligned for fixed $(n/d, T/d^2)$ choice.

complexity $\dim(\mathcal{A})$ is avoidable? Perhaps surprisingly, we answer this question affirmatively through experiments on linear regression. In what follows, during MTL pretraining, we train with $M \rightarrow \infty$ independent sequences per task to minimize population MTL risk $\mathcal{L}_{\text{MTL}}(\cdot)$. We then evaluate resulting Alg on different dimensions d and numbers of MTL tasks T . Figures 4(a,b,c) display the MTL and transfer risks for dimensions $d = 5, 10, 20$. In each figure, we evaluate the results on $T = \{1, 2, 5\} \times d^2$ and the x -axis moves from 0 to $n = 2d$. Each task has isotropic features, noiseless labels and task vectors $\beta \sim \mathcal{N}(0, \mathbf{I}_d)$. Here, our first observation is that, the Figures 4(a,b,c) seem (almost perfectly) aligned with each other, that is, each figure exhibits identical MTL and transfer risk curves.

To further elucidate this, Figure 4(d) integrates the transfer risk curves from $d = 5, 10, 20$ and overlays them together. This alignment indicates that, for a fixed point $\alpha = n/d$ and $\beta = T/d^2$, the transfer risks remain unchanged. Here, n proportional to d can be attributed to linearity, thus, the more surprising aspect is the dependence on number of tasks T since this indicates that, rather than $\frac{\dim(\mathcal{A})}{T}$, the generalization risk behaves like $\frac{d^2}{T}$. In words, rather than model complexity, what matters is the problem complexity d .

Inductive bias is a natural explanation of this behavior: Intuitively, the MTL pretraining process identifies a favorable algorithm that lies in the span of the source tasks $\text{MTL} = (\beta_i)_{i=1}^T$. Specifically, while the transformer model can potentially fit MTL tasks through a variety of algorithms, we speculate that the optimization process is implicitly biased (akin to implicit regularization literature [47, 36, 3]) to an algorithm $\text{Alg}(\text{MTL})$. Such bias would explain the lack of $\dim(\mathcal{A})$ dependence since $\text{Alg}(\text{MTL})$ solely depends on the source tasks. We leave the theoretical exploration of the empirical d^2/T behavior to a future work. However, we provide further discussion on why it is rather surprising both from MTL and transfer learning perspectives.

To this end, let us first introduce the optimal estimator for linear regression with Gaussian task prior. *Minimum Bayes risk:* The optimal prediction in terms of average risk can be described explicitly [41, 28] under prior $\beta \sim \mathcal{N}(0, \Sigma)$ and is given by the optimally-weighted ridge regression solution

$$\hat{\beta} = (\mathbf{X}^\top \mathbf{X} + \sigma^2 \Sigma^{-1})^{-1} \mathbf{X}^\top \mathbf{y}, \quad (6)$$

Here $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$, $\mathbf{y} = [\mathbf{y}_1, \dots, \mathbf{y}_n]^\top \in \mathbb{R}^n$ are the concatenated features and labels obtained from the task sequence and σ^2 is the label noise variance. With this in mind, what is the ideal algorithm $\text{Alg}(\text{MTL})$ based on the (perfect) knowledge of source tasks? Eqn. (6) crucially requires the knowledge of the task covariance Σ and variance σ^2 . Thus, even with the hindsight knowledge that we are dealing with a linear problem, we have to estimate the task covariance from source tasks. This can be done by constructing the empirical covariance $\hat{\Sigma} = \frac{1}{T} \sum_{i=1}^T \beta_i \beta_i^\top$. To make sure that $\hat{\Sigma}$ -weighted LS performs $\mathcal{O}(1)$ -close to Σ -weighted LS, we need a spectral norm control, namely, $\|\hat{\Sigma} - \Sigma\| / \lambda_{\min}(\Sigma) \leq \mathcal{O}(1)$. When $\Sigma = \mathbf{I}_d$ (as in our experiments) and tasks are isotropic, the latter condition holds with high probability when $T = \mathcal{O}(d)$. This is also numerically demonstrated in Figure 8 in the appendix. This behavior is in contrast to the stronger $T \propto d^2$ requirement we observe for ICL. For instance, $T \propto d^2$ is sufficient to ensure the stronger entrywise control $\|\hat{\Sigma} - \Sigma\|_{\ell_\infty} \leq \mathcal{O}(1)$ rather than spectral norm.

Let us now ask the same question for the MTL algorithm: If the transformer perfectly learns the MTL tasks $\text{MTL} = (\beta_t)_{t=1}^T$, it does not actually need $n = d$ samples to perform well on new prompts drawn from source tasks. To see this, consider the following algorithm: Given a prompt, $\text{Alg}(\text{MTL})$ conducts a discrete search over $(\beta_t)_{t=1}^T$ and returns the source task that best fits to the prompt. Thanks to the discrete search space, it is not hard to see that, we need $n \propto \log(T)$ samples rather than $n \propto d$ (also see Figure 8). In contrast, based on Figures 4(a,b,c), MTL behaves closer to $n \propto d$ empirically. On the other hand, $\text{Alg}(\text{MTL})$ implemented by the transformer is rather intelligent: This is because MTL risks for $d \in \{5, 10, 20\}$ are all strictly better than implementing least-squares² and the performance improves as T gets smaller. We leave the thorough exploration of the inductive bias of the MTL training and characterization of $\text{Alg}(\text{MTL})$ as an intriguing future direction.

• **Exploring transfer risk in terms of distance.** Besides drawing source and target tasks from the same $\mathcal{D}_{\text{task}}$, we can also investigate transfer risk in an instance specific fashion. Specifically, the population risk of a new task \mathcal{T} can be bounded as $\mathcal{L}_{\mathcal{T}}(\text{Alg}) \leq \mathcal{L}_{\text{MTL}}(\text{Alg}) + \text{dist}(\mathcal{T}, (\mathcal{D}_t)_{t=1}^T)$. Here, $\text{dist}(\cdot)$ assesses the (distributional) distance of task \mathcal{T} to the source tasks $(\mathcal{D}_t)_{t=1}^T$ (e.g. [4, 21]). In case of linear tasks, we can simply use the Euclidean distance between task vectors, specifically, the distance of target weights $\beta_{\mathcal{T}}$ to the nearest source task $\text{dist}(\mathcal{T}) = \min_{t \in [T]} \|\beta_{\mathcal{T}} - \beta_t\|_{\ell_2}$. In Fig. 5 we investigate the distance of specific target tasks from source tasks and how the distance affects the transfer performance. Here, all source and target tasks have unit Euclidean norms so that closer distance is equivalent to larger cosine similarity. We again train each MTL task with multiple sequences $M \rightarrow \infty$ (as in Fig. 4) and use $T = 20$ source tasks with $d = 5$ dimensional regression problems. In a nutshell, Figure 5 shows that task similarity (based on ℓ_2 distance) is indeed highly predictive of transfer performance across different distance slices (namely $[0, 0.2]$, $[0.2, 0.6]$, $[0.6, 1]$, $[1, 2]$). An intriguing question is identifying the broader conditions under which transfer risk can be rigorously related to the MTL risk, which is further discussed in Appendix B.2.




Figure 5: Transfer risk as a function of distance to the MTL tasks.

5 Insights on Model Selection

In Section 3, we study the generalization error of ICL, which can be eliminated by increasing sample size n or number of sequences M per task. In this section, we will discuss implicit model selection in ICL building on the formalism that transformer is a learning algorithm. Following Figure 2 and prior works [20, 26, 22], a plausible assumption is that, transformer can implement ERM algorithms up to a certain accuracy. Then, model selection can be formalized by the selection of the right hypothesis class so that running ERM on that hypothesis class can strike a good bias-variance tradeoff during ICL.

To proceed with our discussion, let us consider the following hypothesis which states that transformer can implement an algorithm competitive with ERM.

Hypothesis 1 Let $\mathcal{F} = (\mathcal{F}_h)_{h=1}^H$ be a family of H hypothesis classes. Let $\mathcal{S} = (\mathbf{x}_i, \mathbf{y}_i)_{i=1}^n$ be a data-sequence with n examples sampled i.i.d. from \mathcal{D} and let $\mathcal{S}^{(m)} = (\mathbf{x}_i, \mathbf{y}_i)_{i=1}^m$ be the first m examples. Consider the risk³ associated to ERM with m samples over $\mathcal{F}_h \in \mathcal{F}$:

$$\text{risk}(h, m) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathcal{S}^{(m)})} [\ell(\mathbf{y}, \hat{f}_{\mathcal{S}^{(m)}}^{(h)}(\mathbf{x}))] \quad \text{where} \quad \hat{f}_{\mathcal{S}^{(m)}}^{(h)} = \arg \min_{f \in \mathcal{F}_h} \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{y}_i, f(\mathbf{x}_i)),$$

Let $(\varepsilon_{TF}^{h,m}) > 0$ be approximation errors associated with $(\mathcal{F}_h)_{h=1}^H$. There exists $\text{Alg} \in \mathcal{A}$ such that, for any

²Ordinary least-squares achieves the minimum risk for transfer learning ($T = \infty$) however it is not optimal for finite T .

³Since the loss ℓ is bounded by 1, $0 \leq \text{risk}(h, m) \leq 1$ for all m including the scenario $m = 0$ and ERM is vacuous.

$m \in [n], h \in [H]$, $f_{\mathcal{S}^{(m)}}^{\text{Alg}}$ can approximate ERM in terms of population risk, i.e.

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathcal{S}^{(m)})}[\ell(\mathbf{y}, f_{\mathcal{S}^{(m)}}^{\text{Alg}}(\mathbf{x}))] \leq \text{risk}(h, m) + \varepsilon_{\text{TF}}^{h, m}.$$

For model selection purposes, these hypothesis classes can be entirely different ML models, for instance, $\mathcal{F}_1 = \{\text{convolutional-nets}\}$, $\mathcal{F}_2 = \{\text{fully-connected-nets}\}$, and $\mathcal{F}_3 = \{\text{decision-trees}\}$. Alternatively, they can be a nested family useful for capacity control purposes. For instance, Figures 2(a,b) are learning covariance/noise priors to implement a constrained-ridge regression. Here \mathcal{F} can be indexed by positive-definite matrices with linear classes of the form $\mathcal{F} = \{f(\mathbf{x}) = \mathbf{x}^\top \boldsymbol{\beta} \mid \boldsymbol{\beta}^\top \mathbf{F}^{-1} \boldsymbol{\beta} \leq 1\}$.

Under Hypothesis 1, ICL selects the most suitable class that minimizes the excess risk for each $m \in [n]$.

Observation 1 Suppose Hypothesis 1 holds for a target distribution $\mathcal{D}_{\mathcal{T}}$. Let $\mathcal{L}_{\mathcal{T}}^{\star} := \min_{\text{Alg} \in \mathcal{A}} \mathcal{L}_{\mathcal{T}}(\text{Alg})$ be the risk of the optimal algorithm. We have that

$$\mathcal{L}_{\mathcal{T}}^{\star} \leq \frac{1}{n} \tilde{\mathcal{O}}^1 \min_{m=0} \min_{h \in [H]} \{\text{risk}(h, m) + \varepsilon_{\text{TF}}^{h, m}\}.$$

Additionally, denote Rademacher complexity of a class \mathcal{F} by $\mathcal{R}_m(\mathcal{F})$. Define the minimum achievable risk over function set \mathcal{F}_h as $\mathcal{L}_h^{\star} := \min_{f \in \mathcal{F}_h} \mathbb{E}_{\mathcal{D}_{\mathcal{T}}}[\ell(\mathbf{y}, f(\mathbf{x}))]$. Since $\text{risk}(h, m)$ is controlled by $\mathcal{R}_m(\mathcal{F}_h)$ [35], we have that

$$\mathcal{L}_{\mathcal{T}}^{\star} \leq \frac{1}{n} \tilde{\mathcal{O}}^1 \min_{m=0} \min_{h \in [H]} \{\mathcal{L}_h^{\star} + \varepsilon_{\text{TF}}^{h, m} + \mathcal{O}(\mathcal{R}_m(\mathcal{F}_h))\}.$$

Here, ICL adaptively selects the classes $\arg \min_{h \in [H]} \{\mathcal{L}_h^{\star} + \mathcal{R}_m(\mathcal{F}_h) + \varepsilon_{\text{TF}}^{h, m}\}$ to achieve small risk. This is in contrast to training over a single large class $\mathcal{F} = \bigcup_{i=1}^H \mathcal{F}_i$, which would result in a less favorable bound $\approx \min_{h \in [H]} \mathcal{L}_h^{\star} + \frac{1}{n} \sum_{m=0}^{n-1} \mathcal{R}_m(\mathcal{F})$. A formal version of this statement is provided in Appendix B.4. Hypothesis 1 assumes a discrete family for simpler exposition ($|\mathcal{F}| = H < \infty$), however, our theory in Section 3 allows for the continuous setting.

We emphasize that, in practice, we need to adapt the hypothesis classes for different sample sizes m (typically, more complex classes for larger m). With this in mind, while we have H classes in \mathcal{F} , in total we have H^n different ERM algorithms to compete against. This means that VC-dimension of the algorithm class is as large as $n \log H$. This highlights an insightful benefit of our main result: Theorem 2 would result in an excess risk $\propto \frac{n \log H}{nT} = \frac{\log H}{T}$. In other words, the additional $\times n$ factor achieved through Theorem 2 facilitates the adaptive selection of hypothesis classes for each sample size and avoids requiring unreasonably large T .

6 Extending Results to Stable Dynamical Systems

Until now, we have studied ICL with sequences of i.i.d. (input, label) pairs. In this section, we investigate the scenario where prompts are obtained from the trajectories of stable dynamical systems, thus, they consist of dependent data. Let $\mathcal{X} \subset \mathbb{R}^d$ and $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{X}$ be a hypothesis class elements of which are dynamical systems. During MTL phase, suppose that we are given T tasks associated with $(f_t)_{t=1}^T$ where $f_t \in \mathcal{F}$, and each contains n in-context samples. Then, the data-sequence of t 'th task is denoted by $\mathcal{S}_t = (\mathbf{x}_{ti})_{i=0}^n$ where $\mathbf{x}_{ti} = f_t(\mathbf{x}_{t, i-1}) + \mathbf{w}_{ti}$, \mathbf{x}_{t0} is the initial state, and $\mathbf{w}_{ti} \in \mathcal{W} \subset \mathbb{R}^d$ are bounded i.i.d. random noise following some distribution. Then, prompts are given by $\mathbf{x}_{\text{prompt}}^{(i)} := (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_i)$. Let $\mathcal{S}^{(i)} = \mathbf{x}_{\text{prompt}}^{(i)}$, and we can make prediction $\hat{\mathbf{x}}_{ti} = f_{\mathcal{S}_t^{(i-1)}}^{\text{Alg}}(\mathbf{x}_{t, i-1})$. We consider the similar optimization problem as (ICL-ERM).

In order to perform the analysis of generalization error in a dynamical system, we need this system to be stable (which differs from algorithmic stability!). In this work, we use an exponential stability condition that controls the distance between two trajectories initialized from different points [45, 24].

Definition 3 ((C_ρ, ρ) -stability) Denote the state of dynamical system resulting from initial state \mathbf{x}_{t_0} and $(\mathbf{w}_{ti})_{i=1}^m$ by $f_t^{(m)}(\mathbf{x}_{t_0}) \in \mathcal{X}$. Let $C_\rho \geq 1$ and $\rho \in (0, 1)$ be system related constants. We say that the dynamical system for the task t is (C_ρ, ρ) -stable if, for all $\mathbf{x}_{t_0}, \mathbf{x}'_{t_0} \in \mathcal{X}$, $m \geq 1$, and $(\mathbf{w}_{ti})_{i \geq 1} \in \mathcal{W}$, we have

$$f_t^{(m)}(\mathbf{x}_{t_0}) - f_t^{(m)}(\mathbf{x}'_{t_0}) \leq C_\rho \rho^m \|\mathbf{x}_{t_0} - \mathbf{x}'_{t_0}\|_{\ell_2} \quad (7)$$

Assumption 2 There exist \bar{C}_ρ and $\bar{\rho} < 1$ such that all dynamical systems $f \in \mathcal{F}$ are $(\bar{C}_\rho, \bar{\rho})$ -stable.

In addition to the stability of the hypothesis set \mathcal{F} , we also leverage the algorithmic-stability of the set \mathcal{A} similar to Assumption 1. Different from Assumption 1, we restrict the variability of the transformer with respect to Euclidean distance metric, similar to the definition of Lipschitz stability.

Assumption 3 (Algorithmic-stability for Dynamical Systems) Let $\mathcal{S} = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_m)$ be any dynamical sequence with $m \geq 1$ and \mathcal{S}^j be the sequence such that \mathbf{w}_j is replaced by \mathbf{w}'_j ($j = 0$ implies that \mathbf{x}_0 is replaced by \mathbf{x}'_0). As a result, starting with the j 'th until the m 'th index, the sequence \mathcal{S}^j has different samples $(\mathbf{x}'_j, \dots, \mathbf{x}'_m)$. There exists a constant $K > 0$ such that for any \mathcal{S} , $\mathbf{x}'_0 \in \mathcal{X}$, $\mathbf{w}'_j \in \mathcal{W}$, and $j \in [m]$, we have

$$\mathbb{E}_{\mathbf{w}_{m+1}} \mathbb{E}_{\mathbf{x}_{m+1}} \ell(\mathbf{x}_{m+1}, f_{\mathcal{S}}^{Alg}(\mathbf{x}_m)) - \mathbb{E}_{\mathbf{w}_{m+1}} \mathbb{E}_{\mathbf{x}_{m+1}} \ell(\mathbf{x}_{m+1}, f_{\mathcal{S}^j}^{Alg}(\mathbf{x}_{m'})) \leq \frac{K}{m} \sum_{i=j}^m \|\mathbf{x}_i - \mathbf{x}'_i\|_{\ell_2}. \quad (8)$$

Theorem 4 in the appendix justifies this assumption for multilayer transformer architectures. To proceed, we state the main result of this section.

Theorem 3 Suppose the loss function $\ell(\mathbf{x}, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ is L -Lipschitz for all $\mathbf{x} \in \mathcal{X}$ and Assumptions 2&3 hold. Assume \mathcal{X} and \mathcal{W} are bounded by \bar{x}, \bar{w} , respectively. Then, with the same probability, the identical bound as in Theorem 2 Eqn. (4) holds after updating K to be $\bar{K} = 2K \frac{C_\rho}{1-\rho} (\bar{w} + \bar{x}/\sqrt{n})$.

The proof of this result is similar to that of Theorem 2 and can be found in Appendix C. The primary difference occurs when bounding the martingale differences $|X_{t,i} - X_{t,i-1}|$. In the stable dynamical settings, to bound Martingale's increments, we utilize independence of noise and stability of the system instead of using the independence of input and output pairs. One can also obtain an entropy integral bound akin to (5) using a proper variation of the pairwise algorithmic-stability condition in Assumption 1 Eqn. (3).

7 Numerical Evaluations

Our experimental setup follows [20]: All ICL experiments are trained and evaluated using the same GPT-2 architecture with 12 layers, 8 attention heads, and 256 dimensional embeddings. Most of our experiments are based on linear regression problems and linear dynamics. We first provide the details and discussions of Figure 2 and then provide additional experiments aligned with our findings.

- **Linear regression (Figures 2(a,b)).** We consider a d -dimensional linear regression tasks with in-context examples of the form $\mathbf{z} = (\mathbf{x}, y) \in \mathbb{R}^d \times \mathbb{R}$. Given t 'th task β_t , we generate n i.i.d. samples via $y = \beta_t^\top \mathbf{x} + \xi$, where $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I})$, $\xi \sim \mathcal{N}(0, \sigma^2)$ and σ is the noise level. Tasks are sampled i.i.d. via $\beta_t \sim \mathcal{N}(0, \Sigma)$, $t \in [T]$. Results are displayed in Figures 2(a)&(b). We set $d = 20$, $n = 40$ and significantly larger T to make sure model is sufficiently trained and we display meta learning results (i.e. on unseen tasks) for both experiments. In Fig. 2(a), $\sigma = 1$ and $\Sigma = \mathbf{I}$. We also solve ridge-regularized linear regression (with sample size from 1 to n) over the grid $\lambda = [0.01, 0.05, 0.1, 0.5, 1]$ and display the results of the best λ selection as the optimal ridge curve (Black dotted). Recall from (6) that ridge regression is optimal for isotropic task covariance. In Fig. 2(b), we set $\sigma = 0$ and $\Sigma = \text{diag}(1, \frac{1}{2^2}, \frac{1}{3^2}, \dots, \frac{1}{20^2})$. Besides ordinary least squares (Green curve), we also display the optimally-weighted regression according to (6) (dotted curve) as $\sigma \rightarrow 0$. In both figures, in-context learning results (Red) outperform the least squares results (Green) and are perfectly aligned with




Figure 7: Experiments to assess the algorithmic stability Assumption 1. Each figure shows the increase in the risk for varying ICL sample sizes after an example in the prompt is modified. We swap an input example in the prompt and assign a flipped label to this new input, e.g., we move from $(\mathbf{x}, f(\mathbf{x}))$ to $(\mathbf{x}', -f(\mathbf{x}'))$.

optimal ridge/weighted solutions (Black dotted). This in turn provides evidence for the automated model selection ability of transformers by learning task priors.

- Partially-observed dynamical systems (Figures 2(c) & 6).** We generate in-context examples $\mathbf{z}_i = \mathbf{x}_i \in \mathbb{R}^r$, $i \in [n]$ via the partially-observed linear dynamics $\mathbf{x}_i = \mathbf{C}\mathbf{s}_i$, $\mathbf{s}_i = \mathbf{A}\mathbf{s}_{i-1} + \boldsymbol{\xi}_i$ with noise $\boldsymbol{\xi}_i \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_d)$ and initial state $\mathbf{s}_0 = \mathbf{0}$. Each task is parameterized by $\mathbf{C} \in \mathbb{R}^{r \times d}$ and $\mathbf{A} \in \mathbb{R}^{d \times d}$ which are drawn with i.i.d. $\mathcal{N}(0, 1)$ entries and \mathbf{A} is normalized to have spectral radius 0.9. In Fig. 2(c), we set $d = 10$, $r = 4$, $\sigma = 0$, $n = 20$ and use sufficiently large T to train the transformer. For comparison, we solve least-squares regression to predict new observations \mathbf{x}_i via the most recent H observations for varying window sizes H . Results show that in-context learning outperforms the least-squares results of all orders $H = 1, 2, 3, 4$. In Figure 6, we also solve the dynamical problem using optimal ridge regression for different window sizes. This reveals that ICL can also outperform auto-regressive models with optimal ridge tuning, albeit the performance gap is much narrower⁴. It would be interesting to compare ICL performance to a broader class of system identification algorithms (e.g. Hankel nuclear norm, kernel-based, atomic norm [29, 17, 39]) and more broadly understand the extent ICL can inform practically-relevant system identification.

- Stability analysis.** In Assumption 1, we require that transformer-induced algorithms are stable to input perturbations, specifically, we require predictions to vary by at most $\mathcal{O}(1/m)$ where m is the sample size. This was justified in part by Theorem 1. To understand empirical stability, we run additional experiments where the results are displayed in Fig. 7. We study stability of four function classes: linear models, 3-sparse linear models, decision trees with depth 4, and 2-layer ReLU networks with 100 hidden units, all with input dimension of 20. For each class \mathcal{F} , a GPT-2 architecture is trained with large number of random tasks $f \in \mathcal{F}$ and evaluate on new tasks. With the exception of Fig. 2(a), we use the pretrained models provided by [20] and the task sequences are noiseless i.e. sequences obey $y_i = f(\mathbf{x}_i)$. As a coarse approximation of the *worst-case* perturbation, we perturb a prompt $\mathbf{x}_{\text{prompt}}^{(m)} = (\mathbf{x}_1, y_1, \dots, \mathbf{x}_{m-1}, y_{m-1}, \mathbf{x}_m)$ as follows. Draw a random point $(\mathbf{x}'_1, y'_1) \sim (\mathbf{x}_1, y_1)$ and flip its label to obtain $(\mathbf{x}'_1, -y'_1)$. We obtain the adversarial prompt via




Figure 6: Dynamical system experiments. The difference from Fig. 2(c) is that we compare ICL to the optimally-tuned ridge regression with different history windows H .

⁴We note that Kalman filter can be used to achieve optimal state estimation (which leverages the full history i.e. $H \rightarrow \infty$). However it would require the knowledge of the ground-truth dynamics unlike ICL.

$\bar{\mathbf{x}}_{\text{prompt}}^{(m)} = (\mathbf{x}'_1, -\mathbf{y}'_1, \dots, \mathbf{x}_{m-1}, \mathbf{y}_{m-1}, \mathbf{x}_m)$ ⁵. In Figure 7, we plot the test risk change between the adversarial and standard prompts. All figures corroborate that, after a certain sample size, the risk change noticeably decreases as the in-context sample size increases. This behavior is in line with Assumption 1; however, further investigation and longer context window are required to accurately characterize the stability profile (e.g. to verify whether stability is $\mathcal{O}(1/m)$ or not).

8 Conclusions and Future Directions

In this work, we approach prompt-based in-context learning as an algorithm learning problem from a statistical perspective. Through this view, we presented generalization bounds for MTL where the model is trained with T tasks each mapped to a sequence containing n examples. We analyzed generalization bounds for two fundamental scenarios in which the prompts are obtained from a sequence of i.i.d. (input, output) pairs or a trajectory of a dynamical systems. Additionally, we have discussed generalization guarantees on new tasks with respect to excess MTL risk and presented numerical simulations supporting our hypothesis. Finally, we have discussed how in-context learning can select the correct hypothesis set over different function classes by implicitly solving empirical risk minimization over the data provided within the prompt.

There are multiple interesting questions related to our findings: (1) Our bounds are mainly useful for capturing multitask learning risk, which motivates us to study the question: How can we control generalization on individual tasks or prompts with specific lengths (rather than average risk)? (2) We provided guarantees for dynamical systems with full-state observations. Can we extend such results to more general dynamic settings such as reinforcement learning or system identification with partial state observations? (3) Our investigation of transfer learning in Section 4 revealed that transfer risk is governed by the number of MTL tasks and task complexity however it seems to be independent of the model complexity. It would be interesting to characterize this inductive bias during pretraining and characterize exactly what algorithm is learned by the transformer.

Acknowledgements

This work was supported in part by the NSF grants CCF-2046816 and CCF-2212426, Google Research Scholar award, and Army Research Office grant W911NF2110312.

References

- [1] Huggingface pretrained models, 2023.
- [2] Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. What learning algorithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*, 2022.
- [3] Navid Azizan and Babak Hassibi. Stochastic gradient/mirror descent: Minimax optimality and implicit regularization. In *International Conference on Learning Representations*, 2019.
- [4] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.
- [5] Olivier Bousquet and André Elisseeff. Algorithmic stability and generalization performance. *Advances in Neural Information Processing Systems*, 13, 2000.
- [6] Olivier Bousquet and André Elisseeff. Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526, 2002.

⁵To properly verify Assumption 1 one should adversarially optimize $\mathbf{x}'_1, \mathbf{y}'_1$ and also swap the other indices $m > i > 1$.

- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [8] Lisha Chen, Songtao Lu, and Tianyi Chen. Understanding benign overfitting in gradient-based meta learning. In *Advances in Neural Information Processing Systems*, 2022.
- [9] Yuan Cheng, Songtao Feng, Jing Yang, Hong Zhang, and Yingbin Liang. Provable benefit of multitask representation learning in reinforcement learning. *arXiv preprint arXiv:2206.05900*, 2022.
- [10] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- [11] Liam Collins, Aryan Mokhtari, Sewoong Oh, and Sanjay Shakkottai. Maml and anil provably learn representations. *arXiv preprint arXiv:2202.03483*, 2022.
- [12] Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Zhifang Sui, and Furu Wei. Why can gpt learn in-context? language models secretly perform gradient descent as meta optimizers. *arXiv preprint arXiv:2212.10559*, 2022.
- [13] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4):633–679, 2020.
- [14] Tuan Dinh, Yuchen Zeng, Ruisu Zhang, Ziqian Lin, Shashank Rajput, Michael Gira, Jy-yong Sohn, Dimitris Papailiopoulos, and Kangwook Lee. Lift: Language-interfaced fine-tuning for non-language machine learning tasks. *arXiv preprint arXiv:2206.06565*, 2022.
- [15] Simon S Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*, 2020.
- [16] Mohamad Kazem Shirani Faradonbeh and Aditya Modi. Joint learning-based stabilization of multiple unknown linear systems. *arXiv preprint arXiv:2201.01387*, 2022.
- [17] Maryam Fazel, Ting Kei Pong, Defeng Sun, and Paul Tseng. Hankel matrix rank minimization with applications to system identification and realization. *SIAM Journal on Matrix Analysis and Applications*, 34(3):946–977, 2013.
- [18] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- [19] Dylan Foster, Tuhin Sarkar, and Alexander Rakhlin. Learning nonlinear dynamical systems from a single trajectory. In *Learning for Dynamics and Control*, pages 851–861. PMLR, 2020.
- [20] Shivam Garg, Dimitris Tsipras, Percy Liang, and Gregory Valiant. What can transformers learn in-context? a case study of simple function classes. *Neural Information Processing Systems*, 2022.
- [21] Steve Hanneke and Samory Kpotufe. On the value of target data in transfer learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [22] Noah Hollmann, Samuel Müller, Katharina Eggenberger, and Frank Hutter. TabPFN: A transformer that solves small tabular classification problems in a second. *arXiv preprint arXiv:2207.01848*, 2022.
- [23] Weihao Kong, Raghav Somani, Zhao Song, Sham Kakade, and Sewoong Oh. Meta-learning for mixed linear regression. In *International Conference on Machine Learning*, pages 5394–5404. PMLR, 2020.

- [24] Suhas Kowshik, Dheeraj Nagaraj, Prateek Jain, and Praneeth Netrapalli. Near-optimal offline and streaming algorithms for learning non-linear dynamical systems. *Advances in Neural Information Processing Systems*, 34:8518–8531, 2021.
- [25] Vitaly Kuznetsov and Mehryar Mohri. Time series prediction and online learning. In *Conference on Learning Theory*, pages 1190–1213. PMLR, 2016.
- [26] Michael Laskin, Luyu Wang, Junhyuk Oh, Emilio Parisotto, Stephen Spencer, Richie Steigerwald, DJ Strouse, Steven Hansen, Angelos Filos, Ethan Brooks, et al. In-context reinforcement learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*, 2022.
- [27] Yingcong Li, Mingchen Li, M Salman Asif, and Samet Oymak. Provable and efficient continual representation learning. *arXiv preprint arXiv:2203.02026*, 2022.
- [28] Dennis V Lindley and Adrian FM Smith. Bayes estimates for the linear model. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(1):1–18, 1972.
- [29] Lennart Ljung. System identification. In *Signal analysis and prediction*, pages 163–173. Springer, 1998.
- [30] Horia Mania, Michael I Jordan, and Benjamin Recht. Active learning for nonlinear system identification with guarantees. *arXiv preprint arXiv:2006.10277*, 2020.
- [31] Nikolai Matni and Stephen Tu. A tutorial on concentration bounds for system identification. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 3741–3749. IEEE, 2019.
- [32] Andreas Maurer. A vector-contraction inequality for rademacher complexities. In *International Conference on Algorithmic Learning Theory*, pages 3–17. Springer, 2016.
- [33] Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.
- [34] Aditya Modi, Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Joint learning of linear time-invariant dynamical systems. *arXiv preprint arXiv:2112.10955*, 2021.
- [35] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [36] Behnam Neyshabur, Ryota Tomioka, Ruslan Salakhutdinov, and Nathan Srebro. Geometry of optimization and implicit regularization in deep learning. *arXiv preprint arXiv:1705.03071*, 2017.
- [37] Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*, 2022.
- [38] Samet Oymak and Necmiye Ozay. Revisiting ho–kalman-based system identification: Robustness and finite-sample analysis. *IEEE Transactions on Automatic Control*, 67(4):1914–1928, 2021.
- [39] Gianluigi Pillonetto, Tianshi Chen, Alessandro Chiuso, Giuseppe De Nicolao, and Lennart Ljung. Regularized linear system identification using atomic, nuclear and kernel-based norms: The role of the stability constraint. *Automatica*, 69:137–149, 2016.
- [40] Yuzhen Qin, Tommaso Menara, Samet Oymak, Shinung Ching, and Fabio Pasqualetti. Non-stationary representation learning in sequential linear bandits. *IEEE Open Journal of Control Systems*, 2022.
- [41] Dominic Richards, Jaouad Mourtada, and Lorenzo Rosasco. Asymptotics of ridge (less) regression under general source condition. In *International Conference on Artificial Intelligence and Statistics*, pages 3889–3897. PMLR, 2021.

- [42] Tuhin Sarkar and Alexander Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conference on Machine Learning*, pages 5610–5618. PMLR, 2019.
- [43] Yahya Sattar, Zhe Du, Davoud Ataee Tarzanagh, Laura Balzano, Necmiye Ozay, and Samet Oymak. Identification and adaptive control of markov jump systems: Sample complexity and regret bounds. *arXiv preprint arXiv:2111.07018*, 2021.
- [44] Yahya Sattar and Samet Oymak. Non-asymptotic and accurate learning of nonlinear dynamical systems. *Journal of Machine Learning Research*, 23(140):1–49, 2022.
- [45] Yahya Sattar and Samet Oymak. Non-asymptotic and accurate learning of nonlinear dynamical systems. *Journal of Machine Learning Research*, 23(140):1–49, 2022.
- [46] Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Conference On Learning Theory*, pages 439–473. PMLR, 2018.
- [47] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- [48] Yue Sun, Adhyayan Narang, Ibrahim Gulluk, Samet Oymak, and Maryam Fazel. Towards sample-efficient overparameterized meta-learning. *Advances in Neural Information Processing Systems*, 34:28156–28168, 2021.
- [49] Nilesh Tripuraneni, Chi Jin, and Michael Jordan. Provable meta-learning of linear representations. In *International Conference on Machine Learning*, pages 10434–10443. PMLR, 2021.
- [50] Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. *Advances in Neural Information Processing Systems*, 33:7852–7862, 2020.
- [51] Anastasios Tsiamis and George J Pappas. Finite sample analysis of stochastic system identification. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 3648–3654. IEEE, 2019.
- [52] Anastasios Tsiamis, Ingvar Ziemann, Nikolai Matni, and George J Pappas. Statistical learning theory for control: A finite sample perspective. *arXiv preprint arXiv:2209.05423*, 2022.
- [53] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [54] Johannes von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. *arXiv preprint arXiv:2212.07677*, 2022.
- [55] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- [56] Denny Wu and Ji Xu. On the optimal weighted ℓ_2 regularization in overparameterized linear regression. *Advances in Neural Information Processing Systems*, 33:10112–10123, 2020.
- [57] Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*, 2021.

A Stability of Transformer-based ICL

Lemma 1 *Let $\mathbf{x}, \boldsymbol{\varepsilon} \in \mathbb{R}^n$ be vectors obeying $\|\mathbf{x}\|_{\ell_\infty}, \|\mathbf{x} + \boldsymbol{\varepsilon}\|_{\ell_\infty} \leq c$. Then, there exists a constant $C = C(c)$, such that*

$$\|\text{softmax}(\mathbf{x})\|_{\ell_\infty} \leq e^{2c}/n \quad \text{and} \quad \|\text{softmax}(\mathbf{x}) - \text{softmax}(\mathbf{x} + \boldsymbol{\varepsilon})\|_{\ell_1} \leq e^{2c} \|\boldsymbol{\varepsilon}\|_{\ell_1}/n.$$

Proof. Without losing generality, assume the first coordinate is the largest. Using monotonicity of softmax, we obtain $\|\text{softmax}(\mathbf{x})\|_{\ell_\infty} \leq \frac{e^c}{e^c + \frac{1}{n-2}e^{-c}} \leq \frac{e^{2c}}{n}$. Say $\|\boldsymbol{\varepsilon}\|_{\ell_2} = \varepsilon$. As $\boldsymbol{\varepsilon} \rightarrow 0$, we have

$$\lim_{\boldsymbol{\varepsilon} \rightarrow 0} [\text{softmax}(\mathbf{x} + \boldsymbol{\varepsilon}) - \text{softmax}(\mathbf{x})]/\boldsymbol{\varepsilon} = [\text{diag}(\text{softmax}(\mathbf{x})) - \text{softmax}(\mathbf{x})\text{softmax}(\mathbf{x})^\top]\boldsymbol{\varepsilon}/\varepsilon$$

$\|[\text{diag}(\text{softmax}(\mathbf{x})) - \text{softmax}(\mathbf{x})\text{softmax}(\mathbf{x})^\top]\boldsymbol{\varepsilon}\|_{\ell_1} \leq e^{2c} \|\boldsymbol{\varepsilon}\|_{\ell_1}/n$. Integrating the derivative along $\boldsymbol{\varepsilon}$, we obtain the result. \blacksquare

For a matrix \mathbf{A} , let $\|\mathbf{A}\|_{2,p}$ denote the ℓ_p norm of the vector obtained by the ℓ_2 norms of its rows.

Lemma 2 *Let $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_n]^\top$ and $\mathbf{E} = [\boldsymbol{\varepsilon}_1 \dots \boldsymbol{\varepsilon}_n]^\top$ be the input and perturbation matrices respectively. Assume that the tokens $(\mathbf{x}_i, \mathbf{x}_i + \boldsymbol{\varepsilon}_i)$ lie in unit ball i.e. $\|\mathbf{X}\|_{2,\infty}, \|\mathbf{X} + \mathbf{E}\|_{2,\infty} \leq 1$. Let $\mathbf{V}, \mathbf{W} \in \mathbb{R}^{d \times d}$ be the weights of the self-attention layer obeying $\|\mathbf{V}\| \leq 1$ and $\|\mathbf{W}\| \leq 1$. Define the attention outputs $\mathbf{A} = \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\mathbf{X}$ and $\bar{\mathbf{A}} = \text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top)\bar{\mathbf{X}}$. We have that*

$$\|\mathbf{A}\|_{2,\infty}, \|\bar{\mathbf{A}}\|_{2,\infty} \leq 1, \quad \|\mathbf{A} - \bar{\mathbf{A}}\|_{2,1} \leq (2 + 1)e^2 \|\mathbf{E}\|_{2,1}.$$

Proof. First Observe that \mathbf{V} preserves norms i.e. $\mathbf{X}\mathbf{V}$ obeys $\|\mathbf{X}\mathbf{V}\|_{2,\infty} \leq \|\mathbf{X}\|_{2,\infty} \leq 1$ and $\|\mathbf{E}\mathbf{V}\|_{2,1} \leq \|\mathbf{E}\|_{2,1}$.

Next, set $\bar{\mathbf{X}} = \mathbf{X} + \mathbf{E}$ and define attention outputs $\mathbf{A} = \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\mathbf{X}\mathbf{V}$, $\bar{\mathbf{A}} = \text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top)\bar{\mathbf{X}}\mathbf{V}$. Observe that, since softmax applies row-wise to the similarities (e.g. $\mathbf{X}\mathbf{W}\mathbf{X}^\top$), we preserve the feature norms i.e. $\|\mathbf{A}\|_{2,\infty}, \|\bar{\mathbf{A}}\|_{2,\infty} \leq 1$ as advertised.

Now, consider the attention output difference $\mathbf{P} = \bar{\mathbf{A}} - \mathbf{A}$

$$\mathbf{P} = [\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)]\mathbf{X}\mathbf{V} + \text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top)\mathbf{E}\mathbf{V}.$$

For any pairs of tokens, we have $|\mathbf{x}_i^\top \mathbf{W} \mathbf{x}_j| \leq 1$. Using Lemma 1

$$\|\mathbf{P}_2\|_{2,1} = \|\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top)\mathbf{E}\mathbf{V}\|_{2,1} \leq n \|\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top)\|_{\infty} \|\mathbf{E}\|_{2,1} \leq e^2 \|\mathbf{E}\|_{2,1}.$$

Secondly, set $\mathbf{P}_1 = [\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)]\mathbf{X}\mathbf{V}$. We have that

$$\begin{aligned} \|\mathbf{P}_1\|_{2,1} &\leq \|\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\|_{\ell_1} \|\mathbf{X}\mathbf{V}\|_{2,\infty} \\ &\leq \|\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\|_{\ell_1}. \end{aligned}$$

To proceed, let us control the derivative by letting the perturbation $\mathbf{E} \rightarrow 0$ and then integrate the derivative along \mathbf{E} . Clearly, as $\mathbf{E} \rightarrow 0$, the quadratic-terms involving \mathbf{E} disappear and $\|\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\|_{\ell_1}$

$$\leq \|\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\mathbf{X}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\|_{\ell_1} + \|\text{softmax}(\mathbf{X}\mathbf{W}\bar{\mathbf{X}}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\|_{\ell_1}.$$

To bound the latter, consider each row individually, namely pick a row from $\mathbf{X}, \mathbf{X} + \mathbf{E}$ each denoted by the pair $(\mathbf{x}, \mathbf{x} + \boldsymbol{\varepsilon})$. Note that for any cross-product, we are guaranteed to have $|(\mathbf{x} + \boldsymbol{\varepsilon})^\top \mathbf{W} \mathbf{x}_i|, |\mathbf{x}^\top \mathbf{W} \mathbf{x}_i| \leq 1$, $\|\boldsymbol{\varepsilon}^\top \mathbf{W} \mathbf{X}\|_{\ell_1} \leq n \|\boldsymbol{\varepsilon}\|_{\ell_2}$, $\|\mathbf{x}^\top \mathbf{W} \mathbf{E}^\top\|_{\ell_1} \leq \|\mathbf{E}\|_{2,1}$. Applying perturbation bound of Lemma 1, we get

$$\|\text{softmax}((\mathbf{x} + \boldsymbol{\varepsilon})^\top \mathbf{W} \mathbf{X}^\top) - \text{softmax}(\mathbf{x}^\top \mathbf{W} \mathbf{X}^\top)\|_{\ell_1} \leq e^2 \|\boldsymbol{\varepsilon}\|_{\ell_1} \quad (9)$$

$$\|\text{softmax}(\mathbf{x}^\top \mathbf{W} (\mathbf{X} + \mathbf{E})^\top) - \text{softmax}(\mathbf{x}^\top \mathbf{W} \mathbf{X}^\top)\|_{\ell_1} \leq e^2 \|\mathbf{E}\|_{2,1}/n. \quad (10)$$

Adding up all n rows, we obtain (denoting $\boldsymbol{\varepsilon} = \|\mathbf{E}\|$)

$$\lim_{\boldsymbol{\varepsilon} \rightarrow 0} \|\text{softmax}(\bar{\mathbf{X}}\mathbf{W}\bar{\mathbf{X}}^\top) - \text{softmax}(\mathbf{X}\mathbf{W}\mathbf{X}^\top)\|_{\ell_1}/\boldsymbol{\varepsilon} \leq 2 e^2 \|\mathbf{E}\|_{2,1}/\boldsymbol{\varepsilon}.$$

Integrating the derivative along \mathbf{E} , we obtain $\|\mathbf{P}_1\|_{2,1} \leq 2 e^2 \|\mathbf{E}\|_{2,1}$. Cumulatively, we obtain the main claim $\|\mathbf{P}\|_{2,1} \leq (2 + 1)e^2 \|\mathbf{E}\|_{2,1}$. \blacksquare

Lemma 3 (Single-layer transformer stability) Consider the setup of Lemma 2. Let ϕ be a 1-Lipschitz activation function with $\phi(0) = 0$ (e.g. ReLU or Identity). Let $(\mathbf{M}_i)_{i=1}^n \in \mathbb{R}^{d \times d}$ be weights of the parallel MLPs following self-attention. Suppose $\|\mathbf{M}_i\| \leq 1$ and denote the MLP outputs associated to $\mathbf{A}, \bar{\mathbf{A}}$ by $\mathbf{B}, \bar{\mathbf{B}}$. We have that

$$\|\mathbf{B}\|_{2,\infty}, \|\bar{\mathbf{B}}\|_{2,\infty} \leq 1, \quad \|\mathbf{B} - \bar{\mathbf{B}}\|_{2,1} \leq (2 + \epsilon)e^2 \|\mathbf{E}\|_{2,1}.$$

Proof. First note that each row of $\bar{\mathbf{B}}$ is given by $\mathbf{b}_i = \phi(\mathbf{M}_i \mathbf{a}_i)$ thus $\|\mathbf{b}_i\|_{\ell_2} \leq \|\phi(\mathbf{M}_i \mathbf{a}_i)\|_{\ell_2} \leq \|\mathbf{M}_i \mathbf{a}_i\|_{\ell_2} \leq \|\mathbf{a}_i\|_{\ell_2} \leq 1$. Secondly, we can write $\|\mathbf{b}_i - \bar{\mathbf{b}}_i\|_{\ell_2} \leq \|\phi(\mathbf{M}_i \mathbf{a}_i) - \phi(\mathbf{M}_i \bar{\mathbf{a}}_i)\|_{\ell_2} \leq \|\mathbf{M}_i(\mathbf{a}_i - \bar{\mathbf{a}}_i)\|_{\ell_2} \leq \|\mathbf{a}_i - \bar{\mathbf{a}}_i\|_{\ell_2}$. Thus, $\|\mathbf{B} - \bar{\mathbf{B}}\|_{2,1} \leq \|\mathbf{A} - \bar{\mathbf{A}}\|_{2,1}$ and we conclude via Lemma 2. ■

Theorem 4 Consider a transformer composed of L layers with (1) output weights: $\mathbf{H} \in \mathbb{R}^{m \times d}$, (2) self-attention weights: combined key-query weights $(\mathbf{W}_i)_{i=1}^L \in \mathbb{R}^{d \times d}$ and value weights $(\mathbf{V}_i)_{i=1}^L \in \mathbb{R}^{d \times d}$, (3) MLP weights: $(\mathbf{M}_j^{(i)})_{i=1, j=1}^{L,m} \in \mathbb{R}^{d \times d}$ and activation $\phi \in \{\text{ReLU}, \text{Identity}\}$. For some $\epsilon > 0$, assume $\|\mathbf{V}_i\| \leq 1, \|\mathbf{M}_j^{(i)}\| \leq 1, \|\mathbf{W}_i\| \leq 1/2$ and $\|\mathbf{H}\|_{2,\infty} \leq c/m$. Suppose input space is $\mathcal{S} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_m)$ with $\|\mathbf{z}_i\|_{\ell_2} \leq 1$. The model prediction is given as follows

- $\mathcal{S}_{(0)} = \mathcal{S}$. Layer i outputs $\mathcal{S}_{(i)} = \text{Parallel_MLP}_{\mathbf{M}^{(i)}}(\text{Att}_{\mathbf{W}_i, \mathbf{V}_i}(\mathcal{S}_{(i-1)}))$. Here the self-attention layer is given by $\text{Att}_{\mathbf{W}_i, \mathbf{V}_i}(\mathcal{S}) = \text{softmax}(\mathbf{S}^\top \mathbf{W}_i \mathcal{S}) \mathbf{S} \mathbf{V}$ and *Parallel_MLP* applies $f(\mathbf{x}) = \phi(\mathbf{M}_j^{(i)} \mathbf{x})$ on j^{th} token of the *Att* output.
- $\text{TF}(\mathcal{S}) = \mathbf{H}, \mathcal{S}_{(L)}$.

First, this model is properly normalized in the sense that it can attain $\text{TF}(\mathcal{S}) = \pm c$. Second, this model obeys the stability guarantee

$$|\text{TF}(\mathcal{S}) - \text{TF}(\mathcal{S}')| \leq \frac{c}{m} ((1 + \epsilon)e)^\epsilon \|\mathcal{S} - \mathcal{S}'\|_{2,1}.$$

Proof. To see the first claim, let us set $\mathbf{V}_i = \mathbf{M}_{i,l} = \mathbf{I}$ and set all tokens \mathbf{z}_i to be identical i.e. $\mathcal{S} = \mathbf{1}_n \mathbf{z}^\top$. Additionally choose a \mathbf{z} with $\|\mathbf{z}\|_{\ell_2} = 1$ and nonnegative entries. Observe that, thanks to the softmax structure, regardless of \mathbf{W}_i , we have that $\mathcal{S} = \text{Att}_{\mathbf{W}_i, \mathbf{V}_i}(\mathcal{S}) = \text{softmax}(\mathbf{S}^\top \mathbf{W}_i \mathcal{S}) \mathcal{S}$. After attention, MLPs again preserves the tokens i.e. $\phi(\mathbf{M}_{i,l} \mathbf{z}_i) = \mathbf{z}$ for $\phi \in \{\text{ReLU}, \text{Identity}\}$. Thus, after proceeding L layers of this, we find that $\text{TF}(\mathcal{S}) = \langle \mathbf{H}, \mathcal{S} \rangle = \sum_{i=1}^n h_i^\top \mathbf{z}$. To conclude, by setting all tokens $h_i = c\mathbf{z}/m$ or $h_i = -c\mathbf{z}/m$, we obtain the claim. Note that, in general \mathcal{S} can be arbitrary (they don't have to be all same tokens): We can let $\mathbf{W} \rightarrow \infty$ (by allowing a larger ϵ). This way the attention matrix $\text{softmax}(\mathbf{S}^\top \mathbf{W} \mathcal{S}) \rightarrow \mathbf{I}$ so that we end up with the same argument of \mathcal{S} being (almost perfectly) transmitted across the layers and we obtain $\pm c$ output by setting $\mathbf{H} = \mathcal{S}/m$.

To show the stability guarantee, we use Lemmas 2 and 3. Specifically, Lemma 3 guarantees that

- After each layer we are guaranteed to have $\|\mathcal{S}_{(i)}\|_{2,\infty}, \|\mathcal{S}'_{(i)}\|_{2,\infty} \leq 1$.
- After each layer we are guaranteed to have $\|\mathcal{S}_{(i)} - \mathcal{S}'_{(i)}\|_{2,1} \leq (1 + \epsilon)e \|\mathcal{S}_{(i-1)} - \mathcal{S}'_{(i-1)}\|_{2,1}$.

This means that, at the end of the final layer, we have

$$\|\mathcal{S}_{(L)} - \mathcal{S}'_{(L)}\|_{2,1} \leq ((1 + \epsilon)e)^\epsilon \|\mathcal{S} - \mathcal{S}'\|_{2,1}.$$

Collapsing these by the linear \mathbf{H} , we find that $|\text{TF}(\mathcal{S}') - \text{TF}(\mathcal{S})| \leq \underbrace{\|\mathbf{H}, \mathcal{S}_{(L)} - \mathcal{S}'_{(L)}\|_{2,1}}_E \leq \|\mathbf{H}\|_{2,\infty} \|\mathcal{S}_{(L)} - \mathcal{S}'_{(L)}\|_{2,1}$ concluding the proof. ■

A.1 Proof of Theorem 1

Proof. We need to specialize Theorem 4 to obtain the result. Observe that when prompts differ only at the inputs $\mathbf{z}_j = (\mathbf{x}_j, y_j)$ with $\mathbf{z}'_j = (\mathbf{x}'_j, y'_j)$, we have that $\|\mathbf{X}_{\text{prompt}} - \mathbf{X}'_{\text{prompt}}\|_{2,1} \leq 2$. This implies that $|\text{TF}(\mathbf{X}_{\text{prompt}}) - \text{TF}(\mathbf{X}'_{\text{prompt}})| \leq \frac{2}{2^{m-1}} ((1 + \epsilon)e)^D$ for a depth D transformer. Finally, since the loss function ℓ is L -Lipschitz, we obtain the result $K = 2L((1 + \epsilon)e)^D$. ■

B Proofs and Supplementary Results for Sections 3,4,5

B.1 Proof of Theorem 2

Theorem 5 (Theorem 2 restated) *Suppose Assumption 1 holds and assume loss function $\ell(\mathbf{y}, \hat{\mathbf{y}})$ is L -Lipschitz for all $\mathbf{y} \in \mathcal{Y}$ and takes values in $[0, B]$. Let $\mathcal{A}g$ be the empirical solution of (ICL-ERM) and $\mathcal{N}(\mathcal{A}, \rho, u)$ be the covering number of the algorithm space \mathcal{A} following Definition 1&2. Then with probability at least $1 - 2\delta$, the excess MTL risk in (MTL-risk) obeys*

$$R_{MTL}(\mathcal{A}g) \leq \inf_{\varepsilon > 0} \left(4L\varepsilon + 2(B + K \log n) \frac{\log(\mathcal{N}(\mathcal{A}, \rho, \varepsilon)/\delta)}{cnT} \right).$$

Additionally, set $D := \sup_{\mathcal{A}g, \mathcal{A}g'} \rho(\mathcal{A}g, \mathcal{A}g')$ and assume $D < \infty$. With probability at least $1 - 4\delta$,

$$R_{MTL}(\mathcal{A}g) \leq \inf_{\varepsilon > 0} \left(8L\varepsilon + 8(2L + K \log n) \frac{D^{1/2}}{\varepsilon} \frac{\log \log \frac{D}{\varepsilon} \cdot \mathcal{N}(\mathcal{A}, \rho, u)/\delta}{c'nT} + 2(B + K \log n) \frac{\log(1/\delta)}{cnT} \right).$$

Proof. Recall the MTL problem setting of independent (input, label) pairs in Section 2: There are T tasks each with n in-context training samples denoted by $(\mathcal{S}_t)_{t=1}^T \stackrel{\text{i.i.d.}}{\sim} (\mathcal{D}_t)_{t=1}^T$ where $\mathcal{S}_t = \{(\mathbf{x}_{ti}, \mathbf{y}_{ti})\}_{i=1}^n$, and let $\mathcal{S}_{\text{all}} = \bigcup_{t=1}^T \mathcal{S}_t$. We use \mathcal{A} to denote the algorithm set. For a $\mathcal{A}g \in \mathcal{A}$, we define the training risk $\mathcal{L}_{\mathcal{S}_{\text{all}}}(\mathcal{A}g) = \frac{1}{nT} \sum_{t=1}^T \sum_{i=1}^n \ell(\mathbf{y}_{ti}, f_{\mathcal{S}_t^{(i-1)}}^{\mathcal{A}g}(\mathbf{x}_{ti}))$, and the test risk $\mathcal{L}_{\text{MTL}}(\mathcal{A}g) = \mathbb{E}[\mathcal{L}_{\mathcal{S}_{\text{all}}}(\mathcal{A}g)]$. Define empirical risk minima $\mathcal{A}g = \arg \min_{\mathcal{A}g \in \mathcal{A}} \mathcal{L}_{\mathcal{S}_{\text{all}}}(\mathcal{A}g)$ and population minima $\mathcal{A}g^* = \arg \min_{\mathcal{A}g \in \mathcal{A}} \mathcal{L}_{\text{MTL}}(\mathcal{A}g)$. For cleaner exposition, in the following discussion, we drop the subscripts MTL and \mathcal{S}_{all} . The excess MTL risk is decomposed as follows:

$$\begin{aligned} R_{\text{MTL}}(\mathcal{A}g) &= \mathcal{L}(\mathcal{A}g) - \mathcal{L}(\mathcal{A}g^*) \\ &= \underbrace{\mathcal{L}(\mathcal{A}g) - \mathcal{L}_{\mathcal{Z}}(\mathcal{A}g)}_a + \underbrace{\mathcal{L}_{\mathcal{Z}}(\mathcal{A}g) - \mathcal{L}_{\mathcal{Z}}(\mathcal{A}g^*)}_b + \underbrace{\mathcal{L}_{\mathcal{Z}}(\mathcal{A}g^*) - \mathcal{L}(\mathcal{A}g^*)}_c. \end{aligned}$$

Since $\mathcal{A}g$ is the minimizer of empirical risk, we have $b \leq 0$. To proceed, we consider the concentration problem of upper bounding $\sup_{\mathcal{A}g \in \mathcal{A}} |\mathcal{L}(\mathcal{A}g) - \mathcal{L}_{\mathcal{Z}}(\mathcal{A}g)|$.

Step 1: We start with a concentration bound $|\mathcal{L}(\mathcal{A}g) - \mathcal{L}_{\mathcal{Z}}(\mathcal{A}g)|$ for a fixed $\mathcal{A}g \in \mathcal{A}$. Recall that we define the test/train risks of each task as follows:

$$\begin{aligned} \mathcal{L}_t(\mathcal{A}g) &:= \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{y}_{ti}, f_{\mathcal{S}_t^{(i-1)}}^{\mathcal{A}g}(\mathbf{x}_{ti})), \quad \text{and} \\ \mathcal{L}_t(\mathcal{A}g) &:= \mathbb{E}_{\mathcal{S}_t} \mathcal{L}_t(\mathcal{A}g) = \mathbb{E}_{\mathcal{S}_t} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{y}_{ti}, f_{\mathcal{S}_t^{(i-1)}}^{\mathcal{A}g}(\mathbf{x}_{ti})), \quad \forall t \in [T]. \end{aligned}$$

Define the random variables $X_{t,i} = \mathbb{E}[\mathcal{L}_t(\mathcal{A}g) | \mathcal{S}_t^{(i)}]$ for $i \in [n]$ and $t \in [T]$, that is, $X_{t,i}$ is the expectation over $\mathcal{L}_t(\mathcal{A}g)$ given training sequence $\mathcal{S}_t^{(i)} = \{(\mathbf{x}_{tj}, \mathbf{y}_{tj})\}_{j=1}^i$ (which are the filtrations). With this, we have that $X_{t,n} = \mathbb{E}[\mathcal{L}_t(\mathcal{A}g) | \mathcal{S}_t^{(n)}] = \mathcal{L}_t(\mathcal{A}g)$ and $X_{t,0} = \mathbb{E}[\mathcal{L}_t(\mathcal{A}g)] = \mathcal{L}_t(\mathcal{A}g)$. More generally, $(X_{t,0}, X_{t,1}, \dots, X_{t,n})$ is a martingale sequence since, for every t in $[T]$, we have that $\mathbb{E}[X_{t,i} | \mathcal{S}_t^{(i-1)}] = X_{t,i-1}$.

For notational simplicity, in the following discussion, we omit the subscript t from \mathbf{x}, \mathbf{y} and \mathcal{Z} as they

will be clear from left hand-side variable $X_{t,i}$. We have that

$$\begin{aligned} X_{t,i} &= \mathbb{E} \frac{1}{n} \overset{\tilde{\mathcal{O}}}{\sum_{j=1}^n} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i)} \\ &= \frac{1}{n} \overset{\tilde{\mathcal{O}}}{\sum_{j=1}^i} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) + \frac{1}{n} \overset{\tilde{\mathcal{O}}}{\sum_{j=i+1}^n} \mathbb{E} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i)} \end{aligned}$$

Next, we wish to upper bound the martingale increments i.e. the difference of neighbors. Let $\mathcal{S}^{(i:i)} = \mathcal{S}^{(i)} - \mathcal{S}^{(i-1)}$ denote the i 'th element.

$$\begin{aligned} |X_{t,i} - X_{t,i-1}| &= \mathbb{E} \frac{1}{n} \overset{\tilde{\mathcal{O}}}{\sum_{j=1}^n} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i)} - \mathbb{E} \frac{1}{n} \overset{\tilde{\mathcal{O}}}{\sum_{j=1}^n} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i-1)} \\ &\leq \frac{1}{n} \overset{\tilde{\mathcal{O}}}{\sum_{j=i}^n} \mathbb{E} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i)} - \mathbb{E} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i-1)} \\ &\stackrel{(a)}{\leq} \frac{B}{n} + \frac{1}{n} \overset{\tilde{\mathcal{O}}}{\sum_{j=i+1}^n} \mathbb{E} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i)} - \mathbb{E} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i-1)}. \end{aligned}$$

Here, (a) follows from the fact that loss function $\ell(\cdot, \cdot)$ is bounded by B . To proceed, call the right side terms $D_j := |\mathbb{E}[\ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i)}] - \mathbb{E}[\ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Alg}}(\mathbf{x}_j)) \mathcal{S}^{(i-1)}]|$. Denote \mathbf{z}'_ℓ to be the realized values of $\mathbf{z}_\ell = (\mathbf{y}_\ell, \mathbf{x}_\ell)$ given $\mathcal{S}^{(i)}$. Let $\mathcal{S} := (\mathbf{z}'_1, \dots, \mathbf{z}'_i, \mathbf{z}_{i+1}, \dots, \mathbf{z}_j)$ and $\mathcal{S}' := (\mathbf{z}'_1, \dots, \mathbf{z}'_{i-1}, \mathbf{z}_i, \dots, \mathbf{z}_j)$. Observe that, \mathcal{S}' and \mathcal{S} differs in only at i th index and $i < j$, thus, utilizing Assumption 1,

$$D_j := |\mathbb{E}[\ell(\mathbf{y}_j, f_{\mathcal{S}}^{\text{Alg}}(\mathbf{x}_j))] - \mathbb{E}[\ell(\mathbf{y}_j, f_{\mathcal{S}'}^{\text{Alg}}(\mathbf{x}_j))]| \leq \frac{K}{j}. \quad (11)$$

Combining above, for any $n \geq i \geq 1$, we obtain

$$|X_{t,i} - X_{t,i-1}| \leq \frac{B}{n} + \overset{\tilde{\mathcal{O}}}{\sum_{j=i+1}^n} \frac{K}{jn} \leq \frac{B + K \log n}{n}.$$

Recall that $|\mathcal{L}_t(\text{Alg}) - \mathcal{B}_t(\text{Alg})| = |X_{t,0} - X_{t,n}|$ and for every $t \in [T]$, we have $\sum_{i=1}^n |X_{t,i} - X_{t,i-1}|^2 \leq \frac{(B+K \log n)^2}{n}$. As a result, applying Azuma-Hoeffding's inequality, we obtain

$$\mathbb{P}(|\mathcal{L}_t(\text{Alg}) - \mathcal{B}_t(\text{Alg})| \geq \tau) \leq 2e^{-\frac{n\tau^2}{2(B+K \log n)^2}}, \quad \forall t \in [T]. \quad (12)$$

Let us consider $Y_t := \mathcal{L}_t(\text{Alg}) - \mathcal{B}_t(\text{Alg})$ for $t \in [T]$. Then, $(Y_t)_{t=1}^T$ are i.i.d. zero mean sub-Gaussian random variables. There exists an absolute constant $c_1 > 0$ such that, the subgaussian norm, denoted by $\|\cdot\|_{\psi_2}$, obeys $\|Y_t\|_{\psi_2}^2 < \frac{c_1(B+K \log n)^2}{n}$ by [53, Proposition 2.5.2]. Applying Hoeffding's inequality, we derive

$$\mathbb{P} \frac{1}{T} \overset{\tilde{\mathcal{O}}}{\sum_{t=1}^T} Y_t \geq \tau \stackrel{!}{\leq} 2e^{-\frac{c_1 T \tau^2}{(B+K \log n)^2}} \implies \mathbb{P}(|\mathcal{B}(\text{Alg}) - \mathcal{L}(\text{Alg})| \geq \tau) \leq 2e^{-\frac{c_1 T \tau^2}{(B+K \log n)^2}} \quad (13)$$

where $c > 0$ is an absolute constant. Therefore, we have that for any $\text{Alg} \in \mathcal{A}$, with probability at least $1 - 2\delta$,

$$|\mathcal{B}(\text{Alg}) - \mathcal{L}(\text{Alg})| \leq (B + K \log n) \sqrt{\frac{\log(1/\delta)}{c_1 T}}. \quad (14)$$

Step 2: Next, we turn to bound $\sup_{\text{Alg} \in \mathcal{A}} |\mathcal{L}(\text{Alg}) - \mathcal{B}(\text{Alg})|$ where \mathcal{A} is assumed to be a continuous search space. To start with, set $g(\text{Alg}) := \mathcal{L}(\text{Alg}) - \mathcal{B}(\text{Alg})$ and we aim to bound $\sup_{\text{Alg} \in \mathcal{A}} |g(\text{Alg})|$.

Following Definition 2, for $\varepsilon > 0$, let \mathcal{A}_ε be a minimal ε -cover of \mathcal{A} in terms of distance metric ρ . Therefore, \mathcal{A}_ε is a discrete set with cardinality $|\mathcal{A}_\varepsilon| := \mathcal{N}(\mathcal{A}, \rho, \varepsilon)$. Then, we have

$$\sup_{\text{Al } g \in \mathcal{A}} |\mathcal{L}(\text{Al } g) - \mathfrak{L}(\text{Al } g)| \leq \sup_{\text{Al } g \in \mathcal{A}} \min_{\text{Al } g' \in \mathcal{A}_\varepsilon} |g(\text{Al } g) - g(\text{Al } g')| + \max_{\text{Al } g \in \mathcal{A}_\varepsilon} |g(\text{Al } g)|.$$

• We start by bounding $\sup_{\text{Al } g \in \mathcal{A}} \min_{\text{Al } g' \in \mathcal{A}_\varepsilon} |g(\text{Al } g) - g(\text{Al } g')|$. We will utilize that loss function $\ell(\cdot, \cdot)$ is L -Lipschitz. For any $\text{Al } g \in \mathcal{A}$, let $\text{Al } g' \in \mathcal{A}_\varepsilon$ be its neighbor following Definition 2. We have that

$$\begin{aligned} \mathfrak{L}(\text{Al } g) - \mathfrak{L}(\text{Al } g') &= \frac{1}{nT} \sum_{t=1}^n \sum_{i=1}^n \ell(\mathbf{y}_{ti}, f_{S_t^{(i-1)}}^{\text{Al } g}(\mathbf{x}_{ti})) - \ell(\mathbf{y}_{ti}, f_{S_t^{(i-1)}}^{\text{Al } g'}(\mathbf{x}_{ti})) \\ &\leq \frac{L}{nT} \sum_{t=1}^n \sum_{i=1}^n \|f_{S_t^{(i-1)}}^{\text{Al } g}(\mathbf{x}_{ti}) - f_{S_t^{(i-1)}}^{\text{Al } g'}(\mathbf{x}_{ti})\|_{\ell_2} \\ &\leq L\varepsilon. \end{aligned}$$

Since the same bound applies to all data-sequences, we also obtain that for any $\text{Al } g \in \mathcal{A}$,

$$|\mathcal{L}(\text{Al } g) - \mathcal{L}(\text{Al } g')| \leq L\varepsilon.$$

Therefore,

$$\sup_{\text{Al } g \in \mathcal{A}} \min_{\text{Al } g' \in \mathcal{A}_\varepsilon} |g(\text{Al } g) - g(\text{Al } g')| \leq \sup_{\text{Al } g \in \mathcal{A}} \min_{\text{Al } g' \in \mathcal{A}_\varepsilon} \mathfrak{L}(\text{Al } g) - \mathfrak{L}(\text{Al } g') + |\mathcal{L}(\text{Al } g) - \mathcal{L}(\text{Al } g')| \leq 2L\varepsilon. \quad (15)$$

• Next, we turn to bound the second term $\max_{\text{Al } g \in \mathcal{A}_\varepsilon} |g(\text{Al } g)|$. Applying union bound directly on \mathcal{A}_ε and combining it with (14), then we will have that with probability at least $1 - 2\delta$,

$$\max_{\text{Al } g \in \mathcal{A}_\varepsilon} |g(\text{Al } g)| \leq (B + K \log n) \frac{\sqrt{\log(\mathcal{N}(\mathcal{A}, \rho, \varepsilon)/\delta)}}{cnT}. \quad (16)$$

Proof of Eq. (4): Combining the upper bound above with the perturbation bound (15), we obtain that

$$\max_{\text{Al } g \in \mathcal{A}} |g(\text{Al } g)| \leq 2L\varepsilon + (B + K \log n) \frac{\sqrt{\log(\mathcal{N}(\mathcal{A}, \rho, \varepsilon)/\delta)}}{cnT}.$$

This in turn concludes the proof of (4) since $R_{\text{MTL}}(\text{Al } g) \leq 2 \sup_{\text{Al } g \in \mathcal{A}} |\mathcal{L}(\text{Al } g) - \mathfrak{L}(\text{Al } g)|$.

Proof of Eq. (5): To conclude, we aim to establish (5). To this end, we will bound $\max_{\text{Al } g \in \mathcal{A}_\varepsilon} |g(\text{Al } g)|$ via successive ε -covers which is the chaining argument. Following Definition 2, let $D := \sup_{\text{Al } g, \text{Al } g' \in \mathcal{A}} \rho(\text{Al } g, \text{Al } g')$. Define $M := \min\{m : 2^m \varepsilon \geq D\}$, and for any $m \in [M]$, let \mathcal{U}_m denote the minimal $2^m \varepsilon$ -cover of \mathcal{U}_{m-1} , where $\mathcal{U}_0 := \mathcal{A}_\varepsilon$. Since $\mathcal{U}_M \subseteq \mathcal{U}_{M-1} \cdots \subseteq \mathcal{U}_0 \subset \mathcal{A}$, we have $|\mathcal{U}_m| = \mathcal{N}(\mathcal{U}_{m-1}, \rho, 2^m \varepsilon) \leq \mathcal{N}(\mathcal{A}, \rho, 2^m \varepsilon)$ and $|\mathcal{U}_M| \leq \mathcal{N}(\mathcal{A}, \rho, D) = 1$. Let $\text{Al } g^M \in \mathcal{U}_M$ denote the unique algorithm hypothesis in \mathcal{U}_M . We have that

$$\begin{aligned} \max_{\text{Al } g \in \mathcal{A}_\varepsilon} |g(\text{Al } g)| &\leq \max_{\text{Al } g \in \mathcal{U}_0} |g(\text{Al } g) - g(\text{Al } g^M)| + |g(\text{Al } g^M)| \\ &\leq \sum_{m=0}^{M-1} \max_{\text{Al } g \in \mathcal{U}_m} \min_{\text{Al } g' \in \mathcal{U}_{m+1}} |g(\text{Al } g) - g(\text{Al } g')| + |g(\text{Al } g^M)|. \end{aligned} \quad (17)$$

In what follows, we will prove that for any $\text{Al } g, \text{Al } g'$ satisfying $\rho(\text{Al } g, \text{Al } g') \leq u$ ($u > 0$), with high probability, $|g(\text{Al } g) - g(\text{Al } g')|$ is bounded by $\frac{u}{\sqrt{nt}}$ up to logarithmic terms.

Apply similar martingale sequence analysis as in Step 1. This time, we set $X_{t,i} = E[\mathfrak{L}_t(\text{Al } g) - \mathfrak{L}_t(\text{Al } g') | \mathcal{S}_t^{(i)}]$ where we assume $\rho(\text{Al } g, \text{Al } g') \leq u$. Similarly, we have that $X_{t,n} = \mathfrak{L}_t(\text{Al } g) - \mathfrak{L}_t(\text{Al } g')$, and $X_{t,0} = E[\mathfrak{L}_t(\text{Al } g) - \mathfrak{L}_t(\text{Al } g')] = \mathcal{L}_t(\text{Al } g) - \mathcal{L}_t(\text{Al } g')$. Therefore, the sequences $\{X_{t,0}, \dots, X_{t,n}\}, t \in [T]$ are Martingale sequences

with respect to $E[X_{t,i}|S_t^{(i-1)}] = X_{t,i-1}$. We again omit the subscript t for \mathbf{x}, \mathbf{y} and \mathcal{S} in the following and try to bound the difference of neighbors.

$$\begin{aligned}
|X_{t,i} - X_{t,i-1}| &= E \mathfrak{L}_t(\text{Al g}) - \mathfrak{L}_t(\text{Al g}') \mathcal{S}^{(i)} - E \mathfrak{L}_t(\text{Al g}) - \mathfrak{L}_t(\text{Al g}') \mathcal{S}^{(i-1)} \\
&\leq \frac{1}{n} \overset{\textcircled{D}}{E} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_j)) - \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}'}(\mathbf{x}_j)) \mathcal{S}^{(i)} - E \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_j)) - \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}'}(\mathbf{x}_j)) \mathcal{S}^{(i-1)} \\
&\stackrel{(d)}{\leq} \frac{2Lu}{n} + \frac{1}{n} \overset{\textcircled{D}}{E}_{j=i+1} \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_j)) - \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}'}(\mathbf{x}_j)) \mathcal{S}^{(i)} - E \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_j)) - \ell(\mathbf{y}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}'}(\mathbf{x}_j)) \mathcal{S}^{(i-1)} \\
&\stackrel{(e)}{\leq} \frac{2Lu}{n} + \frac{1}{n} \overset{\textcircled{D}}{E}_{j=i+1} \frac{Ku}{j} < \frac{2Lu + Ku \log n}{n}.
\end{aligned}$$

for $i < n$. Here, (d) is from the facts that loss function $\ell(\cdot, \cdot)$ is L -Lipschitzness and $\rho(\text{Al g}, \text{Al g}') \leq u$ by following the same analysis in deriving (15), and (e) follows the Assumption 1. Then we have

$$|X_{t,n} - X_{t,n-1}| \leq \frac{2Lu}{n} < \frac{2Lu + Ku \log n}{n}.$$

Note that $|\mathcal{L}_t(\text{Al g}) - \mathcal{L}_t(\text{Al g}') - (\mathfrak{L}_t(\text{Al g}) - \mathfrak{L}_t(\text{Al g}'))| = |X_{t,0} - X_{t,n}|$ and for every $t \in [T]$, we have $\sum_{i=1}^n |X_{t,i} - X_{t,i-1}|^2 \leq \frac{u^2(2L+K \log n)^2}{n}$. As a result of applying Azuma-Hoeffding's inequality, we obtain

$$P(|\mathcal{L}_t(\text{Al g}) - \mathcal{L}_t(\text{Al g}') - (\mathfrak{L}_t(\text{Al g}) - \mathfrak{L}_t(\text{Al g}'))| \geq \tau) \leq 2e^{-\frac{n\tau^2}{2u^2(2L+K \log n)^2}}, \quad \forall t \in [T].$$

Now let us instead consider $Y_t := g(\text{Al g}) - g(\text{Al g}')$ for $t \in [T]$. Then following proof as in Step 1, we derive

$$P\left(\frac{1}{T} \overset{\textcircled{D}}{Y}_t \geq \tau\right) < 2e^{-\frac{c'nT\tau^2}{u^2(2L+K \log n)^2}} \implies P(|g(\text{Al g}) - g(\text{Al g}')| \geq \tau) \leq 2e^{-\frac{c'nT\tau^2}{u^2(2L+K \log n)^2}}$$

where $c' > 0$ is an absolute constant. Consider the discrete set \mathcal{U}_m with cardinality $|\mathcal{U}_m| = \mathcal{N}(\mathcal{U}_m, \rho, 2^m \varepsilon) \leq \mathcal{N}(\mathcal{A}, \rho, 2^m \varepsilon)$ and its $2^{m+1} \varepsilon$ -cover \mathcal{U}_{m+1} . Applying union bound over \mathcal{U}_m , we have that with probability at least $1 - 2\delta$,

$$\max_{\text{Al g} \in \mathcal{U}_m} \min_{\text{Al g}' \in \mathcal{U}_{m+1}} |g(\text{Al g}) - g(\text{Al g}')| \leq 2^{m+1} \varepsilon (2L + K \log n) \frac{\Gamma \overline{\log(\mathcal{N}(\mathcal{A}, \rho, 2^m \varepsilon)/\delta)}}{c'nT}.$$

Now by again applying union bound, with probability at least $1 - 2\delta$, the first term in (17) is bounded by

$$\begin{aligned}
\overset{\textcircled{D}}{M} \max_{m=0} \min_{\text{Al g} \in \mathcal{U}_m} \min_{\text{Al g}' \in \mathcal{U}_{m+1}} |g(\text{Al g}) - g(\text{Al g}')| &\leq (2L + K \log n) \overset{\textcircled{D}}{M} \frac{\Gamma \overline{\log(M \cdot \mathcal{N}(\mathcal{A}, \rho, 2^m \varepsilon)/\delta)}}{c'nT} \\
&\leq 4(2L + K \log n) \int_{\varepsilon/2}^1 \frac{\Gamma \overline{\log(M \cdot \mathcal{N}(\mathcal{A}, \rho, u)/\delta)}}{c'nT} du. \quad (18)
\end{aligned}$$

Now combining the results of (14), (17) and (18), and following the evidence that $\text{Al g}^M \in \mathcal{U}_M$ is unique, we bound $\sup_{\text{Al g} \in \mathcal{A}_\varepsilon} |g(\text{Al g})|$ as follows, that with probability at least $1 - 4\delta$

$$\sup_{\text{Al g} \in \mathcal{A}_\varepsilon} |g(\text{Al g})| \leq 4(2L + K \log n) \int_{\varepsilon/2}^1 \frac{\Gamma \overline{\log(M \cdot \mathcal{N}(\mathcal{A}, \rho, u)/\delta)}}{c'nT} du + (B + K \log n) \frac{\Gamma \overline{\log(1/\delta)}}{cnT}. \quad (19)$$

Here $D := \sup_{\text{Al g}, \text{Al g}' \in \mathcal{A}} \rho(\text{Al g}, \text{Al g}')$ and $M := \min\{m : 2^m \varepsilon \geq D\}$.

- Combining (15) and (19), we obtain that with probability at least $1 - 4\delta$,

$$\sup_{\text{Al } g \in \mathcal{A}} \mathcal{L}(\text{Al } g) - \mathfrak{L}(\text{Al } g) \leq \inf_{\varepsilon > 0} \left(4L\varepsilon + 4(2L + K \log n) \frac{D^{1/2} \int \frac{\log(M \cdot \mathcal{N}(\mathcal{A}, \rho, u)/\delta)}{c'nT} du + (B + K \log n) \int \frac{\log(1/\delta)}{cnT} \right),$$

where $D := \sup_{\text{Al } g, \text{Al } g' \in \mathcal{A}} \rho(\text{Al } g, \text{Al } g')$ and $M := \min\{m : 2^{m+1}\varepsilon \geq D\} \leq \log \frac{D}{\varepsilon}$.

Applying $R_{\text{MTL}}(\mathfrak{A}g) \leq 2 \sup_{\text{Al } g \in \mathcal{A}} \mathcal{L}(\text{Al } g) - \mathfrak{L}(\text{Al } g)$ completes the proof. \blacksquare

Till now, we consider the setting where each task is trained with only one trajectory. In the following, we also consider the case where each task contains multiple trajectories. To start with, we define the following objective function as an extension of (ICL-ERM) to the multi-trajectory setting.

$$\begin{aligned} \mathfrak{A}g &= \arg \min_{\text{Al } g \in \mathcal{A}} \mathfrak{L}_{\mathcal{S}_{\text{all}}}(\text{Al } g) := \frac{1}{TM} \sum_{t=1}^{\bar{T}} \sum_{m=1}^{\bar{M}} \mathfrak{L}_{t,m}(\text{Al } g) \\ \text{where } \mathfrak{L}_{t,m}(\text{Al } g) &= \frac{1}{n} \sum_{i=1}^{\bar{n}} \ell(\mathbf{y}_{tmi}, f_{\mathcal{S}_t^{(i-1)}}^{\text{Al } g}(\mathbf{x}_{tmi})). \end{aligned} \quad (20)$$

Here, we assume each task $t \in [T]$ contains M trajectories, and $\mathcal{S}_{\text{all}} = \{\{\mathcal{S}_{t,m}\}_{m=1}^M\}_{t=1}^T$ where $\mathcal{S}_{t,m} = \{(\mathbf{x}_{tmi}, \mathbf{y}_{tmi})\}_{i=1}^n$ and $(\mathbf{x}_{tmi}, \mathbf{y}_{tmi}) \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}_t$. Then the following theorem states a more general version of Theorem 2.

Theorem 6 *Suppose the same assumptions as in Theorem 2 hold and let $\mathfrak{A}g$ be the empirical solution of (20). Then, with the same probability, we obtain identical bounds to Theorem 2 by updating T with TM in Equations (4) and (5).*

Choosing $M = 1$ results in the exactly same bound as in Theorem 2.

Proof. Following the same proof steps, and then we derive similar result as (12):

$$\mathbb{P}(|\mathcal{L}_{t,m}(\text{Al } g) - \mathfrak{L}_{t,m}(\text{Al } g)| \geq \tau) \leq 2e^{-\frac{n\tau^2}{2(B+K \log n)^2}}, \quad \forall t \in [T], m \in [M]. \quad (21)$$

Let $Y_{t,m} := \mathcal{L}_{t,m}(\text{Al } g) - \mathfrak{L}_{t,m}(\text{Al } g)$. Since in-context samples are independent, $((Y_{t,m})_{m=1}^M)_{t=1}^T$ are independent zero mean sub-Gaussian random variables, with norm $\|Y_{t,m}\|_{\psi_2} < \frac{c_1(B+K \log n)^2}{n}$. Applying Hoeffding's inequality, we derive

$$\mathbb{P}(|\mathfrak{L}(\text{Al } g) - \mathcal{L}(\text{Al } g)| \geq \tau) \leq 2e^{-\frac{cnMT\tau^2}{(B+K \log n)^2}} \quad (22)$$

where $c > 0$ is an absolute constant. Therefore, we have that for any $\text{Al } g \in \mathcal{A}$, with probability at least $1 - 2\delta$,

$$|\mathfrak{L}(\text{Al } g) - \mathcal{L}(\text{Al } g)| \leq (B + K \log n) \int \frac{\log(1/\delta)}{cnMT}. \quad (23)$$

The result is simply replacing T with MT in (14). It is from the fact that trajectories are all independent no matter they are from the same task or not. By applying the similar analysis, the proof is completed. \blacksquare

B.2 Transfer Discussion from the Lens of Task Diversity

In Section 4, we motivated the fact that transfer risk is controlled in terms of MTL risk and an additive term that captures the distributional distance i.e. $\mathcal{L}_{\mathcal{T}}(\text{Al } g) \leq \mathcal{L}_{\text{MTL}}(\text{Al } g) + \text{dist}(\mathcal{T}, (\mathcal{D}_t)_{t=1}^T)$. The following definition is a generalization of this relation which can be used to formally control the transfer risk in terms of MTL risk.

Definition 4 (Task diversity) Following Section 2, we say that task \mathcal{T} is (ν, ϵ) -diverse over the T source tasks if for any $Alg, Alg' \in \mathcal{A}$,

$$\mathcal{L}_{\mathcal{T}}(Alg) - \mathcal{L}_{\mathcal{T}}(Alg') \leq \frac{1}{T} \sum_{t=1}^T (\mathcal{L}_t(Alg) - \mathcal{L}_t(Alg')) / \nu + \epsilon.$$

Now let us discuss transferability in light of this assumption and Thm 2. Consider the scenario where n is small and $T \rightarrow \infty$. The excess MTL risk will be small thanks to infinitely many tasks. The transfer risk would also be small because larger T results in higher diversity covering the task space. However, if the target task uses a different/longer prompt length, transfer may fail since the model never saw prompts longer than n . Conversely, if we let $n \rightarrow \infty$ and T to be small, although the MTL risk is again zero, due to lack of diversity, it may not benefit transfer learning strongly. Task diversity assumption leads to the following lemma that bounds transfer learning in terms of MTL risk.

Lemma 4 Consider the setting of Theorem 2. Let $\mathcal{A}lg$ be the solution of (ICL-ERM) and assume that target task \mathcal{T} is (ν, ϵ) -diverse over T source tasks. Then with the same probability as in Theorem 2, the excess transfer learning risk $R_{\mathcal{T}}(\mathcal{A}lg) = \mathcal{L}_{\mathcal{T}}(\mathcal{A}lg) - \min_{Alg \in \mathcal{A}} \mathcal{L}_{\mathcal{T}}(Alg)$ obeys $R_{\mathcal{T}}(\mathcal{A}lg) \leq \frac{R_{\text{MTL}}(\mathcal{A}lg)}{\nu} + 2\epsilon$.

Here we emphasize that the statement holds for arbitrary source and target tasks; however the challenge is verifying the assumption which is left as an interesting and challenging future direction. On the bright side, as illustrated in Figures 4&5, we indeed observe that, transfer learning can work with reasonably small T and it works better if the target task is closer to the source tasks.

Proof. Let $\mathcal{A}lg, Alg^*$ be the empirical and population solutions of (ICL-ERM) and let $Alg^\dagger := \arg \min_{Alg \in \mathcal{A}} \mathcal{L}_{\mathcal{T}}(Alg)$.

Then the transfer learning excess test risk of given algorithm $\mathcal{A}lg \in \mathcal{A}$ is formulated as

$$\begin{aligned} R_{\mathcal{T}}(\mathcal{A}lg) &= \mathcal{L}_{\mathcal{T}}(\mathcal{A}lg) - \mathcal{L}_{\mathcal{T}}(Alg^\dagger) \\ &= \mathcal{L}_{\mathcal{T}}(\mathcal{A}lg) - \mathcal{L}_{\mathcal{T}}(Alg^*) + \mathcal{L}_{\mathcal{T}}(Alg^*) - \mathcal{L}_{\mathcal{T}}(Alg^\dagger). \end{aligned}$$

Since target task \mathcal{T} is (ν, ϵ) -diverse over source tasks, following Definition 4, we derive that

$$\begin{aligned} \mathcal{L}_{\mathcal{T}}(\mathcal{A}lg) - \mathcal{L}_{\mathcal{T}}(Alg^*) &\leq \frac{\mathcal{L}_{\text{MTL}}(\mathcal{A}lg) - \mathcal{L}_{\text{MTL}}(Alg^*)}{\nu} + \epsilon = \frac{R_{\text{MTL}}(\mathcal{A}lg)}{\nu} + \epsilon \\ \mathcal{L}_{\mathcal{T}}(Alg^*) - \mathcal{L}_{\mathcal{T}}(Alg^\dagger) &\leq \frac{\mathcal{L}_{\text{MTL}}(Alg^*) - \mathcal{L}_{\text{MTL}}(Alg^\dagger)}{\nu} + \epsilon \leq \epsilon. \end{aligned}$$

Here, since Alg^* is the minimizer of $\mathcal{L}_{\text{MTL}}(Alg)$, $\mathcal{L}_{\text{MTL}}(Alg^*) - \mathcal{L}_{\text{MTL}}(Alg^\dagger) \leq 0$. Then, Lemma 4 is easily proved by combining the above two inequalities. \blacksquare

B.3 Transfer Learning Bound with i.i.d. Tasks

Following training with (ICL-ERM), suppose source tasks are i.i.d. sampled from a task distribution $\mathcal{D}_{\text{task}}$, and let $\mathcal{A}lg$ be the empirical MTL solution. We consider the following transfer learning problem. Concretely, assume a target task \mathcal{T} with a distribution $\mathcal{T} \sim \mathcal{D}_{\text{task}}$ and training sequence $\mathcal{S}_{\mathcal{T}} = (\mathbf{z}_i)_{i=1}^n \sim \mathcal{D}_{\mathcal{T}}$. Define the empirical and population risks on \mathcal{T} as $\mathcal{L}_{\mathcal{T}}(\mathcal{A}lg) = \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{y}_i, f_{\mathcal{S}_{\mathcal{T}}^{i-1}}^{\mathcal{A}lg}(\mathbf{x}_i))$ and $\mathcal{L}_{\mathcal{T}}(Alg) = \mathbb{E}_{\mathcal{S}_{\mathcal{T}}}[\mathcal{L}_{\mathcal{T}}(\mathcal{A}lg)]$.

Then the expected excess transfer risk following (ICL-ERM) is defined as

$$\mathbb{E}_{\mathcal{T}}[\mathcal{R}_{\mathcal{T}}(\mathcal{A}lg)] = \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\mathcal{A}lg)] - \arg \min_{Alg \in \mathcal{A}} \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(Alg)]. \quad (24)$$

Theorem 7 Consider the setting of Theorem 2 and assume the source tasks are independently drawn from task distribution $\mathcal{D}_{\text{task}}$. Let $\mathbf{A}g$ be the empirical solution of (ICL-ERM) and $\mathcal{T} \sim \mathcal{D}_{\text{task}}$. Then with probability at least $1 - 2\delta$, the expected excess transfer learning risk (24) obeys

$$\mathbb{E}_{\mathcal{T}}[\mathcal{R}_{\mathcal{T}}(\mathbf{A}g)] \leq \min_{\varepsilon \geq 0} \left(4L\varepsilon + B \frac{\sqrt{\frac{2 \log(\mathcal{N}(\mathcal{A}, \rho, \varepsilon)/\delta)}{T}}}{T} \right). \quad (25)$$

Proof. Recap the problem setting in Section 2 and let $\text{Al } g^{\dagger} = \arg \min_{\text{Al } g \in \mathcal{A}} \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g)]$. The expected transfer learning excess test risk of given algorithm $\text{Al } g \in \mathcal{A}$ is formulated as

$$\mathbb{E}_{\mathcal{T}}[\mathcal{R}_{\mathcal{T}}(\mathbf{A}g)] = \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\mathbf{A}g)] - \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g^{\dagger})] \quad (26)$$

$$= \underbrace{\mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\mathbf{A}g)]}_{a} - \underbrace{\mathcal{L}_{\mathcal{S}_{\text{all}}}(\mathbf{A}g)}_{\{Z\}} + \underbrace{\mathcal{L}_{\mathcal{S}_{\text{all}}}(\mathbf{A}g)}_{b} - \underbrace{\mathcal{L}_{\mathcal{S}_{\text{all}}}(\text{Al } g^{\dagger})}_{\{Z\}} + \underbrace{\mathcal{L}_{\mathcal{S}_{\text{all}}}(\text{Al } g^{\dagger})}_{c} - \underbrace{\mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g^{\dagger})]}_{\{Z\}}. \quad (27)$$

Here since $\mathbf{A}g$ is the minimizer of training risk, $b < 0$. Then we obtain

$$\mathbb{E}_{\mathcal{T}}[\mathcal{R}_{\mathcal{T}}(\mathbf{A}g)] \leq 2 \sup_{\text{Al } g \in \mathcal{A}} \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g)] - \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\text{Al } g). \quad (28)$$

For any $\text{Al } g \in \mathcal{A}$, let $X_t = \mathcal{L}_t(\text{Al } g)$ and we observe that

$$\mathbb{E}_{t \sim \mathcal{D}_{\text{task}}}[X_t] = \mathbb{E}_{t \sim \mathcal{D}_{\text{task}}}[\mathcal{L}_t(\text{Al } g)] = \mathbb{E}_{t \sim \mathcal{D}_{\text{task}}}[\mathcal{L}_t(\text{Al } g)] = \mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g)].$$

Since $X_t, t \in [T]$ are independent, and $0 \leq X_t \leq B$, applying Hoeffding's inequality obeys

$$\mathbb{P} \left[\mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g)] - \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\text{Al } g) \geq \tau \right] \leq 2e^{-\frac{2T\tau^2}{B^2}}. \quad (29)$$

Then with probability at least $1 - 2\delta$, we have that for any $\text{Al } g \in \mathcal{A}$,

$$\mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g)] - \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\text{Al } g) \leq B \frac{\sqrt{\frac{\log(1/\delta)}{2T}}}{2T}. \quad (30)$$

Next, let $\mathcal{A}_{\varepsilon}$ be the minimal ε -cover of \mathcal{A} following Definition 1, which implies that for any task $\mathcal{T} \sim \mathcal{D}_{\text{task}}$, and any $\text{Al } g \in \mathcal{A}$, there exists $\text{Al } g' \in \mathcal{A}_{\varepsilon}$

$$|\mathcal{L}_{\mathcal{T}}(\text{Al } g) - \mathcal{L}_{\mathcal{T}}(\text{Al } g')|, |\mathcal{L}_{\mathcal{S}_{\text{all}}}(\text{Al } g) - \mathcal{L}_{\mathcal{S}_{\text{all}}}(\text{Al } g')| \leq L\varepsilon.$$

Since the distance metric following Definition 2 is defined by the worst-case datasets, then there exists $\text{Al } g' \in \mathcal{A}_{\varepsilon}$ such that

$$\mathbb{E}_{\mathcal{T}}[\mathcal{L}_{\mathcal{T}}(\text{Al } g)] - \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\text{Al } g) \leq 2L\varepsilon. \quad (31)$$

Let $\mathcal{N}(\mathcal{A}, \rho, \varepsilon) = |\mathcal{A}_{\varepsilon}|$ be the ε -covering number. Combining the above inequalities ((28), (30) and (31)), and applying union bound, we have that with probability at least $1 - 2\delta$,

$$\mathbb{E}_{\mathcal{T}}[\mathcal{R}_{\mathcal{T}}(\mathbf{A}g)] \leq \min_{\varepsilon \geq 0} \left(4L\varepsilon + B \frac{\sqrt{\frac{2 \log(\mathcal{N}(\mathcal{A}, \rho, \varepsilon)/\delta)}{T}}}{T} \right).$$

■

B.4 Model Selection and Approximation Error Analysis

To proceed with our analysis, we need to make assumptions about what kind of algorithms are realizable by transformers. Given ERM is the work-horse of modern machine learning with general hypothesis classes, we assume that transformers can approximately perform in-context ERM. Hypothesis 1 states that the algorithms induced by the transformer can compete with empirical risk minimization over a family of hypothesis classes.

With this hypothesis, instead of searching over the entire hypothesis space $\mathcal{F}_{\text{all}} := \bigcup_{i=1}^H \mathcal{F}_i$, given prompt length m we search over the hypothesis space \mathcal{F}_{h_m} only, and $\text{comp}(\mathcal{F}_m) \leq \text{comp}(\mathcal{F}_{\text{all}})$ where $\text{comp}(\cdot)$ captures the complexity of a hypothesis class.

In Hypothesis 1, we assume that \mathbb{F} is a family of countable hypothesis classes with $|\mathbb{F}| = H$. As stated in Section 5, \mathbb{F} is not necessary to be discrete. The following provides some examples of \mathbb{F} , where the first three correspond to discrete model selection whereas the left are continuous.

- $\mathbb{F}^{\text{sparse}} = \{\mathcal{F}_s : s\text{-sparse linear model}\}$,
- $\mathbb{F}^{\text{NN}} = \{\mathcal{F}_s : 2\text{-layer neural net with width } s\}$,
- $\mathbb{F}^{\text{RF}} = \{\mathcal{F}_s : \text{Random forest with } s \text{ trees}\}$,
- $\mathbb{F}^{\text{ridge}} = \{\mathcal{F}_\lambda : \text{Linear model with parameter bounded by } \|\beta\|_{\ell_2} \leq \lambda\}$ (akin to ridge regression)
- $\mathbb{F}^{\text{weighted}} = \{\mathcal{F} : \text{Linear model with covariance-prior } \beta^\top \Sigma^{-1} \beta \leq 1\}$ (akin to weighted ridge).

To proceed, let us introduce the following classical result that controls the test risk of an ERM solution in terms of the Rademacher complexity [35, 32].

Theorem 8 *Let $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$ be a hypothesis set and let $\mathcal{S} = (\mathbf{x}_i, \mathbf{y}_i)_{i=1}^n \in \mathcal{X} \times \mathcal{Y}$ be a dataset sampled i.i.d. from distribution \mathcal{D} . Let $\ell(\mathbf{y}, \hat{\mathbf{y}})$ be a loss function takes values in $[0, B]$. Here $\ell(\mathbf{y}, \cdot)$ is L -Lipschitz in terms of Euclidean norm for all $\mathbf{y} \in \mathcal{Y}$. Consider a learning problem that*

$$\hat{f} := \arg \min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{y}_i, f(\mathbf{x}_i)). \quad (32)$$

Let $\mathcal{L}^* = \min_{f \in \mathcal{F}} \mathcal{L}(f)$ where $\mathcal{L}(f) = \mathbb{E}[\ell(\mathbf{y}, f(\mathbf{x}))]$. Then we have that with probability at least $1 - 2\delta$, the excess test risk obeys

$$\mathcal{L}(\hat{f}) - \mathcal{L}^* \leq 8L\mathcal{R}_n(\mathcal{F}) + 4B \frac{\log \frac{1}{\delta}}{n},$$

where $\mathcal{R}_n(\mathcal{F}) = \mathbb{E}_S \mathbb{E}_{\sigma_i} [\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \sigma_i^\top f(\mathbf{x}_i)]$ is the Rademacher complexity of \mathcal{F} [35] and σ_i 's are vectors with Rademacher random variable in each entry.

Lemma 5 (Formal version of Observation 1) *Let $\mathcal{L}_{\mathcal{T}}^* := \min_{\text{Alg} \in \mathcal{A}} \mathcal{L}_{\mathcal{T}}(\text{Alg})$ be the optimal target risk as stated in Section 2. Assume that Hypothesis 1 holds, then the approximation error obeys*

$$\mathcal{L}_{\mathcal{T}}^* \leq \frac{1}{n} \min_{m=1}^H \mathcal{L}_h^* + 8L\mathcal{R}_m(\mathcal{F}_h) + \varepsilon_{TF}^{h,m} + \frac{cB}{\sqrt{n}}, \quad (33)$$

where $\mathcal{R}_m(\mathcal{F})$ is the Rademacher complexity over data distribution $\mathcal{D}_{\mathcal{T}}$, and $\mathcal{L}_h^* = \min_{f \in \mathcal{F}_h} \mathbb{E}[\ell(\mathbf{y}, f(\mathbf{x}))]$.

Proof. Let us assume Hypothesis 1 holds for algorithm $\text{Alg} \in \mathcal{A}$. Since $\mathcal{L}_{\mathcal{T}}^*$ is the minimal test loss, we have that

$$\mathcal{L}_{\mathcal{T}}^* \leq \mathcal{L}_{\mathcal{T}}(\text{Alg}) = \mathbb{E}_{S_{\mathcal{T}}} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{y}_i, f_{S^{(i-1)}}^{\text{Alg}}(\mathbf{x}_i)) = \frac{1}{n} \mathbb{E}_{(\mathbf{x}, \mathbf{y}, S^{(i-1)})} \sum_{i=1}^n \ell(\mathbf{y}_i, f_{S^{(i-1)}}^{\text{Alg}}(\mathbf{x})).$$

Then by directly applying Hypothesis 1 we have that

$$\mathcal{L}_{\mathcal{T}}^{\star} \leq \frac{1}{n} \mathbb{E}_{(\mathbf{x}, \mathbf{y})} \ell(\mathbf{y}, f_{\mathcal{S}^{(0)}}^{\text{Alg}}(\mathbf{x})) + \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathcal{S}^{(i)})} \ell(\mathbf{y}, f_{\mathcal{S}^{(i)}}^{\text{Alg}}(\mathbf{x})) \quad (34)$$

$$\leq \frac{B}{n} + \frac{1}{n} \sum_{m=1}^n \min_{h \in [H]} \text{risk}(h, m) + \varepsilon_{\text{TF}}^{h, m}. \quad (35)$$

Here the first term in (35) comes from the fact that loss function is bounded by B , and we assume $\mathcal{S}^{(0)} = \emptyset$, and the second term follows the Hypothesis 1. Next, we turn to bound $\text{risk}(h, m)$. To proceed, let $X^{h, m} := \mathbb{E}_{(\mathbf{x}, \mathbf{y})} \ell(\mathbf{y}, \hat{f}_{\mathcal{S}^{(m)}}^{(h)}(\mathbf{x}))$ be the random variables, where we have $|X^{h, m}| \leq B$. Following Theorem 8, we have that for any $m \in [n], h \in [H]$

$$\mathbb{P} \left[X^{h, m} - \mathcal{L}_h^{\star} - 8L\mathcal{R}_m(\mathcal{F}_h) \geq \tau \right] \leq 2e^{-\frac{m\tau^2}{16B^2}}.$$

The upper-tail bound of the last line implies that there exists an absolute constant $c > 0$ such that

$$\text{risk}(h, m) = \mathbb{E}_{\mathcal{S}^{(m)}} X^{h, m} \leq \mathcal{L}_h^{\star} + 8L\mathcal{R}_m(\mathcal{F}_h) + \frac{cB}{\sqrt{m}}.$$

Combining it with (35) and following the evidence $\sum_{m=1}^n \frac{1}{\sqrt{m}} \leq 2\sqrt{n}$ complete the proof. \blacksquare

C Proof of Theorem 3

Lemma 6 *Suppose Assumptions 2 and 3 hold. Assume input and noise spaces \mathcal{X}, \mathcal{W} are bounded by \bar{x}, \bar{w} . Let $W = (\mathbf{w}_1, \dots, \mathbf{w}_j, \mathbf{w}_{j+1}, \dots, \mathbf{w}_m)$ and $W' = (\mathbf{w}_1, \dots, \mathbf{w}_{j-1}, \mathbf{w}'_j, \mathbf{w}_{j+1}, \dots, \mathbf{w}_m)$. Specifically, the only difference between W and W' is the j 'th term of the sequence. Then, for any $f \in \mathcal{F}$, $\text{Alg} \in \mathcal{A}$, W, W', m , and $j < m$, we have the following:*

$$\mathbb{E}_{\mathbf{x}_{m+1}} \ell(\mathbf{x}_{m+1}, f_{\mathcal{S}^{(m)}}^{\text{Alg}}(\mathbf{x}_m)) \mathbf{x}_0, W - \mathbb{E}_{\mathbf{x}_{m+1}} \ell(\mathbf{x}_{m+1}, f_{\mathcal{S}_j^{(m)}}^{\text{Alg}}(\mathbf{x}'_m)) \mathbf{x}_0, W' < \frac{K}{m} \frac{2\bar{C}_\rho \bar{w}}{1 - \bar{\rho}}.$$

Additionally, if consider the sequences with the initial states different only, then for any $\mathbf{x}'_0 \in \mathcal{X}$, we have

$$\mathbb{E}_{\mathbf{x}_{m+1}} \ell(\mathbf{x}_{m+1}, f_{\mathcal{S}^{(m)}}^{\text{Alg}}(\mathbf{x}_m)) \mathbf{x}_0 - \mathbb{E}_{\mathbf{x}_{m+1}} \ell(\mathbf{x}_{m+1}, f_{\mathcal{S}_j^{(m)}}^{\text{Alg}}(\mathbf{x}'_m)) \mathbf{x}'_0 < \frac{K}{m} \frac{2\bar{C}_\rho \bar{x}}{1 - \bar{\rho}}.$$

Proof. First, let us bound $\|\mathbf{x}_i - \mathbf{x}'_i\|_2$ for every $i = j, \dots, n$. For $i = j$, since \mathcal{W} is bounded by \bar{w} , we have

$$\|\mathbf{x}_j - \mathbf{x}'_j\|_2 = \|f(\mathbf{x}_{j-1}) + \mathbf{w}_j - f(\mathbf{x}_{j-1}) - \mathbf{w}'_j\|_2 \leq 2\bar{w}.$$

For $i > j$, we have the following from Assumption 2:

$$\|\mathbf{x}_i - \mathbf{x}'_i\|_2 \leq \bar{C}_\rho \bar{\rho}^{(i-j)} \|\mathbf{x}_j - \mathbf{x}'_j\|_2 \leq 2\bar{C}_\rho \bar{\rho}^{(i-j)} \bar{w}.$$

Finally, using Assumption 3, we obtain

$$\begin{aligned}
\mathbb{E}_{(\mathbf{x}_{m+1})} \ell(\mathbf{x}_{m+1}, f_{S^{(m)}}^{\text{Alg}}(\mathbf{x}_m)) \mathbf{x}_0, W &- \mathbb{E}_{(\mathbf{x}_{m+1})} \ell(\mathbf{x}_{m+1}, f_{S_j^{(m)}}^{\text{Alg}}(\mathbf{x}'_m)) \mathbf{x}_0, W' \\
&\leq \frac{K}{m} \sum_{i=j}^{\bar{m}} \|\mathbf{x}_i - \mathbf{x}'_i\|_2 \\
&\leq \frac{K}{m} 2\bar{C}_\rho \bar{w} \sum_{i=j}^{\bar{m}} \bar{\rho}^{i-j} < \frac{K}{m} \frac{2\bar{C}_\rho \bar{w}}{1 - \bar{\rho}}.
\end{aligned}$$

To prove the second part of the lemma, similarly we have

$$\|\mathbf{x}_0 - \mathbf{x}'_0\|_2 \leq 2\bar{x} \quad \text{and then,} \quad \|\mathbf{x}_i - \mathbf{x}'_i\|_2 \leq 2\bar{C}_\rho \bar{\rho}^i \bar{x}.$$

Again using Assumption 3, we obtain

$$\begin{aligned}
\mathbb{E}_{(\mathbf{x}_{m+1})} \ell(\mathbf{x}_{m+1}, f_{S^{(m)}}^{\text{Alg}}(\mathbf{x}_m)) \mathbf{x}_0 &- \mathbb{E}_{(\mathbf{x}_{m+1})} \ell(\mathbf{x}_{m+1}, f_{S_j^{(m)}}^{\text{Alg}}(\mathbf{x}'_m)) \mathbf{x}'_0 \\
&\leq \frac{K}{m} \sum_{i=0}^{\bar{m}} \|\mathbf{x}_i - \mathbf{x}'_i\|_2 \\
&\leq \frac{K}{m} 2\bar{C}_\rho \bar{x} \sum_{i=0}^{\bar{m}} \bar{\rho}^i < \frac{K}{m} \frac{2\bar{C}_\rho \bar{x}}{1 - \bar{\rho}}.
\end{aligned}$$

■

Theorem 9 (Theorem 3 restated) *Suppose Assumptions 2 and 3 hold and assume loss function $\ell(\mathbf{x}, \hat{\mathbf{x}}) : \mathcal{X} \times \mathcal{X} \rightarrow [0, B]$ is an L -Lipschitz function for all $\mathbf{x} \in \mathcal{X}$. Let \mathfrak{Alg} be the empirical solutions of (ICL-ERM) under dynamical setting as described in Section 2&6. Then with probability at least $1 - 2\delta$, the excess MTL test risk (MTL-risk) obeys*

$$R_{MTL}(\mathfrak{Alg}) \leq \inf_{\varepsilon > 0} \left(4L\varepsilon + 2(B + \bar{K} \log n) \frac{\log(N(\mathcal{A}, \rho, \varepsilon)/\delta)}{cnT} \right).$$

where $\bar{K} = 2K \frac{\bar{C}_\rho}{1 - \bar{\rho}} (\bar{w} + \bar{x}/\sqrt{n})$.

Proof. We follow the similar way as in the proof of Theorem 2, and the only difference is that we need to consider the dynamical system setting. Therefore, to start with, let us recap the dynamical problem in Section 2&6. Suppose there are T independent trajectories generated by dynamical systems, denoted by $\mathcal{S}_t = (\mathbf{x}_{t0}, \mathbf{x}_{t1}, \dots, \mathbf{x}_{tn})$, $t \in [T]$ where $\mathbf{x}_{ti} = f_t(\mathbf{x}_{t,i-1}) + \mathbf{w}_{ti}$. Here, consider the prediction function $f_{S^{(i)}}^{\text{Alg}} : \mathcal{X} \rightarrow \mathcal{X}$, and let previously observed sequences be $\mathcal{S}_t^{(i)} := (\mathbf{x}_{t0}, \dots, \mathbf{x}_{t,i})$. The objective function in (ICL-ERM) can be rewritten as follows:

$$\begin{aligned}
\mathfrak{Alg} &= \arg \min_{\text{Alg} \in \mathcal{A}} \mathfrak{L}_{S_{\text{all}}}(\text{Alg}) := \frac{1}{T} \sum_{t=1}^{\bar{T}} \mathfrak{L}_t(\text{Alg}) \\
\text{where } \mathfrak{L}_t(\text{Alg}) &= \frac{1}{n} \sum_{i=1}^{\bar{n}} \ell(\mathbf{x}_{ti}, f_{S_t^{(i-1)}}^{\text{Alg}}(\mathbf{x}_{t,i-1})).
\end{aligned} \tag{36}$$

Following the same way as in the proof of Theorem 2, excess MTL risk is bounded by:

$$R_{\text{MTL}}(\mathcal{A} \text{ g}) \leq 2 \sup_{\text{Al g} \in \mathcal{A}} |\mathcal{L}(\text{Al g}) - \mathcal{B}(\text{Al g})|.$$

Step 1: We start with the concentration bound $|\mathcal{L}(\text{Al g}) - \mathcal{B}(\text{Al g})|$ for any $\text{Al g} \in \mathcal{A}$. Instead, define the random variables $X_{t,i} = \mathbb{E}[\mathcal{B}_t(\text{Al g}) | \mathbf{x}_{t0}, (\mathbf{w}_{tk})_{k=1}^i]$ for $i \in [n]$ and $t \in [T]$, that is, $X_{t,i}$ is the expectation over $\mathcal{B}_t(\text{Al g})$ given the filtration of \mathbf{x}_{t0} and $(\mathbf{w}_{tk})_{k=1}^i$. Then, we have that $X_{t,n} = \mathbb{E}[\mathcal{B}_t(\text{Al g}) | \mathbf{x}_{t0}, (\mathbf{w}_{tk})_{k=1}^n] = \mathcal{B}_t(\text{Al g})$. Let $X_{t,0} = \mathbb{E}[\mathcal{B}_t(\text{Al g})]$. Then, for every t in $[T]$, the sequences $\{X_{t,0}, X_{t,1}, \dots, X_{t,n}\}$ are Martingale sequences. Here we emphasize that $X_{t,0} = \mathbb{E}[X_{t,1} | \mathbf{x}_{t0}, \mathbf{w}_{t1}]$. For the sake of simplicity in the notation, we omit the subscript t for \mathbf{x}, \mathbf{w} and \mathcal{S} , and look at the difference of neighbors for $1 < i < n$. Here, given $\{\mathbf{x}_0, (\mathbf{w}_k)_{k=1}^i\}$ implies $\{\mathbf{x}_0, \dots, \mathbf{x}_i\}$ are known.

$$\begin{aligned} |X_{t,i} - X_{t,i-1}| &= \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^{\tilde{n}} \ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^i - \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^{\tilde{n}} \ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^{i-1} \right] \right] \\ &\leq \frac{1}{n} \sum_{j=i}^{\tilde{n}} \mathbb{E} \left[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^i - \mathbb{E} \left[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^{i-1} \right] \right] \\ &\stackrel{(a)}{\leq} \frac{B}{n} + \frac{1}{n} \sum_{j=i+1}^{\tilde{n}} \mathbb{E} \left[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^i - \mathbb{E} \left[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^{i-1} \right] \right] \end{aligned}$$

Here, (a) follows from the fact that loss function $\ell(\cdot, \cdot)$ is bounded by B . To proceed, call the right side terms $D_j := |\mathbb{E}[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^i] - \mathbb{E}[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^{i-1}]|$. For an arbitrary sequence of $(\mathbf{w}'_k)_{k=i+1}^n$, we obtain using the first part of Lemma 6 that

$$\begin{aligned} &\mathbb{E}[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^i, (\mathbf{w}_k)_{k=i+1}^n = (\mathbf{w}'_k)_{k=i+1}^n] \\ &\quad - \mathbb{E}[\ell(\mathbf{x}_j, f_{\mathcal{S}^{(j-1)}}^{\text{Al g}}(\mathbf{x}_{j-1})) \mathbf{x}_0, (\mathbf{w}_k)_{k=1}^{i-1}, (\mathbf{w}_k)_{k=i+1}^n = (\mathbf{w}'_k)_{k=i+1}^n] \leq \frac{K}{j-1} \frac{2\bar{C}_\rho \bar{w}}{1-\bar{\rho}} \end{aligned}$$

Since that is valid for every realization of $(\mathbf{w}_k)_{k=i+1}^n$, we have

$$D_j \leq \frac{K}{j-1} \frac{2\bar{C}_\rho \bar{w}}{1-\bar{\rho}}$$

Combining above, for any $n > i > 1$, we obtain

$$|X_{t,i} - X_{t,i-1}| \leq \frac{B}{n} + \frac{1}{n} \sum_{j=i+1}^{\tilde{n}} \frac{K}{j-1} \frac{2\bar{C}_\rho \bar{w}}{1-\bar{\rho}} < \frac{B}{n} + \frac{K \log n}{n} \frac{2\bar{C}_\rho \bar{w}}{1-\bar{\rho}}$$

If we apply the same step above and use the second part of Lemma 6, we obtain the following bound for $|X_{t,1} - X_{t,0}|$:

$$|X_{t,1} - X_{t,0}| < \frac{B}{n} + \frac{K \log n}{n} \frac{2\bar{C}_\rho (\bar{w} + \bar{x})}{1-\bar{\rho}}$$

Moreover, as the loss function is bounded by B , we have

$$|X_{t,n} - X_{t,n-1}| \leq \frac{B}{n} < \frac{B}{n} + \frac{K \log n}{n} \frac{2\bar{C}_\rho \bar{w}}{1-\bar{\rho}}$$

Note that $|\mathcal{L}_t(\text{Al g}) - \mathcal{B}_t(\text{Al g})| = |X_{t,0} - X_{t,n}|$ and for every $t \in [T]$, we obtain

$$\sum_{i=1}^{\tilde{n}} |X_{t,i} - X_{t,i-1}|^2 \leq \frac{(n-1) \left(\frac{B}{n} + \frac{K \log n}{n} \frac{2\bar{C}_\rho \bar{w}}{1-\bar{\rho}} \right)^2 + \left(\frac{B}{n} + \frac{K \log n}{n} \frac{2\bar{C}_\rho (\bar{w} + \bar{x})}{1-\bar{\rho}} \right)^2}{n^2} \leq \frac{B + K \log n \frac{2\bar{C}_\rho (\bar{w} + \bar{x})}{1-\bar{\rho}}}{n}^2$$

Then the equivalent result in the Step 1 of the proof of Theorem 2 is obtained by updating K to be $\bar{K} = 2K \frac{\bar{C}_\rho}{1-\rho} (\bar{w} + \bar{x}/\sqrt{n})$.

Step 2: Next, we turn to bound $\sup_{\text{AI } g \in \mathcal{A}} |\mathcal{L}(\text{AI } g) - \mathcal{P}(\text{AI } g)|$ where \mathcal{A} is assumed to be a continuous search space. Following the identical analysis as in Step 2 of proof of Theorem 2, by applying a covering argument, we conclude with the the result. ■

D Additional Experiments for Section 4

In Section 4, we discussed how transfer risk seems to require $T \propto d^2$ source tasks. In contrast, constructing the empirical covariance $\hat{\Sigma} = \frac{1}{T} \sum_{i=1}^T \beta_i \beta_i^\top$ can make sure that $\hat{\Sigma}$ -weighted LS performs $\mathcal{O}(1)$ -close to Σ -weighted LS whenever $\|\hat{\Sigma}^{-1}\|/\lambda_{\min}(\Sigma) \leq \mathcal{O}(1)$. In Figure 8, *MTL-ridge* curves with circle markers are referring to the $\hat{\Sigma}$ -weighted ridge regression. As suspected, $T = 3d$ is already sufficient to get very close performance to the optimal weighting with true Σ (black curve). We remark that in Figure 8 $d = 20$, noise variance obeys $\sigma^2 = 0.1$, and linear task vectors β are uniformly sampled over the sphere.

For MTL, Section 4 introduces the following simple greedy algorithm to predict a prompt that belong to one of the T source tasks (aka MTL risk): Evaluate each source task parameter $(\beta_i)_{i=1}^T$ on the prompt and select the parameter with the minimum risk. Since there are T choices, this greedy algorithm will determine the optimal task in $n \cdot \log(T)$ samples⁶. The experiments of this algorithm is provided under *Greedy MTL* legend (square markers). It can be seen that as T varies ($d, 2d, 3d$), there is almost no difference in the MTL risk, likely due to the $\log(T)$ dependence. In short, these experiments highlight the contrast between ideal/greedy algorithms and ICL algorithm implemented within the transformer model.

In Section 4, we also exclusively focused on the setting $M \rightarrow \infty$ i.e. MTL tasks are thoroughly trained. In Figure 9 we consider the other extreme where each task is trained with a single trajectory $M = 1$, which is closer to the spirit of Theorem 2. We set $d = 5$, $n = 10$ and $M = 1$, and vary the number of linear regression tasks T from 2000 to 50000. Not surprisingly, the results show that increasing T helps in reducing the MTL risk. The more interesting observation is that transfer risk and MTL risk are almost perfectly aligned. We believe that this is due to the small M, n choices which would make it difficult to overfit to the MTL tasks. Thus, when $M = 1$, the gap between transfer and MTL risk seems to vanish and Theorem 2 becomes directly informative for the transfer risk. In contrast, as M grows, training process can overfit to the MTL tasks which leads to the split between MTL and transfer risks as in Figure 4.




Figure 8: We display the performance of the *idealized* transfer and MTL algorithms described in Section 4. Unlike ICL experiments, these require $T \cdot d$ tasks.




Figure 9: Comparing MTL and transfer risks when each task has single trajectory ($M = 1$).

⁶Note that, this dependence can be even better if the problem is noiseless, in fact, that is why we added label noise in these experiments.